Dynamic texture spreading: probing the mechanisms of surface interpolation

DANIEL WOLLSCHLÄGER * and FRANZ FAUL

Institut für Psychologie Christian–Albrechts–Universität zu Kiel, Olshausenstraße 62 24098 Kiel, Germany

Abstract—Dynamic texture spreading is a filling-in phenomenon where a colored pattern perceptually spreads onto an area confined by virtual contours in a multi-aperture motion display. The spreading effect is qualitatively similar to static texture spreading but widely surpasses it in strength, making it particularly suited for quantitative studies of visual interpolation processes. We first carried out two experiments to establish with objective tasks that texture spreading is a genuine representation of surface qualities and thus goes beyond mere contour interpolation. Two subsequent experiments serve to relate the phenomenon to ongoing discussions about potentially responsible mechanisms for spatiotemporal integration. With a phenomenological method, we examined to what extent simple sensory persistence might be causally involved in the effect under consideration. Most of our findings are consistent with the idea of sensory persistence, and indicate that information fragments are integrated over a time window of about 100 to 180 ms to form a complete surface representation.

Keywords: Texture spreading; surface completion; filling-in; perceptual grouping; multi-aperture viewing.

1. INTRODUCTION

This article introduces dynamic texture spreading (Fig. 3) [see Note 1], a spatio-chromatic interaction observable with a new variant of multi-slit view displays (see Section 2.1). Not only does it comprise the formation of strong subjective contours but also the vivid illusory perception of surface qualities like color, texture, and depth stratification. It is closely related to other cases of perceptual completion and bears direct relation to color spreading. It is thus helpful to first consider the well-known effects of static color and texture spreading, and in turn introduce dynamic texture spreading itself.

---

*To whom correspondence should be addressed. E-mail: dwoll@psychologie.uni-kiel.de

D. Wollschläger and F. Faul

Color spreading demonstrates how certain spatial configurations can give rise to the perception of surface qualities, such as color or texture, attributed to stimulus regions that are themselves devoid of physical information about these features. Figure 1 shows two recently introduced variants of this effect (Wollschläger *et al*., 2001, 2002). Here, color is seen not only in the physically-colored parts of the display, but also in the regions between them, giving rise to a lightly-colored surface that is bounded by illusory contours. Color spreading can be considered a case of perceptual filling-in (Pessoa *et al*., 1998), a mechanism that has been associated with, for example, the slow peripheral veiling of figures by their backgrounds during prolonged fixation (Troxler fading, Pirenne, 1962; Troxler, 1804), instantaneous filling-in of blind spots and scotomas (Ramachandran, 1992; Ramachandran and Gregory, 1991), neon spreading (Bressan *et al*., 1997), long range color spreading (Pinna *et al*., 2001), texture spreading (Watanabe and Cavanagh, 1991), and color from motion (Cicerone *et al*., 1995) [see Note 2].

Neon color spreading (Fig. 2a) can typically be observed in line drawings where the color of some segments is abruptly changed to a different one. The color of the embedded segments then seems to disperse onto an area confined by virtual contours, producing a subtle, neon-like veil. Some neon-color-spreading displays can be altered to elicit texture spreading when the uniformly-colored embedded line segments leading to neon spreading are replaced with, for example, segments of a stripe texture (Watanabe and Cavanagh, 1991). In these stimuli, a *non-uniform* color pattern perceptually fills in a homogeneous background region (Fig. 2b), again bounded by subjective contours. "Visual phantoms" (Tynan and Sekuler, 1975; Weisstein *et al*., 1977) may also be taken as a rudimentary example of texture spreading in static as well as in motion displays, but they occur only with a very restricted class of one-dimensional textures.

While the aforementioned types of filling-in are highly relevant to the study of completion processes and of interactions between spatial stimulus layout and perceived surface qualities, all of them are somewhat vague and perceptually unstable. This elusive character makes it difficult to reliably study them using quantitative methods, such as scaling or matching procedures, and even more so using objective performance tasks.

Figure 1. (See color plate VII) (a) Static flank-transparency demonstration (Wollschläger *et al*., 2002), including the two-color case. No spreading obtains if the occluding bars are missing. (b) Clipping of one frame from an animated flank-transparency display. The line array is slightly shifted to the right between frames, whereas the virtual ellipse stays put. Small colored flanks are added to the lines in the interior of the virtual ellipse.

Figure 2. (See color plate VII) (a) Neon color spreading in the modified Ehrenstein cross. Color spreading is absent if the embedded segments are not aligned with the inducing lines. (b) Texture spreading in the modified Ehrenstein cross. Texture spreading also vanishes with nonaligned segments.

We believe that our dynamic stimulus changes this situation for texture spreading as it leads to a stable, clearly-defined dispersion of a colored texture. As such, it may provide a useful tool to the experimental study of texture spreading as some recently presented effects did to the study of color-spreading: the watercolor illusion (Devinck *et al.*, 2005; Pinna *et al.*, 2001), color from motion (Cicerone *et al.*, 1995; Ekroll and Faul, 2002), and flank transparency (Wollschläger *et al.*, 2001, 2002) all exhibit much better discernible spreading effects than classical neon-spreading demonstrations.

2. DYNAMIC TEXTURE SPREADING

Dynamic texture spreading can be observed in animated computer displays that simulate a multi-slit view setup (Fahle and Poggio, 1981; Mateeff *et al.*, 1993; Nishida, 2004). These are a generalization of Zöllner's (1862; Rock, 1981) anorthoscopic view apparatus from one to multiple apertures. Suppose you put a disk covered with a stripe pattern on a white background. In front of it, you place a white sheet of paper with black vertical lines painted on it. Seamlessly alongside the lines, you cut narrow rectangular slits, making visible small strips of the background and the textured disk behind (Fig. 3). The animation procedure may be such that the front sheet of paper is translated gradually on a horizontal path, accompanied by a corresponding change in what is visible of the stationary textured shape through the apertures flanking the lines. Alternatively, one may hold the front sheet of paper in a fixed position and horizontally move the striped disk back and forth behind the stationary apertures while tracking the disk with the eyes. At a suitable speed and line density, both animations lead to perceptual integration of non-contiguous image
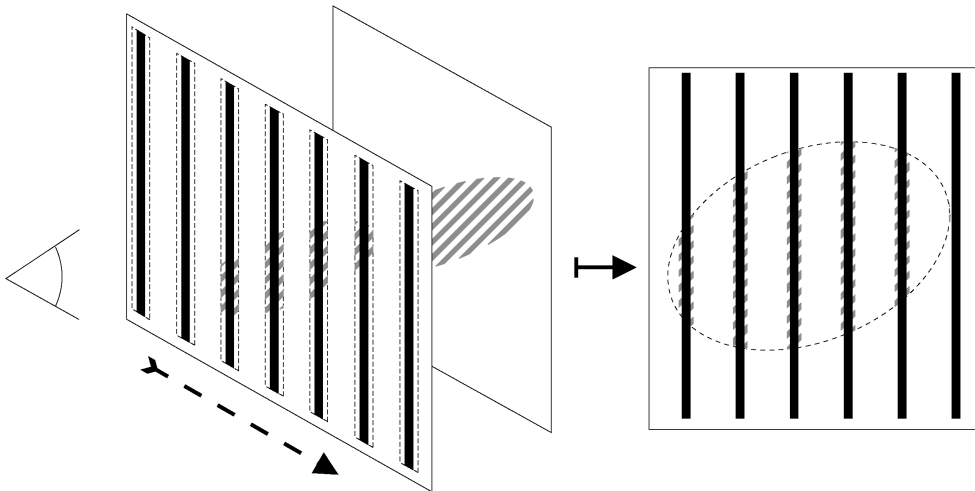

Figure 3. Schematic construction of a dynamic-texture-spreading display (see text). Small striped flanks are added to the lines in the interior of the virtual ellipse.

regions; one has a vivid illusory impression of colored texture in an area defined by crisp subjective contours tracing the outline of the disk.

The described setup differs from standard multi-aperture viewing stimuli in that the latter may have visible or invisible apertures, but they do not incorporate visible lines *between* them. This seemingly subtle distinction has important consequences. Flank transparency generalizes also to static displays (Fig. 1a), allowing for comparisons between static and dynamic perceptual completion with nearly identical stimuli. In addition, it can evoke transparent color spreading even in two-color displays (Fig. 1) — an observation that is currently not handled by models of perceptual transparency (Wollschläger *et al.*, 2001).

2.1. Relation to other phenomena

If the disk is of solid color, the virtual shape also appears uniformly colored, an effect we called flank transparency (Fig. 1b, Wollschläger *et al.*, 2001, 2002) because of the percept's translucent quality, which is similar to the one reported for static neon spreading (Bressan *et al.*, 1997). Flank transparency shares further characteristics with neon spreading, such as a concomitant depth stratification of objects, leading to the perception of the virtual shape as being in front of or behind the background layer. As such, static neon color spreading and texture spreading both have counterparts in the domain of dynamic stimuli, i.e. flank transparency and dynamic texture spreading — each one qualitatively similar but more robust compared to their static analogue.

Filling-in phenomena demonstrate that spatially fragmented input information to the visual system can interact to trigger the perception of complete objects that are defined along the same interdependent perceptual dimensions, such as form, color, motion, and depth, as are "real" objects. In neon color spreading, percepts of glowing and lightly transparent layers lying above the background form in response to only a few embedded segments in a line drawing. This discrepancy between meagre input information available to the visual system in each single frame and a rich final percept is even larger in dynamic texture spreading. In this case, not only the contour but also the object's non-uniform inner structure must be recovered. Yet, in the animated display, it is possible to discern even complex surface layouts like, e.g. a colored Gabor patch (see Section 3.2).

Just as the described stimulus setup is a variant of multi-aperture viewing, it also resembles dynamic occlusion displays using a textured background (Andersen and Cortese, 1989; Bruno and Bertamini, 1990; Hine, 1987; Shipley and Kellman, 1994). These stimuli contain only highly incomplete information about simple objects as they simulate the dynamic accretion and deletion of background texture elements by a moving, and usually invisible mask — the front sheet of paper in the case of dynamic texture spreading.

Even though aperture-viewing and dynamic occlusion displays have been widely used to study perceptual interpolation processes, they appear to form mostly disjoint literature domains (for an exception, see Shipley and Cunningham, 2001). So far,

aperture viewing seems to draw most interest from research into motion perception (Anstis and Atkinson, 1967; Morgan, 1980; Morgan *et al.*, 1982; Parks, 1965; Rock *et al.*, 1987), while dynamic occlusion is more associated with studies of object recognition and contour interpolation (Kellman and Shipley, 1991; Kellman *et al.*, 2001).

Investigations in *both* of the aforementioned research contexts have so far been almost exclusively concerned with aspects of boundary integration. Studies in dynamic occlusion that explicitly addressed a surface component (e.g. Cunningham *et al.*, 1998) focused on the perception of a surface quality *per se*, and not on the formation details of particular modal features associated with that surface, like color, texture or degree of transparency. Visual inspection suggests that all these play a vital role in dynamic texture spreading. Our first two experiments thus serve to establish with objective methods that dynamic texture spreading is indeed the result of a genuine perceptual process that goes beyond mere contour completion, and represents a functionally relevant interpolation of surface features.

2.2. Sensory persistence as a possible explanation for visual interpolation

Many of the cited studies on boundary completion tested whether simple sensory persistence or "retinal painting" is a possible explanation for the observed discrepancy between the fragmented stimulus material and the perception of extended contours. Retinal painting asserts that sensory information persists in some form of retinal or post-retinal storage and can thus be integrated by a simple local mechanism over a certain time-frame. If, for example, in the case of aperture viewing, stimulus fragments are presented such that the complete stimulus is subsequently spread over the retina over time, such temporal integration might be able to automatically recover the extended stimulus just by means of sensory persistence. No complex form of shape interpolation would then be required.

Some authors found strong support for a causal role of retinal painting in the perception of outline shapes under aperture-viewing conditions. Anstis and Atkinson (1967) as well as others (for a review, see Morgan *et al.*, 1982; Shipley and Cunningham, 2001) have demonstrated that the shape distortions observed in aperture viewing closely follow the predictions made by retinal painting.

However, other authors claimed that retinal painting is neither necessary nor sufficient to explain the observed phenomena (Bruno and Gerbino, 1991; Fendrich *et al.*, 2005; Rock *et al.*, 1987; Shimojo and Richards, 1986). Already Zöllner (1862), von Vierordt (1868), and Rothschild (1922) tried to show that anorthoscopic perception can occur even when a steady fixation prevents the image of the object from being spread over the retina over time. With modern eye-tracking equipment, Fendrich *et al*. (2005) recently extended that work and found that boundary integration in anorthoscopic perception can still take place with retinally stabilized images. Anstis (2005) reported that anorthoscopic perception in multi-slit viewing also occurs with rotational movement, which is hard to track because the eyes cannot fully rotate around the axis perpendicular to the sagittal plane. Based on motion-

direction specific noise masking, Nishida (2004) argued that high-level motion analysis is involved in the perception of patterns viewed through multiple slits, and concludes that simple sensory persistence does not constitute a possible underlying mechanism.

A relevant theoretical issue was pointed out by Dixon and Di Lollo (1994). On the one hand, it is necessary for the visual system to integrate stimuli over time to guarantee perceptual continuity; but on the other hand, it must also detect rapid changes and must therefore segregate those stimulus elements that do not belong to one spatiotemporal object. Retinal painting by itself does not provide a mechanism for distinguishing coextensive fragments in the stream of information from disjoint ones, i.e. for deciding when to integrate and when to segregate.

It appears reasonable to conclude that sensory persistence may be only one among several, and possibly more important, mechanisms involved in boundary integration. Indeed, Shipley and Cunningham (2001) explicitly propose the existence of two different representations with different processing mechanisms. They cite evidence from shape perception tasks supporting on the one hand a retinotopic stimulus representation (here referred to as retinal painting), and on the other hand an additional distal representation with an external frame of reference. Yet, this conclusion does not automatically extend to the integration of surface attributes. Surface filling-in may constitute a different perceptual process that may therefore be mediated by different computational mechanisms. Whether surface completion involves or even completely relies on sensory persistence has not yet been explored. Considering that visual completion phenomena encompass a multitude of complex aspects regarding the reconstruction of boundary and surface information, it seems worthwhile to try to separate them and individually test which mechanisms might be responsible for each one. Retinal painting seems a very simple and natural candidate to explain some characteristics of dynamic texture spreading. Therefore, we take sensory persistence as a first starting point in our investigation on dynamic texture spreading. More specifically, one motivation for our experiments was the question whether retinal painting provides a possible explanation for the interpolation of a surface feature like texture.

2.3. Retinal painting applied to dynamic color and texture spreading

Sensory persistence asserts that a complete surface representation can be sustained by discrete, scattered signal elements provided that each point of the represented surface is frequently enough supported by new input signals consistent with that surface. Applied to our stimulus setup, this means that a complete textured object would be perceived as long as the information signaling the presence of a piece of texture is refreshed often enough at each fixed location, regardless of how this refresh frequency is achieved. As the refresh frequency depends on the density and speed of the moving rectangular apertures through which object information becomes available, these two stimulus parameters should be important determinants of the texture-spreading effect.

This consideration is backed by empirical results about the relevance of signal element density for object recognition in animated sparse displays. Shipley and Kellman (1993; Cunningham *et al.*, 1998) found that shape identification improved with more signal elements in their displays used in research on spatiotemporal boundary formation. This result was also reported various times for experiments on dynamic occlusion (Andersen and Cortese, 1989; Bruno and Bertamini, 1990; Bruno and Gerbino, 1991; Shipley and Kellman, 1994) as well as by Mateeff *et al.* (1993) and Bull *et al.* (2003) for multi-aperture stimuli. Fidopiastis *et al.* (2000) showed that the strength of color from motion is modulated by the particular layout of the dot array which provides signal elements for form and color information. Specifically, it was established that the color-spreading effect increases with a denser and more regular dot array. For color from motion (Cicerone *et al.*, 1995) and for flank transparency (Wollschläger *et al.*, 2002), a close dependence on the amount of motion has been observed.

It therefore seems natural to analyze dynamic texture spreading by asking about the specific interplay between the density of the line array and its speed in determining the filling-in effect. This will be the starting point for Experiments 3 and 4.

3. EXPERIMENTS

With Experiments 1 and 2, we intend to show that texture spreading is a functionally relevant effect that involves more complex processing and provides more information than contour interpolation. To this end, we used highly fragmented stimuli in performance tasks whose successful completion required not only boundary information but, ideally, a complete surface representation. Experiments 3 and 4 then derive quantitative information about the texture interpolation process from a phenomenological task and relate the findings to the concept of sensory persistence.

3.1. Experiment 1

In Experiment 1, we used a 2AFC discrimination task to objectively measure the degree to which texture spreading can help to successfully recover the texture composition of a fragmented surface. The stimuli were designed not to contain local features that would be diagnostic for their discrimination. Instead, we wanted to be able to infer from a successful task completion the full perceptual presence of the surface layout. We therefore tried to construct stimuli whose identification rests on features that become apparent only in full view. This, however, proved to be no trivial goal. In particular, two commonly used stimuli, meaningless line drawings and letters, are inappropriate. The former lacks any surface qualities at all, and the latter contains — by the very virtue of its construction — many local cues which aid in its discrimination.

The aforementioned issue reminds us of a general problem when trying to draw inferences from results in performance tasks as they are now commonly used in

the study of perceptual completion: in contrast to purely introspective reports, these tasks often have the advantage of separating genuine perceptual aspects from cognitive reasoning influences (Kellman *et al.*, 2001). But while they can provide an objective measure of the level of some functional relevance of a subjective percept, inferences concerning the qualities of that percept are indirect and based on certain assumptions: Shape discrimination tasks used in studies on boundary interpolation often implicitly presume that a fully completed contour was at the observer's disposal for making a successful shape categorization. To what extent shapes are also discriminable based on local features is seldom discussed. The development of adequate stimulus material thus remains a crucial prerequisite for the study of surface interpolation.

3.1.1. Subjects. Five students from the University of Kiel (two female, three male) participated in Experiment 1, all were naïve as to the purpose of the experiment. We carried out a separate control experiment with five subjects, among them one (JJ) who also participated in the experimental condition, and one of the authors (DW). All subjects had normal color vision and normal or corrected-to-normal visual acuity.

3.1.2. Apparatus and stimuli. The stimuli were displayed on a Sony Multiscan 500 PS 21 inch color monitor with a vertical refresh rate of 100 Hz using the OpenGL graphics library. The timing of the stimulus frames was synchronized with the vertical refresh of the monitor. Color resolution for each pixel was 8 bits for each channel R, G, B. Following a standard procedure (Brainard, 1995), the monitor was calibrated and gamma corrected using a colorimeter (LMT C1210). Observers placed their heads on a chin rest and binocularly viewed the display in a darkened room at a distance of 80 cm.

The background of the display was grey (CIE $x = 0.30$, $y = 0.31$, $L = 38 \, \mathrm{cd \, m^{-2}}$) and subtended a visual angle of 17.4 by 13.3 deg. An array of vertically-oriented black lines (CIE $x = 0.30$, $y = 0.31$, $L < 0.2 \, \mathrm{cd \, m^{-2}}$), 6.75 min of arc wide and 10 deg long, was placed on the background.

The target stimulus was placed in the interior of a virtual disk with a radius of 2.9 deg. Throughout the experiment, only thin stripes (1.35 min of arc wide) of the target stimulus became visible as flanks alongside the black vertical lines.

The target stimuli were non-meaningful images (Fig. 4) which were constructed in the following way: First, the disk was completely covered with randomly placed circles of varying radii. Next, we defined three invisible, non-intersecting Bézier curves with a thickness of 0.25 deg which started and ended at the disk's boundary. As a final step, we colored those circles whose center lay outside the curves green

Figure 4. (See color plate VIII) Examples of the two stimulus classes used in Experiment 1. Circles are colored red if their center falls within one of three imaginary Bézier curves, else green. (a) Target stimulus with a continuous line. (b) Target stimulus without a continuous line.

(CIE $u = 0.16$, $v = 0.47$, $L = 12$) and then those whose center fell within one of the curves red (CIE $u = 0.24$, $v = 0.45$, $L = 12$). Gaussian red-green chromaticity noise as well as Gaussian luminance noise was added to each circle's color to make the curves' boundaries less detectable by simple local edge operators.

Subjects had to discriminate between two classes of the described stimulus. In one class, all three Bézier curves were interrupted once within the disk such that a gap would separate two parts of the curve. In the other class, only two curves were interrupted such that there remained one continuous curve without a gap. The rationale behind the stimulus construction was that in order to verify the presence of a continuous line, the observer must have a complete (i.e. gapless) perceptual representation of that line. The two classes should not be distinguishable based on local features. Since both classes contained continuous line segments as well as gaps between segments, the detection of a continuous segment or of a gap did not suffice to make a reliable discrimination. In addition, the number of circles within the curves did not vary systematically between classes, such that the total amount of redness in a stimulus was not diagnostic for a correct classification.

We designed five different stimuli for each stimulus class. In order to reduce the chance of recognition of a particular exemplar, the target stimulus was randomly rotated on each trial.

We used two "motion conditions": in the first one (subjects JJ, DA), the line array rigidly moved to the right at a speed of 1.13 deg s^{-1}, whereas the virtual target stimulus remained stationary at the center of the screen. In the second one (subjects TG, MH, KR), the line array was stationary, and the virtual target stimulus moved at a speed of 1.13 deg s^{-1}. Subjects were instructed to track the target which made the second condition equivalent to the first one with respect to the resulting retinal image (see Section 2). In the control experiment, neither lines nor the virtual disk moved, providing a static reference case which indicated the achievable performance without temporal integration.

3.1.3. Procedure. We varied the spacing of the vertical lines from 0.25 to 1.1 deg, corresponding to line densities of 80 to 16 lines visible. The dependent variable was the correctness of the subject's answer to the question whether the target stimulus contained a continuous line or not. Each target stimulus class had an equal probability of being chosen on a given trial, as was the case for all five exemplars within one such class.

The spacing of the vertical lines was varied according to a custom adaptive method. We used this in order to minimize the number of trials necessary to reliably reconstruct the function relating independent and dependent variable for a given subject. Unlike methods such as QUEST (Watson and Pelli, 1983), we did not assume any characteristics of the function other than monotonicity, and wanted to reconstruct the function over the whole range of the independent variable, as opposed to only estimating the parameters of a fixed function. Details of the employed algorithm are given in Appendix A.

Each stimulus presentation lasted for 8.4 seconds, after which subjects pressed one of two cursor keys to indicate their classification of the target stimulus. This keystroke automatically started the next trial. Subjects were encouraged to make pauses at their own disposal. Prior to the experiments, all subjects familiarized themselves with the task in 20 training trials which were not included in the analysis. It took subjects from 260 (TG) to 433 (MH) trials to complete the described procedure.

3.1.4. Results. Figure 5a displays the individual results for subject JJ in the motion condition as well as in the static control. The two motion conditions are equivalent with respect to the retinal image and, as expected, did not produce systematically different results. The same was true for the data of individual subjects, which warranted their pooling in the two motion conditions and the static control, respectively. Figure 5b shows the averaged data. For plotting the data, we divided the full parameter range of the independent variable into intervals. The data points represent the mean percent correct of the subject's classification in these intervals. Both curves were visually fitted with psychometric functions [see Note 3].

The discrimination performance is inversely linked to line distance in the motion as well as in the control condition. It continuously varies from near perfect discrimination at low distances to about chance performance at high distances. In the static control condition, the drop in performance with wider line spacing happens faster than in the motion condition, the difference becoming more pronounced with line spacings wider than 0.56 deg. At a line distance of, for example, 0.67 deg, the performance of JJ in the motion condition is still at nearly 90% correct, whereas in the static control condition, it has already dropped to about 65%. For the averaged data, mean performance in the static condition at a line distance of, for example, 0.72 deg is 61% correct, while in the motion condition it is still 79%. The relative performance benefit of the motion conditions over the static control is plotted in Fig. 5c for the averaged data as well as for the fitted psychometric functions. Mean discrimination was up to 30% better in the motion condition than in the static control.

3.1.5. Discussion. The obtained data show that discrimination performance for a fragmented texture pattern objectively benefits from spatiotemporal integration. The performance difference to a static control condition may be seen as a measure for the functionally relevant amount of surface feature information that is gained through integration. The results also imply that the integration of a complex texture does not simply lead to its spatiotemporal averaging or blending into a uniform color, thus setting dynamic texture spreading apart from homogeneous color spreading.

The stimuli were constructed such that the necessary information for a successful discrimination could not come from simple contour integration alone, as was the case with most discrimination tasks reported in the literature, which used line drawings or letters. We were thus able to show that spatiotemporal integration of a
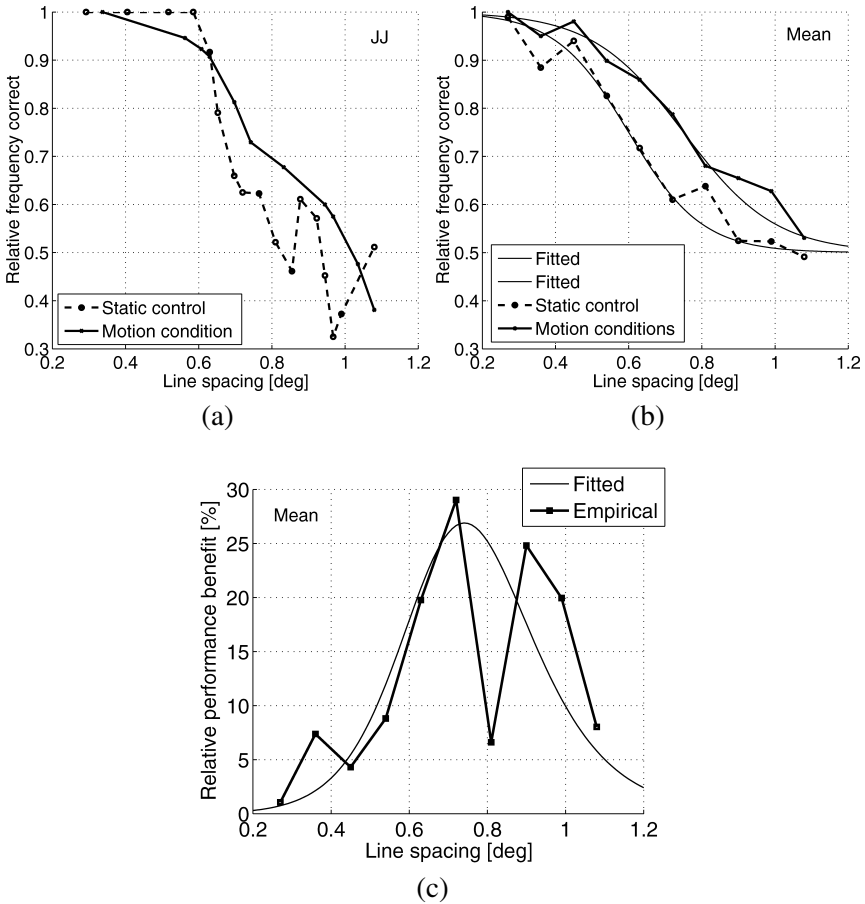
Figure 5. Results from Experiment 1. (a) Individual data from JJ in the motion as well as in the static control condition. (b) Combined results averaged over subjects in the motion as well as in the control conditions. The data was visually fitted with psychometric functions (see text). (c) Mean relative performance benefit of the motion conditions over the static control for the empirical data as well as for the fitted psychometric functions.

surface feature with nonlocal properties, such as a complex texture, can indeed take place.

Our goal was to design a task that, to be successfully accomplished, requires a *complete* surface representation available to the visual system. However, as is evident from the static control condition, information relevant to the discrimination task was, up to a certain degree, also available with incomplete stimuli. One explanation for this finding is that with a line distance smaller than the gaps within a Bézier curve, detection of gaps becomes trivial. Gap information is then available in spatial frequencies low enough to be fully recovered with the available spatial sampling frequency without the need for temporal integration. Another explanation may be required to account for the relatively smooth performance drop

with increasing line spacing in the static control. It is possible that diagnostic local features were still present in the target stimuli to a certain extent. This illustrates the difficulty in designing appropriate stimuli for the objective assessment of the degree to which a complete object representation has formed. Nevertheless, the fact that spatiotemporally integrating the stimulus leads to a measurable performance gain, once line distance becomes larger than curve gaps, means that non-local surface feature information were relevant and could be recovered in the integration process.

3.2. Experiment 2

In Experiment 2, we wanted to use a matching procedure to obtain another objective measure of the subjects' ability to perceptually integrate a fragmented stimulus. To this end, we had subjects match the frequency of a vertically oriented Gabor patch to that of a test Gabor which was shown in a multi-slit view setup with visible slits. We chose a Gabor patch as it constitutes a textured surface without sharp boundaries, thereby reducing the diagnostic value coming from simple boundary integration mechanisms.

3.2.1. Subjects. Five male subjects participated in Experiment 2, among them three students from the University of Kiel who were naïve to the purpose of the experiment, and the authors (FF, DW). Subjects had normal color vision and normal or corrected-to-normal visual acuity.

3.2.2. Apparatus and stimuli. The same technical setup was used as in Experiment 1. We constructed the stimuli such that in the top half of the display, subjects were presented a multi-slit view setup which made partially visible a Gabor patch within the region covered by the slits. The background color was black, the slits were 1.35 min of arc wide, 5.7 deg long and grey (CIE $x = 0.30$, $y = 0.31$, L = 12 cd m^{-2}). The vertically oriented Gabor patch (Fig. 6) of radius 2.9 deg was horizontally centered on the screen. It was equiluminant and varied in its chromaticity according to the Gaussian-windowed cosine-wave on the axis defined by the colors green (CIE $u = 0.16$, $v = 0.47$, $L = 12$ cd m^{-2}) and red (CIE $u = 0.24$, $v = 0.45$, $L = 12$ cd m^{-2}).

The slit array varied in its density. It rigidly moved to the right at a speed of 4.5 deg s^{-1}, whereas the virtual Gabor patch remained stationary. In the bottom half of the display, subjects were presented a fully visible Gabor patch of the same size as the test Gabor. Its chromaticity pattern was identical to that of the test Gabor, but its luminance was reduced to 8 cd m^{-2}. Its luminance was also windowed with a two-dimensional Gaussian to make it blend in with the black background. The match Gabor was adjustable with respect to the frequency of the underlying cosine-wave.

Figure 6. (See color plate VIII) Example Gabor patch used as a target stimulus in Experiment 2. Please note that the colors may be incorrectly reproduced in this print.

The vertical orientation of the Gabor together with the narrow slit width ensured that, in a given frame, no frequency information would be available within each slit. As such, the highest frequency that could be reliably recovered from a static version of the display is given by the Nyquist limit, equal to one half of the slit frequency.

3.2.3. Procedure. The frequency of the test Gabor was set to 0.69, 1.05, 1.39, 1.74, 2.09, 2.79, 3.48 or 4.18 cycles deg^{-1}, corresponding to 4, 6, 8, 10, 12, 16, 20 and 24 cycles in the area of the Gabor. The slit frequency was one of 2.44, 3.14, 3.83 or 4.53 slits deg^{-1}, corresponding to 14, 18, 22 and 26 slits in the area of the Gabor. The two independent variables were combined in a fully crossed 8 × 4 design.

Subjects adjusted the frequency of the match Gabor with the keyboard's cursor keys. Ten repetitions for each setting were carried out, resulting in a total of 8 × 4 × 10 = 320 trials. If a setting could not be made because the Gabor was not seen and its frequency therefore indeterminable, subjects could flag this trial as invalid using the keyboard's space bar.

A separate control condition with a fully visible test Gabor patch was used to yield individual assessments of the adjustment method's inherent inaccuracy. All of the aforementioned frequencies had to be matched by each subject in the control condition, resulting in 8 × 10 = 80 additional trials. The results provided a reference case for the adjustments in the experimental conditions.

3.2.4. Results. Trials that were flagged as invalid, a total of 193, equalling 12% of the trials, were excluded from the following analysis. Since the subjects did not exhibit systematic differences in their frequency settings, we pooled their data. The resulting data are shown in Fig. 7a. The graph depicts the mean relative deviations of subjects' settings from the respective frequency of the test Gabor. The relative deviation is measured in percentage of the test frequency. The graph also contains the respective Nyquist limit for each slit frequency. In addition to the experimental conditions, the graph provides the range of relative deviations in the control condition together with their standard errors.

The settings were in general rather precise: mean relative deviation of the subjects' settings typically ranged from only 10 to 15% of the test frequency. Most importantly however, mean deviations in the experimental conditions were about as large as, or only slightly larger than, in the control condition — with a clear exception of the lowest line density. A general trend for all four slit frequencies was that mean relative deviation was higher for the lowest and highest Gabor frequencies. This also means that relative deviations did not simply increase with Gabor frequency in a monotonic way. With regard to slit frequency, it is apparent that lower slit frequencies are associated with higher mean relative deviations of Gabor frequencies.

Figure 7b depicts the data in a way that makes explicit the relationship between mean relative deviation in cases where the Gabor frequency was below the Nyquist limit imposed by the slit frequency, compared to cases where the slit frequency
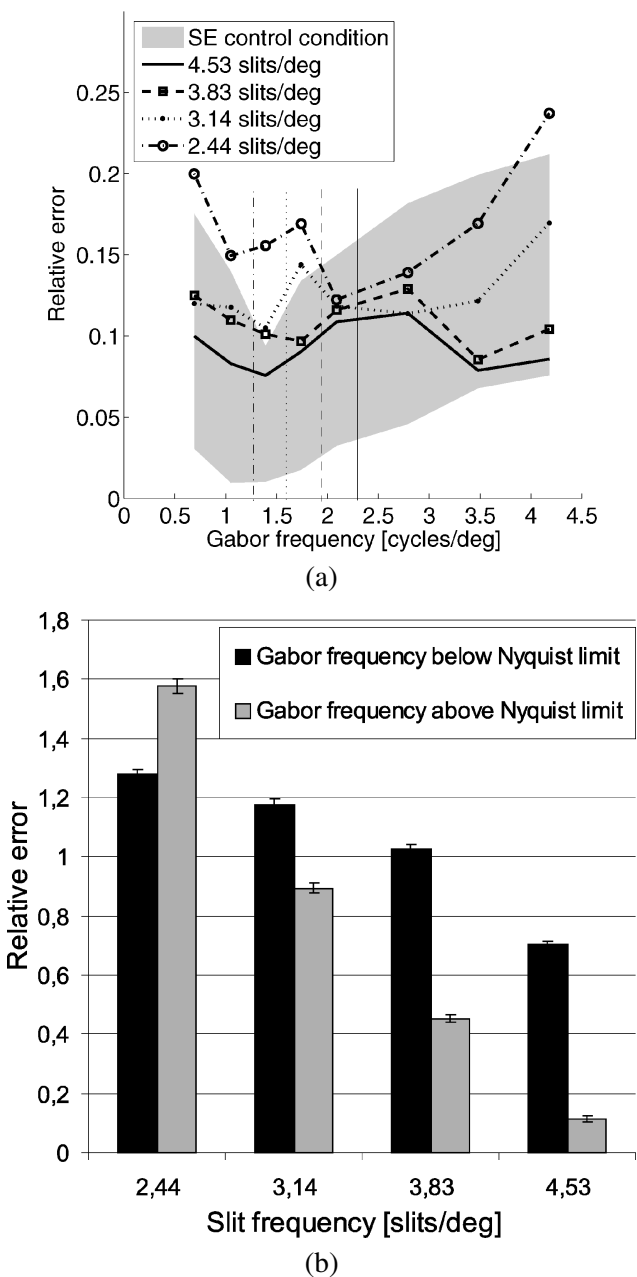
(a)

(b)

Figure 7. Results from Experiment 2. (a) Mean relative deviation of the subject's settings is plotted against the frequency of the test Gabor. Each line represents the settings for one slit frequency. The grey background area shows the range of the subject's settings in the control condition, i.e. the mean relative deviation ±1 SE. Vertical lines indicate the respective Nyquist limits induced by the corresponding slit frequencies. (b) Mean relative deviation of the frequency settings for all tested slit frequencies below and above their respective Nyquist limit. Error bars indicate ±1 SEM.

was above that limit. Contrary to intuition, mean relative deviation was not higher beyond the Nyquist limit; in fact it was lower, with a decreasing difference between them towards the lowest slit frequency. Only here subjects showed worse performance beyond the Nyquist limit. The figure also shows more directly that adjustment inaccuracies generally decreased with higher slit frequencies.

3.2.5. Discussion. Results from Experiment 2 again show that spatiotemporal integration of a texture does not lead to its blending into a uniform color. Instead, the surface structure is preserved as indicated by the fact that the fragmented Gabor pattern's frequency could be reliably recovered well beyond the Nyquist limit, theoretically imposed by the frequency of the slits. In fact, for the two highest slit frequencies, a subject's performance was just as good over nearly all tested frequencies as in the reference case with an unobstructed test Gabor. This means that due to the possibility of temporal integration, surface feature perception was almost perfect in these two cases.

Based on these results, one may ask whether the employed task is at all sensitive to the amount of surface information that is available in the display in each frame. This however follows from the pattern of results which indicates that lower slit frequencies indeed led to higher mean inaccuracies. The analysis of the performance differences with respect to test frequencies below and above the Nyquist limit (Fig. 7b) also suggests a similar conclusion: For the lowest slit frequency, the Nyquist limit had the expected effect of separating better from worse performance. For all higher slit frequency, the Nyquist limit lost this theoretical property, hinting at the fact that there is no resolution limit when spatiotemporal integration is possible (Nishida, 2004). Still, it remains elusive as to why the performance above the Nyquist limit was actually better for the three highest slit frequencies.

3.3. Experiments 3 and 4

Experiments 3 and 4 have complementary designs to address the question to what extent quantitative properties of texture spreading are compatible with the mechanism of sensory integration as delineated in Section 2.2. This is done using a dynamic texture spreading stimulus to study the interplay between the density of the line array and its speed in determining the filling-in effect. Although Experiments 3 and 4 do not directly test how much proper surface information is used to complete the task [see Note 4], results from Experiments 1 and 2 attest that the dynamic texture spreading stimulus indeed provides relevant information about the surface layout. Since Experiments 3 and 4 employ the same stimulus class, there is no reason to assume otherwise for their case.

3.3.1. Subjects. Six subjects (three female, three male) participated in Experiments 3 and 4, all of them students from the University of Kiel who were naïve as

to the purpose of the experiment. Subjects had normal color vision and normal or corrected-to-normal visual acuity. In a previous run of these two experiments with four subjects, we used slightly different parameter ranges for the independent variables. The obtained results did not differ systematically from the ones reported here.

3.3.2. Apparatus and stimuli. The same technical setup was used as in Experiments 1 and 2, with the exception of the monitor's vertical refresh rate, which was set to 75 Hz.

The background of the display was white (CIE $x = 0.30$, $y = 0.31$, $L = 55 \, \mathrm{cd \, m^{-2}}$) and subtended a visual angle of 26.7 by 21.9 deg. An array of vertically-oriented black lines (CIE $x = 0.30$, $y = 0.31$, $L < 0.2 \, \mathrm{cd \, m^{-2}}$), 6.75 min of arc wide and 3.37 deg long, was placed on the background.

In the interior of a virtual ellipse, red (CIE $x = 0.63$, $y = 0.34$, $L = 12.6 \, \mathrm{cd \, m^{-2}}$) striped flanks of 1.35 min of arc width and 2.78 cpd stripes were added parallel to the lines on either side. The virtual ellipse was centered on the screen and had main radii of 1.69 and 1.46 deg (Fig. 3). From trial to trial, its orientation was varied randomly in steps of 45 deg.

A fixation cross was placed on the horizontal center line at an eccentricity of 0.23 min of arc to the disk boundary, as informal observations had indicated that texture spreading was strongest under parafoveal viewing conditions, similar to what has been reported for anorthoscopic perception (Hecht, 1924), and static texture spreading (Watanabe and Cavanagh, 1991). The line array rigidly moved to the right, whereas the virtual disk remained stationary at the center of the screen. The velocity of the moving lines was varied by changing the length of the discrete "jumps" of the lines from one frame to the next. The trailing side of the display was continuously replenished with lines.

3.3.3. Procedure. In both experiments, we used a method very similar to that employed by Haber and Nathanson (1968), and asked subjects to find the threshold for perceiving a complete textured ellipse using an interactive adjustment procedure. Subjects were given two criteria for the perceptual presence of a *complete* ellipse. First, they should be able to discern the striped texture, and second, they should be able to identify the ellipse's orientation, asserting that they saw the complete boundary. The suggested strategy was to first explore the perceptual outcome of the setting over a wide range of the values of the dependent variable, and then to close in on threshold level.

In Experiment 3, the speed of the moving line array served as the independent variable. It was varied in five steps, 3.37, 5.06, 6.75, 8.44, and 10.12 $\mathrm{deg \, s^{-1}}$. The initial value of the line spacing was set to 0.67 deg. Subjects interactively controlled the spacing of the vertical lines in steps of 1 pixel using the arrow keys of the computer keyboard until the illusory ellipse was just noticeable.

In Experiment 4, the spacing of the vertical lines was varied in five steps (0.27, 0.54, 0.81, 1.08, 1.35 deg), corresponding to line densities of 106, 53, 35, 26, and

21 lines visible. The initial speed was set to 6.75 $\deg s^{-1}$. Subjects interactively controlled the speed of the line array in steps of one pixel displacement per frame until they just perceived the illusory ellipse.

The chosen ranges for the two independent variables were derived from results in preliminary studies with the goal of obtaining easily comparable functions in the two experiments relating line density and speed of the line array. That is, the values were selected such that the values for the independent variable in Experiment 4 would be similar to the dependent-variable settings in Experiment 3, and *vice versa*.

Experiments 3 and 4 were carried out separately with the order of the independent variables counterbalanced across subjects. In each experiment, subjects performed 10 repetitions for each of the five conditions in pseudo random order, resulting in a total of 50 trials.

Prior to the experiments, all subjects familiarized themselves with the task of adjusting the speed and spacing of the line array, respectively. There was no time limit for either task — subjects pressed a key to confirm their setting, and continued with the next trial.

3.3.4. Results. Figure 8 displays the combined results for the two experiments for an individual subject as well as averaged across subjects. Line spacing is plotted against the speed of the moving line array. In both experiments, these two variables seem to be linked by a positive, approximately linear relationship. In the following account of the results, we will refer to the data in terms of line density, inversely related to line spacing. It is evident that for both types of threshold determination, the density of the moving line array and its speed are in a compensatory relationship. This means that in Experiment 3, line density could be reduced with higher speeds while maintaining the illusory percept, or, conversely, line density had to be increased with slower speeds to obtain the illusory percept. For Experiment 4, results indicate that with lower line density, speed had to be increased to reach threshold, and with denser line spacing, speed could be reduced.

As intended, the functions relating speed and line spacing in the two experiments have mostly overlapping domains and ranges, and they can thus be easily compared. The functions are separated by a small horizontal offset, but are quite similar in shape. This suggests that the results of the two experiments indeed describe the same relationship, and can be combined into one data set. Yet, the two functions do not coincide exactly. The observed difference between them may be attributed to small shifts of the subjective criterion between experimental sessions.

As is apparent in Fig. 8, there is some inter-subject variability in the absolute settings of the dependent variables. The narrow error bars indicating the SEM however show that subjects' settings were in general quite consistent within conditions.

Based on these data, it is possible to calculate the delay between two consecutive appearances of a flank at a given fixed location within the virtual ellipse at threshold level. This delay between two refreshes of texture information at a stationary point
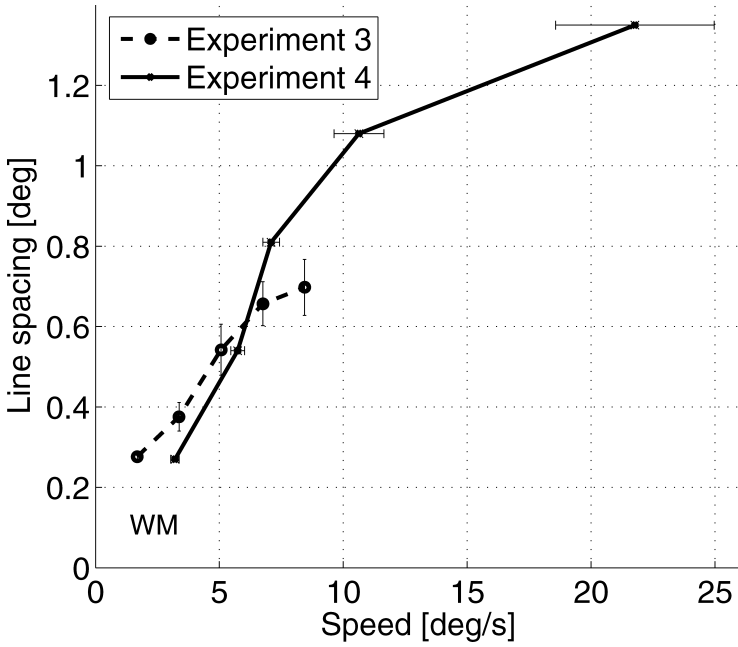
(a)
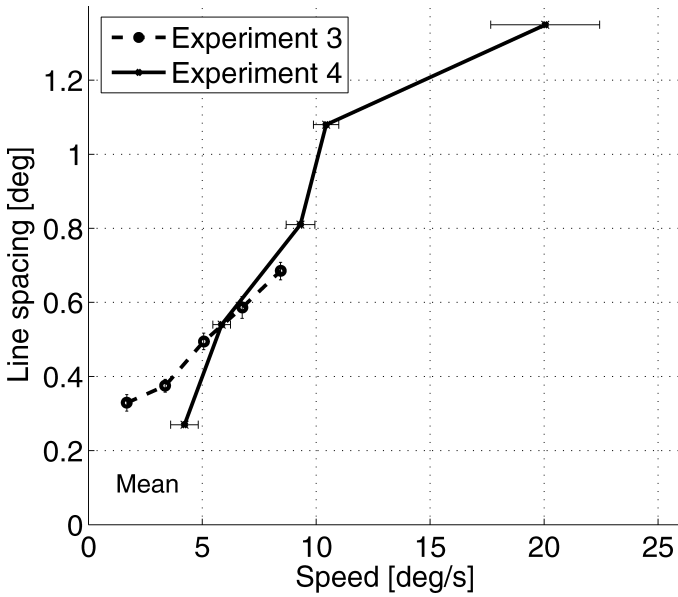
(b)

Figure 8. Results from Experiments 3 and 4. Distance of the moving lines is plotted against the speed of the moving lines. (a) Individual data from WM. Error bars indicate ± 1 SE. (b) Combined results from all subjects. Error bars indicate ± 1 SEM.

is plotted in Fig. 9 against the speed of the moving line array for both experiments. The pattern of results slightly differs between subjects, but in each case the data points varied within a close range from 100 to 180 ms. This consistence of the maximally-tolerated delay is noteworthy because its calculation involves a division operation and is therefore especially sensitive to changes in its input values. The tolerated delay increases non-linearly with very low speeds in Experiment 1, indicating that the line spacing setting never fell below a fixed value. This is to be expected from the fact that a small distance between lines can be bridged by color or texture spreading even in static displays.

3.3.5. Discussion. The results from both experiments show a fair degree of consistency within as well as across subjects. This seems worth noting, considering the peculiar perceptual nature of most illusory percepts on the one hand, and the employed adjustment method of threshold determination on the other hand. Any vagueness in the percept would have necessarily resulted in highly variable threshold settings. Since the individual settings actually stayed in a reasonably close range, one may infer that the percept itself is stable — a prerequisite for any further data analysis with regard to possible perceptual mechanisms.

In line with results presented by Morgan (1980) and Morgan *et al*. (1982), the idea of simple temporal integration through sensory persistence (see Section 2) was able to capture the data from both experiments surprisingly well: the density of the line array and its speed seem to act together in a compensatory manner in determining the strength of the illusory percept. Within the studied parameter range, this means that a change in the strength of the illusory percept due to, for example, a reduction in line density could be compensated by a corresponding increase in the speed of the line array.

That sensory persistence accounts well for our data warranted a further analysis in terms of one of its important parameters: the maximum delay between two input signal refreshes that is tolerated by the visual system while still yielding a complete percept. This delay was not exactly constant over experimental conditions as would be expected from retinal painting, but it stayed in a narrow range from around 100 to 180 ms (Fig. 9), indicating the maximum time window for integration processes. This estimate is remarkably similar to results obtained in a number of other experiments on object recognition: Burr (1980) estimated that information in the visual system is preserved for about 120 ms, similar to the integration timeframe of 150 ms that Morgan *et al.* (1982) found necessary for perceiving a complete circle that is successively defined through single small dots. Studying the integration of local element changes in the perception of subjective contours in spatiotemporal boundary formation, Shipley and Kellman (1994) came to the conclusion that signal integration takes place within 165 ms. In EEG studies on the time course of neural signals in object recognition, Johnson and Olshausen (2003) found recognition-correlated ERP components between 150 and 300 ms. Thorpe *et al.* (1996) reported an EEG component at 150 ms linked to the detection of objects in complex natural scenes.
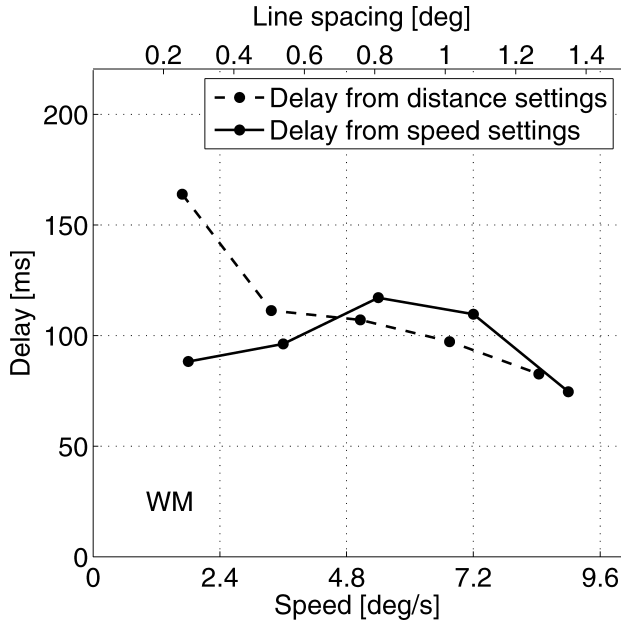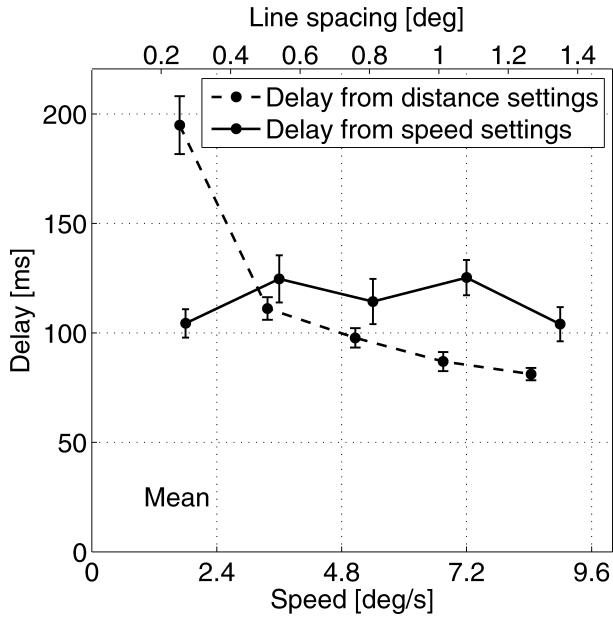
Figure 9. Results from Experiments 3 and 4, re-analyzed in terms of the maximally tolerated delay between two consecutive appearances of a textured flank at threshold level. (a) Individual data from WM. Error bars indicate ±1 SE. (b) Combined results from all subjects. Error bars indicate ±1 SEM.

4. GENERAL DISCUSSION

The results obtained in all four experiments were systematic, reliable, and of sufficient inter-individual consistency. This warrants the use of dynamic texture spreading as a tool to investigate visual completion processes with phenomenological (Experiments 3 and 4) as well as with more rigorous quantitative methods (Experiments 1 and 2). We believe that this possibility constitutes a significant advantage over static color or texture spreading, which are often perceptually unstable and therefore prove elusive to grasp with quantitative methods.

Experiments 1 and 2 served to illustrate that texture spreading allows observers to recover significantly more information about the surface layout of complex and highly fragmented stimuli than just their boundaries. This was done using performance tasks whose successful completion required that a rich perceptual surface pattern representation was at the observer's disposal. Texture spreading brought with it a measurable performance benefit and thus proved to be a functionally relevant effect.

Experiments 3 and 4 then tried to relate the phenomenon under consideration to sensory persistence, a simple perceptual mechanism whose involvement in spatiotemporal integration is frequently discussed (see Section 2.2). The data from these two experiments were in good agreement with a retinal painting account of dynamic texture spreading. In addition, the results of Experiments 1 and 2 may also be interpreted as being compatible with retinal painting. The performance drop associated with a progressive reduction of diagnostic stimulus fragments was gradual rather than abrupt.

Still, some qualitative observations indicate that there are several aspects of the phenomenon that are not accounted for by sensory persistence. First, since sensory persistence and the corresponding integration mechanism are strictly local, the effect should not be sensitive to global changes in the stimulus layout, as long as these changes do not affect the refresh frequency of texture information. However, it can be observed that texture spreading does not occur if the black lines, which do not directly contribute to the sensory persistence of texture cues, are removed from the display (Fig. 1a). While this spatial interaction is not readily explained by retinal painting, it can be naturally integrated into a high-level perspective which emphasizes the global stimulus layout and scene organization. The Bregman–Kanizsa display (Bregman, 1981; Kanizsa, 1979; Nakayama *et al.*, 1989) for example illustrates that fragmented letters are immediately recognized only if the occluding material is also visible. Applying this reasoning to dynamic texture spreading, we may interpret the lines as occluders, signaled by the T-junctions they form with the flanks. Their presence seems to constitute a cue that determines which fragments in the stream of information are classified as coextensive, and are therefore integrated, and which ones are classified as disjoint and duly segregated (Dixon and Di Lollo, 1994).

As a second issue, it is unobvious how to apply the temporal integration scheme to static neon-spreading displays. In this case, it is not plausible to assume that

the gaps between corresponding stimulus elements are temporally joined with eye movements which recurringly spread the complete stimulus onto the retina. However, it should be noted that, in general, the color and texture spreading is a lot weaker in static displays compared to their dynamic counterparts.

Finally, one of the main objections against a major involvement of sensory persistence was that anorthoscopic contour perception can still occur when experimental conditions preclude the possibility that a complete image is temporally spread over the retina. This argues against a critical role of retinal painting in boundary integration. However, apart from the possibility that a retinotopic and a distal representation may co-exist (Shipley and Cunningham, 2001), it does not touch the question whether other perceptual attributes, such as surface filling-in, trace back to sensory persistence. Taken together, this suggests that, unlike certain spatial and figural stimulus aspects, sensory persistence by itself may not be a *necessary* prerequisite for all spreading effects. Nevertheless, when stimulus conditions favor its occurrence, it seems to be able to provide performance-relevant information about surface attributes and to determine crucial characteristics of the percept. In particular, it may be responsible for the particular vividness of the spreading effect in dynamic stimuli.

5. CONCLUSION

Dynamic texture spreading is a filling-in effect arising from motion displays that constitutes another example of the interaction of color, motion and spatial stimulus configuration (Hoffman, 2003). It exhibits a large discrepancy between the small amount of input information available to the visual system in each separate frame and the rich perceptual interpretation of the animated stimulus in terms of a complete object featuring strong illusory contours as well as a genuine surface qualities like a colored texture and depth stratification.

Apart from constituting an interesting phenomenon in its own right, dynamic texture spreading is a particularly robust effect that proved to be amenable to standard psychophysical methods. It may therefore be used as a reliable tool to study models of object-completion processes.

We established that texture spreading is a genuine representation of surface qualities that goes beyond mere contour interpolation. We further related the phenomenon to ongoing discussions about putative mechanisms for spatiotemporal integration: We examined the adequacy of retinal painting as a possible cause for the filling-in effect, a process that has often been proposed as an explanation for anorthoscopic perception. Our data are consistent with this hypothesis and showed that motion and the density of the moving apertures can act together in a compensatory manner in sustaining the texture spreading. This suggests that the interpolation in dynamic texture spreading partly takes place at a rather basic processing stage that temporally integrates sensorily persistent input signals. We discussed several findings which indicate that there are some aspects of dynamic

texture spreading that are not captured by this simple integration mechanism. We thus conclude that an elementary integration process may well play an important role, but is not sufficient to explain dynamic texture spreading in full detail. Further work is required to better isolate the respective contributions of low and high-level mechanisms in dynamic spreading phenomena.

Acknowledgements

NOTES

1. For an online demonstration of the display, please see http://www.psychologie. uni-kiel.de/psychophysik/demos/ts/

2. As was rightly pointed out by one referee, the terminology of filling-in effects is currently quite heterogeneous and does not induce a convincing classification. Until a conceptual clarification of the terms has taken place, we try to adopt the terminology used by the respective original authors.

3. Motion conditions: $0.5/(1 + 8x) + 0.5$, static condition: $0.5/(1 + \exp(-6 + 10x)) + 0.5$.

4. We thank one anonymous reviewer for directing our attention to this issue.

REFERENCES

Andersen, G. J. and Cortese, J. M. (1989). 2-D contour perception resulting from kinetic occlusion, *Perception and Psychophysics* **46**, 49–55.

Anstis, S. (2005). Local and global segmentation of rotating shapes viewed through multiple slits, *J. Vision* **13**, 194–201.

Anstis, S. M. and Atkinson, J. (1967). Distortions in moving figures viewed through stationary slit, *Amer. J. Psychol.* **80**, 572–585.

Brainard, D. H. (1995). Colorimetry, in: *OSA Handbook of Optics*, Bass, M. (Ed.), pp. 26.1–26.54. McGraw-Hill, New York, NY.

Bregman, A. L. (1981). Asking the 'what for' question in auditory perception, in: *Perceptual Organization*, Kubovy, M. and Pomerantz, J. R. (Eds), pp. 99–118. Lawrence Erlbaum, Hillsdale, NJ.

Bressan, P., Mingolla, E., Spillmann, L. and Watanabe, T. (1997). Neon color spreading: a review, *Perception* **26**, 1353–1366.

Bruno, N. and Bertamini, M. (1990). Identifying contours from occlusion events, *Perception and Psychophysics* **48**, 331–342.

Bruno, N. and Gerbino, W. (1991). Illusory figures based on local kinematics, *Perception* **20**, 259–274.

Bull, N. J., Hunter, M. and Finlay, D. C. (2003). Cue gradient and cue density interact in the detection and recognition of objects defined by motion, contrast, or texture, *Perception* **32**, 29–39.

Burr, D. C. (1980). Motion smear, *Nature* **284**, 164–165.

Cicerone, C. M., Hoffman, D. D., Gowdy, P. D. and Kim, J. S. (1995). The perception of color from motion, *Perception and Psychophysics* **57**, 761–777.

Cunningham, D. W., Shipley, T. F. and Kellman, P. J. (1998). The dynamic specification of surfaces and boundaries, *Perception* **27**, 403–415.

Devinck, F., Delahunt, P. B., Hardy, J. L., Spillmann, L. and Werner, J. S. (2005). The watercolor effect: quantitative evidence for luminance-dependent mechanisms of long-range color assimilation, *Vision Research* **45**, 1413–1424.

Dixon, P. and Di Lollo, V. (1994). Beyond visible persistence: An alternative account of temporal integration and segregation in visual processing, *Cognit. Psychol.* **26**, 33–63.

Ekroll, V. and Faul, F. (2002). Perceptual transparency in neon color spreading displays, *Perception and Psychophysics* **64**, 945–955.

Fahle, M. and Poggio, T. (1981). Visual hyperacuity: Spatiotemporal interpolation in human vision, *Proc. Roy. Soc. London B* **213**, 451–477.

Fendrich, R., Rieger, J. W. and Heinze, H.-J. (2005). The effect of retinal stabilization on anorthoscopic percepts under free-viewing conditions, *Vision Research* **45**, 567–582.

Fidopiastis, C., Hoffman, D. D. and Prophet, W. D. (2000). Constructing surfaces and contours in displays of color from motion: the role of nearest neighbors and maximal disks, *Perception* **29**, 567–580.

Haber, R. N. and Nathanson, L. S. (1968). Post-retinal storage? Some further observations on Parks' camel as seen through the eye of a needle, *Perception and Psychophysics* **3**, 349–355.

Hecht, H. (1924). Neue untersuchungen über die Zöllnerschen anorthoskopischen Zerrbilder, *Zeitschrift für Psychologie* **94**, 153–194. (Herausgegeben von F. Schumann: I. Die simultane Erfassung der Figuren)

Hine, T. (1987). Subjective contours produced purely by dynamic occlusion of sparse-points array, *Bull. Psychonomic Soc.* **25**, 182–184.

Hoffman, D. D. (2003). The interaction of colour and motion, in: *Colour — Mind and the Physical World*, Mausfeld, R. J. and Heyer, D. (Eds), pp. 361–378. Oxford University Press, Oxford, UK.

Johnson, J. S. and Olshausen, B. A. (2003). Time course of neural signatures of object recognition, *J. Vision* **3**, 499–512.

Kanizsa, G. (1979). *Organization in Vision*. Praeger, New York, NY.

Kellman, P. J. and Shipley, T. F. (1991). A theory of visual interpolation in object perception, *Cognitive Psychol.* **23**, 141–221.

Kellman, P. J., Guttman, S. E. and Wickens, T. D. (2001). Geometric and neural models of object perception, in: *From Fragments to Objects: Segmentation and grouping in vision*, Shipley, T. F. and Kellman, P. J. (Eds). Elsevier Science Publishers, Oxford, UK.

Mateeff, S., Popov, D. and Hohnsbein, J. (1993). Multi-aperture viewing: perception of figures through very small apertures, *Vision Research* **33**, 2563–2567.

Morgan, M. J. (1980). Analogue models of motion perception, *Phil. Trans. Roy. Soc. London B* **290**, 117–135.

Morgan, M. J., Findlay, J. M. and Watt, R. J. (1982). Aperture viewing: a review and a synthesis, *Quart. J. Exper. Psychol.* **34A**, 211–233.

Nakayama, K., Shimojo, S. and Silverman, G. H. (1989). Stereoscopic depth: Its relation to image segmentation, grouping and the recognition of occluded objects, *Perception* **18**, 55–68.

Nishida, S. (2004). Motion-based analysis of spatial patterns by the human visual system, *Current Biology* **14**, 830–839.

Parks, T. E. (1965). Post-retinal visual storage, *Amer. J. Psychol.* **78**, 145–147.

Pessoa, L., Thompson, E. and Noë, A. (1998). Finding out about filling-in: a guide to perceptual completion for visual science and the philosophy of perception, *Behav. Brain Sci.* **21**, 723–802.

Pinna, B., Brelstaff, G. and Spillmann, L. (2001). Surface color from boundaries: a new 'watercolor' illusion, *Vision Research* **41**, 2669–2676.

Pirenne, M. H. (1962). Light adaptation. I. The Troxler phenomenon, in: *The Eye*, Davson, H. (Ed.), pp. 197–199. Academic Press, London, UK.

Ramachandran, V. S. (1992). Blind spots, *Scientific American* **266**, 85–91.

Ramachandran, V. S. and Gregory, R. L. (1991). Perceptual filling-in of artificially induced scotomas in human vision, *Nature* **350**, 699–702.

Rock, I. (1981). Anorthoscopic perception, *Scientific American* **244**, 145–153.

Rock, I., Halper, F., DiVita, J. and Wheeler, D. (1987). Eye movements as a cue to figure motion in anorthoscopic perception, *J. Exper. Psychol.: Human Perception and Performance* **13**, 344–352.

Rothschild, H. (1922). Untersuchungen über die sog. Zöllnerschen anorthoskopischen Zerrbilder, *Zeitschrift für Psychologie* **90**, 137–166.

Shimojo, S. and Richards, W. (1986). 'Seeing' shapes that are almost totally occluded: a new look at Park's camel, *Perception and Psychophysics* **39**, 418–426.

Shipley, T. F. and Cunningham, D. W. (2001). Perception of occluding and occluded objects over time: spatiotemporal segmentation and unit formation, in: *From Fragments to Objects: Grouping and Segmentation in Vision*, Shipley, T. F. and Kellman, P. J. (Eds), pp. 557–585. Elsevier Science, New York, NY.

Shipley, T. F. and Kellman, P. J. (1993). Optical tearing in spatiotemporal boundary formation: When do local element motions produce boundaries, form, and global motion? *Spatial Vision* **7**, 323–339.

Shipley, T. F. and Kellman, P. J. (1994). Spatiotemporal boundary formation: boundary, form, and motion perception from transformations of surface elements, *J. Exper. Psychol.: General* **123**, 3–20.

Thorpe, S. J., Fize, D. and Marlot, C. (1996). Speed of processing in the human visual system, *Nature* **381**, 520–522.

Troxler, D. (1804). Ueber das Verschwinden gegebener Gegenstände innerhalb unseres Gesichtskreises, in: *Ophthalmologische Bibliothek II*, Himly, K. and Schmidt, J. A. (Eds), pp. 1–119. F. Frommann, Jena, Germany.

Tynan, P. and Sekuler, R. (1975). Moving visual phantoms: a new contour completion effect, *Science* **188**, 951–952.

von Vierordt, K. (1868). *Der Zeitsinn nach Versuchen: §27. Scheinbare Verzerrung bewegter Gegenstände*. Laupp, Tübingen, Germany.

Watanabe, T. and Cavanagh, P. (1991). Texture and motion spreading, transparency and the aperture problem, *Perception and Psychophysics* **50**, 459–464.

Watson, A. B. and Pelli, D. G. (1983). QUEST: A Bayesian adaptive psychometric method, *Perception and Psychophysics* **33**, 113–120.

Weisstein, N., Maguire, W. and Berbaum, K. (1977). A phantom-motion aftereffect, *Science* **198**, 955–958.

Wollschläger, D., Rodriguez, A. M. and Hoffman, D. D. (2001). Flank transparency: transparent filters seen in dynamic two-color displays, *Perception* **30**, 1423–1426.

Wollschläger, D., Rodriguez, A. M. and Hoffman, D. D. (2002). Flank transparency: the effects of gaps, line spacing, and apparent motion, *Perception* **31**, 1073–1092.

Zöllner, F. (1862). Ueber eine neue Art anorthoskopischer Zerrbilder, *Annalen der Physik und Chemie* **117**, 477–484.

APPENDIX A. ADAPTIVE ALGORITHM USED IN EXPERIMENT 1

The algorithm we used in Experiment 1 to determine the tested line distances first divided the complete parameter range into four intervals. It then randomly drew and tested values from each interval until all intervals would contain at least 40 data points. It then calculated the mean percentages correct for each interval and

compared the results of adjacent intervals. If two adjacent intervals would differ by more than 10%, their spanned range would be sub-divided into four intervals. Values from each new interval were then randomly drawn and tested until they all contained at least 40 data points. Mean percentages correct for all existing intervals were (re-)calculated and compared between adjacent intervals. This process of recursive subdivision continued until all differences between the mean percentages correct of adjacent intervals were less than 10%, or until the minimum interval width of 1 px was reached. The method ensured that the sampling rate was increased in areas where the rate of change of the underlying function was high, thus leading to an approximately constant smoothness of the reconstructed curve. Conversely, the sampling rate was lowered and trials saved in areas where the underlying function did not change very rapidly.

The described method may be seen as a breadth-first-search algorithm. The nodes of the search tree correspond to pairs of tested intervals, and the depth of the search tree represents the resolution of tested intervals. The decision criterion used at each node to decide whether to descend into a deeper level of the tree corresponds to the question whether the difference between adjacent intervals was greater than the 10%.
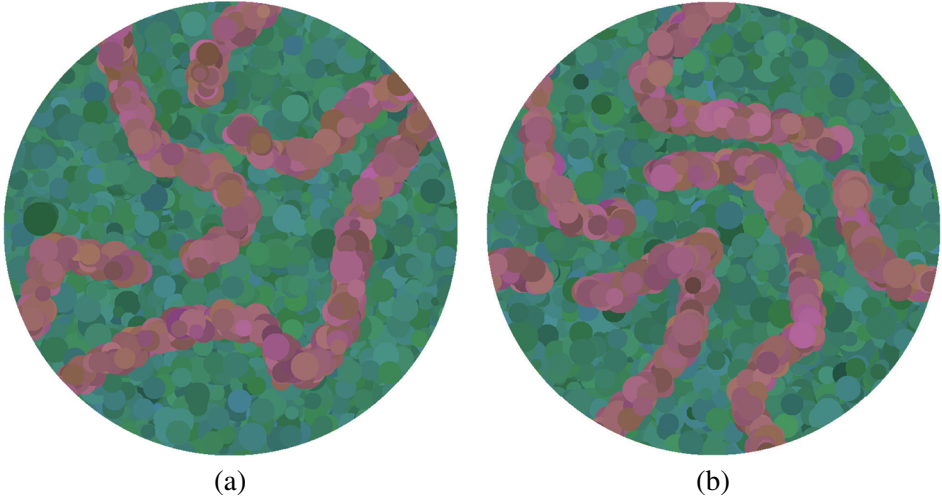
Plate VII
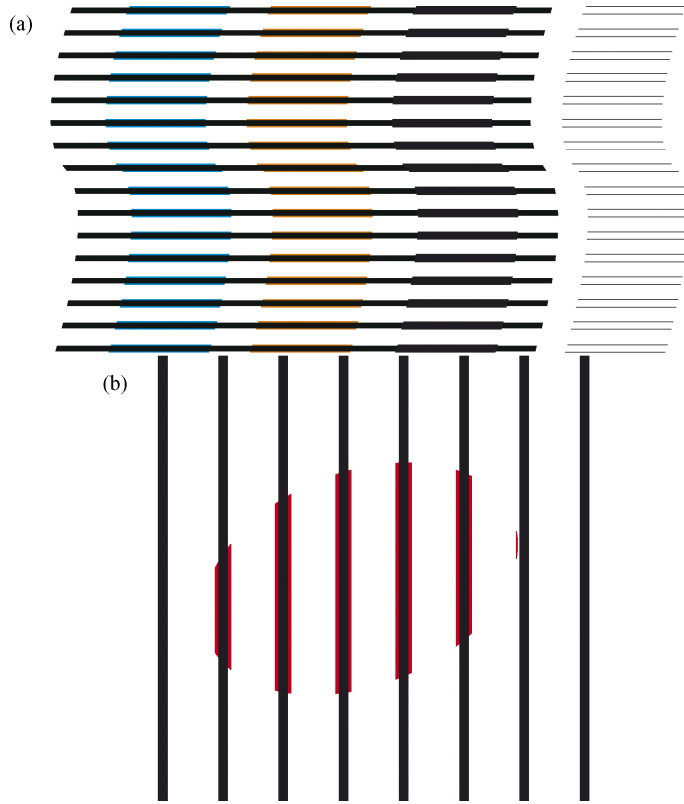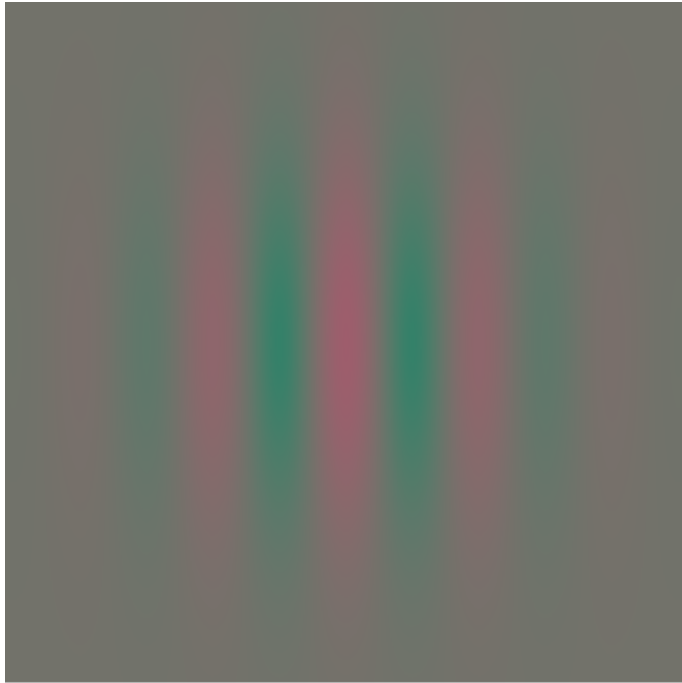

Wollschläger and Faul, Figure 1.
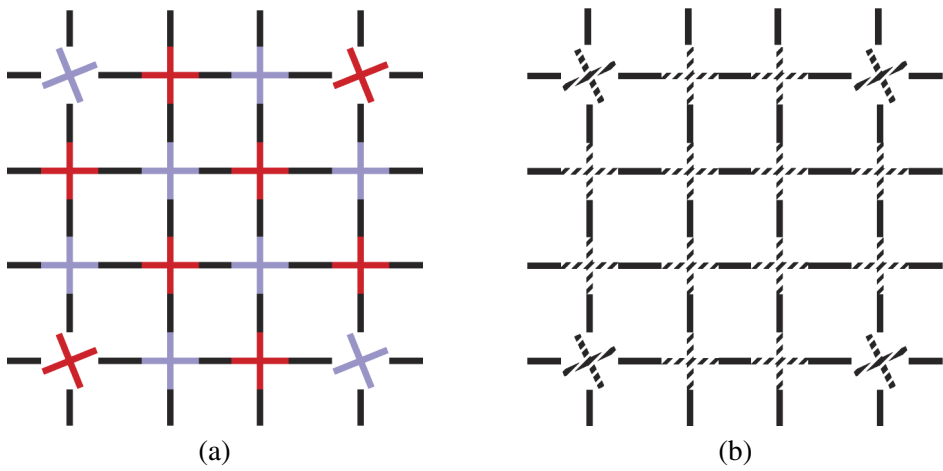
Wollschläger and Faul, Figure 2.

Plate VIII

(a)   (b)

Wollschläger and Faul, Figure 4.

Wollschläger and Faul, Figure 6.