COMBINATION OF ACCRETION/DELETION OF TEXTURE AND OCCLUDING CONTOUR GEOMETRY IN DETERMINING RELATIVE DEPTH

by

ÖMER DAĞLAR TANRIKULU

A dissertation submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Psychology

Written under the direction of

Prof. Manish Singh

and approved by

New Brunswick, New Jersey

October, 2017

ABSTRACT OF THE DISSERTATION

Combination of accretion/deletion of texture and occluding contour geometry in determining relative depth

By ÖMER DAĞLAR TANRIKULU

Dissertation Director:

Prof. Manish Singh

Traditionally, accretion/deletion of texture is considered as a definite cue to ground status. However, accretion/deletion can also arise from self-occlusion due to rotation in depth. When accretion/deletion is interpreted as such, the depth-order switches and the accreting/deleting region is interpreted as being in front rather than behind the adjoining surface. This alternative interpretation of accretion/deletion has been excluded from or ignored in traditional accounts of accretion/deletion. In three studies, we investigated the factors that are crucial for the interpretation of accretion/deletion, and how this influences relative depth judgments. Recent studies (Froyen, Feldman, & Singh, 2013; Tanrikulu, Froyen, Feldman, & Singh, 2016) showed that the geometry of the border influences how accretion/deletion is interpreted. In Study 1, we systematically investigated how these two factors combine to determine relative depth by manipulating the strength of accretion/deletion and a geometric cue to figure/ground (i.e. convexity), and then combining them in various conditions. In Study 2, we investigated which stimulus factors are critical in promoting the rotating-in-front interpretation of accretion/deletion by comparing the stimuli used in recent studies versus those used in traditional studies of accretion/deletion. In the last study, we examined the influence of the speed profile of the accreting/deleting texture on the interpretation of accretion/deletion, and also examined its interaction with the shape of its border. Overall, our results indicate that accretion/deletion should not simply be considered

a "cue to ground status" because it can be interpreted as a surface either in front or behind, depending on the geometry of the occluding contour and the motion profile of the texture. Indeed, we consistently found that static contour geometry can have a greater influence on depth percepts than the motion-based cue of accretion/deletion. This calls for newer accounts to include the geometry of the borders in their models of depth from motion.

Acknowledgements

This research was supported by NSF DGE 0549115 (IGERT: Interdisciplinary Training in Perceptual Science) and by NIH EY021494 (MS, JF). I would like to express my very great appreciation to Professor Manish Singh for his great mentorship throughout the whole time I have been here. I also would like to offer my special thanks to Professor Jacob Feldman, for his many helpful comments and discussions, and to Vicky Froyen, who was the person to whom I run to get help whenever I had a problem at each stage of this project. I am also grateful to Professor Thomas Papathomas and Professor Ahmed Elgammal for his valuable comments on this dissertation. I would like to thank all the members of the Visual Cognition Lab for their helpful discussions. Lastly, I am extremely grateful to all of the faculty members of Rutgers Cognitive Science Center for creating a wonderful and an encouraging academic atmosphere that formed the basis of my motivation for this project and also for the future projects that I am going to be involved in.

Dedication

This dissertation is dedicated to my supportive parents and to my brother.

Table of Contents

List of Tables

List of Figures

1. Introduction

The study of vision might be the oldest research area among the fields of psychology. As they are also common in other areas of psychology, analogies between human brain and artificial devices (e.g. between neurons and telephone wires, between brain and computer) were also very influential in the very early days of vision research. Johannes Kepler (1571-1630) was the first person who drew a connection between the human eye and the camera obscura. His analogy was so prevalent that it influenced theories of vision by other philosophers like Rene Descartes and John Locke (See Baigrie (2002) for a review).

This eye-camera analogy also led the vision research for a long time until it was challenged by J. J. Gibson. Gibson argued that the visual system does not just record static images like a camera, but it obtains information about the spatial layout of the environment by processing the continuous transformation of the optic array resulting from the motion of the objects or of the observer (Gibson, 1966). This allowed researchers to focus more on motion cues and to discover new sources of information for perceiving spatial layout. One prominent example of this is the discovery of accretion and deletion of texture as a cue to perceive occlusion and relative depth at an edge (Gibson, 1966; Gibson et al., 1969). When a translating texture deletes or accretes from a boundary, it is perceived as if it is disappearing or appearing from behind an occluding surface on the other side of the boundary. This in turn generates a percept where the accreting/deleting surface is seen as the background (i.e. behind) and the adjoining surface as figural (i.e. in front).

The first study on the idea of accretion/deletion was done by Michotte, Trines, and Crabbe (1964). In this study, subjects were shown a projected image of a white circular disk on a black background where, starting from one side of it, the disk gradually blackened out. Throughout the display, subjects reported to perceive an unchanging circular disk, which is being covered by another surface. Michotte et al. (1964) named this phenomenon as "screening effect", where the circular shape is still (amodally) perceived, despite the transformation of the shape of the object. Two years later, Gibson (1966) described the display used by Michotte et al. (1964) as “wiping out” or “unwiping” of optical texture. These observations and descriptions were precursors to “the hypothesis of deletion/accretion for edge perception”, which was title of the article by Gibson et al. (1969), in which they described in detail the transformation of the optic array when a surface occludes or dis-occludes another.

According to Gibson et al. (1969), when a part of an object goes out of sight, there are two possibilities that might cause this. One possibility is that that part of the object might go out of existence, (by evaporating or being eaten, etc.). The second possibility is that the object might be occluded by another object. In the latter case, the shape of the object that is being occluded is amodally completed, but not in the former case. Gibson et al. (1969) argued that these two possibilities could be distinguished from each other with the help of optical information. They proposed their hypothesis of accretion/deletion, which describes the optical transformation that distinguishes these two cases (Figure 1.1)

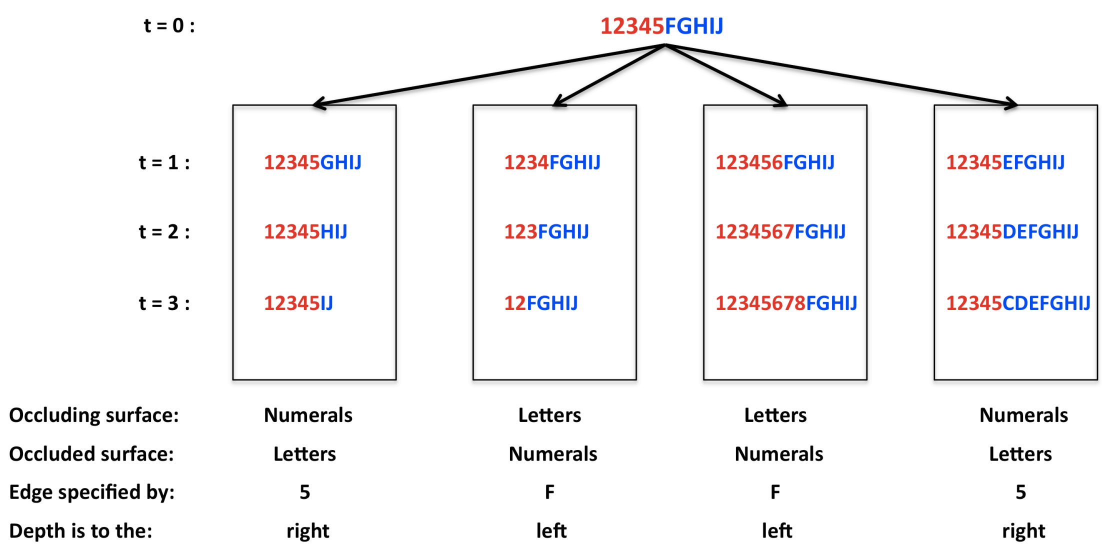

Figure 1.1: A schematic diagram of the accretion/deletion hypothesis proposed by Gibson et al. (1969). The string of symbols corresponds to texture elements on a surface. The letters and numerals do not represent different kinds of elements; they just imply that the image can be divided into two regions. Each box shows a different kind of optical transformation and at the bottom of each box the resulting depth interpretation from that corresponding transformation is shown.

In summary, the hypothesis states that the part of the optic array that is subjected to accretion or deletion corresponds to the surface that is behind because it is being occluded or dis-occluded by the adjacent surface. The part of the optic array that is preserved corresponds to the surface that is occluding the adjacent surface and perceived in front. The location where the dots/elements disappear indicates the edge of the occluding surface.

Around the same year, Kaplan (1969), who was a student of Gibson at the time, experimentally tested the accretion/deletion hypothesis depicted in Figure 1.1. In addition to this hypothesis, Kaplan (1969) also had a second hypothesis that he tested in his study. He predicted that whenever there is accretion/deletion of texture elements on both sides of a border, the region that has a higher rate of accretion/deletion per interval of time would be perceived as behind the region that has a lower rate of accretion/deletion.

In order to test his two hypotheses, Kaplan used a stimulus where a random texture surface is divided into two regions by a vertical motion-defined contour. The texture in the two different regions can move horizontally in opposite direction. When the stimulus is static, it is perceived as a single continuous random textured surface. The contour separating the two regions is visible when there is motion contrast between the two regions. In his first experiment, either one side of the contour had accreting/deleting texture (whereas the other side was static) or both sides of the contour had accreting/deleting texture. In this experiment, he kept the motion-defined contour static. In order to test his second hypothesis, he also manipulated the speed of the accreting/deleting texture. In the second experiment, the contour that separates the two regions also moved either towards left or right. This second experiment allowed him to test his second hypothesis in a little more systematic way because by allowing the contour to move he was able to obtain different levels of relative accretion/deletion rate between the two regions without introducing more speed levels to the experiment.

Based on the subjects' relative depth judgments on the set of stimuli described above, Kaplan (1969) claimed that both of his hypotheses were supported by the data. 98% of the relative depth judgments were correctly predicted by the first hypothesis, whereas 87% of them were correctly predicted by the second hypothesis. He also found a correlation between the relative accretion/deletion rate between the two regions and the success of the predictions of his hypotheses. As the accretion/deletion disparity across the contour increases, the success of the predictions becomes greater. Based on these results he concluded that the rate of accretion/deletion is the crucial factor that determines the relative depth judgments of the observers in such displays.

Since this influential study by Kaplan (1969), accretion/deletion has always been considered as a decisive cue that unambiguously assigns ground status to an image regions (Thompson, Mutch, & Berzins, 1985; Mutch & Thompson, 1985; Niyogi, 1995; Howard & Rogers, 2002; Hegdé, Albright, & Stoner, 2004), and it has been used as such in computational models of depth from motion (Thompson et al., 1985; Mutch & Thompson, 1985; Berzhanskaya, Grossberg, & Mingolla, 2007; Beck, Ognibeni, & Neumann, 2008; Raudies & Neumann, 2010; Barnes & Mingolla, 2013; Layton & Yazdanbakhsh, 2015; Ruda, Livitz, Riesen, & Mingolla, 2015). However, this traditional view is challenged recently by a series of experiments conducted by Froyen et al. (2013) and Tanrikulu et al. (2016).

Froyen et al. (2013) presented a new phenomenon that emerges from the interaction of accretion/deletion cue and geometric cues to figure/ground, which tends to promote figural status, such as convexity, symmetry and parallelism (for a review see Wagemans et al. (2012)). Figure 1.2 shows the display setup used by Froyen et al. (2013) and the corresponding phenomenology. The displays contained alternating light and dark regions with random-dot texture moving horizontally at constant speed, but in opposite directions in alternating regions. Such displays are ambiguous in terms of depth from accretion/deletion since both sides of each border have equal rate of accreting/deleting texture. When Froyen et al. (2013) asked subjects to freely describe what they see when they were shown these multiple-region figure/ground displays, almost all of them reported that they were seeing rotating columns in front of a translating flat sheet in the back. Then the subjects were asked to make relative depth judgments and it was found that the regions that were perceived in front were also perceived as 3D volumes that are rotating in depth in front of a flat surface translating in the opposite direction. There are two main reasons why such an interpretation of these displays is novel and surprising. First of all, the perception of 3D columns rotating in depth was observed even though the speed profile of the dots on the image was constant (instead of cosine), which is technically inconsistent with 3D rotation of rigid objects. Second, rotation in depth was perceived even when asymmetric regions were used. If such asymmetric regions in the stimulus were really 2D projections of 3D rotating columns in depth, then their occluding contours should change their shape continuously. However, the contours were all kept static in the experiments. On the other hand, in spite of all these inconsistencies, the 3D rotating column interpretation of such displays actually allows the visual system to explain the accreting/deleting texture that is seen in front. According to the traditional view of accretion/deletion, the texture appears or disappears due to occlusion by another surface. When one set of regions (dark or light) is seen in front, the appearing/disappearing of texture on these surfaces cannot be explained due to being occluded by another surface in front. Therefore, self-occlusion due to rotation of a 3D volume becomes the only candidate explanation of the accreting/deleting texture. Froyen et al. (2013) also showed that geometric figure/ground cues can resolve this bi-stability depicted in Figure 1.2, by biasing a certain set of regions to be perceived in front. The geometric figure/ground cues used in that experiment were convexity, parallelism and symmetry.

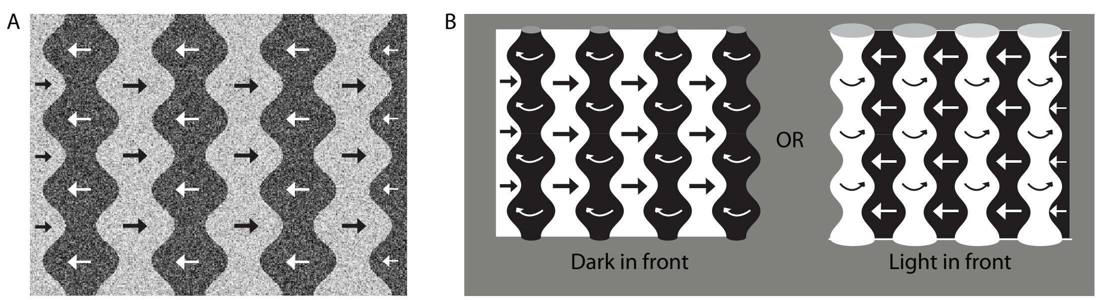

Figure 1.2: (A) The setup of the display used by Froyen et al. (2013) and Tanrikulu et al. (2016): The classical multi-region figure-ground displays had motion in one direction in odd regions and in the other direction in even regions. (B) The corresponding phenomenology: The display could yield one of the two percepts depending on which set of regions was perceived as figural. Either the dark regions were perceived as rotating in front of a light background, which was seen as sliding behind the rotating columns, or vice versa.

However, in the displays used by Froyen et al. (2013), an ambiguity in relative depth is created by introducing accreting/deleting texture on both sides of each border. When this depth ambiguity is resolved by a geometric cue, one set of accreting/deleting regions are perceived to be in front. The visual system *explains* the appearing/disappearing of the dots in this figural region by self-occlusion due to 3D rotation in depth. In order to test whether this interpretation extends to the case where only one side of each border has accreting/deleting texture, Tanrikulu et al. (2016) conducted a similar experiment, including a condition where either the odd or the even regions had static texture. What is crucial here is that according to the conventional definition of accretion/deletion, this should lead to an unambiguous depth order assignment. Tanrikulu et al. (2016) also introduced symmetry and convexity to one set of regions in order to examine the interaction between these geometric cues and accretion/deletion. They included both a "cue-competition" condition, which is shown in the first column in Figure 1.3, in which accretion/deletion (in its traditional sense) and the geometric cues compete against each other. They also included a "cue-cooperation" condition that is shown in the second column in Figure 1.3, in which the two cues were consistent with each other. There were two tasks in the experiment. The first task was to indicate whether the subjects see the indicated region (marked by arrows) in front of its adjacent region. The second task was to indicate whether the subjects see a rotational or a translational motion in the indicated region. In this way, Tanrikulu et al. (2016) would be able to see whether there is a correlation between perceived relative depth and perceived 3D shape of a region. They found that when accretion/deletion was present only in the "convex" regions, these regions were nevertheless perceived as rotating in front on roughly half of the trials (despite the fact such displays contain no ambiguity as far accretion/deletion is concerned). In their second experiment, the motion in the two sets of regions (light or dark) was made incoherent by alternating the direction of motion in convex/symmetric regions such that there would be no motion-based grouping (Figure 1.4). When this grouping effect was eliminated, the proportion of the times the convex regions with accreting/deleting texture were seen in front increased to around 60%. With these results, Tanrikulu et al. (2016) were able to demonstrate that geometric cues (convexity and symmetry) strongly modulates how accretion/deletion is interpreted: Convexity cancels out the effect of accretion/deletion, and symmetry was able to (even though it was lower than convexity) make the region that has accretion/deletion to be perceived as figural on a certain proportion of the trials. The response patterns obtained in the two different tasks were fairly similar, suggesting that the perception of rotation in depth is causally connected to figure-ground interpretation (e.g. if the moving side is interpreted as figural, it is also seen as a 3D rotating column in depth).

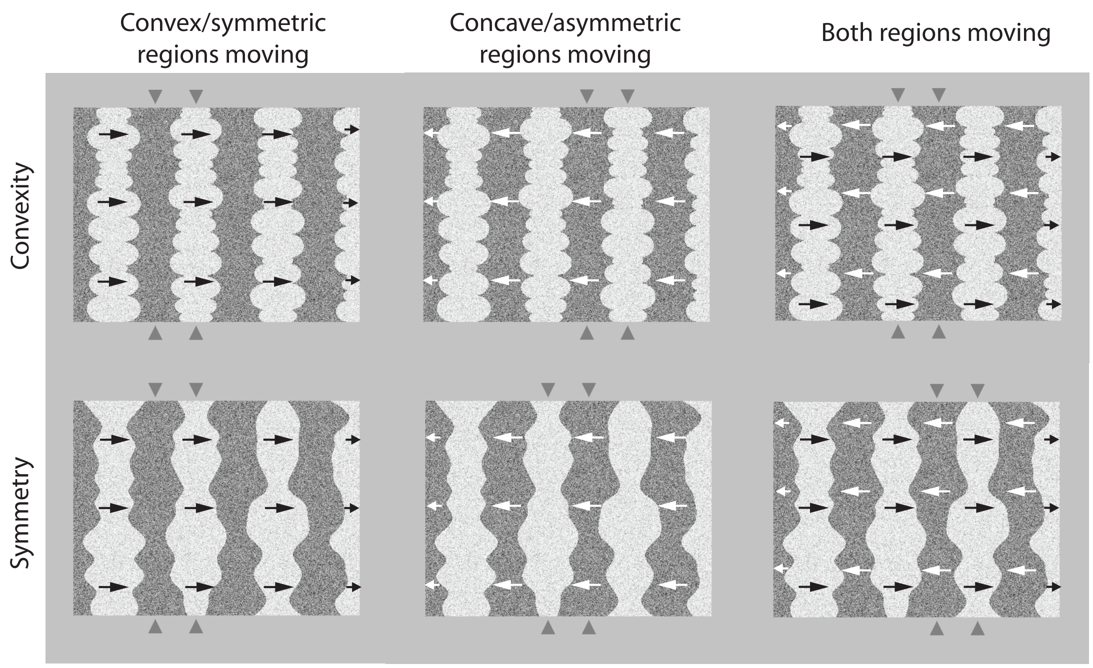

Figure 1.3: The experimental conditions used in Tanrikulu et al. (2016). The horizontal arrows indicate the region that has accreting/deleting texture. The regions either had convex and symmetric boundaries (top row) or just symmetric boundaries (bottom row). Either both regions on the two sides of each boundary had accreting/deleting texture (the third column), or just one side of each border had accreting/deleting texture (the first and the second columns). The first column corresponds to the cue competition condition where the depth indicated by the geometric cues were inconsistent with the depth indicated by accretion/deletion. The second column corresponds to the cue-cooperation condition, where the two cues were consistent with each other.

The results obtained by Tanrikulu et al. (2016), which showed that in certain contexts accreting/deleting regions could be perceived as figural, extends the findings by Froyen et al. (2013) to a case where accretion/deletion should unambiguously assign

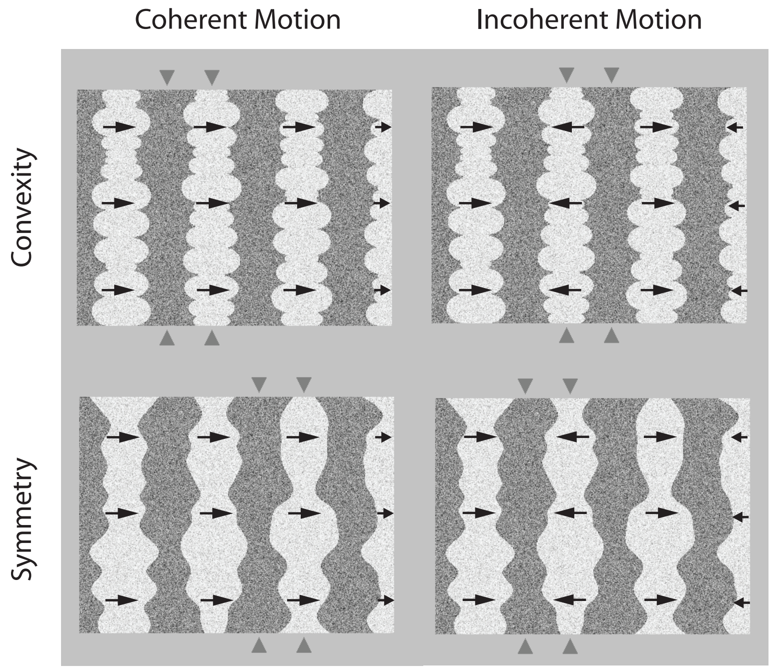

Figure 1.4: The experimental conditions used in the second experiment of Tanrikulu et al. (2016), in order to examine the effect of motion coherence among the accreting/deleting regions.

depth order to surfaces. According to these results, a reconsideration of the traditional account of accretion/deletion as a cue to ground status is required. Exactly what information accretion/deletion conveys depends, rather, on the geometry of the boundary, along with other global factors such as the coherence of the background motion. For example, in the stimuli used by Tanrikulu et al. (2016), it seems that once accretion/deletion of the texture is explained by the visual system as self-occlusion due to rotation, accretion/deletion no longer functions as a cue that indicates the occluded surface. These results calls for a new understanding of accretion/deletion as a cue to relative depth that could account for the results of these recent studies. In order to understand what depth information accretion/deletion conveys, one need to understand the principles and factors that determines how accretion/deletion is interpreted. In the following chapters, three different studies are described, which attempt to achieve this goal by examining this rich interaction between accretion/deletion and the geometry of the occluding contour.

In the first study (Chapter 2), we investigated how a geometric figure/ground cue (i.e. convexity), combines with accretion-deletion by manipulating the strength of each cue and combining them in various conditions. The displays contained alternating light and dark regions with texture (i.e. dots) moving horizontally at constant speed, but in opposite directions in alternating regions. The strength of convexity cue is manipulated by making the negative minima of the contour sharper. Strength of accretion-deletion is manipulated by varying the rate of accretion/deletion in two different ways, either by varying the relative texture density or by varying the overall texture speed in a set of regions.

The second study (Chapter 3) sought to understand the reason for the discrepancy in percepts between the traditional accretion/deletion displays and the rotating-columns displays used by these recent studies (Froyen et al., 2013; Tanrikulu et al., 2016). The rotating-columns displays differ from the traditional accretion/deletion displays in a number of factors, including the presence of figure/ground cues, accretion/deletion on both sides of boundaries, and in the number of distinct motion patches. In the second study, we conducted a series of experiments, in which we systematically manipulated each of these factors in order to determine what factors are actually instrumental in creating the rotating column (i.e. accretion/deletion in front) interpretation.

In all of the studies described above in which rotating-in-front interpretation of accretion/deletion was observed, the accreting/deleting texture had always constant speed in the image. In the third and final study (Chapter 4) we manipulate the image speed profile of the accreting/deleting texture, and examine its interaction with the shape of its bounding contours. We have argued that the way in which accretion/deletion is interpreted determines which depth order accretion/deletion would indicate. In principle, the 2D speed profile of the accreting/deleting texture in the image would be different depending on whether accretion/deletion is generated by a translating surface in behind or by a rotating column in front. This would imply that the speed profile of the accreting/deleting texture should significantly influence the relative depth judgments. In our last study, we test this prediction. We also present a probabilistic model, combining speed profile and contour geometry, that accounts for these findings without explicitly treating accretion/deletion as a cue to ground.

This dissertation is written in a “modular” fashion, in which each chapter is self-contained and written separately of the other chapters.

2. A cue interaction study of depth ordering from accretion-deletion and contour convexity

2.1 Introduction

Our visual system has the capacity to build rich 3-D representations of our environment from complex 2-D retinal images. Even after contours and regions on an image are computed, the visual system has to organize these seemingly unstructured contours and regions. Therefore, one of the crucial steps of visual processing is to determine what regions and contours in the image belongs together as distinct surfaces, objects and other useful perceptual units. An important part of this perceptual grouping involves segmenting images into figural and ground regions. The figural regions own the adjoining contours and they appear to be nearer to the observer than the ground regions, whereas the ground regions are amodally extended behind the figural regions. There are many configural cues which involve the geometry of the contours/regions and tend to promote figural status, such as such as symmetry (Kanizsa & Gerbino, 1976), convexity (Metzger, 1936/2006; Kanizsa & Gerbino, 1976), parallelism (Morinaga, 1941; Metzger, 1936/2006), axiality and part salience (Hoffman & Singh, 1997; Froyen, Feldman, & Singh, 2010) (see Wagemans et al. (2012) for a review).

Besides geometric cues, there are also dynamic cues to figure/ground organization. One important type of dynamic cue is called accretion/deletion of texture. When a texture is accreted or deleted at a boundary, it is perceived as if the accreting/deleting texture gradually appearing or disappearing because it is being occluded by the adjacent surface at the other side of that boundary. Such an interpretation of the texture being accreted/deleted creates a sense of relative depth where the accreting/deleting texture is perceived as being behind the adjacent occluding surface. Accretion/deletion is traditionally considered an unambiguous and decisive cue to ground status (Kaplan, 1969; Gibson et al., 1969; Thompson et al., 1985; Mutch & Thompson, 1985; Niyogi, 1995; Howard & Rogers, 2002; Hegdé et al., 2004) and it has been incorporated as such into computational models of depth from motion (Yonas, Craton, & Thompson, 1987; Berzhanskaya et al., 2007; Beck et al., 2008; Raudies & Neumann, 2010; Barnes

& Mingolla, 2013; Layton & Yazdanbakhsh, 2015; Ruda et al., 2015).

Traditional accounts of accretion/deletion assumed that such an account is the only relative-depth interpretation of accreting/deleting textures. However, there is another possible interpretation of accretion/deletion where the region containing accreting/deleting texture is perceived as figural. Accreting/deleting of texture refers to an event of a moving texture appearing or disappearing at a boundary. The way our visual system interprets this event determines the resulting relative-depth judgment. If the appearing and disappearing of the texture is attributed to being occluded by the adjacent region, then the accreting/deleting region is perceived to be in behind. However, accretion/deletion of texture can also arise from self-occlusion due to rotation of a volumetric object in depth (Figure 2.1). When the accretion/deletion of the texture at a boundary is attributed to self-occlusion, then the accreting/deleting region is perceived to be the figural region that owns that boundary. When this happens, the accreting/deleting regions are also perceived to be 3-D volumes rotating in front of the adjacent regions that are now perceived to be the ground. This rotating-in-front interpretation of accreting/deleting regions have been noted in passing by previous researchers (Kaplan, 1969; Yonas et al., 1987; Royden, Baker, & Allman, 1988) but it has not been incorporated into standard accounts of accretion/deletion as a figure/ground cue.

In these previous studies (Kaplan, 1969; Yonas et al., 1987; Royden et al., 1988), the perception 3-D volumes rotating in depth was observed in spite of the constant speed of the texture, which is inconsistent with 3D rotation of rigid objects (e.g. Ullman, 1979). Motivated by this observation, Thompson and his colleagues (Thompson, Kersten, & Knecht, 1992; Thompson & Painter, 1992) investigated in what ways textural motion around surface boundaries influence structure-from-motion such that translating texture that accretes/deletes at a boundary give rise to the perception of 3-D volumes rotating in depth. They found out that increasing the width of a region that contains accreting/deleting texture (e.g. increasing the width of the red region in Figure **??**) reduces the probability of seeing that region as a 3-D rotating column. To account for their observations, they also proposed a computational model that is based

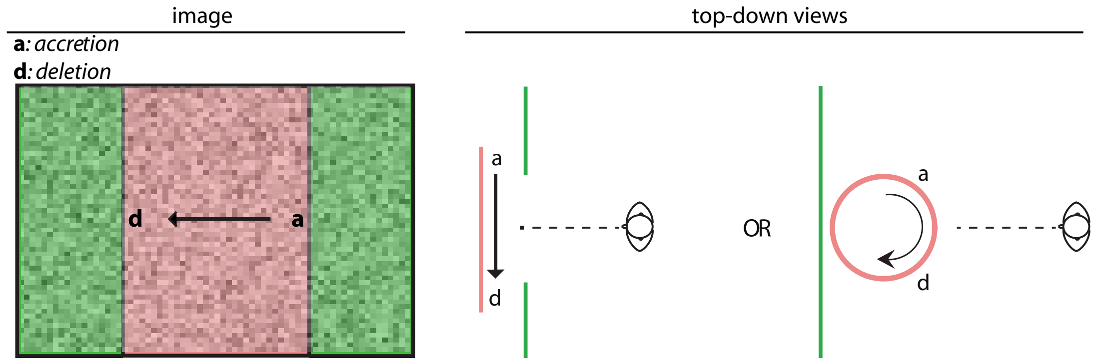

Figure 2.1: The frontal projection of an accreting surface is shown on the left. “a” is the location of texture accretion and “d” is the location of texture deletion. The static surface is depicted in green while the moving surface is depicted in red. On the right, overhead views of the two possible 3D arrangements with different depth-order assignments that are both consistent with the frontal view of the accreting and deleting surface are shown.

on qualitative constraints relating surface shape and depth order to patterns of textural motion around boundaries. However, even though the work done by Thompson et al. (1992) and by Thompson and Painter (1992) were the first studies concerning accreting/deleting textures giving rise to 3D volumes rotating in front, they were mostly focused on the effect of speed profile of textural motion near surface boundaries on structure from motion, not on accretion/deletion and relative-depth per se.

Only recently, there have been a few studies that directly focus on the depth-order ambiguity inherent in accretion/deletion. Kromrey, Bart, and Hegdé (2011) demonstrated that accretion/deletion alone is not sufficient to unambiguously determine relative depth, and information about the occluding border influences the perception of relative depth. When an enclosed region containing accreting/deleting random dot texture was surrounded by a flickering random dot texture, the central region was seen in front even though only the texture in the central region was accreting/deleting. The traditional interpretation of accretion/deletion was favored only when the delineation of the border between the center and the surround regions was made easier by segmentation cues, such as increasing the luminance contrast between the two regions, or making the surrounding region static. They observed that accretion/deleting texture is perceived in front only when the surrounding texture was flickering.

Froyen et al. (2013) demonstrated a striking example of the two interpretations of accretion/deletion illustrated in Figure 2.1. They introduced accretion/deletion on both sides of each boundary in a multi-region figure/ground display (similar to the one shown in Figure 2.2A), and this created a bistable stimulus where either the accreting/deleting texture in the dark or in the light regions were interpreted as in front and rotating in depth (Figure 2.2B). When a geometric cue to figure/ground (e.g. convexity) was introduced to one set of regions, then this ambiguity was resolved and that set of regions were more likely to be perceived as 3-D columns rotating in front, in spite of the constant speed motion of the moving dot texture (which is inconsistent with 3D rotation in depth). In Froyen et al. (2013), there was accretion/deletion on both sides of a boundary, therefore accretion/deletion (in the traditional sense) did not favor any depth order. Tanrikulu et al. (2016) showed that the perception of 3D columns rotating in front also occurs even in displays where only one side of each border contains accreting/deleting texture (i.e. the texture within one set of regions was static). Such a display would be considered an unambiguous stimulus (in terms of relative depth) according to traditional accounts of accretion/deletion. They observed that when the geometric cues of convexity and symmetry were introduced to one set of regions while accretion/deletion was introduced within the other set of regions, the accreting/deleting regions were perceived as 3-D columns rotating in front in roughly half of the trials. They also found that preventing the motion grouping of the moving texture among the accreting/deleting regions by changing the direction of motion in one set of regions further increased the proportion of trials in which the accreting/deleting regions were perceived as rotating in front.

These recent studies challenge the traditional view that accretion/deletion is a decisive cue that unambiguously determines ground status. Partly because of this, traditional accounts of accretion/deletion do not consider the geometry of the border as a relevant factor in determining relative depth. However, as Froyen et al. (2013) and Tanrikulu et al. (2016) demonstrated, the geometry of the border has a strong influence on whether accreting/deleting region is interpreted as being in front or in behind. This influence cannot simple be viewed as competition between geometric cues to

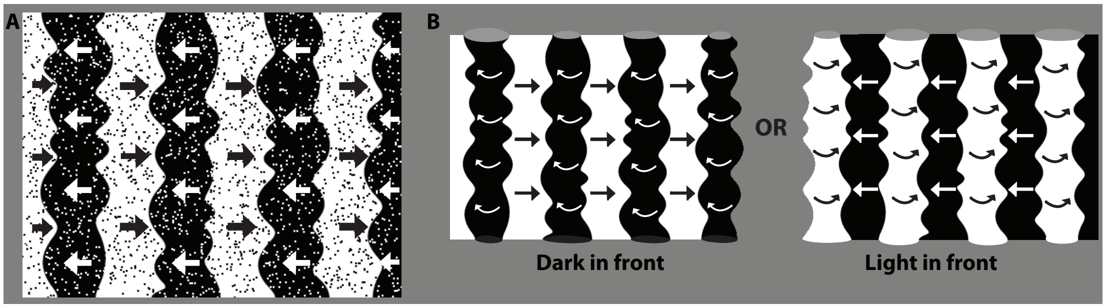

Figure 2.2: Display setup and phenomenology: A. The displays were created by adding motion in one direction to odd regions and in the other direction to even regions in classical figure-ground displays. B. This could yield one of two percepts depending on which one was perceived as figural. The black ones were perceived as rotating in front of a white background which was seen as sliding behind them, or vice versa.

figure/ground and accretion/deletion (in its traditional sense), since accretion/deletion is in fact consistent with both in-front and in-behind interpretations. Once the accretion/deletion of texture is interpreted as self-occlusion due to rotation in depth, accretion/deletion no longer functions as a cue that indicates the occluded surface. Therefore accretion/deletion cannot be said to compete against or cooperate with the geometric cues to figure/ground, but it rather combines with the geometry of the border to determine relative depth. In other words, the ordinal depth information conveyed by accretion/deletion is heavily dependent upon how the accretion/deletion of the texture is accounted for, which is, in turn, influenced by the geometry of the border. Interestingly, this alternative (rotating-in-front) interpretation of accretion/deletion is even strong enough to override the information coming from the constant speed profile of the moving texture, which is, in principle, fully consistent with a translating flat surface, but clearly inconsistent with a 3D rigid surface rotating in depth.

These recent studies by Froyen et al. (2013) and Tanrikulu et al. (2016) showed that the interpretation of accretion/deletion depends on the geometry of border. In this current study, we systematically investigate how these two factors combine to determine relative depth. We manipulated the strength of accretion/deletion and a geometric cue to figure/ground (i.e. convexity), and then combined them in various conditions to examine how they together determine relative depth judgments. This will allow us the examine the contribution of each factor individually and the interaction between them in determining relative depth judgments. More importantly, we will analyze this interaction between the two factors from the perspectives of both the traditional account and our way of understanding accretion/deletion.

According to the traditional accounts of accretion/deletion, what determines the strength of accretion/deletion is the rate of accretion/deletion of the texture elements: As the rate of accretion/deletion increases in a region, the strength of accretion/deletion (in its traditional sense) increases, hence the probability of seeing that region in behind also increases (Kaplan, 1969). We are going to manipulate the strength of accretion/deletion by manipulating the relative rate of accretion/deletion between the two sides of each border. Even though this involves an initial concession to the traditional view of accretion/deletion, by doing this we will be also examining whether this traditional understanding of accretion/deletion is the correct way to understand how accretion/deletion contributes to relative depth interpretation.

Our experiments manipulated the rate of accreting/deleting texture in two different ways. In Experiment 1, we manipulated the relative density of the texture being accreted/deleted on the two sides of a border. In Experiment 2, we manipulated the relative speed of the accreting/deleting texture. In both experiments, we manipulated the degree of convexity introduced to the one side of the border. The borders were made piecewise convex on one side and we manipulated the degree of convexity by varying the part salience of the convex parts of the border (Hoffman & Singh, 1997).

2.2 Experiment 1

In Experiment 1, we looked at how the geometric figure/ground cue of convexity combines with relative density of the accreting-deleting texture to determine relative depth. Both factors were manipulated and combined in various conditions in a figure/ground task. The strength of convexity was manipulated by gradually making the negative minima of curvature sharper on one side. We had three levels of convexity, where the first level corresponded to a boundary which is unbiased in terms of convexity, and the other two levels were referred as “weak convexity” and “strong convexity”. We had five levels of relative texture density for the convex set of regions, while the texture density of the non-convex set of region was fixed throughout. All the alternating light and dark regions included random dots moving horizontally at constant speed but in opposite direction in alternating regions. Subjects indicated the set of regions they saw as a single sheet translating in the background.

2.2.1 Method

Participants

Seven Rutgers University students who were naive to the purpose of the experiment participated in the experiment. All participants had normal or corrected-to-normal visual acuity and were paid for their participation.

Stimuli

The stimulus consisted of eight alternating black and white vertical regions. The height of the stimulus was 6.1^o and its width was 8.6^o. Either the odd or the even regions were given the convexity cue. Here, the term “convex” refers to a boundary that is piecewise convex such that the boundary can be segmented into parts, each of which is convex. The part boundaries are defined by the negative minima of curvature (Hoffman & Singh, 1997).

In this experiment, we used different degrees of convexity in which the convexity of a boundary was manipulated by gradually making the negative minima of curvature sharper on one side. In other words, the relative salience of the part boundaries on the two sides of a boundary was manipulated (Hoffman & Singh, 1997). This is done by creating contours with cubic spline interpolation and then manipulating the sharpness of the negative minima on one side of the contour by shifting the positions of the knots of the spline curve. Figure 2.3 shows the three different levels of convexity.

For each stimulus, seven individual contours were generated where the size (i.e. height and width) of each part of a contour was randomized within a certain

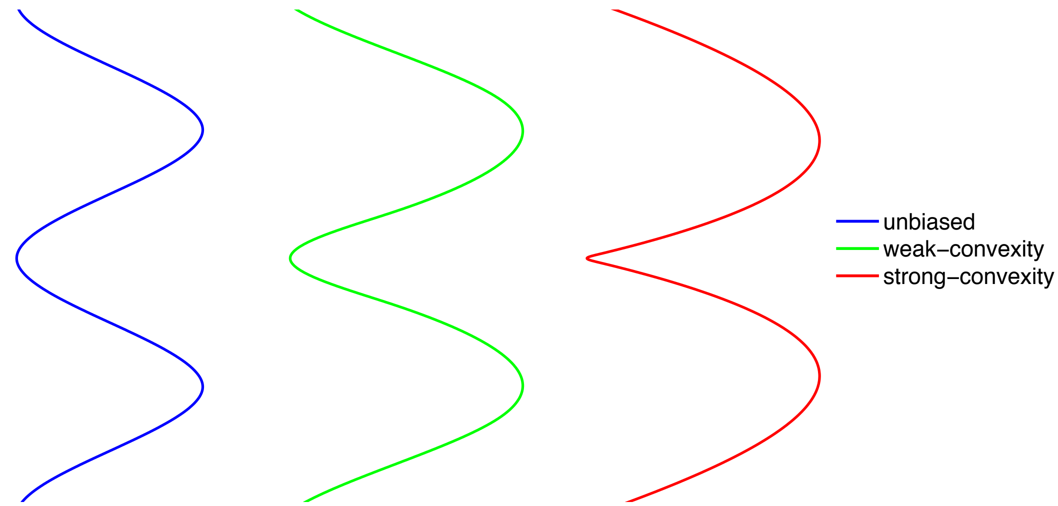

Figure 2.3: An example showing the relation between the sharpness of the negative minima of curvature and the (part-wise) convexity levels used in the experiment. The blue curve is similar to a sine wave where the two sides of the curve is unbiased in terms of convexity. The sharpness of the negative minima is gradually increased in order to obtain the green (weak-convexity) and the red curves (strong-convexity), in which the left side of each curve becomes piecewise convex.

range. Each contour had five and a half parts. Whether a contour starts or ends with a negative minimum or positive maximum of curvature was also randomized. These seven contours were placed on the stimulus such that the area between each consecutive pairs of contours is the same. For this experiment we had three levels of convexity which we refer as “strong convexity”, “weak convexity” and “unbiased” (Figure 2.4). For the unbiased contours the knots were set so that the integral of signed curvature was zero along the boundary.

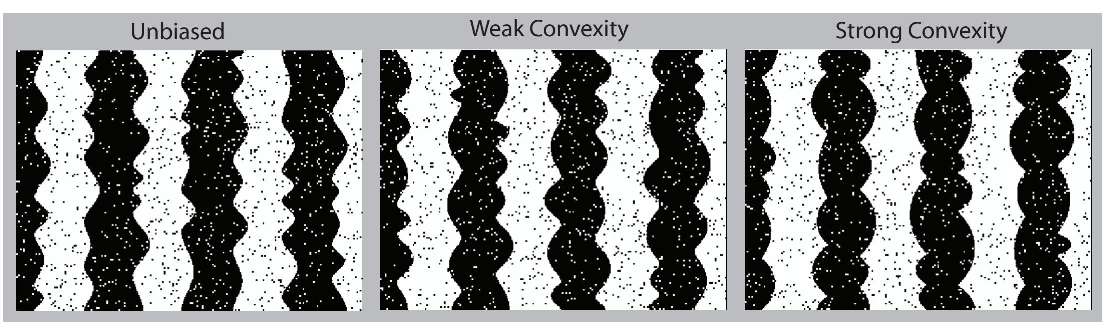

Figure 2.4: An example of the three levels of convexity used in Experiment 1.

To these stimuli, a certain number of dots were placed randomly within each black and white region. White dots were placed in black regions and black dots were placed in white regions. This dotted texture was created by creating a random-dot texture, which is sampled from a binomial distribution. Different parameter values of this binomial distribution were used to control the dot density in different regions. We used five different density levels, which were 1%, 2%, 4%, 8% and 16%. For example, in order to create a black region that includes 4% of white dots, a dot texture was sampled from a binomial distribution with a parameter of $[\theta = 0.04]$. In this way, whether a pixel in a given region would be white or black was determined by the result of this sampling, such that, on average, 4% of the pixels on that region became white while the rest was black. If we want to create a white region that includes 4% of black dots, a dot texture was sampled from a binomial distribution with a parameter of $[\theta = 0.96]$. Stimuli for other density levels were created in the same way but using different θ values. In the next step, the stimuli that were created in this way were resized to double of current size, in order to make the size of a single dot slightly larger (2.5 arcmin by 2.5 arcmin).

In this experiment, we manipulated the relative dot density of the convex regions to the non-convex regions. The dot density of the non-convex regions were fixed at 4%, whereas the dot density of the convex regions could take either of the five possible dot density levels. In this way, we had five levels of relative dot density (i.e. convex/non-convex): 1/4, 1/2, 1, 2 and 4 (Figure 2.5).

As a final step, motion was added to the dot texture of each region. The texture could either move to the right or to the left. For the rightward motion, in each frame t the texture columns $[2, N]$ were taken from texture columns $[1, N-1]$ in frame $t-1$. The color of the pixels (i.e. the texture) in the first column in frame t were resampled from the corresponding binomial distribution. The implementation was the same for the leftward motion. This procedure was repeated at a rate of 48 frames/sec, which resulted in a motion with a speed of 1.8^o/sec.

The odd regions were black with white dots and the even regions were white with black dots in half of the trials, and the opposite in the other half. Hence, the colors of the odd and the even regions were counterbalanced and crossed with other factors. The black and white regions always had textural motion in opposite directions. The

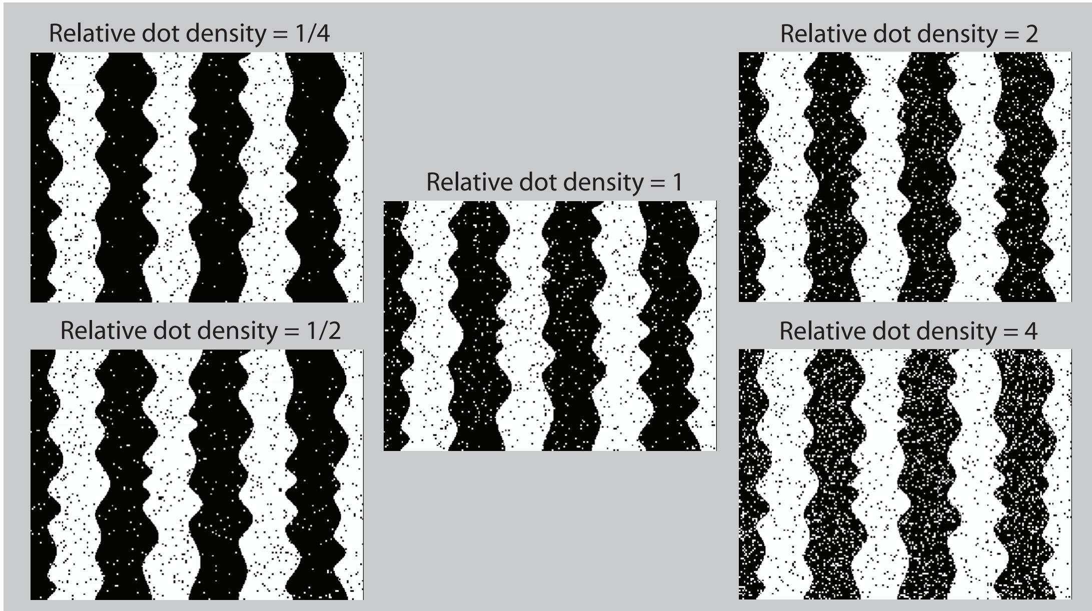

Figure 2.5: An example of the five levels of relative dot density (convex to non-convex regions) used in Experiment 1.

direction of motion was counterbalanced and crossed with other factors. Whether the rightmost part of the stimulus started with a convex or a non-convex region (i.e. the phase of the stimulus) was also counterbalanced and crossed with other factors.

Design and Procedure

The experiment was presented using Psychtoolbox in MATLAB (Brainard, 1997; Kleiner et al., 2007). Subjects sat 76 cm from a 21" CRT monitor (85 Hz, 1280pxl $\times$ 1024pxl) connected to a Windows 7 PC. The task was a forced-choice questions regarding the depth order of the light and the dark regions.

On each trial, subjects were presented with 800 ms of pre-mask, followed by another 800 ms of pre-mask with a fixation cross added on the center of the pre-mask. The masks were created by overlaying randomly generated figure-ground stimuli with multiple regions. The pre-mask was used in order to exert more careful stimulus control by diminishing any potential visual persistence of the previous stimuli. After the pre-mask and the fixations cross disappeared, the experimental displays with moving textures was shown to the subjects for 700 ms. After 700 ms, a post-mask that was identical to the pre-mask was shown for 800 ms. The post-mask was used in order to avoid any potential visual persistence of the stimulus after the termination of the display. Once the post-mask was presented, the subjects were asked to indicate which set of regions (dark or light) was perceived as a single sheet translating in the background. The subjects were instructed to respond with their first percept of depth order. Subjects responded either as "light" or "dark" by using the keyboard. The questions were forced choice.

Subjects ran 960 experimental trials in total split into 8 blocks, i.e. 3(convexity levels: unbiased; weak-convexity; strong-convexity) $\times$ 5(relative dot density of convex to non-convex regions: 1/4; 1/2; 1; 2; 4) $\times$ 2(luminance: dark/bright) $\times$ 2(phase) $\times$ 2(direction of motion: right/left) $\times$ 8(repetitions). All trials were randomized for each subject separately. The experiment was completed in two sessions done in different days with four blocks per session. 60 practice trials were run at the beginning of the first session in order to acquaint the subjects with the stimuli and the task. Another 30 practice trials were run at the beginning of the second session. Each session took approximately 45 minutes for each subject to complete.

2.2.2 Results

A multilevel logistic regression model was fit both to individual and to the aggregate data. The fitted models for each seven subjects and for the aggregate responses are shown in Figure 2.6. A likelihood ratio test on the aggregate data showed that including the main effect for convexity resulted in a significant improvement (LR = 824.54, df = 2, $p < .001$) over an unconditional means model (i.e. containing only an intercept). Addition of the main effect of relative texture density to this model led to also a significant improvement (LR = 46.48, df = 1, $p < .001$). To this model, the main effect of color (i.e whether the part-wise convex regions are dark or light) was added and it was also found to be a significant addition ($LR = 4.28$, $df = 1$, $p < .05$). Besides these main effects, the interaction between color and convexity (LR = 41.53, df = 2, $p < .001$) and between color and relative texture density (LR = 15.59.74, df = 1, $p < .001$) were also significant additions, yielding our final model.

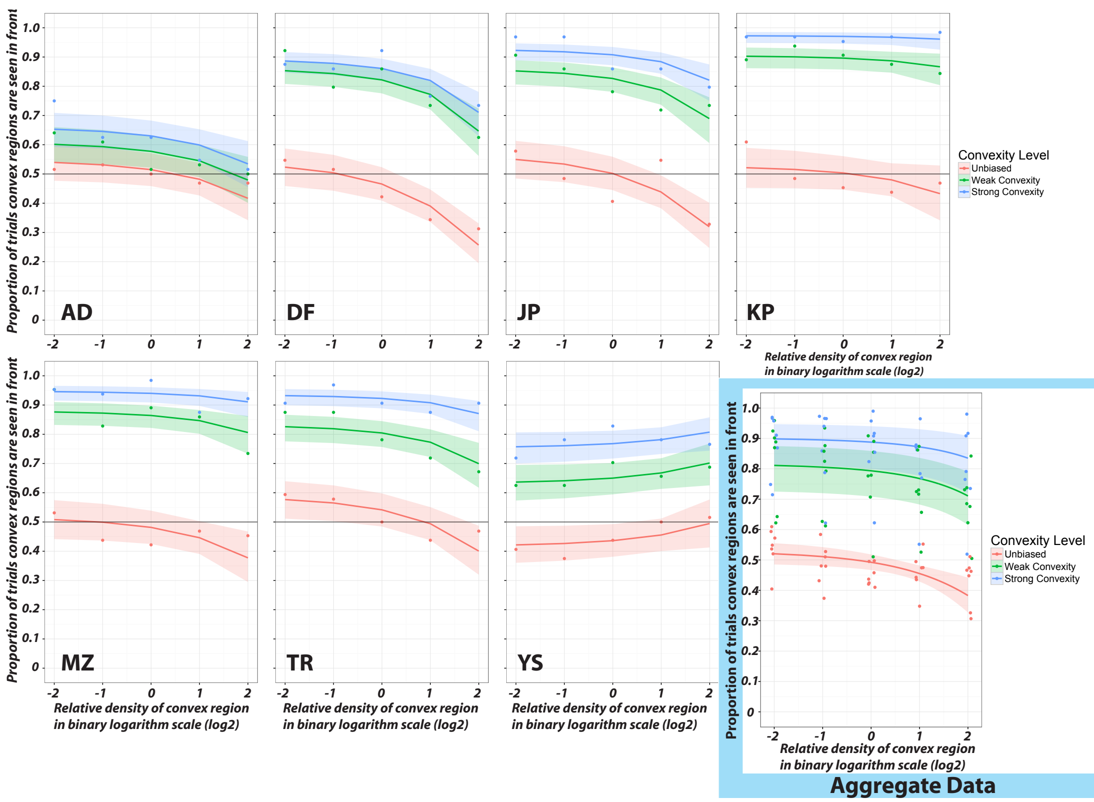

Figure 2.6: The logistic model fitted to the individual and to the aggregate data. The logistic model presented in this Figure includes convexity and texture density as fixed factors (and different participants as a random factor for the aggregate data). The x-axis corresponds to the relative texture density of the convex region to the non-convex region and the y-axis correspond to the proportion of trials in which subjects perceived the convex region in front. Different lines correspond to different convexity levels, and the ribbons around them indicate 95% confidence intervals. Since there is no convex side for the unbiased condition, the plots for the unbiased condition are generated with respect to a reference set of regions. The data points superimposed on the plots indicate the actual proportion of trials convex regions are seen in front for the corresponding condition. The blue lines show the chance level, i.e. where the proportion equals to 0.5.

As expected, increasing the degree of convexity in a set of regions resulted in an increase in the proportion of seeing those regions in front. This can be also seen in Figure 2.6 when the difference between the different lines in each plot (i.e. corresponding to different levels of convexity) are examined. Tukey pairwise comparisons yielded significant differences (for the mean proportion of trials the convex regions seen in front) between all the pair-wise comparisons of different convexity levels. (Unbiased: $M = 0.473$ $SE = 0.011$; Weak-convexity: $M = 0.762$ $SE = 0.043$; Strong-convexity: $M = 0.846$ $SE = 0.046$). When the individual data were examined, the significant effect of convexity was also seen in all seven subjects.

The main effect of relative texture density on the aggregate data was in the opposite direction of the main effect of convexity. As the relative texture density of the convex region increases, the proportion of seeing those convex regions in front decreases. Tukey pairwise comparisons done on the different levels of relative texture density revealed a significant difference between relative texture density of 4 ($M = 0.640$ $SE = 0.035$)and 1/4 ($M = 0.749$ $SE = 0.037$) ($p < .01$), and between 4 and 1/2 ($M = 0.717$ $SE = 0.034$) ($p < .05$). When the individual data were examined, it was seen that among seven subjects, responses from only two subjects (KP and YS in Figure 2.6) were unaffected by relative texture density.

For all individuals, color of the convex regions had a significant effect on their responses. However, the effect of color was not consistent among subjects. While some subjects had a bias towards perceiving light regions in front ($n = 2$), others had a bias towards perceiving dark regions in front ($n = 5$). The effect of color also interacted with relative texture density and convexity level. The effect of color became apparent when the region boundaries were unbiased in terms of convexity (compared to when they are convex), and also when relative texture density of the convex regions was high.

2.2.3 Discussion

In this experiment we examined how the geometric figure/ground cue of convexity and the relative density of accreting/deleting texture combine to determine depth order. As mentioned in the Introduction, while convexity is considered as a strong geometric cue to figural status, accretion-deletion is traditionally considered as a decisive cue to ground status. The main effects we have seen in Experiment 1 were consistent with these expectations: As the level of convexity increased in a region, subjects were more likely to perceive that region in front; whereas as the relative density of the accreting/deleting texture increased in a region, subjects were more likely to see that region in behind. However, when Figure 2.6 is examined it is also seen that the effect of convexity dominated when it combined with accretion-deletion. The proportion of seeing a region in front went below 50% only when the boundary was unbiased in terms of convexity (i.e. the red lines in Figure 2.6) and the relative texture density being accreted-deleted was high. Even when a small degree of convexity was introduced to the boundary (i.e. the green lines in Figure 2.6), the dominant percept became seeing the convex regions as figural, regardless of the relative texture density of the region. However, apart from the two subjects (KP and YS Figure 2.6) who didn't show any effect of relative texture density, the effect of relative texture density can still be seen for each convexity level condition of the other five subjects.

Apart from our main manipulations, the color (dark vs. light) of the convex regions had also an effect on subjects' responses. However, the direction of the effect varied across individuals; while some subjects were more likely to perceive the light regions in front, others were more likely to perceive the dark regions in front. This was observed before in a similar figure/ground study Tanrikulu et al. (2016). The interaction effect between color and our main manipulations showed us that when the figure/ground interpretation is relatively more ambiguous — such as when the boundary is unbiased in terms of convexity or when the relative texture density is high and competing against the convexity cue — then subjects' color biases starts to have an effect on their responses.

According to the traditional view of accretion/deletion, whenever there is accretion/deletion of texture on both sides of a border, the region that contains greater amount of accretion/deletion per interval of time would be perceived as behind the region that has a lower amount of accretion/deletion per interval of time (Kaplan, 1969). This view is consistent with our results from Experiment 1, since the relative density of accreting/deleting texture is directly linked to rate of accretion-deletion disparity between two regions. Alternatively, the effect of texture density observed in this experiment can also be explained without assuming that accretion-deletion is a cue to ground status. As mentioned before, according to our account, accretion/deletion can be interpreted in two different ways. Increasing the density of an accreting/deleting texture that has constant speed profile would provide more information in support of a flat translating surface, which in turn bias subjects to perceive the accreting/deleting texture in behind. This alternative explanation will be further discussed in more detail in the General Discussion section.

2.3 Experiment 2

Another method for manipulating the rate of accretion/deletion on one side of a border is to vary the overall speed of the moving texture. In Experiment 2, we used a similar stimulus (i.e. the alternating light and dark regions included random dots moving horizontally at constant speed but in opposite direction in alternating regions) in order to look at how the degree of convexity of the borders and the relative speed of the accreting/deleting texture on the convex regions combine with each other in order to yield a figure/ground interpretation. Similar to the experimental design in Experiment 1, both factors were manipulated and combined in various condition: We had three levels of convexity (i.e. unbiased in terms of convexity, weak convexity and strong convexity) and five levels of relative texture speed of the convex regions. Subjects indicated the set of regions they perceived as a single sheet translating in the background.

2.3.1 Method

Participants

Seven Rutgers University students who were naive to the purpose of the experiment participated in the experiment. Five of these seven participants were among the subjects who also participated in Experiment 1. All participants had normal or corrected-to-normal visual acuity and were paid for their participation.

Stimuli and Procedure

The stimuli were generated in the same way as in Experiment 1, except that we included different conditions with varying texture speeds, while the texture density was equal (4%) in all regions (the middle image in Figure 2.5).

The texture speed of the non-convex regions was fixed at $1.8^o/s.$, whereas the speed of the texture on the convex regions could have five different values: 0.6, 0.9, 1.8, 3.6 or $5.4^o/s$. This yields five levels of relative texture speed of the convex regions (convex/non-convex): 1/3, 1/2; 1, 2 and 3.

The experimental procedure was exactly the same as in Experiment 1. Subjects ran 960 experimental trials in total split into 8 blocks, i.e. 3(convexity levels: unbiased; weak-convexity; strong-convexity) $\times$ 5(relative texture speed of convex to non-convex regions: 1/3; 1/2; 1; 2; 3) $\times$ 2(luminance: dark/bright) $\times$ 2(phase) $\times$ 2(direction of motion: right/left) $\times$ 8(repetitions). All trials were randomized for each subject separately. The experiment was completed in two sessions done in different days with each session included four blocks. 60 practice trials were run in the beginning of the first session in order to acquaint the subjects with the stimuli and the task. Another 30 practice trials were run in the beginning of the second session. Each session took approximately 45 minutes for each subject to complete it.

2.3.2 Results

A multilevel logistic regression analysis was applied to the subjects' responses. The fitted models for each seven subjects and for the aggregate responses are shown in Figure 2.7. A likelihood ratio test on the aggregate data showed that including the main effect of convexity and relative texture speed were significant improvements over an unconditional means model (i.e. containing only an intercept). (Comparing models that include convexity to the unconditional means model: $LR = 1366.8$, $df = 7$, $p < 0.001$; Comparing model that includes convexity and relative texture speed to the one that only includes convexity: $LR = 303.86$, $df = 5$, $p < 0.001$). Adding the effect of color (i.e. whether the convex region is dark or light) was also a significant expansion to this model ($LR = 277.3$, $df = 1$, $p < 0.001$). In addition to those, interaction effects between convexity and relative texture speed ($LR = 148.61$, $df = 7$, $p < 0.001$), and between convexity and color ($LR = 14.55$, $df = 2$, $p < 0.001$) were also a significant addition, yielding our final model.

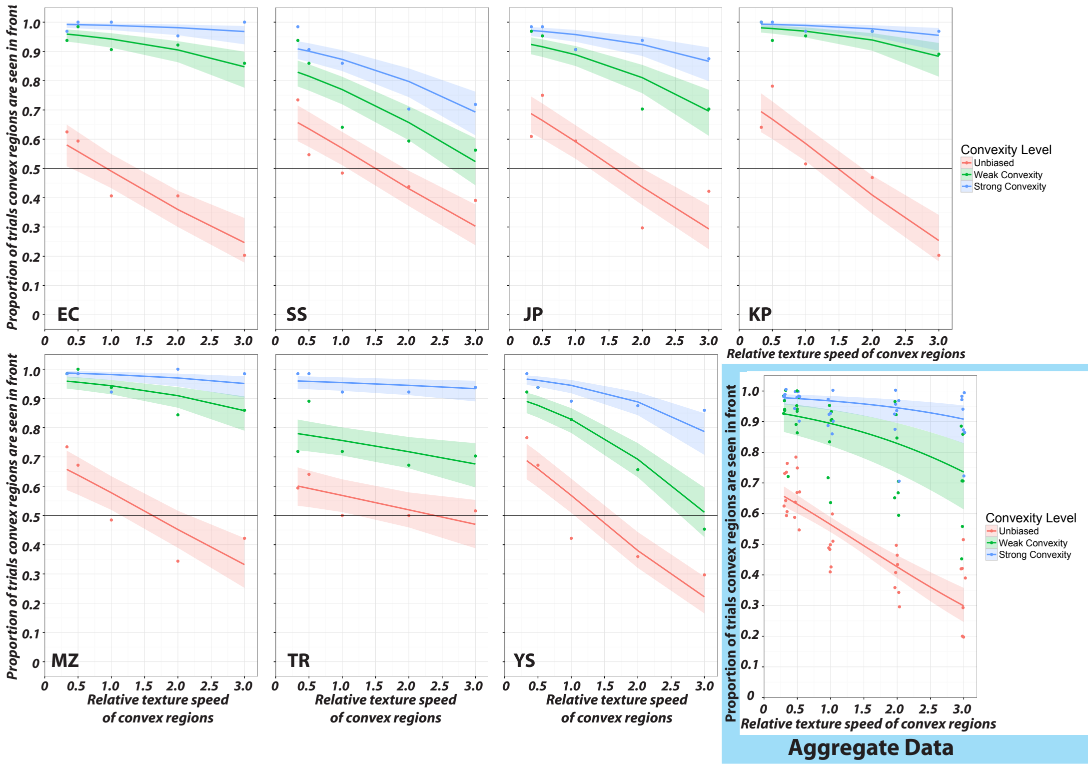

Figure 2.7: The logistic model fitted to the individual and to the aggregate data. The logistic model presented here includes convexity and texture density as fixed factors (and different participants as a random factor for the aggregate data). The x-axis corresponds the relative texture speed of the convex region to the non-convex region and the y-axis correspond to the proportion of trials in which subjects perceived the convex region in front. Different lines correspond to different convexity levels, and the ribbons around them indicate 95% confidence intervals. Since there is no convex side for the unbiased condition, the plots for the unbiased condition are generated with respect to a reference set of regions. The data points superimposed on the plots indicate the actual proportion of trials convex regions are seen in front for the corresponding condition. The blue lines show the chance level, i.e. where the proportion equals to 0.5.

The main effect of convexity was very similar to the one obtained in Experiment 1. As the degree of convexity increased, the proportion of seeing the convex regions in front also increased. This can be also seen in Figure 2.7 from the difference between the different lines in each plot that corresponds to the different level of convexities. Tukey pairwise comparisons revealed that all the pair-wise comparisons among convexity levels were significantly different from each other ($p < 0.001$; Unbiased: $M = 0.515$ $SE = 0.013$; Weak-convexity: $M = 0.838$ $SE = 0.037$; Strong-convexity: $M = 0.939$ $SE = 0.02$). When individual subjects' data were analyzed, it was also seen that effect of convexity was a significant factor in each subject's responses.

It was observed that as the relative texture speed in a region increases, the proportion of seeing that region in front decreases. This can be seen in Figure 2.7 from the negative slope of each plot. Tukey pairwise comparisons revealed that all pairwise comparisons among different levels of relative texture speed were significant ($p < 0.01$) except the differences between relative texture speeds of 0.3 and 0.5, and between 2 and 3 (Relative texture speed of "0.3": $M = 0.86$ $SE = 0.018$; "0.5": $M = 0.858$ $SE = 0.017$; "1": $M = 0.751$ $SE = 0.021$; "2": $M = 0.692$ $SE = 0.03$; "3": $M = 0.658$ $SE = 0.03$). When individual subjects' data were examined, it was seen that all the subjects showed the significant main effect of relative texture speed. When the interaction effect between convexity and relative texture speed was examined, it was seen that the effect of relative texture speed becomes more apparent when the region border was unbiased in terms of convexity. When the individual data was analyzed, it is seen that only two (YS and MZ) out of seven subjects showed this significant interaction.

Except for one subject (KP), color of the convex regions had a significant effect on the responses of the other six subjects. Among those, five of them had a bias towards perceiving dark regions in front, whereas the remaining one subject had a bias for light regions. The effect of color also interacted with the effect of convexity. The effect of color became apparent when the region boundaries were unbiased in terms of convexity (compared to when they are convex). When individual data was analyzed, this significant interaction effect was observed for four individuals among seven.

2.3.3 Discussion

In this experiment, our goal was to examine how the geometric figure/ground cue of convexity combines with relative overall speed of the accreting/deleting texture to yield a figure/ground interpretation. As the degree of convexity increased in a set of regions, subjects were more likely to see those regions in front. As the relative overall speed of the accreting-deleting texture increased in a set of regions, those regions were more likely to be seen in behind. Whenever the two cues combined, convexity seemed to dominate. This is also seen when Figure 2.7 is examined. The green (weak convexity) and the blue (strong convexity) lines in most of the plots, are clearly separated from the orange line (unbiased in terms of convexity), which means that as soon as some degree of convexity is introduced to the border, the dominant percept becomes seeing the convex region in front.

Apart from the main effects of convexity and relative texture speed, these two factors also yielded a significant interaction effect. This interaction can be seen in Figure 2.7. In the graph showing the results for the aggregate data, it is seen that the slope (i.e. the effect of relative texture speed) of the three lines do not seem to be equal. As the convexity level increases (as we go from the orange line to green line and then to blue line), the (absolute value) of the slope decreases. In other words, as we increase the level of convexity, the effect of relative texture speed decreases. The effect of relative texture speed is the greatest when the geometry of the border is unbiased in terms of convexity. This shows that when the geometry of the border does not convey any information about the depth order of the regions, then subjects began relying more on the information coming from the motion of the texture. A similar interaction effect was also observed between color and convexity. When the border is unbiased in terms of convexity, hence provide no information regarding the depth order of the regions, subjects' responses were effected by their individual color biases.

The effect of relative texture speed is also in the expected direction according to the traditional accounts of accretion-deletion. As the rate of accreting-deleting texture is increased on one side of the border, this would bias subjects to perceive the side that has a higher rate of accretion/deletion in the back. However, the effect of relative texture speed can also be explained by our alternative account, which claims that accretion/deletion is not necessarily a cue to ground status, but can be interpreted as either in front or in behind. In a cosine motion profile, which is consistent with 3D rotation in depth, the image speed of the texture elements that are very close to the borders should drop down to zero. Hence, when the overall texture speed is high and constant everywhere on the image, it is easy for our visual system to detect this discrepancy between the constant and cosine speed profiles. Making this discrepancy more apparent by increasing the overall speed of the texture introduces a bias toward to the traditional interpretation of accretion-deletion. These two different accounts for the effect of relative texture speed will be further discussed in the General Discussion section.

2.4 General Discussion

Traditionally accretion/deletion of texture has been considered as a decisive cue to ground status in determining relative depth. Accretion/deletion is thought to arise only from a surface being occluded by another figural region. However, recently this traditional account has been challenged (Froyen et al., 2013; Tanrikulu et al., 2016). Accretion/deletion can also arise from self-occlusion due to rotating in depth, in which accreting/deleting region is interpreted to be figural. It was observed that the geometry of the occluding contour has a strong influence on how accretion/deletion is interpreted. However, how accretion/deletion combines with the geometry of the border have not been studied in a systematic way so far. In this study, we manipulated the strength of accretion/deletion (it its traditional sense) and the degree of convexity of the boundary where the texture accretes/deletes, and investigated how the two factors combine to determine relative depth.

In two experiments, we observed that while increasing the level convexity of a region makes that region more likely to be perceived in front, increasing the density or the speed of the accreting/deleting texture in a region makes that region more likely to be perceived in behind. However, convexity seems to dominate the relative depth interpretations in general. For example, even in the condition in which the speed or the density of the accreting/deleting regions were the highest, introducing weak convexity to those accreting/deleting regions makes those region to be perceived in front most of the time (about 70%).

According to traditional accounts of accretion/deletion, the rate of accreting/deleting texture determines the strength of accretion/deletion. As the rate of accretion/deletion increases in a region, the probability of seeing that region in behind increases (Kaplan, 1969). In our study, we observed that as we increased the speed or the density of accreting/deleting texture in a set of region, the proportion of trials those region perceived to be in behind increased. This result does not directly contradict the traditional view of accretion/deletion. Since accretion/deletion is considered as a cue to being in behind (according to the traditional accounts), increasing the density and the speed of the accreting/deleting texture would make these regions more likely to be perceived in behind, because increasing texture and density also increases the rate of accretion/deletion.

However, as mentioned before, there are recent studies that challenge the traditional view of accretion/deletion. According to these recent studies, accretion/deletion has more than one interpretation, which makes it consistent with both "in-front" and "in-behind" interpretations. Then, according to this view, rate of accretion/deletion can not be the critical factor that directly determines how accretion/deletion is going to influence the relative depth interpretation. Since accretion/deletion is not considered as a cue to being in behind in this view, the rate of accretion/deletion can not account for the observed effect of texture density and speed.

If the rate of accretion/deletion is the critical variable that determines perception of relative depth, then one might expect that the effect of texture density and speed would overlap when expressed as a function of rate of accretion/deletion. Therefore, we calculated the rate of accretion/deletion for all the different texture density and speed conditions from both experiments. Then, we examined the proportion of trials a region is perceived to be in front as a function of rate of accretion/deletion in that region. This allowed us to combine the data from the two experiments and examine whether the rate of accretion/deletion is the critical factor that explains the both observed effects of texture density and speed (Figure 2.8).

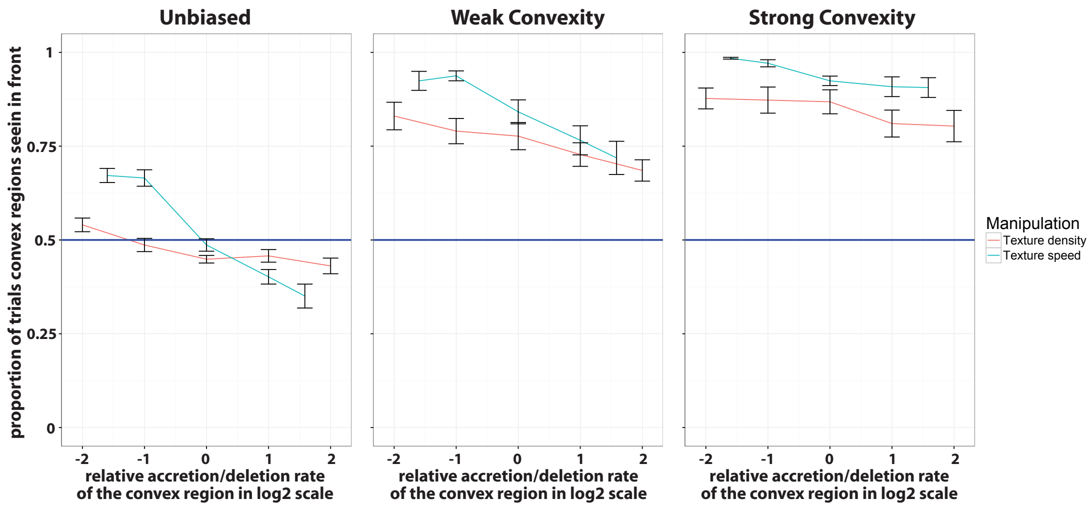

Figure 2.8: Proportion of trials in which the convex regions are perceived in front as a function of relative rate of accretion/deletion in those convex regions (for each convexity level). Error-bars represent $\pm 1SE$ as computed between subjects. The blue line shows the chance level, i.e. where the proportion equals to 0.5. In the unbiased condition, since no region is more convex than the other, the proportions are calculated with respect to a reference set of regions.

Figure 2.8 shows the proportion of trials the convex regions were perceived in front as a function of relative rate of accretion/deletion of the convex regions. The different lines correspond to the two different ways the rate of accretion/deletion is manipulated (i.e. texture density in Experiment 1 and texture speed in Experiment 2). As it is seen from the figure, the effects of texture density and texture speed manipulations (the red and the turquoise line plots) do not line up with each other when they are combined with respect to a common factor, that is the rate of accretion/deletion. In other words, the way the rate of accretion/deletion is manipulated makes a difference in its effect on relative depth judgments. Even though there seems to be a slight negative correlation between relative rate of accretion/deletion in convex regions and the proportion of trials convex regions are seen in front, the difference between the two line plots indicates that there is an influence of texture density and speed that can not be captured by the rate of accretion/deletion.

The effects of density and overall speed of accreting/deleting texture on relative depth judgments can also be accounted by the alternative view of accretion/deletion that was put forward by recent studies (Froyen et al., 2013; Tanrikulu et al., 2016). These studies have shown that accretion/deletion combines with the geometry of the border to determine relative depth. Geometry of the border had an influence on how accretion/deletion is interpreted (i.e. occluded by the adjacent surface vs. self-occlusion due to rotation in depth). It is also highly reasonable to expect that visual information about the speed profile of the accreting/deleting texture would also have an effect on how accretion/deletion is interpreted. If the information coming from the speed profile of the accreting/deleting texture is clearly biasing judgments towards (i.e. is more consistent with) a translating flat surface interpretation, then accreting/deleting regions would be more likely to be perceived as background. If the speed profile of the accreting/deleting texture is more consistent with a rotation-in-depth interpretation, then the accreting/deleting regions would be more likely to be perceived as figural regions. Therefore, not only the geometry of the border but also the motion information would deeply interact with how accretion/deletion is interpreted to determine relative depth.

In our first experiment, we observed that increasing relative texture density in a region makes that region more likely to be perceived in behind. Such a result would be consistent with the explanation summarized above. The speed profile of the dot texture in our experiment was constant, which, according to traditional structure-from-motion models (Ullman, 1979), indicates a translating flat surface. When there are more texture elements that follow a constant speed profile, it means that there is more information available that favors a translating flat surface interpretation of accretion/deletion. This, as a result, favors the "in-behind" interpretation of accretion/deletion. That is why as the relative texture density of an accreting/deleting region is increased, that region would be more likely to be perceived in behind.

This alternative view of accretion/deletion could also account for the result of our second experiment. A 2-D projection of texture elements on a column that rotates in depth would yield a cosine speed profile, where the speed of the texture elements would be approximately zero when they are around the border at which they accrete/delete. Therefore, if the speed of the texture elements that are accreting/deleting is high around the border, this would favor a translating flat surface interpretation of the accreting/deleting texture, rather than "a self-occlusion due to rotation in depth" interpretation. Once accretion/deletion is perceived as a translating flat surface, this would favor the "in-behind" interpretation of accretion/deletion. For that reason, when the overall speed of the accreting/deleting texture in a region is increased, that region would be more likely to be perceived in behind. This is actually what we observed in our second experiment; increasing the speed of the accreting/deleting texture (that was following a constant speed profile) makes that region more likely to be seen in behind.

This alternative account would also suggest an explanation for the discrepancy observed between the effects of two manipulations used in this experiment (Figure 2.8). Even though both effects are explained by the influence of motion information on how accretion/deletion is interpreted, the exact nature of this influence is different in these two types of manipulation. With respect to texture density manipulation, the influence was about adding or removing more available information about the speed profile of the moving dot texture. For the texture speed manipulation, the influence was due to the increased discrepancy between a cosine speed profile and the texture speed around the region borders on the image. This might explain why the effects of the two manipulations were different from each other.

Manipulating density and the overall speed of accreting/deleting texture are two different ways of manipulating rate of accretion/deletion, which is traditionally considered as the critical factor that determines the strength of accretion/deletion as a cue to ground status. However, when Figure 2.8 is examined, it is seen that rate of accretion/deletion can not fully account for the effects of both manipulations. As explained above, this is not surprising if we consider accretion/deletion as a factor that is consistent with both relative depth interpretations (i.e. rotating-in-front vs. translating-in-behind), instead of as a definite cue to be in behind. When certain features of the accreting/deleting texture are manipulated, the influence of that manipulation to relative depth judgments would depend on whether that manipulation makes the dynamic image more consistent with a "rotating-in-front" or with a "translating-in-behind" interpretation of accretion/deletion. The effects of texture density and speed observed in the two experiments are clear examples of this. Increasing the overall speed and the density of the accreting/deleting texture in our displays makes the stimulus more consistent with a translating flat surface interpretation that would yield a constant speed profile (as opposed to a rotation in depth interpretation that would yield a cosine speed profile).

In this study, we examined how accretion/deletion of texture and border convexity combine to determine relative depth judgments by systematically manipulating the strength of each factor. We varied the strength of accretion/deletion by manipulating the rate of accreting/deleting texture elements in two different ways. In Experiment 1, we manipulated the relative density of the accreting/deleting texture on one side of each border. In Experiment 2, we manipulated the relative overall speed of the accreting/deleting texture on one side of each border. However, we found that convexity dominates relative depth judgments when combined with accretion/deletion, even in cases where the degree of convexity introduced to the border is relatively weak. Motion is usually considered to be a strong cue to relative depth and 3D structure. However, our experiments showed that a simple static geometric cue can actually exert a stronger influence than motion.

We also observed that there is an influence of texture density and speed manipulations that can not be fully captured by rate of accretion/deletion. These results suggest that considering accretion/deletion as a factor that is consistent with both "in-front" and "in-behind" interpretations (instead of considering it as a definite cue to being in behind, as traditional accounts do) is a better way to understand how accretion/deletion influences relative depth judgments. In this experiment, we observed that convexity of the border, relative density and relative speed of the accreting/deleting texture has a strong effect on how accretion/deletion is interpreted, which as a result determines how accretion/deletion would influence relative depth judgments. Certainly, more experiments should be done in order to further support this claim, such as manipulating the speed profile of accreting/deleting texture and examining its influence on relative depth judgments. However, our current results suggest that considering accretion/deletion as a definite and unambiguous cue to ground status is not the correct way to understand the influence of accretion/deletion on relative depth.

3. Bridging the gap between standard depth from accretion/deletion and rotating-columns interpretation

3.1 Introduction

One of the early stages of visual perception requires segmenting the retinal image into separate regions and then assigning figure and ground status to those regions. When two adjacent regions in the retinal image share a common border, the visual system has to decide which region owns this common border. The region that owns the border (i.e. figure) is shaped by this border and perceived to be closer to the viewer than the adjacent region. The other region (i.e. background) is perceived to extend amodally behind the figural region. Starting with Rubin (1915/1958) and the Gestalt psychologists, numerous different factors that tend to influence figure/ground assignments have been identified, such as area, enclosure (Rubin, 1915/1958), symmetry (Kanizsa & Gerbino, 1976), convexity (Metzger, 1936/2006; Kanizsa & Gerbino, 1976), parallelism (Morinaga, 1941; Metzger, 1936/2006), lower-region (Vecera, Vogel, & Woodman, 2002), axiality and part salience (Hoffman & Singh, 1997; Froyen et al., 2010) (see Wagemans et al. (2012) for a review).

Some of these factors that tend to influence figure/ground assignments are defined by the static features of an image, such as geometric cues (e.g. convexity, parallelism, symmetry). Others rely on the dynamic features of an image, such as accretion/deletion of texture at region boundaries. Accretion/deletion as a cue to relative depth was first proposed by Gibson et al. (1969) and later experimentally demonstrated by Gibson's student Kaplan (1969). Kaplan (1969) used a display that includes a rectangular randomly textured image divided into two adjacent regions. When the image was static it looked like a single continuous textured surface. When the texture in one of the regions moved horizontally in either direction, a motion-defined contour that separates the two regions became visible. The motion-defined border and one of the regions were static while the texture on the other region was translating either toward left or right with constant speed. As a result, the moving texture progressively accreted or deleted at the boundary. When subjects were shown such a display and asked to indicate relative depth, they perceived the region that had the moving texture behind the static region. In this way, Kaplan (1969) demonstrated that when a moving texture on a region progressively accretes or deletes at a border, that region is perceived to be appearing or disappearing behind the adjacent region at the other side of that border. This leads to the interpretation that the accreting/deleting region is occluded by the adjacent region, and, as a result, is assigned ground status. Kaplan (1969) also tested displays where the border moved horizontally and independently of the moving texture on one side, as well as displays where the texture on both sides of the border was accreting/deleting. He also manipulated the relative speed of the accreting/deleting textures on different sides of the common border. In the end, he concluded that the region that had the higher rate of accreting/deleting texture was consistently more likely to be perceived as behind.

Following the work by Kaplan (1969), accretion/deletion became an essential part of studies that focus on relative depth from motion. Granrud et al. (1984) demonstrated that 5- and 7-months-old infants are also sensitive to accretion/deletion as a cue to being behind. Yonas et al. (1987) observed that when the border separating the two regions is not static, the relative motion of the border and the texture provides reliable information about depth order, independently of the accretion/deletion of the texture. However, Profitt, Bertenthal, and Roberts (1984) demonstrated that accretion/deletion can still determine depth ordering in the absence of this relative-motion cue between the texture and the border. It was also shown that, in certain cases, accretion/deletion can override depth from binocular disparity (Royden et al., 1988; Hildreth & Royden, 2011), from lower-region cue (Royden et al., 1988) and from motion parallax (Ono, Rogers, Ohmi, & Ono, 1988; Hildreth & Royden, 2011). It can resolve ambiguities seen in bi-stable motion displays, such as direction of rotation ambiguity created when viewing parallel projection of an object rotating in depth (Braunstein, Andersen, & Riefer, 1982), or the ambiguity in the overall structure from point-light biological motion displays (Profitt et al., 1984). As a result, the assumption that accreting/deleting surfaces are invariably interpreted as behind began to appear in textbooks (Niyogi, 1995; Howard & Rogers, 2002; Goldstein, 2014), and has been incorporated into computational models of depth from motion (Thompson et al., 1985; Mutch & Thompson, 1985; Berzhanskaya et al., 2007; Beck et al., 2008; Raudies & Neumann, 2010; Barnes & Mingolla, 2013; Layton & Yazdanbakhsh, 2015; Ruda et al., 2015).

However, it is important to note that accretion/deletion can also arise from dynamic self-occlusion when a 3D object rotates in depth (Figure 3.1). For example, imagine that you are looking at a bottle rotating in depth around its vertical central axis. As the bottle rotates, the text on the label of the bottle would be deleted on one side of the bottle, while it would be accreting on the other side. In such cases, the texture would accrete/delete not because it is being occluded by the adjacent region, but because it is being occluded by itself as the bottle rotates in depth. In such situations, the accreting/deleting surface might be incidentally interpreted as rotating in front, contrary to the traditional assumption. This depth reversal has been reported by several researchers (Kaplan, 1969; Thompson et al., 1985; Mutch & Thompson, 1985; Yonas et al., 1987), but it has been generally ignored and has not been seriously considered as a possible interpretation of accretion/deletion as a cue to relative depth.

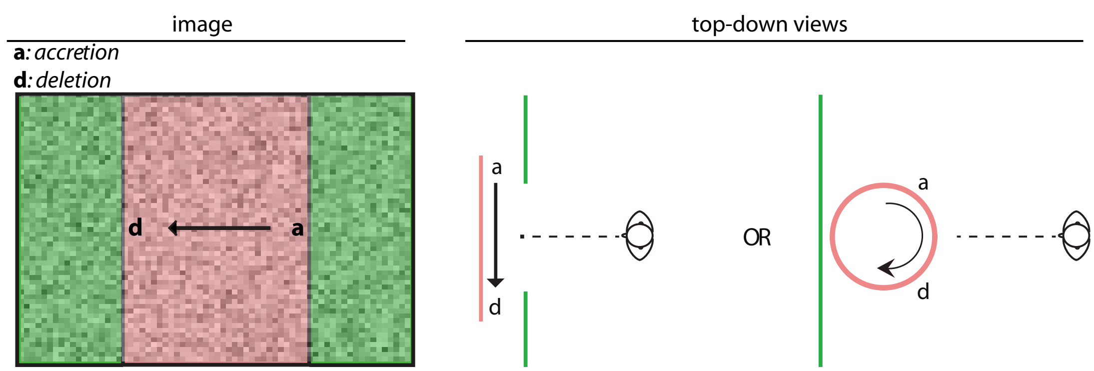

Figure 3.1: The frontal projection of an accreting/deleting surface is shown on the left. "a" is the location of texture accretion and "d" is the location of texture deletion. The static surface is depicted in green while the moving surface is depicted in red. On the right, overhead views of the two possible 3D arrangements with different depth-order assignments that are both consistent with the frontal view of the accreting and deleting surface are shown.

Recently, a series of studies challenged the traditional account of accretion/deletion by focusing on this alternative interpretation of accretion/deletion. By using a display containing alternating light and dark regions with random dot texture moving horizontally at constant speed in opposing direction in alternating regions (e.g. similar to the display shown at the rightmost side of Figure 3.2), Froyen et al. (2013) demonstrated that when accretion/deletion was present on both sides of each border, either the dark or the light regions were perceived to be rotating columns in front, while the other regions were amodally combined into a single large surface translating at the back. When geometric figure/ground cues (e.g. convexity, symmetry, parallelism) are introduced to a set of regions, those regions were more likely to be perceived as rotating columns in front. They concluded that accretion/deletion combines with the geometry of the border to determine relative depth.

Using a similar multi-region figure/ground display, Tanrikulu et al. (2016) demonstrated that the rotating column interpretation of accretion/deletion can be obtained even when one set of regions is static, and accretion/deletion is present on only one side of each border. Since, any ambiguity that could arise from having accretion/deletion on both sides of each border was eliminated, the depth ambiguity observed in Tanrikulu et al. (2016) seemed to have arisen from the ambiguity inherent to the accretion/deletion itself. Hence, they argued that accretion/deletion should not be considered as a definite cue to ground status. In another study, Tanrikulu, Froyen, Feldman, and Singh (2014) investigated how a geometric figure/ground cue (i.e. convexity) and accretion/deletion combine to determine relative depth. They manipulated the strength of each cue and combined them in various conditions. The strength of the accretion/deletion cue was varied by manipulating the rate of accreting/deleting texture elements, which is the factor that Kaplan (1969) and others since have used to determine the strength of accretion/deletion. This was done in two different ways; by varying either the relative density or the overall speed of the accreting/deleting texture in one set of regions. Contrary to the traditional accounts, their analysis indicated that the effect of texture density and speed is not simply mediated by the rate of accretion/deletion.

As summarized above, traditional accounts of accretion/deletion have considered it as a definite cue to being behind. The recent studies outlined above challenged this view and demonstrated that accreting/deleting regions can also be perceived in front, in which accretion/deletion arises because of self-occlusion due to rotation in depth. The different results observed in the traditional studies on accretion/deletion and in the more recent studies created a mystery on the depth information conveyed by accretion/deletion. In this current study, our main goal is to unravel this mystery by investigating which stimulus factors that are only present in these more recent studies but not in the traditional studies are responsible for the rotating column interpretation of accretion/deletion.

The displays used by recent studies differ from the standard accretion/deletion displays along a number of variables. Even though accretion/deletion displays used in the literature vary from one study to another, they all share some common features that can be traced back to the original stimulus used by Kaplan (1969). These traditional displays generally include two rectangular adjacent regions separated by a straight common border. In some cases, the stimulus is divided into a central region and a surrounding region, where the shape of the central region is generally rectangular such that the border at which the texture accretes/deletes is still straight (e.g. Royden et al., 1988; Thompson et al., 1992). In these displays, accretion/deletion is generally introduced to the one side of the border, and the texture on the other is made static. However, the recent studies that focus on the rotating-in-front interpretation of accretion/deletion use multi-region figure/ground displays in which the shape of the borders include some degree of curvature and/or geometric cues to figure/ground. Moreover, in these recent studies accretion/deletion is sometimes introduced on both sides of each border instead of only on one side as used in traditional displays.

These differences between the traditional displays of accretion/deletion and the displays used by recent studies create a gap between these two lines of research. In this study, our goal is to bridge this gap by investigating which stimulus factors are critical in promoting the rotating-in-front interpretation of accretion/deletion. This will also allow us to link the results of these more recent studies to the traditional studies on accretion/deletion, and understand the factors that influence the relative depth information conveyed by accretion/deletion. In order to do this, we will be manipulating each factor that makes the displays used by the recent studies different than the traditional accretion/deletion displays, and examine how these factors influence the interpretation of accretion/deletion. This would also inform future computational models of depth-from-motion about the stimulus factors that cause accretion/deletion to be interpreted as a surface behind or in front — since accretion/deletion is, in principle, consistent with both interpretation.

Figure 3.2 shows the variables that are manipulated in this study. Traditional accretion/deletion displays consist of two regions separated by a straight border in which accretion/deletion of random dot texture is introduced only on one side of the border, while the texture on the other side is static. We manipulate three variables that potentially differentiate traditional displays from rotating-columns displays: the number of regions, geometry of the border and whether accretion/deletion is present on one side or both sides of each border. There are two different ways to manipulate the number of regions in the display. One can either keep the width of a single region, or of the whole stimulus, fixed as the number of regions is manipulated. In Experiment 1, as the number of regions is varied (i.e. 2, 4 or 8 regions), the width of a single region is kept fixed. Hence, the width of the whole stimulus changes with the number of regions in the display. In Experiment 2, the width of the whole stimulus is kept fixed while the number of regions is varied. As a result, the width of a single region changes with the number of regions included in the stimulus. We will be investigating which of these factors are critical in determining whether accreting/deleting regions are interpreted as a surface in behind or in front.

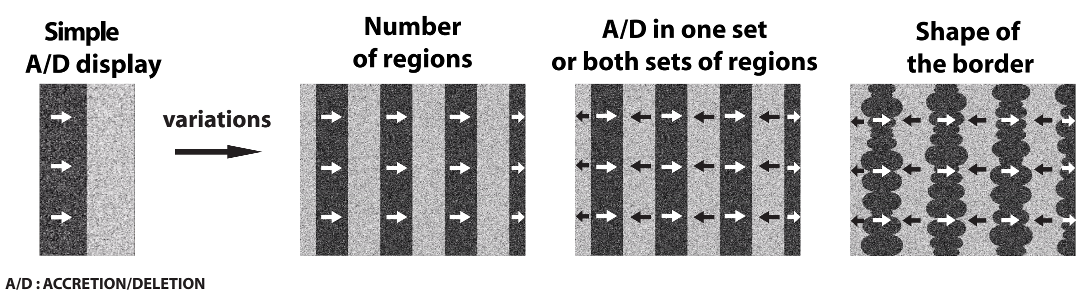

Figure 3.2: A traditional accretion/deletion stimulus and variations introduced to it by the recent studies that focus on the "rotating-in-front" interpretation of accretion/deletion.

In traditional accretion/deletion displays, the regions are often made isoluminant and the borders are only defined by motion contrast. However, it has been shown that the presence of luminance contrast actually promotes the traditional interpretation of accretion/deletion as a surface in the back. Kromrey et al. (2011) demonstrated that when an enclosed central region that includes accreting/deleting random-dot texture is surrounded by a flickering random dot texture, the accreting/deleting central region is perceived to be in front of the surrounding region. However, when the border between the central and the surrounding region is defined by luminance contrast, the traditional interpretation of accretion/deletion is restored and the accreting/deleting central region is perceived to be in behind the surrounding region. They concluded that in order for accretion/deletion (in its traditional sense) to function as a cue to ground status, it requires additional segmentation cues that enhance the delineation of the border at which accretion/deletion occurs. Since our goal is to bridge the gap between the traditional cases of accretion/deletion and the displays that causes rotating-in-front interpretation, in this study we use luminance-defined borders to make sure that accretion/deletion indicates the background surface when the traditional display is used. As a result, if the accreting/deleting region is perceived to be in front, then we would ensure that it is due to the rotating-in-front interpretation of accretion/deletion, and not due to lack of segmentation cues.

3.2 Experiment 1

In Experiment 1, we investigated which factor(s) is (are) critical in promoting the rotating-in-front interpretation of accretion/deletion. We manipulated the variables that make the displays used by the recent studies on accretion/deletion (e.g. Froyen et al., 2013; Tanrikulu et al., 2016) different than the traditional accretion/deletion displays. Traditional accretion/deletion displays generally include two regions separated by a straight border in which accretion/deletion is introduced to one of the regions while the texture on the other region is static. However, in the more recent studies that focus on the rotating-in-front interpretation of accretion/deletion, the displays contain multi-region figure/ground stimulus in which the borders include some degree of curvature and/or geometric cues to figure/ground. Moreover, accretion/deletion can be present on both sides of each border, or only on one side of each border. Therefore, in this experiment we independently manipulated the number of regions in the display, the shape of the borders at which the texture accretes/deletes, and whether accretion/deletion is introduced to one side or both sides of each border. In a forced choice task, subjects indicated which of the two target regions they perceived to be in front of the other one.

3.2.1 Method

Participants

Seven Rutgers University students who were naive to the purpose of the experiment participated in the study. All had normal or corrected-to-normal visual acuity. They were paid for their participation.

Stimuli

The stimulus consisted of alternating dark and light vertical regions presented at the center of the screen surrounded by a gray background (Figure 3.4). The height of the stimulus was 6.1°. The width of the stimulus ranged from 1.2° to 8.6° depending on the number of regions presented in that experimental condition. The colors of the odd and the even regions were counterbalanced and crossed with other factors so that on half of the trials the odd regions were light and the even regions were dark colored (vice versa for the other half). The phase of the stimulus was also counterbalanced and crossed with other factors such that in half of the trials the displays were shown reflected over their vertical axis.

In terms of their geometric properties, three different types of boundaries were used as borders to separate the regions. The boundary could be a straight line, or one that includes curvature, or composed of convex parts. In this study, we will use the term "convex" to refer to a boundary that is piecewise convex such that the boundary can be segmented into parts, in which each individual part is convex. The boundaries of these parts are defined by the negative minima of curvature (Hoffman & Singh, 1997). The type of boundary that includes curvature was made such that it is unbiased in terms of geometrical cues to figure/ground. This "unbiased" boundary was composed of parts that are taken from sinusoids. Both the convex and the unbiased boundaries consisted of five and a half parts, where the width and the height (i.e. frequency and amplitude) of each part were randomized separately within a certain range. The phase of each boundary (i.e. whether it starts with a negative minima or a positive maximum) was also randomized. We created the unbiased boundaries using cubic spline interpolation with 14 control points. In order to create the convex boundary, we made the negative minima points of unbiased boundaries sharp on one side by shifting the positions of the control points of the spline curve (Figure 3.3). Each boundary was generated individually and then placed on the stimulus such that the area between each boundary was the same.

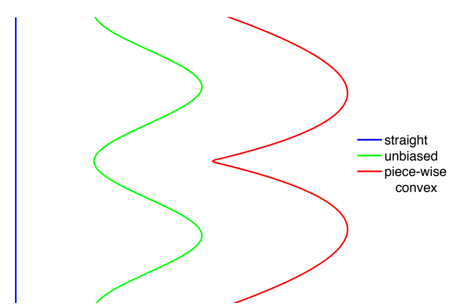

Figure 3.3: The three types of boundaries used in generating the multi-region figure/ground displays. The boundary in the middle is "unbiased" in terms of geometric figure/ground cues. The piecewise convex boundary consists of parts that are convex, where the part boundaries are defined by negative minima of curvature. The sharpness of the negative minima is gradually increased in order to obtain the green (weak-convexity) and the red curves (strong-convexity), in which the left side of each curve becomes piecewise convex.

The regions included random-dot texture that could either be static, move horizontally to left or to right at constant speed. The size of a single dot was 1.5 arcmin by 1.5 arcmin. The motion was coherent for both dark and light regions independently. When both sets of regions (i.e. dark and light) had motion, they were always in opposite directions. For the light regions, the random dot texture was sampled from a beta distribution, in which the probability density function of a beta distribution is:

$$Beta(\alpha, \beta) : p(x|\alpha, \beta) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha, \beta)}$$

where $0 \leq x \leq 1$, shape parameters $\alpha, \beta > 0$ and Beta function, $B(\alpha, \beta)$ is used as a normalization constant with parameters $[\alpha = 2, \beta = 6]$, which resulted in a light texture with sparsely scattered dark pixels. The texture on the dark regions was generated in a similar way where the parameters for the Beta function were $[\alpha = 6, \beta = 2]$. To add motion to this random dot texture, in each frame t the texture columns $[2,\ N]$ were taken from the texture columns $[1,\ N-1]$ in frame $t-1$. The luminance values for the first column were resampled from the corresponding beta distribution. In this way, rightward motion was added to the random dot texture in a region. The leftward motion was added in a similar way, in which the texture columns $[1,\ N-1]$ in frame t were taken from the texture columns $[2,\ N]$ in frame $t-1$. This procedure was repeated at a rate of 22 frames/sec, resulting in a speed of 0.5°/sec. The direction of motion was counterbalanced and crossed with other factors such that in half of the trials the dark regions had leftward motion and the light regions had rightward motion (and vice versa in the other half).

Design and Procedure

Subjects sat 76 cm from a 21-inch CRT monitor (85Hz, 1280pxl x 1024pxl) connected to a Windows 7 PC. The stimulus was presented using Psychtoolbox in Matlab (Brainard, 1997; Kleiner et al., 2007). Subjects were asked forced-choice questions to indicate the region that they perceive to be in front among the two target regions on the display.

The three variables that were manipulated in this experiment were the number of regions, the shape of the borders and whether accretion/deletion is introduced on one or both sides of each border. The stimulus could include two, four or eight regions. The shape of the border could be either straight, curved but unbiased in terms of geometric figure/ground cues, or piecewise convex in one set of regions. Accretion/deletion was introduced either on one side of each border (i.e. only the dark or the light regions had accreting/deleting texture) or on both sides of each border (i.e. all the regions had accreting/deleting texture but the textural motion was in opposite directions in the light and the dark regions). It has been shown that when only one side of each border had accreting/deleting texture, depth-order ambiguity especially arises when accretion/deletion is only introduced to the piecewise convex set of regions (compared to introducing it only to the non-convex set of regions) (Tanrikulu et al., 2016). Therefore, in the condition where accretion/deletion was present only on one set of regions and the borders were convex, accretion/deletion was always introduced to the convex set of regions. Different levels of each variable were crossed with each other, resulting in 18 experimental conditions (Figure 3.4).

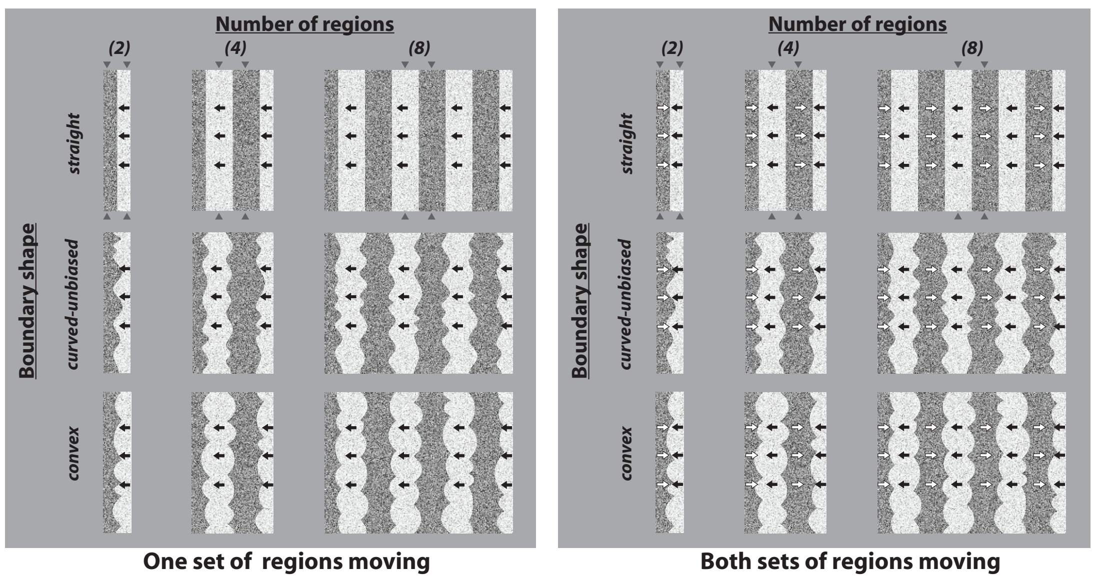

Figure 3.4: The manipulations and single frames of the stimuli from each of the 18 different experimental conditions used in Experiment 1. The arrows on the stimuli were not presented to the subjects. They just indicate the direction of textural motion. The box on the left shows the conditions where only one side of each border had accreting/deleting texture. In those conditions, the textural motion could be introduced either to the dark or to the light regions. In the condition where borders were piecewise convex, the textural motion was always added to the piecewise convex set of regions. The box on the left shows the conditions where both sides of each border had accreting/deleting texture. The triangles appeared above and below of each stimulus as an indication of the two target regions are illustrated at the top row of each box as an example. These triangles appeared for all stimuli in the experiment.

Subjects completed 864 trials in total split into 8 blocks, i.e. 3(number of regions: 2, 4, 8) x 3(border shape: straight, curved, convex) x 2(place of accretion/deletion: one or both sides of borders) x 2(luminance: dark/light) x 2(phase) x 2(direction of motion: left/right) x 6(repetitions). The experiment was completed in two sessions on different days. Each session contained 4 blocks where each block contained 108 trials. All trials were randomized for each subject separately. At the beginning of each session, subjects were given practice trials to get them familiar with the task and the stimulus. Subjects completed 60 practice trials at the beginning of the first session, and 30 practice trials at the beginning of the second session. Each session took approximately 45 minutes to complete.

On each trial, subjects were presented with 800 msec of pre-mask and then 800 msec of pre-mask with a fixation cross added to the center of the mask. The masks were generated by overlaying semitransparent single frames of 10 stimuli randomly chosen among curved-unbiased and convex stimuli that contained eight regions. Following the mask, the experimental display was shown for 2.2 sec. In the last second of the stimulus presentation, two target regions that are adjacent to each other were indicated by black triangles, which appeared at the top and bottom of these two target regions (5 pixels away from the stimulus boundary). The height and the base of a triangle were each 30 arcmin long. Among these two adjacent regions, subjects were asked to indicate the region that they perceived to be in front by using the arrow keys (i.e. left or right). When the stimulus contained two regions, then these two regions were selected as targets. With four regions display, the two full regions in the center were selected as the target regions. For the displays with eight regions, the target regions were determined randomly at the beginning of each trial. (The two half regions at the leftmost and the rightmost of the stimulus were never selected as the target regions.) Since the target regions had to be adjacent to each other, this left us with five different possible pairs of target regions for the eight regions display. Finally, after the stimulus disappeared, a post-mask which is identical to the pre-mask was shown and the subjects pressed the corresponding key to indicate the region that they perceive to be in front. The post-mask stayed on the screen until subject responded. Once the response was made, the next trial started immediately. The subjects were instructed to respond with their initial percept of relative depth in order to avoid any influence of figure/ground reversals (including top-down effects) on the results.

3.2.2 Results

Figure 3.5A presents the results of the condition in which only one side of each border had textural motion, whereas Figure 3.5B presents the results for the condition where both sides had textural motion. Figure 3.5A shows the proportion of trials the regions that had textural motion (i.e. accretion/deletion) were perceived to be in front. For the condition where both sets of regions had textural motion, a reference set of regions was chosen and the bars in Figure 3.5B show the proportion of trials the reference set of regions was perceived to be in front. In the condition where borders were piecewise convex, the convex set of regions were chosen to be the reference set of regions. However, in the conditions where the borders were either straight or unbiased, the two sets of regions become equivalent in terms of both geometric and dynamic figure/ground cues. Therefore the reference set of regions become arbitrary in those conditions. That is why the bars in Figure 3.5B that correspond to the straight and unbiased border conditions (i.e. red and green bars) are around 0.5 (i.e. chance level). This also suggests that curvature by itself does not biases figure/ground interpretation in one direction.

A logistic regression was performed for each subject separately. A likelihood ratio test revealed that including the main effect of place of accretion/deletion was a significant improvement over an unconditional-means model (i.e. containing only an intercept) for each of the seven subjects. ($LR_{max} = 467.65$, $LR_{min} = 46.69$, $df = 1$, $p < 0.001$, see Table 6.1 in Appendix A for more details.) The effect of place of accretion/deletion can be seen when Figure 3.5A and Figure 3.5B are compared to each other. Proportion of seeing an accreting/deleting region in front increases significantly when accretion/deletion is introduced to both sides of each border.

In the following, we are going to analyze the data first for the condition in which only one side of each border had accretion/deletion, and then separately for the condition in which both sides of each border had accretion/deletion.

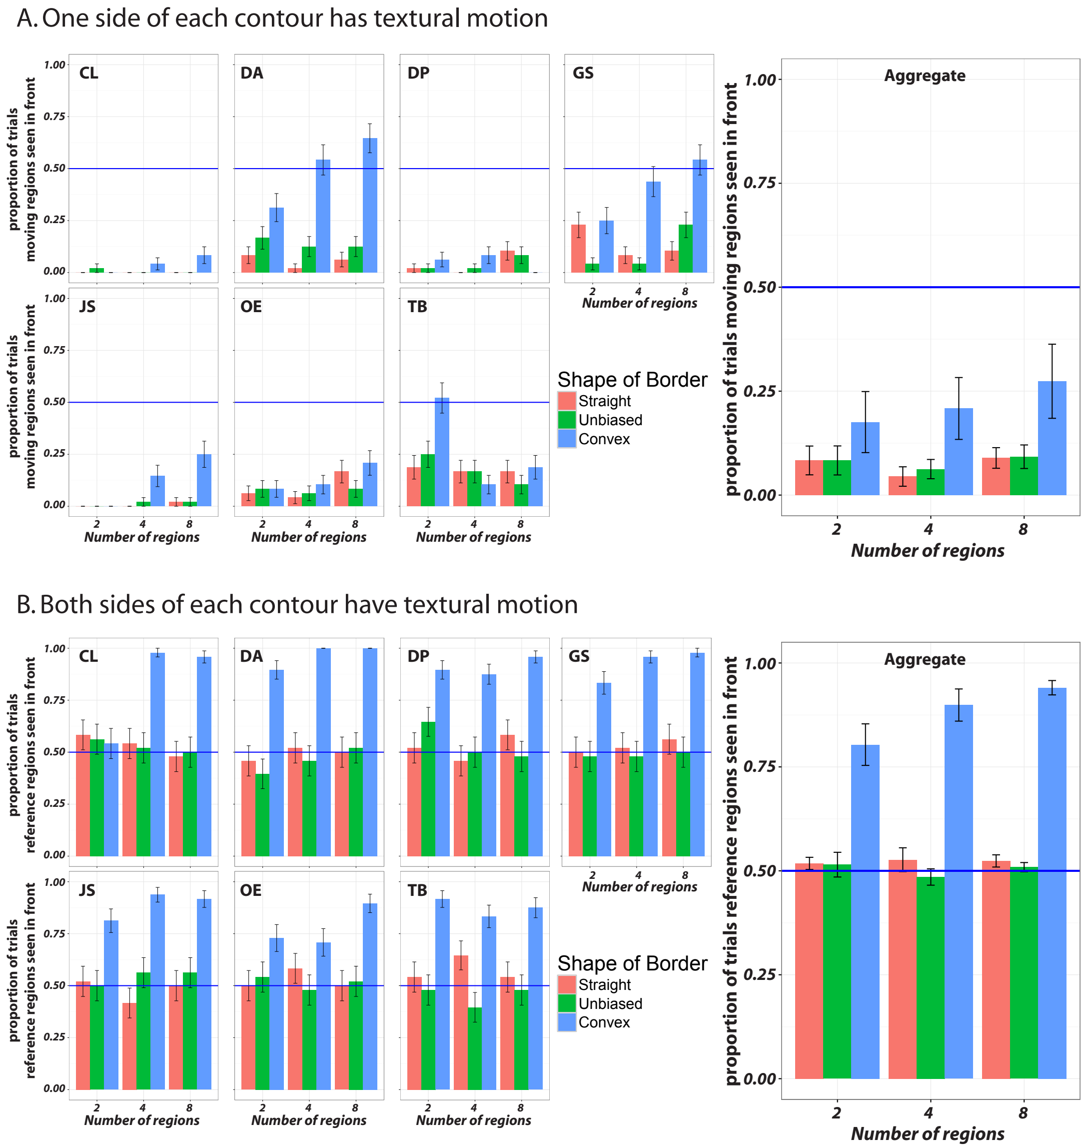

Figure 3.5: Results of Experiment 1. Error-bars represent $\pm 1SE$. The blue line shows the chance level, i.e. where the proportion equals to 0.5. **A.** Proportion of trials the accreting/deleting set of regions perceived to be in front is shown for the condition where only one side of each border had textural motion. **B.** Proportion of trials the reference set of regions perceived to be in front is shown for the condition where both sides of each border had textural motion. When the borders were piecewise convex, the reference region was chosen to be the convex set of regions. In the conditions where the borders were either straight or unbiased, the two sets of regions becomes identical in terms of geometric and dynamic figure/ground cues. In that case, the reference region becomes arbitrary.

Accretion/deletion only on one side of each border

A likelihood ratio test conducted on the logistic regression of each subject's data revealed that the main effect of border geometry was a significant improvement over an unconditional-means model that contains only an intercept for four subjects (i.e. subjects CL, DA, GS and JS) out of seven. ($LR_{max} = 90.01$, $LR_{min} = 9.79$, $df = 2$, $p < 0.01$, see Table 6.2 in Appendix A for more details.) A Tukey test revealed that for all these four subjects the proportion of trials the accreting/deleting regions perceived to be in front was significantly higher in the piecewise convex condition (i.e. blue bars in Figure 3.5A), compared to the straight and unbiased borders (i.e. red and green bars respectively in Figure 3.5A). (When aggregated over seven subjects: straight:$M = 0.072$,$SE = 0.025$; unbiased: $M = 0.079$, $SE = 0.024$; convex: $M = 0.219$, $SE = 0.068$)

The addition of the number of regions factor to these individual logistic models that include the effect of border geometry was found to be a significant addition for four subjects (i.e. subjects GS, JS, OE and TB) out of seven($LR_{max} = 21.27$, $LR_{min} = 6.52$, $df = 2$, $p < 0.05$). When a Tukey test is performed, we found that for three of these four subjects (GS, JS and OE) higher number of regions yielded higher proportions of trials in which accreting/deleting regions were perceived to be in front. However, for one subject (TB) that proportion was significantly higher in the two-regions condition, compared to the four- and eight-regions conditions. (When aggregated over seven subjects: two-regions: $M = 0.114$, $SE = 0.045$; four-regions: $M = 0.105$, $SE = 0.031$; eight-regions: $M = 0.152$, $SE = 0.038$)

When the color factor was also added to these individual logistic models, it resulted in a significant improvement for five subjects (CL, DA, GS, ,JS and TB) out of seven ($LR_{max} = 177.14$, $LR_{min} = 4.10$, $df = 12$, $p < 0.05$). All these five subjects were more likely to perceive the light regions in front compared to the dark regions. (When aggregated over seven subjects: dark: $M = 0.067$, $SE = 0.025$; light: $M = 0.180$, $SE = 0.057$)

For four subjects (DA, DP, GS and TB), adding the interaction factor between border geometry and number of regions to these individual models yielded a significant improvement ($LR_{max} = 17.26$, $LR_{min} = 13.94$, $df = 4$, $p < 0.05$). Among these four subjects, for two of them (DA and GS) the effect of convexity becomes stronger when the number of regions was higher. In other words, these subjects perceived the convex regions to be in front especially when the stimulus had four or eight regions. However, for the other two subjects (DP and TB) it was the opposite; the effect of convexity disappeared (or reversed) as the number of regions increased. (The factors included in the final statistical models of each subject can be found in Table 6.2 in Appendix A.)

Accretion/deletion on both sides of each border

The same logistic regression was applied to each subject's responses for the condition in which both sides of each border had accreting/deleting texture. A likelihood ratio test revealed that border geometry was a significant improvement over an unconditional-means model that includes only an intercept for all seven subjects. ($LR_{max} = 123.45$, $LR_{min} = 27.88$, $df = 2$, $p < 0.001$, see Table 6.3 in Appendix A for more details.) A Tukey test revealed that for all seven subjects, proportion of seeing the reference region in front was significantly higher when the reference region was piecewise convex, compared to the other two conditions (i.e. straight and unbiased). This can be seen in Figure 3.5B by comparing the blue bars (i.e. convex borders) to the red (straight borders) and the green (unbiased borders) bars. (When aggregated over seven subjects: straight:$M = 0.523$,$SE = 0.012$; unbiased: $M = 0.503$, $SE = 0.014$; convex: $M = 0.881$, $SE = 0.024$)

However, addition of the number of regions factor to these individual models did not yield a significant expansion over border geometry for all seven subjects. For all subjects except one (JS), addition of the color factor resulted in a significant improvement. ($LR_{max} = 168.9$, $LR_{min} = 5.02$, $df = 1$, $p < 0.05$). Among these six subjects whose data yielded a significant effect of color, four of them (CL, DA, DP and TB) showed a bias for dark, whereas two of them (GS and OE) showed a bias for light regions. (Aggregated over seven subjects: dark:$M = 0.641$, $SE = 0.056$; light: $M = 0.630$, $SE = 0.058$)

For two subjects (CL and DA), adding the interaction term for border geometry and number of regions to their individual models also resulted in a significant improvement ($LR_{max} = 41.83$, $LR_{min} = 13.25$, $df = 4$, $p < 0.05$). For both subjects, the effect of convexity became very strong (e.g. ceiling effect) when the stimulus contained four or eight regions, but not when it had two regions. For all subjects except one (DP), adding the interaction term between color and border geometry to their individual models yielded a significant expansion of their models ($LR_{max} = 10.04$, $LR_{min} = 6.89$, $df = 2$, $p < 0.05$). Further analysis showed that for all these six subjects the effect of color was observed when the border geometry was straight and/or unbiased, but not when it was piecewise convex. (The factors included in the final statistical models of each subject can be found in Table 6.3 in Appendix A.)

3.2.3 Discussion

Among the three factors manipulated in Experiment 1, the place of accretion/deletion and convexity of the borders seem to be the critical factors that give rise to the rotating-in-front interpretation of accretion/deletion. In the condition where accretion/deletion was introduced to only one side of each border, the reference regions were the set of regions that include accreting/deleting texture. As it is seen from the aggregate graph in Figure 3.5A, the average proportion of seeing these regions in front is about 0.25. However, when accretion/deletion was also added to the other set of regions the average proportion of seeing those reference regions in front increases dramatically to about 0.6 (as seen in the aggregate graph in Figure 3.5B).

The shape of the border had a significant effect on subjects' responses, in that piecewise convex regions were more likely to be perceived as in front. Whether the borders were made straight or curved did not make a difference. However, when piecewise convexity was introduced to one set of regions, those regions were more likely to be perceived in front compared to the non-convex set of regions. The effect of convexity became stronger when accretion/deletion was introduced to both sides of each border (i.e. the difference between the blue bars and the other bars is higher in Figure 3.5B than in Figure 3.5A).

Unlike the other two factors, the number of regions did not have a major influence on how accretion/deletion is interpreted, especially in the condition where both sides of each border had accretion/deletion. Only in the condition where one side of each border had accretion/deletion, four subjects' data yielded a significant effect of number of regions. However the direction of this effect was not systematic across subjects. As the number of regions increased, the proportion of seeing the reference (i.e. accreting/deleting) regions in front increased for three of these subjects, while the proportion decreased for the other subject. The interaction between number of regions and border geometry also showed individual differences. For two of the four subjects, whose data yielded a significant interaction between these two factors, the effect of convexity became stronger with increasing number of regions, whereas this effect was weakened with increasing number of regions for the other two subjects.

In the condition where both sides of each border had accretion/deletion, the reference region becomes arbitrary if the borders were unbiased in terms of any geometrical figure/ground cues. Therefore, the proportions obtained from these those conditions would be, by definition, around chance-level. In order to examine the true influence of number of regions when accretion/deletion is introduced on both sides of each border, we need to focus on the condition where the borders were made piecewise convex. When the effect of number of regions were examined for the condition where accretion/deletion is introduced to both sides of each piecewise convex borders (blue bars in Figure 3.5B), it was observed that number of regions had a significant effect for only two of the subjects. For both subjects, the proportion of seeing the piecewise convex regions in front was higher for the four- and eight-region stimuli, compared to the two-region stimulus. However, this still shows that the effect of number of regions were not as influential as the other two factors that were manipulated.

The direction of the effect of color was also not systematic across subjects. The direction of the effect changed depending on the individual and the experimental condition. For example, in the condition where both sides of each border had accretion/deletion, the effect of color was observed when the shape of the borders were unbiased in terms of convexity. This indicates that the effect of color was especially observed when subjects do not have a strong preference in their relative depth judgments.

3.3 Experiment 2

In Experiment 2, we manipulated the same factors as in Experiment 1, but we changed the way we manipulated the number of regions on the multi-region figure/ground stimulus. In Experiment 1, we kept the width of a single region fixed while varying the number of regions on the stimulus. Therefore, in Experiment 1, as the number of regions increased, the width of the whole stimulus also increased (Figure 3.4). However, in Experiment 2, we kept the width of the whole stimulus fixed while varying the number of regions. As a result, the width of a single region decreased as the number of regions on the stimulus increased, however the overall width of the stimulus did not change.

3.3.1 Method

Participants

Seven Rutgers University students who were naive to the purpose of the experiment participated in the study. All had normal or corrected-to-normal visual acuity. They were paid for their participation.

Stimuli and Procedure

The stimuli were generated in the same way as in Experiment 1 except how the number of regions were manipulated. Instead of adding or removing regions that had fixed widths, the width of the stimulus was kept fixed and then the stimulus was divided to either two, four or eight regions. Other manipulations were done in the same manner as in Experiment 1. In half of the trials, accretion/deletion was applied to only one side of each border, whereas in the other half, it was introduced to both sides of each border. The border could be either straight, curved-unbiased, piecewise convex. Different levels of each factor were crossed with each other (Figure 3.6).

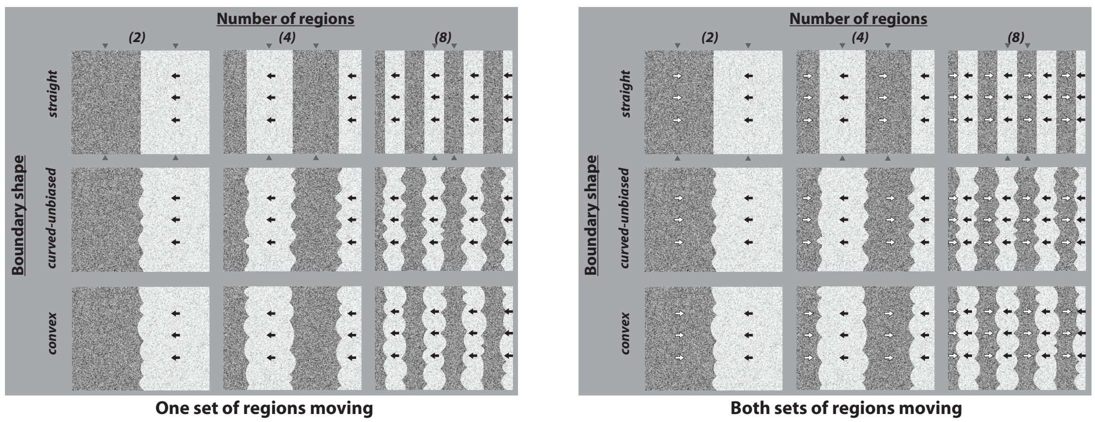

Figure 3.6: The manipulations and single frames of the stimuli from each of the 18 different experimental conditions used in Experiment 2. The arrows on the stimuli were not presented to the subjects. They just indicate the direction of textural motion. The box on the left shows the conditions where only one side of each border had accreting/deleting texture. In those conditions, the accretion/deletion could be introduced either to the dark or to the light regions. In the condition where borders were piecewise convex, the accreting/deleting texture was always added to the piecewise convex set of regions. The box on the right shows the conditions where both sides of each border had accreting/deleting texture. The triangles appeared above and below of each stimulus as an indication of the two target regions are illustrated at the top row of each box as an example. These triangles appeared for all stimuli in the experiment.

The experimental procedure was the same as in Experiment 1. Each subject completed 864 trials in two different sessions. Each session took approximately 45 minutes to complete.

3.3.2 Results

Figure 3.7A shows subjects' responses from the condition where only one side of each border had textural motion. Figure 3.7B shows responses from the condition where both sides of each border had textural motion. In this condition, when the borders were piecewise convex, the reference region was chosen to be the convex set of regions. In the conditions where the borders were either straight or curved-unbiased, the two sets of regions becomes identical in terms of geometric and dynamic figure/ground cues. In that case, the reference region becomes arbitrary. That is why the red and green bars in Figure 3.7B are generally around chance level.

A logistic regression was performed for each subject separately. Likelihood ratio tests revealed that the main effect of place of accretion/deletion was a significant improvement over an unconditional-means model (i.e. containing only an intercept) for each of the seven subjects. ($LR_{max} = 487.06$, $LR_{min} = 157.79$, $df = 1$, $p < 0.001$, see Table 6.4 in Appendix B for more details.) When Figure 3.7A and Figure 3.7B are compared, it can be seen that proportion of seeing the reference region in front increases significantly when accretion/deletion is introduced to both sides of each border, compared to the condition where only one side of each border had accretion/deletion.

In the following, we are going to analyze the data first for the condition in which only one side of each border had accretion/deletion, and then separately for the condition in which both sides of each border had accretion/deletion.

Accretion/deletion only on one side of each border

Likelihood ratio tests performed for each subject's logistic regression revealed that the main effect of border geometry was a significant improvement over an unconditional means model (i.e. contains only an intercept) for all subjects except one (subject DP). ($LR_{max} = 78.745$, $LR_{min} = 15.11$, $df = 2$, $p < 0.001$, see Table 6.5 in Appendix B for

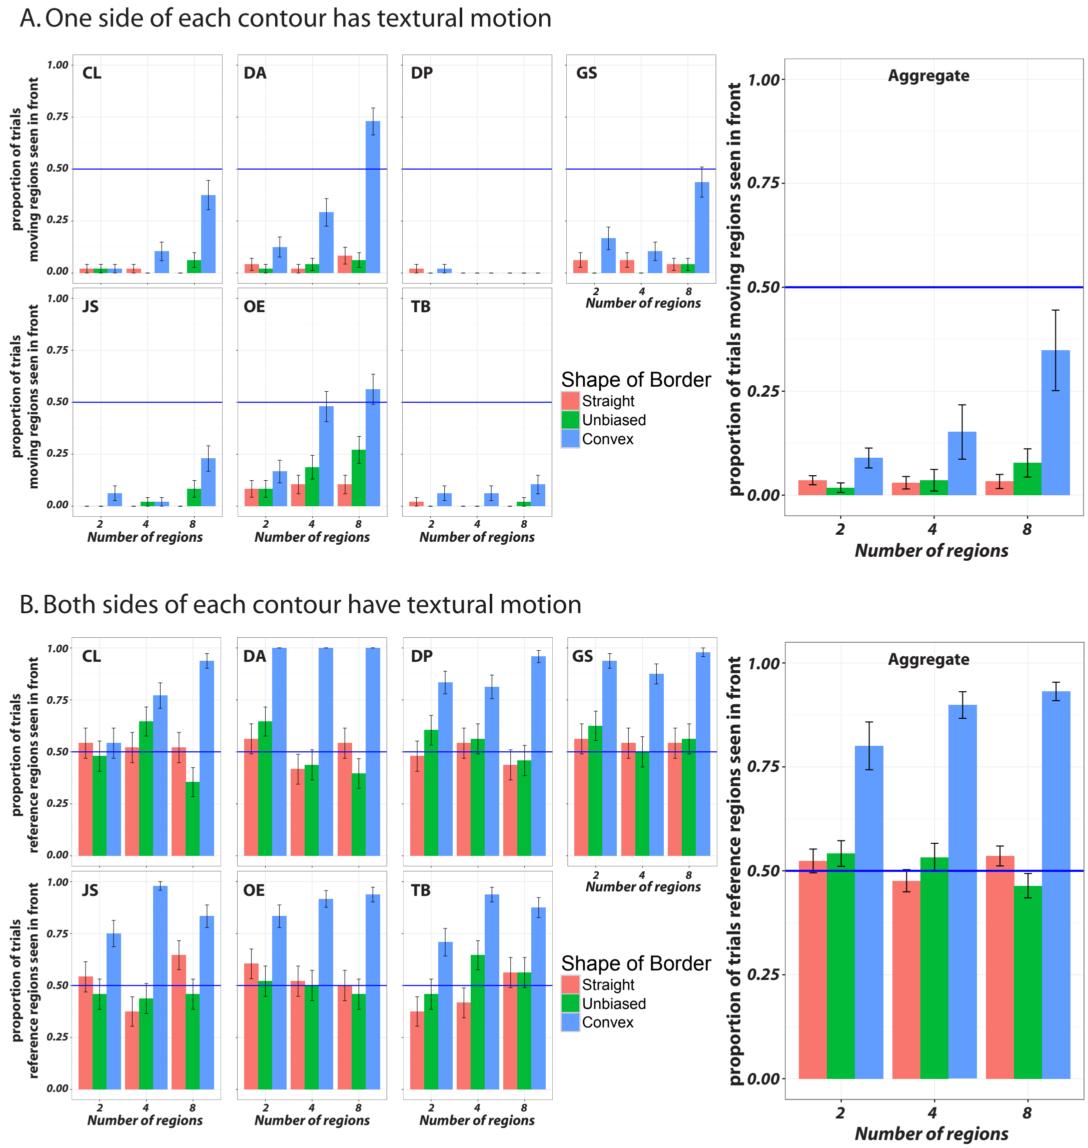

Figure 3.7: Results of Experiment 2. Error-bars represent $\pm 1SE$. The blue line shows the chance level, i.e. where the proportion equals to 0.5. **A.** Proportion of trials the accreting/deleting region perceived to be in front is shown for the condition where only one side of each border had textural motion. **B.** Proportion of trials the reference region perceived to be in front is shown for the condition where both sides of each border had textural motion. When the borders were piecewise convex, the reference regions were chosen to be the convex set of regions. In the conditions where the borders were either straight or curved-unbiased, the two sets of regions becomes identical in terms of geometric and dynamic figure/ground cues. In that case, the reference region becomes arbitrary.

more details.) For these six subjects, Tukey tests revealed that proportion of trials the accreting/deleting regions perceived in front was significantly higher when the borders were piecewise convex (i.e. blue bars in Figure 3.7A), compared to the condition where the border were either straight or curved-unbiased (i.e. red and green bars respectively in Figure 3.7A). (When aggregated over seven subjects: straight:$M = 0.033$,$SE = 0.014$; unbiased: $M = 0.044$, $SE = 0.024$; convex: $M = 0.196$, $SE = 0.057$.)

When the main effect of number of regions was added to the models that include the effect of border geometry, it was found to be a significant addition for five of the subjects (i.e. subjects CL, DA, GS, JS and OE) ($LR_{max} = 38.717$, $LR_{min} = 13.408$, $df = 1$, $p < 0.01$). Tukey tests revealed that for all five subjects the proportion of trials the accreting/deleting regions seen in front was significantly higher when the stimulus contained eight regions, compared to the condition where it contained only two. (When aggregated over seven subjects: two-regions:$M = 0.048$, $SE = 0.014$; four-regions: $M = 0.072$, $SE = 0.034$; eight-regions: $M = 0.153$, $SE = 0.045$.)

Addition of the color factor to these individual models yielded a significant improvement for all subjects except one (subject DP) ($LR_{max} = 57.632$, $LR_{min} = 10.554$, $df = 1$, $p < 0.001$). All the six subjects were more likely to see the light regions in front compared to the dark regions. (When aggregated over seven subjects: dark: $M = 0.044$, $SE = 0.023$; light: $M = 0.138$, $SE = 0.038$.)

Adding the interaction term between number of regions and border geometry to these models yielded a significant improvement only for one subject (subject CL) ($LR = 9.574$, $df = 4$, $p < 0.05$). For this subject, the effect of convexity increases as the number of regions increases. For two of the subjects (subjects JS and TB) addition of the interaction term between border geometry and color also yielded a significant improvement ($LR_{max} = 7.388$, $LR_{min} = 7.316$, $df = 2$, $p < 0.05$). Both of the subjects had light color bias when the borders of the accreting/deleting regions were made piecewise convex. (The factors included in the final statistical models of each subject can be found in Table 6.5 in Appendix A.)

Accretion/deletion on both sides of each border

The same individual logistic regression analyses were also performed for the condition in which both sides of each border had accreting/deleting texture. Likelihood ratio tests showed that the main effect of border geometry was a significant improvement over an unconditional means model (i..e contains only an intercept) for all subjects. ($LR_{max} = 150.75$, $LR_{min} = 23.985$, $df = 2$, $p < 0.001$, see Table 6.6 in Appendix B for more details.) Tukey tests revealed that for all seven subjects, the proportion of trials the reference region perceived in front was significantly higher when the borders were piecewise convex (i.e. blue bars in Figure 3.7B), compared to when the borders were either straight or curved-unbiased (i.e. red and green bars respectively in Figure 3.7B). (When aggregated over seven subjects: straight:$M = 0.512$, $SE = 0.013$; unbiased: $M = 0.513$, $SE = 0.015$; convex: $M = 0.877$, $SE = 0.029$.)

Addition of the main effect of number of regions to these individual models was a significant improvement over border geometry only for one subject (subject TB). ($LR = 63.673$, $df = 1$, $p < 0.001$) A Tukey test revealed that this subject were more likely to perceive the reference region in front when the stimulus had four or eight regions, compared to when it had two regions. (When aggregated over seven subjects: two-regions:$M = 0.622$, $SE = 0.033$; four-regions: $M = 0.636$, $SE = 0.008$; eight-regions: $M = 0.644$, $SE = 0.011$.)

Addition of the interaction term between number of regions and border geometry to these models was a significant improvement only for two subjects (CL and JS) ($LR_{max} = 30.024$, $LR_{min} = 19.995$, $df = 4$, $p < 0.01$). For one of these subjects (CL), there was a significant effect of number of regions when the borders were piecewise convex. When the stimulus had two regions, subject's responses were around chance-level regardless of the shape of the borders. When the number of regions was increased, the subject was more likely to perceive the convex regions in front compared to when the borders were straight or curved-unbiased. The effect of convexity increase with increasing number of regions for this subject. For the other subject (TB), the effect of convexity was stronger when the stimulus had four or eight regions compared to when it had only two regions.

When the main effect of color added to these models, it was a significant improvement for five (i.e. CL, DA, DP, GS and TB) of the seven subjects (LR_{max} = 125.68, LR_{min} = 45.943, df = 1, $p < 0.001$). While three of these five subjects were more likely to see the reference region in front when its texture was dark, the other two subjects had light texture bias. (When aggregated over seven subjects: dark: $M = 0.671$, $SE = 0.058$; light: $M = 0.597$, $SE = 0.065$.) The addition of the interaction term between color and number of region and then the interaction term between color and border geometry were both significant improvements for six of the seven subjects (LR_{max} = 24.946, LR_{min} = 6.898, df = 2, $p < 0.05$) Further analysis revealed that a color bias was observed mainly when the borders were either straight or curved-unbiased but not when they were piecewise convex. However, the interaction between number of regions and color did not show a specific trend. For four of these subjects, color bias increased with increasing number of regions, whereas for the other two subjects, the trend was in the opposite direction. (The factors included in the final statistical models of each subject can be found in Table 6.6 in Appendix A.)

3.3.3 Discussion

As in Experiment 1, the results of Experiment 2 showed that the place of accretion/deletion and convexity of the borders are the most critical factors that contribute to the rotating-in-front interpretation of accretion/deletion. Individual differences were observed for the effect of number of regions. It was a critical factor for some of the individuals depending on whether accretion/deletion is on one or both sides of each border.

The place of accretion/deletion had a significant effect on the relative depth judgments for all the subjects. This effect is clearly visible when the overall height of the bars in Figure 3.7A is compared to the height of the bars in Figure 3.7B. The average proportion of seeing the reference region in front is six times higher in the condition where both sides of each border had accretion/deletion.

The shape of the border had also a significant effect on all subjects' responses.

Subjects were more likely to perceive the piecewise convex regions in front. Whether the border was straight or curved-unbiased did not make a difference in subjects' responses. However, when one side of each border was made piecewise convex, the average proportion of seeing the reference region in front was two times higher than the condition where the border was unbiased (either straight or curved) in terms of convexity. The only exception to this was the responses of subject DP in the condition where only one side of each border had accretion/deletion. This subject perceived the reference (i.e. accreting/deleting) region almost always in behind in this condition regardless of border geometry.

The number of regions was also a critical factor in the condition where only one side of each border had accretion/deletion. However, this effect weakened when accretion/deletion was introduced to both sides of each border. In the former condition, the proportion of trials the reference region seen in front increased three times when the stimulus had eight regions, compared to the two-region stimulus. In this condition, only two subjects didn't show a significant effect of number of regions. However, a floor effect can be seen in the responses of these two subjects (subjects DP and TB in Figure 3.7A). A significant effect of number of regions was observed only for one subject in the condition where both sides of each border had accretion/deletion. However, in order to examine the actual effect of number of regions in this condition we need to look at the responses only in the piecewise convex condition (i.e. only the blue bars in Figure 3.7B), since in the other two border geometry conditions (straight and curved-unbiased) responses would be around chance-level by definition. When we examined the effect of number of regions when the border was piecewise convex, we observed that responses of two more subjects (three subjects in total) show a significant effect of number of regions. However, the number of regions still seemed to be more influential in the condition where only one side of each border had accretion/deletion.

The color of the reference region had again a significant influence on most subjects' responses. However, similar to what we have seen in Experiment 1, the direction of the effect of the color factor and its interaction with other factors was not systematic among subjects.

While all three factors manipulated in Experiment 2 seem to have an effect on how accretion/deletion is interpreted, having accretion/deletion on both sides of each border and making the reference region piecewise convex were the two most critical factors that give rise to the rotating-in-front interpretation of accretion/deletion. Individual differences were observed in the effect of number of regions. While having more regions on the stimulus tend to increase the probability of seeing the reference region in front, this effect weakens when accretion/deletion is introduced to both sides of each border. The effect of number of regions seems to be overshadowed by the strong influence of the other two factors. In other words, having accretion/deletion on both sides of each piecewise convex border already increases the proportion of trials the reference region seen in front significantly, such that the effect of number of regions is overshadowed (e.g. ceiling effect observed for the subject DA in the condition where border was piecewise convex).

3.4 General Discussion

Traditionally, accretion/deletion is considered as a definite cue to ground status. However, recent studies (Froyen et al., 2013; Tanrikulu et al., 2016) challenged this traditional view and demonstrated that accretion/deletion can also be interpreted as a surface in front, where accretion/deletion arises due to rotation in depth. In light of our results, it seems clear that this difference is driven mostly by the presence of accretion/deletion on both sides of the border and the geometry of the occluding contour.

In traditional displays of accretion/deletion, a straight line is generally used as the border at which texture accretes/deletes. In contrast, in the studies that focus on the rotating-in-front interpretation of accretion/deletion, the shape of the border generally includes a curved contour, and in some cases includes geometrical figure/ground cues (e.g. convexity) introduced to one side of the border.

Our results indicate that whether the border is straight or curved does not by itself influence how accretion/deletion is interpreted, unless the border is shaped in such a way as to induce one side to be perceived as figural (e.g. if one side is convex) in which case that side is likely to be perceived as in front and rotating. This effect of convexity was observed across all levels of the other factors that were being manipulated. This indicates that one of the critical factors that creates the discrepancy between the results obtained by traditional studies on accretion/deletion and by the more recent studies is the shape of the border used in those studies. The shape of the border is not generally taken into account in traditional accounts of accretion/deletion. However, our study shows that the shape of the border (i.e. whether the border includes a geometrical figure/ground cue) has a significant influence on whether accreting/deleting region would be interpreted as being in behind or in front.

In traditional displays, accretion/deletion is generally introduced only on one side of a border. However, more recent studies focused also on cases where accretion/deletion is introduced on both sides of a border. Our results demonstrated that an accreting/deleting region is more likely to be interpreted as rotating-in-front when the adjacent region on the other side of the border also has accreting/deleting texture. This effect was also found to be consistent in all levels of the other factors we manipulated. Introducing accretion/deletion on both sides of a border creates an ambiguity in terms of depth order. Since the texture on both sides of the border accretes/deletes, accretion/deletion alone can not determine the the depth order. Other factors, such as a geometric figure/ground cue on one side of the border, resolve this ambiguity. Accreting/deleting texture on the region that is perceived in behind is interpreted as being occluded by the adjacent figural region, which corresponds to the classical interpretation of accretion/deletion. However, since accreting/deleting texture on the figural region can not be attributed to being occluded by its adjacent region, it is interpreted as self-occlusion due to rotation in depth. While rotating-in-front interpretation of accretion/deletion is also observed when accretion/deletion is introduced only one side of a border, introducing it to both sides of the border strongly promotes the rotating-in-front interpretation of accretion/deletion. This is because a situation involving two adjacent accreting/deleting region is rarely perceived as being at the same distance, since this would be a highly non-generic condition. Therefore one of the regions has to be perceived in front of the other, and the accreting/deleting texture on the region that is perceived to be in front is interpreted as rotating column.

The eight-region stimulus used in this study is very similar to the stimulus used in Tanrikulu et al. (2016), who demonstrated that when accretion/deletion was introduced only to the piecewise convex set of regions, the relative depth judgments were around chance level. However, in the current study, when accretion/deletion is only introduced to the piecewise convex set of region, these region were perceived in front only around 30% of the time. The most likely reason for this difference between the results of the two studies is that the piecewise convex regions in Tanrikulu et al. (2016) were also made symmetric, which is another geometric cue for figural status. This additional geometric cue further biased people to see the convex set of regions in front. In the current experiments, all the regions were made asymmetric. In addition to that, Tanrikulu et al. (2016) used a method to create piecewise convex boundaries that is different than the method used here. The boundaries in that experiment had a larger number of convex parts which is a factor that is shown to influence the strength of geometric figure/ground cues (Froyen et al., 2013).

Traditional studies on accretion/deletion generally use displays with two regions separated by a single contour. However, the recent studies focusing on the rotating-in-front interpretation of accretion/deletion use displays that include eight regions. We manipulated the number of regions (two, four and eight) in our displays to investigate its effect of rotating-in-front interpretation of accretion/deletion. However, compared to our other manipulations, the influence of number of regions on accretion/deletion was minor. In Experiment 1, when there is accretion/deletion on one side of each border, the effect of number of regions was largely modulated by individual differences. The effect went completely away when accretion/deletion was introduced on both sides of each border. In Experiment 2, the effect of number of regions were more consistent and visible, but only when one side of each border had accreting/deleting texture. The difference in the effect of number of regions in Experiment 1 and Experiment 2 can also be seen when the graph of the aggregate data in Figure 3.5A and graph of the aggregate data in Figure 3.7A are compared. The proportions obtained from the stimulus that included eight regions were exactly the same in Experiment 1 and 2. However, when the stimulus had two or four regions, the proportion of trials the accreting/deleting region was perceived in front was lower in Experiment 2, compared to the proportions in the corresponding conditions of Experiment 1. Moreover, the difference between the results of the two experiments was higher in the two-region conditions, compared to the difference in the four-region conditions. In other words, the proportion of seeing accreting/deleting region in front clearly decreases as the number of regions decreases in Experiment 2, but not that much in Experiment 1.

The difference between the effect of number of region in Experiment 1 and Experiment 2 is most likely due to the changes in the width of a single region in Experiment 2, which did not happen in Experiment 1. The accreting/deleting texture on the stimulus had a constant speed profile, which is more consistent with a translating flat surface. However, as our experiments show, the accreting/deleting texture can also be interpreted as a 3D figure rotating in depth in spite of that constant speed profile. In other words, when the accreting/deleting region is perceived to be in front. the motion of the accreting/deleting texture is interpreted as having a cosine speed profile, which is consistent with 3D rotation in depth. However, when the width of an accreting/deleting region increases, the discrepancy between the constant speed profile of the accreting/deleting region and the cosine speed profile increases. Due to this increased discrepancy, when the width of an accreting/deleting region is large, it becomes easier for the subject to detect that the accreting/deleting region had actually a constant speed profile, and not a cosine speed profile. Since constant speed profile is more consistent with a translating flat surface (compared to a 3D surface rotating in depth), the accreting/deleting region is more likely to be interpreted as a flat surface being occluded by the adjacent region. Hence, increasing the width of an accreting/deleting region makes it more likely to be interpreted as being in behind. Such an account would explain why the number of regions manipulation was more effective in Experiment 2, compared to Experiment 1. This shows that the effect in Experiment 2 is most likely due to the changes in width of a single region, rather than the number of regions on the stimulus.

Peterson and Salvagio (2008) demonstrated that the effect of convexity as a figural cue increases as the number of regions is increased in the display. They observed this effect regardless of the method they used in manipulating the number of regions (i.e. whether the width of a single region or the width of the whole stimulus was kept fixed as the number of region was varied). In contrast to their results, the effect of number of regions on convexity observed in our study was not as clear as the effect demonstrated by Peterson and Salvagio (2008). The displays and the number of region manipulations they used were similar to the ones we used in this study. The only main difference was the textural motion introduced to certain regions in our displays, which, in turn, seems to be responsible for the discrepancy between their and our results. In Experiment 2, the effect of number of regions were confounded by the varying width of a single region. While it might be true that the width of a single region by itself does not influence the effect of number of regions on the strength of convexity cue (as demonstrated by Peterson and Salvagio (2008)), it would effect how accretion/deletion is going to be interpreted (as explained in the previous paragraph). In Experiment 1, the effect of number of regions was largely modulated by individual differences, which might occur when more than one factor (e.g. convexity and accretion/deletion) are combined to yield a relative depth judgment. For example, Hildreth and Royden (2011) observed high individual differences when people combine binocular-disparity and accretion/deletion in a depth order task. Their claim was that the individual differences were due to the different weights assigned to accretion/deletion by different subjects when combining it with binocular disparity information in making relative-depth judgments.

High individual differences that were observed in the condition where only one side of each border had accretion/deletion (Figure 3.5A and Figure 3.7A) indicate possible problems for the traditional view of accretion/deletion. The accretion/deletion stimulus in that condition is standardly considered as unambiguous by traditional accretion/deletion accounts (i.e. one side of the border had static texture while one the other side the texture is being accreted/deleted). These high individual differences might indicate the ambiguity inherent to accretion/deletion as a cue to relative depth (i.e. accretion/deletion can be interpreted as being in behind or in front). While these individual differences could also be due to assigning different weights to different cues in making relative depth judgments, our experiments also included conditions where the shape of the borders were made unbiased in terms of any geometrical figure/ground cues (i.e. straight and curved-unbiased contours). In those cases, accretion/deletion cue was not being combined with another depth-order cue, however high individual differences on relative depth judgments were still observed on those conditions, which indicates that accretion/deletion is not an unambiguous cue to relative depth as proposed by traditional accounts.

In two experiments, we attempted to identify the critical factors that promote the rotating-in-front interpretation of accretion/deletion. Our results showed that the geometry of the border and whether accretion/deletion is introduced to one or both sides of a border are the two critical factors that have significant effects on how accretion/deletion is interpreted. These results indicate that accretion/deletion is not a definite cue to ground status, and can be interpreted as being in front or in behind depending on the shape of the border and on which regions accretion/deletion is introduced. Traditional accounts and even current computational models of accretion/deletion do not consider the geometry of the border as a key component. However, our results indicate that the geometry of the border should be taken into account whenever accretion/deletion is used in determining the relative depth of surfaces.

4. Combination of speed profile of accreting/deleting texture and occluding contour geometry in determining relative depth

4.1 Introduction

Segmenting the retinal image into separate regions involves assignment of figure and ground status to these regions. This process requires determining which borders are owned by which regions on the image. The region that owns an adjacent border is given the figural status and shaped by this border, whereas the region on the other side of this border (i.e. background) is perceived to be occluded by the figural region and extend amodally behind it. Many visual cues have been identified to tend to promote figural status, such as area, enclosure (Rubin, 1915/1958), symmetry (Kanizsa & Gerbino, 1976), convexity (Metzger, 1936/2006; Kanizsa & Gerbino, 1976), parallelism (Morinaga, 1941; Metzger, 1936/2006), lower-region (Vecera et al., 2002), axiality and part salience (Hoffman & Singh, 1997; Froyen et al., 2010) (see Wagemans et al. (2012) for a review).

There are also dynamic cues to relative depth, where our visual system processes the continuous transformation of the optic array in order to segment the visual scene into figure and ground. One prominent example of this is accretion/deletion of texture (Gibson, 1966; Gibson et al., 1969; Kaplan, 1969). When a moving texture appears (i.e. accretes) or disappears (i.e. deletes) at a boundary, it is perceived as if the textured surface is appearing or disappearing from behind an occluding surface that is on the other side of that boundary. In other words, accretion/deletion of texture is generally viewed as identifying the surface that is being occlud*ed* (i.e. background), and the adjacent region on the other side of the boundary becomes the occlud*ing* surface (i.e. figural).

The prominence of accretion/deletion is illustrated by developmental studies indicating that accretion/deletion is one of the first visual cues to which infants respond in order to perceive shape and relative depth (Granrud et al., 1984; Johnson & Mason, 2002). It has been shown that, in certain cases, accretion/deletion can override depth information from lower-region cue (Royden et al., 1988), motion parallax (Ono et al., 1988; Hildreth & Royden, 2011; Yoonessi & Baker, 2013), and binocular disparity (Royden et al., 1988; Hildreth & Royden, 2011). It can resolve ambiguities in some of the bi-stable motion displays, such as direction of rotation ambiguity created when viewing parallel projection of an object rotating in depth (Braunstein et al., 1982), or the ambiguity in the overall structure from point-light biological motion displays (Profitt et al., 1984). Accretion/deletion has been considered as definite cue to ground status in the literature (Thompson et al., 1985; Mutch & Thompson, 1985; Niyogi, 1995; Howard & Rogers, 2002; Hegdé et al., 2004), and it is also considered as such in computational models of depth from motion (Thompson et al., 1985; Mutch & Thompson, 1985; Berzhanskaya et al., 2007; Beck et al., 2008; Raudies & Neumann, 2010; Barnes & Mingolla, 2013; Layton & Yazdanbakhsh, 2015; Ruda et al., 2015).

In the very first experimental study on depth from accretion/deletion, Kaplan (1969) reported that when there is accreting/deleting texture on both sides of a border, 80% of his subjects reported perceiving some type of rotational motion, as if the accreting/deleting textured surfaces were going around rollers; although he did not further discuss this observation. Some of the later accounts of accretion/deletion (Thompson et al., 1985; Mutch & Thompson, 1985; Yonas et al., 1987) mentioned in passing that accretion/deletion of texture can also result from dynamic self-occlusion where a 3D object rotates in depth (Figure 4.1). In such cases, the accretion/deletion does not necessarily indicate the ground region, since the accretion/deletion of the texture is now due to self-occlusion and not due to being occluded by the adjacent region. In this alternative interpretation, accreting/deleting region can be perceived to be the figural region. However, these traditional accounts of accretion/deletion explicitly noted that their analyses do not apply to cases of accretion/deletion due to rotation in depth. In another study, Royden et al. (1988) also reported that when a rectangular central region that includes accreting/deleting random dot texture is surrounded by static random dot texture, the accreting/deleting central region was perceived as a rotating cylinder in front of the static surrounding region. However, they suggested that this depth reversal of the accreting/deleting region occurs as a result of the cue conflict between accretion/deletion (as a cue being in behind) and shearing motion (as a cue to being in front). In all these studies, even though the alternative interpretation of accretion/deletion is acknowledged, accretion/deletion was always treated as a definite cue that assigns ground status.

Recently, a series of studies challenged this traditional account of accretion/deletion. Froyen et al. (2013) demonstrated that when accretion/deletion is present on both sides of each border on a display that contains alternating light and dark regions with random dot texture moving horizontally in opposing direction in alternating regions (Figure 4.2), either the dark or the light set of regions were perceived to be 3D figures rotating in depth, while the other set of regions were amodally combined into a large translating surface behind the rotating columns. When the borders of one set of regions had geometric figure/ground cues (e.g. convexity, symmetry, parallelism), that set of regions were more likely to be perceived as rotating columns in front. These results indicated that accretion/deletion combines with the geometry of the border in order to determine relative depth, and might not be an unambiguous cue to ground status.

Using a similar multi-region figure/ground display, Tanrikulu et al. (2016) demonstrated that the rotating-in-front interpretation of accretion/deletion can also occur when only one side of each border contains accretion/deletion. Such a display would be considered as unambiguous by traditional accounts of accretion/deletion in which the accreting/deleting region should indicate the background. Tanrikulu et al. (2016) concluded that the depth ambiguity that arises in their displays is due to the ambiguity inherent to accretion/deletion which can either be interpreted as a translating surface in behind or a rotating 3D figure in front depending on the geometry of the border. In another study, Tanrikulu et al. (2014) investigated how accretion/deletion and the geometry of the border (i.e. piecewise convexity) combine to determine relative depth judgments by manipulating the strength of each factor. Convexity cue was manipulated via the relative salience of the part boundaries on the two side of a piecewise convex border (Hoffman & Singh, 1997). They found that increasing the strength of convexity cue on one set of regions made that set of regions more likely to be seen as figural. Accretion/deletion was manipulated by varying the relative rate of accretion/deletion

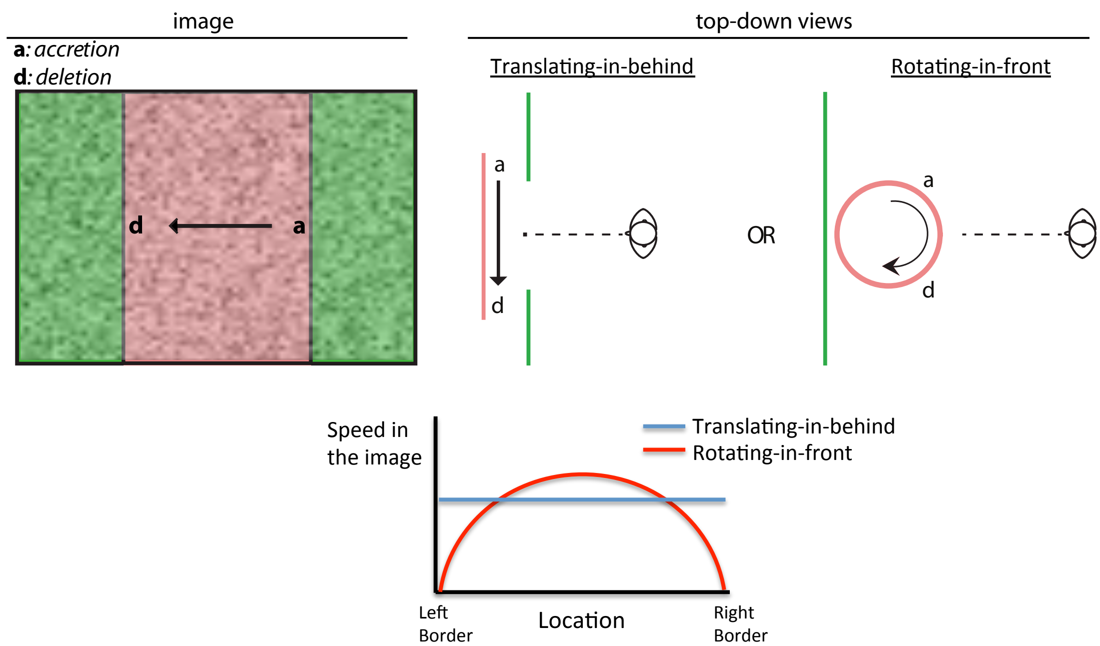

Figure 4.1: The frontal projection of an accreting surface is shown on the left. “a” is the location of texture accretion and “d” is the location of texture deletion. The static surface is depicted in green while the moving surface is depicted in red. On the right, overhead views of the two possible 3D arrangements with different depth-order assignments that are both consistent with the frontal view of the accreting and deleting surface are shown. At the bottom, the two different speed profiles of the moving texture in the image that the two different possible 3D arrangements would create are shown.

(either by changing the relative density or overall speed of accreting/deleting texture). Increasing the relative texture density or texture speed on one set of regions made those regions less likely to seen as figural. However, when the two cues were combined the effect of convexity dominated the relative depth judgments.

In order to identify the critical factors that promote the rotating-in-front interpretation of accretion/deletion, Tanrikulu, Froyen, Feldman, and Singh (2015) focused on the differences between the displays used in standard accretion/deletion studies and the ones used in recent studies (e.g. Figure 4.2) that challenge the traditional accounts of accretion/deletion. By manipulating the factors that make the displays used in these recent studies different from standard accretion/displays, they demonstrated that the geometry of the border and whether accretion/deletion is present on one side or both sides of each border are the two critical factors that significantly influences how accretion/deletion is interpreted.

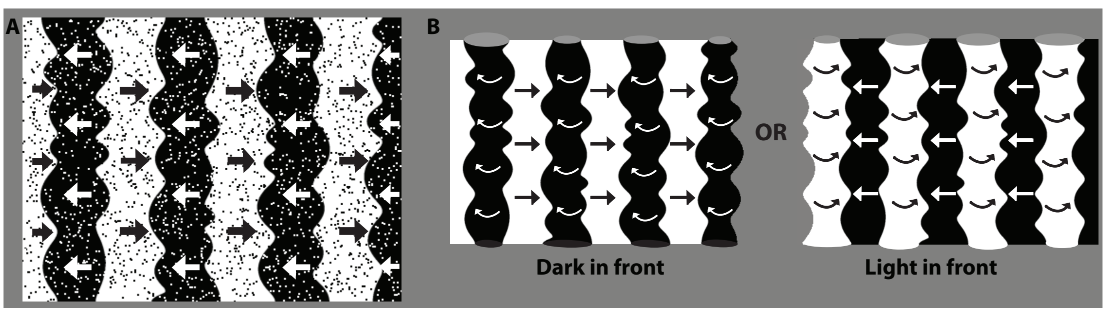

Figure 4.2: Display setup and phenomenology: A. The displays were created by adding motion in one direction to odd regions and in the other direction to even regions in classical figure-ground displays. B. This could yield one of two percepts depending on which one was perceived as figural. The black ones were perceived as rotating in front of a white background which was seen as sliding behind them, or vice versa.

These recent studies on accretion/deletion emphasized the influence of the shape of the border on whether accretion/deletion is interpreted as translating-behind or rotating-in-front, as shown in Figure 4.1. Even though, both of these two possible 3D interpretations will have a 2D projection in which the texture is being accreted and deleted, the 2D speed profile of the accreting/deleting texture will be different in each case. If the accreting/deleting texture is on a flat translating surface, then the 2D image speed of the texture would be constant everywhere inside that accreting/deleting region. If accretion/deletion occurs due to self-occlusion of a rigid 3D column rotating in depth, then the accreting/deleting texture would have a cosine speed profile on the 2D image, in which the speed of the texture elements would be slowest near the boundaries, and fastest near the center of the accreting/deleting region (see the plot shown in Figure 4.1). However, in all the studies summarized above, the rotation-in-depth interpretation arises despite the constant-speed motion of the accreting/deleting texture, which is inconsistent with 3D rotation of rigid objects (e.g. Ullman, 1979). This indicates that the rotating-in-front interpretation of accretion/deletion, which is highly dependent upon the shape of the boundary, is so robust that it makes observers to ignore the inconsistency between the perceived 3D figure and the 2D speed profile of the texture on the image.

Nevertheless, there are couple empirical findings that imply that the speed profile of accreting/deleting texture has an influence on how accreting/deletion is interpreted. In separate experiments, Tanrikulu et al. (2014) observed that increasing the relative density and the speed of accreting/deleting texture in a region makes that region more likely to be perceived as the background. According to the traditional accounts of accretion/deletion, the effect of rate of accretion/deletion would account for these findings. As texture density or speed increases, the rate of accretion/deletion also increases, which in turn will increase the strength of accretion/deletion as a cue to ground status (Kaplan, 1969). However, Tanrikulu et al. (2014) found out that when the data were combined from the two experiments, there was an influence of texture density and speed manipulations that could not captured by rate of accretion/deletion. Tanrikulu et al. (2014) argued that the reason for this result could be attributed to the increasing inconsistency between the image speed profile of the accreting/deleting texture and the cosine speed profile (which is consistent with 3D rotation), as the texture speed and density are increased. In other words, as the speed or the density of the accreting/deleting texture increases, the image becomes more consistent with a translating surface which would be judged to be in behind. In addition to this, it has also been observed that as the width of the accreting/deleting region increases, the probability of interpreting accretion/deleting as a rotating object in front decreases (Thompson et al., 1992; Tanrikulu et al., 2015). These results were also attributed to the influence of texture speed profile on how accretion/deletion is interpreted. In other words, as the width of the accreting/deleting region increases, the image speed profile becomes more inconsistent with a cosine speed profile, which would promote the translating-behind interpretation of accretion/deletion.

Since the geometry of the border has a significant influence on how accretion/deletion is interpreted, then the speed profile of accreting/deleting texture has to combine with the geometry of the border to determine relative depth. It has been shown that the geometry of the boundaries can have an effect on the perceived 3D structure from a dynamic dot displays that are projectively consistent with rotation in depth. Ramachandran, Cobb, and Rogers-Ramachandran (1988) demonstrated that when a dynamic dot display that is projectively consistent with a 3-D vertical cylinder rotating in depth was viewed from a triangular aperture, the perceived structure was a 3-D cone rotating in depth even though the image speed profile of dots were inconsistent with a 3-D rotating cone (i.e. in contrast to a cylinder rotating in depth, an image of a rotating cone would have a speed gradient along its vertical axis)

All these previous studies hints at the possibility that the image speed profile of dynamic texture can combine with the geometry of borders on the image not only to determine the perceived 3D shape but also to determine the relative depth judgments. However, the influence of the speed profile of accreting/deleting texture on perceived depth order has not been investigated yet. In this study, we will be investigating the influence of image speed profile of accreting/deleting texture on relative depth judgments, as well as its combination with the geometry of borders. So far, all studies on accretion/deletion have used dynamic texture that has constant speed profile. By manipulating the speed profile of textural motion in accretion/deletion displays, we will have a better understanding of the relative depth information conveyed by accretion/deletion.

In the following two experiments, we manipulated the image speed profile of accreting/deleting texture to make it more (or less) consistent with either a cosine speed profile (i.e. the red line in the bottom plot of Figure 4.1) or a constant speed profile (i.e. the blue line in the bottom plot of Figure 4.1). We also manipulated the degree of convexity of the border at which the texture accretes/deletes in order to see how texture speed profile combines with the geometry of the border to determine relative depth judgments. In Experiment 1, we used two-region figure/ground displays that are generally used by traditional studies of accretion/deletion. In Experiment 2, we used multi-region figure/ground displays that are used by more recent studies that challenged the traditional accounts of accretion/deletion. Finally, we present a probabilistic model that accounts for our findings without using accretion/deletion as a definite cue to ground status.

4.2 Experiment 1

In this experiment, we investigated how the geometry of the border (i.e. convexity) combines with the image speed profile of the accreting/deleting texture to determine relative depth judgments. We manipulated the degree of convexity of the border separating two regions of a figure/ground display by varying its radius of curvature. We also manipulated the speed profile of the accreting/deleting texture by varying its consistency with 3D rotational motion. In a forced choice task, subjects indicated which of the two regions they perceive in front.

4.2.1 Method

Participants

Six Rutgers University students who were unaware of the purpose of the experiment participated in the study. All had normal or corrected-to-normal visual acuity. All subjects were paid for their participation.

Stimuli

The stimulus consisted of a rectangular area divided into two equal-area regions by a common luminance border. The height of the whole stimulus was 8.6^o and its width was 6.4^o. One of the regions was textured with randomly positioned white dots on a black background, whereas the other region included black dots on a white background ($\sim 70\ dots/deg^2$). The luminance contour could be straight, slightly curved (radius of curvature = 11cm), or highly curved (radius of curvature = 22cm), which are referred as unbiased, weak convexity and strong convexity conditions, respectively (Figure 4.3)

The motion of the dots in the two regions were always horizontal and in opposite direction. The instantaneous 2D image speed of each dot in each frame was determined by a linear combination of a cosine and constant speed profile:

$$S(r) = \alpha C(r) + (1 - \alpha)K$$

where K is the constant speed and $C(r)$ is the cosine speed profile (which is a function of the horizontal location of the dot, r) , and α is the weight assigned to the cosine part of the speed profile. Cosine part of this speed equation was calculated by:

$$C(r) = 2\pi\omega \times cos(0.5\pi\frac{r}{R})$$

where ω is the angular speed, R is the radius of the cross-section and r is the distance of the dot to the central axis (Figure 4.4). For the stimulus we used in this experiment, the rotational axis was chosen to be the left and right stimulus borders for the left and the right region, respectively. In this setup, R corresponds to the distance between the common border of the regions and the left (or the right, depending on which regions the dot is) stimulus border at the vertical location of the dot. r refers to the distance of the dot's current horizontal position to the left (or the right) stimulus border. The constant speed, K is fixed at $1.5^o/sec$. The angular speed, ω, determines the peak speed of the cosine speed profile, and it is set to a value so that either the peak or the average speed of the cosine speed profile is equal to the constant speed. This prevented

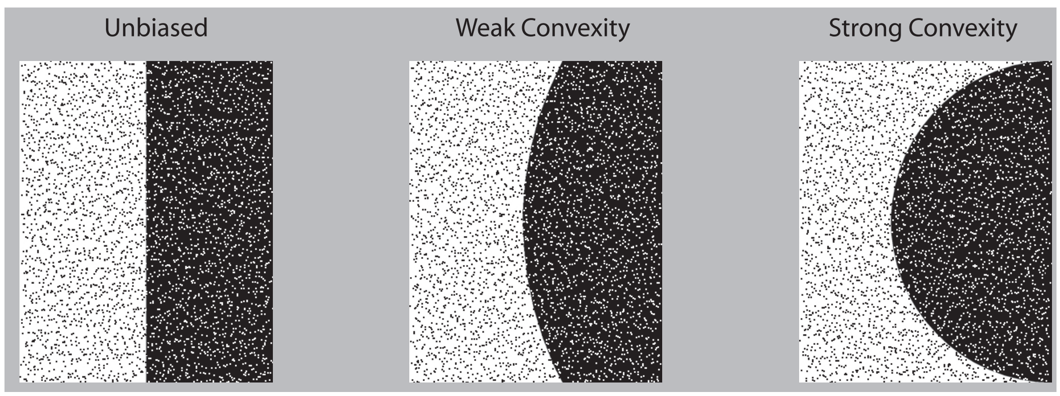

Figure 4.3: The experimental stimuli and the three different convexity levels used in Experiment 1.

subjects to use the average or the peak speed of the textural motion in a region as a heuristic in their judgments. The speed profile of the dots were manipulated by varying the weight assigned to the cosine speed profile of the speed equation (i.e. α), which can be either 0 (i.e. constant speed translational motion), 0.5 or 1 (i.e. fully cosine speed consistent with rotational motion)(Figure 4.5).

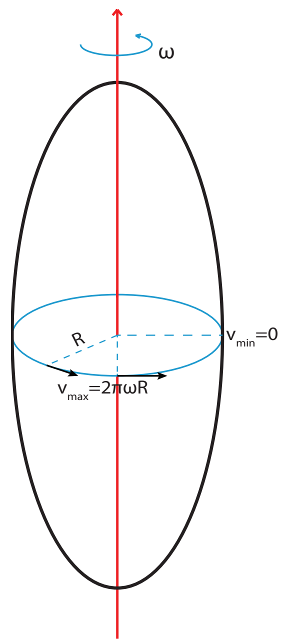

Figure 4.4: Cosine speed profile that is consistent with 3D rotational motion in depth

The size of a single dot was 2.5 arcmin by 2.5 arcmin. The dots had limited lifetimes in order to keep the dot texture uniform over the whole stimulus. In each frame, the location distribution of the dots were checked and if the local texture density of a region of the stimulus is higher than the other parts of the stimulus, then some of the dots are randomly selected from that high density local region to disappear and reappear at the relatively low density local region on the stimulus. This is used in order to prevent subjects to use density gradient of the texture as a cue to 3D shape.

In half of the trials the convex region was black with white dots and the non-convex region was white with black dots, and it was vice versa in the other half. The color of the regions, the direction of motion, whether the left or the right region is convex, and whether the peak or the average of the cosine speed profile equals to the

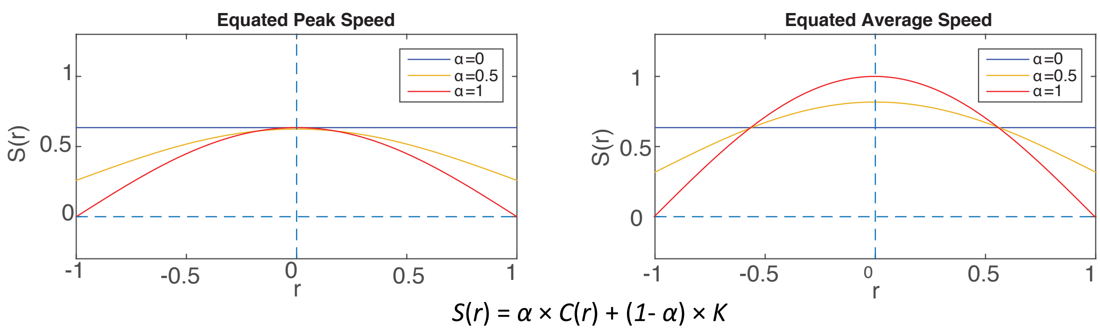

Figure 4.5: The 2D speed profiles of the dot texture motion used in the Experiment 1. The x-axis shows the horizontal location of the dot with respect to the central axis ($r = 0$) and y-axis shows the instantaneous 2D speed of the dot, which is calculated by a linear combination of cosine speed profile ($C(r)$) and a constant speed profile (K). α is the weight given to the cosine part of this linear combination. The three different α values in this experiment and the resulting speed profiles are shown above. In half of the trials, the peak speed of the cosine profile is equated to the constant speed (the figure on the left) and in the other half the average of the cosine profile is equated to the constant speed (the figure on the right).

constant speed were all counterbalanced and crossed with other factors.

Design and Procedure

Subjects sat 76 cm from a 21-inch CRT monitor (85Hz, 1280pxl x 1024pxl) connected to a Windows 7 PC. Subjects' head position was stabilized with a chin rest. The stimulus was presented using Psychtoolbox in Matlab (Brainard, 1997; Kleiner et al., 2007).

There were three conditions regarding the convexity of the border, which were unbiased, weak convexity and strong convexity (Figure 4.3). The speed profile of the two regions (i.e. α values) were manipulated independently. There were three levels of α value for the convex region, and three levels for the non-convex region (0, 0.5 or 1), resulting in nine different conditions regarding the speed profile of the accreting/deleting texture. Different levels of convexity and speed profile were crossed with each other yielding 27 experimental conditions.

On each trial, subjects were presented with 800 msec of pre-mask and then 800 msec of pre-mask with a fixation cross added to the center of the mask. The masks were made from a random dot texture sampled from a beta distribution with parameters $\alpha = 4$, $\beta = 4$ that resulted in a middle gray random dot texture. The fixation cross was positioned such that the subjects were fixated on the common border between the convex and non-convex regions. Since the geometry of the border varies in each trial, the exact location of the fixation cross also slightly varied from trial to trial. The mask was made one and a half times bigger than the size of the stimulus so that it would be more difficult for the subjects to be primed by the position of the fixation cross. Following the mask, the experimental display was shown for 1.2 sec. After the stimulus disappeared, a post–mask which was identical to the pre-mask was shown, and subjects pressed the the corresponding key to indicate whether they perceive the left or the right region in front. They were also given a third "neither" choice that they can use if they absolutely did not perceive any figure-ground separation. As soon as the subject responded, the next trial started immediately. The subjects were also instructed to respond with their initial depth order judgment in order to avoid any influence of figure/ground reversals (i.e. top-down effects) on the results.

Subjects completed 864 trials in total split into 8 blocks, i.e. 3(convexity levels: unbiased, weak, strong) x 9(speed profile) x 2(luminance: dark/light) x 2(phase: convex region on the right or on the left) x 2(direction of motion: left/right) x 6(repetitions). The experiment was completed in two sessions in different days. Each session contained 4 blocks where each block contained 108 trials. All trials were randomized for each subject separately. At the beginning of each session, subjects were given practice trials to get them familiar with the task and the stimulus. Subjects completed 60 practice trials at the beginning of the first session, and 30 practice trials at the beginning of the second session. Each session took approximately 45 minutes to complete.

4.2.2 Results

Figure 4.6 shows the proportion of trials in which the convex region was perceived in front for each subject separately. Subjects rarely used the "neither" response (i.e. neither responses were only 2% of all the responses overall). Therefore, the following analyses were done with the two main relative depth responses of the subjects (i.e. excluding the "neither" responses).

A logistic regression was performed for each subject separately. A likelihood ratio test conducted on these logistic regressions revealed that the main effect of convexity was a significant improvement over an unconditional-means model that contains only an intercept for all six subjects ($LR_{max} = 322.44$, $LR_{min} = 15.24$, $df = 2$, $p < 0.001$). A Tukey test showed that for four of these six subjects (DC, JD, KT and MK) there was a significant difference between the unbiased condition and the two convexity levels. For one subject (JW), all the pairwise comparisons among the convexity conditions yielded significant differences, in which proportion of perceiving a region in front increases as convexity increases. For one subject (LAK), the effect of convexity was in the other direction, where the proportions obtained from the strong convexity condition was significantly lower than the proportions in the other two conditions (i.e. unbiased and weak convexity). The effect of convexity can also be seen in Figure 4.6. Except for the subject LAK, the lines in the leftmost plot of each individual (i.e. unbiased condition) jump up to higher proportion values as soon as convexity is introduced (i.e. the center and the rightmost plots of each individual).

The addition of the α value of the dot motion in the convex region to these individual logistic models that include the effect of convexity was also found to be a significant addition for all six subjects (LR_{max} = 134.3, LR_{min} = 6.17, df = 2, $p < 0.05$). For all six subjects, the proportion of perceiving the convex region in front increased as the α value of the dot motion in that region increased. This strong effect of the α value of the dot motion in the convex region can be also seen in Figure 4.6 from the highly positive slope of most of the lines in each plot.

When the α value of the dot motion in the non-convex region was added to these individual logistic models, it resulted in a significant improvement for five of six subjects (LR_{max} = 227.82, LR_{min} = 27.22, df = 2, $p < 0.001$). For all of these five subjects the proportion of perceiving the convex region in front decreased as the α value of the dot motion in the non-convex region increased. This effect is demonstrated in Figure 4.6 as the separation between the different colored lines in each single plot. The only subject who did not show the effect of the α value of the non-convex region was JD.

The addition of the color factor to these individual logistic models was founded to be significant addition for five of the six subjects (LR_{max} = 128.64, LR_{min} = 7.49, df = 1, $p < 0.01$). Among these five, four subjects showed a light color bias (i.e. more likely to perceive light regions in front compared to dark regions) whereas one of them showed a dark color bias.

For three subjects (JD, KT and MK), adding the interaction factor between convexity and the α value of the non-convex region to these individual models yielded a significant improvement (LR_{max} = 21.01, LR_{min} = 12.78, df = 4, $p < 0.05$). For two of these subjects (JD and MK) the effect of α value in the non-convex region was only seen when the border was unbiased. For subject KT, when the border was unbiased, the proportion of perceiving the convex region in front monotonously decreased as the α value of the non-convex region was increased. However, when the border has either weak or strong convexity, there was no difference between α of 0 and 0.5 on the non-convex region for subject KT. For two subjects (DC and JW), adding the interaction

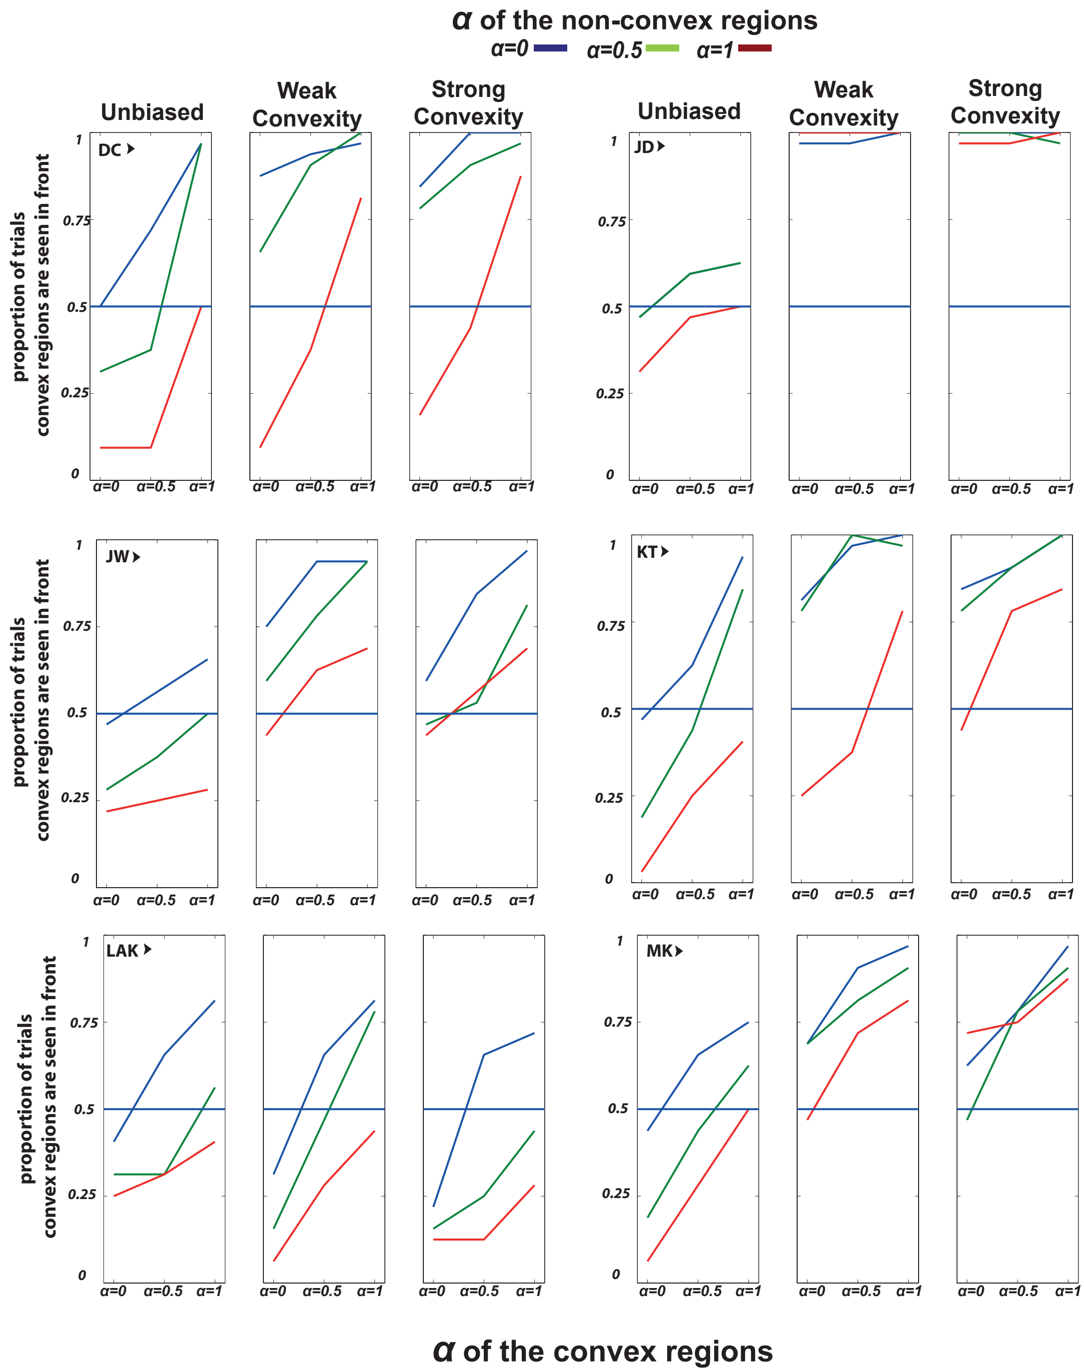

Figure 4.6: Results of Experiment 1: Proportion of trials the convex region perceived in front is shown on y-axis. The blue line shows the chance level, i.e. where the proportion equals to 0.5. The three different plots for each of the six individuals correspond to the three convexity levels used in the experiment. The x-axis corresponds to the alpha levels of the dot motion in the convex region, whereas different lines in each plot correspond to the different alpha levels of the dot motion in the non-convex region.

term between convexity and color also yielded a significant addition to these individual models (LR_{max} = 14.66, LR_{min} = 12.8, df = 4, $p < 0.01$). Both subjects' responses yielded a significant effect of color only when the border was unbiased.

4.2.3 Discussion

Both manipulations in Experiment 1 (i.e. convexity and speed profile of the accreting/deleting texture) had strong influence on subjects' relative depth. Except for one subject, increasing the degree of convexity on one side of the border made that side of the border more likely to be perceived as figural. This was expected given that convexity is a well known cue to figural status.

The crucial manipulation in this experiment was the speed profile of the accreting/deleting texture, which was manipulated independently on each side of the border (i.e. convex and the non-convex region). The speed profile of the moving texture had a strong influence on relative depth judgments for all subjects. We observed that when the speed profile of the moving texture in a region was more consistent with translation (e.g. α=0), then subjects' were more likely to see that region as behind. However, when the speed profile in a region was more consistent with 3D rotation in depth (e.g. α=1) subjects' were more likely to perceive that region in front. This indicates that the speed profile of the accreting/deleting texture influences how accretion/deletion is interpreted (i.e. translating-in-behind or rotating-in-front), which, in turn, influences the perceived depth order.

Even though the α values of both the convex and the non-convex region significantly influenced relative depth judgments, there seems to be slight difference between these two factors in how they interacted with convexity of the border. While the α value of the convex region did not seem to be affected by the convexity of the border, the effect of the α value of the non-convex region interacted to some extent with the effect of convexity for three of the six subjects (JD, KT and MK). For subject JD, the effect of α was only observed when the border was unbiased. This was mainly because of the ceiling effect observed for subject JD in both weak and strong convex conditions. Similarly, for subjects KT and MK, the effect of α of the non-convex region was much more strong when the border is unbiased. This interaction can be seen in Figure 4.6 from the different degree of separation of the colored lines in different convexity levels for these subjects. For these three subjects, convexity seems to cancel out the effect of the speed profile in the non-convex region, but not in the convex region.

Even though both convexity and the speed profile of the accreting/deleting texture influenced subjects' responses, when in competition, it was observed that convexity tends to dominate the depth order judgments. When the plots for the weak and strong convexity conditions in Figure 4.6 is examined, it can be seen that in these conditions the lines are mostly over proportion of 0.5 (except subject LAK). This indicates that the convexity of the border could bias subjects to ignore the projective consistency (i.e. speed profile of the moving texture) in perceiving the 3D structure and layout from motion.

4.3 Experiment 2

The stimulus used in Experiment 1 included two regions separated by a simple single border, which is similar to the stimuli used by traditional studies of accretion/deletion. However, recent studies that have challenged the traditional understanding of accretion/deletion (Froyen et al., 2013; Tanrikulu et al., 2014, 2015, 2016) generally used multi-region figure/ground stimulus (e.g. eight regions), in which the borders were composed of several convex parts. In Experiment 2, we repeated the same experimental design and procedure used in Experiment 1, but this time using a similar stimulus used by these recent studies.

4.3.1 Method

Participants

Six Rutgers University students who were unaware of the purpose of the experiment participated in the study. All had normal or corrected-to-normal visual acuity. All subjects were paid for their participation.

Stimuli and Procedure

The stimulus consisted of eight alternating black and white vertical regions. The height of the stimulus was 6.1° and its width was 8.6°. Either the odd or the even regions were given the convexity cue. In Experiment 2, the term "convex" refers to a boundary that is piecewise convex such that the boundary can be segmented into parts, where each individual part is convex. The part boundaries are defined by the negative minima of curvature (Hoffman & Singh, 1997).

In this experiment, we used different degrees of convexity in which the convexity of a boundary was manipulated by gradually making the negative minima of curvature sharper on one side. In other words, the relative salience of the part boundaries on the two sides of a boundary was manipulated (Hoffman & Singh, 1997). This is done by creating contours with cubic spline interpolation and then manipulating the sharpness of the negative minima on one side of the contour by shifting the positions of the knots of the spline curve. Figure 4.7 shows an example of the different levels of convexity.

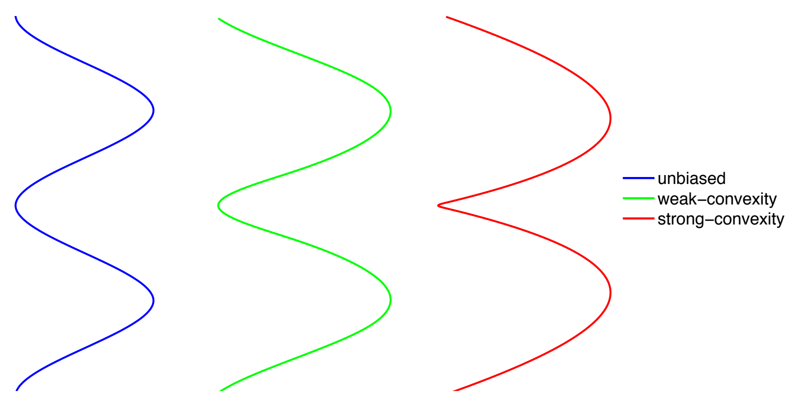

Figure 4.7: An example showing the relation between the sharpness of the negative minima of curvature and the (part-wise) convexity levels used in the experiment. The blue curve is similar to a sine wave where the two sides of the curve is unbiased in terms of convexity. The sharpness of the negative minima is gradually increased in order to obtain the green (weak-convexity) and the red curves (strong-convexity), in which the left side of each curve becomes piecewise convex.

For each stimulus, seven individual contours were generated where the size (i.e. height and width) of each part of a contour was randomized within a certain range. Each contour had five and a half parts. Whether a contour starts or ends with a minima or a maxima of curvature was also randomized. These seven contours were placed on the stimulus such that the area between each contour is the same. For this experiment we had three levels of convexity which we again refer as “strong convexity”, “weak convexity” and ”unbiased” (in terms of convexity) (Figure 4.8). For the unbiased contours the knots were set so that the sum of signed curvature was zero along the boundary.

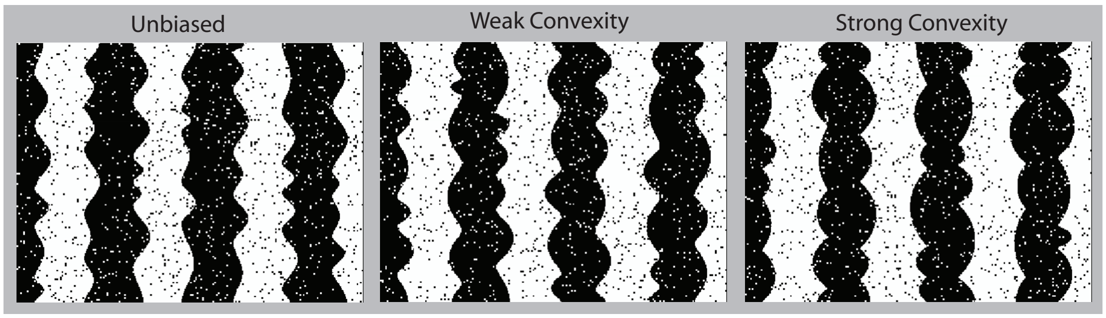

Figure 4.8: An example of the three levels of convexity used in Experiment 2.

The method used to add texture and motion to that texture was the same as the method used in Experiment 1. The motion of the dots in the dark and light regions were always horizontal and in opposite direction. The speed profile of the dots in the convex and and in the non-convex region was manipulated independently. The texture in all the convex regions had the same α value in each trial. In the same way, the α value for the non-convex regions were the same among each other in each trial.

The experimental design and procedure was the same as in Experiment 1 except two differences. In Experiment 2, subjects indicated the set of regions (i.e. light or dark) that they perceived to be amodally completed in behind. Such a question encourages subjects to look at the whole stimulus when responding, instead of focusing on a single region of the stimulus. Secondly, subjects were not given the “neither” option in Experiment 2. In Experiment 1, we have observed that subjects rarely unable to resolve the relative depth ambiguity and use the “neither” option. Therefore, only the two main options (either dark or light set of regions are in behind) were given to the subjects. As in Experiment 1, each subject completed 864 trials in two different sessions for Experiment 2. Each session took approximately 45 minutes to complete.

4.3.2 Results

Figure 4.9 shows the proportion of trials convex regions were perceived in front for each subject separately. A logistic regression was performed for each subject. A likelihood ratio test conducted on these logistic regressions revealed that the main effect of convexity was a significant improvement over an unconditional-means model that contains only an intercept for all six subjects ($LR_{max} = 163.02$, $LR_{min} = 12.86$, $df = 2$, $p < 0.01$). A Tukey test showed that all the pair-wise comparison of convexity levels yielded significant differences for all six subjects, (except between weak and strong convex conditions for subject US). The effect of convexity can also be seen in Figure 4.9. The lines in the leftmost plot of each individual (i.e. unbiased condition) jump up to higher proportion values as soon as convexity is introduced (i.e. the center and the rightmost plots of each individual).

The addition of the α value of the dot motion in the convex regions as a factor to these individual logistic models that already include the effect of convexity was found to be a significant addition for only two subjects (DL and JR) ($LR_{max} = 16.57$, $LR_{min} = 14.73$, $df = 2$, $p < 0.05$). For both of these subjects, the proportion of perceiving the convex regions in front increased as the α value of the dot motion in those regions increased. This effect can be also seen from the plots for subject DL and JR in Figure 4.9 in which the slope of the lines are generally positive for these two subjects.

When the α value of the dot motion in the non-convex regions was added to these individual logistic models, it resulted in a significant improvement for four (DL, JR, TB and US) of six subjects ($LR_{max} = 24.88$, $LR_{min} = 15.54$, $df = 2$, $p < 0.001$). For all of these four subjects the proportion of perceiving the convex regions in front decreased as the α value of the dot motion in the non-convex region increased. This effect is demonstrated in Figure 4.9 as the separation between the different colored lines in the plots of these four subjects.

The addition of the color factor to these individual logistic models was founded to be significant addition for five of the six subjects ($LR_{max} = 725.04$, $LR_{min} = 6.84$,

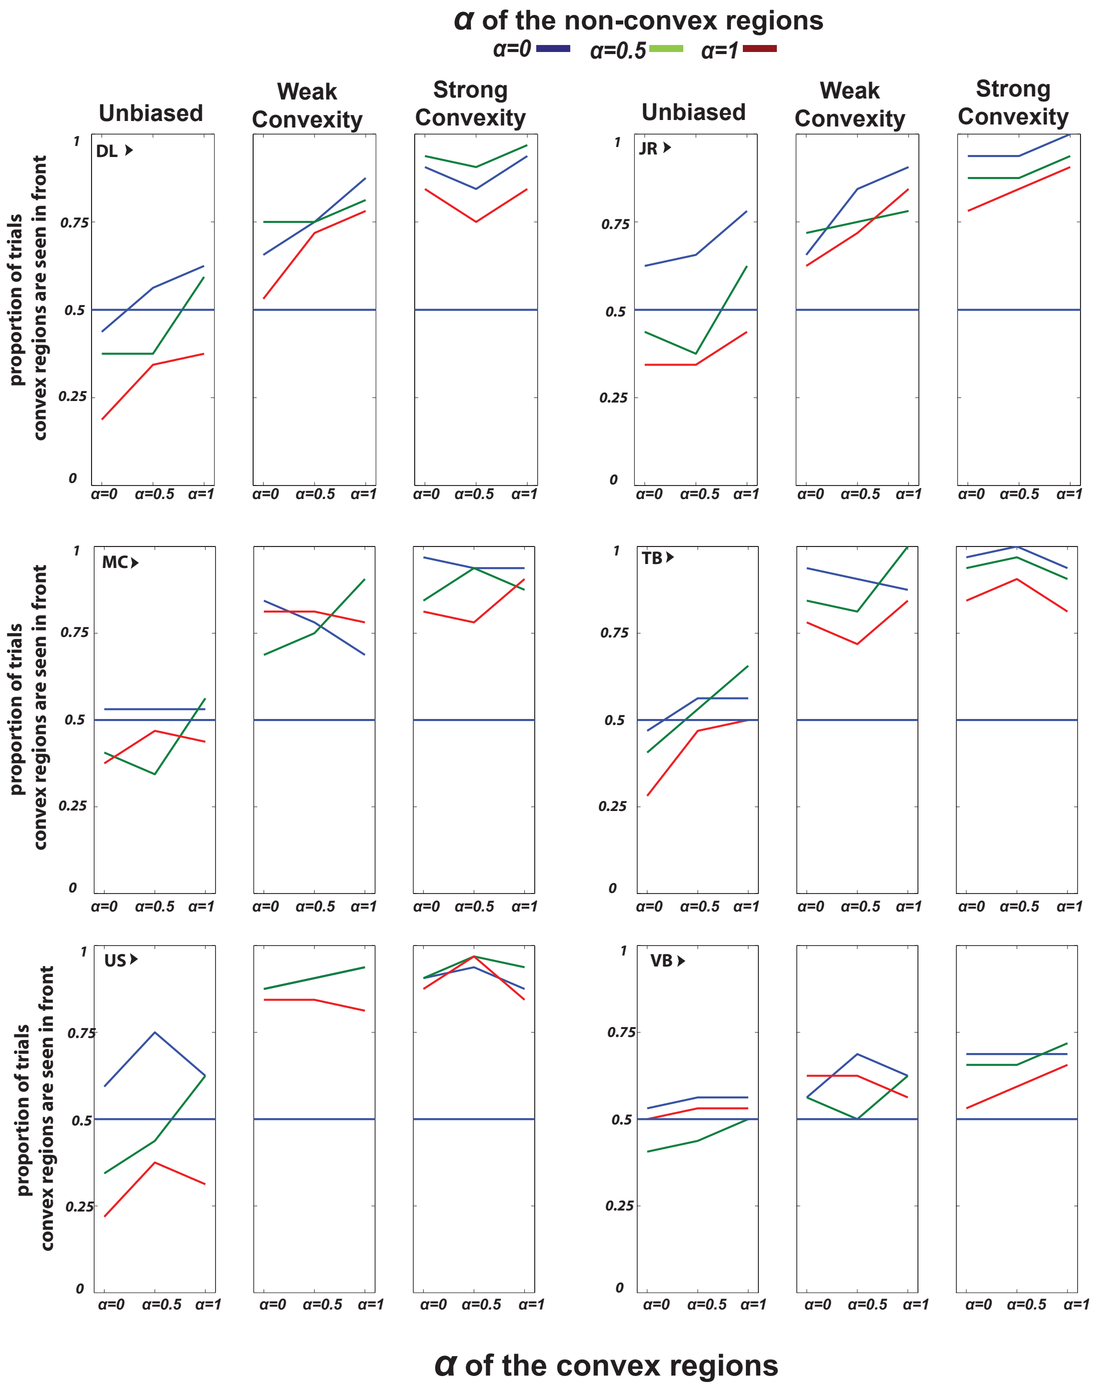

Figure 4.9: Results of Experiment 2: Proportion of trials the convex regions perceived in front is shown on y-axis. The blue line shows the chance level, i.e. where the proportion equals to 0.5. The three different plots for each of the six individuals correspond to the three convexity levels used in the experiment. The x-axis corresponds to the alpha levels of the dot motion in the convex regions, whereas different lines in each plot correspond to the different alpha levels of the dot motion in the non-convex regions.

$df = 1$, $p < 0.01$). Among these five, four subjects showed a dark color bias (i.e. more likely to perceive dark regions in front compared to light regions) whereas one of them showed a light color bias.

For only one subject (TB), adding the interaction factor between convexity and the α value of the convex region yielded a significant improvement ($LR = 15.48$, $df = 6$, $p < 0.05$). For this subject, the effect of the α value of the convex regions was only observed when the border was unbiased in terms of convexity.

4.3.3 Discussion

Among the two factors that were manipulated in Experiment 2, convexity turned out to be the factor that mostly determined relative depth judgments as opposed to the speed profile of the accreting/deleting texture. The responses of all six subjects showed a significant effect of convexity (i.e. as the degree of convexity increases in a set of regions, that set of regions were more likely to be seen as figural). However, there are individual differences with respect to the effect of α manipulation. Responses of only two subjects were significantly influenced by the α value of the texture motion in the convex regions, while responses of four subjects were significantly influenced by the α value in the non-convex regions. The direction of the effect of α manipulation was consistent among these subjects. As the speed profile of the texture became more consistent with a translating surface, subjects were more likely to see that region in behind, whereas when it became more consistent with 3D rotation in depth, subjects were more likely to see that region in front.

Overall, when the convexity and the speed profile of the accreting/deleting texture competes, convexity seems to dominate the relative depth judgments. This can be again seen from the plots of the weak and strong convexity conditions in Figure 4.9. In these conditions the lines are generally over 0.5 (more likely to be seen as in front) regardless of the α value of the convex or the non-convex regions. This again indicates the geometry of the border could bias subjects to ignore the projective inconsistency between the speed profile of the moving texture and the perceived 3D structure.

4.4 Computational Model

A Bayesian model was built to account for our findings without using accretion/deletion as a definite cue to ground status. For the purpose of the model, we focus on the simpler two-region display used in Experiment 1 (which allows us to focus on the contributions of the three relevant variables manipulated, without having to worry about possible interactions across regions in the 8-region displays). The model calculates and assigns probabilities to possible hypotheses of relative depth and 3D structures that could have generated the images used in Experiment 1. The equation used is

$$p(\text{hypotheses}|\text{image}) \propto p(\text{image}|\text{hypothesis}) \times p(\text{hypothesis})$$

where the posterior probability of an hypothesis given an image is proportional to the product of the likelihood of the hypotheses and prior probabilities of these possible hypotheses.

The model uses accretion/deletion as a simplifying assumption that determines the hypothesis space rather than a cue to relative depth. According to this simplifying assumption, if an accretion/deleting region is perceived to be in front, then accretion/deletion has to be explained by self occlusion. Hence, the surface has to be rotating in depth. If an accretion/deleting region is perceived to be behind, then the accretion/deletion has to be explained by a translating flat surface. In other words, the relative depth and the structure of an accreting/deleting surface are dependent upon each other. Being in front is coupled with rotational motion, and being in behind is coupled with translational motion. With this simplifying assumption, given an image that consists of two accreting/deleting regions separated by a common border (like in Experiment 1), then there are two possible 3D layouts that could have generated this image. These are shown in Figure 4.10. The prior probabilities over the two hypotheses are taken to be equal. The idea of using accretion/deletion as a simplifying assumption originates from studies by Thompson et al. (1992) and Thompson and Painter (1992), in which they used accretion/deletion as a qualitative constraint in their computational model of structure from motion. Here, we are applying a similar idea to mainly determine relative depth.

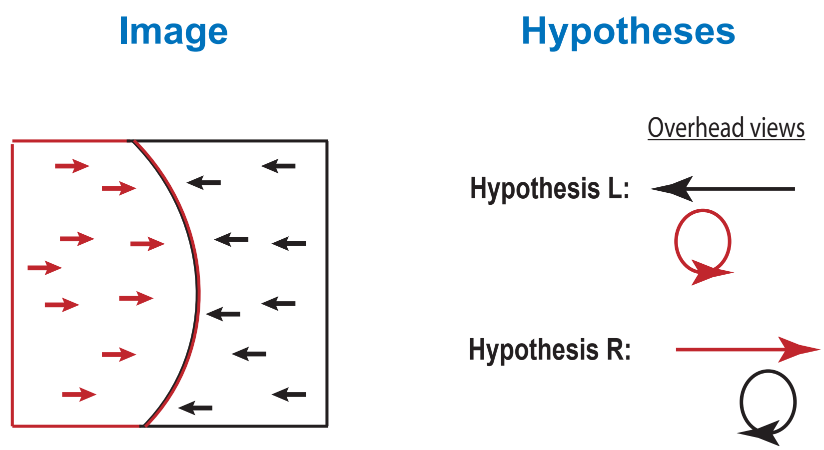

Figure 4.10: The sketch of the image used as a stimulus in Experiment 1, and the two possible 3D layouts that could have generated this image, given the simplifying assumption of accretion/deletion. According to the Hypothesis L, the left region is in from and rotating in depth, whereas the right region is a flat surface translating in behind. It is the opposite for the Hypothesis R, in which the right region is rotating in front of the translating flat surface that is on the left. These are the two hypotheses used by our computational model. See text for more information

The equation for the likelihood term is expanded as shown below.

$$
\begin{aligned}
p(\text{image}|\text{hypothesis}) = p(&\text{speed-profile-convex-side}|\text{hypothesis}) \\
&\times p(\text{speed-profile-nonconvex-side}|\text{hypothesis}) \\
&\qquad\times p(\text{convexity-convex-side}|\text{hypothesis})
\end{aligned}
$$

The model calculates the probability of observing the speed profiles on the convex and the non-convex region separately for a given hypothesis. Then it combines this with the probability of observing a certain degree of convexity given that hypothesis.

An example of how the probability of the speed profile given an hypothesis was calculated is shown in Figure 4.11A. The inputs to the model are the instantaneous speed values and location of each dot on the image. In other words, the model takes a snapshot of the dynamic stimulus. The red data points in Figure 4.11A are example speed and location measurements of the dots on the image. The x-axis corresponds to the location of the dot with respect to the borders of the region the dot is located. The y-axis corresponds to the speed of the dot. The solid black lines represent the two possible hypotheses, which are either translational (on the right) and rotational (on the left) motion. Given a certain noise level, the probability of obtaining image data points from the two generative models are calculated. The speed values of the dots in the image were sampled from a Gaussian distribution centered on the solid black lines on Figure 4.11A (depending on the hypothesis) as follows:

$$
\begin{aligned}
S_{constant} &\sim N(\mu_{constant}, \sigma^2) \\
S_{cosine} &\sim N(\mu_{cosine}, \sigma^2) \\
\mu_{cosine} &= \omega \times cos(r \times 0.5 \times \pi)) \\
\mu_{constant} &= 2\text{pxl/frame}
\end{aligned}
$$

where $S_{constant}$ and S_{cosine} are the speed values of a dot in the image generated from the constant and cosine speed profile hypotheses, respectively. ω was set to a value that would either make the peak or the mean speed of the cosine motion profile equal to the constant speed (i.e. $\mu_{constant}$). r is the distance of the dot from the middle axis of rotation (i.e. mid-point between the two borders of a region at the horizontal cross-section where the corresponding dot is located.). The standard deviation (σ^2) of these Gaussian distributions was treated as a free parameter for each individual.

In order to calculate the likelihood on convexity, we used the joint statistics of depth and convexity that Burge, Fowlkes, and Banks (2010) obtained by analyzing natural image statistics. They found that the convexity-depth distributions follow a power law over depth:

$$p = \text{depth}^k$$

where p is the probability and the value of the constant k depends the convexity value. Figure 4.11B shows an example of this relationship. Degree of convexity is treated as a continuous variable as defined by Fowlkes, Martin, and Malik (2007). The graph

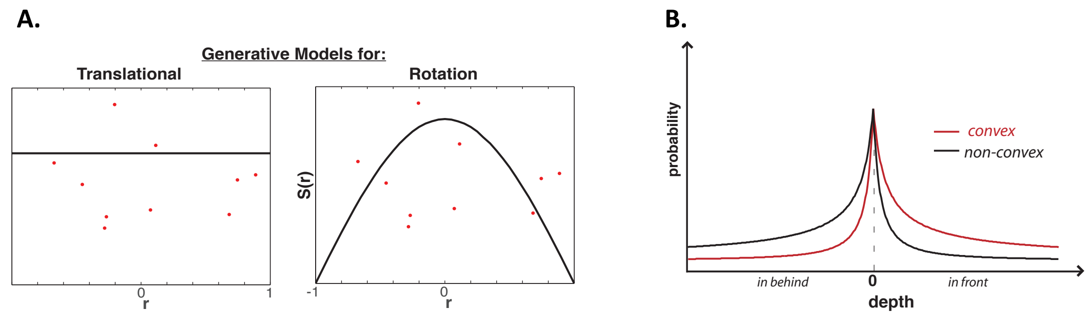

Figure 4.11: A. An example of how speed profile likelihood is calculated. x-axis shows the location of the dot with respect to the axis of rotation. y-axis shows the speed of the dot. Each red dot is an example speed-location measurement of a moving dot on the image. The two plots and the solid black lines in each show the two possible hypotheses regarding the type of motion. Given a noise level, probability of obtaining the red data points from these hypotheses (solid black lines) is calculated. See text for more information. B. Probability of depth difference between two regions given a certain convexity degree for the border that separates the regions. This graph only shows the distribution for one particular value of convexity. The shape of the distribution will slightly change if different convexity values are used.

demonstrates the probability distribution over depth given a particular value of convexity. Given a convex border, the red line corresponds to the convex side of that border, whereas the black line corresponds to the concave side of that same border. That is why the two lines are actually mirror symmetric versions of each other. Since the constant k depends on the convexity value, the shape of the curve in Figure 4.11B would change for each different convexity value. The difference between the depth distributions for the convex and the concave contours (i.e. the separation between the red and the black curve in Figure 4.11B) is called “k-difference”, which is assumed to be a linear function of the convexity value. The slope of this “k-difference”-convexity relation is treated as another free parameter for each individual. Depth is treated as a continuous variable in this equation. Since we are only interested in depth order, we integrate over positive values of depth to calculate the probability of being in front, and over negative values of depth to calculate the probability of being behind.

The relationship in Figure 4.11B gives us the probability of depth given a convexity value. From this we compute probability of a convexity value given the region is in front (or back). Then we combine this with the probability of the speed profile in the convex and the non-convex region (given hypotheses) to obtain the posterior probabilities for each of the two hypotheses shown in Figure 4.10.

Figure 4.12 shows the model predictions superimposed on the individual data from Experiment 1. Figure 4.13 shows the values of the best fitting parameters for the six subjects. The goodness of fit values (R^2) are also shown in Figure 4.12 for each individual. As seen from these figures, the model captures the main effect of the three main factors (i.e. convexity, speed profile of the convex region and of the non-convex region). However, for all subjects (except subject LAK) the model underestimates the proportions especially in the weak-convexity condition. Even though the model captures the main effect of convexity, there is a discrepancy between the model's predictions and subjects' data with respect to the influence of convexity. This can be also seen by comparing the R^2 values of the individuals and the best fitting parameters of these individuals in Figure 4.13. The highest R^2 is obtained from subject LAK who did not show any effect of convexity (the estimated slope of k-difference was 0 for this subject), whereas the lowest R^2 is obtained from subject JD whose responses were strongly influenced by convexity (i.e. the highest estimated slope of k-difference was for this subject).

One possible reason for this observed discrepancy could be due to a similar discrepancy observed by Burge et al. (2010) between their natural image statistics and their psychophysical data. The k-difference value that they obtained from their psychophysical experiment (~ 4.4) was significantly higher than the k-difference value that they obtained from their natural image statistics (~ 0.4). This shows that the influence of convexity on relative depth was much stronger for the subjects compared to what is observed in the natural image statistics. Since our probabilistic model is using the natural image statistics that is obtained by Burge et al. (2010), it is not surprising that our model is underestimating the influence of convexity on relative depth judgments.

However, as mentioned before, it seems that the model is underestimating the influence of convexity especially in the weak-convexity condition. Another possible reason for this might be related to the method used to measure convexity, which is

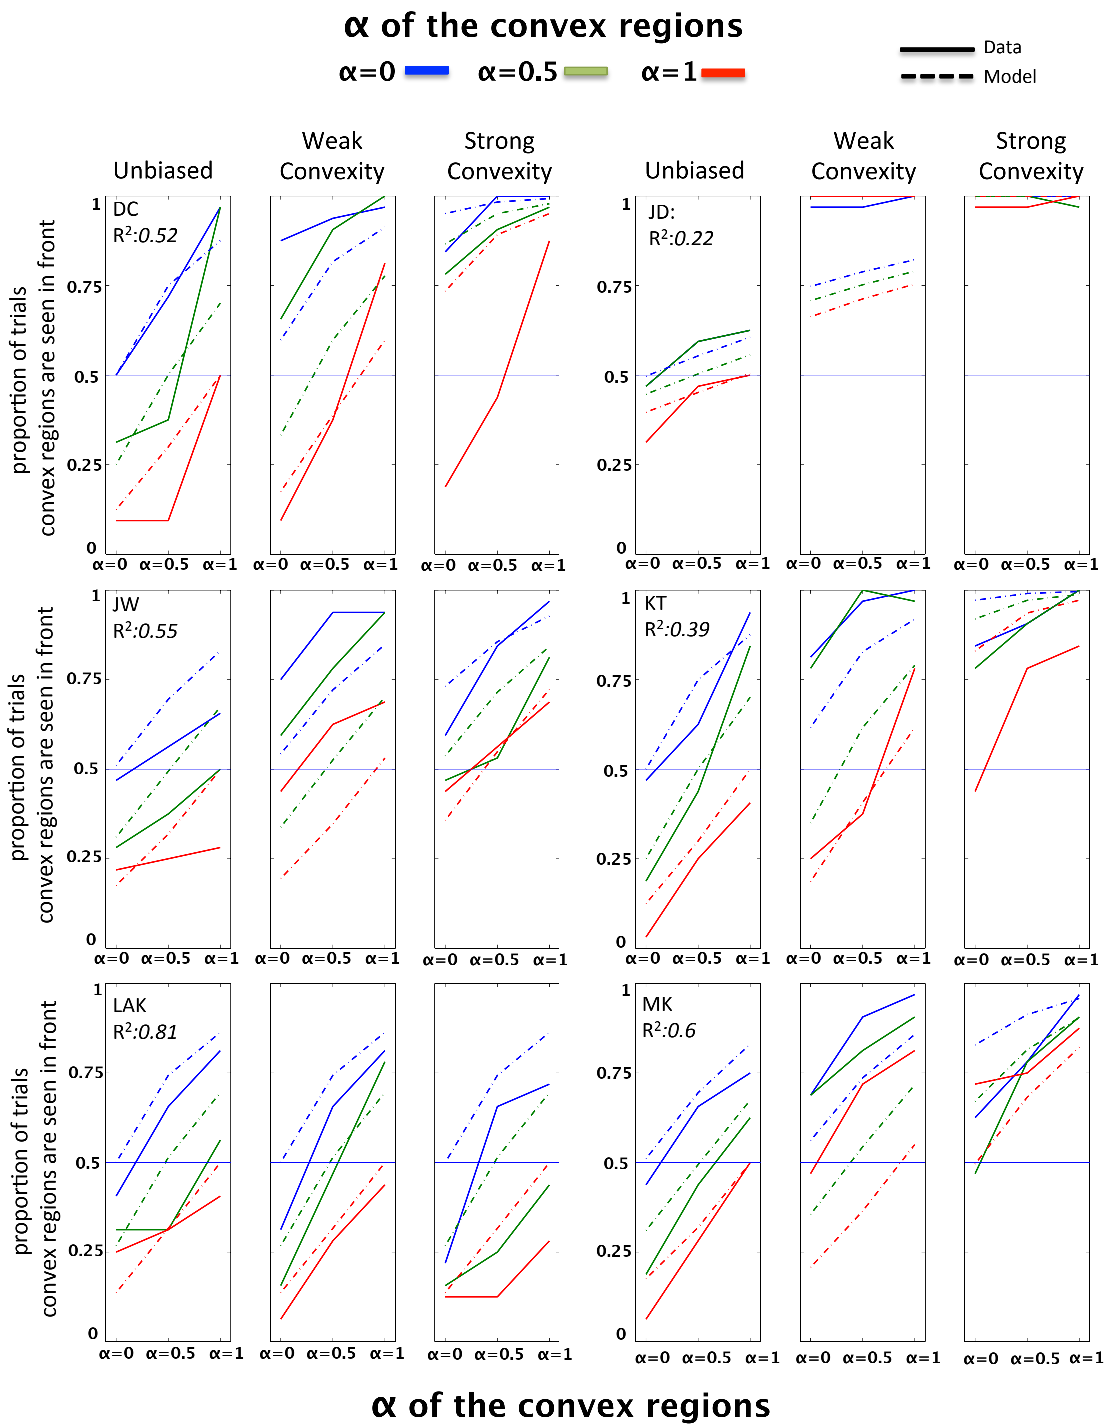

Figure 4.12: The predictions of the model (dashed lines) superimposed on the individual data (solid lines) from Experiment 1. Proportion of trials the convex region perceived in front is shown on y-axis. The blue line shows the chance level, i.e. where the proportion equals to 0.5. The three different plots for each of the six individuals correspond to the three convexity levels used in the experiment. The x-axis corresponds to the alpha levels of the dot motion in the convex region, whereas different lines in each plot correspond to the different alpha levels of the dot motion in the non-convex region. The goodness of fit (R^2) is indicated at the top left of each subject's plot.

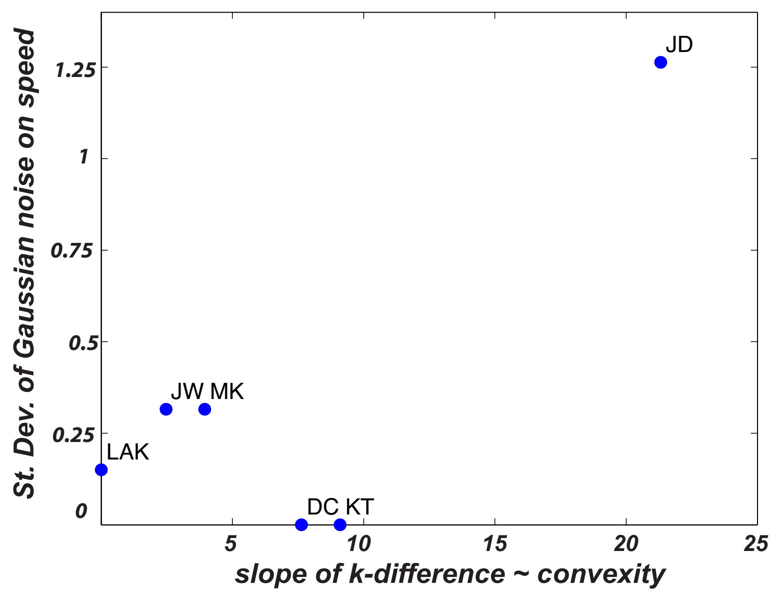

Figure 4.13: The best fitting free parameter values for each subject of Experiment 1.

the method proposed by Fowlkes et al. (2007). This method involves centering a small circular local analysis window on the contour and then calculating the ratio of pairs of randomly sampled dots that can see each other (i.e. that you can draw a line between the two dots that does not intersect the border) inside that analysis window. Then, log ratio of these two ratios obtained from the two regions is calculated and that value is considered to be the relative convexity value of one side of the border (i.e. the side that is used as the numerator in the log ratio). We applied the same method to our stimuli used in Experiment 1 with using an analysis window that will include the whole border. The convexity values we obtained were 0 for unbiased, 0.02 for weak-convexity and 0.15 for strong-convexity. As it is seen from these values, the measured convexity value of the weak-convex stimulus is almost zero, even though weak-convex region perceptually looks very convex. This discrepancy between the measured convexity value of the weak-convex stimulus and how it looks perceptually might be a reason why the model was underestimating the proportions in this condition.

The model was only compared to the data from Experiment 1, but not from Experiment 2. One of the reasons for this is the anomaly observed when the convexity calculation method of Fowlkes et al. (2007) was applied to the borders used in Experiment 2. When the convexity of the piecewise convex region was calculated, negative values of convexity was obtained, which means that the *perceptually* non-convex region were found the be more convex (with respect to method of Fowlkes et al. (2007)) than the *perceptually* convex region. This is due to the sharp cusps at the part boundaries of the border, which yields a high convexity value on the *perceptually* non-convex region. This shows that the method proposed by Fowlkes et al. (2007) can not capture the convexity of piecewise convex contours. Therefore, we limited our comparison of our model only to data from Experiment 1, in which we did not encounter any problems, since the border consists of only one convex part.

4.5 General Discussion

Recent studies (e.g. Tanrikulu et al., 2014, 2016) have argued that accretion/deletion is not a definite cue to ground status, but can be interpreted as indicating either the front or the back surface, depending on how accretion/deletion of the texture is explained. If accreting/deleting of texture is explained as a translating surface being occluded by its adjacent surface, then it would indicate the surface in behind; whereas if it is explained by self-occlusion due to rotation in depth, then it would indicate the surface in front. If this account is correct, then the speed profile of the accreting/deleting texture should have a strong influence on relative depth judgments, since the speed profile of the texture will provide information about the type of motion (translating vs. rotating), which, in turn, support a certain interpretation of accretion/deletion.

We found that the speed profile of accreting/deleting texture has a strong influence on relative depth judgments in both two-region and eight-region figure ground displays. When the speed profile of the accreting/deleting texture is more consistent with 3D rotating in depth, subjects tend to perceive that region more likely to be in front. In the same way, subjects tend to perceive accreting/deleting regions more likely to be in behind if the speed profile of the texture is more consistent with translational motion of a flat surface. This indicates that the depth information conveyed by accretion/deletion depends on how it is interpreted.

We also manipulated the geometry of the border at which the texture was accreting/deleting by varying the degree of convexity of the border. We found that when the speed profile of the accreting/deleting texture was combined with convexity of the border, convexity tends to dominate the relative depth judgments. For example, when the accreting/deleting texture had constant motion profile (which would support translating-in-behind interpretation) inside a highly convex region, subjects were still more likely to perceive that region as a surface rotating in front. In other words, perception of 3D rotation occurs despite the speed profile which is highly inconsistent with the projection of a rigid volumetric object rotating in depth. Moreover, the region boundaries in the stimuli were all asymmetric, which makes it even more inconsistent with 3D rotation in depth. If an asymmetric rigid figure were rotating in depth, then the shape of the borders should change as the figure rotates. However, in our stimuli, the shapes of the borders were fixed and did not change during the presentation of the stimuli. Projective consistency is generally seen as an important factor in perceiving rigid 3D structure from motion. However, our results indicate that projective consistency might not play an important role in structure from motion as previously thought.

The stimuli used in Experiment 1 consisted of two regions separated by a common border, whereas in Experiment 2, the stimuli included eight regions. While the effect of convexity was very strong in both experiments, the effect of speed profile was observed to be stronger in the first experiment compared to the second one. It has been shown that the effect of convexity on figure/ground judgments becomes stronger as the number of regions in the stimuli is increased (Peterson & Salvagio, 2008). This might be the one of the reasons for the difference in the results of the two experiments. However, one important difference between the stimuli used in the Experiment 1 and Experiment 2 is the width of a single region. The width of a region in Experiment 1 was much larger than the width of a single region used in Experiment 2. When the width of the moving texture is larger, it becomes easier to identify what type of speed profile the texture has (i.e. constant vs. cosine). That is probably why the effect of speed profile of the texture was observed in all the subjects in Experiment 1, but not in Experiment 2.

We also built a simple Bayesian model in order to account for the finding of Experiment 1. The important feature of this model is that it does not use accretion/deleting as a definite cue to ground status. Instead, it checks the speed profile of the accreting/deleting texture in order to determine which interpretation (i.e. translating-in-behind or rotating-in-front) of accretion/deletion is more supported by this speed profile. Then, the model combines that information with the convexity of the border to assign probabilities to different depth order hypotheses. The model is able to account for the main effects of convexity and speed profile that were observed in Experiment 1. While the model performs well qualitatively, there are still certain discrepancies between the model's predictions and the psychophysical data. For all the subjects, the model underestimates the influence of convexity on depth order. The reason for this can be understood by examining the discrepancy between the joint statistics of convexity and depth obtained from natural image statistics and from the psychophysical experiments done by Burge et al. (2010). They found that the effect of convexity on relative depth judgments were found to be much stronger than the effect they observed in their natural image statistics. Similar to what Burge et al. (2010) observed, the comparison of our model predictions with the responses we obtained from Experiment 1 indicates a mismatch between the statistics of the natural images and the statistics internalized by the visual system with regards to convexity.

Moreover, our attempt to apply the convexity calculation method used Fowlkes et al. (2007) and Burge et al. (2010) to the stimuli used in Experiment 2 indicates a need for a more definite quantification of convexity in the literature. When the convexity values of the piecewise convex regions from Experiment 2 were calculated using the method by Fowlkes et al. (2007), negative values of convexity were found, which means that according to this convexity calculation methods, the piecewise convex regions were found to be less convex (or more concave) than its adjacent piecewise (perceptually) concave region. This result was found whether the convexity calculation method was applied locally or globally to the whole border. It is possible that a more robust quantification of convexity might be needed in order to close the gap between the joint statistics of convexity and depth in natural image analyses and psychophysical studies.

5. Conclusion

Since its introduction to the literature by Gibson et al. (1969) and Kaplan (1969), accretion/deletion has been an important part of accounts of depth from motion. For example, it is considered as one of the first visual cues to which infants respond to perceive relative depth (Granrud et al., 1984; Johnson & Mason, 2002). It has been even shown that it can override depth information from other important depth cues, such as lower-region cue (Royden et al., 1988), motion parallax (Ono et al., 1988; Hildreth & Royden, 2011; Yoonessi & Baker, 2013), and binocular disparity (Royden et al., 1988; Hildreth & Royden, 2011). Throughout, it has been always considered as a definite cue to ground status.

However, as it has been emphasized here often, accretion/deletion can also arise from self-occlusion due to rotation in depth. When accretion/deletion is interpreted as such, the perceived depth-order switches and the accreting/deleting region is interpreted as being figural. This alternative interpretation of accretion/deletion has been observed before in previous studies (Yonas et al., 1987; Royden et al., 1988), including the very first study that demonstrated accretion/deletion in the literature (Kaplan, 1969). In spite of that, this alternative interpretation of accretion/deletion has not been incorporated in the standard accounts of depth from motion. In this current work , we focused on this alternative interpretation of accretion/deletion, and on its implications for accretion/deletion as a cue to relative depth.

It has been shown that geometry of the boundary has an influence on how accretion/deletion is interpreted (Froyen et al., 2013; Tanrikulu et al., 2016). In order to further investigate this interaction in a more systematic way, in our first study (Chapter 2), we examined how the geometry of the contour (i.e. convexity) and accretion/deletion combine to determine depth-order. We manipulated the strength of both factors, and demonstrated that convexity has a very strong influence on how accretion/deletion is interpreted. We manipulated the strength of accretion/deletion by varying the rate of accretion/deletion in two different ways; either by manipulating the relative texture density or by manipulating texture speed. We found that accreting/deleting regions that have relatively higher density and speed were more likely to be seen at the back. However, when convexity and rate of accretion/deletion combine, the effect of convexity dominated the depth-order judgments. Further analysis of data revealed that the effect of texture density and speed could not simply be predicted by the relative rate of accretion/deletion. We proposed an alternative explanation for the effect of relative texture density and speed. The inconsistency between the constant-speed textural motion and a cosine speed profile (which is consistent with 3D rotation in depth) is greater when texture density or speed is high. Therefore, when texture density or speed is high, it is easier to perceive that the moving texture has a constant speed profile. Hence, that surface is more likely to be perceived as flat and translating in the back. Our alternative explanation for the effect of relative texture density and speed allowed us to account for the effect of rate of accretion/deletion that was demonstrated by Kaplan (1969), without considering accretion/deletion as definite cue to ground status.

In our first study, our results indicated that the geometry of the contour is critical to interpreting accreting/deleting regions in front. However, the displays we used were different than the classical accretion/deletion displays in a number of ways. The classical accretion/deletion displays have generally two regions separated by a straight border and only one side has accreting/deleting texture. However, our displays include multiple-region figure/ground stimuli where the shape of the border is not straight and in some cases accretion/deletion is introduced within both sides of each boundary. In our second study (Chapter 3), we investigated the factors that are most relevant to interpreting accreting/deleting regions in front by manipulating these three factors: number of regions, shape of the border, and whether one or both sides of each boundary contain(s) accretion/deletion. Our results again demonstrated that geometry of the contour is one of the most relevant factors for interpreting accreting/deleting regions in front. This result suggests that accounts of accretion/deletion must include contour geometry as a key component. Having accretion/deletion on both sides of a boundary was also found to be a relevant factor in the interpretation of accretion/deletion, which is expected since having accretion/deletion on both sides introduces more ambiguity to the figure/ground stimulus. However, as it was shown by Tanrikulu et al. (2016), having accretion/deletion on both sides of a boundary is not a necessary factor for perceiving accreting/deleting regions in front. Among these three factors we examined, the number of regions had the smallest effect in determining the interpretation of accretion/deletion.

While the results of our first two studies emphasize the need for depth-from-motion accounts to incorporate the geometry of the occluding contour, our results also suggest that the contour geometry has an influence on shape from motion processing. We have seen that perception of 3D rotating columns heavily depends on the the geometry of the border, and also occurs in spite of the constant-velocity texture motion. These observations raise the question about how the speed profile of texture combines with geometry of the boundary to determine depth-order. In order to answer this question, in the third and final study (Chapter 4), we manipulated the degree of convexity and the speed profile of the accreting/deleting texture (cosine, constant, or intermediate), and then combined them in various conditions to understand how these two factors determine relative depth. Our results demonstrated that the speed profile of the texture has a clear effect on perceived relative-depth. When the speed profile of the texture in a region is closer to a cosine speed profile, that region is more likely to be perceived in front. However, when convexity is introduced, it still dominates relative depth judgements. Accretion/deletion in the convex regions were interpreted as 3D rotating columns in front even in cases where the speed profile of the texture was grossly inconsistent with that interpretation. These results demonstrate that not only standard accounts of depth from motion, but also of shape from motion should incorporate the geometry of the border as a key component. Moreover, we also showed that the speed profile of the accreting/deleting texture has a strong influence on how accretion/deletion is interpreted, which in turn determines which depth order (i.e. front vs. behind) accretion/deletion would indicate.

We also built a probabilistic model to account for our findings from our last study. The model demonstrated that the main effect of convexity and speed profile of accreting/deleting texture on relative depth judgments can be accounted without considering accretion/deletion as a cue to being behind. However, we also observed a certain level of discrepancy between the model's prediction and subjects' responses, which had more general implications than the status of accretion/deletion cue. First, the discrepancy was (to some extent) due to the gap between the natural image statistics and psychophysical studies on the relationship between convexity and depth. This gap indicates a mismatch between statistics of our environment and the statistics internalized by our visual system. The assumption that our visual system internalizes the statistics of our environment is a very prominent and widespread assumption in vision science. This aforementioned gap that were observed in our study (as well as by Burge et al. (2010)) is inconsistent with this assumption. Secondly, when we applied the convexity calculation method proposed by Fowlkes et al. (2007) to the piecewise convex contours used in our last study, we obtained values that are grossly inconsistent with our phenomenology of convexity. This observation calls for a better definition and operationalization of convexity in the vision science literature.

Overall, our findings indicate that accretion/deletion is not a definite cue to ground status, but instead can be interpreted as either figure or ground depending on the geometry of the border at which the texture accretes/deletes. Traditional accounts and computational models of depth from motion do not take the geometry of the border into account when they are integrating accretion/deletion of texture into to their models. Moreover, our results not only urge future models of depth from motion to consider the geometry of the border as an important factor, but also urge future model of 3D structure from motion to take the geometry of the borders in their models.

6. Appendix

6.1 Results of Experiment 1

Table 6.1: Results of the logistic regression analysis for the main effect of place of accretion/deletion. The likelihood ratios (LR) reported here is between a model that includes the main effect of place of accretion/deletion and an unconditional-means model that contains only an intercept. This is done separately for each subject. The numbers reported below the likelihood ratio test are proportion of trials the subject perceive the reference region in front in the corresponding condition. “Single” refers to the condition in which only on side of each border has accretion/deletion, whereas ”Both” refers to the condition in which both sides of each border has accretion/deletion.

CL	**DA**	**DP**	**GS**	**JS**	**OE**	**TB**
LR=467.65	LR=194.96	LR=434.62	LR=196.35	LR=414.01	LR=271.11	LR=181.64
df=1, p<.001	df=1, p<.001	df=1, p<.001	df=1, p<.001	df=1, p<.001	df=1, p<.001	df=1, p<.001
single: 0.016	single: 0.231	single: 0.044	single: 0.218	single: 0.051	single: 0.099	single: 0.206
both: 0.630	both: 0.639	both: 0.657	both: 0.646	both: 0.637	both: 0.606	both: 0.634

Table 6.2: Likelihood ratio results of the multilevel logistic regression fits for the condition in which only one side of each border has accretion/deletion. Regression fits were done separately for each subject. Each factor was added one by one to an unconditional-means model that contains only an intercept. The factors were added in the order shown in the columns of the table, starting from the first column (i.e. border geometry) and ending at the interaction factor (last column). Only the results of the likelihood ratio tests that yielded a significant improvement are shown. The proportions of seeing the reference region in front are also presented under the results of the likelihood ratio test of the corresponding condition

	Border geometry	**Number of Regions**	**Color**	**Number of Regions * Border Geometry**
CL	LR=9.785, df=2, p<.01 Straight: 0 Unbiased: 0.007 Convex: 0.042		LR=4.099, df=1, p<0.05 Light: 0.028 Dark: 0.005	
DA	LR=90.01, df=2, p<.001 Straight: 0.055 Unbiased: 0.129 Convex: 0.5		LR=13.817, df=1, p<0.001 Light: 0.296 Dark: 0.167	LR=13.94, df=4, p<.05
DP				LR=15.321, df=4, p<.01
GS	LR=45.415, df=2, p<.001 Straight: 0.333 Unbiased: 0.295 Convex: 0.667	LR=7.726, df=2, p<.05 Two: 0.174 Four: 0.187 Eight: 0.292	LR=177.14, df=1, p<.001 Light: 0.426 Dark: 0.01	LR=17.261, df=4, p<.01
JS	LR=28.514, df=2, p<.001 Straight: 0.007 Unbiased: 0.014 Convex: 0.132	LR=21.271, df=2, p<.001 Two: 0 Four: 0.056 Eight: 0.097	LR=20.584, df=1, p<.001 Light: 0.093 Dark: 0.009	
OE		LR=6.519, df=2, p<.05 Two: 0.076 Four: 0.069 Eight: 0.153		
TB		LR=16.284, df=2, p<.001 Two: 0.319 Four: 0.146 Eight: 0.153	LR=13.261, df=1, p<.001 Light: 0.273 Dark: 0.139	LR=16.161, df=4, p<.05

Table 6.3: Likelihood ratio results of the multilevel logistic regression fits for the condition in which both sides of each border has accretion/deletion. Regression fits are done separately for each subject. Each factor is added one by one to an unconditional-means model that contains only an intercept. The factors are added in the order shown in the table columns starting from the first column (i.e. border geometry) and ending at the interaction factors (last column). Only the results of the likelihood ratio tests that yielded a significant improvement are shown. The proportions of seeing the reference region in front are also presented under the results of the likelihood ratio test of the corresponding condition

	Border geometry	**Color**	**Border geometry * Number of regions**	**Border geometry * Color**
CL	LR=38.464, df=2, p<.001 Straight: 0.535 Unbiased: 0.528 Convex: 0.826	LR=65.91, df=1, p<.001 Light: 0.458 Dark: 0.801	LR=41.826, df=4, p<.001	LR=10.036, df=2, p<.01
DA	LR=123.45, df=2, p<.001 Straight: 0.493 Unbiased: 0.458 Convex: 0.965	LR=23.639, df=1, p<.001 Light: 0.542 Dark: 0.736	LR=13.254, df=4, p<.05	LR=13.247, df=2, p<.01
DP	LR=70.01, df=2, p<.001 Straight: 0.521 Unbiased: 0.542 Convex: 0.910	LR=28.1, df=1, p<.001 Light: 0.546 Dark: 0.768		
GS	LR=85.173, df=2, p<.001 Straight: 0.528 Unbiased: 0.486 Convex: 0.924	LR=168.9, df=1, p<.001 Light: 0.903 Dark: 0.389		LR=7.712, df=2, p<.05
JS	LR=67.769, df=2, p<.001 Straight: 0.479 Unbiased: 0.542 Convex: 0.889			LR=7.712, df=2, p<.05
OE	LR=27.883, df=2, p<.001 Straight: 0.528 Unbiased: 0.514 Convex: 0.778	LR=5.016, df=1, p<.05 Light: 0.556 Dark: 0.657		LR=6.886, df=2, p<.05
TB	LR=64.32, df=2, p<.001 Straight: 0.576 Unbiased: 0.451 Convex: 0.875	LR=285.12, df=1, p<.001 Light: 755 Dark: 514		LR=8.866, df=2, p<.001

6.2 Results of Experiment 2

Table 6.4: Results of the logistic regression analysis for the main effect of place of accretion/deletion. The likelihood ratios (LR) reported here is between a model that includes the main effect of place of accretion/deletion and an unconditional-means model that contains only an intercept. This is done separately for each subject. The numbers reported below the likelihood ratios are proportion of trials the subject perceive the reference region in front in the corresponding condition. "Single" refers to the condition in which only on side of each border has accretion/deletion, whereas "Both" refers to the condition in which both sides of each border has accretion/deletion.

CL	**DA**	**DP**	**GS**	**JS**	**OE**	**TB**
LR=293.07	LR= 244.79	LR=487.06	LR= 330.9	LR= 352.61	LR= 157.79	LR=394.78
df=1, p<.001	df=1, p<.001	df=1, p<.001	df=1, p<.001	df=1, p<.001	df=1, p<.001	df=1, p<.001
single: 0.069	single: 0.157	single: 0.005	single: 0.102	single: 0.046	single: 0.227	single: 0.030
both: 0.590	both: 0.667	both: 0.632	both: 0.681	both: 0.609	both: 0.644	both: 0.616

Table 6.5: Likelihood ratio results of the multilevel logistic regression fits for the condition in which only one side of each border has accretion/deletion. Regression fits were done separately for each subject. Each factor was added one by one to an unconditional-means model that contains only an intercept. The factors are added in the order shown in the columns of the table, starting from the first column (i.e. border geometry) and ending at the interaction factor (last column). Only the results of the likelihood ratio tests that yielded a significant improvement are shown. The proportions of seeing the reference region in front are also presented under the results of the likelihood ratio test of the corresponding condition

	Border geometry	**Number of Regions**	**Color**	**Number of Regions * Border Geometry**	**Border Geometry * Color**
CL	LR=30.504, df=2, p<.001 Straight: 0.014 Unbiased: 0.028 Convex: 0.167	LR=20.904, df=2, p<.001 Two: 0.021 Four: 0.042 Eight: 0.146	LR=15.432, df=1, p<0.001 Light: 0.028 Dark: 0.111	LR=9.574, df=4, p<.05	
DA	LR=78.745 df=2, p<.001 Straight: 0.049 Unbiased: 0.042 Convex: 0.382	LR=38.717, df=2, p<.001 Two: 0.062 Four: 0.118 Eight: 0.292	LR=24.08, df=1, p<0.001 Light: 0.088 Dark: 0.227		
DP					
GS	LR=44.091, df=2, p<.001 Straight: 0.056 Unbiased: 0.014 Convex: 0.236	LR=13.408, df=2, p<.01 Two: 0.076 Four: 0.056 Eight: 0.174	LR=57.632, df=1, p<.001 Light: 0.009 Dark: 0.194		
JS	LR=22.306, df=2, p<.001 Straight: 0.007 Unbiased: 0.035 Convex: 0.104	LR=16.252, df=2, p<.001 Two: 0.021 Four: 0.014 Eight: 0.104	LR=17.156, df=1, p<.001 Light: 0.009 Dark: 0.083		LR=7.388, df=2, p<.05
OE	LR=40.619, df=2, p<.001 Straight: 0.097 Unbiased: 0.181 Convex: 0.403	LR=21.125, df=2, p<.001 Two: 0.111 Four: 0.257 Eight: 0.312	LR=10.554, df=1, p<.001 Light: 0.167 Dark: 0.287		
TB	LR=15.11, df=2, p<.001 Straight: 0.007 Unbiased: 0.007 Convex: 0.076		LR=11.631, df=1, p<.001 Light: 0.005 Dark: 0.083		LR=7.316, df=2, p<.05

Table 6.6: Likelihood ratio results of the multilevel logistic regression fits for the condition in which both sides of each border has accretion/deletion. Regression fits were done separately for each subject. Each factor was added one by one to an unconditional-means model that contains only an intercept. The factors are added in the order shown in the table columns, starting from the first column (i.e. border geometry) and ending at the interaction factors (last column). Only the results of the likelihood ratio tests that yielded a significant improvement are shown. The proportions of seeing the reference region in front are also presented under the results of the likelihood ratio test of the corresponding condition

	Border geometry	**Number of Regions**	**Color**	**Number of Regions * Border Geometry**	**Border Geometry * Color**	**Number of Regions * Color**
CL	LR=23.985, df=2, p<.001 Straight: 0.528 Unbiased: 0.493 Convex: 0.750		LR=69.802, df=1, p<.001 Light: 0.773 Dark: 0.407	LR=30.024, df=4, p<.001	LR=17.403, df=2, p<.001	LR=23.292, df=2, p<.001
DA	LR=150.75, df=2, p<.001 Straight: 0.507 Unbiased: 0.493 Convex: 1.0		LR=45.943, df=1, p<.001 Light: 0.796 Dark: 0.537		LR=24.946, df=2, p<.001	LR=12.823, df=2, p<.01
DP	LR=57.957, df=2, p<.001 straight:0.486 unbiased:0.542 convex:0.868		LR=125.68, df=1, p<.001 Light: 0.866 Dark: 0.398			LR=11.312, df=2, p<.05
GS	LR=72.986, df=2, p<.001 Straight: 0.549 Unbiased: 0.562 Convex: 0.931		LR=63.673, df=1, p<.001 Light:0.519 Dark:0.843		LR=12.841, df=2, p<.01	LR=12.362, df=2, p<.05
JS	LR=60.983, df=2, p<.001 Straight: 0.521 Unbiased: 0.451 Convex: 0.854			LR:19.995, df=4, p<.01	LR=13.627, df=2, p<.01	LR=6.898, df=2, p<.05
OE	LR=68.324, df=2, p<.001 Straight: 0.542 Unbiased: 0.493 Convex: 0.896				LR=8.968, df=2, p<.05	
TB	LR=52.919, df=2, p<.001 Straight: 0.451 Unbiased: 0.556 Convex: 0.840	LR=63.673, df=1, p<.001 Two: 0.514 Four: 0.667 Eight: 0.667	LR=68.609, df=1, p<.001 Light: 0.440 Dark: 0.792		LR=11.098, df=2, p<.01	LR=7.154, df=2, p<.05

References

Baigrie, B. (2002). The new science: Kepler, galileo, mersenne. In S. Nadler (Ed.), *A companion to early modern philosophy* (p. 45-59). Oxford: Wiley-Blackwell.

Barnes, T., & Mingolla, E. (2013). A neural model of visual figure and ground in dynamically deforming shapes. *Neural Networks*, *37*, 141-164.

Beck, C., Ognibeni, T., & Neumann, H. (2008). Object segmentation from motion discontinuities and temporal occlusion - a biologically inspired model. *PLoS ONE*, *3*(11), e3807.

Berzhanskaya, J., Grossberg, S., & Mingolla, E. (2007). Laminar cortical dynamics of visual form and motion interactions during coherent object motion perception. *Spatial Vision*, *20*(4), 337-395.

Brainard, D. (1997). The psychophysics toolbox. *Spatial Vision*, *10*(4), 433-436.

Braunstein, M. L., Andersen, G. J., & Riefer, D. M. (1982). The use of occlusion to resolve ambiguity in parallel projections. *Perception & Psychophysics*, *31*(3), 261-267.

Burge, J., Fowlkes, C. C., & Banks, M. S. (2010). Natural scene statistics predict how the figure-ground cue of convexity affects human depth perception. *Journal of Neuroscience*, *30*(21), 7269-7280.

Fowlkes, C. C., Martin, D. R., & Malik, J. (2007). Local figure-ground cues are valid for natural images. *Journal of Vision*, *7*(8), 1-9.

Froyen, V., Feldman, J., & Singh, M. (2010). A Bayesian framework for figure-ground interpretation. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, & A. Culotta (Eds.), *Advances in neural information processing systems 23* (Vol. 631-639). Vancouver, British Columbia, Canada: Curran Associates.

Froyen, V., Feldman, J., & Singh, M. (2013). Rotating columns: relating structure-from-motion, accretion/deletion, and figure/ground. *Journal of Vision*, *13*(10).

Gibson, J. (1966). *The senses considered as perceptual systems.* Boston: Houghton Mifflin.

Gibson, J., Kaplan, G., Reynolds, H., & Wheeler, K. (1969). The change from visible to invisible: A study of optical transitions. *Perception & Psychophysics*, *5*, 113-116.

Goldstein, E. B. (2014). *Sensation and perception* (9th ed.). Belmont, CA: Cengage Learning.

Granrud, C. E., Yonas, A., Smith, I. M., Arterberry, M. E., Glicksman, M. L., & Sorknes, A. C. (1984). Infants' sensitivity to accretion and deletion of texture as information for depth at an edge. *Child Development*, *55*(4), 1630-6.

Hegdé, J., Albright, T., & Stoner, G. (2004). Second-order motion conveys depth-order information. *Journal of Vision*, *4*, 838-842.

Hildreth, E., & Royden, C. (2011). Integrating multiple cues to depth order at object boundaries. *Attention, Perception, & Psychophysics*, *73*(7), 2218-2235.

Hoffman, D., & Singh, M. (1997). Salience of visual parts. *Cognition*, *63*(1), 29-78.

Howard, I., & Rogers, B. (2002). *Seeing in depth vol.2 depth perception.* I. Porteous, Toronto.

Johnson, S. P., & Mason, U. (2002). Perception of kinetic illusory contours by two-month infants. *Child Development*, *73*, 22-34.

Kanizsa, G., & Gerbino, W. (1976). Convexity and symmetry in figure-ground organization. In M. Henle (Ed.), *Vision and artifact* (p. 25-32). New York: Springer.

Kaplan, G. (1969). Kinetic disruption of optical texture: The perception of depth at an edge. *Attention, Perception, & Psychophysics*, *6*(4), 193-198.

Kleiner, M., Brainard, M., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What is new in psychtoolbox 3. *Perception*, *36*(14 ECVP Abstract Suppl.).

Kromrey, S., Bart, E., & Hegdé, J. (2011). What the 'moonwalk' illusion reveals about the perception of relative depth from motion. *PLoS ONE*, *6*(6), e20951. doi:10.1371/journal.pone.0020951.

Layton, O. W., & Yazdanbakhsh, A. (2015). A neural model of border-ownership from kinetic occlusion. *Vision Research*, *106*, 64-80.

Metzger, F. (1936/2006). *Laws of seeing.* (L. Spillmann and S. Lehar and M Stromeyer and M. Wertheimer, Trans.): Massachusetts Institute of Technology, Cambridge, MA (Original work published 1936).

Michotte, A., Trines, G., & Crabbe, G. (1964). *Les complements amodaux des structures perceptives (studia psychologia).* Louvain: Publications Universitaires de Louvain.

Morinaga, S. (1941). Beobachtungen über grundlagen und wirkungen anschaulich gleichmäßiger breite. *Archiv für die gesamte Psychologie*, *110*, 309-348.

Mutch, K., & Thompson, W. (1985). Analysis of accretion and deletion at boundaries in dynamic scenes. *Pattern Analysis and Machine Intelligence, IEEE Transactions*, *2*, 133-138.

Niyogi, S. (1995). *Detecting kinetic occlusion.* Massachusetts Institute of Technology, Cambridge, MA.

Ono, H., Rogers, B., Ohmi, M., & Ono, M. (1988). Dynamic occlusion and motion parallax in depth perception. *Perception*, *17*(2), 255-266.

Peterson, M. A., & Salvagio, E. (2008). Inhibitory competition in figure-ground perception: Context and convexity. *Journal of Vision*, *8*(16), 4. Retrieved from + `http://dx.doi.org/10.1167/8.16.4` doi: 10.1167/8.16.4

Profitt, D. R., Bertenthal, B. I., & Roberts, R. J. (1984). The role of occlusion in reducing multistability in moving point-light displays. *Perception & Psychophysics*, *36*(4), 315-323.

Ramachandran, V. S., Cobb, S., & Rogers-Ramachandran, D. (1988). Perception of 3-d structure from motion: The role of velocity gradients and segmentation boundaries. *Attention, Perception, & Psychophysics*, *44*(4), 390-393.

Raudies, F., & Neumann, H. (2010). A neural model of the temporal dynamics of figure-ground segregation in motion perception. *Neural Networks*, *23*(2), 160-176.

Royden, C. S., Baker, J. F., & Allman, J. (1988). Perception of depth elicited by occluded and shearing motions of random dots. *Perception*, *17*, 289-296.

Rubin, E. (1915/1958). Figure and ground. In D. Beardslee & M. Wertheimer (Eds.), *Readings in perception* (p. 194-203). Princeton, NJ: Van Nostrand (Original work published 1915).

Ruda, H., Livitz, G., Riesen, G., & Mingolla, E. (2015). Computational modeling of depth ordering in occlusion through accretion or deletion of texture. *Journal of Vision*, *15*(9), 20.

Tanrikulu, O. D., Froyen, V., Feldman, J., & Singh, M. (2014). A cue interaction study of depth ordering from accretion/deletion and contour convexity. *VSS Conference Presentation*.

Tanrikulu, O. D., Froyen, V., Feldman, J., & Singh, M. (2015). Bridging the gap between standard accretion/deletion and rotating columns. *VSS Conference Presentation*.

Tanrikulu, O. D., Froyen, V., Feldman, J., & Singh, M. (2016). Geometric figure–ground cues override standard depth from accretion-deletion. *Journal of Vision*, *16*(5), 1-15.

Thompson, W., Kersten, D., & Knecht, W. R. (1992). Structure-from-motion based on information at surface boundaries. *Biological Cybernetics*, *66*(4), 327-333.

Thompson, W., Mutch, K., & Berzins, V. (1985). Dynamical occlusion analysis in optical flow fields. *Pattern Analysis and Machine Intelligence, IEEE Transactions*, *4*, 374-383.

Thompson, W., & Painter, J. S. (1992). Qualitative constraints for structure-from-motion. *CVGIP: Image Understanding*, *56*(1), 69-77.

Ullman, S. (1979). The interpretation of structure from motion. In *Proceedings of the Royal Society of London* (Vol. 203, p. 405-426). The Royal Society.

Vecera, S. P., Vogel, E. K., & Woodman, G. F. (2002). Lower region: a new cue for figure-ground assignment. *Journal of Experimental Psychology: General*, *131*(2), 194-205.

Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., & von der Heydt, R. (2012). A century of gestalt psychology in visual perception: 1. perceptual grouping and figure-ground organization. *Psychological Bulletin*, *138*(6), 1172-217.

Yonas, A., Craton, L., & Thompson, W. (1987). Relative motion: Kinetic information for the order of depth at an edge. *Attention, Perception, & Psychophysics*, *41*(1), 53-59.

Yoonessi, A., & Baker, C. L., Jr. (2013). Depth perception from dynamic occlusion in motion parallax: Roles of expansion-compression versus accretion-deletion. *Journal of Vision*, *13*(12), 10. Retrieved from + `http://dx.doi.org/10.1167/13.12.10` doi: 10.1167/13.12.10