Simple assumptions to improve Markov illuminance and reflectance

Yuki Kobayashi, Akiyoshi Kitaoka

Abstract

Murray (2020) recently introduced a novel computational lightness model, Markov illuminance and reflectance (MIR). MIR is a promising new approach that simulates human lightness processing using a conditional random field (CRF) where natural scene statistics of the reflectance and illumination are implemented. Although MIR can account for various lightness illusions and phenomena, it has limitations, such as the inability to predict reverse contrast phenomena. In this study, we improved MIR performance by modifying its inference process, the prior on X-junctions, and that on general illumination changes. Our modified model improved predictions for Checkerboard assimilation, the simplified Checkershadow and its control figure, the influence of luminance noise, and White's effect and its several variants. In particular, White's effect is a partial reverse contrast that is challenging for computational models, so this improvement is a significant advance for the MIR framework. This study showed the high extensibility and potential of MIR, which shows the promise for further sophistication.

Introduction

In lightness/brightness perception research, it is desirable to establish a computational model that precisely predicts human perception (e.g., Kingdom, 2011; Shapiro & Lu, 2011). Although Gestalt theories have made many contributions to the development of this research field (e.g., Bressan, 2006; Gilchrist et al., 1999), computational models (here, we refer to models that output rigorous predictions of human perceptions of any arbitrary image inputs) have also been discussed by many researchers. For instance, spatial filtering models are known as a promising approach that many researchers have examined and discussed (Betz et al., 2015; Blakeslee & McCourt, 1997; Blakeslee & McCourt, 1999; Dakin & Bex, 2003; Economou et al., 2007; Robinson et al., 2007; Zeman et al., 2015).

Although spatial filtering models are known to have high performance for brightness prediction (e.g., Blakeslee & McCourt, 2012), it is not necessarily easy to interpret psychological implications from each computational process in these models (e.g., to assign high weights to higher frequency filter outputs). This is not surprising because they are motivated by physiological processes in the human brain (Blakeslee & McCourt, 1999; Robinson et al., 2007) and aim to account for a low-level processing in brightness perception (Blakeslee & McCourt, 2012). However, if there is a computational model whose processes and structures are intuitively understandable as a representation of the human mind, it will be

useful for quantitative output predictions for arbitrary images and contribute to the qualitative explanation of the perception of lightness/brightness.

Murray (2020)’s novel lightness model, Markov illuminance and reflectance (MIR), is promising as an understandable computational model. MIR is based on a probabilistic model called a conditional random field (CRF). MIR’s CRF design is composed of some prior assumptions motivated by natural scene statistics (e.g., “Illuminance spans a wide range, and lower illuminances are more likely” (Murray, 2020, p. 19)), supposing that humans utilize these statistics to solve lightness ambiguity (Allred & Brainard, 2013; Feldman, 2015; Murray, 2013). Therefore, a large part of the computational processes in MIR has a certain amount of psychological rationality, and the meaning of each process can be relatively straightforwardly interpreted. In MIR, the model output is obtained as an illuminance map (perceived illumination, represented by a lux unit) and a reflectance map (perceived reflectance, or lightness, represented by a percent unit), and thus, each value of output is less arbitrary, unlike the output of spatial filtering models. Murray (2020) also argued that MIR is compatible with the widely known Gestalt theories (Gilchrist et al., 1999) because both share the idea that general lightness perception can be explained by a relatively small number of principles. Indeed, MIR incorporates Gestalt ideas into the CRF, such as “illuminance edges tend to be straighter than reflectance edges” (MIR’s constraint A7) (Adelson, 2000) or “illuminance changes tend to be gradual, not abrupt” (MIR’s constraint A3) (Agostini & Galmonte, 2002b; Land & McCann, 1971). MIR has been reported to be able to account for many lightness illusions and phenomena, suggesting the high probability that it will be a novel pathway in this research field.

In this study, we aimed to extend MIR. We present a modified version of MIR that can account for different types of lightness phenomena. The current version of MIR (Murray, 2020) does not predict reverse contrasts (Nedimović et al., 2021), where target areas neighboring dark regions appear darker than an equiluminant target neighboring bright regions (Agostini & Galmonte, 2002a; Bressan, 2001; Economou et al., 2015). White’s effect is one of the most famous reverse contrasts (White, 1981, 1979). This is a partial reverse contrast in which a gray bar sharing its longer edge with a dark adjacent area appears darker than an equiluminant gray bar sharing its longer edge with a bright adjacent area. Reverse contrast phenomena are a challenging and important issue for lightness/brightness models to explain because it highlights that lightness/brightness perception is more complicated than simply being determined by immediate contrast (Agostini et al., 2020; Economou et al., 2015; Gilchrist, 2006). It is a large leap for computational models to be able to explain White’s effect, as suggested by the fact that the ODOG model, a famous spatial filtering model (Blakeslee & McCourt, 1999), showed its high performance by accounting for White’s effect (see Betz et al., 2015 for a counterargument), which the authors’ preceding model (the DOG model) could not account for (Blakeslee & McCourt, 1997).

We modified some prior assumptions of the original MIR to better reflect natural scenes and improve its performance. The most important improvement was the successful prediction of White’s effect, but our model also includes several improvements. For readers, we made our Python and Julia codes available (supplementary materials will be available after acceptance) to welcome further tests, improvements, and discussions.

Modifications to the original MIR

Full use of the available links between nodes

First, we modified the original MIR’s belief propagation schedule. This is a technical modification of the model’s inference process rather than a modification to the model’s priors. Although the original belief propagation efficiently converges to approximated solutions, it does not fully utilize the available links in the CRF. We added another phase of message passing to utilize all the links in the CRF, and this helped the model to search for a solution more efficiently. See Appendix for details of this modification.

Prior 1 - Consistent X-edges cue illumination changes, but inconsistent ones do not.

The original MIR considers that “X-junctions are evidence for illuminance edges” (Murray, 2020, p. 20), assigning no cost to the illuminance change between two pixels that constitute an X-junction. This rule contributes to the correct predictions of illusions that give shadow impressions, such as Argyle illusion, Snake illusion, and Koffka ring (Adelson, 2000). In the original MIR, this rule was applied to any X-junctions without considering the relationships among luminances of the four pixels.

Although X-junctions do suggest an illumination change, they do not always confirm it. When two colinear luminance edges in an X-junction show the same contrast polarity (Figure 1; hereafter called consistent X-edges), these edges can be considered a cue of illumination changes. However, when two colinear luminance edges in an X-junction show different contrast polarities (Figure 1 (b); hereafter called inconsistent X-edges), they are unlikely to cue an illumination change. If the two colinear luminance edges are only caused by an illumination change, the edges must show consistent contrast polarities. A white–black checkerboard that is uniformly illuminated, for example, consists of many X-junctions caused only by reflectance edges (all of them are inconsistent X-edges), but the original MIR is likely to see them as illumination changes. We put a relatively high cost on inconsistent X-edges ($w_{Xinc} = 50$) while keeping the cost on consistent ones at zero ($w_{Xcon} = 0$). It is widely known that humans take contrast polarities in X-junctions into account when they judge surface characteristics (Kitaoka, 2005).

Figure 1: Examples of X-junctions. Each set of four squares represents an X-junction composed of four pixels. The orange arrows indicate an increasing edge, where a pixel at the arrowhead is brighter than that at the arrow’s root. The blue arrows indicate a decreasing edge, where a pixel at the arrowhead is darker. (a) Examples of consistent X-

edges. Two vertical colinear luminance edges (those between two horizontally aligned pixels) are both increasing (left) or decreasing (right). (b) Examples of inconsistent X-edges. Two vertical colinear luminance edges show inconsistent contrast polarities.

Prior 2 - Scenes with more luminance contrasts contain more illuminance changes.

In the original MIR, the cost for illuminance change in two pixels at a luminance edge (except for two pixels included in an X-junction) was held constant (the weight for Equations 6 and 7 in Murray (2020) was held at $w_1 = 20$). This parameter largely controls the contrast of the model’s illuminance map output. Holding it constant means that the model does not modulate its expectation regarding illumination changes for any images or scenes.

However, information included in an entire image is useful to judge whether a certain luminance edge should be attributed to illumination or reflectance. A scene with high illuminance contrast (i.e., illumination is less homogeneous) also casts an image with a high luminance contrast. Therefore, it is plausible to expect more illuminance change for images with a high luminance contrast. This assumption may seem too simple, but we observed that it plays a good role in predicting many illusions (see Model performance section). We implemented this assumption into the model by modulating parameter w_1 depending on the luminance contrast of the input image.

In our modified model, parameter w_1 was controlled as below. First, as a measurement of an image’s contrast, the average Michelson contrast can be obtained as:

$$MC_{total} = \sum_i \left(\frac{|l_{i1} - l_{i2}|}{l_{i1} + l_{i2}} \right)$$

$$MC_{average} = \frac{MC_{total}}{\sqrt{m^2 + n^2}}$$

where i represents each pair of adjacent pixels in the input image, and l_{i1} and l_{i2} are luminance levels (cd/m^2) of the two pixels constituting i . Therefore, MC_{total} is the sum of the Michelson contrast of all adjacent pixel pairs. The image’s height and width (numbers of pixels) are represented by m and n , respectively, so the denominator of the second equation is the length of the diagonal line of the image. $MC_{average}$ is used as a measurement of an image’s luminance contrast (here, the term “average” is not used in the strict sense) because MC_{total} varies depending on the image size (i.e., MC_{total} increases when an image is up-sampled). Parameter w_1 is determined as $w_1 = w_{Xinc} \times \exp(-MC_{average})$. This exponential function returns 50 ($= w_{Xinc}$) when an image has a zero luminance contrast (i.e., a homogeneous image), and the return value decreases as the image’s luminance contrast increases but never reaches zero.

Model performance

Below, we will show the performance of our modified model by comparing it with Murray (2020)'s original model. One run consisted of five iterations in Murray (2020)'s original implementation, but one iteration in our modified model requires twice as much computation as the original because of the modification of the belief propagation schedule. Therefore, in the following model tests, one run of the modified model consisted of three iterations and that of the original model consisted of six. This makes the model comparison more fair, by making single runs of the two models equally computationally complex.

Main illusory images employed by Murray (2020)

Murray (2020) employed 12 famous illusory figures to examine the model's performance (Figure 2). First, we compared the two models' performances using them. Murray (2020) adopted the best result (with the lowest energy) of ten runs, but we observed minor random fluctuations in the outputs with this number of runs; thus, we decided to take the best result of 30 runs to improve the accuracy.

Table 1 shows the predictions of the two models for 12 images. They are similar except in White's effect and Checkerboard assimilation, which we will discuss later. The modified model predicts illusions for images where the original model also predicts illusions, and the modified model does not predict illusions for images where the original model does not. The modified model does not predict an illusion in Koffka connected while the original does, but Murray (2020)'s experiment on human observers indicated that this figure does not cause a significant illusory effect; thus, the modified model's prediction may be closer to the human perception. The modified model also successfully accounts for differences in illusion magnitudes in the stimuli pairs (Argyle vs. Argyle control, Snake vs. Snake control, Koffka–Adelson and broken vs. Koffka connected, SLC vs. Articulated SLC), for which the original model was reportedly able to account for (Murray, 2020). Moreover, the modified model also predicts the effects of the highest luminance rule, glow, codetermination, and articulation as the original did (these results are shown at <https://>). Therefore, the modified model replicated all the successful results reported by the original model in Murray (2020).

Figure 2: The 12 stimuli used by Murray (2020). Green and red dots indicate the target areas.

Table 1: The modified (M) and original (O) models’ lightness (perceived reflectance) predictions for the target areas in the 12 figures employed by Murray (2020). “Left” and “Right” refer to the target areas indicated in Figure 2 by green and red dots, respectively.

	Left (M)	Right (M)	Left (O)	Right (O)
Argyle	0.68	0.41	0.68	0.41
Long range Argyle	0.68	0.41	0.68	0.41
Broken Argyle	0.41	0.41	0.41	0.41
Snake	0.80	0.48	0.62	0.48
Snake control	0.54	0.54	0.54	0.54
Koffka broken	0.66	0.51	0.66	0.51
Koffka–Adelson	0.66	0.51	0.74	0.51
Koffka connected	0.57	0.57	0.66	0.57
SLC	0.59	0.40	0.59	0.40
Articulated SLC	0.76	0.40	0.76	0.40
White’s effect	0.46	0.41	0.41	0.41
Checkerboard assimilation	0.41	0.41	0.36	0.68

For White’s effect and the Checkerboard assimilation, where the original model fails, we observed improvements in the modified model. The modified model correctly predicts White’s effect and sees clear illuminance stripes (Figure 3). We focused on the prediction of White’s effect in this study; thus, it will be further discussed in a later section.

(a) Modified model’s output

(b) Original model’s output

Figure 3: The outputs for the predictions of White’s effect by the two models.

The modified model also improved regarding the prediction of Checkerboard assimilation (Blakeslee & McCourt, 2004; De Valois & De Valois, 1990). It perceives the same lightness level for the two targets, whereas the original model sees the right target as much lighter. Humans perceive the left one slightly lighter (Murray, 2020), but the effect is dependent on the frequency of the checks (Blakeslee & McCourt, 2004). The modified model’s prediction is obviously closer to human perception, although it may not be complete. This difference between predictions by the two models is largely caused by modification of the belief propagation schedule. Similarly, it also largely reduced the energies (i.e., the model finds a better solution) in prediction for Articulated SLC. In summary, for the stimuli used by Murray (2020), the modified model provided several improvements without any performance deteriorations.

Influence of pixel-wise luminance noise

Betz et al. (2015) examined the ODOG model (Blakeslee & McCourt, 1999) by running it on White-effect figures with narrowband noise. They found that the influences of the narrowband noise on the ODOG model and humans do not match. Testing models with noised images is important to investigate whether the correct prediction can be robustly observed. However, using narrowband noise is difficult in our case because the models assume smaller images; thus, we created images with pixel-wise luminance noise based on Murray (2020)’s SLC figure (Figure 2). First, the 16×16 image was up-sampled to be a 32×32 image so that the influence of each pixel’s noise was moderate. Then, all pixel luminances were independently noised; random and independent samples from a uniform distribution were added to each pixel (the distribution’s expected value was zero, and the upper and lower limits were manipulated). One of the two targets was copied on the other target for both targets to be physically identical. We prepared five conditions by manipulating the magnitude of noise (i.e., the upper and lower limits of the uniform distributions of the noises) from one to five. Although we did not conduct an experiment on human observers, our informal observation confirmed a stable illusory effect even for an SLC image with the most noise (Figure 4 (a)).

We created ten randomly noised images for each noise magnitude condition (one to five); thus, 50 images were obtained. The two models were tested on these images (here, a single test was composed of ten runs, not 30 runs) (Figure 4 (b)). The results revealed that the original model is vulnerable to pixel-wise noise; it predicts virtually no illusion even for an image with the smallest noise. However, the modified model correctly predicts moderate illusions consistently for all conditions. The modified model’s output also shows a trend: the illusion magnitudes decrease with a larger noise, which human observers should observe as well. These results show that the original model’s correct prediction for SLC is less robust and that the modification to the model enhanced its stability against luminance noise. This improvement is mainly because of the modification of the belief propagation schedule.

Figure 4: (a) An example of a noisy image. In this example, the lower and upper limits of the uniform distribution are ± 5 . The SLC effect should be present in most observation environments. (b) Results of the test for the noise’s influence on the model predictions.

Simplified Checkershadow Illusion

We tested the models on images composed of consistent and inconsistent X-edges. We employed Blakeslee & McCourt (2012)’s Fig. 3b (Figure 5 (a) left) because it can be represented using a small simple image. This image can be understood as a simplification of Adelson (1995)’s famous Checkershadow Illusion. As in Adelson (1995)’s original image, the figure is composed of a checkerboard and a simulated shadow (or transparent filter), and lighter tiles in the shadow have a luminance that is physically equal to that of darker tiles out of the shadow. However, humans perceive the former as lighter (and brighter) compared with the latter (Blakeslee & McCourt, 2012). However, the illusory effect becomes much weaker or absent when the shadow is shifted to share its boundary with edges in the checkerboard (Figure 5 (b) left). We tested the two models on these two images (hereafter, we call them SCS, referring to the simplified Checkershadow Illusion and SCS control).

Figure 5 shows the results of the tests. For the SCS, both models correctly predict illusory effects with similar magnitudes. However, the illuminance map of the original model contains speckled noise, which does not match our intuitive illumination impression. The modified model outputs an illuminance map that better matches illumination impression by humans, containing a clear shadowed region. More importantly, the original model predicts an illusion for the SCS control with a magnitude similar to (or stronger than) that for the SCS, while the modified model does not. Thus, the modified model correctly predicts the presence and absence of the illusion for the SCS and SCS control, respectively, but the original model predicts the illusions for both. This difference in performance between the two models is obviously caused by the modification to the prior regarding X-junctions. The original model is likely to interpret inconsistent X-edges in the checkerboard as illuminance changes, whereas the modified model appropriately distinguishes inconsistent and consistent X-edges.

Figure 5: (a) SCS and both model outputs. Both models perceive the difference between the two target area reflectances (tiles with green and red frames), but the modified model perceives a shadowed region more clearly. (b) The SCS control and the models' outputs for it. The modified model does not observe a reflectance difference while the original model does.

Further tests for White's effect

As mentioned, the modified model's most important improvement is its ability to account for White's effect, which has been an important and challenging issue for lightness models and theories (Anderson et al., 2001; Bressan, 2006; Howe, 2005). We further tested our model on other variants of the White-effect figure.

Spehar et al. (1995) reported that White's effect becomes almost absent when the two luminance levels of the stripes are modified to both be brighter (Figure 6 (a) left) or both darker (Figure 6 (a) right) than the gray targets (i.e., the targets are in double decrements or double increments). We examined the model prediction by manipulating the luminance levels of the two surrounding stripes. In this test, both stripes' luminances were manipulated from 10 to 70 cd/m^2 with 2- cd/m^2 intervals (i.e., the brighter ones' luminance range from 12 to 70 cd/m^2 , and the darker ones' range from 10 to 68 cd/m^2). Conditions where the two stripes' luminances become equal or reversed were not included. The target luminance was fixed to that used in Murray (2020)'s original White's effect stimulus, which was 35 cd/m^2 .

The results are shown in Figure 6 (b, c) via heatmaps. Each row indicates the luminance levels of the darker stripes, and each column indicates those of the brighter stripes. The color of each block indicates the magnitude of the predicted illusion; the difference between the predicted lightness of the two targets (left - right). The modified model generally predicts the illusory effect when the target luminance is between the two stripes (an area surrounded by an orange frame; Figure 6 (b)) and predicts greater effects as the contrast between the stripes becomes larger. Moreover, the model rarely predicts the effect when the targets are in double increments or double decrements (outside of the orange frame; Figure 6 (b)). These results show the robustness of the correct predictions for White's effect and its control version (Howe, 2005; Spehar et al., 1995). Figure 6 (c) shows the results of the same test on the original model. The original model's performance for these stimuli is obviously worse than that of the modified one.

(a)

(b)

(c)

Figure 6: (a) A double increment version (top) and a double decrement (bottom) version of the White-effect figure. Both are considered to cause virtually no illusions. (b) Heatmap of the modified model's illusion prediction for the White-effect figure with various stripe luminance conditions. (c) Heatmap of the original model's prediction.

White’s effect is known to be enhanced when the stripe’s frequency becomes higher (Blakeslee & McCourt, 2004; White, 1979). This can be easily tested even on 16×16 sized images. We tested our model on three conditions (Figure 7) where the stripe’s frequencies were 2, 4, or 8 cycles/image. The luminance levels of each region were equal to those in Murray (2020)’s original figure (Figure 2). The results showed that the modified model correctly predicts the stronger effects for a figure of higher stripe frequency (Figure 7). The original model does not predict illusions for all of the three conditions (the results of the original model are shown in the online repository).

Figure 7: Effect of stripe frequency on the modified model’s output. The modified model predicts stronger illusions as the stripe frequency increased.

The modified model also succeeds in predicting human perceptions for some other variants of White’s effect. It has been suggested that White’s effect is enhanced as the target width (more specifically, length of the sides parallel to the stripes) becomes shorter (Blakeslee & McCourt, 2004; Kingdom & Moulden, 1991). The modified model predicts a slight enhancement for narrower targets (Figure 8 (a)). Moreover, Otazu et al. (2008) introduced a White-effect variant in which the effect is largely reduced when targets span three stripes. Via informal observation, we observed that the effect reversed rather than canceled, as shown in

Figure 8 (b). Our model predicts an illusion reversal from White’s original effect. Furthermore, the model correctly predicts no illusion for Figure 8 (c), which was introduced by Howe (2001) as a control of White’s effect. Predicting the absence of the effect in this control figure is challenging for models that consider T-junctions around the targets as the main cause of White’s effect. The present model does not explicitly take into account T-junctions in its processing, and thus, it can correctly predict the elimination of the effect. For the three White-effect variants discussed here, the original model does not predict any illusions (see the online repository for details).

Figure 8: Some variants of the White-effect figure and the modified model’s outputs for them. (a) A White-effect figure variant with narrow targets. The model correctly predicts illusion enhancement (stronger effect than in the original White-effect figure). (b) A White-effect figure with heightened targets. The model predicts a reversal of the effect. (c) Howe (2001)’s variant. The modified model correctly predicts no illusion.

Limitation of the modified model

We have shown several successful predictions by our modified model thus far, but we are also aware of the limitations of our model. First, our model does not predict the illusory effect in Anderson et al. (2001)’s configuration (Figure 9 (a)). Anderson et al. (2001) created this figure to show that White’s effect is recovered when targets are shifted from their positions in Howe

(2001)’s control figure (Figure 8 (c)). Thus, it is challenging for computational models to predict both the effect’s presence in Anderson et al. (2001)’s stimulus and the effect’s absence in Howe (2001)’s one. Indeed, spatial filtering models also fail at this (Robinson et al., 2007). Moreover, our model does not predict a stable illusion for Todorović (1997)’s White-effect variant (Figure 9 (b)). The model sometimes predicts the effect, as in the case of Figure 9 (b), but it is unstable because it is heavily dependent on random seeds (in this paper, we report results from a fixed random seed and showed this successful case in Figure 9 (b)). Our model also does not predict an illusion in Clifford & Spehar (2003)’s zigzag variant of White’s effect (Figure 9 (c))¹, which is reportedly accounted for by some spatial filtering models with certain settings of their parameters (Robinson et al., 2007; Zeman et al., 2015). The original model does not predict correct illusions for any of the three stimuli shown in Figure 9.

Figure 9: Variants of the White-effect figure and the modified model’s outputs for them. (a) Anderson et al. (2001)’s variant. The model does not predict an illusion for it. (b) Todorović (1997)’s variant. Although the model predicts a illusion for the case shown in this figure, it is highly unstable. (c) Clifford & Spehar (2003)’s variant. The modified model does not predict an illusion.

¹ Clifford & Spehar (2003) was a conference presentation; thus, their original figure was not available to us. We designed the stimulus based on Robinson et al. (2007).

Argyle and Snake illusion variants

In this section, we show mixed results in which the modified model sometimes yields better results than the original but sometimes the opposite results are obtained. We thought that the modification to the prior on X-junctions would affect the models’ predictions for variants of Argyle and Snake illusions. The appearance of Argyle and Snake illusions is highly dependent on X-junctions in these images, as is obvious in the “Broken Argyle” and “Snake control” in Figure 2, where the effects are absent when X-junctions were removed. Moreover, contrast polarities in X-junctions are crucial for the occurrence and directions of illusory effects (e.g., Bressan, 2001). Thus, modifying contrast polarities in X-junctions in Argyle and Snake illusion figures is expected to differentiate the original and modified models’ predictions and their appearances to human observers. We created some variants of these illusions (Figure 10) and tested the models on them.

The results are shown in Figure 10. Figure 10 (a), or Argyle control 1, is a variant made from the Argyle shown in Figure 2, by making the darkest regions of X-junctions moderately bright. In this figure, the illusory effect observed in the original Argyle is almost eliminated, and our modified model predicts no illusion. However, the original model still predicts a moderate illusion and sees illuminance stripes (see <https://> for the output maps). When the X-junctions’ brightest regions in Argyle are made darker (Figure 10 (b); Argyle control 2), the illusory appearance is also eliminated, and both models do not predict illusion.

For variants of the Snake illusion, in Figure 10 (c) (Snake control 1), where the darkest regions in X-junctions of the original Snake-illusion figure (Figure 2) are made brighter, the illusory effect becomes weaker but is still observable because of SLC. In this figure, while the original model correctly predicts a weak illusion, the modified model predicts no illusion. For Figure 10 (d) (Snake control 2), where the brightest parts of the X-junctions are darkened, the modified model does not predict the illusion. The original model predicted the effect, but its magnitude is almost the same as or stronger than that for the original Snake illusion, which does not match the subjective strength of the effect.

Figure 10: Variants of Argyle and Snake illusion figures and the models’ predictions. Green and red frames indicate the target areas. (a) Argyle control 1; the darkest pixels in the original Argyle (Figure 2) are moderately brightened. (b) Argyle control 2; the brightest pixels are darkened. (c) Snake control 1; the darkest pixels in the original Snake (Figure 2) are brightened. (d) Snake control 2; the brightest pixels are darkened.

In summary, for variants of Argyle and Snake illusions (Figure 10), the results were mixed. The modified model shows better predictions than the original model in some cases (mainly in variants of the Argyle illusion), and the original model shows better predictions than the modified model in some cases (mainly in variants of the Snake illusion). Many more variants that can differentiate the two models’ predictions can be created; thus, these tests are far from exhaustive. Because further tests and comparisons with human data are needed, we cannot make a definite conclusion here on this issue.

Discussion

In this study, we introduced a modified version of Murray (2020)’s MIR. We modified the original model in the following three aspects: (1) the belief propagation was made more efficient, (2) different potential functions are assigned for inconsistent and consistent X-edges, and (3) the parameter for illuminance changes are flexibly determined based on the total luminance contrast of the input images. The first aspect is a technical improvement of the inference process, and the latter two are modifications to the model’s priors. The priors should reflect observers’ beliefs about scenes, and we believe that the two modifications largely reflect them. These “natural” priors showed significant improvements in the model performance without any deterioration. The main results obtained in this study are summarized in Table 2.

Table 2: Summary of the model tests conducted in this study. Items that were already shown in Table 1 and do not differentiate the two models’ predictions are not included here. “C” refers to correct predictions, “W” refers to wrong ones, and “I” refers to intermediate ones. WE stands for White’s effect. Given that the original model does not predict the original White’s effect, its predictions of the absences of the effect for Spehar’s variant, the target height effect, and Howe’s variant were not considered successful; thus, we marked NA instead of C.

	Modified	Original
Checkerboard assimilation	I	W
White’s effect	C	W
Noise on SLC (Figure 4)	C	W
SCS (Figure 5 (a))	C	I
SCS control (Figure 5 (b))	C	W
Spehar’s WE variant (Figure 6 (a))	C	NA
Frequency’s effect on WE (Figure 7)	C	W
Target width in WE (Figure 8 (a))	C	W
Target height in WE (Figure 8 (b))	I	NA
Howe’s WE variant (Figure 8 (c))	C	NA
Anderson’s WE variant (Figure 9 (a))	W	W
Todorović’s WE variant (Figure 9 (b))	I	W
Clifford’s WE variant (Figure 9 (c))	W	W
Argyle control 1 (Figure 10 (a))	C	W
Argyle control 2 (Figure 10 (b))	C	C
Snake control 1 (Figure 10 (c))	W	C
Snake control 2 (Figure 10 (d))	W	W

Specifically, the main improvements are as follows: (1) the prediction for the Checkerboard assimilation figure is closer to human perception, (2) the prediction of SLC is more stable and immune to luminance noise influence, (3) the predicted lightness and illuminance maps for the SCS and SCS control (Blakeslee & McCourt, 2012) are much closer to human perception, and (4) the model correctly predicts White’s effect. The prediction of White’s effect is challenging and of high importance for lightness models because this effect highlights that lightness cannot be explained solely by the contrast of adjacent areas (Wallach, 1948). White’s effect has long been known as a problem for simple contrast-based theories and has been discussed by many lightness/brightness researchers (Anderson et al., 2001; Betz et al., 2015; Gilchrist, 2006; Howe, 2005; Spehar et al., 1995; Taya et al., 1995). Although there have been various qualitative explanations of White’s effect (e.g., Howe, 2001), not many quantitative (i.e., computational) models have correctly explained it (Blakeslee & McCourt, 1999; Lerer et al., 2021). Therefore, it is a highly important for computational models to incorporate the prediction of White’s effect. We not only showed our model’s ability to predict this effect in a single image but also examined several variants of a White-effect figure. The

elimination of White’s effect in double increment/decrement conditions (Spehar et al., 1995) has rarely been predicted by computational models, but our model succeeds in predicting the presence and absence of the effect for figures whose stripe luminance levels are manipulated (Figure 6). It also predicts the effect’s enhancement caused by higher stripe frequency (Figure 7) and that in narrower targets (Figure 8 (a)). Moreover, our model predicts a reversed effect in a variant in which targets span several stripes (Figure 8 (b)), which matches our informal observation. Additionally, it correctly predicts the absence of the effect in Howe (2001)’s control figure (Figure 8 (c)), which should be predicted to yield an illusion by models that depend on processing of T-junctions (Todorović, 1997). Failures in the prediction of reverse contrast phenomena have been a major weakness of the original MIR (Murray, 2020; Nedimović et al., 2021), but this study proved that it is possible for the MIR framework to overcome this weakness.

Although our model showed significant improvements, there are still limitations. For example, the model is not able to account for Todorović (1997)’s illusion (Figure 9 (b)). This may be caused by the same mechanism as that for White’s effect (Gilchrist, 2006); however, our model failed to predict it. Future models should address this shortcoming by implementing additional priors, such as that on T-junctions. However, we note that none of the computational models currently available have explained all the known variants of White’s effect thus far (Blakeslee & McCourt, 1999; Robinson et al., 2007; Zeman et al., 2015).

Although we believe that our model’s additional two priors are natural (i.e., they largely reflect natural scene characteristics), their further clarification may be possible. X-junctions have been studied and the importance of the contrast polarity has been widely known (Anderson, 1997; Kitaoka, 2005). We simply distinguished X-junctions into two patterns, but a more detailed evaluation of X-junctions can likely be implemented. This may improve the model’s performance for illusions related to transparency or Argyle and Snake illusions. Regarding the prior on the expectation of illumination changes, future studies may need to consider accumulated knowledge of the illumination perception of humans (e.g., Murray & Adams, 2019; Pont & Koenderink, 2007). Although our model simply expects a monotonic increase in illuminance change’s likelihood as the total luminance contrast increases, more complex relationship between the two variables may better reflect the illumination interpretation of humans. Incorporating “natural” priors will enhance MIR’s performance. In this study, we proposed one possible modification to MIR, but it requires further studies and improvements.

To our knowledge, this study is the first extension of Murray (2020)’s original MIR. This study highlights MIR’s high extensibility and potential; the MIR framework welcomes various patterns of modifications based on natural scene statistics and knowledge about human vision. We improved the model’s performance by implementing two additional priors and optimizing the inference. In future studies, perceptual grouping, which is known as an important factor of reverse contrasts (Agostini et al., 2020; e.g., Gilchrist et al., 1999), should be implemented into the model. It may not be simple to represent perceptual grouping in CRFs, but computational methods that represent perceptual grouping, which have been recently presented, may be useful (Froyen et al., 2015; Lezama et al., 2016). These recent developments of computational

methods should be highly beneficial for the MIR framework, computational models, and rigorous understanding of human vision.

Acknowledgments

Y.K. received financial support from the Japan Society for the Promotion of Science (grant number: 20J00606).

Appendix

Open sources

We implemented the original MIR and our modified version in two languages: Python and Julia (<https://>). They are both based on Murray (2020)’s MATLAB code. Both Python and Julia are free, so anyone can test Murray (2020)’s model and our models using them. The Python version may be familiar to more users, but the Julia version runs much faster.

Modification of message passing

Here, we provide the details of the modification we made to the belief propagation. MIR adopts belief propagation with a Bethe cluster graph for inference. MIR’s graph is composed of 1-nodes and 4-nodes; the former represents each node in the CRF hidden layers (i.e., the illuminance of each image pixel), and the latter represents the potential functions for each 4-clique composed of four 1-nodes. Figure 11 (a) shows an illustration of MIR’s cluster graph.

In MIR’s belief propagation, messages (information about the probability of possible states of each node) are passed in vertical and horizontal trees. Message passing in trees is illustrated in Figure 11 (b). In vertical trees, messages start at the 1-nodes in the top row, are passed to neighboring 4-nodes on the southeast (except in the leftmost vertical tree), and then are passed to 1-nodes on the southwest, and these steps are repeated. Once messages arrive at the bottom, backward passing starts in which messages go in the same trees from bottom to top. Likewise, in horizontal trees, messages start at the 1-nodes in the leftmost column, are passed to 4-nodes on the northeast, and then to 1-nodes on the southwest, and these steps are repeated (except for messages in the top horizontal tree). Once messages arrive at the right end, backward passing from right to left starts. For a 16×16 -sized image, messages are passed in 32 trees (i.e., 16 vertical and 16 horizontal trees), and these 32 passing patterns are randomly ordered.

In this belief propagation, most of the 4-nodes, except for those in the rightmost column or the bottom row, do not exchange messages with their neighboring 1-nodes on the southeast, as shown in Figure 11 (b). To utilize all the links, we added one more phase of message passing after the original passing (Figure 11 (c)). In this additional second phase, messages in vertical and horizontal trees go in directions that are different from those in the original phase (Figure 11 (b)), except in the columns and rows on the CRF’s edges. In most of the vertical trees, messages start from the top 1-nodes to the 4-nodes on the southwest, whereas it was to the 4-nodes on the southeast in the original passing phase. Similarly, in most of the horizontal trees,

messages start from the leftmost 1-nodes to 4-nodes on the northeast, while it was to the southeast in the original phase. In the additional phase, also the passing order of the 32 trees is randomized. The additional phase starts after completing the original passing phase.

Figure 11: (a) Illustration of MIR's cluster graph. This example shows a cluster graph of a 4 x 4 image. (b) Message passing in vertical (left) and horizontal (right) trees in the original message passings. Arrows illustrate message directions in forward passing. Messages go in the reverse direction in backward passing. (c) Message passing in vertical (left) and horizontal (right) trees in the additional second phase.

References

- Adelson, E. H. (1995). *Checkershadow Illusion*. <http://persci.mit.edu/gallery/checkershadow>
- Adelson, E. H. (2000). Lightness Perception and Lightness Illusions. In M. Gazzaniga (Ed.), *The New Cognitive Neurosciences* (pp. 339–351). MIT Press.
- Agostini, T., & Galmonte, A. (2002a). Perceptual Organization Overcomes the Effects of Local Surround in Determining Simultaneous Lightness Contrast. *Psychological Science*, 13(1), 89–93. <https://doi.org/10.1111/1467-9280.00417>
- Agostini, T., & Galmonte, A. (2002b). A new effect of luminance gradient on achromatic simultaneous contrast. *Psychonomic Bulletin & Review*, 9(2), 264–269. <https://doi.org/10.3758/BF03196281>
- Agostini, T., Murgia, M., Sors, F., Prpic, V., & Galmonte, A. (2020). Contrasting a Misinterpretation of the Reverse Contrast. *Vision*, 4(4). <https://doi.org/10.3390/vision4040047>
- Allred, S. R., & Brainard, D. H. (2013). A Bayesian model of lightness perception that incorporates spatial variation in the illumination. *Journal of Vision*, 13(7), 1–18. <https://doi.org/10.1167/13.7.18>

- Anderson, B. L. (1997). A Theory of Illusory Lightness and Transparency in Monocular and Binocular Images: The Role of Contour Junctions. *Perception*, 26(4), 419–453. <https://doi.org/10.1068/p260419>
- Anderson, B. L., Todorović, D., & Grossberg, S. (2001). Contrasting theories of White's illusion. *Perception*, 30(12), 1499–1507. <https://doi.org/10.1068/p3012ds>
- Betz, T., Shapley, R., Wichmann, F. A., & Maertens, M. (2015). Noise masking of White's illusion exposes the weakness of current spatial filtering models of lightness perception. *Journal of Vision*, 15(14), 1. <https://doi.org/10.1167/15.14.1>
- Blakeslee, B., & McCourt, M. E. (2012). When is spatial filtering enough? Investigation of brightness and lightness perception in stimuli containing a visible illumination component. *Vision Research*, 60, 40–50. <https://doi.org/10.1016/j.visres.2012.03.006>
- Blakeslee, B., & McCourt, M. E. (1997). Similar mechanisms underlie simultaneous brightness contrast and grating induction. *Vision Research*, 37(20), 2849–2869. [https://doi.org/10.1016/S0042-6989\(97\)00086-2](https://doi.org/10.1016/S0042-6989(97)00086-2)
- Blakeslee, B., & McCourt, M. E. (1999). A multiscale spatial filtering account of the White effect, simultaneous brightness contrast and grating induction. *Vision Research*, 39(26), 4361–4377. [https://doi.org/10.1016/S0042-6989\(99\)00119-4](https://doi.org/10.1016/S0042-6989(99)00119-4)
- Blakeslee, B., & McCourt, M. E. (2004). A unified theory of brightness contrast and assimilation incorporating oriented multiscale spatial filtering and contrast normalization. *Vision Research*, 44(21), 2483–2503.
- Bressan, P. (2001). Explaining Lightness Illusions. *Perception*, 30(9), 1031–1046. <https://doi.org/10.1068/p3109>
- Bressan, P. (2006). The place of white in a world of grays: A double-anchoring theory of lightness perception. *Psychological Review*, 113(3), 526–553. <https://doi.org/10.1037/0033-295X.113.3.526>
- Clifford, C. W. G., & Spehar, B. (2003). Using colour to disambiguate contrast and assimilation in White's Effect. *Journal of Vision*, 3(9), 294–294. <https://doi.org/10.1167/3.9.294>
- Dakin, S. C., & Bex, P. J. (2003). Natural image statistics mediate brightness 'filling in'. *Proceedings. Biological Sciences*, 270(1531), 2341–2348. <https://doi.org/10.1098/rspb.2003.2528>
- De Valois, R. L., & De Valois, K. K. (1990). *Spatial vision*. Oxford University Press. <https://ebookcentral.proquest.com/lib/ritsumei/detail.action?docID=3051914>
- Economou, E., Zdravkovic, S., & Gilchrist, A. L. (2007). Anchoring versus spatial filtering accounts of simultaneous lightness contrast. *Journal of Vision*, 7(12), 2. <https://doi.org/10.1167/7.12.2>
- Economou, E., Zdravković, S., & Gilchrist, A. L. (2015). Grouping Factors and the Reverse Contrast Illusion. *Perception*, 44(12), 1383–1399. <https://doi.org/10.1177/0301006615607118>
- Feldman, J. (2015). Bayesian models of perceptual organization. In *The Oxford Handbook of Perceptual Organization* (pp. 1008–1026). Oxford University Press.

- Froyen, V., Feldman, J., & Singh, M. (2015). Bayesian hierarchical grouping: Perceptual grouping as mixture estimation. *Psychological Review*, 122(4), 575–597. <https://doi.org/10.1037/a0039540>
- Gilchrist, A. L. (2006). *Seeing black and white*. Oxford University Press.
- Gilchrist, A. L., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., Spehar, B., Annan, V., & Economou, E. (1999). An anchoring theory of lightness perception. *Psychological Review*, 106(4), 795–834. <https://doi.org/10.1037/0033-295X.106.4.795>
- Howe, P. D. L. (2005). White's Effect: Removing the Junctions but Preserving the Strength of the Illusion. *Perception*, 34(5), 557–564. <https://doi.org/10.1068/p5414>
- Howe, P. D. L. (2001). A comment on the Anderson (1997), the Todorović (1997), and the Ross and Pessoa (2000) explanations of White's effect. *Perception*, 30(8), 1023–1026. <https://doi.org/10.1068/p3212>
- Kingdom, F. A. A. (2011). Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research*, 51(7), 652–673. <https://doi.org/10.1016/j.visres.2010.09.012>
- Kingdom, F. A. A., & Moulden, B. (1991). White's effect and assimilation. *Vision Research*, 31(1), 151–159. [https://doi.org/10.1016/0042-6989\(91\)90082-G](https://doi.org/10.1016/0042-6989(91)90082-G)
- Kitaoka, A. (2005). *Connecting the X-junction contrast-polarity model with the luminance-based arithmetic model*. 13.
- Land, E. H., & McCann, J. J. (1971). Lightness and Retinex Theory. *Journal of the Optical Society of America*, 61(1), 1–11. <https://doi.org/10.1364/JOSA.61.000001>
- Lerer, A., Supér, H., & Keil, M. S. (2021). Dynamic decorrelation as a unifying principle for explaining a broad range of brightness phenomena. *PLOS Computational Biology*, 17(4), e1007907.
- Lezama, J., Randall, G., Morel, J.-M., & Grompone von Gioi, R. (2016). Good continuation in dot patterns: A quantitative approach based on local symmetry and non-accidentalness. *Quantitative Approaches in Gestalt Perception*, 126, 183–191. <https://doi.org/10.1016/j.visres.2015.09.004>
- Murray, R. F. (2013). *Human lightness perception is guided by simple assumptions about reflectance and lighting* (B. E. Rogowitz, T. N. Pappas, & H. de Ridder, Eds.; p. 865106). <https://doi.org/10.1117/12.2001782>
- Murray, R. F. (2020). A model of lightness perception guided by probabilistic assumptions about lighting and reflectance. *Journal of Vision*, 20(7), 28. <https://doi.org/10.1167/jov.20.7.28>
- Murray, R. F., & Adams, W. J. (2019). Visual perception and natural illumination. *Current Opinion in Behavioral Sciences*, 30, 48–54. <https://doi.org/10.1016/j.cobeha.2019.06.001>
- Nedimović, P., Zdravković, S., & Domijan, D. (2021, August). *Empirical evaluation of computational models of lightness perception*. osf.io/cv9uk
- Otazu, X., Vanrell, M., & Párraga, C. A. (2008). Multiresolution wavelet framework models brightness induction effects. *Vision Research*, 48(5), 733–751. <https://doi.org/10.1016/j.visres.2007.12.008>

- Pont, S. C., & Koenderink, J. J. (2007). Matching illumination of solid objects. *Perception & Psychophysics*, 69(3), 459–468. <https://doi.org/10.3758/BF03193766>
- Robinson, A. E., Hammon, P. S., & Sa, V. R. de. (2007). Explaining brightness illusions using spatial filtering and local response normalization. *Vision Research*, 47(12), 1631–1644. <https://doi.org/10.1016/j.visres.2007.02.017>
- Shapiro, A., & Lu, Z.-L. (2011). Relative Brightness in Natural Images Can Be Accounted for by Removing Blurry Content. *Psychological Science*, 22(11), 1452–1459. <https://doi.org/10.1177/0956797611417453>
- Spehar, B., Gilchrist, A., & Arend, L. (1995). The critical role of relative luminance relations in White's effect and grating induction. *Vision Research*, 35(18), 2603–2614. [https://doi.org/10.1016/0042-6989\(95\)00005-K](https://doi.org/10.1016/0042-6989(95)00005-K)
- Taya, R., Ehrenstein, W. H., & Cavanus, C. R. (1995). Varying the Strength of the Munker—White Effect by Stereoscopic Viewing. *Perception*, 24(6), 685–694. <https://doi.org/10.1068/p240685>
- Todorović, D. (1997). Lightness and Junctions. *Perception*, 26(4), 379–394. <https://doi.org/10.1068/p260379>
- Wallach, H. (1948). Brightness constancy and the nature of achromatic colors. *Journal of Experimental Psychology*, 38(3), 310.
- White, M. (1981). The Effect of the Nature of the Surround on the Perceived Lightness of Grey Bars within Square-Wave Test Gratings. *Perception*, 10(2), 215–230. <https://doi.org/10.1068/p100215>
- White, M. (1979). A New Effect of Pattern on Perceived Lightness. *Perception*, 8(4), 413–416. <https://doi.org/10.1068/p080413>
- Zeman, A., Brooks, K. R., & Ghebreab, S. (2015). An exponential filter model predicts lightness illusions. *Frontiers in Human Neuroscience*, 9. <https://doi.org/10.3389/fnhum.2015.00368>