An exponential filter model predicts lightness illusions

Astrid Zeman [1, 2, 3*], Kevin R. Brooks [3, 4] and Sennay Ghebreab [5, 6]

[1] Department of Cognitive Science, ARC Centre of Excellence in Cognition and its Disorders, Macquarie University, Sydney, NSW, Australia, [2] Commonwealth Scientific and Industrial Research Organisation, Marsfield, NSW, Australia, [3] Perception in Action Research Centre, Macquarie University, Sydney, NSW, Australia, [4] Department of Psychology, Macquarie University, Sydney, NSW, Australia, [5] Cognitive Neuroscience Group, Department of Psychology, University of Amsterdam, Amsterdam, Netherlands, [6] Intelligent Systems Lab Amsterdam, Institute of Informatics, University of Amsterdam, Amsterdam, Netherlands

Lightness, or perceived reflectance of a surface, is influenced by surrounding context. This is demonstrated by the Simultaneous Contrast Illusion (SCI), where a gray patch is perceived lighter against a black background and vice versa. Conversely, assimilation is where the lightness of the target patch moves toward that of the bounding areas and can be demonstrated in White's effect. Blakeslee and McCourt (1999) introduced an oriented difference-of-Gaussian (ODOG) model that is able to account for both contrast and assimilation in a number of lightness illusions and that has been subsequently improved using localized normalization techniques. We introduce a model inspired by image statistics that is based on a family of exponential filters, with kernels spanning across multiple sizes and shapes. We include an optional second stage of normalization based on contrast gain control. Our model was tested on a well-known set of lightness illusions that have previously been used to evaluate ODOG and its variants, and model lightness values were compared with typical human data. We investigate whether predictive success depends on filters of a particular size or shape and whether pooling information across filters can improve performance. The best single filter correctly predicted the direction of lightness effects for 21 out of 27 illusions. Combining two filters together increased the best performance to 23, with asymptotic performance at 24 for an arbitrarily large combination of filter outputs. While normalization improved prediction magnitudes, it only slightly improved overall scores in direction predictions. The prediction performance of 24 out of 27 illusions equals that of the best performing ODOG variant, with greater parsimony. Our model shows that V1-style orientation-selectivity is not necessary to account for lightness illusions and that a low-level model based on image statistics is able to account for a wide range of both contrast and assimilation effects.

Keywords: exponential, filter, model, ODOG, lightness, illusion, contrast, assimilation

1. Introduction

Lightness is the perceived reflectance of a surface, which can vary greatly according to surrounding context, as demonstrated by lightness illusions (see Kingdom, 2011 for a recent review). One clear and well-known example is the Simultaneous Contrast Illusion (SCI), where a gray target patch is perceived as lighter when surrounded

FIGURE 1 | Simultaneous Contrast vs. White's Effect. Albedo of gray target patches in Simultaneous Contrast shift away from background, demonstrating contrast. Targets in White's Effect shift toward surrounding context, demonstrating assimilation. Increasing spatial frequency increases the effect in both cases.
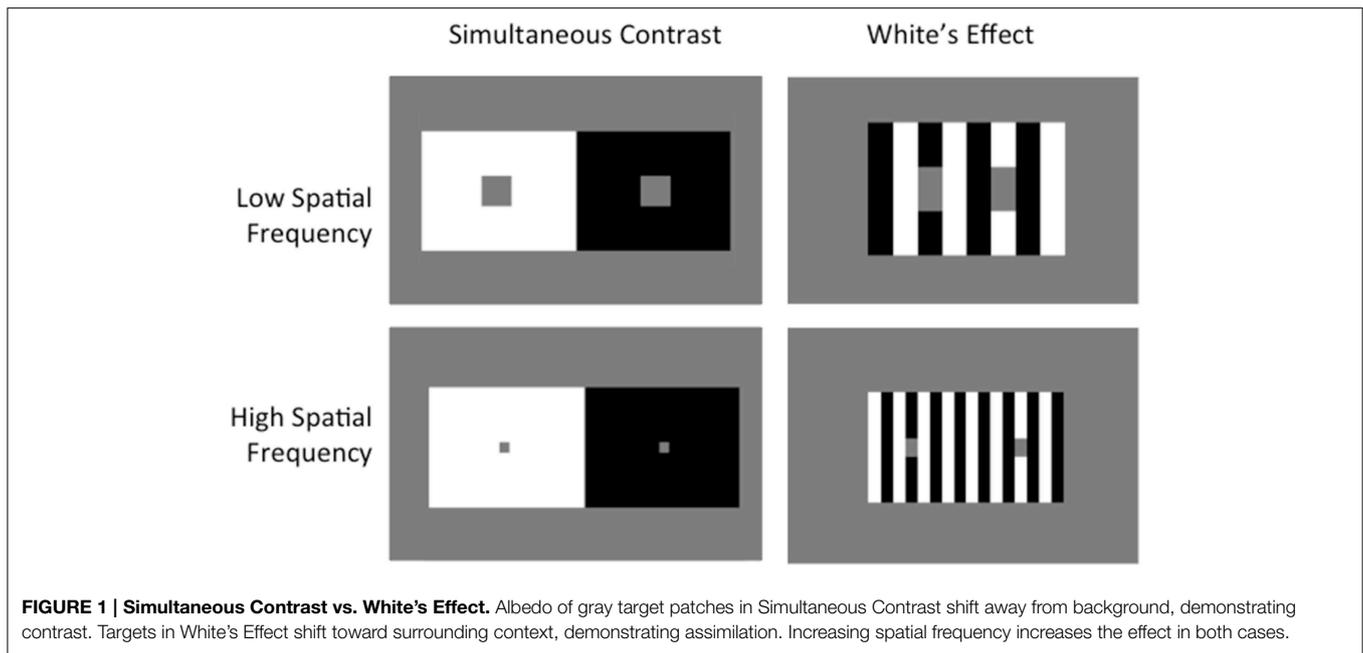
by a black background and darker when surrounded by a white background (Chevreul, 1839) (**Figure 1** left). The SCI demonstrates the contrast phenomenon, where lightness shifts away from surrounding luminance values, luminance being the amount of light that reaches the eye. Under other circumstances, lightness can shift toward the luminance values of bordering areas—a phenomenon known as assimilation[1]. This is effectively demonstrated by a version of White's Illusion (White, 1979), where the test patches are not as wide as they are tall (**Figure 1** right).

Theories that aim to explain lightness illusions can be broadly categorized into low-level and higher-level accounts. Higher-level theories argue that scene interpretation is necessary to account for lightness illusions, where cortical processing of surface curvature, depth, and transparency are known to influence perceived reflectance (Knill and Kersten, 1991). For instance, Schirillo et al. (1990) demonstrated that lightness perception is dependent upon depth cues. Given that depth perception is thought to be a cortical function, higher-level areas must be recruited when perceiving reflectance. In 1999, Gilchrist et al. (1999) established the Anchoring Theory of lightness, where perceived reflectance of a patch is "anchored" to the highest luminance value within the retinal image (global information) and is also "anchored" to luminance values in surface groups that share commonalities such as being situated within the same depth plane (local information). Another notable high-level theory is Anderson (1997)'s Scission Theory, based upon the principle that a visual scene is split into different causal layers of reflectance, transparency, and illumination (the amount of light incident on a surface), to determine the surface properties of

a homogenous area. While these high-level theories are able to offer consistent explanations for a variety of complex lightness phenomena, our aim in this paper is to quantify the performance of low-level models whose computations do not require higher-level scene interpretation. In the interests of providing a succinct quantitative account of a range of lightness phenomena, we apply Occam's Razor, emphasizing the capability of low-level theories to deliver improved modeling precision with greater parsimony.

Low-level theories concentrate on filtering operations and statistical image properties as the key explanation behind many lightness illusions. The main principle underlying low-level theories is that of image reconstruction: that lightness is inferred by reconstructing the most probable source image using filtering operations (Blakeslee and McCourt, 1999; Dakin and Bex, 2003). The filters concerned are considered to reside in early stages of the visual hierarchy such as the retina, LGN, and/or V1. Blakeslee and McCourt (1997) designed a low-level model using a multi-scale array of two-dimensional Difference of Gaussian filters (DOG). The isotropic filters in this model approximated retinal ganglion or LGN single cell function. The DOG model was able to account for the contrast effect shown in the SCI but not the assimilation observed in White's Effect. To account for assimilation, Blakeslee and McCourt (1999) extended this model to include anisotropic filters (oriented difference of Gaussians, or ODOG filters) that were pooled non-linearly. These orientation selective filters best approximate V1 functions, shifting the focus of the model from pre-cortical to cortical operations to account for a larger set of lightness illusions. Shortly after this, Dakin and Bex (2003) introduced an isotropic filter model that reweighted filter outputs using spatial frequency (SF) properties found in image statistics. Using a series of center-surround, Laplacian of Gaussian filters, they demonstrated that low SF structure is an essential ingredient of two well-known lightness illusions: White's Effect and the Craik-Cornsweet-O'Brien Effect (O'Brien,

---

[1]In some cases, target patches have equal bordering white and black areas, making it difficult to distinguish whether a contrast or assimilation effect is predominantly present.
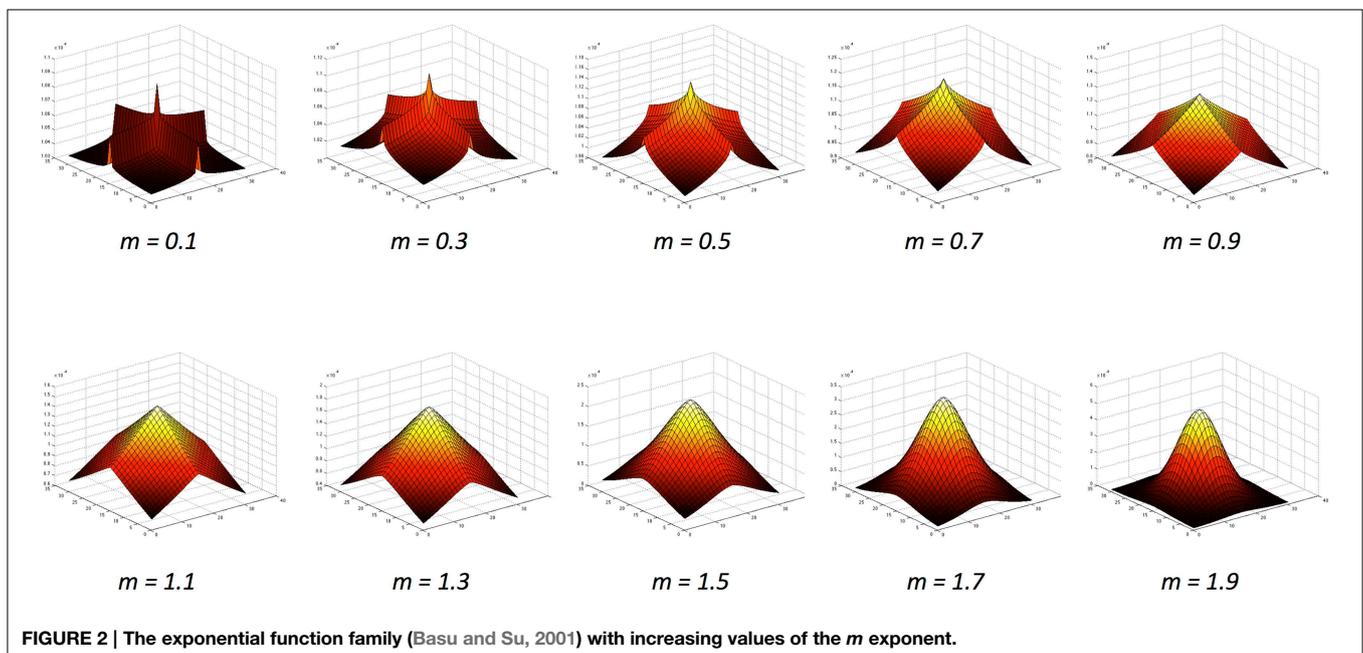
1958; Craik, 1966; Cornsweet, 1970). Dakin and Bex (2003) demonstrated that orientation selective filters were not required to successfully model assimilation effects, and highlighted the importance of weighting or normalization schemes within these low-level models.

Since Dakin and Bex's paper, focus on statistical image properties (Corney and Lotto, 2007) and on post-filtering operations that weight the relative filter outputs (Robinson et al., 2007) has intensified in the context of low-level lightness models. Corney and Lotto (2007) demonstrated contrast and assimilation effects using an approach inspired by image statistics, training an artificial neural network with virtual scenes that possess naturalistic structure. In contrast to Dakin and Bex (2003) who made statistical relationships explicit through weighting operations, Corney and Lotto (2007) trained an artificial neural network to implicitly learn the relationships between images and their underlying statistics. In the same year, Robinson et al. (2007) focused on applying different normalization schemes to improve predictions using the ODOG model. Normalization is commonly used as a weighting scheme to smooth distributions and scale all values to a baseline magnitude (usually 1). Robinson et al. (2007) focused on applying two different normalization schemes to the ODOG model: local normalization of filter outputs (LODOG) and spatial frequency-specific local normalization (FLODOG). In LODOG and FLODOG, parameters of the normalization function (such as normalization window size) were adjusted to produce different model predictions. Robinson et al. (2007) systematically tested ODOG, LODOG, and FLODOG on a catalog of 28 stimuli, 27 of which are known to induce illusions of contrast or assimilation in human observers. While ODOG was able to predict only 13 illusions in the correct direction, the best performing LODOG model was able to predict 18. FLODOG proved the most effective, correctly predicting 24 lightness illusions with an optimal parameter set.

Here we extend the literature using an approach inspired by natural image statistics. As established by Dakin and Bex (2003), the underlying distribution of structural properties present in natural images can greatly influence lightness judgments. Natural images share common underlying statistics, regardless of their origin (Zhu and Mumford, 1997a,b). For example, contrast histograms for natural images are skewed toward lower contrasts and have an exponential tail (Field, 1987; Ruderman and Bialek, 1994). Basu and Su (2001) investigated filters that encode the distribution of contrasts over different spatial frequencies. They concluded that exponential distributions provide a better fit in representing the underlying power distributions of natural images than the Gaussian kernels that have been used in the models described above. By employing exponential filters of different sizes and shapes within a computational model, we represent the profile of contrast statistics present in natural images and observe how these may influence the direction and magnitude of a set of lightness illusions. These filters have x- and y-axis symmetry, ranging from ridged, 'peaky' distributions to flatter, more rounded distributions (illustrated in **Figure 2**).

The exponential filters we explore in this study are offered as another kind of inhibitory mechanism, since the image filtered by the exponential function is subtracted from the original image. As such, this model shares much in common with other filtering approaches, such as ODOG (Blakeslee and McCourt, 1999). Indeed, this filtering approach bears similarity to the extra classical surround model of Ghosh et al. (2006) and is most similar to the filtering approach of Shapiro and Lu (2011), with the exception of the shape of the surround.

While the filters in ODOG (and variants) approximate the functioning of orientation-selective V1 cells, and while Difference or Laplacian of Gaussian filters approximate the operations of isotropic LGN or retinal ganglion cells, exponential filters, not unlike those forming the basis of our model, have been

FIGURE 2 | The exponential function family (Basu and Su, 2001**) with increasing values of the *m* exponent.**

identified in H1 horizontal retinal cells (Packer and Dacey, 2002, 2005). Our model is predominantly motivated by the computer vision literature, where exponential filters have been shown to be excellent edge detectors as well as resilient to noise (Zhu and Mumford, 1997a). The level of biological plausibility in our model is not strongly emphasized, but we do identify possible neurobiological equivalents to the filters that we apply. Geisler (2008) illustrates responses to natural images of a sensor that has a receptive field profile similar to V1, where an exponential function shows a better fit over a Gaussian distribution. While there are parallels here in demonstrating that an exponential fit is better than Gaussian in terms of filter responses, the filters that we apply are not oriented V1-style filters. Therefore, we would not suggest any relationship between our model results and the involvement of cortical neurons.

Our study differs from that of Corney and Lotto (2007) in that we make statistical relationships explicit through filtering and normalization operations, instead of training an artificial neural network to implicitly learn the relationships between images and their underlying statistics. Our method is similar to that of Dakin and Bex (2003), in that we both capitalize on the properties of image statistics to reconstruct the final image. In our method, we employ exponential shape filters that are based on image statistics. In Dakin and Bex (2003), the authors split an image into different spatial frequencies (SFs) using band-pass filters. The distribution of SFs was then reweighted to match that which occurs in natural scenes. In our model and in that of Dakin and Bex (2003), the filters are designed to extract the most salient features while being robust to noise (Basu and Su, 2001). In this way, both of our studies align with the predictive coding principle by Srinivasan et al. (1982)—that by exploiting the spatial correlations of natural scenes, early visual systems are much better able to handle noise in the environment.

In the current study, we set out to investigate how well an exponential model is able to predict human data in response to a large battery of 28 lightness illusions previously used to test ODOG and its derivatives (Blakeslee and McCourt, 1999, 2001, 2004; Blakeslee et al., 2005; Robinson et al., 2007). We apply exponential filters with a range of different shapes and sizes to an input image, with and without normalization of varying spatial extent. The outputs of this model are taken as predictions of lightness both for single filters and for multiple-filter combinations.

2. Materials and Methods

2.1. Stimuli
A standard battery of 28 figures known to produce particular lightness effects was used as a stimulus set in this study (see Robinson et al., 2007). Each stimulus (with the exception of the Benary Cross) involves a pair of uniform, mean luminance target patches, each surrounded by details with the opposite contrast polarity. Stimuli are illustrated in **Figure 3**, reproduced from Robinson et al. (2007). All stimuli are 512 × 512 pixels in size. Each stimulus is listed below in **Table 1** with original sources and comparative results reported for human responses where available. **Table 1** also includes the reported illusion direction by

humans as the patch perceived as the lightest within the image and the corresponding classification of the predominant effect as contrast or assimilation.

The majority of images exhibit assimilation effects, with contrast effects demonstrated by figures *n, o, p, q, w,* and *x*. In some cases, target patches have equal bordering white and black areas, making it difficult to establish whether a lightness effect should be defined as a contrast or assimilation effect (as in stimulus *s*). Stimuli *y* and *z* demonstrate opposing illusion directions for patches with identical bordering surrounds, presenting both contrast and assimilation effects simultaneously. In most cases, illusory effect directions reported in the original articles have been replicated in follow-up studies by Blakeslee and McCourt (used here and in Robinson et al., 2007 for direct strength comparisons). However, due to slight differences in methodology, stimuli *e* and *q* demonstrate discrepancies between the two sets of human data. In these cases, we follow the convention of Robinson et al. (2007) to allow for easy comparison between their models and those described here.

As each stimulus involves 2 (or more) uniform, mean luminance target patches, each surrounded by details with the opposite contrast polarity, the lightness effects observed on these patches are expected to be equal and opposite. Our model's predictions regarding the presence of contrast or assimilation effects are made by taking mean lightness values from the largest rectangular patch inside the bounds of the target areas (matched for size) and subtracting the values for the patch that appears darker from those for the lighter. For stimulus *n*, ("grating induction"), we select rectangular areas that are 26 pixels wide to the left and right of center for our prediction comparison (0.4 of the spatial period of the grating), while maintaining the same patch height as Robinson et al. (2007).

2.2. Model
Our model consists of two-stages: (1) linear filtering using exponential functions (2) non-linear divisive normalization by coefficient of variation. Although the details of each stage may vary, this linear-nonlinear modeling method is commonly used to model physiology (Schwartz and Simoncelli, 2001; Nykamp and Ringach, 2002). Once the two stages of the model have produced lightness values at each pixel location of each target patch, we produce a prediction by calculating the mean difference over the target patches and applying linear scaling. Details of each step in the model and on calculating the comparison metric are described below.

2.2.1. Filtering
The set of exponential filters we apply are taken from Basu and Su (2001). These exponential filters are two dimensional in shape and possess x-symmetry, y-symmetry, and symmetry with respect to the origin. They have unit volume and take the form:

$$g(x) = \frac{1}{K_1} exp^{-K_2|x|^m} \tag{1}$$

where K_1, K_2, and m are all positive constants. The m exponent corresponds to the shape of the filter. The normalization or
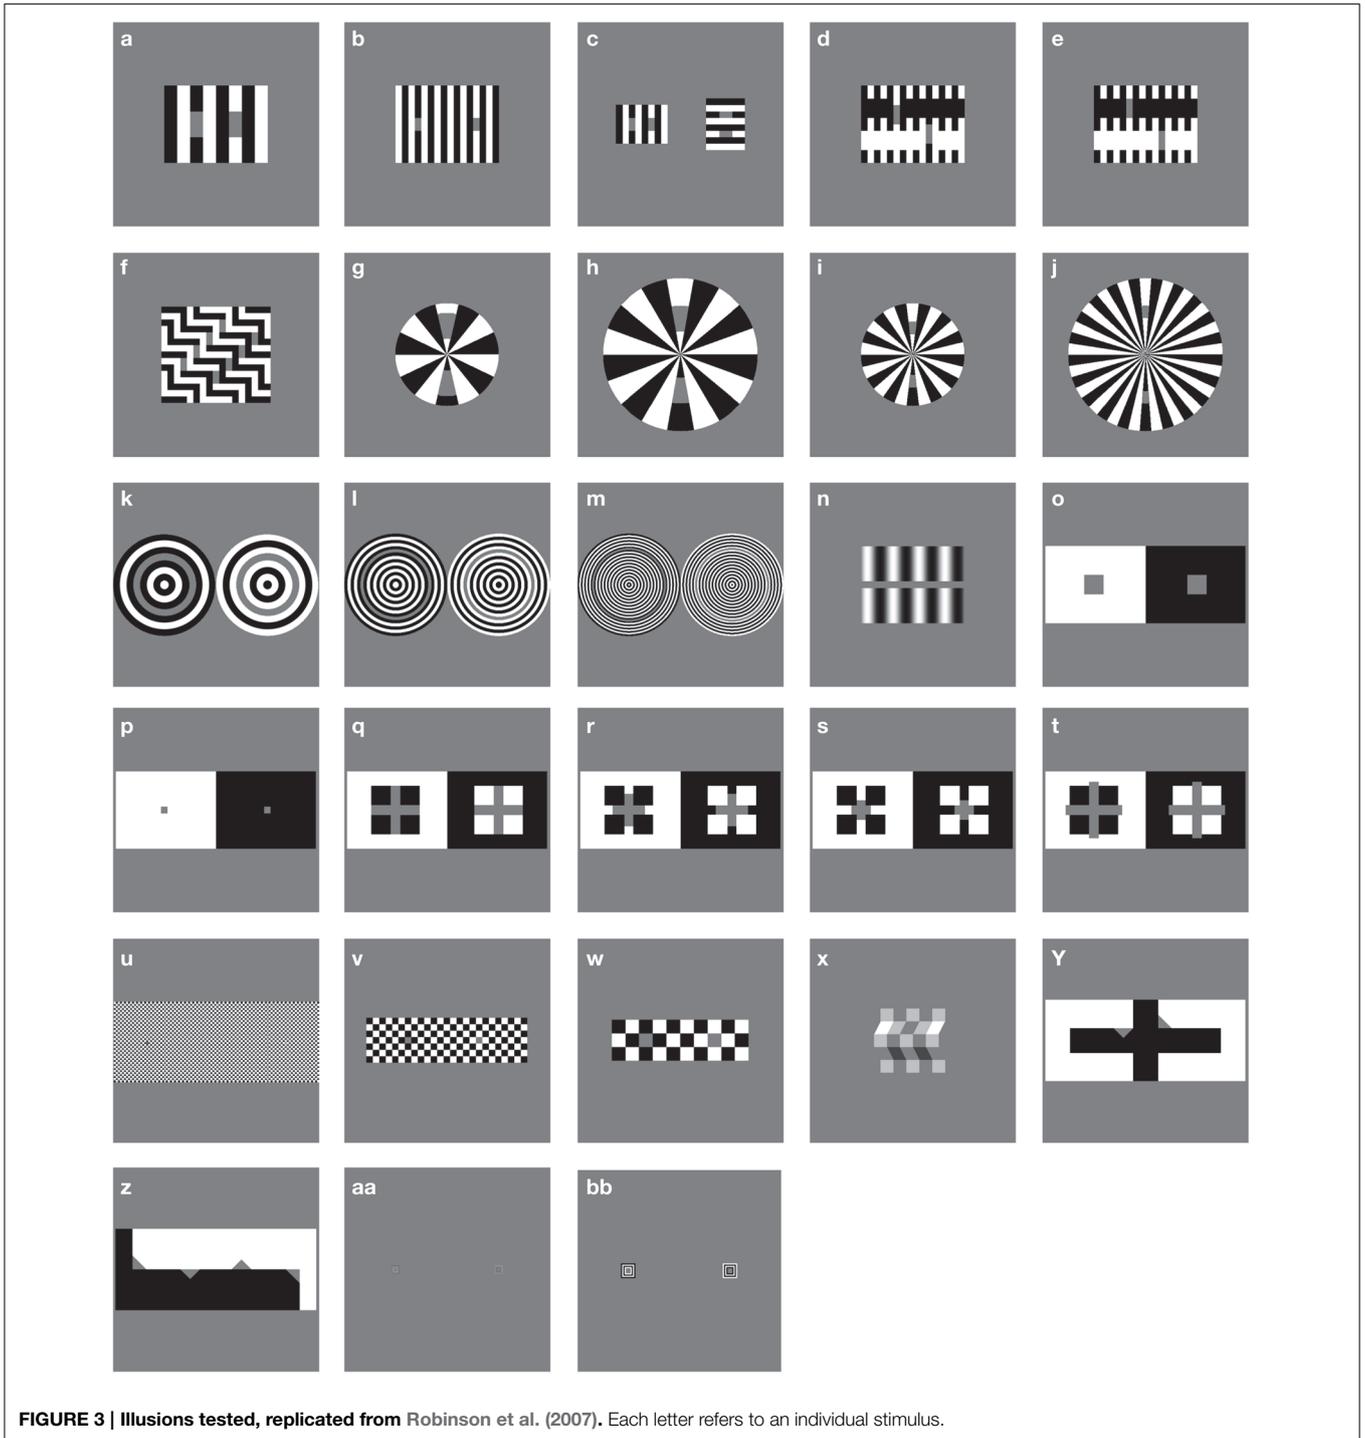
FIGURE 3 | Illusions tested, replicated from Robinson et al. (2007). Each letter refers to an individual stimulus.

scaling factor K_1 is calculated using K_2 and m as follows:

$$K_1 = (1/K_2^{1/m})(1/m)\Gamma(1/m) \qquad (2)$$

where constant K_2 is a function of the variance of $g(x)$, which denotes the size of the filter. $\Gamma(x)$ is the Gamma function defined as:

$$\Gamma(x) = \int_0^\infty t^{(x-1)} exp(-t)dt \qquad (3)$$

Figure 2 illustrates the variety of exponential filter shapes. When m is small, the exponential filter is described as having "high kurtosis," showing a sharper peak with more prominent ridges. When m is large, the exponential filter has "low kurtosis," being flatter and rounder with smoother ridges. A special case is formed

TABLE 1 | Stimuli with original sources, reproduced results (for strength comparison) and illusion direction reported by humans.
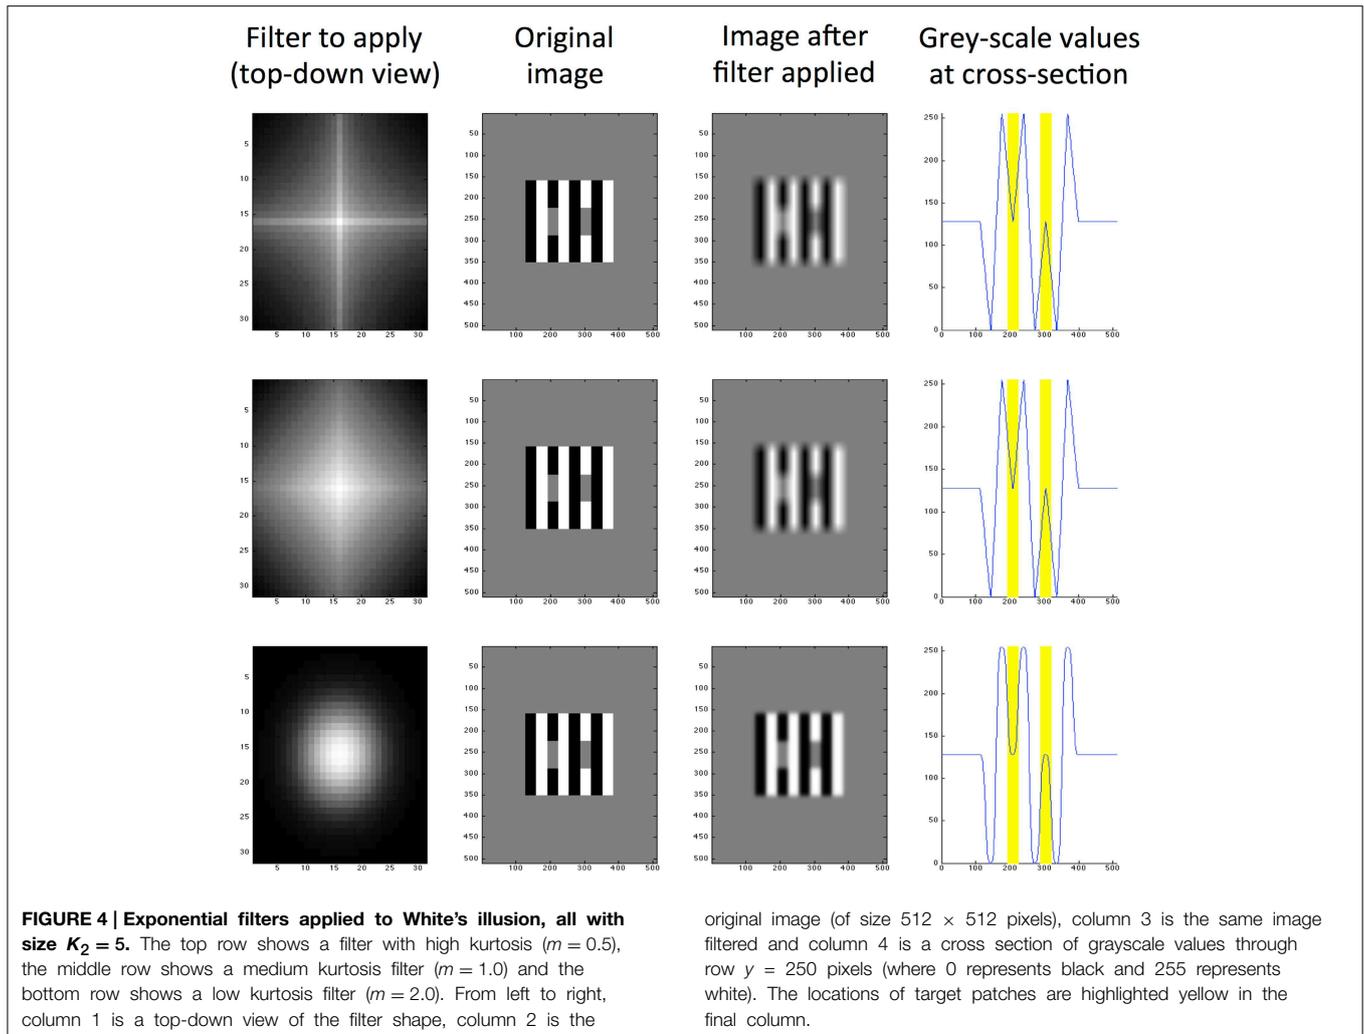
Figure	Original source	Reproduced results	Human Direction	Contrast (C) or Assimilation (A)?
a	White, 1979	Blakeslee and McCourt, 1999	Left	A
b	White, 1979	Blakeslee and McCourt, 1999	Left	A
c	Robinson et al., 2007		Top	A
d	Anderson, 2001	Blakeslee et al., 2005	Right	A
e	Howe, 2001	Blakeslee et al., 2005	No illusion	N/A
f	Clifford and Spehar, 2003		Left	A
g	Anstis, 2003		Bottom	A
h	Anstis, 2003		Bottom	A
i	Anstis, 2003		Bottom	A
j	Anstis, 2003		Bottom	A
k	Howe, 2005		Right	A
l	Howe, 2005		Right	A
m	Howe, 2005		Right	A
n	McCourt, 1982	Blakeslee and McCourt, 1999	Area between black	C
o	Chevreul, 1839	Blakeslee and McCourt, 1999	Right	C
p	Chevreul, 1839	Blakeslee and McCourt, 1999	Right	C
q	Pessoa et al., 1998	Blakeslee and McCourt, 1999	Left (Right in original)	C
r	Todorovic, 1997	Blakeslee and McCourt, 1999	Right	A
s	Todorovic, 1997	Blakeslee and McCourt, 1999	Right	N/A
t	Pessoa et al., 1998	Blakeslee and McCourt, 1999	Right	A
u	De Valois and De Valois, 1988	Blakeslee and McCourt, 2004	Right	A
v	De Valois and De Valois, 1988	Blakeslee and McCourt, 2004	Right	A
w	De Valois and De Valois, 1988	Blakeslee and McCourt, 2004	Left	C
x	Adelson, 1993	Blakeslee and McCourt, 2001	Bottom	C
y	Benary, 1924	Blakeslee and McCourt, 2001	Left	N/A
z_{2-1}	Todorovic, 1997	Blakeslee and McCourt, 2001	Second in 1–2	N/A
z_{4-3}	Todorovic, 1997	Blakeslee and McCourt, 2001	Fourth in 3–4	N/A
aa	Bindman and Chubb, 2004		Left	A
bb	Bindman and Chubb, 2004		Left	A

when $m = \frac{1}{2\sigma^2}$, where the function becomes a Gaussian with added rotational symmetry.

Each filter of a specific size and shape is applied to every pixel within the image. The size of the filter affects the information that is gleaned from an image. Smaller filters (high spatial frequencies or SFs) show better responsiveness but are less resilient to noise. Larger filters (low SFs) blur a lot of information, essentially losing information present in the images, but cope better with noise. There is a trade-off between selecting precise information and having greater resilience to noise, which is where scale selection comes in. The most appropriate filter selection finds the right compromise between these two factors, taking the smallest scale with the most reliable response.

A small amount of Gaussian noise is added to the image (0.1%) before filtering. Adding noise to the image is to avoid divide-by-zero errors when implementing divisive normalization. We are aware of other approaches to avoid divide-by-zero errors, such as adding a constant to the denominator term (Cope et al., 2013).

Responses are then convolved to create a filtered image of the same dimensions as the original input. The filtered convolved image is subtracted from the original image as the final step in processing. We explore a range of different filter shapes and sizes and produce a set of filtered images for every size and shape of filter. We use 10 filter sizes ranging from 5 pixels to 95 in increments of 10. The filter shapes range from 0.1 to 1.9 in increments of 0.2. **Figure 4** illustrates the result of applying three example filters with different shape parameters to White's Illusion. The predictive success of this particular filter size is well-demonstrated for this particular image, regardless of filter shape. The bottom row in **Figure 4** demonstrates a close approximation to the Gaussian filter, which in this case is able to predict the direction and magnitude of White's Effect. This filter differs from the DOG filters used by Blakeslee and McCourt (1997)'s model in two key ways. Firstly, Blakeslee and McCourt use a Difference-of-Gaussian (DOG) filter, rather than an approximate Gaussian pictured here. Secondly, **Figure 4** demonstrates a single filter

FIGURE 4 | Exponential filters applied to White's illusion, all with size $K_2 = 5$. The top row shows a filter with high kurtosis ($m = 0.5$), the middle row shows a medium kurtosis filter ($m = 1.0$) and the bottom row shows a low kurtosis filter ($m = 2.0$). From left to right, column 1 is a top-down view of the filter shape, column 2 is the original image (of size 512×512 pixels), column 3 is the same image filtered and column 4 is a cross section of grayscale values through row $y = 250$ pixels (where 0 represents black and 255 represents white). The locations of target patches are highlighted yellow in the final column.

2.2.2. Normalization (Optional)

After applying a specific contrast filter with shape m and size K_2 to each pixel location in the image, we optionally normalize the filter outputs. Normalization is not only useful in its primary function of constraining the dynamic response range of image filters, but is also beneficial for generating a faithful representation of image contrast. Following (Bonin et al., 2005), at each image location we divisively normalize the linear filter output by the output of a suppressive field, which computes the statistics of filter outputs surrounding the image location of interest. Bonin et al. (2005)'s normalization method, referred to as contrast gain control, is closely related to that found in the LGN and so we apply it here as a biologically plausible method for normalization in pre-cortical areas. In contrast to Bonin et al. (2005), who take the local root-mean-square contrast as the suppressive field, we divide filter responses by the local coefficient of variation. The local coefficient of variation is inversely related to local Weibull statistics and as such is diagnostic of local

image structure. Divisive normalization by the local coefficient of variation amplifies local image contrast. Similarly to Bonin et al. (2005), we compute normalized filter outputs using the following formula:

$$V = V_{max} \frac{g(x)}{c_{50} + c_{local}} \tag{4}$$

where c_{50} determines the strength of the suppressive field, V_{max} is the maximum response of the filter to the image, and $g(x)$ is the filtered response defined above. Finally, c_{local} is the local coefficient of variation:

$$c_{local} = \frac{\sigma}{\mu} \tag{5}$$

c_{local} is calculated based on the mean (μ) and the size of the suppressive field (σ) that is used as one of the parameters in our normalization step. The σ parameter specifies the size of the suppressive field compared to the size of the receptive field. When $\sigma = 1$, the size of the suppressive field is equal to that of the receptive field. When $\sigma = 2$, the size of the suppressive field is twice that of the receptive field.

operation, rather than a bank of filters used by Blakeslee and McCourt (1997).

2.2.3. Analysis Metrics

For each stimulus that we analyze, we take the resultant values (denoted as R) from the filter-only output (step 1) or from normalized output (step 2) with either $\sigma = 1$ or $\sigma = 2$ (as described above, σ represents the size of the suppressive field, as a proportion of the receptive field). We refer to $\sigma = 1$ as short-range normalization, where the suppressive field is the same area as the receptive field. $\sigma = 2$ is referred to as long-range normalization, where the suppressive field is twice the size of the receptive field. Within each image, we compare values over the two areas that have been assigned to be target patches (see Section 2.1). The lighter patch (as established in human experiments) is assigned to be patch A and the darker patch is assigned to be patch B. Mean values are obtained for both target patches before the mean of patch B is subtracted from the mean of patch A. Because patch A is assigned to be the lighter patch, a prediction in the correct direction is indicated by a positive value, whereas an incorrect prediction is negative. A value of zero indicates no difference in patch lightness values and therefore no illusion.

To compare resultant values, we scale the difference between target patches to the strength of White's Illusion for ease of comparison. The magnitude of White's illusion is denoted as R_a. This means that all resultant values are scaled to the strength (or magnitude) of White's illusion. A resultant value of 1 is then interpreted as having identical illusory strength to White's illusion. A value greater than 1 indicates the illusion is stronger than White's, and a value less than 1 (and above 0) indicates the illusion is weaker than White's. Although any stimulus could have been selected for comparative purposes, we follow Robinson et al. (2007)'s convention by selecting stimulus a as our comparative figure.

$$R = (\overline{A} - \overline{B})./|R_a| \qquad (6)$$

We also calculate the difference between model predictions and human results (where available) to quantify how well different model configurations match human data. We do this by subtracting the human result R_{human} from the model result R_{model} for stimuli from a to bb for which human results are available, and calculating the root mean square error (RMS_{error}). The smaller the RMS_{error} value, the better the model matches human data, and the greater the predictive accuracy of the model in terms of illusion magnitude or strength.
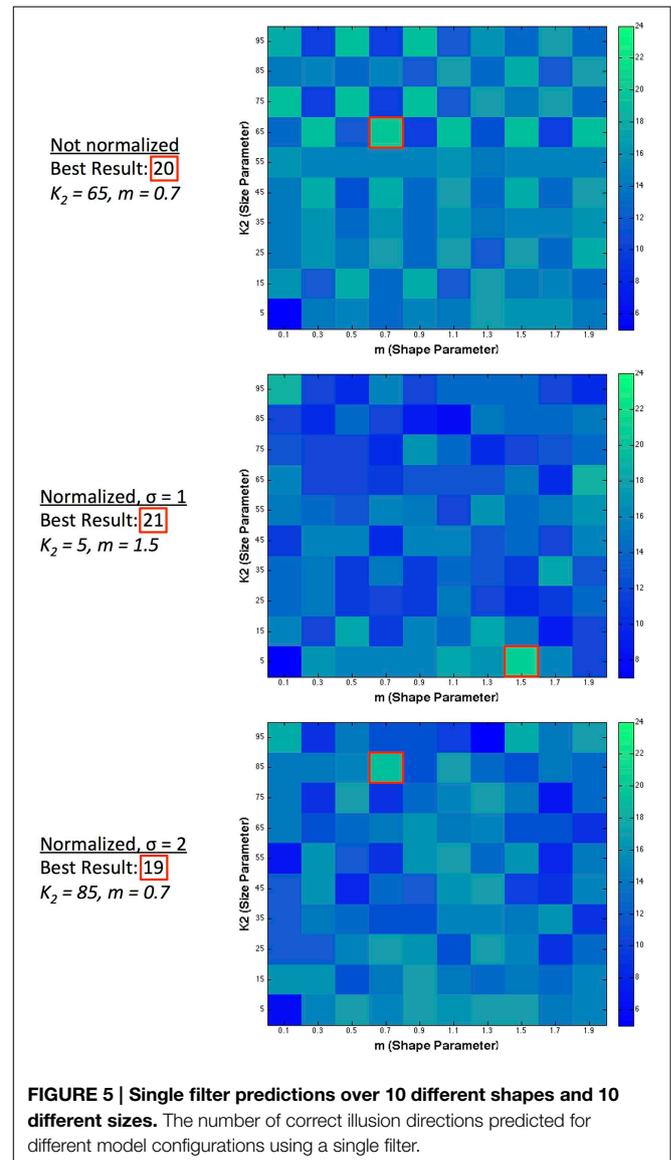
$$RMS_{error} = \sqrt{\frac{1}{n} \sum_{a}^{bb} \left(R_{model} - R_{human} \right)^2} \qquad (7)$$

When combining the outputs of two filters α and β of different sizes or shapes, we simply sum the difference in mean responses to the light and dark patches separately for each filter (removing scaling to figure a):

$$R^{\alpha\beta} = R^{\alpha} + R^{\beta} \qquad (8)$$

3. Results

We assess the performance of our model in two ways: the number of predictions in the correct direction, and also how closely the predicted values match the scaled human data on illusion magnitude. We exclude figure e from our analysis, given that no illusion direction is reported for humans. **Figure 5** illustrates the number of illusion directions correctly predicted (out of a maximum possible of 27) using a single filter over a range of 10 filter shapes and 10 filter sizes. For figure z, there are two predictions, annotated as z_{2-1} and z_{4-3}, for comparing the two left patches and the two right patches in the image, respectively. We take a correct result to be when $(z_{2-1} + z_{4-3})/2 > 0$. RMS_{error} is also calculated using the average over these two comparisons. We show predicted results for various model configurations: with no normalization, and with 2 ranges of local normalization ($\sigma = 1$ and $\sigma = 2$). With no normalization, the highest

Not normalized
Best Result: 20
$K_2 = 65$, $m = 0.7$

Normalized, σ = 1
Best Result: 21
$K_2 = 5$, $m = 1.5$

Normalized, σ = 2
Best Result: 19
$K_2 = 85$, $m = 0.7$

FIGURE 5 | Single filter predictions over 10 different shapes and 10 different sizes. The number of correct illusion directions predicted for different model configurations using a single filter.

number of correct direction predictions made by a single filter was 20 illusions using a large-sized filter with medium kurtosis. With short-range normalization ($\sigma = 1$), the highest number of correct direction predictions made by a single filter was 21 illusions (present in a small-sized filter with high kurtosis). With an increased normalization range ($\sigma = 2$), the best prediction result was slightly lower at 19 out of 28.

Table 2 lists the results for the best performing size and shape filter in terms of the difference of the mean values over the target patches. As mentioned above, we exclude figure e from our results because no illusion is reported in human results. We report values for z_{2-1} and z_{4-3} (in gray) and take the average of these two as our prediction for z, maintaining a single value prediction per illusion. In **Table 2** we also reproduce results from Robinson et al. (2007) for the ODOG, best LODOG, and best FLODOG model alongside human scaled results for direct comparison.

Predictions in the correct direction are shown in bold and tallies of the number of these correct predictions are presented at the bottom. For each model, we also list the RMS_{error} that represents how well the model's predictions match the magnitude of human results.

Table 2 shows that performance was maintained (in terms of number of correct direction predictions) when going from raw filter output to short-range normalized results for single filter predictions. Normalized results provided predictions with much smaller magnitudes of lightness illusions, as we would expect. Across predictions of both direction and magnitude, normalized results with $\sigma = 1$ provided the best predictions for single filters, showcasing the highest number of correct direction predictions (21) and reasonable magnitudes for these predictions (indicated by a substantially reduced RMS_{error} compared to filter-only output). Indeed, in this case RMS_{error} shows an accuracy of

TABLE 2 | Model results for the best single filter with and without normalization alongside ODOG and unscaled human results.

Figure	Shorthand Name	Human Scaled	ODOG	LODOG n = 2	FLODOG n = 2s m = 0.5	Exp model Single filter No norm	Exp model Single filter Norm σ = 1	Exp model Single filter Norm σ = 2
a	WE-thick	1	**1.00**	1.00	1.00	1.00	1.00	−1.00
b	WE-thin-wide	1.1	**2.08**	2.08	2.52	19.36	0.73	1.18
c	WE-dual		−0.30	1.36	1.93	−8.57	0.28	−0.49
d	WE-Anderson	1.54	−0.15	−0.30	−0.43	−1.68	−0.37	0.95
f	WE-zigzag		−0.51	−0.76	**1.26**	55.52	0.42	−1.69
g	WE-radial-thick-small		−0.67	−0.39	0.46	0.52	0.18	0.16
h	WE-radial-thick		−0.41	**0.01**	0.18	0.16	−0.16	1.09
i	WE-radial-thin-small		−0.34	0.21	2.74	2.32	0.28	−1.00
j	WE-radial-thin		−0.22	0.83	3.24	0.52	0.58	1.91
k	WE-circular1		−0.82	−1.04	0.28	1.24	0.22	0.52
l	WE-circular0.5		−0.53	−0.67	1.84	−2.84	0.67	2.40
m	WE-circular0.25		−0.38	−0.49	3.64	−2.15	0.55	−1.30
n	Grating induction	1.49	**2.03**	1.69	0.66	0.20	0.18	−0.30
o	SBC-large	2.72	**4.75**	7.56	3.96	4.01	3.09	0.75
p	SBC-small	4.73	**6.22**	14.94	5.96	8.79	4.52	7.02
q	Todorovic-equal	0.53	−0.36	−0.26	**0.08**	0.19	0.02	1.33
r	Todorovic-in-large	0.57	**0.49**	0.55	0.39	0.03	0.20	−1.00
s	Todorovic-in-small	1.05	**0.80**	0.95	1.08	0.32	0.19	0.80
t	Todorovic-out	0.37	**0.35**	0.38	0.03	0.34	−0.07	1.86
u	Checkerboard-0.16	1.78	**1.10**	0.94	8.03	−0.34	0.33	2.69
v	Checkerboard-0.94	0.68	**0.40**	0.35	−4.89	20.26	−0.19	3.93
w	Checkerboard-2.1	1.36	**0.69**	0.60	−1.48	0.61	0.05	0.77
x	Corrugated Mondrian	2.6	**0.95**	0.91	0.12	12.92	−0.02	2.32
y	Benary cross	2.2	**0.09**	0.06	0.05	−559.23	0.23	−1.94
z_{2-1}	Todorovic benary 1–2	2.86	−0.12	0.55	0.11	−1408.10	0.23	1.77
z_{4-3}	Todorovic benary 3–4	2.28	−0.12	0.58	0.14	1383.70	0.14	7.16
z avg	Todorovic benary average	2.57	−0.12	**0.57**	0.13	−12.20	0.18	4.47
aa	Bullseye-thin		−0.74	−0.35	**0.54**	0.18	0.02	3.31
bb	Bullseye-thick		−0.77	−0.38	**0.07**	1.16	−1.49	1.59
Total correct			13	18	24	20	**21**	19
RMS_{error}			1.29	1.80	2.56	140.59	**1.32**	1.85

The result reported for z is the average of z1 and z2 listed in gray. Bold values indicate predictions in the correct direction. Tallies of correct predictions are presented at the bottom for each model. The most successful single filter result using the exponential filter model has its tally highlighted in bold.

prediction that is matched only by the small values of the ODOG model, which fares considerably less well in terms of number of correct direction predictions (13). The RMS_{error} increased when the normalization range was extended to $\sigma = 2$, where only 19 correct direction predictions were made.

The results presented so far have demonstrated the capability of single filter predictions. We also combined multiple filters to observe the possibility of improving predictive success. **Figure 6** shows the result of combining pairs of filters together, taking a particular size and shape filter and combining it with the best possible match to maximize the number of correct directions predicted. The best result across all environments (normalized and filter-only), for dual filter combinations was 23 correct directions. The best resultant combinations in terms of maximizing the number of correct prediction directions occurred for a number of filter pairings within different environments. In the filter-only environment, the best filter pair combinations
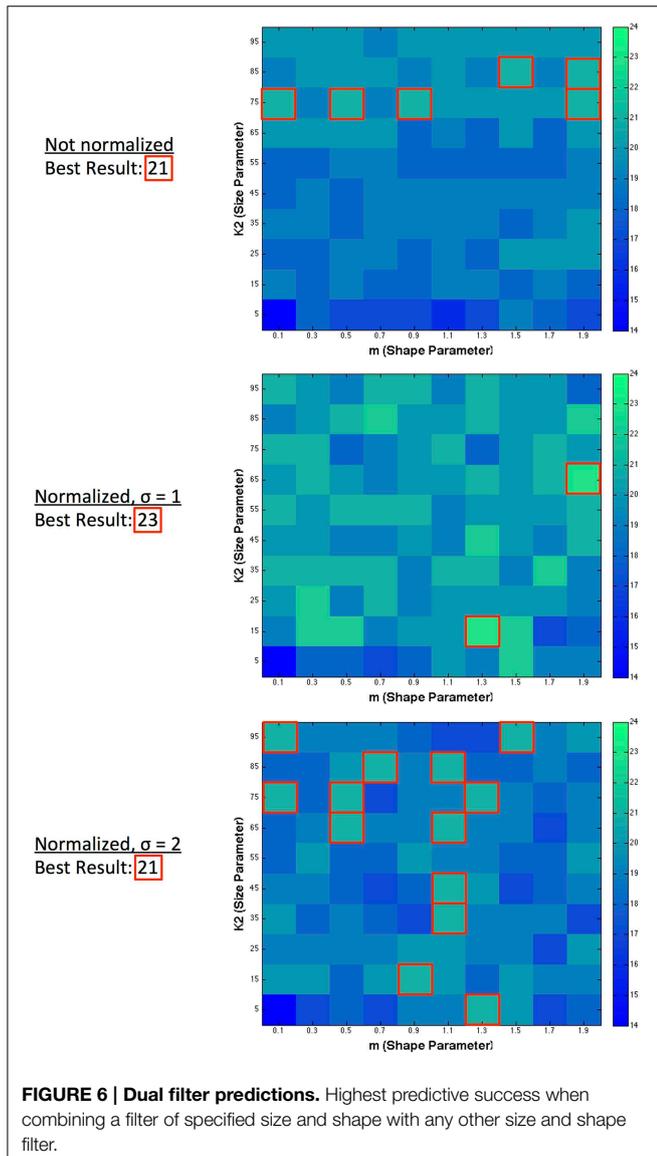
occurred across 6 different large sized filters ranging from high to low kurtosis. For normalized filters with $\sigma = 1$, the best filter pair was with a small sized filter with medium kurtosis and a medium sized filter with low kurtosis. For normalized filters with a larger range of normalization ($\sigma = 2$), the best pairings occurred across a range of filters with medium kurtosis over various sizes, or were large in size and had low to medium kurtosis.

We extended our multi-filter analysis to allow for the combination of any number of size and shape filters to determine whether an optimal combination of multiple filters exists. Using an ordered search sequence over the space of all possible shape and size filter combinations, we found that the maximum predictive success (in terms of illusion direction) that the model was able to achieve was 24 out of 27. This value represents the upper bound of performance of this exponential filter model and was found for the set of short-range normalized filters. **Figure 7** illustrates the four filter combinations that achieve the maximum of 24 correct illusion direction predictions for the exponential filter model. This was found for the set of normalized ($\sigma = 1$) filters. The filters across all four combinations were tallied and the frequency of these is presented on the right. A minimum of 6 filters was required to reach the best prediction as shown in combination 1. These were filters of size $K_2 = (15, 35, 85)$ and shape $m = (0.5, 1.3)$. Combinations 2–4 in **Figure 7** show the other filter combinations for which 24 illusion directions were correctly predicted. We see that a spread of different size and shape filter combinations is required to produce the best predictive performance. Certain filters are found to be informative whereas others are found to be consistently uninformative. Looking at the frequency of specific size and shape filters across all five most successful combinations, we see that filter ($K_2 = 15$, $m = 1.3$) is common across all filter arrangements. It is also evident that the organization of multiple filters is distributed across the parameter space.

4. Discussion

In this study, we applied a series of exponential filters differing in scale and shape to a set of lightness illusions that have previously been tested with Oriented Difference-of-Gaussian (ODOG) filters and associated models. The exponential model far outperforms the early ODOG models, and demonstrates predictive capabilities that match the successes of more recent elaborations of these models—LODOG and FLODOG—that incorporate local normalization post-filtering. Using a single filter, the direction of 21 (out of a possible 27) illusions can be predicted successfully. Using a two-filter combination, the predictive success of the model increases to 23. Extending the model to include any number of combined shape or size filters allows us to define the maximum capability of this model as 24 correct illusion direction predictions. Our results show that a low-level filtering model based on exponential filters can account for a large number of lightness illusions without requiring orientation-selective filters.

Comparing our work to the current literature, we highlight that existing models are restricted to filters of a specific shape (either DOG or LoG). We wanted to explore the

FIGURE 6 | Dual filter predictions. Highest predictive success when combining a filter of specified size and shape with any other size and shape filter.

FIGURE 7 | The four filter combinations that achieve the maximum of 24 correct illusion direction predictions for the exponential filter model. These combinations were found for short-range, normalized filters. The filters across all four combinations were tallied and the frequency of these is presented on the right.

effect of variation in the shape of the filters, which remains fixed in existing models. Our aim was not to emphasize stronger prediction performance, but to investigate whether filters inspired by image statistics can provide predictions on par with current state-of-the-art models. We have shown that this is indeed the case, where Gaussian-shaped filters do not provide the best predictability for the illusion set under all circumstances.

While the 28 stimuli used in this study feature substantial differences, one pertinent respect in which they vary is the induction of contrast or assimilation. Six of our illusions can be classified as predominantly contrast effects, whereas 18 primarily produce assimilation, with 4 illusions unclassifiable (see Section 2.1). Our best single-filter model was able to achieve 5/6 and 13/18 accuracy for contrast and assimilation effects, respectively, showing its ability to deal effectively with both classes of effect.

Among our catalog of illusions there are several sets of images that vary principally in terms of SF. These not only include low and high SF versions of White's Effect (*a* and *b*) and the SCI (*o* and *p*) as highlighted in **Figure 1**. Variations in SF are also seen for radial White's Effect (figures *g* through to *j*), circular configurations of White's Illusion (figures *k*, *l*, and *m*), the Checkerboard illusion (*u*, *v*, and *w*) and Bullseye figures (*aa* and *bb*). In **Table 2** (column 3), we list values of illusion magnitudes where human data is directly comparable with various SF configurations of the same illusion (reproduced from Robinson et al. 2007). Such comparisons are available for White's illusion (*a* and *b*), the SCI (*o* and *p*) and the Checkerboard illusion (*u*, *v*, and *w*). We draw direct conclusions for the performance of our best single-filter model to these figures. For the remaining figures with no directly comparable human data, we make observations based on the general rule that higher spatial frequencies yield greater effects. Our best single-filter model (normalized with $\sigma = 1$) predicts the correct direction of illusion for both high and low SF versions of White's illusions

(stimuli *a* and *b*) and of the SCI (figures *o* and *p*). In the case of the SCI the model can also account for the change in the size of the illusion as a function of SF, successfully predicting a larger effect at higher SF. However, in conflict with the human data, a reduction of the effect at higher SF is predicted for White's illusion. The Checkerboard illusion is an interesting case where the direction of the effect flips from assimilation to contrast for human observers when the visual angle of checkerboard squares is greater than approximately 1° of visual angle. Our best single-filter model is able to successfully account for two out of three illusion directions, with an appropriate increase in magnitude when comparing the lowest (*w*) and highest (*u*) SF versions. Despite an incorrect direction being predicted for figure *v*, the model correctly predicts a reduction in magnitude compared with *u*. Comparing the performance of our model to the best ODOG variants, we see that only ODOG and LODOG are able to account for all variations of correct illusory magnitudes where human data is available, performing with 5/5 correct relative magnitudes (for comparisons $b > a, p > o, u > v, w > v$, and $u > w$). The best performing model in terms of illusion direction, FLODOG, is able to successfully account for 3 out of a possible 5 illusory magnitudes consistent with SF. We conclude that our model is able to surpass that of FLODOG, with 4/5 illusion magnitudes that are commensurate with human data for both high and low spatial frequencies.
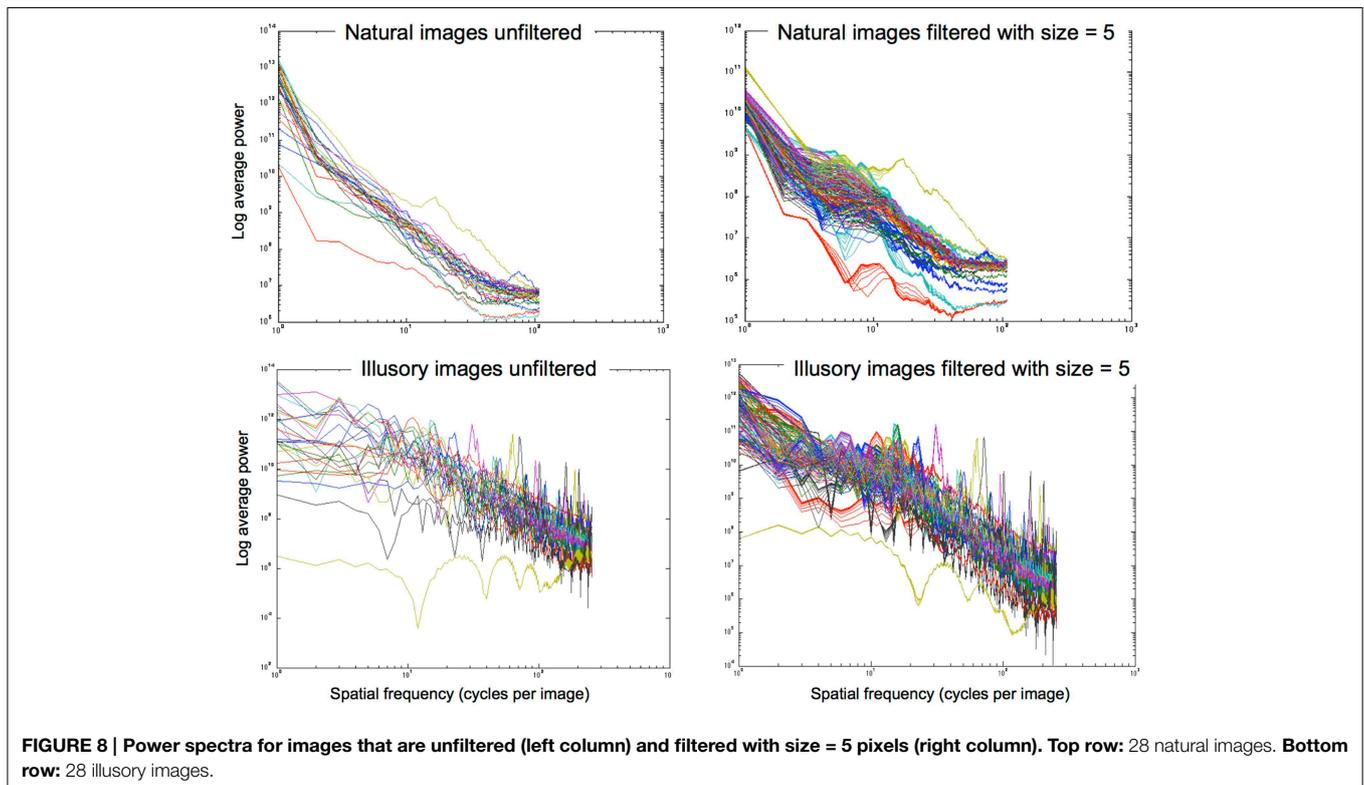
Reflecting on the best performance of the exponential model using a single filter, we note that two particular illusions that were predicted incorrectly—*t* (Pessoa et al., 1998); and *x* (Adelson, 1993)—warrant closer inspection. Stimulus *t* can be said to belong to the family of modified SCI figures from *q* to *t*. Figure *s* is a modified version of figure *o* (conventional SCI), where squares with opposite contrast polarity to the background are overlaid onto the target patch, creating equal boundaries of light and dark. Figures *r*, *q*, and *t* are modified versions of *s* with increasing

crossbar lengths. The spectrum of figure arrangements from *q* to *t* demonstrate changes to figure-ground relationships in terms of object assignment, depth placement and scene segmentation. In figures *q*, *r*, and *s*, the target patch appears to be contiguous with the surrounding white or black regions (as in the SCI: see stimuli *o* and *p*), and is positioned behind black or white square occluders. However, in stimulus *t*, the figure that posed a problem for our most successful single filter model, a quite different depth arrangement is evident, as the target patch now forms a cross that appears to be the most proximal object, and no longer shares the same depth plane as the surround. The exponential model we adopt does not include higher-level information such as depth cues of occlusion. Depth information is also evident in the corrugated Mondrian (figure *x*), providing shadow cues that could be processed by higher cortical levels for lightness judgments. These results may be taken to support suggestions that some illusions may escape successful prediction by low-level mechanisms if their lightness depends on depth relationships (Schirillo et al., 1990).

While the ODOG model and its variants closely approximate the orientation selective operations in V1, exponential filters based on image statistics represent an efficient coding scheme that could be present in pre-cortical areas as early as the retina. The prevailing view in early work with lightness illusions was that they arose from retinal interactions, rather than cortical processing (Cornsweet, 1970; Todorovic, 1997). However, more recent research highlights the influence of higher-level mechanisms on our lightness perception (Adelson, 2000; Gilchrist, 2006). Using our model, we do not prescribe that

filtering mechanisms alone can explain all lightness illusions. Instead, we set out to quantify the gap between what filtering operations can and cannot demonstrate. We propose that our exponential filtering model represents the first stage in a process of operations to estimate lightness. Later operations, such as those responsible for the scission of a scene into its component causal layers (Anderson, 1997) would occur post-filtering and normalization. The anchoring of lightness values to local and global context (Gilchrist, 2006) could occur within normalization operations or post-normalization. In our model's normalization step, the filtered image is first scaled to local responses (using local coefficient of variance) and then to the global maximum response within the image. This provides one of many approximations for the anchoring of lightness values.

The filtering approach we use reshapes contrast distributions toward those that best describe natural images using the exponential filter family. Similarly to Dakin and Bex (2003), we essentially reconstruct an image that represents the most probable naturally occurring source. By redistributing lightness values to more closely reflect the underlying statistical relationships of images within our environment, we can form predictions of perceptual lightness estimates that align with a large array of lightness illusions. **Figure 8** illustrates the power spectra for a set of images that are unfiltered (left column) and filtered (right column) using different shape filters that are all of size 5 pixels. The top row illustrates power spectra for 28 natural images. From these graphs we can see that the power spectra for filtered natural images is quite similar to the power spectra for unfiltered natural images. The bottom row shows the

FIGURE 8 | Power spectra for images that are unfiltered (left column) and filtered with size = 5 pixels (right column). Top row: 28 natural images. **Bottom row:** 28 illusory images.

power spectra for illusory images. The unfiltered images in the bottom left graph show a flatter power spectrum in the lower SFs than the filtered images in the bottom right graph. By applying these exponential filters, we see that they not only push the power spectra of illusory images toward that of natural images, reflecting the properties of image statistics. Applying these filters also boosts low SF information, hypothesized to be a driving factor in the perception of lightness illusions (Dakin and Bex, 2003). Dakin and Bex (2003) find that low spatial frequencies are primarily responsible for the Craik, Cornsweet, and O'Brien (CCOB) illusion that they study. The LoG filters that they apply boost this information when it is not present. From their results, Dakin and Bex (2003) hypothesize that low SF information may drive many illusions.

In a *post-hoc* analysis, we analyse whether filters of a particular shape aid in boosting low SF information, which is postulated by Dakin and Bex (2003) as a driving factor for many illusions. **Figure 9** illustrates the effect of different shape filters on the power distribution of a filtered White's Illusion image. Looking at the left side of the graph, we see that different shape filters have an effect on the low SF distributions. Filters with high kurtosis (those that have a low exponent and a sharper distribution) boost low SFs more than filters with low kurtosis (those that have a high exponent and a flatter distribution). The exponential filters therefore provide a mechanism to boost lower SF information more than Gaussian filters.

We emphasize that this study was conducted to investigate filters that are best able to push the power spectra of images toward that of natural images as well as preserve image structure while being resilient to noise. In earlier work, we showed that a filter size selection model helps in extracting and amplifying local image structure (Ghebreab et al., 2009). This model locally selects the smallest filter (extracting high-frequency information) with a response above a noise threshold (ensuring resilience to noise). In a similar fashion, local selection of filter shape may further enhance local image structure. Instead of performing local scale and shape selection in this paper, we study how different types of filters, varying in size and shape, may explain illusions.

The two-stage process of our model uses exponential filters that allow for efficient coding, followed by divisive normalization to boost shallow edges, promoting faithful representation of salient image features. In this way, the filtering stage of our model relies on the Efficient Coding Hypothesis, a theoretical model of sensory coding in the brain (Barlow, 1961). The Efficient Coding Hypothesis states that sensory information is represented in the most efficient way possible, such that it is closely representative of an organism's natural environment. The Efficient Coding Hypothesis is closely related to the Predictive Coding approach (Srinivasan et al., 1982), which states that the representation of sensory information in a statistically efficient way allows sensory systems to reduce redundancies and also provides greater resilience to noise (Barlow, 1961, 2001). In the specific case of our model, there is ample evidence from Basu and Su (2001) that exponential filters are resilient to many types and intensities of noise. From Dakin and Bex (2003) we see that statistical image representation and noise handling complement one another in understanding and predicting lightness illusions. Alongside
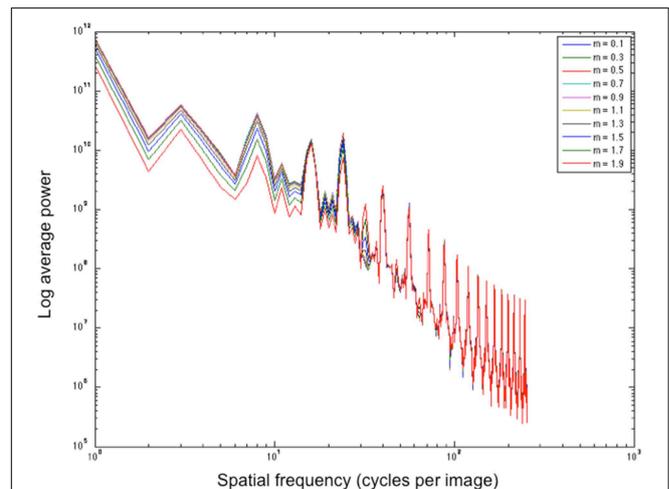

FIGURE 9 | Average power over spatial frequency of different shape filters applied to White's Illusion (figure a). All filters are of size 5 pixels. *m* refers to the exponent.

(Dakin and Bex, 2003), by successfully modeling illusions using properties of image statistics, we support the predictive coding approach proposed by Srinivasan et al. (1982).

In earlier work we showed that globally processing images with filters of different sizes results in scale space image representations that account for different visual phenomena (Ghebreab et al., 2009). We also showed that collapsing scale space representations into a single image representation via local scale selection accounts for even further visual phenomena. This model locally selects the smallest filter (extracting high-frequency information) with a response above a noise threshold (ensuring resilience to noise). In a similar fashion, local selection of filter shape may further enhance local image structure. Instead of performing local scale and shape selection, in this work we first studied if and how different types of filters, varying in size and shape, may explain illusions. We found this is indeed the case. We also tested whether combining different image representations, obtained by globally applying different filters, adds to explaining illusions. The next step in our work would be to determine whether local selection of filter size and shape, based on a model similar to Ghebreab et al. (2009), is able to further explain illusions.

An interesting future direction of study would be to explore additional versions of White's effect, particularly those that have been found to produce an inverted effect (Spehar et al., 1995; Ripamonti and Gerbino, 2001; Spehar et al., 2002). It is well-known that White's effect holds only when the luminance of the two target patches lies between the luminance values of the surrounding gratings (Spehar et al., 1995). Modifying the luminance values of the test patches to double-increments or double-decrements, relative to the gratings, not only drastically reduces the magnitude of illusion, but can also reverse the direction of the illusion from assimilation to contrast (Spehar et al., 1995; Ripamonti and Gerbino, 2001; Spehar et al., 2002). Inverted versions of White's effect have not been successfully

accounted for using (Blakeslee and McCourt, 1999)'s ODOG model, according to Spehar et al. (2002). Testing double-increment and double-decrement versions of White's effect in the exponential filter model may further demonstrate its robustness in accounting for an even larger range of lightness illusions.

Another direction for follow-up work would be to investigate the effects of different types and intensities of noise on human perception of lightness illusions and observe how closely these results are matched by our exponential filter model. Dakin and Bex (2003) show that when introducing different levels of noise into their stimuli, their model maintains a close approximation to human performance. However, ODOG has shown discrepancies in matching human response magnitudes for noisy stimuli (Betz et al., 2014). If the exponential filter model demonstrates results similar to human observers in classifying illusory images with noise manipulations, this would provide further support for predictive coding (Srinivasan et al., 1982).

In summary, our study demonstrates that a filter model based on contrast distribution statistics of natural images is able to account for the direction of 21 out of 27 lightness illusions using a single filter. When two filter combinations are considered, the number rises to 23, with asymptotic performance at 24 for an arbitrarily large combination of filter outputs. We observe the effect of incorporating non-linear divisive normalization, providing a better understanding of the role that contrast gain control provides in the perception of these illusions. While short-range normalization only slightly improves the number of correct direction predictions, it considerably reduces the error in predicting illusion magnitude, measured as RMS_{error}. The exponential filters we employ are not orientation selective, demonstrating that V1-style operations are not required to account for a large number of lightness illusions. Given that these exponential filters could be found as early as the retina, it is possible that the majority of these lightness effects result from pre-cortical operations, leaving only a few to be explained by higher level mechanisms.

Funding

References

Adelson, E. H. (1993). Perceptual organization and the judgment of brightness. *Science* 262, 2042–2044.

Adelson, E. H. (2000). *The New Cognitive Neurosciences, Chapter Lightness Perception and Lightness Illusions, 2nd Edn.* Cambridge, MA: MIT Press.

Anderson, B. L. (1997). A theory of illusory lightness and transparency in monocular and binocular images: the role of contour junctions. *Perception* 26, 419–454.

Anderson, B. L. (2001). Contrasting theories of White's illusion. *Perception* 30, 1499–1501. doi: 10.1068/p3012ds

Anstis, S. (2003). *White's Effect Radial.* Available online at: http://www.cogsci.ucsd.edu/stanonik/illusions/wer0.html

Barlow, H. (1961). "Possible principles underlying the transformation of sensory messages," in *Sensory Communication,* ed W. A. Rosenblith (Cambridge, UK: MIT Press), 217–234. doi: 10.7551/mitpress/9780262518420.003.0013

Barlow, H. (2001). Redundancy reduction revisited. *Network* 12, 241–253. doi: 10.1088/0954-898X/12/3/301

Basu, M., and Su, M. (2001). Image smoothing with exponential functions. *Int. J. Patt. Recogn. Artif. Intell.* 14, 735–752. doi: 10.1142/S0218001401001076

Benary, W. (1924). Beobachtungen zu einem experiment über Helligkeitskontrast. *Psychol. Forsch.* 5, 131–142.

Betz, T., Wichmann, F. A., Shapley, R., and Maertens, M. (2014). Testing the ODOG brightness model with narrowband noise stimuli. *Percept. ECVP Abstr. Suppl.* 43, 164.

Bindman, D., and Chubb, C. (2004). Brightness assimilation in bullseye displays. *Vis. Res.* 44, 309–319. doi: 10.1016/S0042-6989(03)00430-9

Blakeslee, B., and McCourt, M. E. (1997). Similar mechanisms underlie simultaneous brightness contrast and grating induction. *Vis. Res.* 37, 2849–2869.

Blakeslee, B., and McCourt, M. E. (1999). A multiscale spatial filtering account of the white effect, simultaneous brightness contrast and grating induction. *Vis. Res.* 39, 4361–4377.

Blakeslee, B., and McCourt, M. E. (2001). A multiscale spatial filtering account of the Wertheimer-Benary effect and the corrugated Mondrian. *Vis. Res.* 41, 2487–2502. doi: 10.1016/S0042-6989(01)00138-9

Blakeslee, B., and McCourt, M. E. (2004). A unified theory of brightness contrast and assimilation incorporating oriented multi-scale spatial filtering and contrast normalization. *Vis. Res.* 44, 2483–2503. doi: 10.1016/j.visres.2004.05.015

Blakeslee, B., Pasieka, W., and McCourt, M. E. (2005). Oriented multiscale spatial filtering and contrast normalization: a parsimonious model of brightness induction in a continuum of stimuli including white, howe and simultaneous brightness contrast. *Vis. Res.* 45, 607–615. doi: 10.1016/j.visres.2004.09.027

Bonin, V., Mante, V., and Carandini, M. (2005). The suppressive field of neurons in lateral geniculate nucleus. *J. Neurosci.* 25, 10844–10856. doi: 10.1523/JNEUROSCI.3562-05.2005

Chevreul, M. E. (1839). *De la loi du Contraste Simultané des Couleurs et de l'Assortiment des Objets Colorés. Translated into English by C. Martel as The Principles of Harmony and Contrast of Colours, English 2nd Edn, Longman, Brown, Green and Longmans (1855).* Paris: Pitois-Levrault.

Clifford, C. W. G., and Spehar, B. (2003). Using colour to disambiguate contrast and assimilation in white's effect. *J. Vis.* 3:294. doi: 10.1167/3.9.294

Cope, D., Blakeslee, B., and McCourt, M. E. (2013). Modeling lateral geniculate nucleus response with contrast gain control. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 30, 2401–2408. doi: 10.1364/JOSAA.31.000348

Corney, D., and Lotto, R. B. (2007). What are lightness illusions and why do we see them? *PLoS Comput. Biol.* 3:e180. doi: 10.1371/journal.pcbi.0030180

Cornsweet, T. N. (1970). *Visual Perception.* New York, NY: Academic Press.

Craik, K. J. W. (1966). *The Nature of Psychology: A Selection of Papers, Essays and Other Writings by the Late K. J. W. Craik.* Cambridge, UK: Cambridge University Press.

Dakin, S. C., and Bex, P. J. (2003). Natural image statistics mediate brightness 'filling in'. *Proc. Biol. Sci.* 270, 2341–2348. doi: 10.1098/rspb.2003.2528

De Valois, R. L., and De Valois, K. K. (1988). *Spatial Vision*. New York, NY: Oxford University Press.

Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 4, 2379–2394.

Geisler, W. S. (2008). Visual perception and the statistical properties of natural scenes. *Annu. Rev. Psychol.* 59, 167–192. doi: 10.1146/annurev.psych.58.110405.085632

Ghebreab, S., Scholte, H. S., Lamme, V. A. F., and Smeulders, A. W. M. (2009). "A biologically plausible model for rapid natural scene identification," in *Proceedings of Neural Information Processing Systems 2009* (Vancouver, BC), 629–637.

Ghosh, K., Sarkar, S., and Bhaumik, K. (2006). A possible explanation of the low-level brightness-contrast illusions in the light of an extended classical receptive field model of retinal ganglion cells. *Biol. Cybern.* 94, 89–96. doi: 10.1007/s00422-005-0038-4

Gilchrist, A. (2006). *Seeing Black and White*. New York, NY: Oxford University Press.

Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., et al. (1999). An anchoring theory of lightness perception. *Psychol. Rev.* 106, 795–834.

Howe, P. D. L. (2001). A comment on the Anderson (1997), the Todorovic (1997) and the Ross and Pessoa (2000) explanations of white's effect. *Perception* 30, 1023–1026. doi: 10.1068/p3212

Howe, P. D. L. (2005). White's effect: removing the junctions but preserving the strength of the illusion. *Perception* 34, 557–564. doi: 10.1068/p5414

Kingdom, F. A. (2011). Lightness, brightness and transparency: a quarter entury of new ideas, captivating demonstrations and unrelenting controversy. *Vis. Res.* 51, 652–673. doi: 10.1016/j.visres.2010.09.012

Knill, D. C., and Kersten, D. (1991). Apparent surface curvature affects lightness perception. *Nature* 351, 228–230. doi: 10.1038/351228a0

McCourt, M. E. (1982). A spatial frequency dependent grating-induction effect. *Vis. Res.* 22, 119–123, 125–134.

Nykamp, D. Q., and Ringach, D. L. (2002). Full identification of a linear-nonlinear system via cross-correlation analysis. *J. Vis.* 2, 1–11. doi: 10.1167/2.1.1

O'Brien, V. (1958). Contour perception, illusion and reality. *J. Opt. Soc. Am* 48, 112–119.

Packer, O. S., and Dacey, D. M. (2002). Receptive field structure of H1 horizontal cells in macaque monkey retina. *J. Vis.* 2, 272–292. doi: 10.1167/2.4.1

Packer, O. S., and Dacey, D. M. (2005). Synergistic center-surround receptive field model of monkey H1 horizontal cells. *J. Vis.* 5, 1038–1054. doi: 10.1167/5.11.9

Pessoa, L., Baratoff, G., Neumann, H., and Todorovic, D. (1998). Lightness and junctions: variations on white's display. *Investig. Ophthalmol. Vis. Sci. (Suppl.)* 39, S159.

Ripamonti, C., and Gerbino, W. (2001). Classical and inverted white's effect. *Perception* 30, 467–488. doi: 10.1068/p3108

Robinson, A. E., Hammon, P. S., and de Sa, V. R. (2007). Explaining brightness illusions using spatial filtering and local response normalization. *Vis. Res.* 47, 1631–1644. doi: 10.1016/j.visres.2007.02.017

Ruderman, D. L., and Bialek, W. (1994). Statistics of natural images: scaling in the woods. *Phys. Rev. Lett.* 73, 814–817.

Schirillo, J., Reeves, A., and Arend, L. (1990). Perceived lightness, but not brightness, of achromatic surfaces depends on perceived depth information. *Percept. Psychophys.* 48, 82–90.

Schwartz, O., and Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. *Nat. Neurosci.* 4, 819–825. doi: 10.1038/90526

Shapiro, A., and Lu, Z.-L. (2011). Relative brightness in natural images can be accounted for by removing blurry content. *Psychol. Sci.* 22, 1452–1459. doi: 10.1177/0956797611417453

Spehar, B., Clifford, C. W. G., and Agostini, T. (2002). Induction in variants of white's effect: common or separate mechanisms? *Perception* 31, 189–196. doi: 10.1068/p10sp

Spehar, B., Gilchrist, A., and Arend, L. (1995). White's illusion and brightness induction: the critical role of luminance relations. *Vis. Res.* 35, 2603–2614.

Srinivasan, M. V., Laughlin, S. B., and Dubs, A. (1982). Predictive coding: a fresh view of inhibition in the retina. *Proc. R. Soc. Lond. B Biol. Sci.* 216, 427–459.

Todorovic, D. (1997). Lightness and junctions. *Perception* 26, 379–394.

White, M. (1979). A new effect of pattern on perceived lightness. *Perception* 8, 413–416.

Zhu, S. C., and Mumford, D. (1997a). Prior learning and Gibbs reaction-diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.* 19, 1236–1250. doi: 10.1109/34.632983

Zhu, S. C., and Mumford, D. B. (1997b). "Learning generic prior models for visual computation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Los Alamitos, CA: IEEE Computer Society), 463–469.