
Understanding and designing animations in the user interfaces

Rédigée par:
Amira Chalbi Neffati

Dirigée par:
Nicolas Roussel
Fanny Chevalier

*Thèse présentée et soutenue publiquement le 17 avril 2018
pour obtenir le grade de Docteur de l'Université de Lille dans la spécialité Informatique.
École Doctorale Sciences pour l'Ingénieur Lille Nord-de-France*

Composition du jury:

Présidente: Sophie Tison (Université de Lille)

Rapporteurs: Stéphane Conversy (Université de Toulouse-ENAC)
Jason Dykes (City, University of London)

Examineurs: Anastasia Bezerianos (Université Paris-Sud 11)
Benjamin Bach (Edinburgh University)

Abstract

Despite their increasing popularity and omnipresence in modern graphical interfaces, animations are still largely under-comprehended. While prior research and practice provide useful insights about the merits and downsides of animation, it is still unclear what makes a good and effective animation that improves the usability and expressivity of graphical interfaces. The disparity of opinions about the value of animation is mainly due to the fact that most of previous studies have investigated the benefit of adding a particular animation to a particular interface, leaving a notable gap in the deep understanding of the many design aspects that influence the performance of animation. Prior research have also predominantly assessed the value of animation through a narrow empirical angle, which had left several facets of animation unveiled.

This thesis contributes a first constructive step toward better understanding the vast design space of animation and mapping out the various merits of animation that can enrich user interfaces from different perspectives. We first provide a structured view of the roles and drawbacks of animation in user interfaces. We then present the theoretical fundamentals for animation in information visualization. We discuss the main challenges for designing and evaluating animation in dynamic visualizations. Through an empirical study, we investigate the meaning of the Common Fate Law, applied on animation trajectories, in dynamic visualizations. We then introduce a design space that allows a holistic characterization of staged animation and propose an authoring tool to support the prototyping and exploration of staging in visualizations.

Résumé

Malgré leur omniprésence croissante dans les interfaces graphiques modernes, les animations sont encore sous-appréhendées. Bien que la recherche et la pratique fournissent des directives utiles, nous n'avons pas encore une idée claire sur ce qui définit une animation efficace pour les interfaces. La discordance sur la valeur de l'animation est principalement due au fait que la plupart des études précédentes ont étudié l'avantage d'ajouter une animation particulière à une interface particulière, laissant un manque notable dans la compréhension des nombreux aspects qui influencent la performance de l'animation. Ces études ont aussi évalué la valeur de l'animation à travers un angle empirique étroit, ce qui fait que plusieurs facettes de l'animation restent non dévoilées.

Cette thèse contribue un premier pas vers une compréhension plus profonde du vaste espace de conception de l'animation, et vers une meilleure assimilation des différents mérites de l'animation. Nous fournissons d'abord une vue structurée sur les rôles et les inconvénients de l'animation dans les interfaces graphiques. Nous présentons ensuite les fondements théoriques de l'animation pour la visualisation d'information. Nous discutons les principaux défis pour la conception et l'évaluation de l'animation dans les visualisations dynamiques. Nous étudions la signification de la Loi du Destin Commun, appliquée aux trajectoires d'animation, dans les visualisations dynamiques. Nous introduisons ensuite un espace de conception qui permet une caractérisation holistique de l'animation séquentielle, et proposons un outil de création pour permettre le prototypage et l'exploration du "staging" dans les visualisations.

Contents

1	Introduction	1
1.1	Thesis statement and goals	8
1.2	Research contributions	8
2	Roles and drawbacks of animation in user interfaces	10
2.1	The 23 roles of animation in visual interfaces	11
2.1.1	Preliminary qualitative exploration	12
2.1.2	Taxonomy Methodology	13
2.1.3	Taxonomy of animation roles	14
	Keeping in Context	15
	Teaching Aid	20
	User Experience	25
	Data Encoding	30
	Visual Discourse	33
2.1.4	Animation roles and people: summary and discussion	36
2.1.5	General discussion	38
	Coverage of general roles of animation by our taxonomy	38
	Combinations of roles and granularity	39
	What limits the roles of animation?	39
2.2	Animation drawbacks and people	40
2.2.1	Why people do not use animation?	40
2.2.2	Examples of ineffective animation as seen by people	41
2.2.3	Classification of animation drawbacks as seen by people	42
2.3	Conclusion	42
3	Understanding dynamic information visualizations	44
3.1	Information visualization: a simplified model	44
3.2	From abstract data to visual representation: visual encoding	45
3.3	Dynamics of an information visualization	47
3.3.1	Visualization states mapping	48
3.3.2	Characterizing dynamic changes in information visualization	49
3.3.3	A taxonomy of dynamic changes in information visualization	51
	Levels of dynamic change	51
	Data-Driven Change	52
	Visual-Driven Change	53
	Variation of types of change based on the scope of change	53
3.3.4	Manner of a dynamic change	54
3.4	Conclusion	56
4	Animation for information visualization: fundamentals and challenges	57
4.1	Fundamentals of animation	57
4.1.1	Animation parameters	57
4.1.2	Elementary and composite animation	58
4.1.3	Composition based on trajectory	59
4.1.4	Temporal composition	60
4.2	Designing animation for information visualizations	61

4.2.1	Understanding animation parameters	61
	Understanding animation trajectories	61
	Understanding the temporal parameters of animation	62
4.2.2	Defining animation parameters	63
4.3	Evaluating animation in information visualizations	65
4.3.1	Metrics to measure the value of animation	65
4.3.2	Evaluation methodology	67
	Studied animation goals	67
	Evaluation methodology	68
	General user study design considerations	68
	Independent variables	70
	Dependent variables	70
	Tasks	70
	Running a study	71
4.4	People and animation: insights for animation design and evaluation	71
4.5	Conclusion	72
5	Reflecting on the meaning of Common Fate in visualization	74
5.1	Background	74
5.1.1	Perception and Gestalt psychology	75
5.1.2	Gestalt Laws in visualization	75
5.1.3	Perception in visualization	75
5.2	The Gestalt Law of Common Fate	76
5.2.1	Criteria for perceptual grouping	76
5.2.2	Visual variables with grouping	77
5.3	User study background	78
5.3.1	Task rationale	78
5.3.2	Summary of tasks	78
5.3.3	Manipulation of visual variables	79
5.3.4	Design decisions	81
5.3.5	Grouping strength	81
5.4	User study	81
5.4.1	Task generation	82
5.4.2	Attention trials	83
5.4.3	Experimental protocol	83
5.5	Results	84
5.5.1	Data verification	85
5.5.2	Dynamic vs. static groupings	85
5.5.3	Relations between Gestalt Laws	86
5.5.4	Participants feedback	87
5.6	Discussion	87
5.6.1	Implications for dynamic information visualization	88
5.6.2	Limitations	88
5.7	Evaluation scenarios to assess our design implications	89
5.7.1	Make a particular group salient	89
5.7.2	Revealing multiple groups	90
5.8	Conclusion	90

6	Exploring and supporting the design of staged animation	92
6.1	Staged Animation	93
6.1.1	Motivation	93
6.1.2	Limitations and Challenges of Staging	95
6.1.3	Staging in Traditional Animation	96
6.1.4	Animation Choreography: a Model for Staging	96
6.2	Design Space Overview	97
6.3	Design Space: Fixed Dimensions	97
6.3.1	Data Attributes	98
6.3.2	Change Characteristics	98
6.3.3	Observer Task	99
6.4	Design Space: Variable Dimensions	99
6.4.1	Staging Strategy	100
6.4.2	Staging Design	100
	Decomposition	101
	Timing	103
6.5	Values of Animation Choreography	107
6.6	Describing Animation Choreography	114
6.6.1	Staging Hierarchy	114
6.6.2	Scenario: Exploring a Co-authorship Network	115
6.7	Authoring Animation Choreography	117
6.7.1	Design considerations	117
6.7.2	Prototype	119
6.8	Conclusion	120
7	Conclusion & perspectives	122
7.1	Summary	122
7.2	Future work	123
A	Publications	126
B	Generation rules of employees dataset	127
	Bibliography	130

List of Figures

1.1	One of the Chauvet Cave paintings illustrating a scene of hunting. Azéma and Rivère claimed that flickering torches passed over these painted scenes would have heightened onlookers’ sense of seeing live-action stories thanks to the optical phenomenon known as the persistence of vision [Sto; AR12].	1
1.2	Ancient artists at Chauvet Cave superimposed drawings of two stances of a bison to depict the movement of running [Prea].	2
1.3	A painting discovered at Lascaux Cave where three images of horse’s head are superimposed to depict the movement of head tossing [Prea].	2
1.4	A drawing found at the Cave of the Trois-Frères where ancient artists drew two superimposed images of an animal trail to depict the movement of trail’s shaking [Prea].	2
1.5	The Vitruvian Man: a drawing by Leonardo da Vinci done around 1490 depicting a man in two superimposed positions with his arms and legs apart and inscribed in a circle and square [Vit].	2
1.6	The onion skinning tool in Autodesk Maya [Oni].	2
1.7	A freeze-frame drawing representing a lion running from left to right discovered at the La Vache Cave [AR12].	3
1.8	A reconstruction of the The Laugerie-Basse spinning disk by Rivère shows a chamois in different positions on each side: standing on one side of the disc (left image), and lying on the other (right image) [AR12]. When the disk is twirled on a string, the creature appears to move (see animation in [Dis]). . .	3
1.9	Top: a demonstration of a bird and cage thaumatrope [Thaa]. Bottom: Both sides of a bouquet and vase thaumatrope from 1825 [Thab].	4
1.10	(a) Mechanism of the chronophotographic gun, (b) Marey used the gun to study the movement of birds, (c) an “animated photograph” recording 12 movement’s phases of a group of birds [Mara; Marb].	5
2.1	A staged animation to support navigation between directories in a 3D tree-map visualization. When a user opens a directory, (a) the directory begins with a fade-out, (b) followed by a short pause, then (c,d) the new current directory is expanded to fill the space formerly occupied by its ancestor. [BCK05]	16
2.2	Illustration of Role 2: Details of an animated transition between a PDF article and its corresponding LaTeX source code in Gliimpse [DHC11]	17
2.3	Illustration of Role 3: PlanePlotter [Air] visualizes air traffic in real time. . .	18
2.4	Illustration of Role 4: Mnemonic Desktop [BDB06] reveals hidden parts of windows by showing motion trails and a dusty appearance that will fade out as changes are being replayed.	18
2.5	Illustration of Role 5: When the user selects a tool in Google Sketchup’s instructor [Ske], the instructor shows her an animation that explains the basic use of this selected tool.	21
2.6	Illustration of Role 6: GestureBar [Bra+09] shows an illustrative animation to teach the user the gesture to perform in order to delete a shape.	21
2.7	Illustration of Role 7: Educative animation using bouncing arrows to illustrate blood flow in the heart—adapted from “Exploring the Heart-The Circulatory System!” by AboutKidsHealth (Youtube).	22
2.8	Illustration of Role 8: An animation that illustrates how the bubble sorting algorithm works [Vis].	23

2.9	Illustration of Role 9: An animation in NodeTrix [HFM07] that explains step-by-step the change from a node-link diagram (1) to an adjacency matrix (5)	24
2.10	Illustration of Role 10: Soundhound [Sou] mobile app plays an animation of radiating bubbles of changing colors while listening to user’s voice input.	25
2.11	Illustration of Role 11: In the game HayDay [Hay], each time the player wins a star, a animated star jumps from the play scene to the the stars’ counter on the top and a blinking star indicates the increase of the score.	26
2.12	Illustration of Role 12: An animation is used to simulate a curtain effect when sliding a window to reveal underneath content on a mobile phone [Eke17].	27
2.13	Illustration of Role 13: A cartoon-style animation of progress bar to entertain online visitors while the website is loading.	27
2.14	Illustration of Role 14: Ephemeral adaptation progressively reveals content in menus based on history–adapted from [Fin+09].	28
2.15	Illustration of Role 15: A 3D tour in Urbane [Fer+15] to understanding the impact of different building designs.	29
2.16	Illustration of Role 16: Ripple animation effect in the calculator application on Android.	29
2.17	Illustration of Role 17: Animated bubble chart on Gapminder website representing the change of the number of babies per woman and life expectancy over years. Viewers can select specific countries and choose to display trails during animation in order to be able to compare the evolution of these countries. [Gapa]	31
2.18	Illustration of Role 18: To convey randomness, Han et al. [Han+12] used animation to randomly change the pattern of dark icons (representing people suffering from cancer) every 2 seconds	31
2.19	Illustration of Role 19: Polly is a walking prism whose motion conveys a mood [Pol]	32
2.20	Illustration of Role 20: In Mobilistes [Sch+06], different categories of transportation means are animated in different manners. (a) Airplanes fly over the list from the right . (b) Buses get compressed, make a reverse gear to the left and then got inserted in their new position. (c) Cars get compacted, then make a half turn to the left towards their new position.	32
2.21	Illustration of Role 21: Four-stage animation in “Fallen of World War II’ featuring the global evolution between 1940 and 2010 of <i>colonial</i> wars first (purple bars), then <i>interstate</i> wars (green bars) followed by <i>civil</i> wars (yellow bars) and finally <i>civil wars with foreign intervention</i> (orange bars). An additional stage allows a presentation change: a zoom operation for better overview [Fal].	34
2.22	Illustration of Role 22: In this data-story video, animation is used to highlight the content of focus during the narration. [Wea]	34
2.23	Illustration of Role 23: Animation shows how flu spreads in a population of vaccinated (blue) and vulnerable (yellow) children– Adapted from The Guardian [Mea]	35
2.24	The number of participants in our preliminary exploration that mentioned examples of animation illustrating each role in our taxonomy of 23 animation roles.	37
2.25	The number of practionners in our validation interviews that mentioned examples of animation illustrating each role in our taxonomy of 23 animation roles.	38
3.1	The Reference Model pattern, adapted from [HA06]	45

3.2	Our simplified model of information visualization based on the Reference Model pattern [HA06].	45
3.3	Components of a visual representation of data	46
3.4	An excerpt of an imaginary dataset comparing elevation and prices of three categories of accommodation.	46
3.5	A sketch of visualization representing the maximum elevation of three categories of accommodation: house, apartment building and skyscraper.	47
3.6	Overview of the data and visual components based on our simplified model.	47
4.1	A scatterplot-like visualization where both the position, size and color of dots can evolve over time.	58
5.1	User interface of our study, during a DS-SL task. Here, A-B and C-D are pairs of SL-similar objects, whereas A-D and B-C are pairs of DS-similar objects, forming the two concurrent possible groupings.	82
5.2	Summary of the mean grouping strength for all visual variables examined in our study. Error bars are 95% CIs. The leftmost column shows the mean grouping strength per variable, for all comparisons it was involved in. The fact that the CIs of DP and DS fall to the right of the plots shows that these dynamic visual variables have stronger grouping strength than any other tested variables. Each cell in the tabular view shows the mean grouping strength of the corresponding visual variable in row (in orange) vs. that of the visual variable in column (in blue), for this task only.	85
5.3	Mean grouping strength for the dynamic variables (DP, DS, DL) in the context of a conflicting static grouping. Error bars are 95% CIs. That all CIs fall to the right of the halfway mark (0.5) shows that these dynamic variables have stronger grouping strength than static ones.	86
6.1	Example of staging strategies animating the evolution of a population between two years. a) Elements are grouped and animated by gender: male first, then female. b) Elements are grouped and animated by behavior: increasing income first, followed by stable, then decreasing.	93
6.2	The design space of <i>animation choreography</i> . A design is informed by fixed dimensions, which include characteristics of the dataset, changes in the data or (re)presentation, and the task of the observer (end-user) of the visualization. The designer of the staging can vary the variable dimensions, including setting a strategy for staging which informs the decomposition and timing of the staging design.	97
6.3	Staged animation from an initial visualization showing the emergence of antibiotics in the US over years, to a final visualization presenting a list of all antibiotics, showing the original drug (blue square) and its derivatives (gray squares). (1) labels disappear and several blue squares (antibiotics) change to gray (derivatives). (2-3) squares shrink, disperse and organize into groups, aligned horizontally; (4-5) labels and legend appear while squares continue to move to form groups.	102
6.4	The three scheduling alternatives, based on Allen's seven primitive time relationships between two stages [All84].	104

- 6.5 Comparison of a full (top) vs. staged (bottom) animation of Gapminder data (x-axis: fertility rate, y-axis: life expectancy, size: population, color: region). In full animation, all data points evolve simultaneously through their n states. Each frame of the animation shows all countries at the same year of their evolution. Staged animation decomposes by region across n states. The first and last frames of the animation correspond to the initial (year 1800) and final data state (year 2016). Each other frame shows data points at different time of their evolution. For instance, the frame at the second position shows the American countries (in yellow) during their evolution from year 1800 to 2016, whereas all other countries are still displayed with their state in 1800. Subsequent frames show the animation for the rest of the regions. 107
- 6.6 Staged animation of Gapminder data. As in Fig. 6.5, the staging decomposes the animation by region, one at a time. As the status panel at the bottom shows, the “America” and “East Asia & Pacific” regions have already been animated, the “Europe & Central Asia” is being animated and “South Asia” “Sub-Saharan Africa” remain to be animated. 109
- 6.7 Six-stage animation showing the evolution of characters and their relations between 2 episodes (initial) and (6) of a TV-show [Got]: (1) edges disappear, (2) cluster hulls disappear, (3) nodes disappear, (4) remaining nodes (characters common to the two episodes) are repositioned to leave room for appearing nodes (new characters), (5) cluster hulls appear, and finally (6) edges appear to form the final network. 109
- 6.8 Example of a 3-level deep staging hierarchy, decomposing a full animation based on the criteria C_1 (salary value), C_2 (salary change), and C_3 (job sector value). The leaf nodes correspond to the resulting stages S_i , to be animated in order (depth-first traversal). 114
- 6.9 Staging hierarchy for the co-authorship network scenario: (a) one-level staging based on the presence change of authors (Data Attribute Change);(b) one-level staging based on the topic (Data Attribute Value); (c) two-level staging based on the topic, then presence change of authors (Data Attribute Value and Data Attribute Change). 115
- 6.10 Gapminder data displayed in our prototype. The chosen staging decomposes the animation by region, one at a time. Here, the “America” and “East Asia & Pacific” regions have already been animated, the “Europe & Central Asia” region is being animated and the “South Asia” and “Sub-Saharan Africa” remain to be animated. 119
- 6.11 Partial listing for a Gapminder animation staged with two levels, the first one depending on the country region and the second one on the initial fertility rate. 120

List of Tables

2.1	Taxonomy of roles by category. A '✓' mark indicates a role that is present in Baecker and Small's taxonomy [BS90]. A ⚡ mark indicates a role with notable research opportunities.	15
2.2	Classification of animation drawbacks	42
5.1	Summary of the different tasks generated from the comparison of our six visual variables.	79
6.1	Classification of literature based on fixed and variable dimensions of animation choreography	105
6.2	Classification of data narrative examples from work by McKenna et al. [McK+17] based on fixed and variable dimensions of animation choreography	105
6.3	Classification of staged animations collected from data journalism, oral presentation applications and web advertisements	106

Chapter 1

Introduction

*Animation is not the art of DRAWINGS that-move but
the art of MOVEMENTS that-are-drawn.*

— Norman McLaren (1968)

As illustrated in McLaren’s quote above, animation in its broad definition involves all the various ways that have been used to convey different forms of movement. Humans invented many centuries ago various techniques to convey movement and dynamism which have evolved over the years as technology has progressed. A reading through history reveals that our ancestors have started depicting dynamic phenomena by drawing motion as early as they have acquired the capacity to draw. Even more fascinating, many archaeological studies unveil that our ancestors used cave paintings not only to represent movements but also narratives [Azé11; CA05c]. A notable example is the “Grand Panneau” painting at the Chauvet Cave in France that was deeply studied by the two French investigators Marc Azéma and Florent Rivère (see Figure 1.1). Azéma and Rivère claimed that the piece represents a hunting story composed of two main events running from left to right along the decorated wall. The first scene (see left rectangle in the Figure) shows several lions, ears back and heads lowered, stalking a prey. Mammoths and other animals appear nearby. The second scene of the painting (see right rectangle in the Figure) depicts a pride of 16 lions, some drawn smaller than the rest to appear farther away, lunge toward fleeing bison [Sto; AR12].



Figure 1.1: One of the Chauvet Cave paintings illustrating a scene of hunting. Azéma and Rivère claimed that flickering torches passed over these painted scenes would have heightened onlookers’ sense of seeing live-action stories thanks to the optical phenomenon known as the persistence of vision [Sto; AR12].

In addition to representing narratives and telling stories by sequencing events, Azéma and Rivère found that ancient Europeans were using cartoon-like techniques to give observers

the impression that animals were moving across cave walls [Sto; AR12]. Thousands years ago, ancient artists created graphic stories in caves and illusions of movement and action using various techniques of “Stone Age Animation” [Sto]. Azéma and Rivère [AR12] claimed that the main Stone Age animation technique consisted of breaking down movement based on two different processes: the superimposition of successive images, and the juxtaposition of successive images. The two investigators highlighted that “*by these two procedures, prehistoric men foreshadowed one of the fundamental characteristics of visual perception, retinal persistence*”. Stone Age artists primitively applied the technique of “split-action movement” by superimposing two or several successive images of animals at different stances to create the illusion of running (Figure 1.2), head tossing (Figure 1.3) or trail shaking (Figure 1.4). The famous “Vitruvian Man” (Figure 1.5) drawing by Leonardo da Vinci represents a more recent implementation of this ancient superimposition process. Nowadays, cartoon animation artists use a contemporary version of this process known as the “onion skinning” technique (Figure 1.6).




Figure 1.2: Ancient artists at Chauvet Cave superimposed drawings of two stances of a bison to depict the movement of running [Prea].



Figure 1.3: A painting discovered at Lascaux Cave where three images of horse's head are superimposed to depict the movement of head tossing [Prea].



Figure 1.4: A drawing found at the Cave of the Trois-Frères where ancient artists drew two superimposed images of an animal trail to depict the movement of trail's shaking [Prea].




Figure 1.5: The Vitruvian Man: a drawing by Leonardo da Vinci done around 1490 depicting a man in two superimposed positions with his arms and legs apart and inscribed in a circle and square [Vit].




Figure 1.6: The onion skinning tool in Autodesk Maya [Oni].

In addition to superimposition, ancient Europeans used images' juxtaposition to represent movement and add life to their paintings. In this process, positions taken up by the animal successively in a given time period are juxtaposed one after another and turned in the same direction. The most notable illustration of this Stone Age animation, is a drawing found in

the La Vache Cave, showing a lion running from right to left (see Figure 1.7). This primitive principle has evolved over the years and is known nowadays as the freeze-frame technique.




Figure 1.7: A freeze-frame drawing representing a lion running from left to right discovered at the La Vache Cave [AR12].

Azéma and Rivère [AR12] suggested that Stone Age Europeans invented also a kind of optical animation toy. Archaeologists discovered in several sites in France and Spain stone and bone disks, typically with center holes, showing opposing images of sitting and standing animals. Rivère has reproduced these engraved disks and looped strands of animal tendon through the center holes. By twisting these strands, the disks rotate back and forth rapidly enough to make animals appear to be sitting down and standing up. Azéma and Rivère believe that the flickering images engraved on ancient discs represent “Palaeolithic thaumatropes”. A thaumatrope is an optical toy that was popular in the 19th century. It consists of a disk with a picture on each side that is attached to two pieces of string. When the strings are twirled quickly between the fingers the two pictures appear to blend into one due to the persistence of vision. The most common thaumatrope pictures include a bird on one side and a cage on the other, as well as a bouquet of flowers opposed to a vase (Figure 1.9). Azéma and Rivère considered that Palaeolithic thaumatropes “*can be claimed as the earliest of the attempts to represent movement that culminated in the invention of the cinematic camera*”.



Figure 1.8: A reconstruction of the The Laugerie-Basse spinning disk by Rivère shows a chamois in different positions on each side: standing on one side of the disc (left image), and lying on the other (right image) [AR12]. When the disk is twirled on a string, the creature appears to move (see animation in [Dis]).

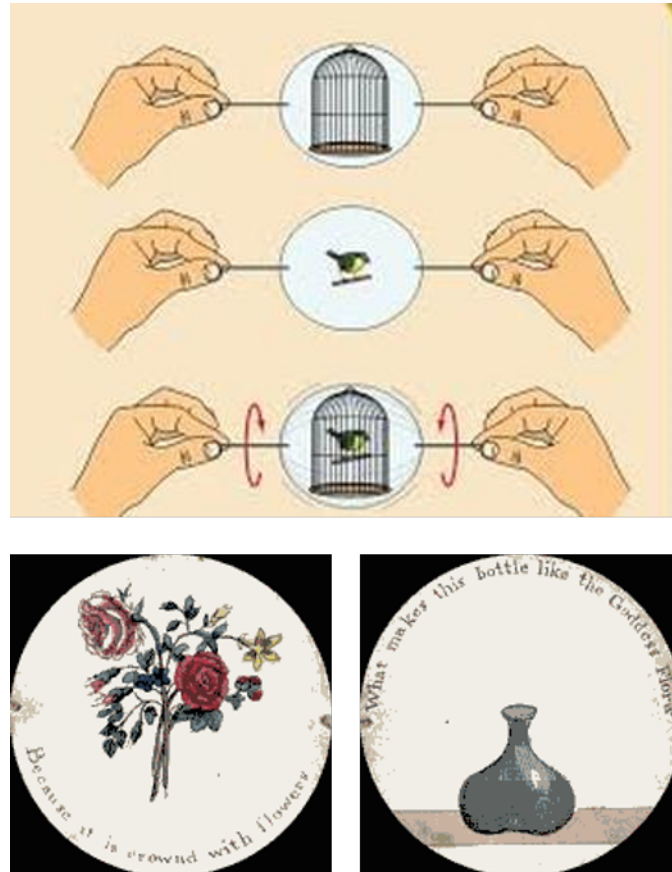


Figure 1.9: Top: a demonstration of a bird and cage thaumatrope [Thaa].
 Bottom: Both sides of a bouquet and vase thaumatrope from 1825 [Thab].

Over the years, humans have continued to explore various techniques to depict motion ranging from simply manipulating hands and using minor mechanics in the shadow play [Sha] and magic lantern [Mag] to introducing more advanced methods boosted by the rise and spread of the Industrial Revolution in Europe and North America. Multiple techniques aimed at generating the illusion that images are moving have been invented, among which the Zoetrope used to show progressive phases of motion based on the principle of stereoscopy, as well as the preeminent Marey's chronophotographic gun made by Étienne-Jules Marey in 1882 (Figure 1.10). This instrument was capable of taking 12 consecutive frames a second, with all the frames recorded on the same picture designated by "animated photograph". Marey's instrument was among the most pioneer works that contributed to the emergence of the cinematography field.




Figure 1.10: (a) Mechanism of the chronophotographic gun, (b) Marey used the gun to study the movement of birds, (c) an “animated photograph” recording 12 movement’s phases of a group of birds [Mara; Marb].

At the beginning of the 20th century, cartoon animation emerged and was produced as animated films in studios using mainly the technique of stop-motion photography applied to hand drawn images to create visual actions. The most notable cartoons of this early period include *Fantasmagorie* (1908), *Gertie the dinosaur* (1914), *Felix the cat* (1919) and *Mickey Mouse* (1928) that shaped the success of the Walt Disney studio. Walt Disney along with many other studios like Warner Brothers, MGM, and Fleischer have achieved the “Golden age” of American animation that has later been widely developed and expanded in both cinema and television during many decades [His].

Around the 1980s, traditional animation has evolved towards computer-generated animation thanks to the tremendous progress of computer-generated imagery. Hand drawing has been replaced by 3D modeling to generate 3D computer animation that got major inspiration from the Disney’s twelve basic principles of animation that were introduced by the Disney animators Ollie Johnston and Frank Thomas in their book *The Illusion of Life: Disney Animation* [JT81] and re-adapted by Lasseter [Las87] to be applied to 3D computer animation. Among these twelve principles we can cite, for example, the *Slow In and Out* effect (alternately named as *Ease In and Out*)—that dictates to slow down the speed of animation near its beginning and its end and accelerate it in between to help essentially with apprehension, and the *Exaggeration* effect—consisting mainly of applying extreme alterations on the physical features (e.g. movement, size, shape, color, etc.) of the the animated visual items to emphasize the essence of realism.

3D computer animation has been progressively expanded to other domains, beyond cartoon animation, to be adopted for example in the domain of computer simulation to replicate natural and physical phenomena [SKP08; MCG03], and in virtual reality applications to mimic realistic behaviors of Humans and physical components of a virtual scene [MTFCY06; KC02]. Researchers in Human-Computer Interaction have started very early to explore how animation can be used in interactive systems for the main purpose of improving the usability of the graphical user interfaces especially in terms of effectiveness and user satisfaction [Sta93].

In the context of graphical user interfaces, an animation consists of a sequence of images that, presented to the observer in a sufficiently fast pace, gives the illusion of motion. Animation usually involves a series of graphic transformations (e.g. movement, color change, resize, change of shape etc.) on a number of graphical objects across successive states of an interactive system evolving over time. The animated transition is a particular class of animation which consists of making the abrupt visual change between *two* successive states more fluid and continuous, usually by interpolating between these two states. Animation involves a set of parameters including the animation trajectory, start time, duration, and pacing. The trajectory defines how an animated visual component moves and the visual aspect that

it takes in terms of shape, size, color and so forth, during the animation. Duration is the length of time that an animation takes from the start time, and is generally measured in seconds. The pacing determines how the interpolated values are distributed over the duration of animation—i.e. values spacing. The less intermediate values are spaced, the slower is the animation and vice versa. For example, linear pacing makes the intermediate values spread out evenly over the duration and generates a constant-speed animation. The same animation with a “SlowIn” (“EaseIn”) pacing would begin slowly, and speeds up as the time progresses.

As I am writing these lines, animation represents an omnipresent component of user interfaces. It is used in all modern operating systems and platforms and can be seen on different devices ranging from mobile, desktop, tabletop to wall-sized displays. Animation is developed to enrich multiple software applications including games, websites, interactive applications, and information visualizations. Users frequently come across and manipulate multiple types of animation. They get for example an animated visual cue when turning on their phones to inform them how to unlock the screen, and a bouncing animated icon notifies them when receiving a new message. When using their computers, animated transitions allow them to smoothly navigate between windows, and they may use an animated tutorial to learn how to use an new application. Users also commonly use animation in their oral presentations to engage the audience or to better explain their ideas.

The wide range of application domains in which animation can be used and the various purposes that it can achieve has certainly expanded its usefulness and multiplied its benefits for users. However, applying animation in various contexts of use (i.e. different domains of application, user profiles, devices, etc.) has in return raised many issues highlighting principally that animation can not always be effective and achieve the same purposes in the same way. This high disparity in the value of animation in user interfaces has consequently led to discrepant opinions between those endorsing the use of animation and those dissuading from using it. This controversy leads us to question: why, despite all the decades that researchers and practitioners have spent designing, studying and using animations, we still face all these contradictions about its value? What makes animation still challenging to design and evaluate?

To answer these questions, let us put ourselves in an animation designer’s shoes and follow all the stages of creating an animation in a user interface. That would help us decipher what are the questions that an animation designer is facing to come up with an effective animation. When a designer thinks about using an animation in her graphical interface, she often starts by asking this main question: **“Q1: Should I use or not use animation? and why?”**.

Although animation has been studied and used since many years, we do not have yet definitive or decisive answers to these questions and we still observe many controversies and contradictory results about the merits of animation. For example, while some previous works proved that animation can facilitate learning visualizations [Pla+12], it is considered in other studies as not effective for learning [Tve+07] and for transferring knowledge [HDS02] compared to other ways of teaching. It is also seen as not helpful for some types of visualizations such as process and algorithm visualization [TMB02a; HDS02], which is paradoxical to Stasko [SK96] and Beacker’s [BS90] findings of the possible benefits of using animation in these kinds of visualizations. However, these controversies can reasonably be justified when we take into consideration the difference in their particular contexts. As for any other interface’s usability feature, optimizing an animation to satisfy a subset of usability criteria for a given context of use does not imply that it will remain effective when considering different usability criteria or changing the context. Since our broad question (Q1) concerns the general possible benefits and drawbacks that animation can have in a user interface, we will be able to find only some partial elements of answers by compiling the findings issued from various studies dealing with animation.

But, what if we try now to narrow the scope of our question and contextualize it better? In practice, an animation designer often works with animation in a particular context of use, that

can be leveraged to inform the relevance of design and use of animation, and suggest many opportunities for applying it, as it can also impose certain constraints on that design and use. Both of these opportunities and constraints should be adjusted in accordance to the goal that an animation is intended to achieve within a given user interface. Hence, the second question that an animation designer should ask: “**Q2: For what purpose can I use animation?**”.

Prior work in literature has attempted to answer this question by providing different views of animation roles in order to guide designers through a better understanding of the different goals for which animation could be used in user interfaces. Notably, Baecker and Small’s taxonomy of animation [BS90] shed light, twenty eight years ago, on the different roles that animation could play in user interfaces. To the best of our knowledge, Baecker and Small’s taxonomy remains the most recent general classification of the roles of animations in user interfaces. Providing such a structured overview of animation roles in user interfaces not only helps the designer deciding the most relevant purposes that animation can achieve in her particular context of use, but it can also enlarge her knowledge about the general goals that animation can achieve in graphical user interfaces as well as inspire alternative uses of animation for future needs. This thesis contributes a revision of Baecker and Small’s taxonomy of animation roles in light of 28 years of animation design and use.

Once the designer sets the goal that animation is aimed to achieve, the next logical question that arises is: “**Q3: How I can achieve my goal using animation?**” Being successful in achieving the targeted goal using an animation depends essentially on the way how this animation is designed. Let us break down this process: to design an animation, the designer needs to i) understand the different animation parameters and ii) decide how to define them. Most of prior works in literature deal with animation as a whole where all parameters are defined in accordance to a specific context and a particular purpose. Only a few works focus on studying particular animation parameters, yet still evaluate such parameters in specific contexts of use and do not provide a holistic knowledge about their design. Examples include Dragicevic et al.’s [Dra+11] study of the pacing functions, or Du et al.’s [Du+15] study of position trajectories. Both works focus their study on the effectiveness of animation for visual tracking while other aspects such as the value of animation as a cognitive aid or for visual appeal are also important to consider.

Hence, the animation parameters remain under-explored and still not deeply understood. Moreover, to decide how to set these parameters, the designer needs to refer back to a set of design guidelines and practices. Although we can find in literature many such guidelines, they are either overly general—making it not so trivial to transpose them when deciding one’s own animation parameters (e.g. what does the apprehension principle imply for animation trajectory?), –or inversely overly specific to a particular use case (e.g. animation duration in many research works is defined according to the characteristics of a particular interface and/or use). Furthermore, many of the general design guidelines are generated from empirical studies where animation is evaluated as a whole. That makes these guidelines questionable because we can not easily isolate the impact of each parameter and how much it contributes to the results of such evaluations. To overcome these limitations that challenge both the understanding and the definition of animation parameters, the designer needs first to understand more in depth animation parameters and get acquainted by design guidelines that cover the various animation’s goals she wants to fulfil and that are more adapted to the different animation parameters, ideally derived from empirical studies studying each of these parameters in isolation.

Let us suppose that the designer designed the animation in such a way that it should effectively achieve the targeted goal. She needs now to answer the following question: “**Q4: How can I assess that my goal has been achieved?**”. There are different metrics that allow to measure the usefulness and effectiveness of animation. Animation is a complex visual composition and as such it has different facets and can be seen through different lenses

and measured using different metrics. To assess the value of animation in a thorough and informative way, the designer needs to study it from different perspectives. However, the evaluation of animation in literature has been concentrated around a subset of evaluation metrics such as those related to multiple objects tracking. Consequently, many metrics remain still under-used. Examples include, for instance, the metrics allowing to measure the semantic value of animation and its impact on the understanding and interpretation of information, as well as those allowing to assess its aesthetic value. We argue that applying a limited range of evaluation metrics does not provide a sufficiently informative knowledge about the value of animation. The three animation aphorisms by Tversky illustrate very well our claim: “Seeing is not perceiving, Perceiving is not understanding, Showing is not explaining” [Tve+07].

1.1 Thesis statement and goals

In this thesis, we argue that if we want to move forward towards better designed and more effective animation, we first need to step back in order to analyze, synthesize and structure the knowledge that prior research has provided about the merits and downsides of animation for user interfaces as well as its design and evaluation. We argue that if we do not proceed regularly to realize this work of revising and restructuring our understanding of animation, we will keep on creating isolated instances of animation and our vision of animation will remain narrowed by our exclusive contexts of application. Hence, we will pass over multiple lessons that can be learned from each others experience. We believe that the very noticeable lack in literature in such works of synthesis and reflection about animation represents one among the principal reasons of the multiple unresolved challenges of the design and evaluation of animation. In this perspective, this thesis aims at providing a structured understanding of the use, design and evaluation of animation. Such structured understanding would provide a formal and common ground that would facilitate carrying purposeful and constructive discussions in the future. It would also help spotting gaps in the current state of the art exploring animation and paves the path to relevant future research opportunities.

After providing a structured overview of the general benefits and drawbacks of animation in user interfaces. We explored its design and evaluation methods in the context of a particular application domain, namely, information visualization. In this context, we studied more in depth the design and evaluation of particular animation parameters (i.e. trajectories and staging). Information visualization domain represents a rich use case that allows us to explore different facets of animation including the functional (e.g. usability), semantic (e.g. expressivity), and aesthetic facets. It also represents a relevant demonstration of the possible harmful impacts that a poorly designed animation can have on the readability, comprehensibility and interpretation of the animated content. Such destructive effects of animation can considerably reduce the value of animated visualizations, violate their meaning and impair the process of exploration and sense making. We do consider that using animation in the field of information visualization has to be studied with much more care and depth. In this perspective, among the multiple parameters of animation, we chose to focus our exploration on the trajectory and staging parameters because we believe that they incarnate a great potential in terms of the readability and expressivity of visualization that have not yet been explored and hence represent relevant research perspectives.

1.2 Research contributions

This thesis makes the following contributions to the fields of Human-Computer Interaction and Information Visualization:

1. A structured view of the animation goals in user interfaces through a taxonomy of 23 roles, organized in 5 categories, illustrated by examples borrowed from academic sources and our observations in the wild. This taxonomy additionally highlights both the positive and negative facets of each category of roles, and discusses the findings of a qualitative exploration that we conducted with people to investigate their use of animation and their perception of its positive and negative impacts as well as a series of validation surveys that we conducted with practitioners to validate our taxonomy.
2. A structured view of the drawbacks of animation that we constructed based on the findings of our qualitative exploration mentioned above. We suggested a classification of the drawbacks of animation perceived and experienced by the participants of our exploration using a set of factors characterizing both the design and use of animation in user interfaces.
3. A theoretical framework to help animation designers gain a deep understanding of dynamic information visualizations. We proposed a set of dimensions to characterize dynamic changes in visualizations that allow animation designer to answer the main following questions: what makes a visualization change? what are the precise components that are affected by a given dynamic change? and how these components change?
4. A discussion of animation fundamentals and challenges in the context of information visualization. We start by formalizing the basics of animation for dynamic visualizations. We then provide a structured view of the design and evaluation of animation where we discuss the limitations of the current state of the art exploring the animation parameters and providing design guidelines. We also suggest a list of metrics to measure the value of animation, discuss the general methods of animation evaluation and point out the different challenges that researchers encounter when assessing the effectiveness of animation and the limitations in some approaches they have adopted.
5. A crowdsourced graphical perception study of the meaning of Common Fate for animated transition in visualization. In this experiment, we asked participants to make perceptual judgements on a series of trials involving four graphical objects under the influence of conflicting static and dynamic visual factors used in conjunction. Our results yield the following rankings for visual grouping: motion > (luminance, size, dynamic size); dynamic size > (dynamic luminance, position); and dynamic luminance > size. We discuss the implications of our findings for the design of dynamic information visualizations.
6. A design space that allows a holistic characterization of staged animation across several visual states of a dynamic visualization. We surveyed and characterized literature and practice using our design space, thereby identifying the limitations of existing staging approaches and highlighting novel opportunities. To support the design of staged animation, we provided a paradigm for authoring staging alternatives based on varying the decomposition and timing parameters, and demonstrated our approach with a concrete scenario. We discussed the design considerations of an authoring tool for staged animation and described our prototype.

Chapter 2

Roles and drawbacks of animation in user interfaces

Just as there are good and bad uses of color, so there are good and bad uses of animation. Inappropriately applied, animation will seem childish and drive users away. But sensibly applied, it can make an interface more graceful and enjoyable to use.

— Thomas and Calder [TC01]

In this chapter, we start by focusing on understanding the roles of animation and asking the following preliminary question: “Why and for what purposes animation has been used?” Answering this question will help us gain an overview of the contexts of application in which animation has been used and the goals it has been targeted to achieve. The literature and practical state of the art of animation include hundreds of works that have been realized across many decades. The very early uses of animation in graphical user interfaces involved using animation to animate the change of viewports during navigation [RMC91a] as well as to illustrate algorithms [BM86]. Since then, it has been widely expanded to include for instance directing attention of users to points of interest [MCS01], improving the understanding of dynamic processes [MA92] and time-varying systems [Sta93]. Hence, considering the great number of goals that animation has been used for, if we attempt to investigate our question without any structured guidance, our task proves to be not so trivial to accomplish.

In this perspective, some prior research proposed different taxonomies to characterize the uses of animation in user interfaces in general, while others were interested in the use of animation in more specific contexts such as information visualization. “Animation at the Interface” by Baecker and Small [BS90] is the most pioneer taxonomy and was proposed twenty eight years ago. In this seminal paper, the authors went beyond the early use of animation to illustrate processes and algorithms, and exposed the fundamental idea that animation can help make interfaces more understandable. Their paper described a set of scenarios, mostly hypothetical at the time, illustrating that animation could help users track objects of interest, choose what to do and see how to do it, but also get feedback on what just happened, or see summaries of what happened in the past. Later works provided several elements to characterize animation in the user interfaces. Partial characterization of animation include for example the seven functions of animation suggested by Novick et al. [NRW11]: signal different context, signal different value, signal different status, signal importance or urgency, signal different function, signal referent (pointing), and signal salience. The Novick et al.’s taxonomy provides a set of relevant communicative functions of animation in user interfaces, but it does cover only a fraction of the rich spectrum of its possible functions. Other works focused on characterizing the functions and use of animation in more particular contexts of application including for instance the domain of computer based instructions for which Weiss et al. [WKM02] suggested five functions of animation: cosmetic function, attention gaining function, motivation function, presentation function and clarification function. The two latter functions were particularly explored in depth as ways for enriching the expressivity and readability of dynamic information visualizations. In this context, several works characterized the use of animation for information visualization. One among the most relevant examples is the work by Bartram [Bar97b] listing potential applications of motion in visualizations such as using motion to communicate data relationships or for representing change.

Another example is the taxonomy by Heer and Robertson [HR07] presenting transition types between data graphics in the context of statistical visualizations. They identify seven transitions between data graphics based on operators such as view transformation, filtering or data schema change. Building upon this work, Fisher [Fis10] suggested a taxonomy describing the types of change that can occur in visualizations and distinguishes six types of transitions related to data and view transformations. Both taxonomies focus on a single specific role of animation, namely, support the transitioning between different views or states of a system. Both works list tasks and situations in which animated transitions may be appropriate, but do not explore the other roles animation can play in graphical user interfaces. Besides the taxonomy of transitions, Fisher [Fis10] discusses the context in which the animation are used in visualization, distinguishing between presentation (i.e. the data is well known to the presenter and the viewer is passive), and data exploration where the data is still unknown, but where the viewer has full control over the interactions. While he does not explicitly discuss roles of animation per se, Fisher still stresses that animation can be designed for different purposes (presentation vs. exploration). In summary, there have been many efforts towards classifying and describing how animation are used, especially in data visualization, but we do observe the lack of recent works surveying and providing reflections about the general uses of animation in user interfaces.

In this chapter, we propose a taxonomy that builds upon several of the prior efforts and expands on them by providing several novel roles of animation for graphical user interfaces (GUIs) in general. We enrich our discussion of the animation roles as well as their benefits and drawbacks with findings from our preliminary qualitative exploration of animation as used and seen by people. This exploration allowed us to highlight to which extent people are aware of the various roles of animation, for what purposes they use it and for what reasons they refrain from using it, and how they perceive its value. For each category of roles, we highlight the most relevant caveats of roles and discuss research opportunities. Highlighting the caveats of animation roles reveals its negative facets that range from role-specific drawbacks to more general ones involving mainly temporal and cognitive complexity. We will discuss in the next chapter the drawbacks of animation based on the findings of our preliminary exploration.

2.1 The 23 roles of animation in visual interfaces

In this chapter we revisit the roles of animation in graphical user interfaces, twenty eight years after Baecker and Small [BS90] seminar paper. We started by conducting a preliminary exploration through an online survey and a series of semi-structured interviews to investigate how people use animation and perceive its value. We then reviewed academic publications and commercial systems, and interviewed 20 professionals of various backgrounds. Our insights led to a taxonomy of 23 roles of animation that provides an expanded set of roles played by animation in interfaces today for keeping in context, teaching, improving user experience, data encoding and visual discourse. We illustrate each role with examples from practice and research and discuss most notable findings from our preliminary exploration and the practitioners' interviews. This expanded description of roles aims at gaining a global understanding of roles that animation has been used to accomplish during the last decades and to exhibit the evolution and the expansion of these roles.

We use the broad definition of animation by Betrancourt and Tversky [BT00]: “computer animation refers to any application which generates a series of frames, so that each frame appears as an alteration of the previous one, and where the sequence of frames is determined either by the designer or the user.” We chose this definition as it captures the essence of animation as a technique, while not specifying what the animation is supposed to convey or support. It also embeds both passive and active use.

2.1.1 Preliminary qualitative exploration

We based our preliminary qualitative exploration on an online survey and a series of semi-structured interviews that aimed at investigating and understanding a) why and b) for what purposes people use or discard animation, c) how they use it and d) to which extent they are aware of both its beneficial and harmful aspects. We also asked participants about positive and negative impressions they kept about animation instances that they have seen as part of, for example, presentations, interactive visualizations, mobile apps, websites, and so forth.

The online survey contained 14 questions (including 3 demographic questions) that were structured in three parts. The first part was dedicated for the demographic questions. In the second part, we focused on the use of animation in the context of oral presentations. We chose to focus on this particular use of animation because it arguably represents the most common use of animation for people who are not expert in designing animation and who are using it simply as a tool in their scholarly, academic or more generally professional presentations. We asked participants to give us their motivation and rationales behind using or discarding animation in their presentations and provide some examples of the types of animation they most often apply. The questions asked in this first part are as follows:

- Do you use animation in your oral presentations?
- If no, why? Do you have any particular reasons for not using or discarding animation?
- If yes, why do you use animation? for what purposes?
- Could you cite some types of animation that you often use (e.g. display text paragraph by paragraph, enlarge image, etc.)?

In the third part of the survey, we wanted to explore the eventual positive and negative experience that participants had with animation instances that they have encountered in various contexts (e.g. oral presentations, web sites, mobile apps, etc.). We varied the questions in such a way that both the most known benefits and drawbacks of animation were covered.

The questions asked in this second part are as follows:

- Could you remember an animation that you saw somewhere (in an oral presentation, a web site, mobile app, etc.) that you *liked*? Could you describe it?
- Could you remember an animation that you saw somewhere (in an oral presentation, a web site, mobile app, etc.) that *helped you understand an information or an idea*? Could you describe it?
- Could you remember an animation that you saw somewhere (in an oral presentation, a web site, mobile app, etc.) that *made you have an emotional reaction*? Could you describe it?
- Could you remember an animation that you saw somewhere (in an oral presentation, a web site, mobile app, etc.) that *convinced you about an idea*? Could you describe it?
- Could you remember an animation that you saw somewhere (in an oral presentation, a web site, mobile app, etc.) that *caused you frustration*? Could you describe it?
- Could you remember an animation that you saw somewhere (in an oral presentation, a web site, mobile app, etc.) that *confused you*? Could you describe it?
- Could you remember an animation that you saw somewhere (in an oral presentation, a web site, mobile app, etc.) that *misled you*? Could you describe it?

The two last parts of the survey allowed us to understand how people perceive and judge the value of animation both as designers and as end users.

Our survey was created and hosted on GoogleForms [Gfo] and spread through social media and emails. Overall, 33 participants answered to our survey. 21 were female and 12 were male, and their age ranged from 20 to 35 years. Participants were students (19/33), software engineers (8/33), computer science teachers (3/33) and researchers in HCI and information visualization (2/33). When asked about whether they use or not animation in their oral presentations, most of the participants had a positive answer (78.8%, 26/33).

The semi-structured interviews were comprised of two parts. In the first part, we asked participants about the different contexts of application in which they use animation. Most of them use animation especially in their oral presentations, whereas some others used animation also in interactive applications, information visualizations and websites. Questions were inspired from those of the second part of the online survey and were as follows:

- Do you use animation?
- If no, why? Do you have any particular reasons for not using or discarding animation?
- If yes, in which particular contexts do you use animation? for what purposes?
- Could you cite some types of animation that you often use?

In the second part, we were interested to explore how participants design their own animation in order to know if they founded their choices on design rationales and if they communicated certain meanings through these choices. We asked participants about the factors that they take into consideration when designing their animation (e.g. category of audience, type and content of animated information, duration of oral presentation, etc.) and how they set the different animation parameters such as duration, speed, temporal sequencing and so forth. When possible, participants presented to us examples of animations that they designed and explained the rationales behind the choice of each animation parameter.

Overall, we interviewed 23 participants. 14 were male and 9 were female, and their age ranged from 26 to 42 years. Participants were students (11/23), software engineers (3/23), computer science teachers (3/23), administrative assistants (3/33), researchers in HCI and information visualization (2/23) and game designers (1/23). When asked about whether they use or not animation in their oral presentations, most of the participants had a positive answer (73.91%, 17/23).

We will discuss the findings of our preliminary exploration in three steps: (i) In this chapter, we will start by presenting *why* and *for what purposes* our participants use animation, and the positive impressions that we found they have about it. These insights will be presented at the end of the description of each corresponding role as the answers provided by participants did not cover all the roles of our taxonomy. We will also analyze the coverage of our taxonomy's roles in order to give some preliminary insights into the degree of *popularity/recognition* and *usefulness* of each role according to respondents. (ii) In this same chapter, we will discuss *why* our participants refrain from using animation and what are the negative impressions they have about it. (iii) Finally, we will discuss *how* our respondents design animation in the fifth chapter of this thesis 4.

2.1.2 Taxonomy Methodology

Our methodology to identify the 23 roles of animation in interfaces, consisted of a multi-step process as follows:

Review. We first searched for related taxonomies discussed in the introduction of this chapter. Then, we surveyed animation examples we were most familiar with as researchers in both

HCI and information visualization (e.g. reviewing Bartram [Bar97b] useful and extensive list of visualization examples, up to 1997). We broadened our search with a literature review of animation techniques in HCI and information visualization by searching on the ACM Digital Library and the website *keyvis.org* for papers containing the words “animation” or “animated transition” in the list of authors’ keywords. We also collected interactive examples that comply with Bertancourt and Tversky’s definition [BT00] that we encountered in websites, mobile apps, and other commercial products.

Build and iterate. We derived an initial taxonomy by grouping animation techniques into roles, and discussing among the collaborators on this research how these roles differ and how to group them in higher-level categories. We systematically compared what we found with the scenarios proposed by the original Baecker and Small’s review [BS90] to guarantee an all-inclusive correspondence, and considered additional examples to see how they would fit in the taxonomy. This whole process helped us refine: we removed minor roles (very few instances) and hybrid roles (covering two or more existing roles). We illustrated some of these hybrid roles in the discussion, and better fleshed out the remaining ones.

Validate. We also validated and fine-tuned our final list of roles based on insights gained from interviewing 20 practitioners and researchers of various backgrounds. We conducted semi-structured interviews where we asked these professionals to think of at least five examples of animation use, without further guidance since we did not want to bias their answers. Only once they completed this list, we asked them to familiarize themselves with our taxonomy, and try to identify the roles for each of the examples they provided. We also invited the interviewees to provide feedback on our list and suggest missing roles, if any. We will use the findings of these validation interviews to illustrate the description of roles with some examples provided by the practitioners. Similarly to the preliminary qualitative exploration, the answers of practitioners did not cover all the roles of our taxonomy. We will analyze the coverage of roles to provide some prior indications about the degree of *prevalence* and *visibility* of the different roles in user interfaces, as well as their level of *memorability* for practitioners. Among the practitioners that we interviewed 15 were male and 5 were female with varied expertise in the domain of HCI and information visualization. They were students at the MSc and PhD level (6/20), researchers (5/20), software engineers (5/20), assistant professors (2/20) and UX designers (2/20). We did not collect information about interviewees age.

2.1.3 Taxonomy of animation roles

In this section, we review all the roles, organized in 5 categories (see Table 2.1). For each role we provide a short definition along with representative examples. A subset of the animation we surveyed is available at: <http://animations.recherche.enac.fr/>, a companion website to this work for the reader to experience the animation we refer to.

Table 2.1: Taxonomy of roles by category. A '✓' mark indicates a role that is present in Baecker and Small's taxonomy [BS90]. A ⚡ mark indicates a role with notable research opportunities.

Categories	Roles
KEEPING IN CONTEXT	1. Staying oriented during navigation 2. Supporting tracking during layout changes ✓ 3. Maintaining up to date 4. Replaying history, summarizing ✓ ⚡
TEACHING AID	5. Affordance and preview ✓ ⚡ 6. Training, demonstration by example ✓ 7. Explaining how a system works 8. Illustrating an algorithm ✓ ⚡ 9. Teaching a new representation of information
USER EXPERIENCE	10. Hooking the user ⚡ 11. Keeping the user engaged ⚡ 12. Providing visual comfort and aesthetics ⚡ 13. Making activity and progress visible ✓ 14. Revealing/Hiding Content 15. Providing a virtual tour 16. Providing feedback of input mechanism
DATA ENCODING	17. Revealing data relationships ⚡ 18. Conveying uncertainty or randomness 19. Encoding emotion 20. Encoding object attribute ⚡
VISUAL DISCOURSE	21. Supporting a narrative ⚡ 22. Highlight content 23. Persuading and convincing ⚡

Keeping in Context

One of the most prevalent categories of animation roles is to keep users aware of the system state, as well as keep them oriented in the information space during changes and manipulations. We include four roles in this category:

Role 1: Staying oriented during navigation

Navigating from an initial point to a target point in an information space (by which we mean a location, or a particular view), was traditionally accomplished by suddenly jumping to the end point. Animation is now commonly used to make navigation feel smooth instead of jarring. Indeed, animation was proved to help users preserve their mental map when navigating in the information space [BB99; HFM07; DHC11]. Archambault et al. [APP11] considered that the animation potential of preserving mental map in the context of graph visualizations consists in the fact that it helps users keeping track of nodes and enhances the readability of the dynamic graphs. Bederson et al. [BB99] considered that it resides also in the fact that it allows users to learn spatial position and relationships between data items by helping them maintain object constancy during navigation.

More generally, mental map preservation methods that have been applied in the context of dynamic visualizations consist of attempting to keep the same global structure or shape of the animated visualization and minimize the overall movement of visual items during the animation. That can be achieved through the *rigid motion* that involves four types of motion: translation, rotation, reflection, and glide reflection, and that has been used by several

works dealing with animation in dynamic information visualizations [FE02; EDF08a; HTC09]. Beyond preserving mental map of information, the effectiveness of animation in maintaining users oriented during navigation is more generally due to its potential to help with visual tracking that will be highlighted in the next role.

Let us now discuss some illustrative examples: After typing “Paris, France” in the Google Maps search box, an animation is used to smoothly zoom and pan the map as if flying over the Earth to see the area around Paris. Such type of animation is a cornerstone component of Zoomable User Interfaces (ZUIs) [Bed11]. When the information space is more complex to navigate—as when users need to navigate from one branch of a large organization chart to another—a staged animation breaks down the animation into multiple steps: e.g. SpaceTree [PGB02] 1) shrinks the branches that get in the way of the new layout; 2) re-centers the trimmed tree so that the new branch fits on the screen, and 3) grows the branch out of the new focus point (see a similar example in Figure 2.1).

In the preliminary exploration, participants did not mention any instance of animation that can be classified under the animation role “staying oriented during navigation”. However, in the validation interviews, some practitioners (4/20) enumerated several examples of animation under this role such as the inertia effect in automatic scrolling on mac OS and the animated zooming in Google Maps.




Figure 2.1: A staged animation to support navigation between directories in a 3D tree-map visualization. When a user opens a directory, (a) the directory begins with a fade-out, (b) followed by a short pause, then (c,d) the new current directory is expanded to fill the space formerly occupied by its ancestor. [BCK05]

Role 2: Supporting tracking during layout change

When the layout of a display changes, it is important that users understand what happens to objects of interest. Animation facilitates the perception of objects’ constancy in a dynamically changing interface [Pal99; RMC91a], hence it helps users track the visual objects and more easily recognize which object is which during the animation. Lavoisier’s maxim stating that, in real life, “Nothing is lost, nothing is created, everything is transformed” has been re-appropriated and applied to user interfaces design to improve understandability while navigating and interacting. For example, a window progressively shrinks into an icon in the taskbar after minimization while other taskbar’s elements move to make room for it. Tracking of elements whose aspect greatly varies over time may result in intricate animation, as is the case in Glimpse [DHC11] that smoothly animates between markup source code (e.g. LaTeX or wiki) and the rendered document (see Figure 2.2). With Glimpse, users see plain text morph into formatted text in the specified fonts or colors, and markup instructions for an image insertion morph into the image itself. Robertson et al. [RMC91a] argued that the effectiveness of preserving object constancy can drastically compensate its time consuming effect. Indeed, object constancy allows users to track relations between animated items without thinking about it, which eliminates the need to spend additional time re-assimilating the interface once the layout change is completed.



Figure 2.2: Illustration of Role 2: Details of an animated transition between a PDF article and its corresponding LaTeX source code in Glimpse [DHC11]

When we asked participants in our preliminary exploration about the main purposes for which they use animation, some of them (4/33 in the survey, 5/23 in the interviews) reported using animation to help observers transitioning between two consecutive slides, two different states of a visualization or two different visualizations. In the validation interviews, several practitioners (10/20) enumerated many instances of animation to support tracking such as the animation of window minimization on OS X, the animated transition between an image and its corresponding histogram in Histomages [CDH12], as well as the animation in One-Click-To-OneNote [One] feature allowing to users to understand what happens to their note when they create a new one. An animation shows the paper folding and disappearing in the OneNote logo to convey the location where users should look for it next time.

Role 3: Maintaining up to date

Dynamic systems have a constant activity, and may prompt the user of the changes in real time or sporadically by animating objects on the screen to keep her up to date about situations. Indeed, motion, which is one of the most used variety of animation in graphical interfaces, is very effective at attracting attention and can be easily perceived in peripheral vision [Pal99; BWC03], hence it has been commonly used as an effective way of notifying users to points of interest in the interface.

A simple example is a clock whose minute hand moves to show the current time. Monitoring applications, such as PlanePlotter [Pla] that visualizes air traffic in real time, also use animation to reflect the most up to date streaming data (see Figure 2.3). Animation is also used as a notification mechanism, to keep the user aware of updates from applications that are not currently visible on the screen: for instance, Moticons [BWC03] are moving icons on the periphery of the display that signal new activity of an application, e.g. a mailbox icon that bounces when a new message arrives. Dyck et al. [Dyc+03] describe how games use animation for “calm notification” of events, as, for example, in Warcraft III, notifications of new achievements appear and fade away, while animated red concentric circles and arrows point users to the location of an event taking place. The authors asserted the ease of animation to indicate direction, location and priority in such a manner that enhances its effectiveness to convey messages without interrupting the workflow of players.

In the preliminary exploration, participants did not mention any instance of animation that can be classified under the animation role “maintaining up to date” which seems to make sense since this role is not applicable in the context of oral presentations that represents the major focus of our exploration. On the other hand, in the validation interviews, two practitioners (2/20) mentioned some instances of animation for maintaining up to date among which the notorious “wizz effect” in MSN which makes the messenger window vibrate to catch the attention that a message has been received. Another example is the head of a character displayed on the bottom in the Doom game [Doo].



Figure 2.3: Illustration of Role 3: PlanePlotter [Air] visualizes air traffic in real time.

Role 4: Replaying history, summarizing

Animation is commonly used to replay history of events that occurred in a dynamic system, to reduce change blindness, or to allow the user to go back in time. Diffamation [Che+10] helps a viewer track how a document has evolved over time. It uses smooth transitions to show added text appearing slowly. Similarly, text that was removed progressively shrinks until it disappears fully. The remaining of the text, corresponding to content that did not change between two versions of the document, moves in trajectories that preserve the visual document structure while minimizing the visual flow. Mnemonic Rendering [BDB06] uses animation to “flashback” what happened to pixels while not being visible in the screen, because they were occluded by another window or in a minimized window; or when the viewer was looking away (see Figure 2.4). The history of pixels is recorded for the whole time they remain hidden or unseen. Once in the field of view again, a “fast-forward” animation unfolds this history until pixels’ color reflect the current state.

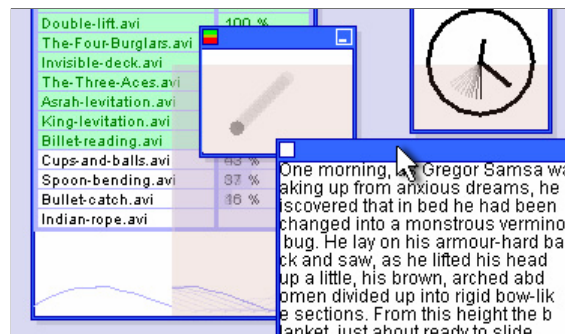


Figure 2.4: Illustration of Role 4: Mnemonic Desktop [BDB06] reveals hidden parts of windows by showing motion trails and a dusty appearance that will fade out as changes are being replayed.

During our preliminary exploration as well as validation interviews, the respondents did not mention any example of animation playing the role of “replaying history, summarizing”.

The caveats of roles and research opportunities of keeping in context category

We introduced four animation roles under this category presented as follows: 1. *staying oriented during navigation*, 2. *supporting tracking during layout changes*, 3. *maintaining up to date*, and 4. *replaying history, summarizing*. In the following, we discuss the main caveats and research opportunities of these roles.

The design of animation can have a drastic impact on its potential for supporting orientation during navigation. The animation designer Val Head argued that mismatched direction and speed of animation are highly disorienting and can even trigger discomfort and sickness for users with particular physiological conditions including for example epilepsy and vestibular

disorder [Hea17]. Relevant examples of such problematic animation include moving background objects at a different speed than foreground objects in a parallax animation, or moving the window view in a different direction than the direction in which the user is scrolling.

The core concept of mental map preservation that has been majorly adopted in literature, consists of keeping the same global structure and minimizing the movement of visual items composing the animated visualization. Based on this concept, one can assume that animation instances that impact a large set of visual items or generate a considerable quantity of movement can probably hinder mental map preservation and may hence prove not to be effective to keep users oriented during navigation. However, minimizing the overall movement during animation is not always an evident and successful design alternative, as highlighted for instance by Saffrey and Purchase [SP08] who demonstrated that restricting node movement can be harmful to diagram readability because it can generate significant nodes' overlap. Deeper exploration is needed in the future to assess the benefits and costs of animation for mental model preservation considering the map preservation methods that have not been yet studied.

The effectiveness of animation for visual tracking depends on multiples factors related both to its design and the characteristics of the interface within which it is used. According to research works on perception, these factors include for example the number of distractors [Pyl04], the visual objects' speed [AF07], occlusions between objects [AF07], and the motion paths of the objects being tracked [SY06]. For example, Dragicevic et al. [Dra+11] studied the impact of the animation's pacing on the visual tracking performance and found that slow in/slow out is the most effective in supporting multiple object tracking compared to other pacing techniques namely, constant speed, fast in/fast out, and adaptive speed. Another relevant study by Chevalier et al. [CDF14] demonstrated that staggered animation has a negligible and sometimes even negative impact on multiple object tracking performance.

Visual tracking of animated items across different states of an interface evolving over time requires a visual comparison between these items. Cottam et al. [CLW12] distinguished two categories of visual comparison: identity based and non-identity based. Identity-based comparison requires to recognize exactly which item is which during and after the animation while the non-identity comparison focuses more on the global distribution of the evolving items. The authors asserted that dynamic changes in an interface make the identity of the changing items highly fragile and hinder their identification. This claim unveils an important limitation in the animation's potential for visual tracking that was highlighted by Heer and Robertson [HR07] who argued that, during animation, object constancy can be abused if an object is transformed into a completely unrelated object, establishing a false relation between the initial and final state of an animated object. Furthermore, the effectiveness of animation for visual tracking depends highly on the quality of the matching (alternately named mapping) approach that is applied to make the conjunction between the visual objects at the beginning and the end of an animation as highlighted for instance by Villegas et al. [VEF15]. Future research in this space will help uncover the importance of all these factors on the effectiveness of animation for visual tracking and provide more informative guidelines on the most effective way to include these guidelines in the design of an animation.

The effectiveness of animation as a signaling and notification medium not only depends on its design characteristics— including for example the type of the animated effect (e.g. blast animation was found to be more effective for notification then scrolling and fading in effects [McC00]) and the speed and amplitude of movement [War+92]— but also on the size of the animated items (e.g. icon, banner, etc.) [BS90], its position and distance with reference to the view field of the user [War+92]. Furthermore, animation for notification can be problematic for various reasons. Some studies proved that animation can attract attention without considerable distraction [War+92], while others argued that it can distract the users from their primary focus and interrupt the task they are performing when receiving

the notification [McC+03]. The distracting effect of animation not only depends on the design of animation but also on the level of engagement and the attentional load of the user tasks [BWC03] (the less engaging is the task, the more distracting can be the animated notification).

More generally, the tradeoff between the power of animated notification in attracting and directing attention and its undesirable distracting effect should be considered with care when applying it within a user interface [MC00]. Furthermore, animated notification can also be annoying and visually ineffective if it is applied on many visual items or triggered repetitively in the interface [BWC03; SW95; Woo95; Gil12]. Conversely, in other studies, animation repetition proved to attract more effectively the visual attention of users towards the animation content. A relevant example is the work by Lee et al. [Lee+15] who found that web animation signaling to the users the existence of ads led to the ad's avoidance behavior when first seen, but after repetitive exposures they induce positive user attitude through the mere exposure effect. They also found that the recognition rate of the animated ads increased continuously with the increase of the number of repetitions.

The number of distinct animation instances used within a user interface can also influence the value of animation for notification. For instance, blinking animation is used in supervisory control systems as a primary visual cue for alarm conditions to attract and direct visual attentions. In this context, studies suggested a limited coding granularity of 4 [Gil12] or 5 [War+92] flashing frequencies. Beyond this limit, users start find blinking excessively annoying and visually ineffective because too many items are flashing simultaneously. The over-flashing not only reduces effective alarm information but also makes the displays visually disturbing, distracting users from effectively perceiving the needed information from other representations [BWC03].

Leveraging animation to replay history and summarize insights is perhaps the most under-explored role in this roles' category. Animation proves rather compelling and promising for collaborative applications used over a long period of time (e.g. in visual analytics to describe insight provenance). For instance, the recent KTGraph work [Zha+18] uses playback animation to help collaborators investigate histories of analysis in the context of an asynchronous collaboration to analyze a collection of documents. However, animation designed for such applications should be defined with care because it can generate mis-interpretation and confusion with reference to the portrayed information. A relevant example is the work of Diffamation by Chevalier et al. [Che+10] where authors refrained from staging the animation of edits between text revisions because sequencing editorial operations may be interpreted as the real chronological sequence of document editions.

Teaching Aid

Animation can be used to foster discovery and learning. Educative animation are widespread, in all domains, and for all ages. We list five roles in this category:

Role 5: Providing affordance and preview

Animation can be used to attract users' attention to widgets or methods that may otherwise be hard to discover or understand. Appropriate graphic design might reveal the affordance of a widget and its functionality, but animation can make it more noticeable and better explain how to use it, or offer a preview of what the tool can do. Indeed, the movement has been recognized in psychology and in the theory of affordance by Gibson as a great medium of communication [Gib14].

Some mobile phones use animation to prompt users to perform a sliding gesture to unlock the device. A website might use an animated arrow at the bottom of the screen to

encourage users to scroll when nothing else encourages them to do so. Google Sketchup's instructor [Ske] (see Figure 2.5) and proxemics interaction [VB04] both provide contextual assistance by featuring a preview of a tool in use or a gesture to perform in a short looping animation. While demonstrating what a tool can do (Sketchup's instructor), or how to select a new tool (proxemics), the contextual aids do not, however, provide guidance on how to use the tool to make it works (see the next role on Training, demonstration by example).

In the preliminary exploration, participants did not mention any instance of animation that can be classified under the animation role “providing affordance and preview” which seems to make sense since this role is not applicable in the context of oral presentations that represents the major focus of our exploration. In the validation interviews, several practitioners (6/20) cited various examples of animation providing affordance among which the famous “Clippy” also known as Microsoft’s “Jar-Jar”, the “pull to refresh” animation on Facebook and Twitter mobile apps, as well as the information panel moving in out of screen edge on smartphones to help users figure out the invocation method (swipe from edge).




Figure 2.5: Illustration of Role 5: When the user selects a tool in Google Sketchup's instructor [Ske], the instructor shows her an animation that explains the basic use of this selected tool.

Role 6: Training, demonstration by example

Animation is commonly used as tutorials of how to complete a task through demonstration by example. For instance, GestureBar [Bra+09] demonstrates the gesture which can be performed to activate a function, for example demonstrating a rubbing gesture when users put their cursor on the eraser icon (see Figure 2.6). Unlike examples in the previous role (i.e. Google Sketchup's instructor and proxemics interaction), the contextual aid of GestureBar consists in demonstrating how to use the tool, rather than giving a preview of what it can do. ToolClips [GF10] adds short narrated demonstrations to the traditional text descriptions of tooltips. Juggling Lab [Jug] allows users to author juggling patterns by specifying a “score sheet”, and watch the resulting animation of a stickman juggling colored balls.




Figure 2.6: Illustration of Role 6: GestureBar [Bra+09] shows an illustrative animation to teach the user the gesture to perform in order to delete a shape.

In the preliminary exploration, participants did not mention any instance of animation accomplishing the role of “training, demonstration by example”. In the validation interviews, practitioners suggested examples of animation for training. For example, some practitioners (5/20) mentioned several example of animation used to train users. Relevant examples include the driving simulator simulating the driving experience through virtual reality animation

and animated tutorials that people use to learn how to sketch. One practitioner reported that animated tutorials to explain how new devices work are among the most useful roles of animation for her (e.g. animation on Surface devices to help users learn left, right, top, and bottom swipes). However, she argued that after seeing the animation once, she would never access the animation again since it becomes useless.

Role 7: Explaining how a system works

Animation is often used in educative applications and media to explain how a system works. Examples of animation that explains various phenomena and complex systems can be found in abundance, e.g. they can explain how the heart works by showing ventricular beats driving and regulating blood flow [Abo17] (see Figure 2.8), or how a mechanical watch works [Wat] to depict how energy is transferred from the barrel to the wheels and to the clock hands. These educative dynamic illustrations are often integrated within a larger communication medium for storytelling (see Visual Discourse roles' category bellow).




Figure 2.7: Illustration of Role 7: Educative animation using bouncing arrows to illustrate blood flow in the heart—adapted from “Exploring the Heart-The Circulatory System!” by AboutKidsHealth (Youtube).

In our preliminary exploratory, some participants (2/33 in the survey, 5/23 in the interviews) reported using animation to simulate particular phenomena, like for example the wind circulation around the wings of the plane or the movement of the planets in the solar system, and also to reproduce particular processes or mimic a certain behavior of a system (e.g. how a computer display gets refreshed, how a haptic hardware prototype works, or how a position's prediction algorithm operates). One participant used the animation in the context of her work as a game designer to simulate scenarios of problems that can occur in a game play and propose solutions.

More generally, many participants in our preliminary exploration (10/33 in the survey, 6/23 in the interviews) reported using animation to better structure and explain the information they are presenting and to make their ideas clearer. They acknowledged that animation helps the audience understand and make sense of the content of their slides. Some among them (4/16) considered animation as an effective pedagogic and didactic medium. One participant insisted on the fact that animation gives more freedom and flexibility to the way she explains information.

From the answers of all participants in the preliminary inquiry, we observed that the effectiveness of animation for helping to understand information depends essentially on the content that it depicts and the effect that it has on the comprehensibility of this content. When we asked participants to cite examples of animation that helped them to understand information, they mentioned for instance: simulations of processes or systems such as an animation that explains the Krebs Cycle or how the STP network protocol works, animated transitions that help making sense of the dynamic changes in the interface like for example animating the evolution of the values of a bar chart over time or the color change of items in

a website, animation that structures information such as an animation that presents step by step the outline of a presentation or an animation that starts by presenting a problem and then progressively reveals steps towards the solution.

In the validation interviews, one practitioner (1/20) cited the example of animation triggered at first use of an application that makes all the visual interface gray except the useful parts and displays explanation bullets on top of these parts [Tan].

Role 8: Illustrating an algorithm

Animation can be used to illustrate a mathematically defined simulation, or an algorithm. One of the pioneer use of animation described by Baecker in 1990 [BS90], was to explain how an algorithm works. Sorting out Sorting [Bae98], shows how elements to be sorted are progressively rearranged by different sorting algorithms. The website Setosa [Set] uses a series of animation to explain the logic of various mathematical concepts. For instance, to explain Markov chains, a ball travels through a graph representation of the chain, from node to node on the chain, conveying how decisions are made to choose which edge to follow based on probabilities associated with the edges.

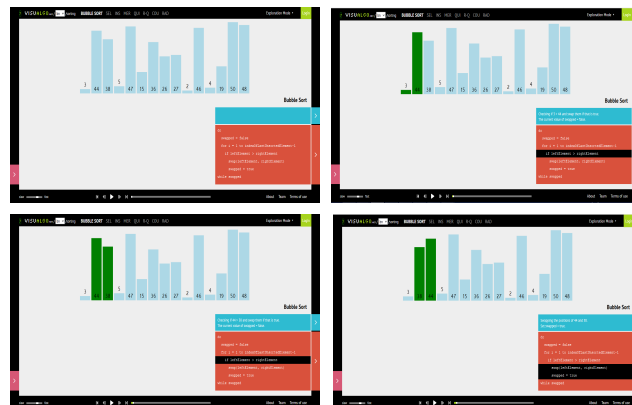


Figure 2.8: Illustration of Role 8: An animation that illustrates how the bubble sorting algorithm works [Vis].

In the preliminary exploration, participants did not mention any instance of animation used to illustrate algorithms. In the validation interviews, one practitioner (1/20) mentioned the series of animation created by Mike Bostock for teaching data structures and algorithms. The animations show how algorithms transform the data step by step [Bos18].

Role 9: Teaching a new representation of information

Animation can be used to help users learn a novel spatial layout or a new visual encoding altogether by demonstrating how items are transformed (grouped, aggregated or divided, change visual properties), and moved to their final location. Animation is used in Twinlist [Pla+12], to help physicians compare two separate lists of medications. Animation is used to teach users how the items in the two lists are reorganized and merged based on similarities and difference, therefore explaining a new unfamiliar layout. Animation is also used in NodeTrix [HFM07] to explain how matrix representations of graphs correspond to the more familiar node-link diagrams (see Figure 2.9). Once learning has taken place, the animation may not be needed anymore and its use discontinued.




Figure 2.9: Illustration of Role 9: An animation in NodeTrix [HFM07] that explains step-by-step the change from a node-link diagram (1) to an adjacency matrix (5)

Respondents in both our preliminary exploration and validation interviews did not mention any instance of animation used to teach a new representation of information.

Caveats of roles and research opportunities of teaching aid category

We introduced five animation roles under this category presented as follows: 5. *affordance and preview*, 6. *training, demonstration by example*, 7. *explaining how a system works*, 8. *illustrating an algorithm* and 9. *teaching a new representation of information*. In the following, we discuss the main caveats and research opportunities of these roles.

Because the outcomes of existing empirical studies are mixed, more research to evaluate the value of animation as a teaching aid is needed [TMB02a]. Understanding when to use the technique, what are the important components and how to streamline the authoring of effective animation are still open questions. Affordances and teaching by demonstration remain important roles as the HCI field strives to make interfaces more “natural” [Lee+15]. When menus and buttons with labels are not present, discoverability of what is possible becomes an issue for new users [Coc+15]. The affordance power of animation resides mainly on the information that it conveys through a set of graphical transformations. These transformations should be carefully defined to prevent the risk of false interpretations of affordance [Vau97].

The fact that animation is perceptually salient makes its use sometimes problematic in the context of teaching. Indeed, Lowe [Low99] explored how students extract information from an animated content during a complex visual learning and revealed that much of the information extracted was perceptually salient rather than thematically relevant which makes the effectiveness of animation as a learning aid questionable. Previous research works have also demonstrated that animation does not help to recall the conveyed information in a learning context [FS96].

Research exploring how animation can be complemented by other modalities, such as narration [MA91] or interactivity [Low03] for education is also a promising direction. Evaluating the value of interactive animated material such as Introduction to Machine Learning [M11] and research on authoring tools to facilitate the generation of such material are other opportunities for the field.

Animation used to explain how a system works or to illustrate an algorithm, if not carefully designed, can lead to misunderstanding and false interpretations of the conveyed information. For example, the speed of animation may give a false indication about the speed of a system or the value of the run-time of an algorithm. This issue was highlighted by one participant in our preliminary exploration who uses animation to explain how a network of sensors works. She insisted that she chooses very carefully the speed of animated arrows illustrating communication flow between sensors not to induce false interpretations about the real speed of the network communication.

User Experience

Animation is often described as visually pleasant, playful and engaging, all qualities that improve the user’s experience when interacting with a system. Animation is now considered a core component of UI design. In Google Android’s Material Design principles [Mat] for instance, “material motion” is defined to “describe spatial relationships, functionality, and intention with beauty and fluidity”. The Windows UX design guidelines [Uwp] state that “Purposeful, well-designed motion brings your app to life and makes the experience feel crafted and polished [...] The quality of the experience depends on how well the app responds to the user, and what kind of personality the UI communicates.”. We find five roles for animation as a tool to enhance user experience:

Role 10: Hooking the user

Animated introductions can make a website or application attractive and memorable, possibly “hooking” users and encouraging them to explore further. Websites often use animation for fun, to appeal or intrigue users, resulting in an amusing, playful, enjoyable experience raising their interest. For example, the Christmas Express website [MM17] starts with a man pedaling on a bike, then opens a full landscape of turning wheels and falling snowflakes and music. The mobile application Soundhound [Sou] displays radiating bubbles of changing colors while waiting for users to start recording sound (see Figure2.10).




Figure 2.10: Illustration of Role 10: Soundhound [Sou] mobile app plays an animation of radiating bubbles of changing colors while listening to user’s voice input.

In our preliminary exploration, when analyzing the answers for the question prompting participants to cite examples of animation that they liked, we observed that their appreciation depends on various factors including the design of the animation (e.g. energetic, well structured, attractive, coherent transitions), the content and goal of the animation (e.g. an animation showing an exploded view of an iPhone, animation to focus on relevant parts, animation that shows progressively the information), and the impact of the animation on their understanding of the animated information (e.g. “an animation that made it easy for me to understand the presentation topic”, informative animation). Overall, 14 participants (10/33 in the survey, 4/23 in the interviews) cited examples of animation that play the role of hooking the user.

In the validation interviews, some practitioners (6/20) mentioned some examples of animation aimed at hooking the user such as the icon “Settings” which turns with the drop down menu in Android.

Role 11: Keeping the user engaged

Because animation attracts a viewer’s attention, it can be used to keep users engaged (or get their attention back) when an activity requires a substantial attention time. Video games such as HayDay [Hay] (see Figure 2.11) or student testing environments leverage animation to keep users engaged in the game, or as a reward mechanism after a substantial effort or lengthy problem. This role is also important in presentation and report tools such as PowerPoint, which contains animation that authors can insert to retain the attention of the audience, or re-engage it. This role is most useful when content is lengthy and the audience is passive. For example, many data-driven videos [Ami+15] and documentaries contain animated charts in addition to the voice narration, in the hope that viewers will watch the materials till the end.



Figure 2.11: Illustration of Role 11: In the game HayDay [Hay], each time the player wins a star, an animated star jumps from the play scene to the the stars’ counter on the top and a blinking star indicates the increase of the score.

Many participants in our preliminary exploration (9/33 in the survey, 4/23 in the interviews) reported using animation to engage the audience as they asserted that it adds life, rhythm and dynamism to the slides and makes the presentation more attractive, aesthetic and especially less boring. In our validation interview, some practitioners (4/20) mentioned various examples of animation to engage users such as the “funny waiting animations” on websites.

Role 12: Providing visual comfort and aesthetics

Animation allows for smooth transitions, which can be experienced as more visually comfortable than jarring, abrupt transitions. When properly done, animation can also have true aesthetic qualities. Examples abound: icons explode into windows, product cards flip to show more details as if on the back of a card, or pixels fly like particles blown by the wind as found in Histomages [CDH12]. Markus Ekert, a motion designer, used animation to simulate a curtain effect revealing content under the main window of a touch mobile phone [Eke17] (see Figure 2.12). Metaphors of the real physical world are often used, staggered animation, which consist of introducing a small delay in the motion start time of elements, aim to recreate the feel of ballet choreography [CDF14].




Figure 2.12: Illustration of Role 12: An animation is used to simulate a curtain effect when sliding a window to reveal underneath content on a mobile phone [Eke17].

Some participants in our preliminary exploration (4/33 in the survey, 3/23 in the interviews) reported using animation for aesthetic purposes. Examples they cited include transitioning between an interactive application's windows using a fade in/fade out effect and progressive display of bullet lists. In our validation interview, many practitioners (8/20) mentioned few examples of animation used for cosmetic purpose such as animations that generate the effect of deepness using the parallax effect and animated bouncing lists.

Role 13: Making activity visible

Animation is used as a visual feedback to track progress, and a distractor to fight boredom when the system is busy. Since the hourglass cursor introduced in Windows 1.0 (1985) and NeXTSTEP's spinning wait wheel (1989), animated icons have been used in place of the cursor to indicate that the system is busy working (presumably) on an operation. Progress bars show the remaining time, or percentage of completion before the process is completed. The Little Big Room website [Big] (see Figure 2.13), for example, uses fun cartoon-style animation not only to keep online visitors entertained while waiting, but also to set the tone and give a hint of the navigation metaphor on the website. Countdowns are used to prepare the user for an imminent event, e.g. when PhotoBooth [Pho] is about to snap a picture or when Camtasia [Cam] is about to start screen casting.




Figure 2.13: Illustration of Role 13: A cartoon-style animation of progress bar to entertain online visitors while the website is loading.

In the preliminary exploration, participants did not mention any instance of animation that can be classified under the animation role “making activity visible” which seems to make sense since this role is not applicable in the context of oral presentations that represents the major focus of our exploration. In our validation interview, many practitioners (8/20) mentioned some examples of animation for making activity visible such as the “3-2-1” countdown animation in Photo Booth and the loading animations (e.g. the circle spinner, or line loader on almost any loading UI). One practitioner considered that such animations give users valuable feedback while they wait for something to happen, and without them, “chaos would ensue”.

Role 14: Revealing/hiding content

Content may not be visible at all times, but made visible through hovering or clicking, and animation are used to reveal or hide content in a transient manner. Examples include side panels swooshing from the edge of mobile device screens, drop-down menus that expand smoothly on desktop applications or balloon tooltips that appear progressively when the cursor hovers over icons. Ephemeral adaptation [Fin+09] is an advanced technique that relies on prediction to progressively reveal features in an interface: predicted items appear immediately, while remaining items gradually fade in (see Figure 2.14).

In the preliminary exploration, participants did not mention any instance of animation used to reveal or hide content which seems to make sense since this role is not applicable in the context of oral presentations that represents the major focus of our exploration.

In our validation interview, some practitioners (4/20) mentioned some examples of animation to reveal and hide content. Relevant instances include collapsed forms and comments as well as expand and collapse animations accompanying some interactions like opening an email. One practitioner explained that this type of animation implies a physical nature to the item: “It feels like I am opening a letter when I do this. That gives me a little thrill!”.



Figure 2.14: Illustration of Role 14: Ephemeral adaptation progressively reveals content in menus based on history—adapted from [Fin+09].

Role 15: Providing a virtual tour

Animation can help users experience a phenomena from the physical world and experiment with alternatives. When architects work on a building, animation of the virtual world [Fer+15] can provide a tour of the building and show how they are integrated in the landscape (see Figure 2.15). Users shopping for a new car can experience in the driver seat in various situations. While this can be achieved with still pictures, the dynamic experience is more realistic and visceral with animated tours. Similarly, complex molecules can be observed from different angles, and travel through the galaxy simulated (i.e. World Wide Telescope [Tel].)




Figure 2.15: Illustration of Role 15: A 3D tour in Urbane [Fer+15] to understanding the impact of different building designs.

Respondents in both our preliminary exploration and validation interviews did not mention any instance of animation used to provide a virtual tour.

Role 16: Providing feedback of input mechanism

Animation can be tightly coupled with direct user input, to convey a particular input metaphor, the state of the input device, or the state of the device sensors. Performing a slide or swipe gesture on a long document on a tablet triggers the auto scroll of the document, whose speed progressively decreases (simulating friction) until the document stops moving. Parallax effects, often seen in phones, respond to the orientation of the input device, captured by accelerometers. Touch devices sometimes also use a ripple animation to confirm input (see Figure 2.16).




Figure 2.16: Illustration of Role 16: Ripple animation effect in the calculator application on Android.

Three participants in our exploratory interviews (3/23) reported using animation in the user interface of their applications by adding an animated feedback to icons to inform users that their action has been taken into account.

In our validation interviews, many practitioners (13/20) enumerated several instances of animation used to provide feedback of input mechanics such as animated buttons on online forms providing feedback when user clicks on them, the bounce animation to indicate scrolling

boundary, and the “defocusing effect” to indicate an overlay.

Caveats of roles and research opportunities of user experience category

We introduced seven animation roles under this category presented as follows: 10. *hooking the user*, 11. *keeping the user engaged*, 12. *providing visual comfort and aesthetic*, 13. *making activity and progress visible*, 14. *revealing/hiding Content*, 15. *providing a virtual tour* and 16. *providing feedback of input mechanism*. In the following, we discuss the main caveats and research opportunities of these roles.

Heer and Robertson [HR07] warned that the engaging nature of animation can increase interest, but can also be used to make misleading information more attractive or maybe frivolous, hence serving as a form of temporal “chart junk” [TS85]. The attractive and engaging power of animation combined with the selective nature of Human’s visual processing, can be used to cognitively bias the users, narrow their exploration of the animated content and limit their interest only to the most pleasing and attractive parts of the information.

Animation is a complex visual creation that, while it has commonly proved to be a source of aesthetic and visual comfort for most of viewers, it can subjectively be unappreciated by others. Beyond the simple disappreciation and dissatisfaction, animation effects intended originally for aesthetic and visual comfort can have more serious and harmful impacts on some users with particular physiological condition, as is the case for instance for people with vestibular disorders that are seriously affected with special effects of animation such as fake zooms, parallax, or sliding [Hea17].

Data Encoding

Information visualization encodes data attributes with visual variables such as position, color, size or shape. In dynamic visualizations, variables may change over time, reflecting changes of the underlying data attributes. Visual variables are not necessarily static: motion can also be used to encode a data attribute. Animation, as a means to encode data characteristics, can be broken down into four roles as follows:

Role 17: Revealing data relationships

Animation is commonly used to show how data changes and evolves over time, but also over other non-temporal dimensions such as age or income. A famous example is Gapminder [Gapa] (see Figure 2.17). Countries are placed on a scatterplot based on their birth rate and life expectancy statistics, and animation reveals how the world has changed over two centuries. Users can simply play the animation or control the advance step by step and see the circles for each country move over time. Trails can be left behind the data points to cope with memorability issues associated with ephemeral content. Cartographic applications are another example, e.g. to show how rainfall location affects forest growth [Blo+99a], demonstrating the effectiveness of animation in conveying causality relationships between data items [WNB99].




Figure 2.17: Illustration of Role 17: Animated bubble chart on Gapminder website representing the change of the number of babies per woman and life expectancy over years. Viewers can select specific countries and choose to display trails during animation in order to be able to compare the evolution of these countries. [Gapa]

Two participants in our exploratory interviews (2/23) reported using animation to enforce comparison. One participant for example used animation to help comparing the evolution of multiple animated bars in a bar chart. In our validation interviews, two practitioners (2/20) mentioned some examples of animation that help revealing data relationship such as the flipping animation [Fli].

Role 18: Conveying uncertainty or randomness

Animation is used to represent randomness or uncertainty by moving objects on the screen in a random way over the range of possible values. In the work by Han et al. [Han+12], cancer risk was first represented as a fixed array of 100 icons, with a number of them dark (e.g. 9 for a 9% chance), but some users wrongly interpreted the display and believed that they have a 0% or 100% chance of developing cancer (see Figure 2.18). Animation was then used to enhance the representation of randomness by changing the position of the dark icons every 2 seconds. Similarly, animation has been used to indicate uncertainty by wiggling the position of points or boundaries on maps or other visualizations [ESG97].

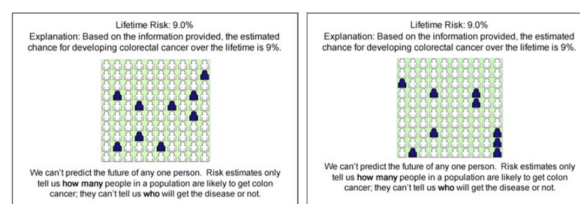


Figure 2.18: Illustration of Role 18: To convey randomness, Han et al. [Han+12] used animation to randomly change the pattern of dark icons (representing people suffering from cancer) every 2 seconds

Respondents in both our preliminary exploration and validation interviews did not mention any instance of animation that can be classified under the animation role “conveying uncertainty or randomness”.

Role 19: Encoding emotions

Particular motions can be used to convey stereotypic emotions. The fact that motion is a good proxy for emotion has been known for long, and extensively employed in traditional cartoon animation: “When you take something that’s inert, and through motion, give it life, make it appear to be alive, living, breathing thinking and having emotions, that’s animation.” [Las87]. Perlin has explored different emotion expressions with Polly, a prism that walks around a surface. Different behaviors, like feeling dejected, can be assigned, which affects the prism’s gait [Pol] (see Figure 2.19). This small application very nicely illustrates how motion can affect perceived emotion. Yet, it is just a prism!



Figure 2.19: Illustration of Role 19: Polly is a walking prism whose motion conveys a mood [Pol]

In our exploratory survey, two participants (2/33) mentioned examples of animation used to express emotions such as “joy” and “sadness”. One of these two participants argued that the potential of animation to express emotion relies a lot on the content it conveys.

In our validation interviews, two practitioners (2/20) mentioned few examples of animation used to express emotions such as the google inbox empty inbox sunrise animation.

Role 20: Encoding object attribute

Motion can be used to encode distinct attributes of visual objects and distinguish different data categories. It can then be used for filtering and brushing [BW02]. When reordering a priority traffic list with Mobilistes [Sch+06], list items are moved using an animated trajectory that is specific to the type of means of transportation (bus, car or plane) (see Figure 2.20). In motion pointing [FEG09], each visual item of the interface such as a push-button or radio-button has a driver, e.g. an animated point displayed on top of it and playing a unique cyclic elliptical movement. To select a specific item, the user has to imitate the motion of its driver using the input device and mimicking that motion (i.e. shape of trajectory, frequency, phase, amplitude, etc.). In the context of cartography, DiBiase [DiB+92] demonstrated the effectiveness of animation to express data in various forms using three types of “dynamic variables”: scene duration, rate of change between scenes, and scene order.




Figure 2.20: Illustration of Role 20: In Mobilistes [Sch+06], different categories of transportation means are animated in different manners. (a) Airplanes fly over the list from the right. (b) Buses get compressed, make a reverse gear to the left and then got inserted in their new position. (c) Cars get compacted, then make a half turn to the left towards their new position.

In addition to the visual variables of static maps, animated maps are composed of three basic design elements or "dynamic variables" –scene duration, rate of change between scenes, and scene order. The dynamic variables can be used to emphasize the location of a phenomenon, emphasize its attributes, or visualize change in its spatial, temporal and attribute dimensions. In combination with static maps, graphs, diagrams, images and sound, animation enhances analysts' ability to express data in a variety of complementary forms.

In our preliminary exploration, participants did not mention any example of animation used to encode object attribute. In our validation interviews, two practitioners (2/20) mentioned examples of animation that encodes data attributes among which an animation to encode acoustic properties in 3D virtual scenes.

Caveats of roles and research opportunities of encode data category

We introduced four animation roles under this category presented as follows: 17. *revealing data relationship*, 18. *conveying uncertainty or randomness*, 19. *encoding emotion* and 20. *encoding object attribute*. In the following, we discuss the main caveats and research opportunities of these roles.

Heer and Roberston [HR07] warned that the animation potential of supporting perception of causality and intentionality should be employed with caution because incorrect interpretations of causality may mislead more than inform. Although animation has proved efficiency in encoding relationships between data items, its ephemeral and transient nature complicates the task of comparing the behavior of visual items being animated in different phases of a global animation [Mun14].

While techniques such as Gapminder are popular, the value of animated displays compared to static one (such as small multiples) is still unclear [Rob+08]. In fact, because of its strong perceptual draw [Bar97b], animation might prevent users from perceiving other elements of the display. However, it also offers unique opportunities to encode complex and subjective information such as convey musical impressions or Human moods and feelings.

The fact that animation encodes emotions can be problematic if the emotions expressed by animation are inappropriate and contradictory to the nature and the content of the animated information (e.g. one of the participants in our exploratory interviews was furious about an example of visualization about diseases that contained joyful animation and considered it as unrespectable). Furthermore, the emotional power of animation may hinder the objectivity of users in their interpretation and judgment of the animated context and can be even abused and converted to a manipulative power.

Visual Discourse

Animation is a common tool for storytelling and visual discourse, hence the abundance of examples. We identified four roles of animation for discourse purpose, as follows:

Role 21: Supporting a narrative

Animation supports narrative by providing progressive disclosure of arguments and transition smoothly between them [SH10]. Pitch interactive recently published a powerful online narrative visualization documenting every drone strike carried out in Pakistan [Dro]. Animation is first used to introduce the topic and progressively disclose the elements of the display. Then, the series of drone attacks periods is interspersed with pauses and commentaries. It includes clickable links where animation is used to transition from current display to the next. Scrolling is also commonly used to activate a transition to a next step of the story like for example in the New York Times' article "How the Recession Reshaped the Economy, in 255

Charts” [Rec]. “Fallen of World War II” [Fal] is an interactive documentary that examines the human casualties in WWII. The animation follows a linear narration, but it allows viewers to pause during key moments to interact with the charts and explore further (see Figure 2.21).




Figure 2.21: Illustration of Role 21: Four-stage animation in “Fallen of World War II” featuring the global evolution between 1940 and 2010 of *colonial* wars first (purple bars), then *interstate* wars (green bars) followed by *civil* wars (yellow bars) and finally *civil wars with foreign intervention* (orange bars). An additional stage allows a presentation change: a zoom operation for better overview [Fal].

Surprisingly, only one participant out of 23 in our exploratory interviews reported using animation to tell a story. In our validation interviews, two practitioners (2/20) mentioned examples of animation that supports narration including for instance the scrollytelling effect—a data-driven storytelling technique that consists of unfolding data stories as the reader scrolls down or up [Sto+16].

Role 22: Highlighting content

Motion has a strong attention-drawing power, and can be used to highlight or emphasize content. For example, an animation where a set of bars in a chart jump when the narrator is discussing them is used in the data video depicting wealth inequality in America [Wea] (see Figure 2.22). Animation featuring icons moving along a line chart can also serve to emphasize an increasing or decreasing trend. Animation of objects or text blinking, jumping or getting larger is also commonly featured in presentation tools such as PowerPoint or Keynote to emphasize a point.



Figure 2.22: Illustration of Role 22: In this data-story video, animation is used to highlight the content of focus during the narration. [Wea]

In our preliminary exploration, many participants (10/33 in the survey, 5/23 in the interviews) used animation to guide the attention of the audience during oral presentations. Most of them insisted on the usefulness of animation to direct the focus of the audience in accordance to the speech. In our validation interviews, some practitioners (4/20) mentioned examples of animation to highlight content such as loop effect in the Apple dock and the fading in/out effect used in video players to indicate focus on the video.

Role 23: Persuading, convincing

Animation can be used to make a particular point across with the intent to persuade or convince its viewer. For instance, the Guardian published a story aimed at convincing people to vaccinating their children by illustrating how disease spreads on hypothetical populations of children whether they are vaccinated or not. An infected person regularly shows up and enters in contact with a child: a red circle flies over and bumps into a child who may or may not contract the disease. A vulnerable child will automatically turn red at the contact of the red circle, and the flu will spread to all surrounding vulnerable children. Different populations show the impact of vaccination rates conveying that immunizing a child does not only protect the child, “immunization is contributing to the control of the disease in the population” [Mea] (See Figure 2.23).

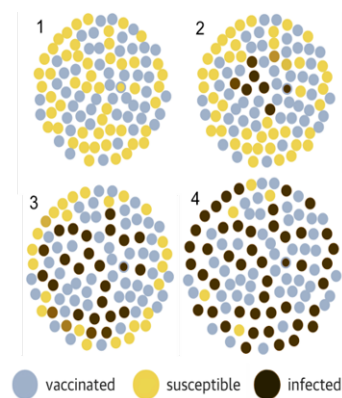


Figure 2.23: Illustration of Role 23: Animation shows how flu spreads in a population of vaccinated (blue) and vulnerable (yellow) children– Adapted from The Guardian [Mea]

In our preliminary exploration, when analyzing the responses for the question where we asked participants to cite examples of animation that convinced them about an idea, we observed that the persuasive impact of animation on viewers can depend on various factors such as the design of the animation (e.g. appropriately structured and without exaggeration), and the content and goal of the animation (e.g. animation that expresses cause and effect or conclusion or focuses on relevant part of the slide, well targeted animation).

Caveats of roles and research opportunities of visual discourse category

We introduced four animation roles under this category presented as follows: 21. *supporting a narrative*, 22. *highlight content* and 23. *persuading and convincing*. In the following, we discuss the main caveats and research opportunities of these roles.

Many parameters come into play when designing an animation for storytelling such as number of steps in a narrative, the timing, the mapping of the animation with the story content (e.g. semantically relevant choice of pictograms and motions). Much research is needed to study the impact of each design component. Like every other technology, animation can be misused to fill a nefarious role. For example, instead of attracting users’ attention to a new or important item it can be used to distract (focusing attention on the wrong thing as a strategy). It can be used to imply some magic effect that does not exist, or explain how things work in the wrong way. Research on deception [Pan+15] is certainly important to inform viewers and prevent designers to (un)intentionally use animation for such purpose.

The persuasive power of animation can be harmfully used as a way to mentally manipulate opinions and judgments [Nor14]. The manipulative power of animated ads has been deeply studied in the domain of media and communication. Researchers revealed, for example,

that people get very often interested in ads mainly due the influence of the attractive animations [SHK14]. However, they find very often difficulties to explain the rational reasons behind their interest and even to recall the most relevant content of these ads. Saul et al. [KD97] demonstrated that the persuasive power of animation can be used to facilitate the judgment of jurors but it can also mislead them toward false judgments.

2.1.4 Animation roles and people: summary and discussion

When reporting the findings from both our preliminary exploration and validation interviews, we gave details about the number of participants and practitioners that mentioned each role of our taxonomy in their responses. We will present below a summary of the coverage statistics and discuss briefly what we think are the possible interpretations of these results.

The chart in Figure 2.24 shows the number of participants in our preliminary exploration that mentioned animation instances illustrating each animation role in our taxonomy. This number can be used as a score to measure the popularity as well as the estimated and experimented usefulness of the different animation roles among the population of our exploration. The top three most popular and useful animation roles in the opinion of our participants are in the decreasing order: role 7 (explaining how a system works), role 22 (highlight content) and role 10 (hooking the user). This finding can be further supported by the fact that we found that participants employ majorly the following types of animation: appearance/disappearance of content (e.g. display text line by line, fade in/out of images, abrupt appearance and division effect), entry/exit of content (e.g. abrupt entry of content, “parade effect” to animate text entry), content emphasis (e.g. enlarge images, focus zoom effect on parts of the slide, blinking and flashing effects, “whirlwind” and rotation effect). Such observation does not seem surprising for us since these roles are the most relevant and useful for oral presentations that were the main focus of our preliminary investigation.

In opposition to the high-score roles, we can observe the high number of roles with very low scores of popularity and usefulness: 12 roles with a zero score (e.g. role 1, role 5, role 13, role 18), role 21 with a score equal to 1, and role 17 and 19 with a score equal to 2. The low score of most of these roles (role 1, 5, 13, 17, 18) seem logical and understandable because, as we have already mentioned before, these roles are not applicable in the context of the oral presentation that was the major focus of our preliminary exploratory. However, what seems really intriguing for us is the fact that people are still not taking advantage from the great potential of animation for telling stories and expressing emotions represented respectively by role 21 and role 19. That explains why we believe that these later roles incarnate relevant research opportunities to investigate in the future (see Table 2.1). Overall, we believe that deeper exploration in the future based on these observations would help us explain what makes a wide range of animation merits still not recognized and mis-understood by people who do not have particular knowledge and expertise about animation.




Figure 2.24: The number of participants in our preliminary exploration that mentioned examples of animation illustrating each role in our taxonomy of 23 animation roles.

Let us now have a reading through the coverage of animation roles in the answers of practitioners during the validation interviews. The chart in Figure 2.25 shows how many practitioners provided examples of animation related to each role in our taxonomy. One major observation that we can have at a first glance at the chart is that the level of expertise and knowledge about animation does make a notable difference in terms of the level of popularity and usefulness of animation roles. While popularity and usefulness of animation roles in the first chart were concentrated around a limited number of very high level scores, we can see here that they are more largely distributed between roles, with still relatively important disparity in score but very fewer number of zero-score roles (5 roles vs 12 roles in the first chart).

The top three most popular and useful animation roles in the opinion of our practitioners are in the decreasing order: role 16 (providing feedback of input mechanism), role 2 (supporting tracking during layout changes) and equally role 12 and 13 (providing visual comfort and aesthetics, making activity and progress visible). Since most of our practitioners are experts in HCI, this finding seems logical if we recall that the general goals of animation for keeping users in context and enhancing their experience with interactive systems have been the most studied goals along many decades.

Among the roles with very low score of popularity and usefulness we can cite the main following roles: role 4, role 15, and role 18 with a score of zero, role 8 with a score of 1, role 17 and role 20 with a score of 2. If we refer back to the Table 2.1 summarizing our taxonomy roles, we can find out that most of these roles (role 4, 8, 17, 20) represent important research opportunities, and that is partly due to the fact that they remain still not sufficiently studied, understood and exploited.




Figure 2.25: The number of practitioners in our validation interviews that mentioned examples of animation illustrating each role in our taxonomy of 23 animation roles.

Our reading through all these findings raised several additional questions about the ubiquity and prevalence of animation in our society and its visibility in the multiple interfaces that everyone of us is interacting with daily: Are all users able to detect all the animations that designers may hide in the websites or mobile apps they are using? What makes users not be aware and not recognize the same roles of animation although they often deal with the same set of user interfaces? Another relevant question concerns the memorability of animation: Users encounter dozens of new animation instances everyday. So, what makes them remember only a very limited subset of all these animations? We believe that answering these questions can help better understand the relation between people and animation.

2.1.5 General discussion

Like all taxonomies, our list of 23 roles organized in 5 categories attempts to be complete yet simple. Below we outline limitations and suggest areas for further refinement and validation.

Coverage of general roles of animation by our taxonomy

Amongst the 5 collaborators of this taxonomy, we gathered over 200 examples of animation to build and refine our 23 roles. We could assign all of our examples to at least one role in our list. To further assess whether our list was comprehensive, we asked 20 colleagues with different backgrounds (HCI researchers and practitioners such as user experience and interaction designers, UI developers) to provide examples they could think of. Overall, we collected 95 additional examples (accessible on our companion website: <http://animations.recherche.enac.fr/>). We could fit 98% of these examples within our taxonomy. The remaining (two examples) were thoroughly discussed as they dealt with animation reflecting user input, such as the parallax effect resulting from holding a smartphone at different angles. We initially opted to exclude these effects as they are very tightly coupled with

interaction, making it debatable whether we considered them animation. However, since several interviewees included such examples, we added the role 16: “providing feedback of input mechanism”, covering our entire corpus. Our corpus did not cover examples of animation beyond GUIs such as tangible interfaces (e.g. physical visualizations) or augmented reality interfaces. While these examples are beyond the scope of our taxonomy, we believe many of our roles would still be relevant. We hope future work will expand our classification and reveal specific roles of animation to such interfaces.

Combinations of roles and granularity

The roles listed in the taxonomy are not orthogonal and we encountered many examples where animation fulfilled multiple roles. For example, animation in Histomages [CDH12] and NodeTriX [HFM07] are used during the first encounter of the technique as a teaching aid to help users understand a mapping between a familiar representation (image and node-link diagram respectively) and a new representation of the same information (histogram and matrix). However, once users understand this mapping, the animation can aid them to keep in context, and support the tracking of object during successive layout changes. These animation instances might also be combined with others focusing on improving the user experience. We originally had many more roles which we eliminated as they seemed to be either a composite of multiple other roles, or rather application-specific and rarely encountered. For example, “confirming selection” was dropped as a combination of “making activity visible” and “maintaining up to date”. While we strived to reach optimal granularity for the roles (e.g. not representing a handful of examples but not representing several hundred either), it still varies substantially. While several roles could certainly be subdivided further, our intent was to provide a reasonable number for practical use and highlight novel ones since Baecker and Small [BS90], to inform and inspire future research on animation.

What limits the roles of animation?

The discussion of the caveats of roles revealed many negatives facets of animation and shed light on multiples factors that may have a considerable impact on the effectiveness of animation which can go up to turning its roles to be more detrimental than beneficial. These factors involve mainly the design characteristics of animation, the context of use where it is applied, and the phase of the interaction’s process with the user interface during which it is shown. More generally and beyond these role-specific drawbacks, studies about animation in literature have spotted multiple drawbacks of animation.

These drawbacks are mainly presented in research works as limitations and “side effects” of animation instances that they propose. Only few works discussed in a more elaborated way these limitations including mainly the work by Tversky et al. [TMB02a] explaining why animation can be problematic for learning and the work by Fisher [Fis10] discussing some major downsides of animation for data exploration. We believe that as animations designers need a structured view of animation roles, they do also need a structured and elaborated view of its limitations and drawbacks. Such structured view would help spot the major factors causing animation to be problematic and highlight relevant design considerations. It can also shed light on relevant research questions to explore in the future.

During this thesis, we attempted we take a step towards a deeper and more structured understanding of animation drawbacks. In the following, we present the findings that we extracted from our preliminary exploration aiming at understanding why people refrain from using animation and the negative experiences that they had with animation instances that they have encountered.

2.2 Animation drawbacks and people

Prior literature points to various downsides of animation. Many studies consider animation as time consuming and argued that it can make interaction and data exploration slower [Fis10; BB99]. Tversky [TMB02a] argued that animated visualizations are harder to perceive compared to their static counterparts. Animated visualizations are also considered less accurate than static ones [Fis10]. Animations can also be deceptive [ZS03], and irritating especially when users lose tracked objects [Fis10]. If poorly designed, animations are considered to be potentially misleading and can violate the semantics of the underlying data [HR07].

In our preliminary exploration, respondents provided various reasons why they refrain from using animation in their oral presentations and user interfaces in general and enumerated multiple examples of animation that gave them negative impressions. We will present both the perceived and experienced drawbacks of animation by participants. We will then propose a classification of the main drawbacks of animation mentioned by our participants.

2.2.1 Why people do not use animation?

In our preliminary exploration, participants reported that they refrain from using animation for various reasons. We break them down into categories as follows:

Animation is useless

A fair number of participants (10/33 in the survey, 5/23 in the interview) qualified animation as “not professional”, “gadget feature”, “needless for explanation”, “waste of time”, “just for wow”, “without interest”, “heavy when used just for aesthetics”, “corny”, and “childish”.

Animation can be misleading

Some participants (3/33 in the survey, 1/23 in the interview) reported that animation is misleading since it can sometimes convey wrong information. One participant cited the example of an animated visualization where a fish eye lens moves over pushing consequently the data items apart. She argued that even if the animation effect is aimed essentially to reduce visual clutter, it might cause a misunderstanding of the data.

Animation can be confusing and disturbing

A few participants (2/23 in the interview) considered some examples of poorly designed animation as confusing and disturbing. Such examples include for example Prezi [Preb] animation with excessive zoom in/out effects and animation with complex motion trajectories. Another example is the use of animation effects not suited for text such as translation and rotation. A participant found that these instances of “nasty” text animation hurt the eyes, cause discomfort and make the animated interface hard to follow.

Animation can be distracting

Several participants (3/33 in the survey, 2/23 in the interview) reported that they refrain from using animation because they consider that it distracts drastically the audience and disrupts their concentration on the principal information.

Animation is time consuming

In the context of oral presentations, some participants (2/23 in the interview) explained that they refrain from using animation because it consumes time to the detriment of the speech and the follow-up discussion.

Animation design is tedious and time consuming

Several participants (3/33 in the survey, 5/23 in the interview) reported that they can not design appropriate animation and manage it easily. They also reported that well-designed animation requires many tweaking and takes a lot of time.

2.2.2 Examples of ineffective animation as seen by people

When asked about examples of animation that they found ineffective, participants provided multiple examples of “bad animation” that we classify as follows:

Animation that they found confusing

Examples of animation that participants considered confusing include for instance: 3D transition in Prezi that goes in all the directions, animation without clear subject or purpose, animation without any link to the animated content, exaggerated animation with excessive movement or bouncing effects, poorly-structured animation where for example animated items are not well arranged in space.

Animation that they found disturbing

Examples of animation that some participants found disturbing included: animation combined with annoying sound effects, animation composed of repetitive movements (e.g. objects are moving back and forth continuously), very fast animation, intrusive pop-up animation that interrupt the navigation on websites without any prior notification, obtrusive animation that hinders the visibility of the main information (e.g. pop-up animation on mobile applications that completely shift the text out of the screen), repetitive and excessive animation (abused use of animation, e.g. every single piece of text, image, slide appears with an animation, objects moving in all directions up/down, left/right). More particularly, some participants reported that all “heavy” Prezi animations make them feel dizzy and “nauseous”.

Animation that they found misleading

Participants found animation were misleading, causing them to misunderstand the conveyed information mainly because they had difficulty to decipher the information or meaning that animation were intended to convey (e.g. “with animated colors of button hovered zones, I did not understand which color encodes which state (enabled/disabled)”), or because some aspects of the animation were not appropriately designed (e.g. poor temporal structuring, trajectories with random directions).

Animation that they disliked

In general, participants disliked the “funky” and “superfluous” “purposeless” animation that does not contribute to the understanding of the underlying information. They also did not appreciate animation that affects negatively the readability of information (e.g. very stylish and exaggerated animation in the style of for instance the “dissolution”, “checkerboard” or “explosion” transitions, or “whirlpool trajectory” in PowerPoint) and the fluidity of the navigation (e.g. very slow or very quick animation or complex Flash animation). Moreover, participants disliked animation that is without a clear or logical scenario, incoherent animation effects between similar types of information, animation used in an inappropriate context (e.g. joyful animation in a visualization talking about the sad topic of diseases). In the context of information visualizations, one participant disliked the animation because she considers that it is cognitively demanding and it destructs continuously the mental map that she constructs about the visual layout, the organization and the distribution of the visual items composing a visualization.

2.2.3 Classification of animation drawbacks as seen by people

We compiled a list of the main keywords that participants used to qualify the negative aspects of animation and we classified them based on a set of dimensions characterizing both the design and use of animation in user interfaces (Table 2.2). The main purpose of this classification is to initiate reflections in the future about a taxonomy to characterize the general drawbacks of animation in user interfaces.

Table 2.2: Classification of animation drawbacks

General design characteristics	Flashy Fancy Interrupted Difficult to design Not fluid
Location in the interface (Animation parameter)	Badly placed in such a way that it hides the main information
Trajectory (Animation parameter)	Jitter Jerk (bouncing movement)
Duration and speed (Animation parameter)	Slow Sluggish Too fast Abrupt Lengthy Unresponsive
Structuring (Animation parameter)	Visually cluttered Messy Overwhelming Information overload
Nb. animations (Frequency of animation)	Too many
Nb. repetitions (Frequency of animation)	Repetitive
Comprehensibility	Difficult to understand Not clear
Usefulness	Unnecessary Lag Time consuming In-your-face effect (intrusive) Unresponsive Try-hard Screen tearing Consumes graphical resources
Impact on interaction	Non-interactive interface due to playing animation Annoying Fatigue Unpleasing when too long Hurts my eyes
Impact on users	Distracting Obtrusive makes the text parts useless Users remember more the animation than the information itself
Impact on visibility/saliency/recall of information	Dominant upon the main information Animation can mutilate the purpose of a speech or hinder its comprehension Cognitive overload
Impact on Understanding Information	Misleading Valueless Childish Useless Superfluous overload Useless, only for aesthetics Not always relevant Use in inappropriate context

2.3 Conclusion

In the first part of this chapter, we have revisited the original classification of animation roles by Baecker and Small [BS90] in light of 28 years of work in HCI in academia and in industry. Our goal is to highlight the growing opportunities for the use of animation. A common language about the roles of animation will foster renewed discussions about animation, and in turn encourage a broader scope of empirical research assessing their merits and limitations. For example, we now include the use of animation to engage, persuade or even fight boredom, commonly used in practice but rarely studied in the literature. Such roles will need to be investigated in more depth as research on storytelling [KM13; Lee+15] and visualization literacy advances [Boy+14; Alp+17]. Furthermore, we highlight the different roles that animation can play to encode data and that also remain in their turn under-explored in the

literature. We believe that animation incarnate a considerable semantic value that gives it a interesting expressive power that can inspire and enrich future works especially in the field of information visualization.

While the first part of this chapter exhibited the functional and useful side of animation, the second part shed light on its downsides. In this latter part, we have discussed the drawbacks of animation based on the findings of our preliminary exploration and classified the main drawbacks of animation as seen by our participants based on a list of dimensions characterizing both its design and use. We believe that such structured view on the downsides of animation represents a constructive step towards elaborating a taxonomy of animation drawbacks that would contribute to gaining a more complete and deeper knowledge about the various disparate facets of animation in user interfaces.

Animation can be seen a double-edged sword that can be beneficial if well designed and used in a suitable context, or can be rather harmful when poorly designed or not appropriately used. Hence, both the design and the context in which animation is used together represent the main key factors that will determine whether an animation is useful and effective.

In the remainder of this thesis, we set out to better understand the potential of animation through a focused exploration of the definition, and potential of animation, applied to information visualization. This scope not only provides a propitious research field to explore the most promising categories of animation roles that yet remain over-looked (e.g. data encoding, visual discourse), but is an excellent use case to study in depth the negative and harmful impact that animation can have on the perception and understanding of visual interfaces.

Along the following chapters of this thesis, we will present the results of our exploration of the use of animation in the context of information visualization, starting with an introductory chapter about this application domain.

Chapter 3

Understanding dynamic information visualizations

As explained in the introduction of this thesis, animated transition is a particular class of animation which consists of making the abrupt visual change between *two* successive states more fluid and continuous, usually by interpolating between these two states. Animated transition represents the most prevalent implementation and use of animation in the context of information visualizations.

We argue that a crucial step toward an effective design of animated transition in dynamic visualizations is to understand their fundamental constructing components with the interactions between them, and figure out the various changes that such visualizations can undergo. Such understanding would provide the animation designer the necessary knowledge she needs about the main characteristics of a given dynamic visualization that would both inform and constrain the design of animation.

An information visualization includes essentially visual objects representing the data in a specific representation space (e.g. 2D or 3D space, with axes or without axes). When that visualization changes over time, it comprises additional aspects such as change of states, visual transformations and so forth, that portray its *dynamic behavior* in time. Based on previous works, we discuss these different aspects and propose a conceptual model to describe them. Most of taxonomies characterizing the dynamic change of information visualizations that we can find in literature focus mainly on answering the main following question: *(Q1) What are the causes of dynamic changes in visualization?* Beyond figuring out what can trigger a dynamic change in a given visualization, the animation designer needs also a finer knowledge about: *(Q2) What components change precisely in this visualization? (Q3) and how they change?*

To answer these two latter questions, we will discuss in this chapter a set of dimensions to characterize dynamic changes in visualization: *when* (time period of dynamic change), *who* (scope of dynamic change), *what* (type of dynamic change) and *how* (manner of dynamic change). To explore more in depth the *what* and *how* aspects, we will present a taxonomy of dynamic changes and discuss different metrics to characterize the manner of dynamic changes. The knowledge about the possible dynamic changes in a visualization (Q1), the visualization's components that are affected by each change (Q2) and the manner in which these components change (Q3) are the main input elements of the design process of animation for dynamic visualizations.

3.1 Information visualization: a simplified model

The *Reference Model* suggested by Heer and Agrawala [HA06] provides a design pattern for building information visualizations that separates *data models*, *visual models*, *views*, and *interactive controls*. The *data model* includes the abstract data visually represented by visual attributes (e.g. position, size, color etc.) composing *the visual model*. The *view* represents the graphical display of the visual model, and the *interactive control* manages the user input and triggers updates to the different levels of the system (Figure 3.1).

We build upon this model and we propose a *simplified model* where we consider that an information visualization is composed of two main components: a *data component* including both the data source and dataset, and a *visual component* that comprises both the visual representation of this data and its graphical display. We also take into consideration the user's

interaction with the visualization but we do not study it as a separate component (Figure 3.2). We do not present that simplified model as an alternative to the Reference Model [HA06], but rather an abstraction of that model where we focus on the high level of data and visual aspects of a visualization.




Figure 3.1: The Reference Model pattern, adapted from [HA06]




Figure 3.2: Our simplified model of information visualization based on the Reference Model pattern [HA06].

3.2 From abstract data to visual representation: visual encoding

Abstract data consists of a collection of *data objects* that have a set of attributes called *data attributes*. The process of *visual encoding* (alternatively known as visual mapping) involves mainly defining *the visualization type* (e.g. chart, network, map, etc.) and transforming the data objects into a *visual representation*.

Engelhardt [Eng] introduced the “ingredients” of a visual representation of data which are: *objects*, *spaces*, and *properties*, where the objects are graphical objects presenting visual components at different levels of the visual representation. Thus, an entire graph visualization can be considered as a graphical object and the nodes composing it correspond as well to graphical objects. The space represents the surface containing the graphical objects, and the properties include all the visual variables characterizing these objects such as color, size, position and so forth. Following the same approach, we consider that the *visual representation* is constructed by *visual objects* and their corresponding *visual properties* within a *representation space* (see Figure 3.3).

Each *data object* is associated to a *visual object* by encoding its *data attributes* into *visual attributes*. Visual objects are visual marks characterized by a position (i.e. spatial variable) and a set of retinal variables defined initially by Bertin [Ber83] as shape, size, color hue, color value, color intensity and texture. Later research works proposed additional visual variables such as motion [Car03b] and transparency [Mac92]. The retinal properties define the *appearance* or *aspect* of the visual object. It is relevant to recall that not all the data attributes of a data object are necessarily visually encoded. Similarly, not all the visual attributes are associated to data attributes. To better explain this idea, let us take a concrete example. The

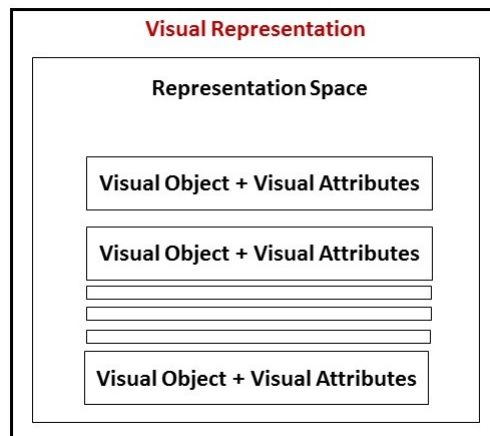


Figure 3.3: Components of a visual representation of data

Figure 3.4 represents an excerpt of an imaginary dataset comparing the elevation and m^2 price of three categories of accommodation namely: house, apartment building and skyscraper. The Figure 3.5 illustrates a sketch of a visualization representing a part of the data. Both the color hue and shape of visual marks encode the data attribute “accommodation category”, the x-position and size are not mapped to any data attribute and the height of visual marks encodes the maximum data value of elevation level of the corresponding accommodation category. The m^2 price is not associated to any visual variable the visualization.

Overall, the visual encoding consists essentially of visually representing data through visual marks. However, the visual encoding is constrained by multiple factors among which for instance the purpose and the type of visualization. The purpose of a visualization defines for example what piece of information we want to present to viewers and what insights we want them to gain. Deciding these aspects guides many choices such as the set of data attributes to encode visually. Let us consider that we have a dataset of various information about countries: weather, security, transport, housing, education, health, income, imports and exports. If we want to inform viewers about the tourism in a country, we would more likely visualize data about weather, security and transport for example. However, a visualization representing the life quality of that country would rather represent data about, say health, housing, education, and income. Similarly, the type of visualization can constrain the visual encoding by, for example, imposing the use or non-use of some visual variables as it is the case for maps where the (x,y) position is a fixed visual attribute (we need to put visual objects somewhere in the plane!), and bar charts where bars cannot take different shapes (a bar is a bar!).

Accommodation category	Min elevation (m)	Max elevation (m)	Min price m^2 (euro)	Max price m^2 (euro)
House	4	8	100	3000
Apartment building	10	20	1000	8000
Skyscraper	40	800	100M	800M

Figure 3.4: An excerpt of an imaginary dataset comparing elevation and prices of three categories of accommodation.




Figure 3.5: A sketch of visualization representing the maximum elevation of three categories of accommodation: house, apartment building and skyscraper.

The visual objects are contained in the *representation space* (introduced by Card et al. [CMS99a; Car03a] as *spatial substrate*). This graphical space can be either non metric and thus not contain axes, or it can have a metric structure and be characterized by a defined category of axes. Card et al. [CMS99a] defined four types of axes: unstructured axis (when there is no axis), nominal axis (when a region is divided into sub-regions), ordinal axis (when the ordering of these sub-regions is meaningful) and quantitative axis (when a region has a metric). This axis classification was equivalently defined by Engelhardt [Eng] as axis, ordered lineup, and unordered lineup. Card et al. [CMS99a] mentioned also that axes can be either linear or radial.

In reference to our *simplified model* of the information visualization (Figure 3.2), the data component includes the data objects with their data attributes while the visual component consists of the visual representation including the visual objects, their visual attributes and the representation space where these objects are displayed (Figure 3.6).

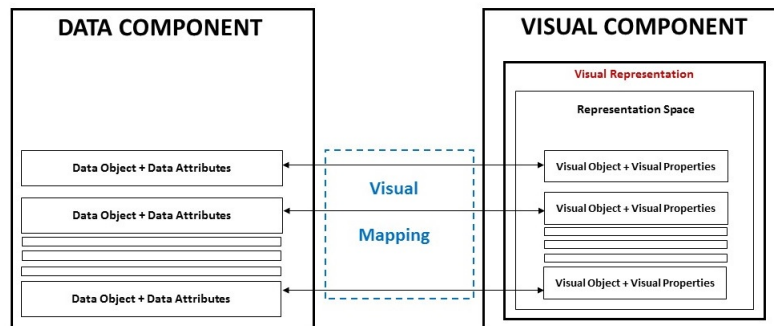


Figure 3.6: Overview of the data and visual components based on our simplified model.

Among the main characteristics of data, we highlight *data dimensionality* and *data cardinality* (respectively introduced by North [Nor06] as data schema and quantity). The data dimensionality is defined by the list of data attributes of data objects composing a dataset, whereas, the data cardinality specifies the number of data objects composing the dataset.

3.3 Dynamics of an information visualization

An information visualization typically portrays data objects, represented by visual objects with visual attributes, in a specific representation space. Together, the representation space,

the visual objects contained in it and their visual attributes compose a *visual representation* of data at a specific instant and from a specific perspective.

Interactive visualizations dynamically convey information through a series of states over time. Visualization states are discrete events that are either automatically generated based essentially on the current flow of information to be visualized (e.g. dynamic visualization of streaming data, clock data, sensor update), or triggered in response to users' actions consisting of, for instance, specifying the information of interest (e.g. filtering), changing the organization of data (e.g. sorting) or navigating through different levels of the visualization (e.g. zoom in/out).

Following our approach of structuring the information visualization in data and visual components, we consider that the visualization state includes two aspects: *data state* and *visual state*. The *data state* at a specific instant t is defined by the set of data objects present at that instant and the value of their corresponding data attributes. Whereas, the *visual state* at a specific instant t is presented by the set of visual objects present at that instant with their position and visual aspect, and the representation space in which they are arranged.

The change of the visualization from an initial state S_0 to a new state S_1 involves always a transformation of the visual state. However, it does not necessarily and always portray a data state change. When the visual component changes, the change can be due exclusively to an visual transformation or to a change in the data component. In contrast, a change of the data component, generates usually a transformation of the visual component, but sometimes that visual change may not be visible like for example when the visual objects corresponding to the changing data objects are outside the viewport.

3.3.1 Visualization states mapping

Let us consider two successive visualization states: S_0, S_1 defined by the object sets O_0 and O_1 respectively. When the visualization changes from S_0 to S_1 , O_0 will be transformed into O_1 through a set of operations that can be described by the *mapping relations* (alternatively designated as relations of equivalence [Che+10] and matching [VEF15]) between the two states. The possible mapping relations that can exist between two states of a visualization can be summarized as follows:

- One-to-One mapping: an object in S_0 reappears and persists in S_1 . It can remain unchanged during the transition from S_0 to S_1 or can have some transformations of its properties (its visual attributes change over time)
- One-to-None: an object existent in S_0 disappears in S_1
- None-to-One: an object non existent in S_0 appears in S_1
- One-to-Many: one object in S_0 gives birth to many objects in S_1
- Many-to-One: many objects in S_0 merge onto one object in S_1

If the mapping relations between all the elements of the sets O_0 and O_1 are of the type *one-to-one*, S_0 and S_1 share the same sets of objects with possible differences in properties between initial and final objects. Otherwise, S_0 and S_1 share only a subset of objects. Depending on the type of visualization state (i.e. data or visual state), each relation of mapping can be generated through particular types of dynamic changes in the data or visual component of the visualization as we will discuss in details in the next section.

Generating the mapping relations in a determinant way requires a fine-grained *identity mapping* [CLW12] for which we need to know which objects in the initial state S_0 corresponds exactly to which objects in the final state S_1 . However, such knowledge is not always given and can rather be generated or calculated using different methods to construct the mapping

relations including for examples applying the operations of matrices' inverse and composition to match between the code views and its corresponding document view in Gliimpse [DHC11], calculating the equivalences of positions of a given word across a document revisions based on the width and height of the document's page [Che+10], and using combinatorial optimization algorithms to perform the matching between animated elements in such a way that minimizes the sum of distances or the maximum distance between these elements [VEF15].

3.3.2 Characterizing dynamic changes in information visualization

Let us now refer back to the three main questions that we presented at the beginning of this chapter:

- Q1 What are the causes of dynamic changes in visualization?
- Q2 What components change precisely in this visualization?
- Q3 How these components change?

When surveying the literature, we found two main works that answered these questions in the specific context of graph visualization. Ahn et al. [APS14] suggested a task taxonomy for temporal network evolution. They define the temporal network evolution tasks as user-initiated operations over network entities and their properties to achieve specific goals. They described the tasks based on three dimensions: entity, property, and temporal feature. The entities are the objects of interest in the data analysis and can be: node/link, group, or network. The property of an entity changes over time. Properties are defined as the characteristics of the entities and include both structural properties and domain properties, which can be compared over time. The temporal features are defined as the predicates that change the state of the entities and the properties over time. Analysts need to identify temporal features in addition to the entities and the properties, so that they can answer questions about network evolution.

In GraphDiaries [BPF14], Bach et al. structured their task taxonomy for dynamic networks based on 3 dimensions: *when*, *where*, *what*. The *when* represents the time dimension and can include a particular time (snapshot), two times, a period, and all times. Furthermore, in addition to asking questions about "when" in the tasks, the authors took into consideration additional attributes such as how often, how fast, how long, in what order (inspired from [Mac04]) during, and starts (inspired from [All84]). The *where* dimension represents the topological structures composing the network such as: nodes and their attributes, links, graph or sub-graph, cluster, group, etc. Finally, the *what* dimension includes the type of change and the behavior of network elements.

Building upon these two task taxonomies, we propose to characterize dynamic changes in information visualizations based on the 4 following dimensions: *when*, *who*, *what*, and *how*.

When: defines the *time period* of a dynamic change.

Dynamic changes can happen between two discrete timesteps that can be or not be consecutive (2 time steps), or across a series of consecutive or non consecutive timesteps (n time steps). This dimension allows us to explore the following question: "*Between which time steps or across which time steps a given dynamic change happened?*"

Who: defines the *scope* of a dynamic change.

The scope specifies which entity in the visualization was affected by a dynamic change. The dynamic change can involve either a *single item* or a *set of items*. A set of items is a group of

items sharing one or several common *static* visual attributes or *dynamic behaviors* [Lee+06; VBW17]. This dimension allows us to explore the following question: “Which item or set of items was affected by a given dynamic change?”

What: defines the *type of a dynamic change*.

The type of a dynamic change specifies what happened to a visual item or a set of items and what changed in it. We will present below a taxonomy describing the different types of change that a dynamic visualization can undergo and we will also discuss the variation of these changes depending on the scope they are targeting.

Many previous works investigated the dynamics of information visualizations and classified their changes in various taxonomies. Heer and Robertson [HR07] proposed a taxonomy of transition types between data graphics in the context of statistical dynamic visualizations. They identified 7 different types of transitions: *view transformation*, *substrate transformation*, *filtering*, *ordering*, *timestep*, *visualization change* and *data schema change*. Based on Heer and Robertson’s taxonomy, Fisher [Fis10] suggested a taxonomy describing the types of change that can occur within a visualization. He distinguished 6 types of transitions that can describe visualization changes including both data and view transformations: *change of the view*, *change of charting surface*, *filter the data*, *reorder the data*, *change the representation* and *change of the data*.

Heer and Shneiderman later introduced a taxonomy for interactive dynamics for visual analysis [HS12]. This taxonomy focuses on the tools allowing creating visualizations, interacting with them and collaborating around them. It consists of 12 task types grouped into three high-level categories: (1) data and view specification (visualize, filter, sort, and derive), (2) view manipulation (select, navigate, coordinate, and organize), and (3) analysis process and provenance (record, annotate, share, and guide). Heer and Shneiderman’s taxonomy does not study the dynamic aspect of visualizations, but it provides an overview of the possible interaction operations that users can have with a visualization and that can trigger a change in this visualization.

Cottam et al. [CLW12] presented a taxonomy for dynamic visualizations of streaming data, that aims at investigating the effect of the change of visual variables on the comparison of data states and the identification of the different graphical objects changing over time. They classify dynamic visualizations using a two dimensional matrix defined by spatial and retinal properties. Spatial variables may be dynamic in two ways: changing existing values and changing the number of visual elements (i.e. adding or deleting). Retinal values change over time as the attributes they are encoding evolve. This taxonomy gives a good overview of the data-driven changes of a dynamic visualization. However, as the authors argued, it is also significant to study and highlight the difference between “user-requested” (that we can alternatively call visual-derived) and “data-derived” changes, which was not explored in their work. Additionally, the authors investigated only the changes that affect the data component of a visualization while multiple changes that concern the visual component are relevant to explore as we can recognize for all the taxonomies presented above.

All these different taxonomies give to the animation designer several levels of answers to the above first question: What are the causes of dynamic changes in visualization? Once the designer acquires the knowledge she needs about the different changes a dynamic visualization can undergo, she needs to concretely know what are the visual and data components affected by each particular change (Q2), and how these components changed (Q3). The knowledge about the fine-grained changes that visualization components undergo is the major input that the designer needs to effectively define the different animation parameters including trajectories, pacing function, start time, and so on.

Whatever is the high-level dynamic change the animation designer wants to portray, she will always handle the same set of dynamic changes of the visualization data and visual components at a low-level. For example, during a filtering change or a timestep change, some visual elements will disappear, some will persist and others will appear. We argue that one of the most relevant knowledge about dynamic changes from high-level perspective is whether a given dynamic change is driven by the data or visual component. This knowledge can be very informative and even sometimes decisive for the design decisions that the designer would have to take when designing an animation. We assume that in the case of a visual-driven change, the animation designer can have more flexibility in choosing among design alternatives mainly because there is arguably less risk of violating the meaning of data and inducing wrong interpretations compared to the case where the designer is animating data-driven changes. When animating data-driven changes, the designer may restrict their design choices and put reserves on some design options. For example, they may refrain from using complex trajectories and stick to the standard linear ones.

How: defines the *manner of a dynamic change*.

We will discuss below a set of metrics allowing to characterize the manner of a dynamic change including: the direction of change, value of change, shape of change, rate of change and cardinality of change.

3.3.3 A taxonomy of dynamic changes in information visualization

As we suggested in Section 3.1, an information visualization is structured in two components: the data component and visual component. A visual change can be driven by the data or by the visual component. Hence, we distinguish two types of dynamic change in an information visualization: *data-driven change* and *visual-driven change*. When the data changes, the corresponding visual objects necessarily change. However, this consequent visual transformation is not always perceivable as it is the case, for instance, when the visual components are displayed outside the viewport. When the visual component changes, it does not affect the data.

For each category of change (data-driven and visual-driven), we will start by enumerating the different types of high-level changes issued from the taxonomy of Heer and Robertson [HR07] that belong to the category. We will then discuss the low-level changes at different levels of granularity of the corresponding component (data and visual component) and highlight for each data change its associated visual change.

Levels of dynamic change

The data component is composed of a set of data objects and their corresponding data attributes. We can define two levels of *data-driven change*: *data attribute level* and *data object level*. The visual component is mainly composed by visual objects and their visual variables and the representation space in which these objects are displayed. The representation space is in its turn composed by a set of *charting components* that ensure the framing the visualization and include for example the axes and their labels, the legends, the titles, and so forth. Hence, we define 3 levels of *visual-driven change*: *visual variable level*, *visual object level*, and *charting components level*.

Data-Driven Change

There are different reasons that can lead to the change of data. The most evident case is the change of data source. However, in the context of this chapter, we focus on situations where the data source does not change. In that case, if we deal with real time data like, for instance, streaming data, sensor update, clock, stock market information and so on, then the dataset changes continuously over time. More generally, when dealing with temporal data, a timestep change causes generally a data change. Less commonly, data can change due to users' actions as it is the case for example in file explorers of operating systems or in database management systems where users can for example add, delete and edit files or database rows. With reference to Heer and Robertson's [HR07] taxonomy, the data-driven changes include the timestep change and the data schema change.

At a low level, a data-driven change involves various types of transformations: change of data value (modify the value of one or several data attributes of one or several data objects), change of data schema (show or hide one or many data attributes of the data objects), change of the data cardinality (change the number of data objects by adding or removing objects). With reference to the levels of dynamic change detailed above, we distinguish two levels of data-driven changes: *data attribute change* and *data object change*:

1. Data Attribute Change

- Evolve data attribute: this transformation consists of changing the value of the data attribute. It generates the *change of data value* and induces the evolution of the corresponding visual attribute if the data attribute is visually encoded. The evolution of data attribute can be described through a One-to-One mapping between the initial and final state of the data attribute.
- Hide data attribute: this transformation allows the navigation through the hierarchy of data dimensions in the context of an hierarchical data schema. It consists of a zoom-out operation in the data schema to show less details. In the context of OLAP (Online Analytical Processing) this operation is called *roll-up* [Aal13]. This transformation can be described by a Many-to-One mapping, between the initial and final data states. Following an operation of merge, many visual objects merge onto one visual object.
- Show data attribute: this transformation is the reverse of *hide attribute*. It consists of zooming-in the data hierarchy to show more details. In the context of OLAP this operation is called *drill-down* [Aal13]. This transformation can be described by a One-to-Many mapping between the initial and final data states. Following an operation of split, a visual object split onto several visual objects.

2. Data Object Change

- Evolve data object: this transformation consists of the combination of all the changes of data attributes characterizing a data object. This transformation induces a change in the *position* or the *aspect* of the visual object representing the data object. This transformation can be described by a One-to-One mapping between the initial and final state of the data object.
- Create data object: this transformation consists of adding a new data object. It induces the *change of data cardinality*. This transformation induces the creation and appearance of a visual object representing the data object and represents a change in the *existence* of this object. This transformation can be described by a None-to-One mapping between the initial and final data states.

- Delete data object: this transformation consists of deleting a data object generating also the *change of data cardinality*. This transformation induces the deletion and disappearance of the associated visual object and represents also a change in the *existence* of this visual object. This transformation can be described by a One-to-None mapping between the initial and final data states.

Visual-Driven Change

With reference to Heer and Robertson [HR07] taxonomy, the visual-driven changes include the view transformation, substrate transformation, filtering, ordering, and visualization change. All these different changes involve 3 types of low-level visual-driven changes: *visual attribute change*, *visual object change*, and *charting components change*.

1. Visual Attribute Change

- Evolve visual attribute: this transformation consists of changing the value of the visual attribute. This operation does not involve any data change. Many examples illustrate this transformation such as changing color of nodes in a graph, or highlighting a subset of points in a scatterplot during a brushing operation as well as growing or shrinking the size of visual objects during a zoom in/out operation. This transformation can be described by a One-to-One mapping between the initial and final state of the visual attribute.
- Change of visual variable's type: this transformation consists of replacing a visual variable by another one like for instance encoding a data attribute in size instead of color. That generates the *change of visual mapping*. This transformation can be described by a One-to-One mapping between the old and new visual variable.

2. Visual Object Change

- Evolve visual object: this transformation consists of the combination of all the changes of visual variables characterizing a visual object. This transformation can alter the *position* and the *aspect* of the visual object. This transformation can be described by a One-to-One mapping between the initial and final state of the visual object.
- Show/Hide visual object: these transformations consist of altering the *presence* or *visibility* of the visual object. These changes occur mainly during the operations of filtering, pan and zoom, and the change of viewport. This transformation can be described respectively by a None-to-One or a One-to-None mapping between the initial and final visual state.

3. **Charting components change:** This change corresponds to the substrate transformation in the taxonomy by Heer and Robertson [HR07] and includes for example the transformation of axes (e.g. axis rescale, change from linear to log scale, etc.), the change of the projection system (e.g. change from a Mercator projection to a globe [Fis10]) as well as the bifocal and graphical fisheye distortions [HR07].

Variation of types of change based on the scope of change

When we read through the different types of changes detailed above, we can observe that all of these changes involve two major categories of change: *persistence* and evolution. In the following, we will explain how these two categories of change vary depending the scope of change (i.e. single item or set of items).

Persistence:

The definition of the *persistence* change depends on the *scope* of that change.

- *single* item: a single item can appear, disappear, or evolve.
- *set* of items: in the same approach of Ahn et al. [APS14], we take into consideration the aggregation of the change affecting individual items and its impact on the sets they belong to. Hence, a set of items can **grow** (as a result of one or many items appearance), **contract** (as a result of one or many items disappearance), **appear** (if all items appear at once) or **disappear** (in reference to the birth and death transformations in the work of Ahn et al. [APS14])

Evolution:

The definition of the *evolution* depends on the *scope* that change.

- *single* item: the evolution of a single item can involve the change of one or many of its visual attributes.
- *set* of items: the evolution of a set of items is the combination of the evolution of single items composing it. Thus, in the same way as above, this evolution can involve the change of one or many visual attributes. However, in reality, each visual item in the set can have a different evolution behavior for each of its visual attributes. Hence, in a set level, it makes more sense to focus on the evolution of only one visual attribute across all the items composing that set.

3.3.4 Manner of a dynamic change

The manner of change can be characterized using the following criteria: direction of change, value of change, shape of change, rate of change, cardinality of change.

Direction of change: defines the difference between the final and initial values of each data attribute. At a first level, the direction of change indicates if an attribute is “Stable” or “Changing” by answering this first question: “Are the final and initial values equal or different?”. Then, if the answer is rather that these two values are different, comes the second level question: “How are they different?”. The answer of this question depends on the type of the data attribute. If it is *categorical* or *ordinal*, we can characterize the direction of change based on the final value or/and the initial value of that attribute. In that way, we can get different variations of direction of change such as:

- From first value a to \forall final value
- From \forall first value to final value a
- From value a to value b

When the data attribute is *quantitative* the direction of change is defined by the sign (positive or negative) of the difference value between the final and initial values (Δt), and can be as follows:

- Stable: when the final value is equal to the initial ($\Delta v = 0$)
- Increase: when the final value is bigger then the initial $\Delta v > 0$)

- Decrease: when the final value is smaller than the initial $\Delta v < 0$)

We can use the direction of change metric to characterize only dynamic changes happening between 2 timesteps.

Value of change: in the case of 2 timesteps, the value represents the quantity or amount of the change and is defined by the absolute value of dt ($|\Delta t|$). In the case of an n data states change, the value of change can be either reduced to the value of change between the initial and final data states, or can be alternatively defined based on an aggregation rule (e.g. sum, average, minimum, maximum).

Shape of change: when we plot the evolution of a data attribute across n time steps on a graph, we can get a shape that describes the variation of the direction of change over time. This shape depends on the *monotonicity* of the change across the time steps in the following way [APS14]:

- Monotonous change: a change can keep the same direction of change over time. Thus, it can have the following shapes of change:
 - All the time increasing
 - All the time decreasing
 - All the time stable

These three shapes represent the “primitive” shapes whose concatenation/composition/alternation gives more composite shapes representing the non monotonous change as detailed below.

- Non Monotonous change: a change can have different directions of change over time. Thus, it can have the following shapes of change:
 - Concatenation-Convergence: a converging shape represents a data attribute that first increases or decreases then remains stable
 - Concatenation-Divergence: a diverging shape represents a data attribute that first remains stable then increases or decreases
 - Alternation-Fluctuating
 - Alternation-Regular
 - Peak or valley: “an entity property increases or decreases abruptly and then returns to its earlier value.”

Rate of change: defines the rate of change (“the amount of change in a given time period” [APS14]). It can have the following values: acceleration and deceleration. The rate of change can be used to characterize only n timesteps changes.

Cardinality of change: this aspect defines how many attributes or items are concerned with the change. This aspect was essentially inspired from the study of Robertson et al. [Rob+08] about the effectiveness of animation in trend visualization where they gave to the users some tasks dealing with the cardinality such as the following question: “Which continent had the most significant increase in GDP per capita (i.e. the continent with largest percentage of countries with significant increases in GDP).” However, it is convenient to mention that the term of cardinality of change can characterize another aspect of the dynamic change as it is the case of TreeVersity by Gomez et al. [Gg+13] where cardinality of change defines the *direction of change* in our case. The cardinality of change can be used to characterize both 2 and n timesteps changes.

3.4 Conclusion

In this chapter, we presented the theoretical basics that allow the animation designer to understand and characterize the dynamic changes in information visualizations. Such knowledge represents a crucial input that informs and constrains the design of animated transition in visualizations. The animation designer needs now to figure out the basics to design an effective animation and measure its effectiveness for dynamic visualizations. In the next chapter, we will discuss the fundamentals and challenges of designing animation for information visualizations.

Chapter 4

Animation for information visualization: fundamentals and challenges

There is no particular mystery in animation...it's really simple, and like anything that is simple, it is about the hardest thing in the world to do.

— Bill Tylda at the Walt Disney Studio, June 28, 1937 [Gra37]

The previous chapter allows the animation designer to answer the main following question: “What can I animate in a dynamic visualization?”. Now, the designer needs to know how to design an effective animation for a given dynamic visualization. To do so, she needs to gain a good understanding of the different animation parameters and define them in the most optimal way. Once the animation is created, the designer needs to assess to which extent this animation is useful and effective. We start this chapter by explaining the fundamentals of animation and we follow up with a discussion of both the basics and challenges of designing and evaluating animation in dynamic visualizations.

4.1 Fundamentals of animation

Animation is a complex visual construction that involves several parameters, namely: trajectory, start time, duration and pacing. When designed for information visualizations, the complexity of animation gets considerably incremented by the complexity of the dynamic changes that it is conveying: animation designers have often to combine and synchronize several animated transformations within the same animation. Such complexity is managed through several *composition* approaches that we will discuss in this section.

4.1.1 Animation parameters

As explained in the previous chapter, visual items (i.e. visual objects or charting components) composing a dynamic visualization are characterized by a set of visual properties including their position, size, shape, color, and so forth. Animation applies a set of graphical transformations to these visual items during a given interval of time. When an animation designer thinks about creating an animation, she asks three main questions: i) What are the transformations that I should apply to each visual item? ii) When should I start and finish the animation of each visual item? iii) At which pace should the animation progress?

The first question concerns the spatial path along which each visual item will move, and the aspect it will have during the animation. The trajectory of animation describes the evolution of a visual property (e.g. position, shape, size, color, etc.) and represents the list of values that this property takes during animation. While the concept of trajectory generally evokes movement, i.e. the displacement of a visual element in the physical space of the interface, mainly the 2D plane, it also describes how each visual variable evolves in its own definition space, such as the color space, the size space, the texture space, and so on [Sta90]. Thus, a trajectory can describe in a more general way the transformations that a visual element can undergo in terms of a change of its position or its appearance over time with reference to a fixed reference [Sta90; STV06]. It can be described by a list of coordinates of the visual property in the corresponding space of values. There are two types of trajectory: *linear* trajectory and *non linear* trajectory that we will qualify as advanced trajectory in the remainder of this chapter. Linear trajectory represents the simplest and most standard way to interpolate between two

given values, and takes the form of a straight line in the definition space (e.g. position space, color space, etc.), whereas the so-called advanced trajectory corresponds to more particular forms of trajectory (e.g. curvilinear, circular) and results in a non linear interpolation.

The second and third questions (see ii and iii above) concern the *temporal parameters* of animation that include: the start time, duration and pacing function [HS93]. The pacing function describes the rhythm of evolution of a visual property and represents the speed at which the property changes from one value to another during the animation. Various alternative of pacing functions have been used in both literature and practice including for example the constant speed, slow in-slow (also known as ease in-ease out), and fast in-fast out.

4.1.2 Elementary and composite animation

With reference to our simplified model proposed in Section 3.1, animation concerns the visual component of a dynamic visualization. Hence, it can be seen from different perspectives according to the different levels of this component, namely: visual variable, visual object and charting component. Within each of these levels, we can work with two levels of granularity for animation: *elementary* and *composite*.

Depending on the level at which we decide to work, an elementary animation can be an animation of a visual variable, a visual object or a charting component of the representation space. A composite animation is the composition of many elementary animations. We distinguish two categories of composite animation: *homogeneous* and *hybrid*. An homogeneous composite animation involves the same type of elementary animations: elementary visual variable animation, elementary visual object animation, or elementary charting component animation. An hybrid composite animation combines elementary animations of visual objects and charting components

To better understand these two categories, let us have a look at the Figure 4.1. It shows a scatterplot-like visualization composed of two axes (charting components), four dots (visual objects) characterized with (x,y) position, shape, size, and color (visual variables). Let us now imagine some scenarios of dynamic changes that these different visual components can undergo and think about different ways to animate these changes.




Figure 4.1: A scatterplot-like visualization where both the position, size and color of dots can evolve over time.

Let us take for example the red dot and suppose that during a dynamic change of the visualization between two timesteps, this dot moved to the top left corner, got bigger in size and its color changed to blue. As animation designers, we can choose to animate these transformations using three sequential elementary animations, like for example: position animation, followed by size animation and finally color animation. We may also opt for playing all these elementary animations simultaneously. In both cases, since all these elementary animations have the same type (i.e. visual variable animation), we designate this composite animation as homogeneous. In addition, since all these elementary animations concern the visual variables of the same visual object (our initially red dot), we call this instance of composition as *intra-visual object* composition.

Now, what if we have three dots (e.g. the green, blue and gray) that got relocated and changed to a smaller size? We can use different ways of *inter-visual object* composition to generate an homogeneous composite animation synchronizing the changes of all these dots. The simplest case is to animate both the position and size changes of all the three dots at the same time. We can also choose to animate dots in sequence, like for instance: animate the blue, gray and finally the green dot with the position and size of each dot getting animated simultaneously. Another possible scenario could be to animate first the position of all the three dots, and then animate their size.

Finally, let us imagine that the axes of the visualization got rescaled from linear to logarithmic scale. Following this rescale, all the four dots should be relocated. How can we animate these changes? We can first animate the axes' rescale and then shift the positions of dots. Alternately, to portray finer details of change, we can first fade out the old axes labels, rearrange the axes ticks following a log scale, fade in the new axes labels, and finally relocate the four dots. All these examples represent instances of an hybrid composite animation.

Let us recapitulate and formalize. Intra-visual object composition generates an homogeneous composite animation involving two or several visual variables of the same visual object. Inter-visual objects composition generates an homogeneous composite animation involving two or several visual objects. These two concepts were inspired from the concepts of intra-object and inter-object animations proposed by Mirlacher et al. [MPB12]. Intra-charting components composition generates an homogeneous composite animation involving two or many visual variables of the same charting component. Inter-charting component composition generates an homogeneous composite animation involving two or several charting components. There are two approaches of animation composition: composition based on trajectory, and composition based on temporal parameters—i.e. temporal composition. A composite animation can be generated by applying one of these two approaches or by combining them.

All the scenarios that we described above represent only a subset of the possible composition alternatives that we can generate within this very simple four-dot visualization. So, imagine how this composition will become incrementally complicated as we start dealing with more complex visualizations and more diverse dynamic changes. Thus the need to better understand the aspects of animation composition in the context of information visualizations.

4.1.3 Composition based on trajectory

Composition based on trajectory consists of manipulating the trajectories of visual variables characterizing the visual elements (i.e. visual objects or charting components). The composition based on trajectory can be seen as a way of structuring the visual scene of a dynamic visualization in terms of the spatial arrangement between the visual elements and their relative visual aspect (e.g. their size, color, shape, etc.). It defines the relationship between the individual changes of position and aspect of the visual elements during the animation.

The general concept of structuring the evolution of position and aspect of visual elements composing a dynamic visual scene have been studied in various application domains including for example the multimedia application, dance choreography, and animation.

In the context of multimedia applications, the composition has been exclusively based on the position and focused on the spatial ordering and overlapping features of the multimedia objects (e.g. images, videos, text, etc.) [VTS96]. Spatial composition aims at representing three aspects: the topological relationships between the objects [EF91] (e.g. disjoint, meet, overlap, equal, etc.), the directional relationships between the objects [PT97] (e.g. left, right, above, above-left etc.) and the distance characteristics between the objects (e.g. outside 5cm, inside 2cm etc.).

In dance practice, the “composition in space” involves not only the spatial configuration of dancers and the way how their movements are organized and structured, but also their postures (stances) and gestures on the stage [Cal+89].

Most relevant to this thesis is the domain of animation. The composition based on animation trajectory has consisted majorly of manipulating the trajectories of position. Notable examples include the trajectory bundling by Du et al. [Du+15] that consists of grouping the position trajectories of visual objects during animated transition based on their spatial proximity. Another relevant example is the differentiated transitions in Mobiliste [Sch+06], where the position trajectories of animated items are modified relatively to each others to avoid crossing and overlap. Both these instances of trajectory-based composition employ curved spatial paths (i.e. advanced trajectories) for the same purpose of reducing visual clutter and facilitating visual tracking. Another interesting approach of composition consists of the rigid motion animation that is generated through the combination of the following animated affine transformations: translation, rotation, scaling and shearing. As we previously explained in Section 2.1.3, rigid motion is used mainly to maintain the same global shape the animated visualization and minimize the overall movement of visual items during the animation [FE02; EDF08a; HTC09].

Beyond research literature, we can also find in practice some good illustrations of composition based on trajectories. For instance, the Google Material Design recommends to animation designers in one of the motion choreography guidelines [Cho]: “When multiple elements remain visible during a transition, only the most important ones should be included. Some elements may disappear during the transition but reappear once the transition completes, if they are too distracting during the transition itself.” This example illustrates an interesting approach of composition that employs differently the position trajectories.

4.1.4 Temporal composition

Temporal composition consists of manipulating the temporal parameters of animation (i.e. start time, duration and pacing function) to generate a composite animation. It allows to define the triggering and finishing order of the elementary animations as well as their relative pacing.

Similarly to the composition based on trajectory, temporal composition has been studied in various domains of application. In multimedia applications, temporal composition concerns the temporal relationships between the various events happening in a multimedia application. The Allen’s algebra of temporal intervals [All84] provides a descriptive framework of the possible temporal relationships that can exist between two time intervals including for instance: equal, overlap, start, during, finish, before and meet (and their inverse).

In the dance choreography, the temporal composition (i.e. composition in time) allows to sequence the stances (i.e. position and postures) of dancers across the beats (i.e. time units composing a scene of dance).

Temporal composition for animation was early discussed in the pioneer work of path-transition paradigm by Stasko [Sta90] where the author introduced 3 ways of generating a composite transition namely: concatenate (i.e. execute multiple animated transitions sequentially), iterate (i.e. execute the same animated transition repeatedly) and compose (i.e. execute multiple animated transitions in parallel). Later works focused on studying particular instances of temporal composition including mainly staged and staggered animation [Pla+12; HR07; BPF14; CDF14]. Mirlacher et al. [MPB12] briefly discussed the temporal composition of inter-object and intra-object animations from the perspective of user interface engineering.

4.2 Designing animation for information visualizations

When designing an animation to achieve a particular goal, the animation designer mainly needs to understand in depth the different animation parameters involved in its definition, and decide these parameters in order to fulfil the targeted goal in the most effective way.

4.2.1 Understanding animation parameters

In literature, most of the prior works study animation as a whole and mainly approach animation parameters as a set of variables or criteria instantiated with respect to certain constraints (e.g. context of use) to meet a precise goal. Hence, they do not discuss in depth which animation's goals can be supported with which design instances of each animation parameter. Although some previous research works focused on studying particular animation parameters, they did explore these parameters within a narrow spectrum of context of use and to achieve only a subset of the possible goals that animation parameters can support. In addition, most of these works adopt a single design approach and do not discuss the spectrum of design alternatives of animation parameters. Such narrowed scope do not allow us to gain a holistic knowledge about these parameters that we need in order to be able to use animation in a wider range of contexts and apply it to achieve more various goals. We argue that a proper understanding of animation parameters requires discerning all the different goals they can support, and gaining an informative overview of their various design alternatives.

Understanding animation trajectories

When surveying the literature in the chase of gaining a sufficient understanding of animation trajectories, the animation can encounter the following issues:

Focus on a subset of visual variables

Although some works in literature adopted the universal definition of movement or motion as qualifying the change of any visual property [STV06; Sta90; Bar97a], movement in its more restrictive and exclusive definition referring to the change of position has been so far the most predominately studied and the more deeply understood instance of animation trajectory [Du+15; WL94; Sch+06; SY06], compared to all the rest of the rich set of visual properties: size [Sch+06], shape [Sch+06], luminance [SB01a], and texture [Rom+18]. Even if these predominant design choices prove to be reasonable and appropriate in the specific contexts in which they were adopted, such a conventional approach ignores a rich space for alternatives, which still remains under-explored.

Focus on a subset of animation goals

Previous works focused on a limited subset of the possible animation goals that animation trajectories can support, and were concentrated mainly around supporting visual tracking. In this perspective, Du et al. [Du+15] studied the impact of a particular instance of position trajectory, namely, bundled trajectory on the effectiveness of animation for visual tracking.

Some works focused on studying the limitations of position trajectories for visual tracking. For example, Franconeri et al. [FJS10] proved that the visual tracking performance depends only on the spatial spacing of visual objects and concluded that as long as the distribution of object spacing is maintained constant, people could reliably track an unlimited number of objects as fast as they could track a single object, while Sukanuma and Yokosawa demonstrated that the tracking accuracy is impaired when visual objects of interest share similar

position trajectories with distractors [SY06]. Only few works explored different goals of animation trajectory beyond visual tracking, and studied for instance how different types of animation trajectory can be used to encode information and convey a meaning (e.g. position [WL94]; position, size and shape [Sch+06], texture [Rom+18]), or express emotions (e.g. position [CSR12]).

Limited exploration of design alternatives

The trajectory of animation can be mainly characterized by its form and have either simple rectilinear pattern (linear trajectory) or have a more sophisticated forms that can be for instance curvilinear or circular (advanced trajectory). By varying the form of the trajectory of values, we can obtain a wide range of alternatives for designing animation. However, in practice, a review of the literature indicates that the design approaches have been majorly reduced to linear trajectories, while very few and narrowed design alternatives applying advanced trajectories have been so far explored.

The exploration of advanced trajectories of values has been aimed mainly at enhancing the readability of animated interfaces and better supporting visual tracking. The most relevant examples include the use of 3D rotations to generate a more intuitive and meaningful animation in multidimensional data visualizations [FFT88; RMC91a; EDF08b; HTC09], the use of arc trajectories to generate a smoother and less confusing animation and reduce visual clutter [Las87; Yee+01], as well as the use of curved trajectories to enhance the visual tracking performance [Du+15; Sch+06; DHC11].

Some few works explored the potential of advanced trajectories to improve the expressivity of the interfaces such as using advanced trajectories to encode information [Sch+06; Sir00] or imitate the realistic behaviors of solid objects in user interfaces [CU93]. Observing all these different works unveils the fact that the exploration of advanced trajectories concerned mainly the trajectory of a single visual property, not surprisingly, position.

Understanding the temporal parameters of animation

The temporal parameters of animation concern generally the way how the animation is temporally structured and the pace following which it progresses. Temporal parameters are: start time, duration and pacing function. Varying the start time and duration generates two major techniques of temporal composition, namely: staging and staggering which is a particular class of staging. As for trajectories, temporal parameters remain still not well studied in literature and animation designers may encounter mainly the following issues when attempting to understand them:

Focus on a subset of animation goals

The animation duration have been mainly explored within the goal of enhancing the readability [BB99], and visual comfort of animation [DHC11], or facilitating visual tracking [HR07]. Similarly, the animation staging have been mainly applied to enhance the visual tracking during the animation and help understanding the animated information [HR07; GVM12; Pla+12]. It has been narrowly explored as a narrative tool [Fra+14; Pla+12].

Similarly, a reading through the literature, reveals that animation pacing has not been sufficiently studied and understood. Dragicevic et al. [Dra+11] studied the impact of various techniques of animation pacing (constant speed, slow-in/slow-out, fast-in/fast-out, adaptive speed) on the visual tracking, and found that slow in/slow out was the most effective technique supporting visual tracking, while Feria [Fer12] demonstrated that increasing the speed of

movement of visual objects hinders the tracking accuracy and decreases the maximum number of targets people can track [Fer13].

Limited exploration of design alternatives

Most of the previous works adopted constant duration of animation in their interfaces, while some other suggested adapting the duration in accordance to various criteria including for example the quantity of change during the animation [Che+10], or the degree of importance or interest of a given animation [Pla+12; GVM12].

Prior works studying staged animation applied a very narrow range of design alternatives consisting majorly of the “disappear, move, appear” general approach and did not discuss in depth the design space of staged animation [BPF14; ZKS11; PGB02; HR07; FE02].

The exploration of advanced pacing curves in literature consisted many of applying the slow in/slow out pacing mainly to provide aesthetic fluidity [Las87], facilitate the visual tracking [DHC11; Pla+12] and improve spatial and temporal predictability [HR07].

After exploring and analyzing the knowledge that prior works have provided to help an animation designer understand the different animation parameters, we realized that we still need to gain a deeper understanding of animation parameters and explore how these parameters can be leveraged to achieve the different animation goals by discussing and studying various design alternatives. Once a designer assimilates the different design parameters that she should take into consideration and adjust to create an animation, she needs now to decide how to define these different parameters in the most effective way.

4.2.2 Defining animation parameters

Once an animation designer gains a good understanding of animation parameters and gets an overview of the various alternatives among which she can chose to define an animation, she needs now to decide what precise design approach to pick from this wide range of alternatives. This design decision can be informed and guided with a set of design guidelines and recommendations.

Even though we can find in the literature a wide set of various design guidelines for animation, we observed some limitations that do not help an animation designer get sufficiently informative and useful instructions about how to decide animation parameters. The main limitations can be classified as follows:

Guidelines focus on specific animation’s goals or parameters

Some guidelines are dedicated to achieve specific animation goals or focus on instructing the design of specific animation parameters.

At the very early stage of using animation in the context of user interfaces, Lasseter suggested adapting and applying the twelve principles of cartoon animation in the context of 3D computer animation. His pioneer work inspired many works that readapted some of these principles to be used in the context of user interfaces. For example, Chang and Ungar suggested three main principles of animation in user interfaces: solidity, exaggeration, reinforcement [CU93]. These guidelines are aimed essentially at conveying realism and making the interfaces more lively by giving indication about the speed, material and weight of some elements on the screen. Building upon these principles, Thomas [TC01] suggested for principles of animation: attachment, reluctance, smoothness, anticipation that were mainly targeted to reinforce direct manipulation in user interfaces. Similarly, to guide the creation of animated presentations, Zongker provided a set of guidelines that can be summarized as follows: use parameterization, treat animations as models, build slides hierarchically [ZS03].

However, in practice, most of the animation goals as discussed in the Chapter 2 go often beyond the perspective of increasing realism or simulating reality in the interfaces. As argued by Liddle [Lid16], focusing design guidelines on reality offers a limited perspective of the communicative possibilities of animation. Hence, these design guidelines can not be used to inform the design of animation to achieve different goals other than the particular goal of realism. Some guidelines, despite mostly generic, made particular focus on some different animation goals such as visual tracking [GVM12; BPF14], visual aesthetic [FE00], data encoding [Sch+06] and narration [Pla+12].

While all the above guidelines focused on a subset of animation goals, other ones were concentrated rather on a subset of animation parameters. In the context of information visualization, Heer and Robertson [HR07] suggested a set of design recommendations with precise design instructions for specific animation parameters including for example using slow in/slow out to maximize predictability, using translation and divergence (expand/contract) motions instead of rotation to generate animations that are simpler to understand as well as using staging for complex transitions. However, most of the other design recommendations that they suggested (e.g. respect semantic correspondence, avoid ambiguity, group similar transitions) do not provide hints about what precise design decisions to take for the rest of animation parameters. Some other works focused on specific parameters such as trajectories [Sch+06; FE02] and staging [BPF14; Pla+12; FE02; GVM12]

Specific guidelines generated from general studies

Many general guidelines that we can find in literature are generated from empirical explorations studying animation as a whole. Hence, even if they provide instructions about particular animation parameters, these instructions can be arguable since we are not sure to which extent the conclusions about separate animation parameters that were not studied in isolation can be valid. However, narrow focus of design guidelines does not seem the only problems that an animation designer can encounter. Overly general guidelines can also be problematic.

Guidelines are overly generic

Many of the design guidelines that we can find in literature are very general and high level. Relevant examples include the congruence and apprehension principles [TMB02a], where congruence recommends matching of the animation's content and format to that of the information to be conveyed, and apprehension requires the ease and accuracy of perceiving the information presented by the animation.

Similarly, Stasko [Sta93] suggested four design guidelines inspired from the principles of traditional animation. Appropriateness dictates that the operation or process should be represented according to the user's mental model and system entities. Smoothness dictates avoiding jerky and wildly varying animations because they are difficult to follow. Duration and control dictates varying the duration and control with the type of animation. Finally, demonstrations dictates that the demonstration of unit operations such as selection should be short (not more than a few seconds). Stasko also recommended that animating continuous processes with a clocktime correspondence should be kept faithful to the clocktime. When animation is used as explanation, he recommended that the user should be allowed to control the rate and replay. Finally, moderation prescribes judicious application of animation: "too much is overdone and too cute".

These design guidelines do not explain concretely how to translate their requirements into design decisions and they do not provide precise instructions about how to specifically choose the different parameters of animation to achieve a particular goal.

Guidelines are analogous and intersecting, yet still limited

When exploring design guidelines, we were intrigued by the high interference and redundancy between them. On one side, we recognize that iterating and building upon prior work is a relevant research approach. But on the other side, we believe that it is more relevant to remedy to the limitations that we still observe in the design guidelines of animation rather than extensively re-studying the current guidelines. Relevant examples include analogies between congruence [TMB02a] and appropriateness [Sta93], non-ambiguity [Che+10] and avoid ambiguity [HR07], as well as between apprehension [TMB02a] and smoothness associated with duration & control [Sta93].

Some guidelines are contradictory

Deciding animation parameters could become much more challenging and complicated when an animation designer faces contradictory guidelines. Relevant examples include the arc trajectories that most of the guidelines recommend applying [Bab17; Arc] whereas some others recommend avoiding [Wil17].

4.3 Evaluating animation in information visualizations

Animation design gets inspiration from various fields (e.g. psychology, visual perception, cognitive science), and its use is generalized across many application domains (e.g. graphical user interfaces, information visualizations, games, etc.). Hence, it seems logical and necessary that the evaluation of animation should be done considering different perspectives and should be adapted to the various goals that it can achieve within the different domains of application. That implies for example applying different evaluation approaches and metrics and using different tasks according to the goal of animation and the context in which it is intended to be used. However, our reading through the evaluation methodology that has been adopted to assess the value of animation, revealed many gaps with reference to these requirements.

In this section, we will start by providing a structured view on the set of metrics that can be used to measure the value of animation. We then discuss the general methods that have been adopted to assess the value of animation in user interfaces, point out their limitations, and suggest alternative evaluation approaches.

4.3.1 Metrics to measure the value of animation

Animation can be seen as any other feature of a user interface and hence its value can be assessed using the usability metrics allowing to measure to which extent an animation is effective with reference to the role it is intended to fulfill within a user interface. Beyond this general usability metrics, we need various metrics that allow to evaluate the animation from different perspectives and that can be leveraged to assess its value for specific goals or contexts of use. Such metrics include mainly: comfort, perception, cognition and subjective quality metrics.

Usability metrics

Frøkjær et al. [FHH00] defined three main metrics that can be used to measure the usability of a user interface: effectiveness, efficiency, and satisfaction. In the context of this section, we classify the satisfaction under an independent category of evaluation metrics, that we designate with *subjective quality*. Hence, we will consider only the effectiveness and efficiency metrics under the usability category and we will discuss later the satisfaction metric.

When applied to animation, effectiveness is the accuracy and completeness with which users achieve certain goals using an animated interface. Indicators of effectiveness include quality of solution (or more generally the quality of outcome of user's interaction with the animated interface) and error rates. It is important to analyze not only the rate of errors but also the kind of errors [SP98; BB99]. The dropout rates can be also used to measure how effective is an animation in enhancing user experience.

When applied to animation, efficiency is the relation between (1) the accuracy and completeness with which users achieve certain goals and (2) the resources expended in achieving them using an animated interface. Indicators of efficiency include task completion time and learning time. The completion time indicates how long does it take to carry on tasks using the animated interface. It can also give us an idea about the speed of users' performance in performing tasks with the animated interface.

These generic metrics can be adapted to assess the value of animation to achieve the different animation goals that we presented in the Chapter 2. For example, to evaluate how animation is efficient for supporting navigation, we can use the metric of navigation speed to measure how fast they can return to where they have been before during the navigation [BB99]. As Bederson recommended [BB99], experimenters should carefully measure the time spent on animation with reference to the overall completion time to gain a clear understanding of the impact that animation had on the interaction time which represent an important factor if we refer back to our classification of animation drawbacks 2.2.

The learning time indicates how long does it take for users to learn how perform particular tasks within the animated interface. With reference to our taxonomy of animation roles, the learning time should be used as a primary indicator to measure the efficiency of animation in achieving the goals that are tightly related to learning a user interface such as: affordance and preview, training and demonstration by example, and teaching a new representation of information.

Furthermore, the number of times the participant has to replay the animation (one or the other way) is also a good indicator of the effectiveness of the animation. Arguably, the more the replays, the least the animation is capable of conveying changes. Similarly, to evaluate impact on user experience studies can measure viewing time.

Comfort metrics

Comfort metrics allow to measure the impact of animation on users' visual and physical comfort and can mainly include: visibility, readability, physical demand and effort, dizziness, nausea, vertigo, and headache. All these different metrics can be used in subjective satisfaction questionnaires to measure how comfortable were users when seeing the animation. But as Tversky argued: "seeing is not perceiving" [Tve+07], which means that the fact that users were able to see the animation in a comfortable and appropriate way does not imply that they were able to correctly perceive it and extract the information that it conveys. Hence, in addition to comfort metrics, we do need perception metrics.

Perception metrics

Perception metrics allow to measure the impact of animation on user's perceptual system and mainly include: visual detection, visual identification, visual distraction, visual tracking performance, disorientation, change blindness, and visual salience. For example, for affordance and demonstration by example, user studies can verify that animations improve discoverability of a feature and/or its execution [Wat+10].

The perception metrics allow to measure how well users were able to perceive the animated content and detect any eventual negative perceptual impact that an animation can have on

users. But, again as Tversky argued: “perceiving is not understanding” and “showing is not explaining” [Tve+07], which means that the fact that users were able to perceive and visually follow an animation does not guarantee that they were able to understand the information conveyed by it and extract an appropriate knowledge from it. Hence, we need cognition metrics to assess the cognitive impact of animation on users.

Cognition metrics

Cognition metrics allow to measure the impact of animation on the user’s cognitive system and mainly include: cognitive overload, mental demand, decision quality, comprehensibility, learning, memorability (recall), impact on spatial memory (reconstruction ability [BB99], cognitive map), visual momentum. For example, for affordance and demonstration by example, user studies can verify if new knowledge that users learn about an interface’s features persist over time. Memorability could be a useful metric when dealing with animation for persuasion. The visual momentum can be used to measure how visual scene transitions impact the users ability in extracting and integrating information across multiple displays. Studies on maintaining up to date and replaying histories can focus on errors being made during information recall.

Subjective quality

Subjective metrics allow to measure the impact of animation on user’s subjective attitude and mainly include: satisfaction, abandonment (i.e. do user stay and even click on something or do they leave?), interest, frustration, annoyance, deception, and boredom.

Users’ satisfaction can be measured by attitude rating scales such as SUMI [KC93], and subjective preference. Subjective metrics can be used to evaluate the impact of data encodings. Questionnaires can also investigate social perceptions, asking participants to describe impressions and positive/negative sentiments generated by the animated objects.

4.3.2 Evaluation methodology

In this section, we will discuss the evaluation methodology adopted in the literature to assess the value of animation.

Studied animation goals

Empirical studies in literature have focused only on a subset from the wide range of animation goals that we presented in Chapter 2. Several animation roles remain still over-looked and narrowly understood while some others have been extensively explored. Indeed, although animation is largely used in practice to improve user experience, this category of animation roles has been very barely studied in research. Aesthetics is an extremely challenging and over-looked topic in HCI. Research on how to improve quality and better evaluation methodologies will be instrumental in moving the field forward. Animation for data encoding has not been in its turn sufficiently explored. Similarly, since visual discourse and storytelling are relatively recent style of interfaces in HCI [KM13], there is limited work on evaluation.

Conversely, under the keeping in context category, the general animation goals of “supporting visual tracking” and “maintaining orientation during navigation” have been extensively studied in diverse contexts of use, while for instance the “replaying history, summarizing” remains still over-looked. Similarly, animation as teaching aid has been in general largely explored, but the “affordance and preview” goal is still under-explored while it represents a promising research perspective with its wide spread use by UX designers for web and mobile

user interfaces. Similarly, the “illustrating an algorithm” goal has emerged as a relevant research opportunity mainly because more works are focusing on novel approaches to visualize the behavior of machine learning and artificial intelligence algorithms.

Evaluation methodology

Animation has mainly been evaluated based on task-based quantitative studies. Alternative quantitative approaches can be more adapted for certain animation’s goals. For example, to evaluate the impact of animation on user experience, studies can measure interactions at the first encounter, during or immediately after the animation, using click log or tracking mouse and eye movement. Moreover, A/B testing can facilitate incremental improvements.

We recognize the importance of assessing the value of animation from quantitative perspectives but we argue that it does not cover sufficiently the rich spectrum of the different merits that an animation can have within a user interface. For example, for animation aiming at supporting tracking during layout changes, qualitative studies asking users to describe what happened to target objects in real scenarios can provide important insights. For animation to enhance user experience, more qualitative studies are required for assessing visual comfort and aesthetics, or evaluating virtual tours or making background activities visible.

For animation to provide visual comfort and aesthetics, studies assessing first impressions of infographics [HRC15] might provide some inspiration. Collaborative research with web and graphic design practitioners is needed. Studies on maintaining up to date and replaying histories can capture how many insights users perceived.

General user study design considerations

There are two challenging aspects that we can encounter when designing a user study to evaluate an animation that we will designate respectively with: i) isolation of animation parameters’ effects, and ii) equivalence between tested conditions. These two aspects will decide both the validity and the generalizability of results issued from a given user study.

Isolation of animation parameters’ effects

As [APP11] argued: “All experiments are limited by their parameters, and empirical results are generalizable only within the scope of those parameters.”. Animation evaluation, as the evaluation of any other feature in a user interface, follows this rule. Hence, it is important to recognize what are the findings of a given empirical study that can be generalized, and to which extent this generalization can be possible.

Animation is a combination of multiple parameters and each separate parameters can have a considerable impact of its effectiveness and performance. When evaluating animation as a whole (instead of focusing on varying a particular animation parameters), it is not trivial to isolate the effect that each parameter had on the performance of animation. Hence, findings and conclusions issued from the empirical study remain valid only within the scope of very similar instances of animation since the change of one or several parameters can dramatically change the result (straight path vs. curved path may dramatically improve readability during the animation). In this perspective, Jones [JS00] argued that research findings about the effectiveness of animation for learning are varied and inconsistent partly due to the large number of variables involved in the design of animation that generates inconclusive results and induces high variances across studies.

Equivalence between tested conditions

The equivalence between tested conditions is mainly discussed in the case of comparison between an animated versus a static interface. Comparison between animation and a static counterpart is a commonly applied evaluation approach. For example, controlled experiments studying the effectiveness of animation for visual tracking generally compare the animation condition to a static control condition. In this context, in the case of staged animation for instance, the static counterpart is to show the exact same number of stages through a comic strip, or a slide show.

Similarly, to evaluate teaching aids, a common study design is to compare conditions with and without animation. For visual discourse, controlled studies comparing versions with and without animation can reveal aspects such as comprehension, attention or engagement, assessing if people changed their behavior by designing specific pre-test and post-test.

When evaluating an animated interface versus a static interface, the main goal is to evaluate the benefits of an instance of animation as compared to its static counterpart. Fair comparison between the two conditions requires ensuring that they are equivalent from different perspectives. Tversky [TMB02a] argued that the major problem in several empirical studies evaluating animation for learning is the incomparable content between animated and static graphics.

In the context, Castro-Alonso et al. [CAAP16] introduced many different kinds of biases that can occur when evaluating animated versus static visualizations to measure the effectiveness of animation as a learning aid. We will present below the most relevant ones to the context of this thesis:

Appeal bias means that animated and static visualizations have different quality of visual elements. An appeal bias is observed when certain attractive features are included in one of the visualizations only, generally in the animations. For example, in static visualization, elements are in black and white and in the animated visualization, elements are colored.

Variety bias means that animated and static visualizations have different quantity of visual elements (e.g. arrows vs. no arrows). This bias was also discussed by Tversky et al. [TMB02a] who highlighted the fact that many empirical studies comparing static and animated visualizations were biased due to non equivalence in the quantity of information.

Number bias A number bias is generated when one of the formats is presented in more images than the other. This is typical when many simultaneous statics images are compared to a single animation.

Size bias A size bias is observed when one of the depictions is noticeably larger than the other. This can cause the larger visualization to be easier to study than the smaller one, as evidence can be found showing that size variations leads to differences in perceptibility of the elements presented [SL08].

Interaction bias An interaction bias is produced when the images are not matched in their features of interactivity. As observed by [TMB02a], generally animations give more interaction to learners than static pictures.

More generally, the most challenging aspect when comparing static and animated interfaces is deciding the equivalence between the two interfaces because it is not always trivial to find the correspondence between parts. For example, what could be the static counterpart of visual emphasis effects such as color highlighting, fade in/out, or afterglow effects?

Independent variables

The choice of independent variables should be mainly in accordance with the animation goal. For animation for keeping in context, many variables are important to study including for example the animation speed, number of stages in a staged animation and their order. Other independent variables can include: level of difficulty (the amount of changes occurring in the view, for example, the number of elements to track, the number of steps to look at). We argue that it is essential to test animation for different levels of difficulty, so as to better understand the break points of one or the other instance (static vs. animation), and/or animation parameters (straight path, for example, may be as beneficial as curved path when no much changes occur in the view). The variety of examples, and justification of why such examples are challenging for the techniques will make evaluations much more informative. For visual discourse, many parameters come into play such as number of steps in a narrative, the timing, the mapping of the animation with the story content (e.g. semantically relevant choice of pictograms and motions). Much research is needed to study the impact of each variable.

Dependent variables

Most of the empirical studies in literature define the error and time as the main dependent variables. Additional dependent variables can include subjective satisfaction. The answer can be ascertained by interview or by written surveys that include satisfaction scales (e.g. Likert scale) and space for free-form comments. In general, questionnaires [HSA08] and interviews help tease apart subjective impressions and identify possible causes of problems. For example, in the work by Han et al. [Han:11] participants were presented with cancer risk estimates using representations varying in how the randomness is expressed, then asked to rate their perceived cancer risk, level of uncertainty, and worry.

Subjective satisfaction questionnaires can capture possible drawbacks (e.g. dizziness, distraction, boredom, etc.). Questionnaires can measure subjective ratings about the perceived value of the animation.

Tasks

Low level tasks used in user studies include essentially general multi-object tracking (MOT) tasks. For example, to measure how animations support keeping track of changes, controlled experiments with motion object tracking (MOT) tasks can point at a set of objects before the animation, and ask users to locate these objects at the end of the animation. Other studies use revisitation tasks. For example, to evaluate navigation support, controlled experiments using revisitation tasks are useful to capture if users know where they are in the information space. However, revisitation and MOT tasks are most often used in a target-agnostic context (e.g. without visually highlighting objects of interests), which may not be representative of reality, so evaluations of the role and true value of animations in more realistic scenarios is needed.

Other tasks include spotting context-specific changes (e.g. changes between text document revisions, changes in a dynamic network) that are especially used for animation to keep in context. Less commonly, problem solving tasks are used in some studies [Pla+12]. Several prior works insisted on the need to use higher-level tasks and real world scenario for more informative evaluations [BPF14].

High level tasks can include questions that require analytical work and reasoning and focus on asking for insights based on open-ended questions rather than precise answers. They can also include knowledge discovery tasks [DiB+92; BBL12] While using such tasks allows a more realistic and hence informative assessment of animation's effectiveness, it does also involve some challenges. Indeed, giving higher level tasks to the users can induce additional

cognitive load on them. In addition, insight-based evaluation remains among the most arguable and challenging evaluation approaches [Nor06].

Running a study

To evaluate teaching aids, a common study design is to compare conditions with and without animations, using a between-subject design (where one group sees the animation and the other one not) [TMB02a]. These studies often administrate pre-test (to capture pre-existing knowledge) and post-test, featuring comprehension questions. Questionnaires [HSA08] and interviews help tease apart subjective impressions and identify possible causes of problems. Questionnaires can also investigate social perceptions, asking participants to describe impressions and positive/negative sentiments generated by the animated objects.

4.4 People and animation: insights for animation design and evaluation

During our preliminary exploration described in details in the Section 2.1.1, we gained valuable insights that we believe can inform the design and evaluation of animation for information visualizations. We classified these insights as follows:

What guides or constrains the design of animation by people

When designing their own animation, participants reported being essentially informed and constrained by i) the context of use (e.g. mobile vs. desktop application, touch vs. mouse input, professional vs. personal use) and ii) the type of the animated content (e.g. image vs. text, landscape photos vs. persons' photos). For example, one participant explained that when she designs animation for personal use, she uses a lot of animation instances and more funny effects. However, when she creates animated presentations for professional purposes, she uses fewer animations and more sober and serious effects.

Semantics in animation designed by people

While some participants reported basing the design of their animation on personal subjective preferences, many others mentioned expressing certain semantics when defining the different parameters of animation, mainly: type of animation, trajectories and temporal parameters.

Type of animation

We will enumerate below the main rationales that participants advanced about their choice of animation types.

- *One Section-One type of transition*: "I use the same type of transition between slides in the same section. When I change section, I change the type of transition. It helps me to inform people that I'm going from one level to another in my presentation. It also helps people to navigate through the organization of my presentation."
- *Enforce link between slides*: "I only use transitions when there is a relationship (semantic link) between the slides. Transitions can have a semantic value (e.g. explosion, drop (make it fall from the top) the slide) that could help communicate a certain idea".
- *Break link between slides*: "Sometimes, I want to break the link between two slides and emphasize the fact that these two slides talk about two completely different things. In this context, the transition serves to me as the "dissolve" effect used in the film making."

- “I choose the animation style depending on the tone that I need to convey”
- “I Used to use fade in/out but my client requested the panels to fly into the screen because they wanted it to feel a bit more concrete and not as if it comes out of the air”
- “Animation styles can be sometimes associated with some cultural meanings specifically in movies. Hence, some styles can create nostalgia as they are references to some movies for the past.”

Trajectory

We present below some details on how participants reported choosing their animation trajectory.

- “In the user interface of my experiment, I used an animation of the “X” icon which means that there is an error. The icon is animated by a movement of right-left and left-right to mimic the movement of the head when expressing negation or denial. For me it does not make sense to animate the “X” from up-down and down-up (as if you say no and you move your head by an affirmation)”.
- “I animated the appearance of the arrows of a diagram from left to right to express a chronological sequence.”

Timing parameters

We present below some details on how participants reported choosing their temporal parameters.

- *Delay*: “I insert a delay after each question and answer to give the impression of a “Question-Answer” situation. I used a longer delay in the last question with a “Yes” answer to add a “drama and suspense effect” and emphasize more the last idea because it is the most important to me”
- *Pacing*: “I used the ease-in/out to make the animated items seem as if they are pulled by hand.”
- *Pacing*: “To my opinion, animation easing functions are just like fonts. Changing the easing changes the style of your presentation to be more or less “serious”. Some easing function give more the effect of cartoon animations. The linear and ease in/out pacing communicate a “serious” meaning of the animation (not a “cartoonish” meaning).”

A reading through all these findings highlights that most of our participants, despite not being expert in animation, are aware of the semantic and expressive potential of animation and they do exploit it when designing their own animation. That makes us further criticize the fact that these aspects are still over-looked in the literature, and insist on the need to study them more in depth in the future.

4.5 Conclusion

We explained the general fundamentals of animation and presented a theoretical formalization to characterize the composition of animation in dynamic information visualizations. Composing animation based on trajectories and temporal parameters is a complex and rich process that generates a wide spectrum of design alternatives that are yet to be deeply explored in the future.

The complexity of dynamic information visualizations and the richness of changes they can undergo, described in the previous chapter of this thesis, combined with the complexity of animation and its high functional and expressive power highlighted in our taxonomy of animation roles, raises important design and evaluation challenges but also generates numerous benefits and opportunities. When surveying the design and evaluation of animation in literature, we were intrigued by the fact that prior research has been investigating animation from a narrow perspective with reference to all the merits that it can have for information visualizations:

- Design spaces of animation parameters have not yet been deeply explored.
- Animation design for information visualization has not yet exploited all the visual variables that can considerably enrich the value of animation.
- The most relevant categories of animation goals for information visualization, namely: data encoding and visual discourse are still over-looked.
- Animation evaluation is still highly narrowed to the multiple object tracking (MOT) perspective. Let us refer back to the three animation aphorisms by Tversky: “Seeing is not perceiving, Perceiving is not understanding, Showing in not explaining” [Tve+07]. Tversky’s claim further convinces us that focusing only on visual tracking does not provide a complete view of both the benefits and drawbacks that animation can induce in visualizations.

Considering all these observations, we decided to focus the major work of this thesis on two main research directions. Our first research direction was founded on the two following motivations: exploring animation trajectories of under-studied visual variables, and assessing the effectiveness of these trajectories for animation goals beyond visual tracking. Hence, we decided to investigate the effectiveness of animation trajectories of position, size and luminance to encode visual grouping in dynamic visualizations (Chapter 5).

Our second research motivation was to explore in depth the temporal composition of animation; highlight the fact that beyond its prevalent use as a perceptual aid for visual tracking, it incarnates a great semantic and expressive potential to encode data and tell stories about it; shed light on the multiple empirical lenses through which we can evaluate temporal composition beyond MOT (Chapter 6).

Chapter 5

Reflecting on the meaning of Common Fate in visualization

The Gestalt Law of Common Fate (LCF) [Kof22] is an example of a widely known guideline for designing animations where visual elements that move with the same velocity (i.e. same speed and same direction) are perceived as sharing the same “fate”, and thus belong to the same group. The Law of Common Fate is also the only of the five Gestalt Laws that deals with dynamic (e.g. animated, time-changing) properties; the others all concern static instances of grouping in visual perception [Wag+12].

Although the Gestalt Laws—including LCF—were derived from perceptual psychology experiments in the early 1900s at the “Berlin School” of psychology [Kof22], only a few isolated examples of application to *dynamic* visualizations have been explored [Blo+99b; FH02; WB05]. This represents an opportunity for visualization research to delve deeper into human perception for the purposes of optimizing animated transitions. For example, better knowledge of the automatic grouping of animated objects may suggest ways to structure animated transitions so that their complexity is decreased and they become easier to perceive and overview. Furthermore, while most examples of LCF use visual elements with identical position trajectories, the philosophical meaning for a “common fate” of objects engaged in joint motion is not necessarily restricted to velocity [SB01b]. Rather, a general interpretation of “common fate” would merely imply shared dynamic behavior between multiple objects so that they are perceived to be under the influence of the same physical process [ABL98]. Such shared behaviors include growth and shrinkage (size) and darkening and brightening (luminance). Given this background, it is useful to ask ourselves how the visual grouping arising from common fate is influenced by such dynamic behaviors, and how these factors interact with each other when used in conjunction. Answering these questions may shed light on possible additional operators that can be used to add structure to animated transitions in interfaces.

In this chapter, we study these intricacies of the Gestalt Law of Common Fate by means of a large-scale crowdsourced graphical perception experiment involving more than 100 respondents performing perceptual grouping tasks. Our experiment was designed to compare three static visual factors (position, size, and luminance) and three dynamic visual factors (velocity, luminance change, and size change) in trials where four graphical objects were grouped by two properties at a time—two pairs were grouped based on one factor and two other pairs based on another. This enabled us to not only study the individual grouping strength of each visual factor but also to rank the factors in order of their relative grouping strength. We briefly discuss how these findings can be used to inform the design of animated transitions for the purposes of reducing cognitive load and improving user perception.

5.1 Background

Perceptual psychology has a long and distinguished history in the grand endeavor to understand the human brain, and its applications to graphical perception is of particular importance to visualization and statistical graphics. Here we give a general overview of relevant work in these research areas.

5.1.1 Perception and Gestalt psychology

Perception is defined as the innate sensory components of the human cognitive system that are pre-conscious and used to represent and understand the environment, and *visual perception* is the perceptual component dealing with sight. As the most important of the senses, the human visual system has evolved over millions of years to allow individuals to distinguish, identify, and track objects in their vision [Jer73]. Unlike the brain, which even with today’s technology is difficult to observe and decipher, the sensory organs are straightforward, which led to much of early cognitive psychology experiments focusing on perception. Visual perception is no exception.

Much of the seminal work on visual perception was conducted in the early 1900s by the so-called “Berlin School” of experimental psychology. This eventually led to the development of *Gestalt psychology* [Kof22], a theory of mind based on a holistic view of human visual perception where the sum of the perceived “gestalt” is qualitatively different than its component parts, and in effect has an identity of its own. One key practical outcome of Gestalt psychology was the development of the law of *prägnanz* (German, *pithiness*) [Kof22], which can be operationalized into the so-called “Gestalt laws” [Wag+12]: examples include the Law of Proximity, which states that objects at close distances are perceptually grouped, or the Law of Similarity, which states that objects with similar visual appearance are grouped together. Analogously, the Law of Common Fate—incidentally, the only Gestalt law dealing with dynamic settings—states that objects with the same movement are perceptually grouped together.

5.1.2 Gestalt Laws in visualization

The Gestalt Laws are an important component of visual perception that researchers have attempted to leverage for more efficient visual communication. Early work in cartographic animation assumes that common fate is more generally valid for objects that change together, such as blinking, even if no formal evaluation is reported [Blo+99b]. Ware and Bobrow used motion as a mechanism to highlight a subgraph of interest in a larger graph. While they initially found that motion dominates hue for highlighting [WB04a], their most recent studies suggest that motion and hue can be used in conjunction for the highlighting of two different entities simultaneously [WB05].

Friedrich and colleagues [FH02; NF02] successfully applied the law of common fate to make subgraphs apparent when transitioning from one layout to another to preserve the viewer’s mental map. The goal was to find an animation of the subgraph of interest that would be interpreted by the brain as movement of three-dimensional objects, using affine transform to decompose the motion into a series of translation, rotation, scaling, and shearing. They found that the law of common fate not only holds for objects moving in the same direction, but also for objects which move in any structured way.

5.1.3 Perception in visualization

Statisticians have long been interested in the perception of their charts. Examples of pioneering works on this topic include those of Eells [Eel26], Croxton [CS32; CS27], and Peterson et al. [PS54]. Recent studies have extended these works to interactive visualization. Representative examples include Haroz and Whitney’s work on attention capacity [HW12], and Harrison et al.’s crowdsourced ranking of visualizations based on correlation [Har+14]. Heer and Bostock [HB10] have also demonstrated that graphical perception studies can be crowdsourced for a significantly lower cost than laboratory experiments with a negligible loss of precision.

5.2 The Gestalt Law of Common Fate

The dominant interpretation of the Gestalt Law of Common Fate (LCF) is that the concept of “common fate” solely refers to the visual grouping of elements moving in a *coherent motion*, i.e. with the same speed and direction. One way to intuitively explain this phenomenon is that the moving objects that are visually grouped are seen as under the influence of a factor that causes them to move with the same trajectory.

However, this simplistic interpretation is not the only one. Admittedly, Wertheimer, one of the founders of Gestalt psychology, did use moving objects with identical velocity as an illustrating example in his original German manuscript [Wer38]. However, as noted by Sekuler and Bennett in 2001 [SB01b], he also included a passage on broader interpretations of the concept of common fate that never appeared in the English transcript: “*The principle [of common fate] applies to a wide range of conditions; how wide, is not discussed here.*”

Biased by the belief that Gestaltists only had motion in mind when developing LCF, subsequent studies in psychology have mostly focused on investigating the limits of figure-ground segmentation under variations of motion coherence [SS04; Utt+00; LB01]—a few exceptions include studies on dynamic luminance [ABL98; SB01b] and its application to cartographic animation [Blo+99b] and graph visualization [WB05; NF02]—which may explain why the simplified and incomplete version of the law has become prevalent. As stated by Brooks in a recent survey on perceptual grouping: “*Although common fate grouping is often considered to be very strong, to my knowledge, there are no quantitative comparisons of its strength with other grouping principles.*” [Bro14]

Given this background, we formulate two distinct research questions that we focus on in this work:

RQ1 *Does the Law of Common Fate extend to other dynamic visual variables, such as dynamic luminance or size?* We believe, as Wertheimer suggests, that the idea of common fate is not restricted to velocity alone.

RQ2 *What is the relation between the (extended) Law of Common Fate and other Gestalt Laws?* As the only Gestalt law dealing with dynamic settings, and given the perceptual urgency of motion [TMB02b], we are interested in the relation between LCF and other grouping laws.

To answer these two questions, we discuss criteria that may have an impact on perceptual grouping and identify visual variables that obey these criteria.

5.2.1 Criteria for perceptual grouping

As is clear from the above treatment on Gestalt psychology, perceptual grouping of visual objects arises from relations between *visual variables*. Identifying such visual variables was one of the fundamental advances of early work in visualization; for example, Bertin [Ber83] lists seven visual variables, and Cleveland and McGill [CM85] list ten. However, it is not feasible for us to study all of these visual variables, and besides, not all of them have the same potential for exhibiting perceptual grouping. Here we describe our selection criteria.

Associativity Visual variables that support grouping are often described as *associative* in the literature. It is worth noting, however, that this term has often been misunderstood by the community. As Carpendale points out [Car03c], there are discrepancies between the notion of associativity as defined by Bertin [Ber83, p. 48] and that which is usually understood [Rei11]. For Bertin, a variable is associative if objects can be grouped across other

variables despite changes in that one. In contrast, Carpendale’s definition of associativity refers to perceptual grouping power. This explains the differences between Carpendale’s and Bertin’s classifications of associative variables.

As our study focuses on visual grouping, we adopted Carpendale’s definition, which provides us with one dynamic (motion) and eight static (position, size, shape, luminance, color, orientation, grain, texture) associative variables.

Ordered transitions Since our focus is on common fate, our second criterion of selection pertains to the dynamic aspects of the above listed associative variables. Among these variables, there are several for which it is difficult to describe a dynamic behavior and decide a transition. For instance, working with shapes, textures or color hue, we would have many options to choose from as people (including us) have no clear intuition of how one should transition from one value to another.

Thus, in the spirit of keeping the study as simple as possible, we focused on variables that, in addition to being associative, are *ordered*, i.e. a change in these variable can be perceptually interpreted as increasing or decreasing. This allows for deterministically interpolating between values for both increasing transitions (the value of the visual variable grows), or decreasing transitions (the value shrinks). From the list of associative variables given by Carpendale [Car03c], only position, luminance, and size are ordered.

5.2.2 Visual variables with grouping

We focus on the static and dynamic versions of the three visual variables satisfying our criteria as follows:

Static variables Static visual variables are invariant over time and can thus have no impact on perception of common fate. Instead, including these factors allows us to answer RQ2 on the relation between LCF and other Gestalt laws.

- *Static position (SP)*: The geometric position of objects was ranked as the most perceptually accurate [Ber83; CM85], and thus it follows that two elements in close proximity are grouped. This phenomenon is known as the Law of Proximity.
- *Static size (SS)*: According to the Law of Similarity, elements with the same size will presumably be grouped together; furthermore, Bertin [Ber83] names size as the second most perceptually accurate visual variable, whereas Cleveland and McGill [CM85] rank area as number five.
- *Static luminance (SL)*: By the same Law of Similarity, elements with the same color are grouped together. Bertin ranks it at position five, and Cleveland and McGill rank it as “color saturation” at six for perceptual accuracy.

Dynamic variables Dynamic visual properties represent behavior that changes over time, which means that they may exhibit common fate effects. These factors allow us to answer RQ1 on whether the concept of common fate extends beyond mere object motion.

- *Dynamic position (DP)*: The canonical example of the Law of Common Fate: objects moving with the same speed and direction are perceived as belonging to the same group.
- *Dynamic size (DS)*: This is a property we want to investigate. Are visual objects growing or shrinking in the same way perceived as belonging to the same group?

- *Dynamic luminance (DL)*: This is another property we want to investigate. Are visual objects becoming brighter or darker in the same way perceived as belonging to the same group? Prior studies already provide evidence that DL enables visual grouping [SB01b; Wag+12; WB05]. However, their setups did not allow for also investigating the relation to other visual variables, both static and dynamic.

5.3 User study background

The goals of our study are to (i) determine whether the LCF extends to visual variables beyond motion, and to (ii) determine the relative grouping strength of LCF and other Gestalt Laws. Here we present the design rationale for the study.

5.3.1 Task rationale

One way to empirically estimate the grouping power of visual variables is to present viewers with a set of objects of which subsets share similar values for one or several of the variables under study, and test whether these values influence the perception of groups among the objects. Such perception implies grouping capability for the corresponding visual variable.

In our study, we chose to give participants perceptual tasks where four graphical objects were grouped by two properties at a time so as to create two orthogonal possible groupings, and ask participants which emerging groups they perceive. In other words, we make two visual variables compete, and record which one—if any—coincide with the participant’s answer, and hence influenced her grouping perception.

From a visual variable’s grouping power perspective, any answer to the above question falls into one of the three following categories: (i) the participant’s grouping coincides with that dictated by the first visual property, and we can assume that the corresponding visual variable thus has the highest grouping strength for this task; (ii) she grouped the objects based on the second, competing property, so we assume that the other visual variable has the highest grouping strength for this task; or (iii) none of the above (i.e., she grouped differently or did not find any groups at all), in which case none of the two variables can be said to have a grouping power for this task.

Our focus being on common fate, we are primarily interested in tasks where dynamic variables are involved, and hence on animated transitions implementing these dynamic behaviors. However, for the sake of experimental data completeness, we also tested static variables against each other, and our perceptual tasks also included static visualizations.

By making a dynamic visual variable compete against any of the static variables whose grouping power is established (i.e., by the Law of Proximity or the Law of Similarity), we can quantitatively measure the grouping power of the Law of Common Fate—in our case, restricted to motion, dynamic luminance, or dynamic size specifically. The more the cases where participants deviate from the Laws of Proximity and Similarity in favor of the dynamic property, the stronger the evidence that the associated dynamic visual variable has perceptual grouping power, and subsequently the stronger the evidence that the Law of Common Fate applies to this variable (RQ1). The relative grouping strength between each variable is directly measurable from tasks comparing pairs of non-conflicting visual variables (RQ2).

5.3.2 Summary of tasks

Table 5.1 summarizes all of the possible pairwise comparisons for our six visual variables. Out of the 36 cells, we do not consider self-comparisons (diagonal), nor do we count duplicates (i.e. SP vs. DL is the same as DL vs. SP); these thus are grayed out. We also discard any

pairwise comparison where a dynamic visual variable competes against its static counterpart (i.e. SS vs. DS). The reason is to avoid conflicts: having orthogonal groups bound to the same visual variable would necessarily break the notion of similarity at a point during the animation, making such cases difficult to interpret.¹

This leaves out 12 distinct pairwise comparisons that form our set of tasks for the study: DP-DS, DP-DL, DP-SS, DP-SL, DS-DL, DS-SP, DS-SL, DL-SP, DL-SS, SP-SS, SP-SL and SS-SL.

	DP	DS	DL	SP	SS	SL
DP	—	DS-DP	DP-DL	—	DP-SS	DP-SL
DS	DP-DS	—	DS-DL	DS-SP	—	DS-SL
DL	DP-DL	DS-DL	—	DL-SP	DL-SS	—
SP	—	DS-SP	DL-SP	—	SP-SS	SP-SL
SS	DP-SS	—	DL-SS	SP-SS	—	SS-SL
SL	DP-SL	DS-SL	—	SP-SL	SS-SL	—

Table 5.1: Summary of the different tasks generated from the comparison of our six visual variables.

5.3.3 Manipulation of visual variables

Let *visual property* henceforth denote a specific value for a visual variable. To create the above tasks, where objects are grouped by similar visual properties, we manipulate static properties (i.e. position, size, and luminance) as well as dynamic behaviors (i.e. changes in position, changes in size, and changes in luminance). Through these manipulations across objects, we can manipulate the relation of *similarity*—in the most general sense of the term—between objects to create distinct groups of objects sharing similar visual properties.

Here, we propose a generalization of the definition of *similarity* in a particular visual variable’s definition space for both the static and dynamic aspects.

General notation In the following, we use \mathcal{S} to refer to the set of visual objects in a task. For a given object A in \mathcal{S} , $V_A(t)$ refers to the value of a visual variable at time t of the animation, and $\Delta V_A(t_{i-1}, t_i)$ denotes the difference of values for A between time t_{i-1} and t_i (i.e. $\Delta V_A(t_{i-1}, t_i) = V_A(t_i) - V_A(t_{i-1})$), where the increment between times t_{i-1} and t_i corresponds to one step at the finest observable temporal resolution.

More specifically, let $P_A(t)$, $S_A(t)$ and $L_A(t)$ refer to the position, size, and luminance of the object A at time t of the animation. Object sizes and coordinates are homogenous, i.e. normalized by the visualization dimensions to the range $[0, 1]$. Similarly, object luminance is normalized to $[0, 1]$.

Similarity and similar behavior Visual objects A and B are similar with respect to a visual variable V at time t if their difference is below a threshold:

$$|V_A(t) - V_B(t)| \leq \tau_V$$

In the static case, the notion of similarity for two objects A and B directly refers to the Law of Proximity for position, and the Law of Similarity for size and luminance. In other

¹For example, comparing SP and DP would mean that two objects grouped by static position would only be in proximity during a single point in the trial, e.g. at the beginning or end; they would become separated (and thus no longer proximate) by varied dynamic positions (velocities) during the rest of the trial.

words, these static situations correspond to the special cases in the above definition where $P_A(t)$ and $P_B(t)$ are constant over time (i.e. static position), $S_A(t)$ and $S_B(t)$ are constant over time (i.e. static size), and $L_A(t)$ and $L_B(t)$ are constant over time (i.e. static luminance).

What the Law of Common Fate suggests, is that even if objects are not similar at any time t , the fact that they behave similarly is a factor for perceptual grouping. Put differently, this means that the difference in their *variations* across time is below a certain threshold. Formally, visual objects A and B *behave similarly* between t_{i-1} and t_i if:

$$|\Delta V_A(t_{i-1}, t_i) - \Delta V_B(t_{i-1}, t_i)| \leq \theta_V$$

Applying the above definitions in the context of our visual variables during an animated transition, we have:

- A and B SP-similar if: A and B similar in position, $\forall t$ during the transition;
- A and B SS-similar if: A and B similar in size, $\forall t$ during the transition;
- A and B SL-similar if: A and B similar in luminance, $\forall t$ during the transition;
- A and B DP-similar if: A and B behave similarly in position, $\forall t$ during the transition;
- A and B DS-similar if: A and B behave similarly in size, $\forall t$ during the transition; and
- A and B DL-similar if: A and B behave similarly in luminance, $\forall t$ during the transition.

We can operationalize these rules to create groupings for any of the above visual variables by ensuring both that (1) objects that are to be grouped are indeed similar (within some tolerance), and that (2) there exist no other object in \mathcal{S} that is similar to the objects in the group. Also note that these rules do not apply generally across all situations, but only in the context of our experiment. In general, similarity is highly contextual. For example, two objects with identical luminance will not be perceived as similar if one is placed on a darker background and the other on a lighter background.

Neutrality To control for perceptual processes and confounding effects, all objects in \mathcal{S} should be theoretically *neutral*, i.e. they should all be similar in all aspects (both static and dynamic). For simplicity and to guarantee perceptual grouping neutrality, we use a set of *static* and *identical* visual objects as a default set. It is only when testing the effect of visual variables on grouping that we modify these specific object properties to create distinct groups, as described above.

The only exception for neutrality is position, since it does not make sense to have objects overlap.² In fact, we cannot either enforce equidistance between all possible pairs of objects for sets of more than three objects. Dot lattices are commonly used in psychology experiments that study the laws governing grouping by proximity [Bro14]; however, we chose to avoid too much regularity in object arrangement since this can also play a role in grouping by proximity [SK06].

Any positioning strategy deviating from the above options will necessarily introduce a small bias for a set of more than three objects. To minimize the odds of spatial proximity that may occur by uniform random positioning, we used a similar approach as Bridson's Poisson-disc sampling [Bri07], which results in a balanced spatial distribution by adding a constraint on the spatial position of each object relative to the closest neighbor as follows:

²Perfect position similarity (i.e. $\tau_P = 0$) would entail all objects on the exact same position.

each object must be located within a distance range of $[d_{min}, d_{max}]$ from its closest neighbor (see also Bostock’s illustration of uniform random vs. Poisson-disc sampling [Bos14]). The smaller the difference between d_{min} and d_{max} , the more regular the objects’ arrangement.

5.3.4 Design decisions

We made several decisions when designing our experiment based on pilot testing and the above theoretical framework.

Choice of animation Because we primarily study the impact of dynamic changes on perceptual grouping, our main focus when testing dynamic variables lies in what happens during the animated transition itself, and nothing more. We want to prevent any bias that may be caused by the exposure to the first frame (i.e. the initial static state) or the end frame (i.e. the final static state). Hence, for the tasks involving dynamic changes, we prompt the participants with the same animation that loops continuously, with short interruption—a white screen—between two loops. In other words, the visual objects are never static.

Choice of object number Our task consists of mapping conflicting visual variables to different subsets within the set of visual objects, and then asking participants which perceptual grouping is strongest. We chose to use four objects grouped by two visual variables in our task. The motivation for this choice is the fact that four objects allow for setting up conflicting groups, and also conform to research in object tracking, which states that most people can track up to four moving objects [CA05a]. Practically speaking, given four objects, A , B , C , and D , we assign the pairs (A, B) vs. (C, D) to be similar in visual variable V , and the orthogonal pairs (A, C) vs. (B, D) to be similar in V' .

Four objects organized in groups of two yields three possible grouping choices.³ For simplicity, we give all three groupings as multiple-choice options in a trial. Two groupings relate to the two respective visual variables (V and V'). The third choice, however, has no meaning, and indicates that a participant perceived the strongest grouping from a non-existent similarity. If this meaningless choice is favored by participants, this may mean there was a confound, or none of the visual variables involved causes perceptual grouping.

5.3.5 Grouping strength

For any given trial confronting two visual variables V and V' , we define the *grouping strength* of V and V' as binary scores—preferred vs. not preferred. Depending on the participant’s answer, one of the two variables is preferred over the other, resulting in a grouping strength of 1 for this former and 0 for the latter; or none of these variables guided the participant’s choice, and both grouping strengths are set to 0.

5.4 User study

Based on the above background, we designed a crowdsourced user study to investigate our research questions (Figure 5.1). The survey can be found at <http://bit.ly/common-fate>.

The following describes how we generated the tasks in order to manipulate the presence of groups of similar objects with respect to the visual variables to compare.

³Combinations of N objects in groups of $n = \binom{N-1}{n}$.




Figure 5.1: User interface of our study, during a DS-SL task. Here, A-B and C-D are pairs of SL-similar objects, whereas A-D and B-C are pairs of DS-similar objects, forming the two concurrent possible groupings.

5.4.1 Task generation

A task consists of a set of four labeled visual objects $\mathcal{S} = \{A, B, C, D\}$, two visual variables V and V' being compared, and the associated groupings G_V and $G_{V'}$. Each grouping is a partition of \mathcal{S} into two subsets of two objects, such that $G_V = \{G1_V, G2_V\}$, $G_{V'} = \{G1_{V'}, G2_{V'}\}$, $G1_V \cup G2_V = \mathcal{S}$, $G1_{V'} \cup G2_{V'} = \mathcal{S}$ and $G_V \neq G_{V'}$.

Visual objects are generated as circular shapes with a default size (i.e. diameter) of 0.05 (recall that all measures of distance are normalized between 0 and 1), filled out with a default luminance of 0.5 and outlined with a black stroke. Default positions are randomly generated using a procedure similar to Poisson-disc sampling [Bri07]. Each object is labeled with its letter, uppercased, using a black bold font and white outline to make it readable in all variations of luminance filling.

These default values are manipulated for each object with respect to the visual variables being confronted in the task. To guarantee that the groups are not ambiguous, we tested the most favorable conditions for each visual variable by adopting a strategy that aims to maximize the similarity intra-groups (i.e. by setting as small as possible value for τ_V), while maximizing the inter-group distance.

Manipulations of SP: When SP is under testing, we randomly generate positions so as to enforce SP-similarity within the groups $G1_{SP}$ and $G2_{SP}$ (we set $\tau_P = 0.2$), while maintaining distance between the two groups (we set a minimum inter-group distance of 0.5).

Manipulations of SS: When SS is tested, we use pre-defined size for $G1_{SS}$ (small: diameter set to 0.025), and $G2_{SS}$ (large: diameter set to 0.075). This corresponds to $\tau_S = 0$, and a minimum inter-group distance of 0.05 (in size).

Manipulations of SL: When SL is tested, we use pre-defined luminance for $G1_{SL}$ (dark: luminance set to 0), and $G2_{SL}$ (light: luminance set to 1). This corresponds to $\tau_{SL} = 0$ and a minimum inter-group distance of 1.

Manipulations of DP: When DP is tested, we randomly generate the velocity of $G1_{DP}$

in any direction, at a speed > 0.02 to guarantee a noticeable motion within the window constraints. We set the velocity of $G2_{DP}$ to the same speed, but in an orthogonal direction.

Manipulations of DS: When DS is tested, we make objects in $G1_{DS}$ shrink linearly from the default size (diameter = 0.05) to a smaller size (diameter = 0.025), and objects in $G2_{DS}$ grow linearly from the default size (diameter = 0.05) to a larger size (0.075).

Manipulations of DL: When DL is tested, we make objects in $G1_{DL}$ darken linearly from the default luminance (0.5) to a darker one (i.e. 0), and objects in $G2_{DL}$ lighten linearly from the default luminance to a lighter one (i.e. 1).

We also apply constraints on trials to prevent objects from overlapping or leaving the screen. Depending on the task, the resulting visualization is either static (i.e. SP-SS, SP-SL, SS-SL), or a two-second infinitely-looping animation (i.e. all other tasks, since they involve a dynamic variable).

5.4.2 Attention trials

Although crowdsourced graphical perception experiments are a powerful tool [HB10], special care is required for large-scale studies [KCS08]. A typical approach is to add attention trials—i.e. trials where an obviously correct answer exists—to filter responses of insincere workers.

We designed such attention trials as trials similar to the regular ones where, instead of confronting two visual variables against each other, we make the two visual variables work together. For instance, in an attention trial involving SL and DS, one pair of objects would be colored white and shrinking, and the other two objects would be colored black and growing in size. This results in a single meaningful grouping defined by both visual variables, and participants who chose other grouping options should be regarded as potentially insincere.

Three types of attention trials were implemented using the following combinations: SS-DP, SP-DL, and SL-DS, each of which was repeated twice. It should be noted that these are virtually inseparable from regular trials, so workers are unlikely to be able to game the study. The mere act of trying to determine whether a trial is an attention one requires sufficient effort that its purpose has been fulfilled. Furthermore, several comments from our participants indicate that our attention trials were not distinguishable from the regular ones.

5.4.3 Experimental protocol

Procedure Participants were first asked to give informed consent. Participants were then given a calibration guide: we asked them to maximize the browser window and zoom the viewport to make the experiment interface visible and avoid scrolling.

Participants were given an explanation of the task based on an example of four animated objects. They were instructed that their task consists in selecting the pairwise grouping of these objects that most fit their intuition among three proposed groupings—presented in a list of radio buttons. Participants were enforced to pick one grouping before moving to the next trial. They could change their mind until pressing the “Next” button. Participants were instructed to disregard completion time, and to regularly take breaks.

Participants were explicitly instructed to “not spend too much time thinking since there is no correct or wrong answer”, and were strongly encouraged to stick to their first impression. Finally, they were warned that they would be given attention tests, and that we may reject their responses based on these tests. After being introduced to the task, participants were asked to practice on a set of three pre-defined trials.

Participants were then asked to accomplish the 42 trials constituting the experiment. Depending on whether the task involved the manipulation of a dynamic variable or not, participants were shown either a static visualization, or a two-second animated transition that

played automatically and looped infinitely until a grouping answer was submitted. Participants had no control over the animation. No feedback was given on submitted answers. Participants did not have the ability to return to previous trials. Their overall advancement was visible on a progress bar (Figure 5.1).

After finishing all trials, participants filled out a demographic survey and were given the opportunity to provide freeform comments. The whole study lasted 20 minutes on average.

Participants We posted the experiment on [Amazon Mechanical Turk](#). 100 participants (46 female) responded. Participants were all [Turk Masters](#) and were paid \$2 for their time.

Apparatus and setup The experimental survey was created on the [Qualtrics](#) platform. Custom JavaScript code built using [D3](#) was inserted to generate animations. Since the study was crowdsourced using the Amazon Mechanical Turk platform, we did not have control over the study apparatus, but we logged information about the screen resolution and the web browser used during the study. All participants used Google Chrome or Mozilla Firefox browsers as advised, at a screen resolution ranging from 320×480 pixels (iPhone) to 1920×1200 pixels.

Experimental design Our independent variables were the visual variables DP, DS, DL, SP, SS and SL combined into pairs to form the 12 tasks described above. Each type of task was repeated three times for each participant. Our experiment factors were as follows:

$$\begin{array}{r}
 100 \text{ crowdsourced participants} \\
 \times 12 \text{ TASKS (pairs of DP, DS, DL, SP, SS, SL)} \\
 \times 3 \text{ repetitions} \\
 \hline
 3,600 \text{ trials (excluding practice trials)}
 \end{array}$$

Trials were grouped into blocks of six, and each block was followed by one attentional trial. The order of trials, regular and attention ones, was randomized across participants. For each trial, we recorded the visual variables V and V' , their grouping strength (i.e. the binary score of whether the variable was preferred), and the completion time.

Hypotheses

- H1 *The Law of Common Fate extends to dynamic luminance and dynamic size.* This prediction follows from our discussion above as well as from Wertheimer himself [[Wer38](#)] and prior studies on correlated modulations of luminance as a figure-ground segmentation factor [[SB01b](#); [Wag+12](#); [WB05](#)].
- H2 *The grouping strength of Common Fate is higher than Proximity.* Animation has high urgency—to the point that it must be carefully controlled [[TMB02b](#)—so we think it will outperform even the Law of Proximity.
- H3 *The grouping strength of Proximity is higher than Similarity.* Our experimental setup will allow us to study even static confrontations, and we think proximate objects will exhibit stronger grouping than visually similar ones.

5.5 Results

Due to growing concerns in various research fields over the limits of null-hypothesis significance testing for reporting and interpreting experimental results [[Cum13](#); [DCH14](#)], we base all our analyses on estimation, i.e. effect sizes with confidence intervals [[CF05](#)]. This




Figure 5.2: Summary of the mean grouping strength for all visual variables examined in our study. Error bars are 95% CIs. The leftmost column shows the mean grouping strength per variable, for all comparisons it was involved in. The fact that the CIs of DP and DS fall to the right of the plots shows that these dynamic visual variables have stronger grouping strength than any other tested variables. Each cell in the tabular view shows the mean grouping strength of the corresponding visual variable in row (in orange) vs. that of the visual variable in column (in blue), for this task only.

approach also aligns with the latest recommendations from the APA [APA10], and has been successfully applied in prior work evaluating animated transitions [CDF14].⁴

5.5.1 Data verification

Before running analyses, we performed a series of verifications to check the Turkers' sincerity and filter out invalid data. We counted the number of failed attention trials. Out of 6 trials, only 8 participants failed 1 or 2 of such trials. Interestingly, several participants commented explicitly on the attention trials, e.g. "I didn't see the attention checks. That makes me nervous that I missed them," and "Despite watching with a paranoid intensity I never noticed any of the attention checks." This confirms our assumption that Turk Masters are sincere, and that our experimental data is valid.

We also looked at answers that did not correspond to one or the other visual variables involved in the trial, which can be seen as incoherent groupings. This data helps us verify if participants understood the instructions correctly, as well as spot potential problems with our experiment. Out of the 100 participants, 21 participants picked an incoherent choice once (out of the 42 trials), 5 participants picked an incoherent choice twice, and a single participant picked such an option 9 times in total. Given a total of 3,600 trials, this gives us confidence that our results are sound.

5.5.2 Dynamic vs. static groupings

In our study, we first set out to determine whether the Law of Common Fate extends to dynamic luminance or size. Figure 5.3 shows the overall grouping strength of all three *dynamic* variables (one per row) when competing against *static* variables only (see Figure 5.2 for the detailed per-task results). The grouping strength of each dynamic visual variable was assessed by aggregating all trials where the dynamic visual variable was in competition with a static visual variable (i.e. aggregation of 2 tasks \times 3 repetitions = 6 data points per variable per participant), yielding a total of 100 data points for each of the DP, DS, and DL variables. Point estimates and 95% confidence intervals were computed using bootstrapping [KG13]. In the figure, higher values (to the right) mean higher grouping strength for the dynamic visual variables. Overall, the grouping strength of a variable can be understood as the probability that the variable will be most influencing perceptual grouping over another, conflicting visual

⁴More information on statistics practices for HCI research in [Dra15].

grouping cues. Hence, a grouping strength above the critical threshold 0.5 in the figure indicates a higher influence of the dynamic variable on grouping over that of the tested static variable.




Figure 5.3: Mean grouping strength for the dynamic variables (DP, DS, DL) in the context of a conflicting static grouping. Error bars are 95% CIs. That all CIs fall to the right of the halfway mark (0.5) shows that these dynamic variables have stronger grouping strength than static ones.

Overall, our data confirms the Law of Common Fate for motion (DP). Furthermore, the data leaves no doubt as to the effect of DL and DS on perceptual grouping for the type of tasks we tested, since all point estimates and associated CIs fall in the right half of the plot (i.e. above 0.5%). In other words, grouping based on the dynamic behavior was chosen in more cases than cases it was not chosen.

5.5.3 Relations between Gestalt Laws

Our second research question concerned the relationship and rankings between the different Gestalt laws. Figure 5.2 shows a detailed summary of the results of our experiment. The leftmost column of the figure shows the overall grouping strength of each visual variable for the study. The grouping strength was assessed by aggregating all trials where the corresponding visual variable was involved (i.e. aggregation of 4 tasks \times 3 repetitions = 12 trials per variable per participant) across participants, yielding a total of 100 data points for each of the visual variables. A higher value again corresponds to a stronger grouping power.

While it is clear from the plots (leftmost column) that DP and DS have stronger grouping power overall (CIs above 0.5) than DL, SS, and SL (CIs below 0.5), no strict order can be established between all six visual variables.

In the right part of Figure 5.2 (columns 2-7), a tabular view similar to Table 5.1 summarizes all comparison tasks (one per cell). Each cell corresponds to the comparison of one visual variable (one per row) against another visual variable (one per column). The effect of each visual variable on grouping was assessed by performing a contrast, i.e. aggregating all trials corresponding to that comparative task for each of the two visual variables involved.

This detailed view allows us to refine our comparative assessment. Passing all visual variables in review (i.e. read the table by columns, for instance), we can make all possible pairwise comparisons. Regarding grouping by dynamic behavior, DP has a clear observable effect on perceptual grouping when in competition with DL, SS, or SL. DS also exhibits a strong effect when competing against DL or SP. DL has stronger grouping power than SS (and by symmetry, a lower grouping power than DP and DS). As for static variables, SP has no strong power compared to DS. It is, however, a fair competitor to DL, as well as SS and SL. SS is clearly not a strong grouping factor when put against visual variables DP and DL; but as strong of a factor as SP and SL. Finally, SL has no strong effect on grouping compared to DP, however, it has as strong of a grouping power as DL, SP and SS.

5.5.4 Participants feedback

Reviewing the open comments that participants entered at the end of study yielded additional insights. Several Turkers made assumptions about the purpose of the study: “*it appeared as if you were trying to appear which factors influenced grouping*”, sometimes even leading to self-reflection on their approach of visual grouping, e.g. “*never really thought before about whether I associate connections more between colors, or size, or movement, etc.*” Their observations not only point out the noticeable engagement of these participants, but may also reveal the quality and relevance of our study.

The grouping difficulty diverged between participants. Some found the task easy to do, while many others felt it was hard to decide a grouping. Reviewing comments of the participants who struggled reach a decision, we discern that pairing was challenging mainly because they found different possible ways to group objects: “*It was sometimes hard to determine groups because you could associate them together in a few different ways.*” The conflicting—both valid—options made them hesitant as to what visual factor would yield the best group: “*It was hard to decide which characteristic made it the best pairings*”, “*It was not always easy to decide. I tried to prioritize...*” These comments reinforce the idea that perceptual grouping of several groups—where none of the visual variable primes over the other—is possible in some cases. The forced choice as imposed in our experiment seems to have conflicted with what these participants perceived.

Some of the participants went beyond giving general comments and explained the grouping strategy they adopted to do the tests. Several comments confirmed the traditional Law of Common Fate, e.g. “*I tended to group those together that were moving in the same direction,*” “*Anytime the motion was the same I pretty much felt that was the highest criteria for what was a group*”, with an interesting mention of causality “*I went with what points seemed to be moving in common or influencing each others actions*”; or a more generalized definition of common fate: “*I would look for ones that were doing the same thing—such as changing size and group those.*” Other systematical approaches to grouping were described, including rankings by attribute (static and dynamic confounded): “*I grouped [...] the same colored circle together first and foremost, and after that I went by size.*” Overall, as suggested by our quantitative results, no strict order between the different variables seems to emerge.

5.6 Discussion

Our work was guided by the desire to explore the meaning of the seemingly straightforward concept of “common fate” in Gestalt psychology. Is that term merely shorthand for “visual elements that move in the same way,” or does it suggest a deeper and more fundamental meaning where the visual elements are behaving in the same way, regardless of what that behavior is? Put differently, we were interested in learning which dynamic behaviors beyond common velocity of two visual objects would make them be perceived as grouped by an observer. Our secondary objective was to determine how this dynamic grouping interacts with static visual properties such as the proximity, luminance, and size of graphical objects. The results from our crowdsourced user study allow us to begin answering these questions and discuss how they generalize to other settings.

As shown in Figure 5.3, DP, DL, and DS present a clear benefit compared to their static equivalents (SP, SL, and SS). Yet, as illustrated by Figure 5.2, DL is also clearly outperformed by DP and DS while there is no clear distinction between these two. In the real world, objects rarely grow or shrink in our field of view, except when moving closer or further. It may well be the case that, when compared to DL, DS benefits from our ability to perceive such real-world motions: “The human brain is specialized to recognize and interpret certain kinds

of structured movements, especially motions that have a physical, an organizational, or a biological interpretation” [NF02], a phenomenon that aligns with prior studies exploring motion decomposition [FH02; NF02]. The unclear situation between DP and DS may also result from our difficulty to isolate one particular motion-based common-fate group from another [LF11].

Overall, our study confirms that the Law of Common Fate applies—at least to some extent—to dynamic luminance and dynamic size, and that these variables can be used for making a group visually salient in a dynamic visualization. Our results also suggest several cases where two visual variables support conjunction search, which enables emphasis of two groups at the time. This is, for instance, the case for DS and SL; a combination that has been put to test in graph visualization [FH02]. Below, we discuss how our results can be interpreted and leveraged for the design of dynamic visualizations.

5.6.1 Implications for dynamic information visualization

While further investigation is necessary to generalize our results in more realistic scenarios, our empirical study on the effects of variable conjunction on perceptual grouping is an initial step towards more fleshed out design guidelines for dynamic information visualizations and animated transitions. Here, we discuss several situations where our study can be useful to make design decisions.

Choosing from visual variables to emphasize groups depends on two main factors. First, one should not use the same pair of visual variables for making groups perceptually emerge when a visualization is intended to treat several groups as equivalent, and, on the contrary, when a single group should be made more salient than any other searchable groups. Second, the range of possibilities for visual variables to be used for perceptual grouping is limited by the encoding of the visualization. In other words, if luminance is already mapped to an attribute, altering this encoding for grouping purposes may be perceived as a change in the very attribute’s values. It is hence preferable, whenever possible, to have recourse to another visual variable for grouping.

Making a single group salient. Perceptual grouping can be used, for instance, to make the current selection of elements pop out. In such cases, where the intent is to focus attention on a single group, a visual variable with strong grouping power should be preferred. For instance, in a dynamic visualization where objects have a fixed size (i.e. size is not mapped to any attribute), a static, but meaningful position (i.e. SP encodes data attribute), and a luminance that changes over time (i.e. DL encodes values); a sub-group of these elements has most chance to pop out if encoded using DS, while not interfering with the current visual encoding.

Revealing several groups. When it is important to enable conjunction search—i.e., when several groups should be identifiable at the time—a set of visual variables that have similar grouping strength when used together—i.e. none of the variables dominates the others for grouping—should be preferred. For instance, imagine an animated transition between two scatterplots, where position only encodes data. During the transition, dots will inevitably move from their initial position to their final position. To make two groups emerge in such case (recall that DP, SP are not available), one can choose to use either the pair SS/SL for revealing two groups, or DS/DL. Other configurations, such as SS/DL are more likely to create perceptual unbalance: DL dominating SS for grouping, the group encoded with SS is prone to be unnoticed.

5.6.2 Limitations

While our results provide evidence that there is more to the Law of Common Fate than is traditionally understood, there are limitations to our study that should be investigated

to examine the role of dynamic visual properties in perceptual grouping. Our study is characterized by a careful control of potential confounding effects to isolate the effect of our independent variables on grouping. By making this choice, we favored internal validity over external validity [CPT82]. Future studies should be conducted to confirm that our results hold in more realistic scenarios. We discuss aspects of this below.

Perceptual distortion. Prior literature has evidenced the influence of particular variables on the perception of other variables, when used in conjunction. For instance, when size and luminance are both used to encode data, interpreting values can be challenging since “[...] size severely influences the perception of lightness, and lightness barely has any influence on size perception” [LMW10]. Further, motion is known to draw the viewer’s attention, resulting in a moving item to have higher chances to be attended by the viewer. As a result, any variable encoded on moving items is also more likely to be noticed. In our study, it is therefore difficult to assert whether size or luminance changes were noticed because of the saliency of the variable itself or rather because they were made more noticeable due to the associated position change.

Larger datasets. One of the main limitations of our study lies in the low number of data points that we considered. A more general experimental design would consider larger datasets, test various scenarios making the number of groups and group size vary, and introduce distractors.

Finer granularity. In order to guarantee observable effects, we enforced best conditions for each visual variable to create distinct groups (i.e., maximize the inter-group distance, and maximize the intra-group similarity). Other possible scenarios include a wider range of possible values as initial and final states, making the perceptual identification of groups more difficult (i.e. shades of gray instead of white or black).

Beyond position, size, and luminance. To keep the experimental design manageable, we focused on a subset of the visual variables that likely play a role in perceptual grouping. We discussed the rationale behind our choice of variables explored in this work; other variables such as texture, grain, orientation may, in their dynamic behavior, play a role in perceptual grouping. These should be considered in future work.

Richer interpolations. There are many ways to interpolate between an initial and a final value in a visual variable’s space. In this work, we choose a monotonous, linear interpolation in an ordered space, progressively transitioning to a “more” or a “less” value. In such spaces, other transition strategies could be considered, such as, e.g., different rhythmic patterns (i.e. blinking for luminance, flickering motion), exaggerations (i.e. overshoot), etc. These are factors that may well also fall under the purview of Common Fate.

5.7 Evaluation scenarios to assess our design implications

We present our reflection about the possible approaches that can be adopted in the future to assess how the Law of Common Fate, applied to dynamic position, size and luminance, can be employed for emphasizing a particular group or making several groups equally salient in a dynamic visualization.

5.7.1 Make a particular group salient

As discussed in the Section 5.6.1, the three studied visual variables (i.e. position, size, luminance) can be used to make a group more salient than the rest of groups in a given visualization. Here, we discuss evaluation scenarios that would allow to assess how dynamic position, size and luminance can influence the visual salience of the dynamic behavior that they are encoding. A relevant research question to investigate is the impact of each of

these dynamic visual variables on the perception of change trends in an animated dynamic visualization. An empirical study can ask participants about the main trend that they are able to see in an animated visualization, which means detecting the most predominant dynamic group in the visualization. The main goal of this task would be to capture the *first impression* that participants have when observing the animated visualization once, and what dynamic behaviors they were capable to detect at first glance. This evaluation approach would allow to test the power of the visual variables in bringing attention to dynamic groups, and hence their effectiveness for emphasis and visual prompting.

Another relevant research question to explore is the impact of the value (size of effect) and cardinality (numeracy) of change (defined in Section 3.3.4) of a given visual variable (i.e. position, size, or luminance) on its effectiveness in making a group salient. For example, one can assume that the larger, or the more important is the change the more visible and salient would be its corresponding group. Similarly, the more visual elements affected by a given change, the more visible and salient would be the group composed by these elements.

5.7.2 Revealing multiple groups

Several scenarios can be used to evaluate the potential of the visual variables position, size and luminance to reveal the existence of multiple groups in a dynamic visualization. Relevant examples from the literature include evaluating the perception of relationships between two different subgraphs in a dynamic graph visualization. Ware and Bobrow [WB04b] studied how two subgraphs can be highlighted differently within the same graph visualization and proposed the set of the following visual queries that should be supported by the highlighting techniques: detecting whether a node belongs to a subgraph A, the intersection of subgraphs A & B, or neither A nor B. In their user study, the authors used the following single task: “*Do highlighted subgraphs A and B intersect (i.e. do they have nodes in common)?*” The list of possible tasks can be further extended using the tasks that have been used in literature when studying disjoint groups. Relevant tasks include for example asking about the number and the size of groups (e.g. small, large) [VBW17]. Asking about the number of groups would not only allow us to assess the visual conjunction between the three visual variables that we studied in our experiment (i.e. position, size, luminance), but it would also enable to measure the level of salience of each visual variable: if only a subset of groups is detected, then the visual variables encoding these groups can be considered more salient than the visual variables used to encode the groups that the viewer failed to detect. If all the groups are detected, the level of salience can be measured by the chronological order in which the participant detected each group or the number of animation replay required to detect each group. In the experiment, participants should be asked about both the number and the description of each detected group (e.g. items shrinking, items moving to the left, etc.).

5.8 Conclusion

In this chapter, we have presented results from a crowdsourced user study for a graphical perception task involving four visual objects with varying static and dynamic visual properties. By pairing these four objects in two orthogonal ways (A-B and C-D vs. A-C and B-D, for example), we were able to assign conflicting visual behaviors to the two sets of pairs. The rationale behind our study was to be able to determine which visual behavior resulted in stronger perceived grouping than other behaviors. Ultimately, our results indicated that the Law of Common Fate is not restricted to mere motion, and that dynamic luminance and dynamic size also have strong grouping power. However, we were unable to establish an absolute ranking between the dynamic visual variables. On the other hand, we found that all

three dynamic visual variables had stronger grouping strength than the static ones—with the exception of DL, which has as close grouping strength as that of SS—which suggests that the Law of Common Fate is perceptually stronger than the Laws of Similarity and Proximity.

Our empirical study represents a constructive step towards a deeper exploration of the expressive power of animation trajectories for portraying various semantic relationships between data elements in dynamic visualization such as grouping, causality, hierarchy as so forth. As a first step toward further study of the role of dynamic visual variables on perception of groups in visualizations, we have presented different ways to apply our findings on visual grouping to use animation as a medium for emphasizing a particular group within a visualization, or to support visual conjunction of several groups. We also discussed actionable directions for the evaluation of the value and effectiveness of animation for both visual emphasis of a single group, and for revealing the existence of several groups of items within a visualization.

Chapter 6

Exploring and supporting the design of staged animation

In this chapter, we present a design space that allows a holistic characterization of *animation choreography*—authoring staged animation across several visual states of a dynamic visualization. Staged animation offers the possibility to direct viewers’ attention to important changes, to isolate and clarify changes happening simultaneously, and to tell rich and evocative stories with visualization. While many instances of staged animation can be found in the literature and in practice, the design space to characterize or guide the design of staged animation has neither been deeply explored nor formalized. To fill this gap we propose a design space articulated around two main facets: i) the visualization and data context within which the staging takes place; ii) the strategies (simplify object tracking, emphasize changes) and parameters (decomposition, timing) employed to generate the staged animation. We survey and characterize literature and practice using our design space, thereby identifying the limitations of existing staging approaches and highlighting novel opportunities. To support the design of staged animation, we provide a paradigm for authoring staging alternatives based on varying the decomposition and timing parameters, and demonstrate our approach with a concrete scenario.

Animation is commonly used in interactive visualization applications to allow observers to gradually track and understand changes that occur on a display (see Chapter 2). For instance, imagine that you want to study the evolution of inhabitants of a city district over several years, in terms of gender, age, job sector and income. You may visualize this data as a scatterplot (x-axis: age, y-axis: income) where data points are colored by gender, and smoothly animate the changes of position to learn about the progression of the residents.

In general, dynamic visualizations such as this are designed to animate all data elements at once. All men and women of different ages be they unemployed or retired, students with moderate stable incomes, or lawyers with widely fluctuating incomes, will evolve at the same time during the animation. Such an animation would surely give us a general overview of how these people have evolved but would not allow us, for example, to easily decipher the behavior of individual groups emerging from the whole, or detect the different trends describing their evolution because of the high number of changes happening at the same time on the display.

To help the observer better track the behavior of particular groups of individuals (e.g. *How did men’s income evolve compared to that of women’s?*) or follow the different changes happening (e.g. *Who had a stable or decreasing income? Who had a major salary rise?*), the animation can be broken down into a set of *stages*, where only a sub-set of elements are animated at a time to facilitate tracking. For instance, animating all of the data points representing men first then followed by women would help us get a gist at the evolution of these sub-groups in isolation (see Fig. 1-a). An alternative staging would allow us, for instance, to see in sequence what part of the population had a raise, then those whose salary remained stable, followed by the residents whose income decreased overall (see Fig. 1-b). And we can imagine to further decompose the sequence to more deeply study additional aspects of these sub-groups. For instance, a finer grained staging would allow us to see in sequence the evolution of each change in income (rise, stable, decrease) by current job sector, starting with finance, followed by business, then other sectors, and ending with retail.

In contrast to *full animation* that consists of animating all elements at the same time, *staged animation* consists of decomposing the animation into a set of stages, creating different

structured views of a dynamic visualization that can help surface various aspects of the data. The decomposition can be based on sub-groups (e.g. women vs. men), or be based on the dynamic changes or behaviors of elements (e.g. stable vs. decreasing vs. increasing salary). In either case, the information can be already visually encoded (e.g. gender is mapped to color) or not (e.g. job sector is not visualized in our example), offering a rich set of possibilities for structuring the animation.

In practice, staged animation is used to clarify dynamic behaviors, to direct attention to changes of interest, and to communicate relationships supporting rich narratives with visualizations. It has been the focus of several research studies [Pla+12; ZKS11; GVM12; May01; Put11; Kha+13; Spa+12], and practitioners have also leveraged this type of animation to tell stories in data journalism [M11; Hel; Got], add structure to oral presentations [Key; Ppt], or even improve visual appeal for more compelling web advertisements [Ski; Jd1; Vn1; Met].

Both literature and practice offer numerous examples of staged animation. However, these instances mainly consist of isolated point designs that instantiate staging to fit a particular context of use. Prior research work has proposed elements of characterization of animated transitions pertinent to staging [HR07; BPF14], but globally, the design space of staged animation has not yet been systematically explored nor formalized. Our research moves beyond point designs by synthesizing and extending previous research and practice to present a holistic design space for *animation choreography*—authoring staged animation across states of a dynamic visualization. We use our design space to describe and compare existing approaches to staged animation, and to generate new ones.




Figure 6.1: Example of staging strategies animating the evolution of a population between two years. a) Elements are grouped and animated by gender: male first, then female. b) Elements are grouped and animated by behavior: increasing income first, followed by stable, then decreasing.

6.1 Staged Animation

In this section, we discuss the motivation for a theoretical framework, as well as the challenges and limitations of staging. We ground our discussion with research in traditional animation, cognitive psychology, and visualization.

6.1.1 Motivation

Staged animation has been investigated as a way to overcome perception and usability challenges with full animation, in particular using a process of dividing animations into steps for visualizations.

Challenges with animation

The use of animation in visualization has been aimed at alleviating the cognitive load associated with the tracking and understanding of multiple simultaneous changes occurring when, for example, visualizing the temporal evolution of a dataset [Rob+08; Che+10], hiding or revealing

structures in a dataset on user interactions [RMC91b] or transitioning between multiple representations for a dataset [HFM07]. Animation attempts to achieve this generally through a smooth interpolation of the visual attributes of data items during the transitions. A smooth interpolation, however, is not always sufficient to support the detection and apprehension of the changes, when there are too many of such changes happening at the same time [Rob+08].

Research in cognitive psychology and perception reveals that some of the main obstacles to comprehension of changes stem from the limitations of our perceptual capabilities in terms of focused attention [SC99] that may lead to change blindness [CCD03], and the limitations of our working memory in terms of our capacity to integrate, process and recall information [Har07; KO11; SMP98]. Not only is the number of objects that we can track simultaneously particularly low [CA05b; Lig01] and our visuospatial attention influenced by clutter and object saliency [Ros99], but the extraction, processing and integration of information over time is also difficult when multiple dynamic behaviors concurrently occur in a display [Sch01; Low04; WM13]. The later issue is qualified by Lowe [Low04] as the “overwhelming effect” of animation. The perceptual difficulties of tracking too many objects, and/or too many changes simultaneously are further exacerbated by the fact that presenting all the information at the same time also misaligns with the way we tend to reason, work and communicate most effectively, in sequential small steps rather than everything at once [Tve11; Zac+07].

Staging in visualization: divide to conquer

Staging, in visualization, refers to the decomposition of an animation into steps. By presenting information in steps, the resulting *staged animation* can help cope with the perceptual and cognitive difficulties associated with visually tracking and making sense of complex dynamic visualizations.

Researchers have explored the potential of staging to help observers better apprehend a wide variety of changes in a dynamic visualization. Staging has been used to maintain object constancy when transitioning between visual representations [HR07; FP02; EDF08b] as well as between layouts for a given visual representation [Pla+12; GEY12]. For example, Dynavis [HR07] employs staging to transition from a bar chart to a scatterplot in steps: first changing the visual marks from bars to small circles maintaining their y-position, then moving the dots horizontally as the x-axis reveals the new dimension to form the scatterplot. Staging is also a common technique when animating during temporal navigation in datasets, for instance, to clarify what elements disappear from one data step to the next, what elements remain and how they evolve, and finally what elements appear [FE02; BPF14]. Staging has also been used in various applications, such as when browsing tree visualizations [PGB02; FP02; MDB04; MB05] to reveal or hide nodes in the hierarchy on user interaction, and when navigating through versions of a given diagram to support differencing and merging [ZKS11; Got]. We discuss prior literature in greater detail in Section 6.5.

The strategy of breaking down an animation into stages has been used to satisfy the *segmentation principle* [May01] that prescribes decomposition of instructions in the context of multimedia learning [Put11; Kha+13; Spa+12], and the *do one thing at a time principle* recommended for more understandable animated presentations [ZS03].

Overall, prior empirical studies have evidenced the potential of staged animation as a perceptual and cognitive aid in some contexts. Staged animation has proven to be effective in helping the visual tracking of visual items during animation [ZKS11; BPF14; Pla+12], and distinguishing the different changes to be detected during the evolution of dynamic visualizations [ZKS11; BPF14; GVM12]. It has been also found to be less demanding, reducing the cognitive load [HR07; GVM12], and acknowledged as potentially effective for learning visualizations [BCK05; HFM07; RM15]. In most of these works, staged animation

has also been subjectively reported as more visually pleasant than its single-stage animation counterpart.

6.1.2 Limitations and Challenges of Staging

Previous research highlights the potential of staging for more understandable and compelling dynamic visualizations. However, more research is necessary to better understand the positive and negative aspects of staged animation from both the designer’s and the observer’s perspective. We discuss costs, constraints and drawbacks highlighted by prior studies, as well as the challenges associated with the study of staging in the present research context.

Limitations and drawbacks

Prior research found that staging can be limited by the fact that it extends the overall animation duration, making it frustrating for the observer [Pla+12; HR07; ZKS11; BPF14]. In particular, in two studies, participants found staging useful to learn a novel visualization [Pla+12], but complained about the time it takes for the animation to unfold after they became acquainted with the visualization. The temporal sequencing that staging produces can also be misleading and confusing in some cases [ZKS11; RM13]. For instance, participants in a study on dynamic graphs had difficulties distinguishing between animation stages (e.g. designed frames placed between two data timesteps), and the actual data timesteps [RM13].

These findings provide interesting insights on the impact of staged animation on observers, and the difficulties faced by designers. However, prior research has only studied particular instances of staging, and focused evaluation mainly on time and error measures for some tracking tasks as well as qualitative impressions from observers [ZKS11; Pla+12; GVM12; HR07; BPF14; BCK05]. A more systematic approach to design and evaluation would help develop comprehensive knowledge on staging, its various dimensions, and their effects on and beyond time and accuracy performance, such as interpretation, power for insight discovery, or support for narrative.

A rich, but ill-defined design space

Designing and studying staged animation is a challenging process because the design space is rich, but also ill-defined. Standard animation already involves numerous parameters such as trajectories, interpolation, and pacing. Decomposition into stages extends the range of possibilities with additional aspects to handle, including the criterion to decompose the animation as well as the ordering and scheduling of the stages.

We also found that the terminology around staging varies widely; the visualization community has called staged animation: “multi-step”, “step-by-step” [Pla+12], “stage-by-stage” [EDF08b], and “serialized animation” [RPD09]. These all refer to the same core concept of decomposing an animation into stages. Similarly, “stages” have alternatively been called “steps” [PGB02; Pla+12; GVM12; FE02; RPD09; Fra+14], “transitions” [Pla+12], “sequences” [Pla+12], “phases” [ZKS11; MDB04; MKH12] and “frames” [Pla+12; MKH12]. When contrasted with staged animation, single-stage animation has been designated as “full” [Pla+12], “all at once” [ZKS11], “continuous” [Tve11], “direct” [HR07], and “linear” [GVM12] animation.

The lack of a settled terminology and formal characterization of staged animation makes it difficult for designers to comprehend the range of possibilities to make appropriate decisions, and for researchers to build a research agenda. A holistic design space would help clarify and stabilize the terminology about the different aspects defining staged animation, contributing a “considerable value for a theoretical model” [Shn03]. Beyond the evident fact that a standard terminology would tremendously facilitate search for existing work on the topic — currently a

non-trivial task due to diverse vocabulary — a formal and common ground terminology would help designers carry on more easily meaningful and constructive discussions in the future.

6.1.3 Staging in Traditional Animation

In the domain of traditional animation, staging (one of the 12 animation principles [Las87]) has a meaning rooted in theater and cinema, as it refers to the way of directing the audience’s attention to the most important part of the scene, via the use of light and shadow, angle and position of the camera, or particular placement of a character or an action in the frame [JT95].

Nagel et al. [NPD16] recognized this theatrical value, qualifying staging for visual analysis as “a carefully choreographed process of breaking up a complex whole into its component parts and purposefully preparing the manner of their appearance”. In their work, they use staging to guide the observer from discovering and enjoying a public visualization (i.e. hooking the user, and teaching a visualization [Che+16]) to later making sense of the underlying data and extracting insights.

We draw inspiration from the definition in traditional animation, and argue that beyond the temporal value of staging in visualization allowing it to reduce the number of objects and changes to track simultaneously, we can also take advantage of its theatrical value making it a powerful tool that can be used to emphasize changes of interest, express relationships, and communicate information from different perspectives. We introduce and explore the notion of *animation choreography* as an expressive method for decomposing and structuring a dynamic visual scene.

6.1.4 Animation Choreography: a Model for Staging

In this work, we address the lack of theoretical foundations for staged animations and propose a design space to establish a descriptive, comparative and generative common ground [BL00] guiding their design and broadening the perspectives for their evaluation.

Previous models

Prior models in animation have been proposed by the visualization research community. Most relevant to our work is the taxonomy proposed by Heer and Robertson [HR07], that outlines the different transition types between statistical data graphics representing the same dataset, along with design principles for creating animated transition between these representations. They posit staging as one of the design guidelines to alleviate visual complexity during a transition, and evaluate instances of staging, but do not discuss in depth the design space of staged animation. We build on this taxonomy and propose a holistic design space that systematically covers the varied aspects to characterize staging.

Other works include brief discussions about particular parameters of staging and design considerations in specific contexts. These include discussions about alternative ways of breaking down the animation based on the types of changes, the structure of the data and connections between data items [GVM12; BPF14], and the duration of stages [PGB02; BPF14]. We cover, generalize and extend these aspects in our consolidated design space.

Defining the design space

The proposed design space (Section 6.2) characterizes the factors involved along two main categories of dimensions: i) the factors specific to the context in which the staged animation takes place (Section 6.3), and ii) the factors that the animation designer has control over to decompose and schedule a staged animation for visualization (Section 6.4). We use our design space to analyze prior literature and practice relevant to staged animation, and point to

limitations of existing staging approaches as well as highlight novel opportunities (Section 6.5). We provide directions on how to leverage our design space to support the design of staged animation, introduce a novel paradigm for authoring staging alternatives and demonstrate our approach based on a concrete scenario (Section 6.6).

6.2 Design Space Overview

To design a staged animation, the animation designer has to work with a set of given elements, including data (e.g. data about a population in 2016 and 2017 as in our previous example), a visualization context (e.g. a scatterplot of the data in 2016, and a scatterplot of the data in 2017, as in Fig. 6.1), and the observer’s task (e.g. *is there a difference in evolution between English and French speakers’ income?*). Given this input, the designer needs to make choices and take decisions to create a staged animation that complies with the associated constraints, applies to the data and visualization context, and best support the observer’s task.

Our design space (Fig. 6.2) is articulated around two categories of dimensions that capture such given input (fixed dimensions), and design choices to create a staging (variable dimensions).

The *fixed dimensions* represent what is given to an animation designer. These dimensions characterize the data and visualization context within which staging will take place, and the tasks an observer needs to perform using the visualization. The *variable dimensions* describe what the staging designer decides in terms of strategies and parameters employed to generate a staged animation. After the designer chooses a staging strategy defining the main purpose of decomposing the animation, they have access to design parameters (i.e. decomposition, timing) to implement a staging that addresses the strategy appropriately.

There are complex dependencies and interactions between all of these dimensions, each of which has different impacts on the design and quality of a staged animation. In the following sections, we detail both the fixed and the variable dimensions, share our thoughts on how these can be used in practice to inform animation design, and discuss how these dimensions may constrain or influence design decisions.




Figure 6.2: The design space of *animation choreography*. A design is informed by fixed dimensions, which include characteristics of the dataset, changes in the data or (re)presentation, and the task of the observer (end-user) of the visualization. The designer of the staging can vary the variable dimensions, including setting a strategy for staging which informs the decomposition and timing of the staging design.

6.3 Design Space: Fixed Dimensions

Fixed dimensions characterize the immutable data and visualization context within which the design of staging must be realized. Fixed dimensions thus comprise of questions being posed about the nature of the data attributes (Section 6.3.1), the changes that the data and the visual

(re)presentation undergo (Section 6.3.2), and the purpose for designing the staged animation (Section 6.3.3).

6.3.1 Data Attributes

The data attributes dimension characterizes inherent attribute properties, namely, type and value. Data attributes can be of type categorical (unordered), ordinal or quantitative (both ordered) [Mun14].

6.3.2 Change Characteristics

The change characteristics dimension describes the changes that the data and the visualization undergo. This dimension encompasses questions including: Does the data and/or visualization undergo any changes? What type of changes? Across how many states?

Each change in a dynamic visualization can be categorized as being either a representation change (e.g. change from a line chart to a bar chart), a presentation change (e.g. view change after a zoom operation), or a data attribute change (e.g. fluctuating data values). Representation, presentation and data attribute changes could occur in isolation, or a data attribute change could induce a presentation change. For instance, an axis re-scale may be necessary to keep all elements visible in a scatterplot, if relocating data points to their new positions make them fall beyond the axis limits. Finally, these changes could be characterized as resulting over 1, 2 or n (> 2) distinct data states.

Data states

The underlying data can be characterized as exhibiting 1, 2 or n distinct data states. 1 data state accounts for stable data and implies that the changes observed are either representation or presentation changes (see paragraph below). For varying data (e.g. time-varying data) on the other hand, observed changes in data attribute values can be described as occurring either over 2 or n (> 2) data states. We separate 2 and n because, as discussed later, there are important implications for the design of staging.

Representation and presentation changes

The visualization, i.e. both representation and presentation [CMS99a; CM01], can either remain stable or evolve over time. Evolving representation accounts for changes of the mapping from the data to a structure that can be displayed visually, including visualization change and data schema change [HR07]. For instance, change from a scatterplot representation to a bar chart representation. Evolving presentation accounts for changes of the display of the representation, including view transformation (e.g. axis rescale), filtering and ordering operations [HR07]. In either case, these changes impact the way data is represented and displayed, but not the underlying data itself.

Data attribute changes

A data attribute can either be stable or evolving over time. Evolving data attributes include both attributes undergoing a presence (appear/disappear) and value change over timesteps [HR07]. Presence change affects the number of visual items, whereas attribute value change would lead to changes in visual properties (e.g. position, size) to which the attribute is mapped.

6.3.3 Observer Task

The observer task dimension focuses on what task the visualization should support for the observer. A staged animation can be designed to: (1) *understand* a visualization and the data it represents, (2) *verify* hypotheses and answer questions about the data, or (3) *present* data and knowledge about the data to an audience.

Observers use a visualization mainly to analyze or present the data under study [CMS99b]. When analyzing data, observers may first have to get acquainted with a visualization if they encounter it for the first time. They can also carry out an exploratory analysis to get an overview of a dataset and gain a high level understanding of the distribution, attribute values and evolution of data items. The *understand* task category therefore encompasses activities supporting both “understanding the visual representation” and “understanding the data” through discovery (i.e. generate hypotheses [BM13]).

In a more advanced phase of data analysis, observers form specific hypotheses and questions about the data, and thus engage in confirmatory analysis tasks. The *verify* task category encompasses activities around investigating a question about a data item, group of items, attribute, trend, or other specific characteristics of items or changes (i.e. verify hypotheses [BM13]).

In the context of data presentation, the observers’ goal is to *present* a story about the answers and knowledge they acquired about the data, or walk the audience through a series of cognitive operations [BM13]. For instance, we know that only two people in our example have seen their salary decrease in 2017 (see Fig. 6.1). A possible narrative could emphasize how a large portion of the population have seen their situation improved or remaining stable, whereas only two unfortunate persons made less money in 2017.

The observer’s task dimension closely relate to the “why” dimension in Brehmer and Munzner’s taxonomy [BM13]. In their taxonomy, *discover* encompasses both the generation and verification of hypotheses. In our work, we separate these two aspects since designing staging for a general exploratory purpose is highly different from designing staging where the hypothesis or question is known. Deciding on a specific staging strategy implies making a choice as to the lens through which the data will be presented. For instance, staging by population gender would be a poor choice if it is known that the observer is interested in studying salary evolution. Similarly, staging by population gender only shows one of the many facets of the data, possibly making it more difficult for the observer to discover patterns and trends involving other data attributes. In the context of exploratory analysis, the designer must be wary of not biasing the visualization (and staging) too much toward a particular aspect of the data.

Further, Brehmer and Munzner’s taxonomy assumes that the observer is familiar with the visualization they work with, and therefore does not include aspects pertaining to understanding a visual representation. We include this aspect in our taxonomy, along with the discovery phase under the *understand* task.

The kind of tasks the observer needs to perform with the visualization determines the most appropriate staging strategy to support them and tailors, in turn, the design choices defining both the decomposition and timing.

6.4 Design Space: Variable Dimensions

The variable dimensions outline the various editorial decisions an animation designer can make in the process of designing a staged animation, guided by constraints imposed by the fixed dimensions. These decisions include staging strategies a designer can adopt to support the observer task and determination of appropriate decomposition and timing parameters.

6.4.1 Staging Strategy

Staging can be used for different purposes: an animation designer can either choose a staging strategy that simplifies object tracking, or a staging strategy that emphasizes changes. A staging strategy focusing on simplifying object tracking attempts to augment the perception of certain visual items by reducing the number of distractor items and minimizing occlusion. For example, when animating evolution of a population, decomposing the animation into stages based on gender (see Fig. 6.1-a) would help reduce the number of data items to track at the same time. In contrast, a strategy that emphasizes changes aims to augment the perception of each type of change that the data and/or the visual representation and presentation undergo, by isolating each change in a dedicated step of the animation. For example, to observe how salaries evolve over time, a staging strategy that separates increase, decrease or no change in salary into stages helps emphasize the different types of temporal trends (see Fig. 6.1-b).

While the observer task is what the end user wants to accomplish, the staging strategy is how the animation designer gets there. That is, the strategy used is informed by the observer task. For example, if designed for the purpose of helping observers to understand the data (i.e. help generate hypotheses), a staging strategy for simplifying object tracking may suffice. In contrast, if the task to be supported by the animation is that of verifying a specific hypothesis, then the staging strategy could be tailored to emphasize changes in relevant attributes, even at the detriment of other details. Lastly, for the present task, when the animation is employed as a means for presenting insights about the data through narration, a curated staging strategy which draws attention to these insights and presents them as a story would be appropriate [HD11; SH10]. This storytelling feature is known in narrative visualization literature as guided emphasis [HD11]. Note that the link between the observer task and staging strategy is not fixed, and the expertise and insight of the designer should be acknowledged in making strategy choices for each specific instance of staging.

Distinguishing the two staging strategies may lead the reader to assume that a staging aimed at simplifying object tracking does not always help with emphasizing changes, and conversely a staging that emphasizes changes does not necessarily reduce the number of visual items to track at a given stage. This assumption is partially valid for some cases. In the staging by Plaisant et al. [Pla+12], the animation separating elements into two lists consists of a sequence of stages based on similarities between elements: identical, unique, and similar. This strategy does not emphasize changes in the visualization. Similarly, in the staged animation between a scatterplot and bar chart by Heer and Robertson [HR07], where the first stage moves points to their x-coordinates and updates the x-axis, and the second stage morphs the points into bars, we can observe that separating position and shape changes does not reduce the number of items to track at each stage.

However, in most cases these two strategies are not mutually exclusive. As an illustration, the staged dynamic network in GraphDiaries [BPF14] animates first the disappearing nodes, then the remaining nodes, and finally the appearing ones between two timesteps. This staging does not only separate and highlight the major changes happening in the network, but it also reduces the number of items to follow during the animation.

6.4.2 Staging Design

Once a staging strategy has been determined, the staging design process comprises of two phases: (1) decomposition of the animation into stages and (2) timing of the stages.

Decomposition refers to the alternatives a designer can choose from when deciding upon the criteria used to break the animation down into multiple stages. For example, when animating salary trends over time for a population, a designer can choose to decompose

animation into stages based on the job sectors or the different trends of salary evolution to support different types of temporal analyses.

The timing parameters of staging comprise of ordering and scheduling. Ordering describes the order in which the stages will be played, whereas the scheduling defines the relative temporal relationships, duration, and delays of the stages.

The choice of decomposition and timing parameters is highly — but not entirely — steered by the staging strategy selected by the animation designer, and thus the fixed dimensions of our design space as we will detail in the remainder of this section.

Decomposition

Staging consists of decomposing an animation into steps. Naturally, a critical design decision therefore lies in choosing what criteria will govern the decomposition of data items to animate into subgroups. Depending on the changes the visualization undergo, and the goal that the staging is aimed to address, decomposing an animation into stages can be based on: (1) data attribute value, (2) data attribute change or (3) (re)presentation change.

Decomposition based on data attribute value

An animation can be decomposed into stages based on the values of a stable or evolving data attribute.

Stable data attributes are attributes that are not subject to changes over data timesteps. These include for instance categorical attributes such as name and numerical attributes such as identification numbers. Decomposing based on such stable attributes consists of grouping data items by value (e.g. male vs female), and associating each sub-group to an animation stage.

Decomposition can also be based on the value of an evolving data attribute. In this case, the designer must decide on a single data timestep as a reference for the decomposition, i.e. the initial step, the final step or any intermediary step. For instance, since salaries can fluctuate over the years, a decomposition based on “high salaries” (e.g. salaries > 80,000) requires a choice of which data point to use for the grouping: most recent salary, oldest salary, or anytime in between.

Decomposition based on data attribute change

Stages can also be defined by decomposing the data items to animate based on the changes in data attribute values over time.

In the case of an evolving data attribute, data items can be grouped based on a particular evolution pattern (defined by a predicate). For instance, increasing, decreasing and stable evolution could form the basis of staging based on salary change (see Fig. 6.1-b). Note that the nature and complexity of evolution patterns the designer can choose from depend on the data attribute type (categorical, ordinal, quantitative), and the number of data states (2 or more). In particular, if the data attribute is quantitative, then an evolution pattern can be defined based on the value, direction, rate, and shape of change (i.e. monotonic, non-monotonic, converging, fluctuating, etc.) [APS14]. For example, a “zigzag” pattern where value drops then increases can only be defined on an ordinal or quantitative attribute, that evolves over $n > 2$ data states.

Decomposing data items based on their values at a specific data state, on their overall evolution between initial and final data state, on a more elaborate evolution pattern over several data states, or a combination of the above, offers a wide range of possibilities to create various staging strategies to highlight particular groups of data items and/or behaviors.

Data attribute change also encompasses change in presence, since in some cases, data items can appear, disappear or simply evolve across data states. For instance, a common staging strategy between 2 data states consists of first animating data items that disappear (reducing

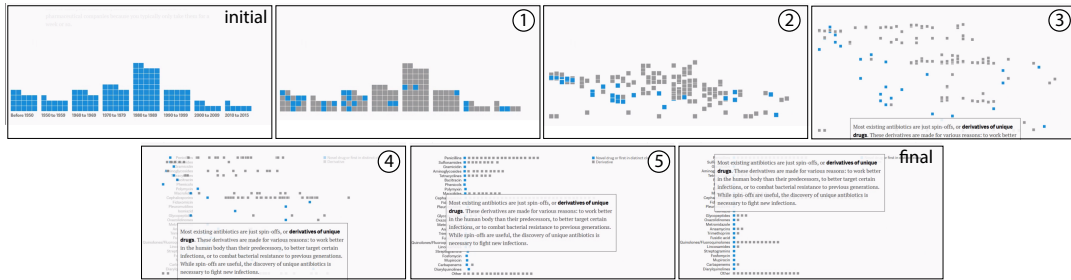


Figure 6.3: Staged animation from an initial visualization showing the emergence of antibiotics in the US over years, to a final visualization presenting a list of all antibiotics, showing the original drug (blue square) and its derivatives (gray squares). (1) labels disappear and several blue squares (antibiotics) change to gray (derivatives). (2-3) squares shrink, disperse and organize into groups, aligned horizontally; (4-5) labels and legend appear while squares continue to move to form groups.

visual clutter), then making all remaining data items evolve to their target values (helping focus on evolution between the initial and final data states), and finally adding all new data items (helping apprehend novel items in their context without other distraction). A descriptive example of this decomposition approach is the three-stage animation in GraphDiaries [BPF14] animating node and edge removal first, followed by node repositioning and finally node and edge addition.

Note that, as with evolving values, when additional intermediary data states are involved, considering presence change across n (> 2) data states introduces complexity since data items could non-uniformly appear and disappear across the states. In this case, the designer could either consider a simple overall pattern between the initial and final state, or specify more elaborate presence evolution patterns such as “disappear no more than twice across the n states”, or “appear after year 2012, and stay in at least until 2015”.

Decomposition based on (re)presentation change

When the representation or presentation undergo changes, staging can be applied to animate these changes in steps. For instance, if the user decides to map the x-axis to job sector instead of age in our example of Figure 6.1, a staging based on representation change could first remove the labels from the x-axis (corresponding to the age dimension), then move the data points horizontally so that they are organized by job sector, then finally reveal the labels of the new axis (corresponding to the job sector dimension). In this case, the data is not undergoing any change, only the framing elements of the visual representation (axes, labels, grid, etc.) are concerned.

We can find in the literature and in practice examples of this decomposition approach. For instance, when navigating between two scatterFigures of the the same data, Scatterdice [EDF08b] employs a three-state animation consisting of extrusion to 3D to represent, on a z-axis, the new dimension to display, then rotation of the cube to present the cube face with the new dimension, and finally projection to 2D to form the new scatterplot. Another example from data journalism is shown in Fig. 6.3.

Impact of fixed dimensions on decomposition approaches

The choice of an appropriate decomposition criteria depends highly on characteristics of the data attributes and the visualization as well as the type of changes they undergo. The number of these changes can also have a considerable impact on the appropriateness of the decomposition criteria. This assumption was advanced, for instance, by Zaman et al. [ZKS11] who argued that the staging strategy emphasizing changes (“disappear, evolve, appear”), adopted in their DARLS system, could become counterproductive when the data items undergo a large number

of changes. To address this, they suggested a finer decomposition that further breaks down the changes into smaller sets. However, it is important to warn that finer staging does not necessarily induce a better and more comprehensive animation as it was pointed out for example in DynaVis [HR07] where the extreme “do one thing at a time” staging proved to be overly complex generating unnecessary motion and causing confusion to the users. While an increase in the number of changes challenges the potential of a staging decomposition approach to “emphasize changes”, larger cardinalities of visualizations involving additional visual items to animate calls into question the extent to which a given decomposition approach can still “simplify object tracking” as revealed by Heer and Robertson [HR07]. An example of decomposition alternatives to address this limitation is suggested in the multi-step animation by Plaisant et al. [Pla+12] where the authors animated in the early design all identical drugs at once, but ended up splitting the stage further by animating one pair of drugs at a time to make the animation less overwhelming.

Observer task indirectly impacts the choice of a decomposition criteria through the chosen staging strategy. It can also take a considerable part in the specification of the decomposition by determining for example which changes to emphasize to answer a specific question, or what aspects in the visualization will compose the sequence of a staged visual story. The staging scenario in Section 6.6.2 further illustrates how the observer task impacts the choice of the decomposition criteria.

Timing

Timing discusses alternatives with regards to the temporal parameters of staging, namely, ordering and scheduling.

Ordering

Animation stages can be ordered to: reduce visual complexity, favor the visibility of changes and data properties, or express a narration.

Visual complexity is partly characterized by the number of data items and visual clutter resulting from such high cardinalities [Oli+04]. Hence, to reduce visual complexity, stages should be ordered such that stages reducing the number of visual items to be tracked or lessening the visual clutter gain precedence over other stages. Ordering in this manner supports the simplify object tracking staging strategy. Illustrative examples include animating the disappearing visual items in the first stage [HR07; BPF14; Got; Ger; Fil] and starting with the trim or collapse operations in order to reduce overlapping between items [PGB02; GVM12; EL00; BCK05]. In more particular cases, this kind of ordering is intrinsically induced by a decomposition approach explicitly aiming at simplifying object tracking [FE02; EDF08b].

To favor the perception of changes in data attribute value or (re)presentation, stages can be ordered based on the degree of interest or importance of the changes. This ordering alternative is in accordance with the staging strategy that emphasizes changes, and was adopted for instance by Guilmaine et al. [GVM12] who recommended keeping the changes that are more likely to be important to the user visible toward the end of the animation. Similarly, to favor the perception of specific data attribute values, stages can be ordered based on the degree of interest or importance of these values. In addition, staging order could also be inferred based on the inherent order in the attribute values, if any (e.g. from higher to lower value). For instance, Plaisant et al. [Pla+12] animated drugs in the order that corresponds to how users would order them in real world.

Lastly, stages could be ordered to make narrative sense instead of focusing on data and (re)presentation characteristics, as the sequencing of information can be used to convey

different meaning in narration [Hul+13]. This approach was adopted by Franklin et al. in their narrative system for medical treatment exploration [Fra+14].

Scheduling

When scheduling animation stages, a designer needs to define the duration of each stage, and whether and how the different stages overlap in time, i.e. by defining delays before each stage. Possible scheduling alternatives include simultaneous, parallel, and sequential scheduling — more finely described by the time relationships described by Allen’s [All84] interval algebra (Fig. 6.4).

Several strategies can be considered when defining the duration of stages: fixing the same duration for all of the stages, or tailoring the duration of each stage based on different rationales such as reflecting the importance of the stage [Pla+12; GVM12], reinforcing apprehension [Pla+12], and adjusting the duration relative to the number of items undergoing changes [Pla+12; GEY12] or relative to the overall amplitude of change the items undergo in a stage [GVM12; BPF14; BCK05].

Delays can be introduced to better direct the attention of the observers, by helping with apprehension [Pla+12; GEY12; BCK05] and reinforcing the separation of changes so as to reduce cognitive load as aimed by the staged animation [BPF14; Fis10]. Delays can also be employed to provide emphasis and support narration when designing for the “present” observer task [Pla+12].

Duration and delay together define the scheduling of the stages in relation to one another. On the one end of the spectrum, simultaneous scheduling (*equal*) implies that stages are played at the same time, and for the same duration, which does not result in a staged animation. On the one end of the spectrum, sequential scheduling (*meet, before*) implies that a stage is effectuated only when a prior stage has finished its execution fully. The delay further defines whether a pause is introduced between the stages (*before*), which can further reinforce separation of stages. In between these two extremes, consecutive stages can be scheduled in parallel (*overlap, during, start, finish*). Parallel execution of animation stages is typically employed to reduce the overall duration of the animation compared to a sequential scheduling. Guilmaine et al. also specifically discussed the usefulness of the *during* scheduling to enhance the visibility of changes [GVM12].

Note that our design space is inclusive of staggered animation. Staggered animation is a particular class of animation that consists of adding an incremental delay in start times across elements [CDF14; HR07; Mil; Bub; Fil; Inc]. Staggering can be thought of as an extreme case of staging, where each data item is a group (i.e. decompose by data element id) and items are scheduled so that every pair of consecutive data items satisfies an *overlap* time relationship. Staggered animation is often presumed suitable to simplify object tracking [HR07], but has been found unhelpful for the task [CDF14]. Staggering is, however, a popular animation effect in visualization for its aesthetic appeal [CDF14; Bre+17] which explains its wide use in practice, especially in data journalism and web advertising [Ger; Iot; Hap; Jd1; Vn1].




Figure 6.4: The three scheduling alternatives, based on Allen’s seven primitive time relationships between two stages [All84].

Table 6.3: Classification of staged animations collected from data journalism, oral presentation applications and web advertisements

Ref	Fixed Dimensions							Variable Dimensions							Scheduling (Interleaving)		
	Data states			(Re)pres. change	Data Att. change	— Observer Task —			Strategy		— Decomposition —			— Ordering —			
	1	2	n			Understand	Verify	Present	SOT	EC	DAV	DAC	RPC	RVC		FVC	EN
Happy Forecast-Postcodes [Hap]	•			•		•			•		•		•	•			meet,overlap,equal
Happy Forecast-Regions[Hap]	•			•		•			•		•		•	•			meet,overlap,equal
Motion Graph-BubbleToBar [Mot]	•			•		•			•		•		•	•			meet,overlap
Motion Graph-BarToLine [Mot]	•			•		•			•		•		•	•			meet,overlap
EPFL Data [Epl]	•			•		•		•	•		•		•	•	•	•	overlap
FW4 Visualization [Fw4]	•			•		•			•		•		•	•			overlap
Small cities [Sma]	•			•		•			•		•		•	•			meet,equal,overlap
SandDance-Dimensions Change[San]	•			•		•			•		•		•	•			overlap
SandDance-Facets[San]	•			•		•			•		•		•	•			overlap
Fallen of WWII-Show[Fal]	•			•		•		•	•		•		•	•	•	•	overlap
Fallen of WWII-Hide [Fal]	•			•		•			•		•		•	•	•	•	overlap
Fallen of WWII-By country [Fal]	•			•		•			•	•	•		•	•	•	•	meet,equal,overlap
Fallen of WWII-Years[Fal]	•			•		•			•	•	•		•	•	•	•	meet,equal,before,overla
Fallen of WWII-Categories [Fal]	•			•		•			•	•	•		•	•	•	•	meet,equal,overlap
Robin Hood-Dots[Rob]	•			•		•			•		•		•	•			overlap
Robin Hood-Timesteps[Rob]	•	•		•	•	•		•	•		•		•	•	•	•	meet
German Elections-Between Candidates [Ger]	•			•		•			•		•		•	•			meet,overlap
German Elections-Overview [Ger]	•			•		•			•		•		•	•	•		meet,equal,overlap
German Elections-One Candidate [Ger]	•			•		•		•	•		•		•	•	•		meet,equal,overlap
German Elections-Timeline [Ger]	•	•		•		•			•		•		•	•			meet,equal,overlap
Explain Histogram-Barchart [Exp]	•			•		•			•		•		•	•			overlap
Explain Histogram-Dots[Exp]	•			•		•			•		•		•	•			overlap
Explain Histogram-Sort[Exp]	•			•		•			•		•		•	•			overlap
Global Ad Spend [Glo]	•			•		•		•	•		•		•	•	•		meet,overlap
Health Law-Dots [Heab]	•			•		•			•		•		•	•			overlap,start
Health Law-Linechart [Heab]	•			•		•		•	•		•		•	•	•	•	meet,overlap
Health Law-Vertical [Heab]	•			•		•			•		•		•	•			meet,overlap
Longevity of No Diggity [Lon]	•			•		•			•		•		•	•			overlap
Future of Hits [Hit]	•			•		•			•	•			•	•	•		meet,equal
Life on Earth [Lif]	•			•		•			•		•		•	•			meet,overlap
CIA world browser [Cia]	•			•		•			•		•		•	•			meet,overlap
Map your moves [Mov]	•			•		•			•		•		•	•			meet
One Year of Ebola [Ebo]	•			•		•			•		•		•	•			meet,overlap
Pop lyrics [Lyr]	•			•		•		•	•		•		•	•		•	meet,equal
How is Life [Howa]	•			•		•			•		•		•	•	•		meet,equal,overlap
Income taxes [Ine]	•			•		•			•		•		•	•			overlap
Use of time [Use]	•			•		•			•		•		•	•			overlap
First time [Fir]	•			•		•			•		•		•	•			overlap
Languages similarities [Sim]	•			•		•		•	•		•		•	•		•	meet,equal,overlap
Miles Wikipedia [Mil]	•			•		•			•		•		•	•			meet,overlap
Packed bars [Pac]	•			•		•			•		•		•	•			meet,equal,overlap
Paris attack [Par]	•			•		•			•		•		•	•			overlap
Game of Thrones-Seasons [Got]	•	•		•	•	•		•	•		•		•	•	•	•	meet
Game of Thrones-Filtering[Got]	•			•		•			•	•			•	•	•	•	meet
How much say [Howb]	•			•		•			•		•		•	•			meet
Best tennis season [Ten]	•			•		•			•		•		•	•	•	•	meet,overlap
Free Willy and Flipper by Numbers [Wil]	•			•		•			•	•	•		•	•	•	•	meet,overlap
US lead the world [Us1]	•			•		•			•		•		•	•			overlap
Craft beer [Cra]	•			•		•			•		•		•	•			meet,overlap
Gapminder-Human Development Trends [Gapb]	•			•		•	•	•	•		•		•	•	•	•	meet,overlap,during
Full Metal Jousting [Met]	•			•		•			•		•		•	•			overlap
Keynote chart animation [Key]	•			•		•			•	•			•	•	•	•	meet,equal
PowerPoint chart animation [Ppt]	•			•		•			•	•			•	•	•	•	meet,equal
JD & Co [Jd1]	•			•		•			•		•		•	•			overlap
VN Star [Vn1]	•			•		•			•		•		•	•			overlap
Heart Health in China [Heaa]	•			•		•			•		•		•	•			overlap
Nike SB Dunk [Nik]	•			•		•			•		•		•	•			overlap
Avec qui skier? [Ski]	•			•		•			•		•		•	•			overlap




Figure 6.5: Comparison of a full (top) vs. staged (bottom) animation of Gapminder data (x-axis: fertility rate, y-axis: life expectancy, size: population, color: region). In full animation, all data points evolve simultaneously through their n states. Each frame of the animation shows all countries at the same year of their evolution. Staged animation decomposes by region across n states. The first and last frames of the animation correspond to the initial (year 1800) and final data state (year 2016). Each other frame shows data points at different time of their evolution. For instance, the frame at the second position shows the American countries (in yellow) during their evolution from year 1800 to 2016, whereas all other countries are still displayed with their state in 1800. Subsequent frames show the animation for the rest of the regions.

6.5 Values of Animation Choreography

We demonstrate the descriptive, comparative and generative value [BL04; BL00] of our design space through a classification of over 70 examples of staging approaches curated from the visualization literature (Table 6.1) and from practice (Tables 6.2–6.3). The classification helps identify common approaches to the design of staged animation, permits comparison across specific instances, and points out limitations and opportunities for new uses of staging. We discuss our curation and classification methodology, followed by the main insights gained from the analysis of our corpus, leading to novel design alternatives and new research opportunities.

Methodology

We first compiled a set of common keywords used to refer to staged animations (see Section 6.1.1) through a survey of research articles pertaining to “animation” and “animated transition”. Using these keywords as search terms, we curated papers on staging from academic digital libraries (ACM, IEEE) and Google Scholar, finding 18 articles that employ (and in some cases, evaluate) staged animation. We then looked for examples of staged animation use in practice. We reviewed the corpus of 80 visual data stories curated by McKeena et al. [McK+17] and retained the 15 that used staged animation. We broadened our search and collected 42 examples from established news media and online design sources. In total, we curated 75 distinct instances of staged animations (several of the above employ multiple staging approaches). We included all examples from practice and additional unclassified examples in a companion website (<https://stagedanimation.github.io/>).

Through careful consideration of the rationales driving the design of staged animation instances and their context of use by two independent coders, we determined the fixed and variable dimensions of these staging instances guided by the design space. Analyses of the coded literature and practice instances led to a set of insights (I) about existing approaches to staging design and raised new research opportunities (RO) discussed below.

I1: Focus on staging over 1 and 2 data states

Staged animation instances from both literature (see Table 6.1) and practice (see Tables 6.2–6.3) have mainly been used to present changes over 1 or 2 data states. GraphDiaries [BPF14] is the only example of staging that presents data changes across n data states. However, the animation design consists of staging the changes *pairwise between consecutive data states*. The resulting animation of the evolution of data across n data states then consists of a chain composed by the repetition of the same staged animation over 2 consecutive data states. This noticeable prevalence of staging over 1 and 2 data states unveils the first following research opportunity.

ROI: Value of staging over n data states

The research community has yet to explore the value of a staging approach applied over n data states. Supporting focus on the evolution of subsets of the data could help with visual tracking, and facilitate identification and comparison of temporal trends.

Figure 6.5 (bottom) shows an example of a staging design where the different stages correspond to subsets of the data (i.e. regions of the world) that are animated to reflect their evolution over the n states. The resulting animation first show the evolution of the American countries only (yellow circles) *from beginning to end*, that is, from 1800 to 2016 while all other countries remain at their initial state in 1800. Then, “East Asia and Pacific” countries (red circles) are animated over their n values, from 1800 to 2016, while American countries already reached their final state (2016), and the rest are still at their initial state (1800). This allows to focus on the evolution of each region, one at a time, and contrasts with the GraphDiaries example, where the same staging is applied n times between consecutive data states.

We believe that applying staging over n data states can be beneficial in various cases, as for example to isolate and compare temporal trends, or support visual presentation and storytelling focusing on the evolution of different groups. More generally, when authoring a presentation or a narrative discourse, a designer may want to show sequentially the overall evolution of each subset of the items. For instance, with reference to the Gapminder example, she may want to tell the audience how each continent evolved across years in terms of fertility rate and highlight the main trends characterizing this evolution. We believe that achieving this goal would be more appropriately supported using a staging approach that animates subsets of the data across their n states of evolution.

This approach, however, should be considered with care, as it is breaking the concept of the global timing into a sequence of relative timelines, which may affect the temporal comprehension during the animation (i.e. some data may already have been animated to their final state, whereas other data points are still presented at their initial state and remain to be animated). In other words, at any given frame of the animation (except from the first and last frames), the contextual cues are incorrect. That is, it does not make sense to examine the positions of American countries at 2016 against the backdrop of other regions at 1800. To cope with this problem, and avoid the issue of “uninformative animation” [HR07], it is important to provide cues for understanding the current visual state of the data points. For example, in the prototype we developed for exploring staged animation designs (Fig. 6.10), we use a timeline that explicitly conveys the scheduling of the different groups of points.

Our design space allows us to identify staging designs that have not yet been explored. Further research is needed to explore staging strategies in this particular data context, to clarify the real benefits of applying staging across n data states, and possible solutions to the above problem of temporal inconsistencies.




Figure 6.6: Staged animation of Gapminder data. As in Fig. 6.5, the staging decomposes the animation by region, one at a time. As the status panel at the bottom shows, the “America” and “East Asia & Pacific” regions have already been animated, the “Europe & Central Asia” is being animated and “South Asia” “Sub-Saharan Africa” remain to be animated.



Figure 6.7: Six-stage animation showing the evolution of characters and their relations between 2 episodes (initial) and (6) of a TV-show [Got]: (1) edges disappear, (2) cluster hulls disappear, (3) nodes disappear, (4) remaining nodes (characters common to the two episodes) are repositioned to leave room for appearing nodes (new characters), (5) cluster hulls appear, and finally (6) edges appear to form the final network.

I2: Limited exploration of decomposition approaches

We observed a wide spectrum of decomposition approaches for staging in the wild, ranging from the standard “disappear, evolve, appear” to those involving various types of data attributes [MI1; Bac; Jd1] and emphasizing changes in data attribute values [Got; Rob; Ger] (e.g. see Fig. 6.7), representation [Mot; San] or presentation [Hap; Fw4; 3d1; Ten]. In contrast, and despite of the use of staging in varied contexts, previous research work has adopted a limited range of decomposition approaches. More particularly, “disappear, evolve, appear” and its context-specific variants, such as “trim, reposition, expand” and “collapse, permute, expand”, are the most prominent approaches [BPF14; ZKS11; PGB02; GVM12; EL00]. In most of the other cases, staging follows a decomposition based on (re)presentation changes (e.g. [FE02; EDF08b]) whereas for a few others the staging reflects decomposition based on data attribute value (e.g. [Pla+12; RPD09; HR07]). The widespread use of this approach when designing staged animation for diverse visualization types, such as networks, trees, and diagrams, seems reasonable for various reasons.

The staging strategy guiding the prevalent “disappear, evolve, appear” decomposition approach emphasizes data changes in general by isolating the persistence and evolution of data elements. In addition, the stages’ ordering ensures that disappearance/trim/collapse of visual items are effectuated before evolution, leading to a decrease in the number of visual items to keep track of and thus the visual complexity during the evolution stage. Thus, this decomposition approach is commonly used when generating a staging that aims to support “understand” tasks, such as, visually tracking specific items during the animation [HR07; GVM12], tracking addition or deletion of nodes and edges [ZKS11; GEY12], and identifying the change trend for a group of nodes [BPF14].

This decomposition approach has also been used to support “verify” tasks, such as, finding groups of nodes undergoing specific changes [BPF14], and comparing and reconciling lists of drugs based on their similarities [PGB02]. However, we suspect that this decomposition approach may not be sufficient if the observer needs to answer more specific or complex questions, or wants to present the insights acquired through data analysis as an elaborated narrative. Additionally, this predominant decomposition approach does not transfer well to different types of visualizations, or different types of changes. Thus, the exploration of alternative approaches for decomposition to support a variety of observer tasks, for different visualizations as well as different data characteristics is essential and represents a relevant research opportunity.

RO2.1: Decomposing in various contexts, for different tasks

In many cases the prevalent approach of decomposition may not be effective, nor even possible because of the variation of the context or tasks to be supported. It might be more beneficial in such cases to tailor the staging strategy and design to the context and the task at hand. That leads us to believe that alternative staging approaches need to be explored to generate more opportunities.

Our proposal is partly validated by some previous works that adopted different approaches to decompose the animations into stages. A first relevant illustration is an animation where the items are staged into groups based on the Euclidean distance between them even though the visualization involves appearance and disappearance changes [RPD09]. Instead of presence change decomposition, Reitz et al. [RPD09] define stages between 2 states of a dynamic hierarchical network based on node position (data attribute value). Each stage thus comprises of nodes that are within a fixed Euclidean distance from each other between the two states. Their motivating staging strategy favored simplification of tracking of visual items and hence, the chosen decomposition approach reduces the number of elements to track while also restricting the movement to a limited area of the visualization.

A second example is the multi-step animation by Plaisant et al. [Pla+12] designed to support comparison of items in two lists of drugs. The usual decomposition approach based on presence change is not applicable to their case. Instead, to support the observer task of answering how the lists differ and to narrate the process of triage and collation between the two lists, they adopt a staging strategy focused on simplifying tracking of the list comparison process. They decompose items into stages based on their similarity and differences, resulting in 3 stages: identical, unique, and similar. The last stage is followed by a representation change that compacts the 3 lists generated by the previous stages.

To further motivate our proposal, let us get back to our initial scenario about the district's inhabitants. To gain an idea of the district attractiveness, we can start by staging following the standard “disappear, evolve, appear” approach to see specifically how many residents left, stayed or moved to live in it. Let us focus now on the residents who stayed in the district. Investigating the evolution of their income relative to their gender allows us to get an overview of the social situation of men and women. To achieve this, we can further decompose the “evolve” stage based on gender (e.g. men first then women, as illustrated in Fig. 6.1-a).

We may then ask more specific questions about the district such as the evolution of the number of residents by job sector to identify the most important job sectors in the district, and get an idea of the residents' qualifications. We further discuss the potential of novel decomposition approaches in a detailed scenario in Section 6.6.2.

From the above preliminary scenario, we can deduce that staging can be parametrized in terms of decomposition criteria (e.g. decomposing by presence and by gender) and their relations (i.e. decomposing by presence then by gender results in a different staged animation than decomposing by gender then by presence). To support the authoring of staged animation considering these aspects, we introduce a novel paradigm (discussed in Section 6.6.1) that models these aspects.

RO2.2: Many staging approaches for the same visualization

Most of the prior works employing staging adopt a single decomposition approach for a given visualization. But, in reality, we very often have several questions about the same visualization and analysis would be supported by seeing the data from different perspectives and at different levels of granularity, as already highlighted in the population scenario above. Hence, it can be interesting to have different choices of staging to explore different aspects of the same visualization. The most relevant example from literature that inspired this opportunity is the work by Guilmaine et al. [GVM12] where the authors suggested different ways of decomposing the same animated visualization: staging, level-by-stage, stage-by-level, hierarchical and hybrid. That leads us to ask further: how would the variation of the decomposition approach influence the reading and interpretation of a given visualization? This idea is partially inspired by the work of Pennington and Hastie [PH92] demonstrating that the variation of grouping approaches used in a narrative affects considerably the interpretation and judgment of the portrayed information by the observer. More concretely, their study showed that grouping court evidence by sub-stories led to more confident and unanimous decisions among jurors over evidence that is presented haphazardly (e.g. with grouping based on motives rather than temporal proximity).

I3: Limited exploration of timing parameters

Timing parameters, i.e. ordering and scheduling of the stages, are critical aspects of staging, especially when considering empirical evidence that narrative visualizations' sequencing can affect considerably the understanding and recall of information [Tho77; Hul+13]. However,

they remain under explored, as already pointed out in prior work [HR07; BPF14]. We discuss below different facets of these parameters.

Ordering

The most common ordering approach in previous research work aims at reducing visual complexity during the animation stages. Some exceptions include ordering to favor visibility of changes [GVM12] and specific data properties [Pla+12], or ordering to express a narration [Pla+12; Fra+14]. Interestingly, two previous works that order stages to favor visibility of change or data properties present contradictory recommendations for stage ordering. Guilmaine et al. [GVM12] recommend ordering the stages so that changes of interest are effectuated towards the end of the animation while the stages comprising of less important changes are executed at the start. Conversely, Plaisant et al. [Pla+12] consider the first stage as being the most important part of the staged animation. This contradiction is in fact not so surprising if we consider the diverging *primacy* and *recency* effects [Jr62]. Given a list of items to remember, the primary principle says that we tend to assume that items that are at the beginning of the list of items are of greater importance or significance, whereas the recency principle stipulates the opposite. More generally, the principle of “serial position effect” [Jr62] points out the tendency of an observer to recall the first and last items in a series best, and the middle items worst. This highlights that the effect of ordering is still poorly explored and understood in the context of staged animation.

In light of the above discussion, and given the yet under-explored variants of ordering approaches, we see an opportunity to further investigate the impact that different orders may have on interpretation of data, when it can support an argument, a point of view, or a narrative, or in the contrary, when it hinders the detection of existing relationships. We further discuss these two points below in the form of research opportunities.

RO3.1: Ordering can reveal hidden relations in data

Gershon suggested the animation of segmented components as a technique facilitating the exploration of fuzzy data [Ger92]. Although this animation technique was not presented in the paper as an instance of staged animation, we argue that its findings can be applied consistently to the design of staging. In this context, playing the animated sequence of components from high to low values of temperature then from low to high revealed respectively expanding and shrinking structures in the visualization. The authors affirmed that this type of animation can reveal the value, location and shapes of structures embedded in the data.

In the context of uncertainty visualization in fluid flows, Lodha et al. [Lod+96] suggest a system that proposes different alternatives of visually encoding uncertainty, including the animation of uncertainty glyphs along the path of the fluid particles. In this context, they propose two options: “ranked animations” and “priority sequencing”. The first option allows users to specify some criteria to rank the glyphs on the traces of the particle movement, resulting in stages presented in the order of this ranking. The second option, allows to show the animated glyphs in order of highest to lowest uncertainty. The authors argue this “priority sequencing” alerts the users to areas of high inaccuracy. Interestingly, they argue that it also identifies areas of equal uncertainty that are spatially apart by presenting them immediately one after one another.

In the map visualization by DiBiase et al. [DiB+92], a sequence of dynamic temperature maps was originally ordered in the chronological order of months. The authors reordered these maps based on the value of temperature from the months with least variation to the ones with the most variation of temperature. The duration of the animation of each map increases with the magnitude of the temperature variation. This re-ordering revealed a trend of maximum

uncertainty that had never been noticed before. By analogy, the following main question can be derived from all the above examples: how would the variation of ordering influence the perception and understanding of possible relations between the information being presented in the separate stages?

In a study examining in part the impact of sequential animation interactivity on decision making, Gonzalez et al. [GK97] found that users who had the choice to change the order of an animated sequence of images did better in terms of decision making because they had more flexibility in exploration, although the animation with fixed order proved to be easier to operate and less cognitively taxing. This specific study on image sequences leads to the need for a general investigation into the value of interactive (re)ordering in staged animation.

RO3.2: Ordering can suggest non-existing relations in data

In Diffamation, Chevalier et al. [Che+10] refrained from staging the animation of edits between text revisions because sequencing editorial operations may be interpreted as the real sequence of document editions. It is thus relevant to ask the following question: could the ordering of stages lead to misinterpretations, such as mistakenly perceive a chronological order, or other temporal relationships where there are not?

Furthermore, the *causality* as presented by Michotte [Mic63] considers that a moving object that contacts an initially stationary target is responsible for the subsequent movement of that target. If we transpose this principle to the context of staged animation, we may ask: could any causality relation be inferred between two consecutive stages?

Scheduling

The vast majority of prior research work that we surveyed used only non-overlapping stages (*meet* and *before*) Less commonly, animation stages were played in parallel. The *overlapping* scheduling approach was mainly used in the context of staggered animation [HR07; CDF14].

The above scheduling choices seem reasonable relatively to the contexts where they were adopted, but only cover a small subset of the possible scheduling alternatives, with a clear predominance of the *meet* scheduling approach. The rich space of alternatives remain to be explored. The variation of the scheduling could help better understand data or answer specific questions about it, by presenting it from different perspectives. These aspects are partly covered by the following research opportunity.

RO3.3: Scheduling to support data comparison

When we need to compare information conveyed by separate stages, a parallel scheduling is likely to be more effective than temporally separating the stages, especially when we refer back to the rule of thumb “Eyes beat memory” of Munzner [Mun14] asserting that “using our eyes to switch between different views that are visible simultaneously has much lower cognitive load than consulting our memory to compare a current view with what was seen before”. More questions arise when we think about the Gestalt Law of Synchrony, which suggests that visual items that change at the same time will tend to be perceived as belonging to the same group [PL98]. So, would playing stages in sequence or in parallel make the relations between data elements in separate stages more or less visible and easily perceived? Future research is necessary to better understand the role of scheduling in staged animation.




Figure 6.8: Example of a 3-level deep staging hierarchy, decomposing a full animation based on the criteria C_1 (salary value), C_2 (salary change), and C_3 (job sector value). The leaf nodes correspond to the resulting stages S_i , to be animated in order (depth-first traversal).

6.6 Describing Animation Choreography

The design space identifies the staging parameters to be specified and can thus be leveraged to guide the design of tools for authoring staged animation alternatives. We introduce the *staging hierarchy* paradigm to support authoring of staged animation. We then present a concrete scenario that illustrates the various insights and research opportunities discussed in the previous section, and demonstrates the usability of our framework.

6.6.1 Staging Hierarchy

Deciding how a full animation should be broken down into stages involves the specification of the decomposition criteria, and their relations. We introduce the *staging hierarchy* paradigm for specifying staging designs. Nodes at the different levels correspond to decomposition rules (i.e. “the items that satisfy X ”) to be applied successively to construct the different stages of the animation S_i (see Fig. 6.8).

More specifically, each level in the hierarchy corresponds to a decomposing criterion, for instance in Figure 6.8, C_1 corresponds to Decompose by Data Attribute Value of the *salary*. Each node at a given level of the hierarchy corresponds to a decomposition rule of the associated criterion (i.e. $C_{1,a}$: “the items that satisfy *initial salary* $\leq 60,000$ ”, $C_{1,b}$ “the items that satisfy *initial salary* $> 60,000$ ”). Adding subsequent levels implies further decomposition, for example, items with lower *salary* ($C_{1,a}$) are further decomposed based on C_2 , Decompose by Data Attribute Change of the *salary*, declined in three rules: rise ($C_{2,a}$), stable ($C_{2,b}$) and drop ($C_{2,c}$). Varying the number of criteria affects the granularity of the staging. Further, the order of the nodes at each level specifies the order in which the different decomposition rules will be considered, for example, the first level described in Figure 6.8 generates a staging where lower salary employees are animated first, followed by higher salary ones. Permuting the order of salary nodes $C_{1,a}$ and $C_{1,b}$ reverses the stages order into first high salary employees, then low salary ones.

Each node also has timing parameters, i.e. delay and duration, specified relative to the parent node.

The staging hierarchy fully describes a staged animation: the final stages correspond to the leaf nodes, the order and scheduling of the stages is obtained by a depth-first traversal of the tree (order) and combination of timing parameters (scheduling).

Note that each level should result in a partition of the items contained in the parent node, otherwise some items may not be animated. If the set of rules for a criterion is not inclusive of the data items in the parent node, a dummy node “the rest” should be added to guarantee that all items are eventually animated. We qualify these remaining items as *leftovers* and we will talk about them more in details in the remainder of this chapter.




Figure 6.9: Staging hierarchy for the co-authorship network scenario: (a) one-level staging based on the presence change of authors (Data Attribute Change); (b) one-level staging based on the topic (Data Attribute Value); (c) two-level staging based on the topic, then presence change of authors (Data Attribute Value and Data Attribute Change).

Note also that the hierarchy offers a lot of flexibility in that the rules at a specific level need not be consistent across. For instance, the items corresponding to people with an initial lower salary ($C_{1,a}$), and who got a raise ($C_{2,a'}$) are further decomposed by *jobsector*: within, then without the finance job sector ($C_{3,a''}$, $C_{3,b''}$), whereas people with an initial higher salary ($C_{1,b}$) who got a raise ($C_{2,a'}$) are decomposed based on the retail job sector ($C_{3,c''}$, $C_{3,d''}$).

6.6.2 Scenario: Exploring a Co-authorship Network

The following scenario illustrates the value of applying different staging strategies to support the exploration of evolving data. The corresponding staging hierarchy is depicted in Fig. 6.9.

Mike comes across a time-varying co-authorship network dataset detailing collaborations between visualization researchers per year, from 1995 to 2015. To gain initial insight into the data, Mike starts by developing an animated slide show of the dynamic network, using one node-link diagram per year (data states) whose nodes correspond to authors and edges encode collaborations. Node size is mapped to the number of papers published by an author at a given time, while edge thickness represents the number of papers two authors have collaborated on at that time. Mike indicates authors that undertake research on the same topic at a given time using node fill color, and generates a global layout for the network across all the data states to ensure node position stability between timesteps.

Mike plays the full animation to start his exploration. However, the network undergoes simultaneously multiple types of changes, including variations in the authors and their collaboration patterns, evolution of latent research areas and authors per research area, as well as changes in number of papers per author and pair-wise collaborations. Mike feels overwhelmed, as keeping track of all of the changes is difficult, if not impossible. In addition, multiple nodes and edges undergo these changes simultaneously at any given time, which impedes object tracking. To address these issues, he decides to break down the full animation into multiple stages, focusing on authors' presence.

Disappear, evolve, appear. Mike decides to employ the standard “disappear, evolve, appear” (Data Attribute Change) decomposition approach (similar to GraphDiaries [BPF14]), applied to each pair of successive years to see how the network evolves from year to year (Fig. 6.9 (b)). The resulting animation comprises of a series of staged animation across the n data states, each staged animation between two subsequent years comprising of 1) more transparent disappearing items (nodes and edges), 2) interpolating visual attributes (i.e., size, color) of the evolving items, and 3) less transparent appearing items. Mike is aware that this decomposition approach is popular for dynamic networks analysis, since the staging strategy here emphasizes changes to the network topology in general by isolating the disappearance, evolution, and appearance change. In addition, since the stages are ordered such that disappearing elements are removed first followed by evolution of the stable elements, the number of objects to track and consequently the visual complexity is reduced.

Exploring the data through the staged animation, Mike gets specifically interested in identifying what the main research topics are and how each changes over time. However, with the current decomposition, multiple topic clusters undergo a given type of change simultaneously making it difficult to track the evolution of a single topic.

Focus on topics. To answer his question about evolution of individual topics over the years, Mike decides on a staging strategy that simplifies tracking of elements belonging to a single topic. He thus decomposes the simple full animation into stages based on research topic, this time across the n states, generating one stage per research topic (Fig. 6.9 (b)). The resulting animation comprises 1) animation across all n years, of the elements whose topic value is equal to Topic 1, *at least once across the n states*, 2) similarly with Topic 2, and so forth. This staging strategy and design better support trend analysis of individual topics (does interest in a topic grow, shrink, remain stable or unstable over time? when the topic had the most number of researchers?). However, this approach may be misleading, since the decomposing rule across the n states may not result in a partition of the elements. For instance, an author that switched topic will satisfy more than one rule on topic value, and thus be eligible for more than one stage, which raises an interesting challenge: should we allow the same element to belong to different animation stages, and then have the same node be animated from beginning to end several times (but this makes parallel staging such as meet and overlap time relationships impossible, unless the visual element is duplicated)? Or should we ignore elements that are already covered by a previous stage, and therefore not faithfully represent data satisfying a criteria? These are some of the many challenges when coping with n data states. Since Mike's focus is on the evolution of topics, he relaxes the constraint on animating every element just once, and allow items to revert back to their initial state whenever necessary.

Order topics by recent interest. To favor the visibility of general interest in topic over time, and specifically the most recently popular topics, Mike further decides to order the topics based on decreasing number of researchers interested in the topic in 2015. For each topic, Mike decides to further emphasize the change in interest between successive years. He adds a second level of staging that employs the “disappear, evolve, appear” decomposition approach (Fig. 6.9 (c)).

This multi-level staging generates the following stages: (I) for Topic 1, for each pair of subsequent years: (I.1) disappear (nodes and edges leaving Topic 1 in the following time step fade out), (I.2) evolve (nodes and edges that remain in the Topic 1 change in size, etc), (I.3) appear (nodes that enter the Topic 1 fade in). Then (II) for Topic 2, for each pair of subsequent years: (II.1) disappear, (II.2) evolve, (II.3) appear, and so on for the remaining topics.

Focus on trends. Using the latest staged animation, Mike continues with his analysis of topic trends (shape of change) and notices that two of the topics (1 and 3) exhibit a continual increase in interest over the years while the interest in Topic 2 first fluctuates then stabilizes over time. To favor the visibility of the similarity of the trends, Mike changes the ordering of the stages so that the topics that have similar trends (Topics 1 and 3) are effectuated consecutively followed by the more varied trends (Topic 2).

Mike now knows that both Topics 1 and 3 show a steady increase in interest over time, whereas Topic 2 has received varied attention over the years. The staged animations he produced, however, make comparison between topics difficult since they are animated in sequence. Mike needs to take note of trends and collate this information for multiple topics. Mike notes that he will have to further explore with various ordering and scheduling to facilitate comparison.

6.7 Authoring Animation Choreography

We implemented an authoring tool that allows to test staging alternatives within a dynamic bubble chart visualization. We will start by discussing the design considerations that we encountered and thought about during the implementation process of our tool. We will then provide a description of our tool.

6.7.1 Design considerations

The main design decision that we adopted for implementing our authoring tool concerned the following aspects: animation timing, generation of animation sequence, and the management of staging leftovers. We will discuss these aspects in details below.

Staged animation timing

As explained above, defining the timing of staged animation involves setting the value of the start and stop time of the stages, their duration and the eventual delays between them. Different approaches can be adopted to define all these different parameters. The first approach and most naive one consists of fixing a global duration of the animation and then divide this allocated time on the number of stages to get the duration of stages and the eventual delays between them. When adopting this approach, we can end up with some stages where little change is happening and other ones where too much change is happening in a very short time. That may generate an excessively slow animation that can be boring and less engaging or inversely a very fast animation that can be frustrating and difficult to follow. To remedy to this limitation, a more adapted approach can consist of partitioning the duration time between stages based on a set of criteria such as for example the quantity or amplitude of changes happening at every stage, the degree of importance or interest of each stages as discussed in detailed in the Section 6.4.2 and then insert, if possible, delays between these stages.

A third approach follows the reverse way by starting first with defining the duration of stages and the delays between them so that the final global duration of animation corresponds to the sum of all the stages' durations and the delays. A possible limitation of this approach is that we can end up having a long global duration of animation that cannot be adapted for certain constraints related for example to the overall timing of interaction with the animated visualization [HR07].

In our tool, we adopted a different timing approach. We define t_{Start} , t_{Stop} (i.e. the duration) and delay by a *referential timing*: recursive timing of children in reference to the parent and to the global animation. The global animation duration is set by default to 5 seconds and can be customized via the graphical interface. With reference to the staging hierarchy described in Section 6.6.1, we normalize the animation start and stop times of the parent level duration between 0.0 and 1.0. The start and stop times and the delays of the the underlying levels are defined relatively to the parent level duration. Figure 6.11 illustrates an example where the first node of the second level starts at the same time as the first parent level and finished at the middle of the parent level duration corresponding to a stop time equal to 0.5. The second node of the second level is animated after a delay of 0.05 with reference to the stop time of the first node's animation and its animation will last until the end of the parent level duration. This relative timing approach allows to set both "positive" and "negative" delays where a given stage can start n seconds after or before the end of the previous stage respectively.

Generation of animation sequence

The animation sequence can be generated following different approaches. One possible approach is to calculate a global ranking of all the visual objects composing the visualization. The ranking index will define the order following which every object will be animated. Objects sharing the same value of rank will be animated simultaneously and belong, hence, to the same animation stage. In our authoring tool we adopted a different approach. The *CompoundAnimation* module calculates the start and stop time for each visual object with reference to the global animation.

Managing leftovers

We define *leftovers* as the visual objects that do not fit to the any of the staging criteria defining each stage. There are various reasons for which leftover can exist. These reasons can be unintentional or intentional. Incidentally, an animation designer can forget to consider the characteristics of a subset of visual items when setting the staging predicates, which leaves them without any corresponding staging group. Conversely, she may want to filter out the leftover items in purpose because they are not relevant to the information that she wants to convey or the story she wants to tell. This case can be seen as a sort of “implicit” or “programmed” filtering.

In these two cases, leftovers will end up not being shown in the final staged animation. In another different case, an animation designer may just want to give less focus on the leftover items because they are less relevant to the information she wants to convey or the story she wants to tell. In that case, leftovers can be animated in an additional stage at the end of the corresponding stages or at the end of the whole staged animation. For example, let us refer back to the population scenario presented at the beginning of this chapter. Let us consider the following extended list of job sectors: health, education, industry, finance, business and retail. Imagine that the designer wants to design a staged animation that portrays separately the evolution of salary of men and women with a particular emphasis of men and women working in the three following job sectors: health, education and industry. One relevant staging approach can be to show in a first stage the evolution of men working in these three job sectors, followed by a stage animating the rest of men working in the rest of job sectors (e.g. finance, leisure, agriculture, etc.), then animate the women working in the three job sectors, and finally animate all the rest of women working in the rest of job sectors. Alternatively, if she is mainly interested in enforcing the comparison between the evolution of men and women in these three job sectors compared to the rest of sectors, a more convenient staging approach can be to animate in the first stage the men working in the three job sectors, and the second stage the women working in the three job sectors, and finally animate all the rest of men and women working in the remaining job sectors.

In all the cases, the animation designer should be warned/notified about the existence of leftovers and should be able to identify them easily. Identifying leftovers involves knowing their exact number and being able to locate them in the visualization interface. Various options of visual feedback can be applied to help the designer identify the leftovers. A possible option can be to display, before the beginning of the staged animation, special and distinctive visual marks in the initial locations of the leftover items, or/and in their final locations after the animation ends (see examples in Figure 6.11). A second option can be to alter the original visual aspect of the leftover items to make them easily noticeable such as highlighting them in a specific color, shrinking their size, reducing their luminance or making them transparent. In our current prototype, we display the number of visual items at each stage and the number of the eventual leftovers.



Figure 6.10: Gapminder data displayed in our prototype. The chosen staging decomposes the animation by region, one at a time. Here, the “America” and “East Asia & Pacific” regions have already been animated, the “Europe & Central Asia” region is being animated and the “South Asia” and “Sub-Saharan Africa” remain to be animated.

6.7.2 Prototype

We have started implementing the above concepts in a prototype written in [Python3](#) using the [PyQt](#) bindings for the [Qt](#) toolkit. In its current state, this prototype allows to select a dataset, a renderer for this dataset, and a staged animation. The animation can then be played for a specified duration by clicking on a button. The user can alternatively scrub a time slider back and forth to closely inspect it.

Datasets can be randomly generated at runtime or loaded from JSON files. Our custom JSON format specifies the Python class to use for data pieces and then describes each of them with a list of static attributes and a list of states giving the value of other attributes at specific points in time. For example, our JSON file for [Gapminder](#) data describes countries with a static name (e.g. Zimbabwe), code (e.g. ZW) and region (e.g. Sub-Saharan Africa) and varying attributes such as fertility rate, population and life expectancy. Renderers are Python classes that can visually represent a particular kind of data by mapping its static and dynamic attributes to visual ones, e.g. the fertility rate and life expectancy to a 2D position, the population to a circle size and the region to a label (Figure 6.10).

Axes are automatically added to the view as well as a background white rectangle that materializes the 2D bounding box for the chosen dataset and renderer. The view can be

```

{
  "description": "America",
  "start": 0.0, "stop": 0.16666666,
  "criterion": "lambda p: p.region=='America'",
  "levels": [
    {
      "description": "Fertility rate > 6",
      "start": 0.0, "stop": 0.5,
      "criterion": "lambda p: p.initial.frate > 6"
    },
    {
      "description": "Fertility rate <= 10",
      "start": 0.5, "stop": 1.0,
      "criterion": "lambda p: p.initial.frate <= 6"
    }
  ]
},

// ... similar code for the other regions

```

Figure 6.11: Partial listing for a Gapminder animation staged with two levels, the first one depending on the country region and the second one on the initial fertility rate.

panned and zoomed at any time using the mouse. When animating, our prototype linearly interpolates attribute values over time before mapping them to visual attributes, if possible (e.g. for numeric values), or falls back to a nearest approximation otherwise.

Staged animations are described in a custom JSON format that follows the hierarchical approach described above. , a criterion to select the objects to which it applies, and a list of sub-levels (Figure 6.11). Criteria are given as Python functions that take a data piece and return a boolean value that can depend on any combination of its attributes values at any time, e.g. at the first time for which it is specified in the dataset (`p.initial.population`), at the last (`p.final.population`) or at any other particular time (e.g. `p.at(2016).population`).

6.8 Conclusion

The holistic design space describing the composition of staged animations forms the core contribution of this chapter. The proposed design space not only provides the necessary theoretical grounding to understand and contextualize the varied staging approaches seen in previous works, but also enables discussion on trade-offs between these designs. In particular, our work is the first to initiate the discussion around the possibilities of staging across n data states to emphasize particular changes or data characteristics, and the challenges associated with dealing with multiple data states.

The animation choreography design space sheds light on opportunities for exploring novel staging approaches that go beyond facilitating visual tracking, to support exploration of dynamic data as well as rich narratives through animated visualization. Our discussion on staging calls for empirical evaluation to assess the benefits and caveats of staging in various contexts. Identification and formalization of the parameters of staging design paves the way for a more systematic evaluation of the effects of the different staging parameters on comprehensibility and memorability of the data for varied tasks. Discussion on staging an animation across 2 or n states also presents design trade-offs (temporal inconsistency vs. reduced visual complexity) that need to be empirically tested and further explored.

Staging, when considered in the context of theater and cinema, comprises the notion of emphasizing elements of interest through the use of light, motion, and so forth. Thus, another aspect to consider when designing animation stages is that of drawing attention to visual items active in a given stage (belonging to that stage). Generally, visual items that are not active in a stage are held steady while the ones that are active are animated. However, when active items of a stage do not undergo any changes, distinguishing them from the non-active items becomes difficult. A possible solution is to use highlights or opacity of the items to draw attention. While this addresses the problem of differentiating active from non-active items, active items that do change would still be more visually salient than those that do not. In addition, the technique is only applicable if these visual properties do not already encode other data attributes. Thus better ways to achieve this need to be explored. Another interesting question that the introduction of attention-focusing techniques brings to staging is that of their use for priming viewers to expect changes the items undergo during the stage. It would be interesting to explore how priming affects comprehension and memorability.

Another direction of research is how to enrich animation with advanced interpolation (as opposed to linear interpolation) to better support tracking when stages are scheduled to intersect. This exploration can be guided for instance by the trajectory bundling work by Du et al. [Du+15], or the study about temporal distortion for animated transitions by Dragicevic et al. [Dra+11].

When designing intermediate frames in an animation, whether staged or not, the view states corresponding to these frames are artificial and are not truthful to data since they are an interpolation of the true data values between the initial and final state of a transition. When adding staging to the mix, it is of interest to explore how the true and interpolated data states can be differentiated (especially in the case of n states staging).

The scenarios illustrating varied staging approaches based on the design space also evoke the need for the development of an authoring tool to experiment with different staging approaches in a seamless manner to encourage exploration of alternatives. We have implemented a prototype that supports the specification of staged animation in a custom JSON format that follows the hierarchical approach described in Section 6.6.1. While our initial implementation allows exploration of design alternatives, a graphical user interface supporting the authoring of staging remains to be explored.

In conclusion, the research road map to better understand staged animation leaves a lot of questions for our community to investigate. Our work presented in this chapter addresses a gap in the literature, and we hope that this theoretical foundation and reflections would initiate researchers and designers to further explore animation choreography in creative and expressive ways.

Chapter 7

Conclusion & perspectives

Animation has been used over many decades, but it still remains not deeply understood. Research and practice communities have not yet converged to decisive conclusions as for the merits and downsides of animation in graphical user interfaces. We believe that there are two major reasons why animation is subject to discrepancies and conflicting opinions. The first reason lies in that animation has been predominantly approached and studied as design points instantiated to achieve particular goals and finely tailored to particular contexts. There is a notable lack of literature on deep theoretical and reflexive works about animation. The majority of research works about animation are based more on reusing, inheriting and incrementing previous animation design approaches, and less focus on discussing and testing the unexplored design alternatives of animation. The second reason is the fact that animation has mainly been seen and assessed through a narrow empirical angle. Many questions about the value of animation for user interfaces have not yet been posed or investigated. Many facets of animation are still to be unveiled.

This thesis, represents a first constructive step towards filling these two major gaps and contributes a deep understanding of the roles and drawbacks of animation in user interfaces; fundamentals for animation in information visualizations and a discussion of the main challenges for designing and evaluating animation in this context; and an in-depth exploration of two main animation parameters, namely: animation trajectories and staging.

7.1 Summary

Roles and drawbacks of animation in user interfaces. We first introduced a taxonomy of animation roles that revisited the pioneer taxonomy by Beacker and Small [BS90] (Chapter 2). Our taxonomy contains 23 roles classified into 5 categories of roles: keeping in context, teaching aid, user experience, data encoding, and visual discourse. We discussed the rationales making animation support each role of our taxonomy and illustrated each role with examples from literature and practice. For each category of roles, we discussed the caveats and research opportunities.

We enriched the discussion with the findings that we extracted from an preliminary exploration that we conducted with people using animation, and a series of interviews that we conducted with practitioners to validate our taxonomy. These two studies allowed us to explore people's awareness about animation roles and their perception of (un)usefulness, as well as further understand the reasons behind (un)popularity of animation.

Building on our analysis of the drawbacks of animation perceived and experienced by participants of our preliminary exploration, we further proposed a classification of the main animation drawbacks suggested by our participants. We believe that this discussion presents a constructive step towards the elaboration of a taxonomy for animation drawbacks in the future.

Understanding dynamic changes in information visualization. We introduced theoretical basis allowing an animation designer to understand the constructing components of a dynamic visualizations. We proposed a taxonomy of dynamic changes that provides the

designer with an overview of the set of dynamic changes that she can animate within a visualization. We also suggested additional dimensions to characterize dynamic changes in visualization, including the period of change, scope of change and manner of change. We discussed later in the Chapter 6 how all these dimensions can be employed to enrich the design of staged animation.

Animation for information visualization: fundamentals and challenges. We introduced a theoretical formalization to characterize animation in the context of dynamic visualizations. We discussed the challenges that an animation designer encounters when designing and evaluating animated visualizations and we enriched this discussion with insights that we gained from our preliminary exploration described in Chapter 2.

Reflecting on the meaning of Common Fate in visualization. We conducted a perceptual study to explore the meaning of Common Fate for trajectories of position, size and luminance in dynamic visualization. This evaluation allowed us to extract a set of design implications that can guide future design of animation in dynamic visualizations.

Exploring and supporting the design of staged animation. We presented a design space that allows a holistic characterization of animation choreography—authoring staged animation across several visual states of a dynamic visualization. We surveyed and characterized literature and practice using our design space, thereby identifying the limitations of existing staging approaches and highlighting novel opportunities. To support the design of staged animation, we provided a paradigm for authoring staging alternatives based on varying the decomposition and timing parameters, and demonstrate our approach with a concrete scenario. We discussed the design considerations of an authoring tool for staged animation and described our prototype.

7.2 Future work

This thesis contributes an in-depth discussion about the rich spectrum of the design space of animation. Our work sheds light on many gaps in the literature, and therefore opens up to several research opportunities that we discuss in the following:

Study the variations and adaptation of animation roles across contexts of use. Animation is applied to various contexts of use ranging from local applications to web-based platforms, running on desktop computers, mobile devices, tabletops and wall-sized devices, for individual or collaborative use. As the context of use change, the different roles that an animation can play may not hold, and the possible drawbacks that it may induce may not be the same from one context to another. For example, one of the participants in our preliminary exploration reported that she encounters real challenges in designing scalable and responsive animation that has to be adapted for both small and large display form factors. She argued that responsive animation is “*way harder than responsive web layout*” and that “*even many responsive-minded web developers still struggle scaling animations*”. We argue that future works should explore the main characteristics of each context of use and study how each context impacts the effectiveness of animation roles. Such exploration would allow us to provide guidelines helping animation designers to adapt their design to various situations.

Elaborate a taxonomy of animation drawbacks. Although there have been multiple works that attempted to provide a structured view of the animation benefits, there have not been similar endeavors for the animation drawbacks. We argue that constructing a clear map of the downsides of animation is as important and crucial as understanding its merits. A taxonomy surveying and classifying the major drawbacks of animation in literature and practice will not only help designers better design animation, but will also help raising questions pertaining to the design and usability of animation, that are important to explore in the future.

Evaluate the impact of staging decomposition approaches on data exploration and analysis. We started to explore some evaluation strategies to study how the staging decomposition approaches can influence the understanding and interpretation of data. When surveying the literature about staged animation in dynamic visualizations, we observed that prior works used mainly low level visual tracking tasks to assess the effectiveness of staged animation. Some higher level tasks included comparing two lists of items or two versions of a diagram or detecting particular dynamic behaviors in a dynamic network. We chose to generate our future user study tasks based on the dimensions that we proposed in the Chapter 3 following a similar approach adopted by Bach et al. when evaluating GraphDiaries [BPF14]. Our principal idea is to generate a synthetic dataset where we “dissimulate” a set of insights that participants should be able to discover using a staged animation. This dataset describes the evolution of employees in a company based on a set of rules that we defined (see Appendix B). The major challenge that we have so far encountered is to find the “optimal” staging decomposition strategy for each task and objectively justify our choice. We and our collaborators on this research work attempted to brainstorm staging decomposition strategies for a set of tasks, and we found out that we were not able to converge toward the same strategy. This challenge can be solved by ranking the “estimated” (e.g. through brainstorming) or the “measured” (e.g. through an empirical test) effectiveness of staging strategies for solving a given task and picking the one with the highest score.

Support animation prototyping in dynamic visualizations through direct manipulation and sketching. One of the participants that we interviewed in our preliminary exploration said: *“For me designing animations is like sculpting: adding something here, removing something there, and so it is an organic process to finally end up with the end result”*. Animation design is an iterative process that requires a lot of tweaking and testing. More intuitive and flexible ways to prototype animation are needed. There are many relevant research questions to investigate in this perspective.

First, we need to explore and understand how animation designers tend to describe and represent a given animation. Expert animation designers generally employ different vocabularies to verbally describe animation to their customers, and use various ways to represent it visually: hand gestures, hand-drawn sketches, storyboards, quick prototyping tools to generate low fidelity and high level mockups such as CodePen [Cod], Adobe After Effect [Aft], Balsamiq [Bal] and Framer [Fra]. Exploring these different methods will certainly provide us with a set of useful description paradigms and prototyping techniques that can be implemented in a novel animation prototyping tool for dynamic visualizations. This future exploration can find inspiration in several prior works about sketch-based interfaces [WL17; DCL08].

Second, an effective animation prototyping tool that can support the authoring of animated visualization would also pave the way for the investigation of dynamic exploration of visualization mappings based on temporality. Through an accessible language, an analyst could easily experiment with different alternatives of animation within a dynamic visualization, to explore the data, enhance the visibility of the whole dynamic changes, highlight particular components

or changes in the visualization, and eventually create a story about data. Animation designers and end users do not necessarily share the same mental representation of a given animation, and they may read, describe and represent it in very different ways, which is an important aspect to explore.

In conclusion, our general understanding of animation is still rather incomplete. This thesis contributes an initial step toward mapping out the vast design space of animation, and the different purposes animation can play in user interfaces. Through our thorough analysis, we revealed several opportunities that remained to be explored, and pointed to limitations in current approaches studying animation. We hope our work will inspire future research, and serve as a solid basis for a more systematical review and exploration of the rich and expressive medium that animation is, towards more usable, engaging, and meaningful interfaces.

Appendix A

Publications

Papers in International Journals or Conferences

1. **Amira Chalbi**, Hrim Mehta, Fanny Chevalier, Christopher Collins, Nicolas Roussel, and Fanny Chevalier
Animation Choreography: A Design Space for Staged Animation
IEEE Transactions on Visualization and Computer Graphics (*under revision*)
2. Fanny Chevalier, Nathalie Henry Riche, Catherine Plaisant, **Amira Chalbi**, and Christophe Hurter
Animations 25 Years Later: New Roles and Opportunities.
In Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '16), Paolo Buono, Rosa Lanzilotti, and Maristella Matera (Eds.). ACM, New York, NY, USA, 280-287.

Posters

1. Hrim Mehta, **A. Chalbi**, Fanny Chevalier, and Christopher Collins
DataTours: A Data Narratives Framework
Proc. of IEEE Conf. on Information Visualization (InfoVis), Posters, 2017.

Doctoral Consortium

Amira Chalbi

Comprendre et Concevoir les Animations dans le Contexte des Interfaces Graphiques
IHM 16: Rencontres Doctorales-28ème Conférence Francophone sur l'Interaction
Homme-Machine, Oct 2016, Fribourg, Suisse

Outreach Activities

Dans la peau d'un robot

article in le blog binaire and in the Inria official website inria.fr

Banned From the US? There's a Robot for That

interview in WIRED

Mon Histoire avec l'Informatique

presentation at "L codent, L créent"

Appendix B

Generation rules of employees dataset

The dataset describes the evolution of employees in an SME (Small and Medium-sized Enterprise) between 1950 and 2015. The company was created in 1950. The generation rules are as follows:

Number of employees (E)

Conditions: E in [20,249]

Evolution: evolving; some evolution scenarios:

- at year_x: decrease of E
- at at year_x: important increase of E

Name (N)

Conditions: N is Unique; $N.length = 4$

Evolution: stable

Birth Date (BD)

Conditions:

- $RD(RecruitmentDate) - BD \geq 20$ (so that the employee has the minimum required age when he joins the company)
- $RD - BD \leq 50$ (the company does not recruit employees above the age of 50)

Evolution: stable

A (Age)

Conditions:

- $Min(A) = 20$ years
- $Max(A) = 65$ years

Evolution: evolving

- every year, $A = A + 1$
- If $A = 65$ years, the employee retires

G (Gender)

Conditions: Gender distribution = 60% are male and 40% are female

Evolution: stable

Recruitment Date (RD)

Conditions: RD depends on the distribution of E (number of employees) over the years.

Evolution: stable

Pre-work experience (Pre-WE)

Work experience before joining the company.

Conditions: Pre-WE in [0,10]

Evolution: stable

Post-work experience (Post-WE)

Work experience since joining the company.

Conditions:

- $\text{Min}(\text{Post-WE}) = 0$ years
- $\text{Max}(\text{Post-WE}) = 35$ years
- If $\text{Post-WE} = 35$ years, the employee retires

Evolution: evolving

- at 1950: all employees with $\text{Post-WE} = 0$
- every year, $\text{Post-WE} = \text{Post-WE} + 1$
- at year_x : $\text{Post-WE} = \text{year}_x - \text{RD}$

Qualification (Q)

Conditions: Q in { 'Bac+2', 'Bac+3', 'Bac+5', 'Bac+8', 'Other' }

Evolution: evolving.

Department (D)

Conditions: D in { 'Direction', 'Procurement', 'HR', 'Finance', 'Production', 'Marketing', 'R&D' }

Evolution: evolving; some evolution scenarios:

- at year_x : departement_x is created
- at year_x : departement_x and departement_y are merged
- at year_x : departement_x was split into departement_1 and departement_2

Department Hierarchy (D_{Hie})

Conditions: D_{Hie} in {1,2,3}.

Evolution: evolving.

Bibliography

- [3d1] *A 3-D View of a Chart That Predicts The Economic Future: The Yield Curve*. 2017. URL: <https://www.nytimes.com/interactive/2015/03/19/upshot/3d-yield-curve-economic-growth.html> (visited on 08/03/2017).
- [Aal13] Wil MP van der Aalst. “Process cubes: Slicing, dicing, rolling up and drilling down event data for process mining”. In: *Asia-Pacific Conference on Business Process Management*. Springer. 2013, pp. 1–22.
- [ABL98] David Alais, Randolph Blake, and Sang-Hun Lee. “Visual features that vary together over time group together over space”. In: *Nature Neuroscience* 1.2 (1998), pp. 160–164.
- [Abo17] Aboutkidshealth. *Exploring the Heart - The Circulatory System!* 2017. URL: <https://www.youtube.com/watch?v=-s5iCoCaofc&t=185s> (visited on 11/01/2017).
- [AF07] George A Alvarez and Steven L Franconeri. “How many objects can you track?: Evidence for a resource-limited attentive tracking mechanism”. In: *Journal of vision* 7.13 (2007), pp. 14–14.
- [Aft] *CodePen official website*. 2018. URL: <https://www.adobe.com/products/aftereffects.html> (visited on 01/01/2018).
- [Air] *PlanePlotter at AirNav System*. 2018. URL: <https://www.youtube.com/watch?v=iw0KfYRppqd4> (visited on 01/01/2018).
- [All84] James F. Allen. “Towards a general theory of action and time”. In: *Artificial Intelligence* 23.2 (1984), pp. 123–154. ISSN: 0004-3702. DOI: [http://dx.doi.org/10.1016/0004-3702\(84\)90008-0](http://dx.doi.org/10.1016/0004-3702(84)90008-0). URL: <http://www.sciencedirect.com/science/article/pii/0004370284900080>.
- [Alp+17] Basak Alper et al. “Visualization literacy at elementary school”. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM. 2017, pp. 5485–5497.
- [Ame] *How Americans Die*. 2017. URL: <https://www.bloomberg.com/graphics/dataview/how-americans-die/> (visited on 08/03/2017).
- [Ami+15] Fereshteh Amini et al. “Understanding data videos: Looking at narrative visualization through the cinematography lens”. In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM. 2015, pp. 1459–1468.
- [APA10] APA. *The Publication Manual of the American Psychological Association*. 6th. Washington, DC, 2010. URL: <http://www.apastyle.org/manual/>.
- [APP11] Daniel Archambault, Helen C Purchase, and Bruno Pinaud. “Animation, small multiples, and the effect of mental map preservation in dynamic graphs”. In: *Visualization and Computer Graphics, IEEE Transactions on* 17.4 (2011), pp. 539–552.

- [APS14] Jae-wook Ahn, Catherine Plaisant, and Ben Shneiderman. “A Task Taxonomy for Network Evolution Analysis”. In: *IEEE Trans. Vis. Comput. Graph.* 20.3 (2014), pp. 365–376. DOI: [10.1109/TVCG.2013.238](https://doi.org/10.1109/TVCG.2013.238). URL: <http://dx.doi.org/10.1109/TVCG.2013.238>.
- [AR12] Marc Azéma and Florent Rivère. “Animation in Palaeolithic art: a pre-echo of cinema”. In: *Antiquity* 86.332 (2012), pp. 316–324.
- [Arc] *Motion-Material Motion*. 2017 (accessed December 31, 2017). URL: <https://material.io/guidelines/motion/material-motion.html#material-motion-how-does-material-move>.
- [Azé11] Marc Azéma. *Préhistoire du cinéma*. Errance, 2011.
- [Bab17] Nick Babich. *Functional Animation In UX Design: What Makes a Good Transition?* 2017 (accessed December 31, 2017). URL: <https://uxplanet.org/functional-animation-in-ux-design-what-makes-a-good-transition-d6e7b4344e5e>.
- [Bac] *Why Infectious Bacteria Are Winning*. 2017. URL: <https://qz.com/576057/why-infectious-bacteria-are-winning/> (visited on 08/03/2017).
- [Bae98] Ronald M. Baecker. “Sorting Out Sorting: A Case Study of Software Visualization for Teaching Computer Science”. In: *Software Visualization*. MIT Press, 1998. Chap. 24, 369–381.
- [Bal] *Balsamiq official website*. 2018. URL: <https://balsamiq.com/> (visited on 01/01/2018).
- [Bar97a] L. Bartram. “Can motion increase user interface bandwidth in complex systems?” In: *1997 IEEE International Conference on Systems, Man, and Cybernetics. Computational Cybernetics and Simulation*. Vol. 2. 1997, 1686–1692 vol.2. DOI: [10.1109/ICSMC.1997.638254](https://doi.org/10.1109/ICSMC.1997.638254).
- [Bar97b] Lyn Bartram. “Perceptual and interpretative properties of motion for information visualization”. In: *Proceedings of the 1997 workshop on New paradigms in information visualization and manipulation*. ACM. 1997, pp. 3–7.
- [BB99] Benjamin B. Bederson and Angela Boltman. “Does Animation Help Users Build Mental Maps of Spatial Information?” In: *Proceedings of the 1999 IEEE Symposium on Information Visualization*. INFOVIS ’99. Washington, DC, USA: IEEE Computer Society, 1999, pp. 28–. ISBN: 0-7695-0431-0. URL: <http://dl.acm.org/citation.cfm?id=857189.857656>.
- [BBL12] Ilya Boyandin, Enrico Bertini, and Denis Lalanne. “A Qualitative Study on the Exploration of Temporal Changes in Flow Maps with Animation and Small-Multiples”. In: *Comp. Graph. Forum* 31.3pt2 (June 2012), pp. 1005–1014. ISSN: 0167-7055. DOI: [10.1111/j.1467-8659.2012.03093.x](https://doi.org/10.1111/j.1467-8659.2012.03093.x). URL: <http://dx.doi.org/10.1111/j.1467-8659.2012.03093.x>.
- [BCK05] T. Bladh, D.A. Carr, and M. Kljun. “The effect of animated transitions on user navigation in 3D tree-maps”. In: *Information Visualisation, 2005. Proceedings. Ninth International Conference on*. 2005, pp. 297–305. DOI: [10.1109/IV.2005.122](https://doi.org/10.1109/IV.2005.122).

- [BDB06] Anastasia Bezerianos, Pierre Dragicevic, and Ravin Balakrishnan. “Mnemonic Rendering: An Image-based Approach for Exposing Hidden Changes in Dynamic Displays”. In: *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology*. UIST '06. Montreux, Switzerland: ACM, 2006, pp. 159–168. ISBN: 1-59593-313-1. DOI: [10.1145/1166253.1166279](https://doi.org/10.1145/1166253.1166279). URL: <http://doi.acm.org/10.1145/1166253.1166279>.
- [Bed11] Benjamin B Bederson. “The promise of zoomable user interfaces”. In: *Behaviour & Information Technology* 30.6 (2011), pp. 853–866.
- [Ber83] Jacques Bertin. *Semiology of Graphics*. University of Wisconsin Press, 1983. ISBN: 0299090604.
- [Big] *Little Big Room*. 2018. URL: <http://www.little-big-room.com/> (visited on 01/01/2018).
- [BL00] Michel Beaudouin-Lafon. “Instrumental Interaction: An Interaction Model for Designing post-WIMP User Interfaces”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '00. The Hague, The Netherlands: ACM, 2000, pp. 446–453. ISBN: 1-58113-216-6. DOI: [10.1145/332040.332473](https://doi.org/10.1145/332040.332473). URL: <http://doi.acm.org/10.1145/332040.332473>.
- [BL04] Michel Beaudouin-Lafon. “Designing Interaction, Not Interfaces”. In: *Proceedings of the Working Conference on Advanced Visual Interfaces*. AVI '04. Gallipoli, Italy: ACM, 2004, pp. 15–22. ISBN: 1-58113-867-9. DOI: [10.1145/989863.989865](https://doi.org/10.1145/989863.989865). URL: <http://doi.acm.org/10.1145/989863.989865>.
- [Blo+99a] Connie Blok et al. “Visualization of relationships between spatial patterns in time by cartographic animation”. In: *Cartography and Geographic Information Science* 26.2 (1999), pp. 139–151.
- [Blo+99b] Connie Blok et al. “Visualization of relationships between spatial patterns in time by cartographic animation”. In: *Cartography and Geographic Information Science* 26.2 (1999), pp. 139–151. URL: <http://dx.doi.org/10.1559/152304099782330716>.
- [BM13] M. Brehmer and T. Munzner. “A Multi-Level Typology of Abstract Visualization Tasks”. In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2376–2385. ISSN: 1077-2626. DOI: [10.1109/TVCG.2013.124](https://doi.org/10.1109/TVCG.2013.124).
- [BM86] Ronald Baecker and Aaron Marcus. “Design principles for the enhanced presentation of computer program source text”. In: *ACM SIGCHI Bulletin*. Vol. 17. 4. ACM. 1986, pp. 51–58.
- [Bos14] Mike Bostock. *Visualizing Algorithms*. <http://bost.ocks.org/mike/algorithms/> [Accessed March 2015]. 2014.
- [Bos18] Mike Bostock. *Visualizing Algorithms*. 2018. URL: <https://bost.ocks.org/mike/algorithms/> (visited on 01/01/2018).
- [Boy+14] Jeremy Boy et al. “A principled way of assessing visualization literacy”. In: *IEEE transactions on visualization and computer graphics* 20.12 (2014), pp. 1963–1972.

- [BPF14] Benjamin Bach, Emmanuel Pietriga, and Jean-Daniel Fekete. “GraphDiaries: Animated Transitions and Temporal Navigation for Dynamic Networks”. In: *Transactions on Visualization and Computer Graphics (TVCG)* 20.5 (May 2014), pp. 740–754. DOI: [10.1109/TVCG.2013.254](https://doi.org/10.1109/TVCG.2013.254). URL: <https://hal.inria.fr/hal-00906597>.
- [Bra+09] Andrew Bragdon et al. “GestureBar: Improving the Approachability of Gesture-based Interfaces”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’09. Boston, MA, USA: ACM, 2009, pp. 2269–2278. ISBN: 978-1-60558-246-7. DOI: [10.1145/1518701.1519050](https://doi.org/10.1145/1518701.1519050). URL: <http://doi.acm.org/10.1145/1518701.1519050>.
- [Bre+17] M. Brehmer et al. “Timelines Revisited: A Design Space and Considerations for Expressive Storytelling”. In: *IEEE Transactions on Visualization and Computer Graphics* 23.9 (2017), pp. 2151–2164. ISSN: 1077-2626. DOI: [10.1109/TVCG.2016.2614803](https://doi.org/10.1109/TVCG.2016.2614803).
- [Bri07] Robert Bridson. “Fast Poisson Disk Sampling in Arbitrary Dimensions”. In: *ACM SIGGRAPH Sketches*. 22. 2007. URL: <http://doi.acm.org/10.1145/1278780.1278807>.
- [Bro14] Joseph L. Brooks. “Traditional and New Principles of Perceptual Grouping”. In: *Oxford Handbook of Perceptual Organization* (2014). URL: <http://tinyurl.com/PerGroup>.
- [BS90] R Baecker and I Small. “Animation at the Interface”. In: *The Art of Human-Computer Interface Design*. Addison-Wesley, 1990.
- [BT00] Mirielle Bétrancourt and Barbara Tversky. “Effect of computer animation on users’ performance: a review/(Effet de l’animation sur les performances des utilisateurs: une sythèse)”. In: *Le travail humain* 63.4 (2000), p. 311.
- [Bub] *Bubble to Bust to Recovery*. 2017. URL: <https://www.bloomberg.com/graphics/dataview/bubble-to-bust-to-recovery/> (visited on 08/03/2017).
- [BW02] Lyn Bartram and Colin Ware. “Filtering and Brushing with Motion”. In: *Information Visualization* 1.1 (Mar. 2002), pp. 66–79. ISSN: 1473-8716. DOI: [10.1057/palgrave/ivs/9500005](https://doi.org/10.1057/palgrave/ivs/9500005). URL: <http://dx.doi.org/10.1057/palgrave/ivs/9500005>.
- [BWC03] Lyn Bartram, Colin Ware, and Tom Calvert. “Moticons: detection, distraction and task”. In: *International Journal of Human-Computer Studies* 5.58 (2003), pp. 515–554. URL: <http://www.sciencedirect.com/science/article/pii/S1071581903000211>.
- [CA05a] Patrick Cavanagh and George A. Alvarez. “Tracking multiple targets with multifocal attention”. In: *Trends in Cognitive Sciences* 9.7 (2005), pp. 349–354. URL: <http://dx.doi.org/10.1016/j.tics.2005.05.009>.
- [CA05b] Patrick Cavanagh and George A. Alvarez. “Tracking multiple targets with multifocal attention”. In: *Trends in Cognitive Sciences* 9.7 (2005), pp. 349–354. ISSN: 1364-6613. DOI: <http://dx.doi.org/10.1016/j.tics.2005.05.009>. URL: <http://www.sciencedirect.com/science/article/pii/S136466130500149X>.
- [CA05c] Jean Clottes and Marc Azéma. *Les Félines de la grotte Chauvet*. Seuil, 2005.

- [CAAP16] Juan C Castro-Alonso, Paul Ayres, and Fred Paas. “Comparing apples and oranges? A critical look at research on learning from statics versus animations”. In: *Computers & Education* 102 (2016), pp. 234–243.
- [Cal+89] Thomas W. Calvert et al. “Composition of multiple figure sequences for dance and animation”. In: *New Advances in Computer Graphics*. Springer, 1989, pp. 245–255.
- [Cam] *Camtasia official website*. 2018. URL: <https://www.techsmith.com/video-editor.html> (visited on 01/01/2018).
- [Car] *Bloomberg Carbon Clock*. 2017. URL: <https://www.bloomberg.com/graphics/carbon-clock/> (visited on 08/03/2017).
- [Car03a] Stuart Card. “The Human-computer Interaction Handbook”. In: ed. by Julie A. Jacko and Andrew Sears. Hillsdale, NJ, USA: L. Erlbaum Associates Inc., 2003. Chap. Information Visualization, pp. 544–582. ISBN: 0-8058-3838-4. URL: <http://dl.acm.org/citation.cfm?id=772072.772109>.
- [Car03b] M. S. T. Carpendale. *Considering Visual Variables as a Basis for Information Visualisation*. Tech. rep. Calgary, AB: University of Calgary, 2003. URL: http://pharos.cpsc.ucalgary.ca/Dienst/UI/2.0/Describe/ncstrl.ucalgary_cs/2001-693-16.
- [Car03c] M. Sheelagh T. Carpendale. *Considering Visual Variables as a Basis for Information Visualisation*. Tech. rep. 2001-693-16. Calgary, AB: University of Calgary, 2003. URL: <http://tinyurl.com/Carpendale>.
- [CCD03] Kirsten Cater, Alan Chalmers, and Colin Dalton. “Varying rendering fidelity by exploiting human change blindness”. In: *Proceedings of the 1st international conference on Computer graphics and interactive techniques in Australasia and South East Asia*. ACM, 2003, pp. 39–46.
- [CDF14] Fanny Chevalier, Pierre Dragicevic, and Steven Franconeri. “The Not-so-Staggering Effect of Staggered Animations on Visual Tracking”. In: *IEEE Transactions on Visualization and Computer Graphics (TVCG/Proc. of Infovis ’14)* 20.12 (Nov. 2014), pp. 2241–2250. DOI: 10.1109/TVCG.2014.2346424. URL: <https://hal.inria.fr/hal-01054408>.
- [CDH12] Fanny Chevalier, Pierre Dragicevic, and Christophe Hurter. “Histomages: Fully Synchronized Views for Image Editing”. In: *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology*. UIST ’12. Cambridge, Massachusetts, USA: ACM, 2012, pp. 281–286. ISBN: 978-1-4503-1580-7. DOI: 10.1145/2380116.2380152. URL: <http://doi.acm.org/10.1145/2380116.2380152>.
- [CF05] Geoff Cumming and Sue Finch. “Inference by eye: confidence intervals and how to read pictures of data”. In: *American Psychologist* 60.2 (2005), p. 170. URL: <http://dx.doi.org/10.1037/0003-066X.60.2.170>.
- [Che+10] Fanny Chevalier et al. “Using Text Animated Transitions to Support Navigation in Document Histories”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’10. Atlanta, Georgia, USA: ACM, 2010, pp. 683–692. ISBN: 978-1-60558-929-9. DOI: 10.1145/1753326.1753427. URL: <http://doi.acm.org/10.1145/1753326.1753427>.

- [Che+16] Fanny Chevalier et al. “Animations 25 Years Later: New Roles and Opportunities”. In: *Proceedings of the International Working Conference on Advanced Visual Interfaces*. ACM. 2016, pp. 280–287.
- [Chi] *The World According to China*. 2017. URL: <https://www.nytimes.com/interactive/2015/07/24/business/international/the-world-according-to-china-investment-maps.html> (visited on 08/03/2017).
- [Cho] *Motion Choreography*. 2018. URL: <https://material.io/guidelines/motion/choreography.html#> (visited on 01/01/2018).
- [Cia] *CIA world browser*. 2017. URL: <http://archive.stefaner.eu/projects/relation-browser/relationViewer.swf> (visited on 09/01/2017).
- [CLW12] Joseph A. Cottam, Andrew Lumsdaine, and Chris Weaver. “Watch This: A Taxonomy for Dynamic Data Visualization”. In: *IEEE Conference on Visual Analytics Science and Technology (VAST)*. 2012.
- [CM01] M. S. T. Carpendale and Catherine Montagnese. “A Framework for Unifying Presentation Space”. In: *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology*. UIST '01. Orlando, Florida: ACM, 2001, pp. 61–70. ISBN: 1-58113-438-X. DOI: 10.1145/502348.502358. URL: <http://doi.acm.org/10.1145/502348.502358>.
- [CM85] William S. Cleveland and Robert McGill. “Graphical Perception and Graphical Methods for Analyzing Statistical Data”. In: *Science* 229 (1985), pp. 828–833. URL: <http://www.jstor.org/stable/1695272>.
- [CMS99a] Stuart K. Card, Jock D. Mackinlay, and Ben Shneiderman, eds. *Readings in Information Visualization: Using Vision to Think*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1999. ISBN: 1-55860-533-9.
- [CMS99b] Stuart K. Card, Jock D. Mackinlay, and Ben Shneiderman, eds. *Readings in Information Visualization: Using Vision to Think*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1999. ISBN: 1-55860-533-9.
- [Coc+15] Andy Cockburn et al. “Supporting novice to expert transitions in user interfaces”. In: *ACM Computing Surveys (CSUR)* 47.2 (2015), p. 31.
- [Cod] *CodePen official website*. 2018. URL: <https://codepen.io/> (visited on 01/01/2018).
- [Com] *The Dark Side of Guardian Comments*. 2017. URL: <https://www.theguardian.com/technology/2016/apr/12/the-dark-side-of-guardian-comments> (visited on 08/03/2017).
- [CPT82] Bobby J Calder, Lynn W Phillips, and Alice M Tybout. “The Concept of External Validity”. In: *Journal of Consumer Research* 9.3 (1982), pp. 240–244. URL: <http://www.jstor.org/stable/2488620>.
- [Cra] *Craft beer*. 2017. URL: <https://pudding.cool/2017/04/beer/> (visited on 09/01/2017).
- [CS27] Frederick E. Croxton and Roy E. Stryker. “Bar charts versus circle diagrams”. In: *Journal of the American Statistical Association* 22.160 (1927), pp. 473–482. URL: <http://www.jstor.org/stable/2276829>.

- [CS32] Frederick E. Croxton and Harold Stein. “Graphic Comparisons by Bars, Squares, Circles, and Cubes”. In: *Journal of the American Statistical Association* 27.177 (1932), pp. 54–60. URL: <http://www.jstor.org/stable/2277880>.
- [CSR12] Alhadi Chafi, Loris Schiaratura, and Stéphane Rusinek. “Three patterns of motion which change the perception of emotional faces”. In: *Psychology* 3.1 (2012), p. 82.
- [CU93] Bay-Wei Chang and David Ungar. “Animation: From Cartoons to the User Interface”. In: *Proceedings of the 6th Annual ACM Symposium on User Interface Software and Technology*. UIST ’93. Atlanta, Georgia, USA: ACM, 1993, pp. 45–55. ISBN: 0-89791-628-X. DOI: 10.1145/168642.168647. URL: <http://doi.acm.org/10.1145/168642.168647>.
- [Cum13] Geoff Cumming. *Understanding the New Statistics: Effect Sizes, Confidence Intervals, and Meta-Analysis*. Routledge, 2013.
- [DCH14] Pierre Dragicevic, Fanny Chevalier, and Stéphane Huot. “Running an HCI experiment in multiple parallel universes”. In: *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (alt.CHI)*. 2014, pp. 607–618. URL: <http://dx.doi.org/10.1145/2559206.2578881>.
- [DCL08] Richard C Davis, Brien Colwell, and James A Landay. “K-sketch: a ’kinetic’ sketch pad for novice animators”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM. 2008, pp. 413–422.
- [DHC11] Pierre Dragicevic, Stéphane Huot, and Fanny Chevalier. “Glimpse: Animating from markup code to rendered documents and vice versa”. In: *Proceedings of the 24th annual ACM symposium on User interface software and technology*. ACM. 2011, pp. 257–262.
- [DiB+92] David DiBiase et al. “Animation and the Role of Map Design in Scientific Visualization”. In: *Cartography and Geographic Information Systems* 19.4 (1992), pp. 201–214. DOI: 10.1559/152304092783721295. URL: <http://www.tandfonline.com/doi/abs/10.1559/152304092783721295>.
- [Dis] *Ancient kids’ toys have been hiding in the archaeological record*. 2018. URL: <https://www.sciencenews.org/article/ancient-toys-kids-archaeological-record> (visited on 01/01/2018).
- [Doo] *Doom game official website*. 2018. URL: <https://doom.com/> (visited on 01/01/2018).
- [Dra+11] Pierre Dragicevic et al. “Temporal Distortion for Animated Transitions”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’11. Vancouver, BC, Canada: ACM, 2011, pp. 2009–2018. ISBN: 978-1-4503-0228-9. DOI: 10.1145/1978942.1979233. URL: <http://doi.acm.org/10.1145/1978942.1979233>.
- [Dra15] Pierre Dragicevic. *HCI Statistics without p-values*. Research Report RR-8738. Inria, June 2015, p. 32. URL: <https://hal.inria.fr/hal-01162238>.
- [Dro] *Out of Sight, Out of Mind*. 2017. URL: <http://drones.pitchinteractive.com/> (visited on 11/01/2017).

- [Du+15] Fan Du et al. “Trajectory Bundling for Animated Transitions”. In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI ’15. Seoul, Republic of Korea: ACM, 2015, pp. 289–298. ISBN: 978-1-4503-3145-6. DOI: [10.1145/2702123.2702476](https://doi.org/10.1145/2702123.2702476). URL: <http://doi.acm.org/10.1145/2702123.2702476>.
- [Dyc+03] Jeff Dyck et al. “Learning from Games: HCI Design Innovations in Entertainment Software.” In: *Graphics interface*. Vol. 2003. 2003, pp. 237–246.
- [Ebo] *One Year of Ebola*. 2017. URL: <http://cdn.wfp.org/2016/ebola/> (visited on 09/01/2017).
- [EDF08a] Niklas Elmqvist, Pierre Dragicevic, and Jean-Daniel Fekete. “Rolling the Dice: Multidimensional Visual Exploration using Scatterplot Matrix Navigation”. In: *IEEE Transactions on Visualization and Computer Graphics* 14.6 (2008), pp. 1141–1148. URL: <http://doi.ieeecomputersociety.org/10.1109/TVCG.2008.153>.
- [EDF08b] Niklas Elmqvist, Pierre Dragicevic, and Jean-Daniel Fekete. “Rolling the Dice: Multidimensional Visual Exploration using Scatterplot Matrix Navigation”. In: *IEEE Transactions on Visualization and Computer Graphics* 14.6 (Nov. 2008). Best Paper Award, pp. 1141–1148. DOI: [10.1109/TVCG.2008.153](https://doi.org/10.1109/TVCG.2008.153). URL: <https://hal.inria.fr/hal-00699065>.
- [Eel26] Walter C. Eells. “The relative merits of circles and bars for representing component parts”. In: *Journal of the American Statistical Association* 21.154 (1926), pp. 119–132. URL: <http://www.jstor.org/stable/2277140>.
- [EF91] Max J Egenhofer and Robert D Franzosa. “Point-set topological spatial relations”. In: *International Journal of Geographical Information System* 5.2 (1991), pp. 161–174.
- [Eke17] Markus Ekert. *Curtain Effect on Meek*. 2017. URL: <https://dribbble.com/shots/1467386-Curtain-effect-on-Meek> (visited on 11/01/2017).
- [EL00] Peter Eades and Mao Lin Huang. “Navigating Clustered Graphs Using Force-Directed Methods”. In: *Journal of Graph Algorithms and Applications* 4.3 (2000), pp. 157–181. DOI: [10.7155/jgaa.00029](https://doi.org/10.7155/jgaa.00029).
- [Eng] Yuri Engelhardt. “Objects and Spaces: The Visual Language of Graphics”. English. In: *Diagrammatic Representation and Inference*. Ed. by Dave Barker-Plummer, Richard Cox, and Nik Swoboda. Vol. 4045. Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 104–108. ISBN: 978-3-540-35623-3. DOI: [10.1007/11783183_13](https://doi.org/10.1007/11783183_13). URL: http://dx.doi.org/10.1007/11783183_13.
- [Epf] *EPFL Data Monolith*. 2017. URL: <https://vimeo.com/194404753?from=outrro-embed> (visited on 09/01/2017).
- [ESG97] Charles R Ehlschlaeger, Ashton M Shortridge, and Michael F Goodchild. “Visualizing spatial data uncertainty using animation”. In: *Computers & Geosciences* 23.4 (1997), pp. 387–395.
- [Exp] *Explain Histogram*. 2017. URL: <http://tinlizzie.org/histograms/> (visited on 09/01/2017).
- [Fal] *Fallen of WWII*. 2017. URL: <http://www.fallen.io/ww2/> (visited on 09/01/2017).

- [FE00] Carsten Friedrich and Peter Eades. “The Marey graph animation tool demo”. In: *International Symposium on Graph Drawing*. Springer, 2000, pp. 396–406.
- [FE02] Carsten Friedrich and Peter Eades. “Graph Drawing in Motion”. In: *J. Graph Algorithms Appl.* 6.3 (2002), pp. 353–370. URL: <http://www.cs.brown.edu/publications/jgaa/accepted/2002/FriedrichEades2002.6.3.pdf>.
- [FEG09] Jean-Daniel Fekete, Niklas Elmqvist, and Yves Guiard. “Motion-pointing: target selection using elliptical motions”. In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 2009, pp. 289–298.
- [Fer+15] Nivan Ferreira et al. “Urbane: A 3d framework to support data driven decision making in urban development”. In: *Visual Analytics Science and Technology (VAST), 2015 IEEE Conference on*. IEEE, 2015, pp. 97–104.
- [Fer12] Cary S Feria. “The Effects of Distractors in Multiple Object Tracking are Modulated by the Similarity of Distractor and Target Features”. In: *Perception* 41.3 (2012). PMID: 22808583, pp. 287–304. DOI: 10.1068/p7053. eprint: <http://dx.doi.org/10.1068/p7053>. URL: <http://dx.doi.org/10.1068/p7053>.
- [Fer13] Cary S. Feria. “Speed has an effect on multiple-object tracking independently of the number of close encounters between targets and distractors”. In: *Attention, Perception, & Psychophysics* 75.1 (2013), pp. 53–67. ISSN: 1943-393X. DOI: 10.3758/s13414-012-0369-x. URL: <http://dx.doi.org/10.3758/s13414-012-0369-x>.
- [FFT88] Mary Anne Fisherkeller, Jerome H Friedman, and John W Tukey. “PRIM-9, an interactive multidimensional data display and analysis system”. In: *Dynamic Graphics for Statistics* (1988), pp. 91–109.
- [FH02] Carsten Friedrich and Michael E Houle. “Graph drawing in motion II”. In: *Graph Drawing*. Vol. 2265. Lecture Notes in Computer Science. Springer, 2002, pp. 220–231. URL: http://dx.doi.org/10.1007/3-540-45848-4_18.
- [FHH00] Erik Frøkjær, Morten Hertzum, and Kasper Hornbæk. “Measuring usability: are effectiveness, efficiency, and satisfaction really correlated?” In: *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM, 2000, pp. 345–352.
- [Fil] *Film Dialogue from 2,000 Screenplays, Broken*. 2017. URL: <https://pudding.cool/2017/03/film-dialogue/> (visited on 08/03/2017).
- [Fin+09] Leah Findlater et al. “Ephemeral adaptation: The use of gradual onset to improve menu selection performance”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2009, pp. 1655–1664.
- [Fir] *First time*. 2017. URL: <http://flowingdata.com/2017/02/23/the-first-time/> (visited on 09/01/2017).
- [Fis10] Danyel Fisher. “Animation for Visualization: Opportunities and Drawbacks”. In: *Beautiful Visualization*. Microsoft, 2010, pp. 329–352.
- [FJS10] SL Franconeri, SV Jonathan, and JM Scimeca. “Tracking multiple objects is limited only by object spacing, not by speed, time, or capacity”. In: *Psychological science* 21.7 (2010), pp. 920–925.

- [Fli] *Animation and Data Visualization with Vue*. 2018. URL: <https://x-team.com/blog/vue-animation-data-visualization/> (visited on 01/01/2018).
- [Foo] *Diary of a Food Tracker*. 2017. URL: <https://www.nytimes.com/interactive/2015/11/17/health/wiredwell-food-diary-super-tracker.html?mtrref=www.google.fr&gwh=32815B4C90B2CB6406F42E9F&gwt=pay> (visited on 08/03/2017).
- [FP02] Jean-Daniel Fekete and Catherine Plaisant. “Interactive Information Visualization of a Million Items”. In: *Proceedings of the IEEE Symposium on Information Visualization (InfoVis’02)*. INFOVIS ’02. Washington, DC, USA: IEEE Computer Society, 2002, pp. 117–. ISBN: 0-7695-1751-X. URL: <http://dl.acm.org/citation.cfm?id=857191.857732>.
- [Fra] *Framer official website*. 2018. URL: <https://framer.com/> (visited on 01/01/2018).
- [Fra+14] Lyndsey Franklin et al. “TreatmentExplorer: An interactive decision aid for medical risk communication and treatment exploration”. In: *Interacting with Computers* 28.3 (2014), pp. 238–252.
- [FS96] Pete Faraday and Alistair Sutcliffe. “An Empirical Study of Attending and Comprehending Multimedia Presentations”. In: *Proceedings of the Fourth ACM International Conference on Multimedia*. MULTIMEDIA ’96. Boston, Massachusetts, USA: ACM, 1996, pp. 265–275. ISBN: 0-89791-871-1. DOI: 10.1145/244130.244225. URL: <http://doi.acm.org/10.1145/244130.244225>.
- [Fw4] *FW4 Visualization*. 2017. URL: <https://uclab.fh-potsdam.de/fw4/en/vis/> (visited on 09/01/2017).
- [Gapa] *Gapminder*. 2017. URL: <http://www.gapminder.org/> (visited on 11/01/2017).
- [Gapb] *Human Development Trends*. 2005. URL: <https://www.gapminder.org/downloads/human-development-trends-2005/>.
- [Ger] *German Elections 2017*. 2017. URL: <http://www.2q17.de/#/merkel/11.08.2017> (visited on 09/01/2017).
- [Ger92] Nahum D. Gershon. “Visualization of Fuzzy Data Using Generalized Animation”. In: *Proceedings of the 3rd Conference on Visualization ’92*. VIS ’92. Boston, Massachusetts: IEEE Computer Society Press, 1992, pp. 268–273. ISBN: 0-8186-2896-0. URL: <http://dl.acm.org/citation.cfm?id=949685.949733>.
- [GEY12] S. Ghani, N. Elmqvist, and J. S. Yi. “Perception of Animated Node-Link Diagrams for Dynamic Graphs”. In: *Comput. Graph. Forum* 31.3pt3 (June 2012), pp. 1205–1214. ISSN: 0167-7055. DOI: 10.1111/j.1467-8659.2012.03113.x. URL: <http://dx.doi.org/10.1111/j.1467-8659.2012.03113.x>.
- [GF10] Tovi Grossman and George Fitzmaurice. “ToolClips: an investigation of contextual video assistance for functionality understanding”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2010, pp. 1515–1524.
- [Gfo] *Google Forms*. 2018. URL: <https://www.google.com/forms/about/> (visited on 01/01/2018).

- [Gg+13] John Alexis Guerra-gómez et al. *TreeVersity: Interactive Visualizations for Comparing Two Trees with Structure and Node Value Changes*. 2013.
- [Gib14] James J Gibson. *The ecological approach to visual perception: classic edition*. Psychology Press, 2014.
- [Gil12] Walter Gilmore. *The user-computer interface in process control: a human factors engineering handbook*. Elsevier, 2012.
- [GK97] Cleotilde González and George M. Kasper. “Animation in User Interfaces Designed for Decision Support Systems: The Effects of Image Abstraction, Transition, and Interactivity on Decision Quality”. In: *Decision Sciences* 28.4 (1997), pp. 793–823. ISSN: 1540-5915. DOI: [10.1111/j.1540-5915.1997.tb01332.x](https://doi.org/10.1111/j.1540-5915.1997.tb01332.x). URL: <http://dx.doi.org/10.1111/j.1540-5915.1997.tb01332.x>.
- [Glo] *Global Ad Spend*. 2017. URL: http://economist.columnfivemedia.com/?_ga=2.50495150.643356665.1500114279-486756203.1500114279 (visited on 09/01/2017).
- [Got] *Game of Thrones*. 2017. URL: <https://interactive.twitter.com/game-of-thrones/#?episode=1> (visited on 09/01/2017).
- [Gra37] Don Graham. *transcripts of action analysis class with Bill Tytla at the Walt Disney Studio*. 1937.
- [Gro] *How Different Groups Spend Their Day*. 2017. URL: <http://www.nytimes.com/interactive/2009/07/31/business/20080801-metrics-graphic.html?mtrref=www.google.fr> (visited on 08/03/2017).
- [GVM12] David Guilmaine, Christophe Viau, and Michael J. McGuffin. “Hierarchically Animated Transitions in Visualizations of Tree Structures”. In: *Proceedings of the International Working Conference on Advanced Visual Interfaces*. AVI ’12. Capri Island, Italy: ACM, 2012, pp. 514–521. ISBN: 978-1-4503-1287-5. DOI: [10.1145/2254556.2254653](https://doi.org/10.1145/2254556.2254653).
- [HA06] Jeffrey Heer and Maneesh Agrawala. “Software design patterns for information visualization”. In: *Visualization and Computer Graphics, IEEE Transactions on* 12.5 (2006), pp. 853–860.
- [Han+12] Paul KJ Han et al. “Representing randomness in the communication of individualized cancer risk estimates: effects on cancer risk perceptions, worry, and subjective uncertainty about risk”. In: *Patient education and counseling* 86.1 (2012), pp. 106–113.
- [Hap] *Happy Forecast*. 2017. URL: <http://www.thehappyforecast.com/SW/> (visited on 09/01/2017).
- [Har+14] Lane Harrison et al. “Ranking Visualizations of Correlation Using Weber’s Law”. In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 1943–1952. URL: <http://doi.ieeecomputersociety.org/10.1109/TVCG.2014.2346979>.
- [Har07] Mark Harrower. “The cognitive limits of animated maps”. In: *Cartographica: The International Journal for Geographic Information and Geovisualization* 42.4 (2007), pp. 349–357.
- [Hay] *HayDay*. 2017. URL: <http://haydaygame.com/> (visited on 11/01/2017).

- [HB10] Jeffrey Heer and Michael Bostock. “Crowdsourcing graphical perception: using Mechanical Turk to assess visualization design”. In: *Proceedings of the ACM Conference on Human Factors in Computing Systems*. 2010, pp. 203–212. URL: <http://dx.doi.org/10.1145/1753326.1753357>.
- [HD11] Jessica Hullman and Nick Diakopoulos. “Visualization Rhetoric: Framing Effects in Narrative Visualization”. In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (Dec. 2011), pp. 2231–2240. ISSN: 1077-2626. DOI: 10.1109/TVCG.2011.255. URL: <http://dx.doi.org/10.1109/TVCG.2011.255>.
- [HDS02] Christopher D Hundhausen, Sarah A Douglas, and John T Stasko. “A meta-study of algorithm visualization effectiveness”. In: *Journal of Visual Languages & Computing* 13.3 (2002), pp. 259–290.
- [Heaa] *Future of Heart Health in China*. 2017. URL: http://d2nd1faw866zkf.cloudfront.net/uploads/swf/GE_china_IA_R10_Trans.swf (visited on 09/01/2017).
- [Heab] *Health Law Success*. 2017. URL: <https://www.nytimes.com/interactive/2014/05/19/health/rating-a-health-laws-success.html?mcubz=3www.washingtonpost.com/graphics/national/white-death/> (visited on 09/01/2017).
- [Hea17] Val Head. *Designing Safer Web Animation For Motion Sensitivity*. 2017 (accessed December 31, 2017). URL: <https://alistapart.com/article/designing-safer-web-animation-for-motion-sensitivity>.
- [Hel] *Fewer Helmets, More Deaths*. 2017. URL: <https://www.nytimes.com/interactive/2014/03/31/science/motorcycle-helmet-laws.html> (visited on 08/03/2017).
- [HFM07] N. Henry, J. Fekete, and M.J. McGuffin. “NodeTrix: a Hybrid Visualization of Social Networks”. In: *Visualization and Computer Graphics, IEEE Transactions on* 13.6 (2007), pp. 1302–1309. ISSN: 1077-2626. DOI: 10.1109/TVCG.2007.70582.
- [His] *The History of Animation*. 2018. URL: <https://history-of-animation.webflow.io/> (visited on 01/01/2018).
- [Hit] *The long future of Hits*. 2017. URL: <http://poly-graph.co/timeless/future.html> (visited on 09/01/2017).
- [Howa] *How is Life*. 2017. URL: <http://www.oecdbetterlifeindex.org/#/14111511111> (visited on 09/01/2017).
- [Howb] *How much did everyone say*. 2017. URL: <https://pudding.cool/2017/08/the-office/> (visited on 09/01/2017).
- [HR07] Jeffrey Heer and George G. Robertson. “Animated Transitions in Statistical Data Graphics”. In: *InfoVis 2007*. InfoVis Conference, 2007. URL: <http://research.microsoft.com/apps/pubs/default.aspx?id=167169>.
- [HRC15] Lane Harrison, Katharina Reinecke, and Remco Chang. “Infographic aesthetics: Designing for the first impression”. In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM. 2015, pp. 1187–1190.

- [HS12] Jeffrey Heer and Ben Shneiderman. “Interactive Dynamics for Visual Analysis”. In: *Commun. ACM* 55.4 (Apr. 2012), pp. 45–54. ISSN: 0001-0782. DOI: [10.1145/2133806.2133821](https://doi.org/10.1145/2133806.2133821). URL: <http://doi.acm.org/10.1145/2133806.2133821>.
- [HS93] Scott E. Hudson and John T. Stasko. “Animation Support in a User Interface Toolkit: Flexible, Robust, and Reusable Abstractions”. In: *Proceedings of the 6th Annual ACM Symposium on User Interface Software and Technology*. UIST ’93. Atlanta, Georgia, USA: ACM, 1993, pp. 57–67. ISBN: 0-89791-628-X. DOI: [10.1145/168642.168648](https://doi.org/10.1145/168642.168648). URL: <http://doi.acm.org/10.1145/168642.168648>.
- [HSA08] Jan Hartmann, Alistair Sutcliffe, and Antonella De Angeli. “Towards a theory of user judgment of aesthetics and user interface quality”. In: *ACM Transactions on Computer-Human Interaction (TOCHI)* 15.4 (2008), p. 15.
- [HTC09] Christophe Hurter, Benjamin Tissoires, and Stéphane Conversy. “Fromdady: Spreading aircraft trajectories across views to support iterative queries”. In: *IEEE transactions on visualization and computer graphics* 15.6 (2009), pp. 1017–1024.
- [Hul+13] J. Hullman et al. “A Deeper Understanding of Sequence in Narrative Visualization”. In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2406–2415. ISSN: 1077-2626. DOI: [10.1109/TVCG.2013.119](https://doi.org/10.1109/TVCG.2013.119).
- [HW12] Steve Haroz and David Whitney. “How Capacity Limits of Attention Influence Information Visualization Effectiveness”. In: *IEEE Transactions on Visualization and Computer Graphics* 18.12 (2012), pp. 2402–2410. URL: <http://doi.ieeecomputersociety.org/10.1109/TVCG.2012.233>.
- [Inc] *Income Taxes You Would Pay in Each State*. 2017. URL: <http://flowingdata.com/2017/04/25/taxes-state/> (visited on 09/01/2017).
- [Iot] *The Internet of Things*. 2017. URL: <http://www.informationisbeautiful.net/visualizations/the-internet-of-things-a-primer/> (visited on 08/03/2017).
- [Jd1] *JD & Co*. 2017. URL: <https://jdand.co/> (visited on 09/01/2017).
- [Jer73] Harry J. Jerison. *Evolution of the Brain and Intelligence*. Academic Press, 1973. URL: <http://www.jstor.org/stable/2741455>.
- [Jr62] Bennet B. Murdock Jr. “The Serial Position Effect of Free Recall”. In: *Journal of Experimental Psychology* 64.5 (1962), p. 482.
- [JS00] Sara Jones and Mike Scaife. “Animated diagrams: An investigation into the cognitive effects of using animation to illustrate dynamic processes”. In: *International Conference on Theory and Application of Diagrams*. Springer, 2000, pp. 231–244.
- [JT81] Ollie Johnston and Frank Thomas. *The Illusion of Life: Disney Animation*. Disney Editions, 1981.
- [JT95] Ollie Johnston and Frank Thomas. *The Illusion of Life: Disney Animation*. Hyperion, 1995.
- [Jug] *Juggling Lab*. 2017. URL: <http://jugglinglab.sourceforge.net/> (visited on 11/01/2017).

- [KC02] Young-Min Kang and Hwan-Gue Cho. “Complex Deformable Objects in Virtual Reality”. In: *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*. VRST '02. Hong Kong, China: ACM, 2002, pp. 49–56. ISBN: 1-58113-530-0. DOI: [10.1145/585740.585749](https://doi.org/10.1145/585740.585749). URL: <http://doi.acm.org/10.1145/585740.585749>.
- [KC93] Jurek Kirakowski and Mary Corbett. “SUMI: The software usability measurement inventory”. In: *British journal of educational technology* 24.3 (1993), pp. 210–212.
- [KCS08] Aniket Kittur, Ed H. Chi, and Bongwon Suh. “Crowdsourcing User Studies with Mechanical Turk”. In: *Proceedings of the ACM Conference on Human Factors in Computing Systems*. 2008, pp. 453–456. URL: <http://dx.doi.org/10.1145/1357054.1357127>.
- [KD97] Saul M Kassin and Meghan A Dunn. “Computer-animated displays and the jury: Facilitative and prejudicial effects.” In: *Law and Human Behavior* 21.3 (1997), p. 269.
- [Key] *Animating Tables and Charts in Keynote*. 2017. URL: <https://macmost.com/animating-tables-and-charts-in-keynote.html> (visited on 09/01/2017).
- [KG13] Kris N Kirby and Daniel Gerlanc. “BootES: An R package for bootstrap confidence intervals on effect sizes”. In: *Behavior Research Methods* 45.4 (2013), pp. 905–927. URL: <http://dx.doi.org/10.3758/s13428-014-0508-5>.
- [Kha+13] Aïmen Khacharem et al. “Using segmentation to support the learning from animated soccer scenes: An effect of prior knowledge”. In: *Psychology of Sport and Exercise* 14.2 (2013), pp. 154–160. ISSN: 1469-0292. DOI: <http://dx.doi.org/10.1016/j.psychsport.2012.10.006>. URL: <http://www.sciencedirect.com/science/article/pii/S1469029212001161>.
- [KM13] Robert Kosara and Jock Mackinlay. “Storytelling: The next step for visualization”. In: *Computer* 46.5 (2013), pp. 44–50.
- [KO11] Menno-Jan Kraak and Ferjan Ormeling. *Cartography: visualization of spatial data*. Guilford Press, 2011.
- [Kof22] Kurt Koffka. “Perception: An introduction to the Gestalt-theorie”. In: *Psychological Bulletin* 19 (1922), pp. 531–585.
- [Las87] John Lasseter. “Principles of Traditional Animation Applied to 3D Computer Animation”. In: *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '87. New York, NY, USA: ACM, 1987, pp. 35–44. ISBN: 0-89791-227-6. DOI: [10.1145/37401.37407](https://doi.org/10.1145/37401.37407). URL: <http://doi.acm.org/10.1145/37401.37407>.
- [LB01] Sang-Hun Lee and Randolph Blake. “Neural synergy in visual grouping: When good continuation meets common fate”. In: *Vision Research* 41.16 (2001), pp. 2057–2064. URL: [http://dx.doi.org/10.1016/S0042-6989\(01\)00086-4](http://dx.doi.org/10.1016/S0042-6989(01)00086-4).

- [Lee+06] Bongshin Lee et al. “Task Taxonomy for Graph Visualization”. In: *Proceedings of the 2006 AVI Workshop on BEyond Time and Errors: Novel Evaluation Methods for Information Visualization*. BELIV '06. Venice, Italy: ACM, 2006, pp. 1–5. ISBN: 1-59593-562-2. DOI: [10.1145/1168149.1168168](https://doi.org/10.1145/1168149.1168168). URL: <http://doi.acm.org/10.1145/1168149.1168168>.
- [Lee+15] Bongshin Lee et al. “More than telling a story: Transforming data into visually shared stories”. In: *IEEE computer graphics and applications* 35.5 (2015), pp. 84–90.
- [LF11] Brian R. Levinthal and Steven L. Franconeri. “Common-fate grouping as feature selection”. In: *Psychological Science* 22.9 (2011), pp. 1132–1137. URL: <http://dx.doi.org/10.1177/0956797611418346>.
- [Lid16] Daniel Liddle. “Emerging Guidelines for Communicating with Animation in Mobile User Interfaces”. In: *Proceedings of the 34th ACM International Conference on the Design of Communication*. ACM. 2016, p. 16.
- [Lif] *Life on Earth*. 2017. URL: <http://www.whitevinyldesign.com/gatesnotes/howmuchlife/> (visited on 09/01/2017).
- [Lig01] Nancy J Lightner. “Model testing of users’ comprehension in graphical animation: The effect of speed and focus areas”. In: *International Journal of Human-Computer Interaction* 13.1 (2001), pp. 53–73.
- [LMW10] Jing Li, Jean-Bernard Martens, and Jarke J. van Wijk. “A Model of Symbol Size Discrimination in Scatterplots”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '10. Atlanta, Georgia, USA: ACM, 2010, pp. 2553–2562. ISBN: 978-1-60558-929-9. DOI: [10.1145/1753326.1753714](https://doi.org/10.1145/1753326.1753714). URL: <http://doi.acm.org/10.1145/1753326.1753714>.
- [Lod+96] Suresh K. Lodha et al. “UFLOW: visualizing uncertainty in fluid flow”. In: *Proceedings Visualization '96*. 1996, pp. 249–254.
- [Lon] *Quantifying the Longevity of No Diggity*. 2017. URL: <http://poly-graph.co/timeless/nodiggity.html> (visited on 09/01/2017).
- [Low03] Richard K Lowe. “Animation and learning: selective processing of information in dynamic graphics”. In: *Learning and instruction* 13.2 (2003), pp. 157–176.
- [Low04] Richard Lowe. “Interrogation of a dynamic visualization during learning”. In: *Learning and instruction* 14.3 (2004), pp. 257–274.
- [Low99] Richard K Lowe. “Extracting information from an animation during complex visual learning”. In: *European journal of psychology of education* 14.2 (1999), pp. 225–244.
- [Lyr] *Are pop lyrics getting more repetitive*. 2017. URL: <https://pudding.cool/2017/05/song-repetition/> (visited on 09/01/2017).
- [MA91] Richard E Mayer and Richard B Anderson. “Animations need narrations: An experimental test of a dual-coding hypothesis.” In: *Journal of educational psychology* 83.4 (1991), p. 484.
- [MA92] Richard E Mayer and Richard B Anderson. “The instructive animation: Helping students build connections between words and pictures in multimedia learning.” In: *Journal of educational Psychology* 84.4 (1992), p. 444.
- [Mac04] Alan M. MacEachren. *How Maps Work - Representation, Visualization, and Design*. Guilford Press, 2004, pp. I–XIII, 1–513. ISBN: 978-1-57230-040-8.

- [Mac92] Alan M MacEachren. “Visualizing uncertain information”. In: *Cartographic Perspectives* 13 (1992), pp. 10–19.
- [Mag] *Magic Lantern*. 2018. URL: https://en.wikipedia.org/wiki/Magic_lantern (visited on 01/01/2018).
- [Mara] *Étienne-Jules Marey*. 2018. URL: https://en.wikipedia.org/wiki/%C3%89tienne-Jules_Marey (visited on 01/01/2018).
- [Marb] *Étienne-Jules Marey*. 2018. URL: <http://www.ctie.monash.edu.au/hargrave/marey.html> (visited on 01/01/2018).
- [Mat] *Material motion*. 2018. URL: <https://material.io/guidelines/motion/material-motion.html> (visited on 01/01/2018).
- [May01] R.E. Mayer. *Multimedia Learning*. Cambridge University Press, 2001. ISBN: 9780521787499. URL: https://books.google.fr/books?id=yMJ9o-w_6WEC.
- [MB05] Michael J McGuffin and Ravin Balakrishnan. “Interactive visualization of genealogical graphs”. In: *Information Visualization, 2005. INFOVIS 2005. IEEE Symposium on*. IEEE. 2005, pp. 16–23.
- [MC00] Paul P Maglio and Christopher S Campbell. “Tradeoffs in displaying peripheral information”. In: *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM. 2000, pp. 241–248.
- [McC+03] D Scott McCrickard et al. “Establishing tradeoffs that leverage attention for utility: empirically evaluating information display in notification systems”. In: *International Journal of Human-Computer Studies* 58.5 (2003), pp. 547–582.
- [McC00] D Scott McCrickard. “Maintaining information awareness in a dynamic environment: Assessing animation as a communication mechanism”. PhD thesis. College of Computing, Georgia Institute of Technology, 2000.
- [MCG03] Matthias Müller, David Charypar, and Markus Gross. “Particle-based Fluid Simulation for Interactive Applications”. In: *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. SCA '03. San Diego, California: Eurographics Association, 2003, pp. 154–159. ISBN: 1-58113-659-5. URL: <http://dl.acm.org/citation.cfm?id=846276.846298>.
- [McK+17] S. McKenna et al. “Visual narrative flow: Exploring factors shaping data visualization story reading experiences”. In: *Computer Graphics Forum (EuroVis '17)* 36.3 (2017).
- [MCS01] D. McCrickard, R. Catrambone, and J. Stasko. “Evaluating animation in the periphery as a mechanism for maintaining awareness”. In: *INTERACT 2001*. 2001, pp. 148–156.
- [MDB04] Michael J McGuffin, Gord Davison, and Ravin Balakrishnan. “Expand-ahead: a space-filling strategy for browsing trees”. In: *Information Visualization, 2004. INFOVIS 2004. IEEE Symposium on*. IEEE. 2004, pp. 119–126.
- [Mea] *Measles outbreak*. 2017. URL: <https://www.theguardian.com/society/ng-interactive/2015/feb/05/-sp-watch-how-measles-outbreak-spreads-when-kids-get-vaccinated> (visited on 11/01/2017).
- [Met] *Full Metal Jousting*. 2017. URL: http://d2nd1faw866z kf.cloudfront.net/uploads/swf/HC_Jousting_960x600_4.swf (visited on 09/01/2017).

- [Mic63] A. Michotte. *The perception of causality*. Methuen's manuals of modern psychology. Basic Books, 1963. URL: <https://books.google.fr/books?id=Ae59AAAAMAAJ>.
- [Mil] *How Miles is Referenced in Wikipedia*. 2017. URL: <https://pudding.cool/2017/03/miles/> (visited on 09/01/2017).
- [Min] *At Top Colleges, an Admissions Gap for Minorities*. 2017. URL: <http://www.nytimes.com/interactive/2013/05/07/education/college-admissions-gap.html?mtrref=www.google.fr&gwh=C08D42BC1D2BC022AE27B31BD9A383C9&gwt=pay> (visited on 08/03/2017).
- [MKH12] D. Mashima, S. Kobourov, and Y. Hu. “Visualizing Dynamic Data with Maps”. In: *IEEE Transactions on Visualization and Computer Graphics* 18.9 (2012), pp. 1424–1437. ISSN: 1077-2626. DOI: 10.1109/TVCG.2011.288.
- [MI1] *A visual introduction to machine learning*. 2017. URL: <http://www.r2d3.us/visual-intro-to-machine-learning-part-1/> (visited on 08/03/2017).
- [MM17] William Holcman (Merci-Michel). *L'atelier des lettres de Noel*. 2017. URL: <http://xmasexpress.merci-michel.com/fr/> (visited on 11/01/2017).
- [Mot] *Motion Graph*. 2017. URL: <http://well-formed.eigenfactor.org/projects/motionGraph/> (visited on 09/01/2017).
- [Mov] *Map your moves*. 2017. URL: <http://moritz.stefaner.eu/projects/map%20your%20moves/> (visited on 09/01/2017).
- [MPB12] Thomas Mirlacher, Philippe Palanque, and Regina Bernhaupt. “Engineering Animations in User Interfaces”. In: *Proceedings of the 4th ACM SIGCHI Symposium on Engineering Interactive Computing Systems*. EICS '12. Copenhagen, Denmark: ACM, 2012, pp. 111–120. ISBN: 978-1-4503-1168-7. DOI: 10.1145/2305484.2305504. URL: <http://doi.acm.org/10.1145/2305484.2305504>.
- [MTFCY06] Nadia Magnenat-Thalmann, Alessandro Enrico Foni, and Nedjma Cadi-Yazli. “Real-time Animation of Ancient Roman Sites”. In: *Proceedings of the 4th International Conference on Computer Graphics and Interactive Techniques in Australasia and Southeast Asia*. GRAPHITE '06. Kuala Lumpur, Malaysia: ACM, 2006, pp. 19–30. ISBN: 1-59593-564-9. DOI: 10.1145/1174429.1174432. URL: <http://doi.acm.org/10.1145/1174429.1174432>.
- [Mun14] Tamara Munzner. *Visualization analysis and design*. CRC Press, 2014.
- [NF02] Keith V. Nesbitt and Carsten Friedrich. “Applying Gestalt principles to animated visualizations of network data”. In: *Proceedings of the International Conference on Information Visualisation*. 2002, pp. 737–743.
- [Nik] *Nike SB Dunk*. 2017. URL: https://www.nike.com/us/en_us/e/sb-dunk/ (visited on 09/01/2017).
- [Nor06] Chris North. “Information visualization”. In: *Handbook of Human Factors and Ergonomics, Third Edition* (2006), pp. 1222–1245.
- [Nor14] Gareth Norris. “The influence of computer generated animations on juror decision making”. In: *Digital Evidence & Elec. Signature L. Rev.* 11 (2014), p. 46.

- [NPD16] Till Nagel, Christopher Pietsch, and Marian Dörk. “Staged analysis: From evocative to comparative visualizations of urban mobility”. In: *Proc. IEEE Visualizaton (Arts Program)* (2016), pp. 23–30.
- [NRW11] David Novick, Joseph Rhodes, and Wervyn Wert. “The Communicative Functions of Animation in User Interfaces”. In: *Proceedings of the 29th ACM International Conference on Design of Communication*. SIGDOC ’11. Pisa, Italy: ACM, 2011, pp. 1–8. ISBN: 978-1-4503-0936-3. DOI: [10.1145/2038476.2038478](https://doi.org/10.1145/2038476.2038478). URL: <http://doi.acm.org/10.1145/2038476.2038478>.
- [Oli+04] Aude Olivia et al. “Identifying the perceptual dimensions of visual complexity of scenes”. In: *Proceedings of the Cognitive Science Society*. Vol. 26. 26. 2004.
- [One] *OneNote official website*. 2018. URL: <https://www.onenote.com/clipper> (visited on 01/01/2018).
- [Oni] *Onion Skin in Maya*. 2018. URL: <http://lesterbanks.com/2017/09/onion-skin-maya-viewport-extension/> (visited on 01/01/2018).
- [Ope] *How the U.S. and OPEC Drive Oil Prices*. 2017. URL: <https://www.nytimes.com/interactive/2015/09/30/business/how-the-us-and-opec-drive-oil-prices.html> (visited on 08/03/2017).
- [Pac] *Packed barchart*. 2017. URL: <https://www.youtube.com/watch?v=a10djve7Bt8> (visited on 09/01/2017).
- [Pal99] S.E. Palmer. *Vision Science: Photons to Phenomenology*. A Bradford book. BRADFORD BOOK, 1999. ISBN: 9780262161831. URL: <http://books.google.co.uk/books?id=IEl4QgAACAAJ>.
- [Pan+15] Anshul Vikram Pandey et al. “How deceptive are deceptive visualizations?: An empirical analysis of common distortion techniques”. In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM. 2015, pp. 1469–1478.
- [Par] *Paris Attack*. 2017. URL: <http://googletrends.github.io/parisattacks/> (visited on 09/01/2017).
- [PGB02] C. Plaisant, J. Grosjean, and B.B. Bederson. “SpaceTree: supporting exploration in large node link tree, design evolution and empirical evaluation”. In: *Information Visualization, 2002. INFOVIS 2002. IEEE Symposium on*. 2002, pp. 57–64. DOI: [10.1109/INFVIS.2002.1173148](https://doi.org/10.1109/INFVIS.2002.1173148).
- [PH92] Nancy Pennington and Reid Hastie. “Explaining the Evidence: Tests of the Story Model for Juror Decision Making”. In: 62 (Feb. 1992), pp. 189–206.
- [Pho] *Magic Mirror Photobooth - Unique, Fun, Magical*. 2017. URL: <https://www.youtube.com/watch?v=A0giDYRXEOg> (visited on 11/01/2017).
- [PL98] SE Palmer and DJ Levitin. “Synchrony: A new principle of perceptual grouping”. In: *39th Annual Meeting of the Psychonomic Society, Dallas, TX*. 1998.
- [Pla] *PlanePlotter*. 2017. URL: <http://www.coaa.co.uk/planeplotter.htm> (visited on 11/01/2017).
- [Pla+12] Catherine Plaisant et al. *Multi-Step Animation to Facilitate the Understanding of Spatial Groupings: the Case of List Comparisons*. Tech. rep. 2012.
- [Pol] *Ken Perlin*. 2017. URL: <http://mrl.nyu.edu/~perlin/experiments/polly/> (visited on 11/01/2017).

- [Ppt] *Animate individual elements of a PowerPoint chart*. 2017. URL: <http://www.techrepublic.com/blog/microsoft-office/animate-individual-elements-of-a-powerpoint-chart/> (visited on 09/01/2017).
- [Prea] *Des cinéma durant la préhistoire?* 2018. URL: <https://www.futura-sciences.com/sciences/videos/cinemas-durant-prehistoire-102/> (visited on 01/01/2018).
- [Preb] *Prezi official website*. 2018. URL: <https://prezi.com/> (visited on 01/01/2018).
- [PS54] Lewis V. Peterson and Wilbur Schramm. “How accurately are different kinds of graphs read?” In: *Educational Technology Research and Development* 2.3 (June 1954), pp. 178–189. URL: <http://dx.doi.org/10.1007/BF02713334>.
- [PT97] Dimitris Papadias and Yannis Theodoridis. “Spatial relations, minimum bounding rectangles, and spatial data structures”. In: *International Journal of Geographical Information Science* 11.2 (1997), pp. 111–138.
- [Put11] Dian Kemala Putri. “Comprehension of complex animation : cueing, segmentation and 2D / 3D presentations”. Theses. Université de Bourgogne, Nov. 2011. URL: <https://tel.archives-ouvertes.fr/tel-00795026>.
- [Py104] Zenon Pylyshyn. “Some puzzling findings in multiple object tracking: I. Tracking without keeping track of object identities”. In: *Visual cognition* 11.7 (2004), pp. 801–822.
- [Rec] *How the Recession Reshaped the Economy, in 255 Charts*. 2017. URL: <https://www.nytimes.com/interactive/2014/06/05/upshot/how-the-recession-reshaped-the-economy-in-255-charts.html> (visited on 11/01/2017).
- [Rei11] Andreas Reimer. “Squaring the Circle? Bivariate Colour Maps and Jacques Bertin’s Concept of Disassociation”. In: *Proceedings of the International Cartographic Conference* (2011), pp. 3–8. URL: <http://tinyurl.com/Reimer2011>.
- [RM13] S. Rufiange and M. J. McGuffin. “DiffAni: Visualizing Dynamic Graphs with a Hybrid of Difference Maps and Animation”. In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2556–2565. ISSN: 1077-2626. DOI: [10.1109/TVCG.2013.149](https://doi.org/10.1109/TVCG.2013.149).
- [RM15] P. Ruchikachorn and K. Mueller. “Learning Visualizations by Analogy: Promoting Visual Literacy through Visualization Morphing”. In: *IEEE Transactions on Visualization and Computer Graphics* 21.9 (2015), pp. 1028–1044. ISSN: 1077-2626. DOI: [10.1109/TVCG.2015.2413786](https://doi.org/10.1109/TVCG.2015.2413786).
- [RMC91a] George G. Robertson, Jock D. Mackinlay, and Stuart K. Card. “Cone Trees: Animated 3D Visualizations of Hierarchical Information”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’91. New Orleans, Louisiana, USA: ACM, 1991, pp. 189–194. ISBN: 0-89791-383-3. DOI: [10.1145/108844.108883](https://doi.org/10.1145/108844.108883). URL: <http://doi.acm.org/10.1145/108844.108883>.

- [RMC91b] George G Robertson, Jock D Mackinlay, and Stuart K Card. “Cone trees: animated 3D visualizations of hierarchical information”. In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM. 1991, pp. 189–194.
- [Rob] *Robin Hood*. 2017. URL: <http://povertytracker.robinhood.org/#spring2014> (visited on 09/01/2017).
- [Rob+08] George G. Robertson et al. “Effectiveness of Animation in Trend Visualization”. In: *IEEE Transactions on Visualization and Computer Graphics* 14.6 (2008), pp. 1325–1332. URL: <http://doi.ieeecomputersociety.org/10.1109/TVCG.2008.125>.
- [Rom+18] Hugo Romat et al. “Animated Edge Textures in Node-Link Diagrams: a Design Space and Initial Evaluation”. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI ’18. Montréal, Canada: ACM, 2018.
- [Ros99] Ruth Rosenholtz. “A simple saliency model predicts a number of motion popout phenomena”. In: *Vision research* 39.19 (1999), pp. 3157–3163.
- [RPD09] Florian Reitz, Mathias Pohl, and Stephan Diehl. “Focused animation of dynamic compound graphs”. In: *Information Visualisation, 2009 13th International Conference*. IEEE. 2009, pp. 679–684.
- [San] *SandDance*. 2017. URL: <https://sanddance.azurewebsites.net/BeachPartyApp/BeachPartyApp.html> (visited on 09/01/2017).
- [SB01a] Allison B Sekuler and Patrick J Bennett. “Generalized common fate: Grouping by common luminance changes”. In: *Psychological Science* 12.6 (2001), pp. 437–444.
- [SB01b] Allison B. Sekuler and Patrick J. Bennett. “Generalized common fate: grouping by common luminance changes”. In: *Psychological Science* 12.6 (Nov. 2001), pp. 437–444. URL: <http://pss.sagepub.com/content/12/6/437.abstract>.
- [SC99] Daniel J Simons and Christopher F Chabris. “Gorillas in our midst: Sustained inattentive blindness for dynamic events”. In: *Perception* 28.9 (1999), pp. 1059–1074.
- [Sch+06] Céline Schlienger et al. “Les transitions visuelles différenciées : principes et applications”. In: *IHM 2006, 18ème Conférence Francophone sur l’Interaction Homme-Machine*. Montréal, Canada, Apr. 2006, pp 59–66. DOI: [10.1145/1132736.1132744](https://doi.org/10.1145/1132736.1132744). URL: <https://hal-enac.archives-ouvertes.fr/hal-01021747>.
- [Sch01] Wolfgang Schnotz. “Sign systems, technologies, and the acquisition of knowledge”. In: *First International Seminar on Using Complex Information Systems*. Pergamon Press. 2001, pp. 9–29.
- [Set] *Setosa website*. 2017. URL: <http://setosa.io/#/> (visited on 11/01/2017).
- [SH10] E. Segel and J. Heer. “Narrative Visualization: Telling Stories with Data”. In: *IEEE Transactions on Visualization and Computer Graphics* 16.6 (2010), pp. 1139–1148. ISSN: 1077-2626. DOI: [10.1109/TVCG.2010.179](https://doi.org/10.1109/TVCG.2010.179).
- [Sha] *Shadow play*. 2018. URL: https://en.wikipedia.org/wiki/Shadow_play (visited on 01/01/2018).

- [SHK14] Jaana Simola, Jukka Hyönä, and Jarmo Kuisma. “Perception of visual advertising in different media: from attention to distraction, persuasion, preference and memory”. In: *Frontiers in psychology* 5 (2014).
- [Shn03] Ben Shneiderman. “Theories for Understanding Information Visualization”. In: *The craft of information visualization: readings and reflections* (2003), pp. 349–349.
- [Sim] *Similarities Between Languages*. 2017. URL: <http://www.beautifulinenglish.com/> (visited on 09/01/2017).
- [Sir00] Stephane Sire. “La collaboration directe. Un paradigme d’interaction pour le travail collaboratif assisté par ordinateur”. PhD thesis. 2000.
- [SK06] Lars Strother and Michael Kubovy. “On the surprising salience of curvature in grouping by proximity”. In: *Journal of Experimental Psychology: Human Perception and Performance* 32.2 (2006), pp. 226–234. URL: <http://dx.doi.org/10.1037/0096-1523.32.2.226>.
- [SK96] John T. Stasko and Colleen M. Kehoe. “Using Animations to Learn about Algorithms: An Ethnographic Case Study”. In: Technical Report GIT-GVU-96-20 (1996).
- [Ske] *SketchUp Help Center*. 2017. URL: <https://help.sketchup.com/ru/article/3000081> (visited on 11/01/2017).
- [Ski] *Avec qui allez-vous skier?* 2017. URL: <http://dataveyes.com/#!/fr/projects/intersport> (visited on 09/01/2017).
- [SKP08] Shinjiro Sueda, Andrew Kaufman, and Dinesh K. Pai. “Musculotendon Simulation for Hand Animation”. In: *ACM Trans. Graph.* 27.3 (Aug. 2008), 83:1–83:8. ISSN: 0730-0301. DOI: [10.1145/1360612.1360682](https://doi.org/10.1145/1360612.1360682). URL: <http://doi.acm.org/10.1145/1360612.1360682>.
- [SL08] Wolfgang Schnotz and Richard K Lowe. “A unified view of learning from animated and static graphics”. In: *Learning with animation: Research implications for design* (2008), pp. 304–356.
- [Sma] *Small cities, large parliaments*. 2017. URL: <http://work.interactivethings.com/nzz-swiss-maps/exekutive-und-legislative.html> (visited on 09/01/2017).
- [SMP98] John Sweller, Jeroen J. G. van Merriënboer, and Fred G. W. C. Paas. “Cognitive Architecture and Instructional Design”. In: *Educational Psychology Review* 10.3 (1998), pp. 251–296. ISSN: 1573-336X. DOI: [10.1023/A:1022193728205](https://doi.org/10.1023/A:1022193728205). URL: <http://dx.doi.org/10.1023/A:1022193728205>.
- [Sou] *SoundHound Inc.* 2017. URL: <https://soundhound.com/soundhound> (visited on 11/01/2017).
- [SP08] Peter Saffrey and Helen Purchase. “The mental map versus static aesthetic compromise in dynamic graphs: a user study”. In: *Proceedings of the ninth conference on Australasian user interface-Volume 76*. Australian Computer Society, Inc. 2008, pp. 85–93.
- [SP98] Ben Schneiderman and Catherine Plaisant. *Designing the user interface*. Don Mills, Ontario: Addison Wesley, 1998.

- [Spa+12] Ingrid A.E. Spanjers et al. “Explaining the segmentation effect in learning from animations: The role of pausing and temporal cueing”. In: *Computers & Education* 59.2 (2012), pp. 274–280. ISSN: 0360-1315. DOI: <http://dx.doi.org/10.1016/j.compedu.2011.12.024>. URL: <http://www.sciencedirect.com/science/article/pii/S0360131511003411>.
- [SS04] Frank Stürzel and Lothar Spillmann. “Perceptual limits of common fate”. In: *Vision Research* 44.13 (2004), pp. 1565–1573. URL: <http://dx.doi.org/10.1016/j.visres.2003.12.019>.
- [Sta90] John T. Stasko. “The Path-transition Paradigm: A Practical Methodology for Adding Animation to Program Interfaces”. In: *J. Vis. Lang. Comput.* 1.3 (Sept. 1990), pp. 213–236. ISSN: 1045-926X. DOI: [10.1016/S1045-926X\(05\)80007-2](https://doi.org/10.1016/S1045-926X(05)80007-2). URL: [http://dx.doi.org/10.1016/S1045-926X\(05\)80007-2](http://dx.doi.org/10.1016/S1045-926X(05)80007-2).
- [Sta93] John Stasko. “Animation in user interfaces: principles and techniques”. In: *User interface software*. JohnWiley and Sons, 1993, pp. 81–101.
- [Sto] *Stone Age art gets animated*. 2018. URL: <https://www.sciencenews.org/article/stone-age-art-gets-animated> (visited on 01/01/2018).
- [Sto+16] Charles D Stolper et al. “Emerging and recurring data-driven storytelling techniques: Analysis of a curated collection of recent stories”. In: *Microsoft Research, Washington, USA* (2016).
- [STV06] Agnès SAULNIER, Jérôme THIEVRE, and Marie-Luce VIAUD. “La perception du mouvement dans la visualisation : le cas des graphes”. In: *Revue d’Interaction Homme-Machine*. Vol 7 Num 2. 2006.
- [SW95] Nadine B Sarter and David D Woods. “How in the world did we ever get into that mode? Mode error and awareness in supervisory control”. In: *Human factors* 37.1 (1995), pp. 5–19.
- [SY06] Mutsumi Suganuma and Kazuhiko Yokosawa. “Grouping and trajectory storage in multiple object tracking: Impairments due to common item motions”. In: *Perception* 35.4 (2006), pp. 483–495.
- [Tan] *TAN Inria-Histomages*. 2018. URL: <http://tan.lille.inria.fr/histomages/?lang=fr> (visited on 01/01/2018).
- [TC01] Bruce H. Thomas and Paul Calder. “Applying cartoon animation techniques to graphical user interfaces”. In: *ACM Transactions on Computer-Human Interaction* 8.3 (Sept. 2001), pp. 198–222. URL: <http://dx.doi.org/10.1145/502907.502909>.
- [Tel] *World Wide Telescope*. 2017. URL: <http://www.worldwidetelescope.org/webclient/> (visited on 11/01/2017).
- [Ten] *The best Tennis Season Ever*. 2017. URL: <http://ig.ft.com/sites/novak-djokovic-the-best-tennis-season-ever/> (visited on 09/01/2017).
- [Thaa] *Thaumatrope*. 2018. URL: <http://dceart.weebly.com/thaumatrope.html> (visited on 01/01/2018).
- [Thab] *Thaumatrope*. 2018. URL: <https://en.wikipedia.org/wiki/Thaumatrope> (visited on 01/01/2018).

- [Tho77] Perry W. Thorndyke. “Cognitive structures in comprehension and memory of narrative discourse”. In: *Cognitive Psychology* 9.1 (1977), pp. 77–110. ISSN: 0010-0285. DOI: [https://doi.org/10.1016/0010-0285\(77\)90005-6](https://doi.org/10.1016/0010-0285(77)90005-6). URL: <http://www.sciencedirect.com/science/article/pii/0010028577900056>.
- [TMB02a] Barbara Tversky, Julie Bauer Morrison, and Mireille Bétrancourt. “Animation: Can It Facilitate?” In: *Int. J. Hum.-Comput. Stud.* 57.4 (Oct. 2002), pp. 247–262. ISSN: 1071-5819. DOI: [10.1006/ijhc.2002.1017](https://doi.org/10.1006/ijhc.2002.1017). URL: <http://dx.doi.org/10.1006/ijhc.2002.1017>.
- [TMB02b] Barbara Tversky, Julie Bauer Morrison, and Mireille Bétrancourt. “Animation: can it facilitate?” In: *International Journal of Human-Computer Studies* 57.4 (2002), pp. 247–262. URL: <http://dx.doi.org/10.1006/ijhc.2002.1017>.
- [TS85] Edward R Tufte and Glenn M Schmieg. “The visual display of quantitative information”. In: *American Journal of Physics* 53.11 (1985), pp. 1117–1118.
- [Tve+07] B. Tversky et al. “Enriching Animations”. In: *Learning with Animation: Research Implications for Design*. Ed. by R. Lowe and W. Schnotz. 2007. ISBN: 978-0521617390.
- [Tve11] Barbara Tversky. “Visualizing Thought”. In: *Topics in Cognitive Science* 3.3 (July 2011), pp. 499–535.
- [Use] *Most Common Use of Time, By Age and Sex*. 2017. URL: <https://flowingdata.com/2015/11/30/most-common-use-of-time-by-age-and-sex/> (visited on 09/01/2017).
- [Usl] *Why US lead the world*. 2017. URL: https://pudding.cool/2017/03/incarceration/index.html?utm_content=buffer3adbb&utm_medium=social&utm_source=twitter.com&utm_campaign=buffer (visited on 09/01/2017).
- [Utt+00] William R Uttal et al. “Motion and shape in common fate”. In: *Vision Research* 40.3 (2000), pp. 301–310. URL: [http://dx.doi.org/10.1016/S0042-6989\(99\)00177-7](http://dx.doi.org/10.1016/S0042-6989(99)00177-7).
- [Uwp] *Motion for UWP apps*. 2018. URL: <https://docs.microsoft.com/fr-fr/windows/uwp/design/motion/index> (visited on 01/01/2018).
- [Vau97] Leslie Carlson Vaughan. “Understanding Movement”. In: *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. CHI ’97. Atlanta, Georgia, USA: ACM, 1997, pp. 548–549. ISBN: 0-89791-802-9. DOI: [10.1145/258549.259028](https://doi.org/10.1145/258549.259028). URL: <http://doi.acm.org/10.1145/258549.259028>.
- [VB04] Daniel Vogel and Ravin Balakrishnan. “Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users”. In: *Proceedings of the 17th annual ACM symposium on User interface software and technology*. ACM. 2004, pp. 137–146.
- [VBW17] Corinna Vehlow, Fabian Beck, and Daniel Weiskopf. “Visualizing group structures in graphs: A survey”. In: *Computer Graphics Forum*. Vol. 36. 6. Wiley Online Library. 2017, pp. 201–225.
- [VEF15] Javier Villegas, Ronak Etemadpour, and Angus Graeme Forbes. “Evaluating the perception of different matching strategies for time-coherent animations.” In: *Human Vision and Electronic Imaging*. 2015, p. 939412.

- [Vis] *VisAlgo*. 2017. URL: <https://visualgo.net/en> (visited on 11/01/2017).
- [Vit] *Vitruvian Man*. 2018. URL: https://en.wikipedia.org/wiki/Vitruvian_Man (visited on 01/01/2018).
- [Vn1] *VN Star*. 2017. URL: <http://vnstar.com.ua/en/> (visited on 09/01/2017).
- [VTS96] Michalis Vazirgiannis, Yannis Theodoridis, and Timos Sellis. “Spatio-temporal composition in multimedia applications”. In: *Multimedia Software Development, 1996. Proceedings., International Workshop on*. IEEE. 1996, pp. 120–127.
- [Wag+12] Johan Wagemans et al. “A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure–ground organization”. In: *Psychological Bulletin* 138.6 (2012), pp. 1172–1217.
- [War+92] Colin Ware et al. “Moving icons as a human interrupt”. In: 4 (Oct. 1992), pp. 341–348.
- [Wat] *How a mechanical watch works*. 2017. URL: <https://www.youtube.com/watch?v=uGcoIue1Bs8> (visited on 11/01/2017).
- [Wat+10] Gareth Watson et al. “Do dynamic work instructions provide an advantage over static instructions in a small scale assembly task?” In: *Learning and Instruction* 20.1 (2010), pp. 84–93.
- [WB04a] Colin Ware and Robert Bobrow. “Motion to support rapid interactive queries on node–link diagrams”. In: *ACM Transactions on Applied Perception* 1.1 (2004), pp. 3–18. URL: <http://doi.acm.org/10.1145/1008722.1008724>.
- [WB04b] Colin Ware and Robert Bobrow. “Motion to support rapid interactive queries on node–link diagrams”. In: *ACM Transactions on Applied Perception (TAP)* 1.1 (2004), pp. 3–18.
- [WB05] Colin Ware and Robert Bobrow. “Supporting visual queries on medium-sized node–link diagrams”. In: *Information Visualization* 4.1 (2005), pp. 49–58. URL: <http://dx.doi.org/10.1057/palgrave.ivs.9500090>.
- [Wea] *Wealth Inequality in America*. 2017. URL: <https://www.youtube.com/watch?v=QPKKQnijsM> (visited on 11/01/2017).
- [Wer38] Max Wertheimer. “A Source Book of Gestalt Psychology”. In: Originally published in German in 1923, translated by W.D. Ellis. Routledge & Kegan Paul, 1938. Chap. Laws of organization in perceptual forms.
- [Wil] *Free Willy and Flipper by Numbers*. 2017. URL: <https://pudding.cool/2017/07/cetaceans/> (visited on 08/03/2017).
- [Wil17] Issara Willenskomer. *UI Animation Principles: Disney is Dead*. 2017 (accessed December 31, 2017). URL: <https://medium.com/ux-in-motion/ui-animation-principles-disney-is-dead-8bf6c66207f9>.
- [WKM02] Renée E Weiss, Dave S Knowlton, and Gary R Morrison. “Principles for using animation in computer-based instruction: theoretical heuristics for effective design”. In: *Computers in Human Behavior* 18.4 (2002), pp. 465–477. ISSN: 0747-5632. DOI: [https://doi.org/10.1016/S0747-5632\(01\)00049-8](https://doi.org/10.1016/S0747-5632(01)00049-8). URL: <http://www.sciencedirect.com/science/article/pii/S0747563201000498>.

- [WL17] Tom White and Ian Loh. “Generating Animations by Sketching in Conceptual Space”. In: *19th International Conference on Computational Creativity, ICC3, Atlanta*. 2017.
- [WL94] Colin Ware and Serge Limoges. *Perceiving data displayed through oscillatory motion*. University of New Brunswick, Faculty of Computer Science, 1994.
- [WM13] Andreas Wutz and David Melcher. “Temporal buffering and visual capacity: the time course of object formation underlies capacity limits in visual cognition”. In: *Attention, Perception, & Psychophysics* 75.5 (2013), pp. 921–933.
- [WNB99] Colin Ware, Eric Neufeld, and Lyn Bartram. “Visualizing causal relations”. In: *Proceedings of IEEE Information Visualization*. Vol. 99. 1999.
- [Woo95] David D Woods. “The alarm problem and directed attention in dynamic fault management”. In: *Ergonomics* 38.11 (1995), pp. 2371–2393.
- [Yee+01] Ka-Ping Yee et al. “Animated Exploration of Dynamic Graphs with Radial Layout”. In: *Proceedings of the IEEE Symposium on Information Visualization 2001 (INFOVIS’01)*. INFOVIS ’01. Washington, DC, USA: IEEE Computer Society, 2001, pp. 43–. ISBN: 0-7695-1342-5. URL: <http://dl.acm.org/citation.cfm?id=580582.857705>.
- [Zac+07] Jeffrey M Zacks et al. “Event perception: a mind-brain perspective.” In: *Psychological bulletin* 133.2 (2007), p. 273.
- [Zha+18] J. Zhao et al. “Supporting Handoff in Asynchronous Collaborative Sensemaking Using Knowledge-Transfer Graphs”. In: *IEEE Transactions on Visualization and Computer Graphics* 24.1 (2018), pp. 340–350. ISSN: 1077-2626. DOI: [10.1109/TVCG.2017.2745279](https://doi.org/10.1109/TVCG.2017.2745279).
- [ZKS11] Loutfouz Zaman, Ashish Kalra, and Wolfgang Stuerzlinger. “The Effect of Animation, Dual View, Difference Layers, and Relative Re-layout in Hierarchical Diagram Differencing”. In: *Proceedings of Graphics Interface 2011*. GI ’11. St. John’s, Newfoundland, Canada: Canadian Human-Computer Communications Society, 2011, pp. 183–190. ISBN: 978-1-4503-0693-5. URL: <http://dl.acm.org/citation.cfm?id=1992917.1992947>.
- [ZS03] Douglas E. Zongker and David H. Salesin. “On Creating Animated Presentations”. In: *Symposium on Computer Animation*. Ed. by D. Breen and M. Lin. The Eurographics Association, 2003. ISBN: 1-58113-659-5. DOI: [10.2312/SCA03/298-308](https://doi.org/10.2312/SCA03/298-308).