

A psychophysical model to control the brightness and key-to-fill ratio in CG cartoon character lighting

Pisut Wisessing
Trinity College Dublin
wisessip@tcd.ie

Douglas W. Cunningham
BTU Cottbus-Senftenberg
douglas.cunningham@b-tu.de

Katja Zibrek
Trinity College Dublin
kzibrek@tcd.ie

Rachel McDonnell
Trinity College Dublin
ramcdonn@tcd.ie

ABSTRACT

Lighting is a commonly used tool to manipulate the appearance of virtual characters in a range of applications. However, there are few studies which systematically examine the effect of lighting changes on complex dynamic stimuli. Our study presents several perceptual experiments, designed to investigate the ability of participants to discriminate lighting levels and the ratio of light intensity projected on the two sides of a cartoon character's face (key-to-fill ratio) in portrait lighting design. We used a standard psychophysical method for measuring discrimination, typical in low-level perceptual studies but not frequently considered for evaluating complex stimuli. We found that people can easily differentiate lighting intensities, and distinguish between shadow strength and scene brightness under bright conditions but not under dark conditions. We provide a model of the results, and empirically validate the predictions of the model. We discuss the practical implications of our results and how they can be exploited to make the process of portrait lighting for CG cartoon characters more consistent, such as a tool for manipulating shadow while maintaining the level of perceived brightness.

CCS CONCEPTS

• **Computing methodologies** → **Perception**; Animation; Visibility.

KEYWORDS

virtual character, psychophysics, rendering, animation, visibility

ACM Reference Format:

Pisut Wisessing, Katja Zibrek, Douglas W. Cunningham, and Rachel McDonnell. 2019. A psychophysical model to control the brightness and key-to-fill ratio in CG cartoon character lighting. In *ACM Symposium on Applied Perception 2019 (SAP '19)*, September 19–20, 2019, Barcelona, Spain. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3343036.3343127>

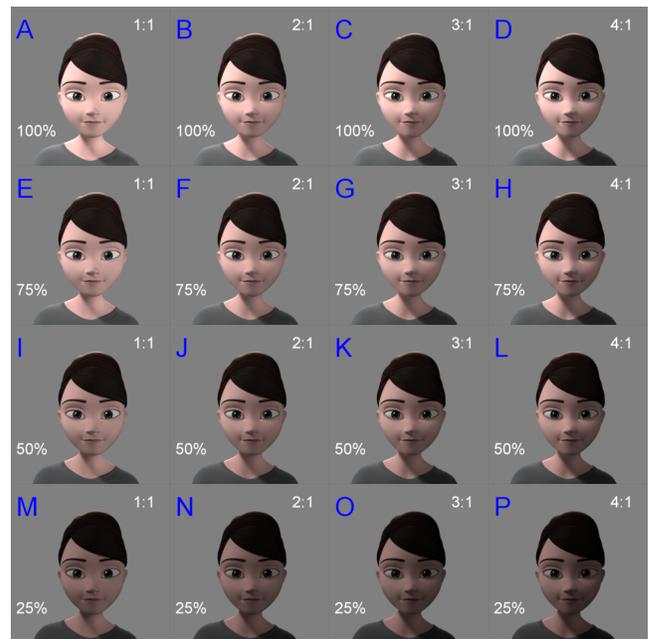


Figure 1: Still images taken from the 16 movies, rendered in different key light intensity and key-to-fill ratios.

1 INTRODUCTION

Lighting is an essential component of every 3D computer generated scene containing virtual characters. In addition to changing overall brightness, a strong perceptual impact can be created by changing the strength of shadows, which is achieved through changes in key-to-fill-ratio (KTFR). KTFR is the relationship between the illumination due to the main lighting source (i.e., key light) and the light that illuminates the shadows, i.e., fill-light. This is typically

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SAP '19, September 19–20, 2019, Barcelona, Spain

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6890-2/19/09...\$15.00

<https://doi.org/10.1145/3343036.3343127>

achieved using a three-point-lighting setup (Appendix, Figure 3). However, little is known on how we perceive the overall brightness of a character under different shadow conditions.

There has been a long history of investigation in perceptual psychology, psychophysics, and vision science, focusing on the relationship between the physical changes in illumination intensity levels on the one side and the perception of either surface reflectance (lightness) or surface illumination (brightness) on the other side (for recent reviews please see [Gilchrist 2013; Kingdom 2011]). This led to the discovery of essential concepts, such as Weber-Fechner's and Steven's law [Fechner 1966; Stevens 1957] which showed that the perception of illumination intensity is nonlinearly related to the actual physical illumination. In addition, the work led to the discovery of lightness constancy and color constancy, which assert that the perceived appearance of objects remains relatively constant even under large variations in the lighting conditions [Adelson 2000; Foster 2011].

On the other hand, there is mounting evidence that people do not fully discount illumination when perceiving surface reflectance. Logvinenko and Maloney [2006], for example, investigated the relationship between illumination and perceived surface lightness. They presented participants with two arrays of Munsell chips (while the two arrays were identical, each array had chips of different albedos) under different lighting conditions. The participants were asked to rate the similarity of pairs of chips (one from each array). The resulting distance matrix was submitted to a multidimensional scaling (MDS) analysis to recover the perceptual space underlying surface lightness. They found that while albedo (surface reflectance) was the primary determinant of surface lightness, the intensity of the illumination (i.e., the shading) also had an effect. That is, people do not fully discount the effect of shading in lightness perception.

Our study uses the same approach to model the perception of overall light intensity (i.e., scene brightness) and shadow depth. To do this, we use a more sophisticated stimulus—a CG character—and a more realistic, three-point lighting setup, which is a practical illumination scheme used in film and photography. Our stimulus is also animated, since the goal is to study lighting effects in the context of computer animation. We also had the observers rate the similarity of the entire scene rather than small patches within the scene. By collecting similarity values for a complex scene under many different lighting conditions, we can use MDS to automatically recover the critical perceptual dimensions involved in perceiving the illumination in a scene.

Our main contributions are: we conducted an exploratory psychophysics experiment to study the proximity structure (relative locations of the animations in perceptual space) of perceived scene lighting intensity and shadow depth for animated cartoon characters. We also derived a log-polar parametric model connecting lighting changes and the perceptual dimensions of scene brightness and shadow depth. We then validated this model with a perceptual experiment. We believe our method is a first step towards developing an improved lighting tool that allows shadows on a cartoon character's face to be adjusted without altering the overall perceived brightness of the character.

2 RELATED WORK

The human visual system has developed an incredible – and very complex – ability to account for contextual information, illumination information, object geometry, material properties, and scene layout in order to isolate the albedo of a surface [Gilchrist 2013]. One key finding demonstrating this ability to isolate the separate sources of image intensity changes, as mentioned above, is lightness constancy. The explanation of this phenomenon has been a major challenge in visual science [Brainard 2003; Gilchrist et al. 1999]. There have been many empirical attempts to map people's ability to discriminate brightness and lightness levels (see, e.g., [Fechner 1966; Gilchrist 2013; Kingdom 2011; Stevens 1957]). In general, once the contributions of the different sources of image space intensity changes are determined, they are kept separate. In addition to lightness constancy, descriptions of a scene provide evidence for this: people tend to not mention illumination effects such as shadows and shading when describing a scene [Kardos 1934].

Although light intensity is largely discounted in the perception of a surface's albedo, it is still processed and used for additional perceptual tasks. For example, changes in brightness across the surface of an object that are due to illumination and not texture, i.e., shadows, are an important factor in acquiring shape information [Mingolla 1983] and have long been used in computer graphics and vision to retrieve or present the 3D shape of an object [Bruckstein 1988]. Additionally, face perception has been shown to be affected by shadows [Hill and Bruce 1996; Johnston et al. 1992].

Stylized shadows have been used to enhance the perception of shape, detail, and depth of 3D objects [Rusinkiewicz et al. 2006; Šoltészová et al. 2011]. Shadows also create a difference in the spectral reflection of the surface of an object, but the perceived color tends to remain that of the illuminated part of the object, such as in the “checkershadow illusion” described by Adelson [2011]. However, the effect of the shadow will be perceived differently according to lighting conditions, as studies analyzing the perceived magnitude of luminance differences (contrast) under different lighting conditions show [Pamir and Boyaci 2016; Peli et al. 1991; Van Nes and Bouman 1967]. While contrast is known to be constant across middle range luminance levels, this constancy breaks down in darker lighting conditions, with contrast perception decreasing under 8 cd/m^2 , according to Peli et al. [1991]. Even in the middle range, people do not fully discount the effect of lighting [Logvinenko and Maloney 2006]. In summary, lighting perception involves complex interactions between changes in physical lighting, scene layout, object geometry, the observer's position, and the specifics of the human visual system (to mention just a few factors) in order to preserve consistency of object properties.

Other work in the field of computer graphics has used psychophysics to propose novel perceptual models for computer generated objects. Examples include material reflectance [Vangorp et al. 2007], surface gloss and transparency [Cunningham et al. 2007], local adaptation [Vangorp et al. 2015], and more recently material appearance [Lagunas et al. 2019].

The majority of this work has focused on still images. This is surprising, since motion is an essential aspect of natural visual scenes. One study on color constancy [Werner 2007] showed that

the synergistic integration of color and motion signals is an important mechanism for improving color identification. Therefore, with added motion, color constancy improves. More recently, perception of material properties have been studied on dynamic stimuli such as cloth [Bi and Xiao 2016], liquids [van Assen and Fleming 2016], and optical flow characteristics [Doerschner et al. 2011].

3 STIMULI

We chose a female CG character, Mery¹, for the experiment because of her cartoon-like appearance typically seen in animated movies (see Figure 1) and detailed, controllable facial rigs enabling an animator to create high-quality animation. Mery was lit in a well-known, three-point lighting setup [Millerson 1991] consisting of a key light (the primary source illuminating the character), a fill light brightening up the shadow side of the character and a back light separating the character from the background. The direction of each light is shown in Appendix, Figure 3 (Top).

We followed the conventional 3D animation production pipeline and rendered the contribution from the key and fill light separately and later combined them for the desired **Key-to-Fill Ratio (KTFR)**, the proportion of light intensity projected on the key side and fill side of the character face. Note that it is valid to render images under separate light sources and then add them together since radiance can be summed together [Nicodemus et al. 1992]. KTFR is normally used as a measure of contrast in cinematography and photography. We also adopted the conventional *expose to the right* technique, illuminating the character with the maximum possible intensity before overexposing, to create the 100% key intensity and 1:1 KTFR lighting condition (Appendix, Figure 3 (Bottom)). After that, we reduced the key and fill contributions in the gamma-corrected image space to create 4 levels of key intensities: 100%, 75%, 50% and 25%, and 4 levels of KTFR: 1:1, 2:1, 3:1 and 4:1.

In studio portrait lighting, the background is typically in solid color lit independently from the subject (no shadow interaction). To help reduce the influence of the background in our experiment, the character was rendered in front of a 18% gray background, believed to be the middle gray according to the *Zone System* [Adams 1948] and perceived to be the midway between black and white in CIELAB color space [Stone 2003]. Note that the use of a gray for the background is relatively common in psychophysics, but so is black or even speckled. Note that the surface with the highest intensity (and thus the white anchor [Gilchrist et al. 1999]) was the eyes of the character. Moreover, one eye was in the shadow and one was not, allowing a direct comparison.

All lights used in our renders were white so the white point temperature takes on that of the monitor, Dell UP2713H monitor, which was calibrated to 100% sRGB color gamut, 6500k white point, and 80 cd/m^2 brightness. The experiment was conducted in a completely dark room to reduce the interference of outside light.

4 EXPERIMENTAL DESIGN

We rendered 16 animations (Figure 1) of the Mery character talking naturally for 2 seconds with a neutral expression (key light intensity: 100%, 75%, 50%, 25% x 4 KTFR: 1:1, 2:1, 3:1, 4:1).

The participants (7 females, 8 males, aged 23-50, avg. age 35) were asked to rate the dissimilarity of the two movie clips (which had no audio) on a scale of 0 to 16. Note that while rating scales with more than 9 points probably do not add resolution, they also generally do not decrease accuracy either [Cunningham and Wallraven 2011]. Participants started with a training session during which they were shown example pairs of identical stimuli (A-A, P-P) and informed that this was a dissimilarity of 0, as well as example pairs of extreme difference (A-P, P-A) and told that this dissimilarity was a 16. Next, all 256 possible pairs were displayed in a random order. Still image thumbnails of an A-P pair were displayed beside the 16 end of the rating scale at all times. The movies were looped continuously until the participant made a decision. After each rating, the participant was shown a blank screen with the middle-gray background and a white fixation point in the middle screen. Participants pressed the mouse button to start the next trial. The size of each movie in the pair was 10.3 x 10.3 degrees of visual angle. The pair was displayed 3.43 degrees of visual angle apart. The experiment lasted approximately 15-20 minutes.

4.1 Multidimensional Scaling (MDS)

Most perceptual methods for determining the functional relationship between stimulus and perception require the experimenter to know the perceptual dimension or dimensions that are being studied. Richardson [1938] (see also [Torgerson 1952]) introduced a new class of method, referred to as MDS, to directly determine from a set of simple pairwise similarity ratings the number of perceptual dimensions involved, as well as their scale values. MDS takes an $N \times N$ dissimilarity matrix of N stimuli and returns the coordinates of the N stimuli in a d -dimensional space that best fits the dissimilarity matrix (for more, see [Cox and Cox 2001; Cunningham and Wallraven 2011]). MDS assumes Euclidean distances. Since we use Likert scales, it is not clear that our similarity structure is, in fact, Euclidean. Thus, we used a non-metric MDS which uses the ordinal relationships to obtain scaled proximities.

All analyses were done in Statistica 7.1 and Matlab 2019a software packages and the computational details can be found in Borg [2012].

5 RESULTS

The ratings were averaged across participants to create the 16 x 16 dissimilarity matrix. Scree plots, Kaiser criteria, parallel analysis, and the optimal coordinates all suggest that two dimensions are sufficient to explain the dissimilarity in the matrix. The resulting MDS-derived perceptual similarity space (which has a stress value of 0.036, indicating a good fit) is presented in Figure 2 (Top), and Table 1 in the Appendix.

As can be seen by analyzing the positions of the blue data-points in the figure, the four key intensities for each KTFR (e.g., A, E, I and M) vary strongly and consistently along *Dimension 1*. Analyses show that the average pixel value (in any single frame of the animation) correlates nearly perfectly with *Dimension 1* ($r^2 = 0.95$). It seems that *Dimension 1* represents scene brightness.

The effect of varying KTFR, known to control the contrast caused by shadows, for a given key intensity seems to be roughly orthogonal to the effect of key intensity. This, combined with a close

¹<http://www.meryproject.com>

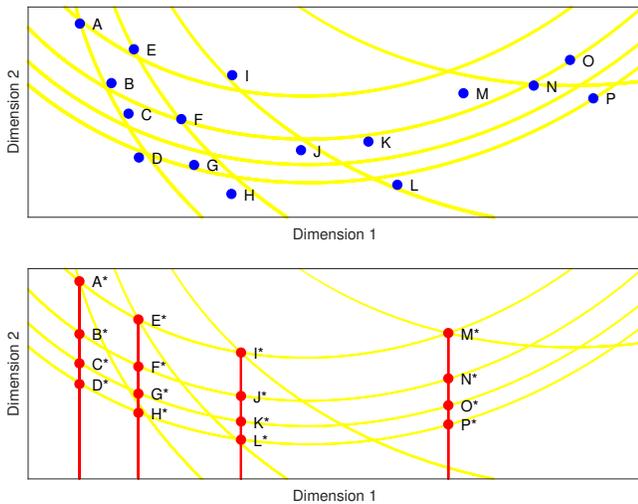


Figure 2: Top: The blue data-points, representing the different levels of key intensity and KTRF are shown in the MDS perceptual space along the two dimensions. The yellow lines represent the proposed parametric log-polar model that best fits the data. **Bottom:** The red lines show the location in the perceptual space of the stimuli with ‘adjusted’ brightness.

examination of the stimuli, suggests that *Dimension 2* relates to the perception of shadow strength. One critical insight from the MDS space is that the effect of KTRF changes as a function of key intensity. When the key illumination is high (100%), participants appear to perceive the “true” scene brightness (A, B, C, D). This is consistent with them (mostly) “discounting” the effect of shadows. Note that slight deviation from vertical at the highest key intensity might be due to the fact that as the KTRF decreases (at a given key level), the average pixel intensity will decrease slightly. As the key intensity decreases, participants seem to lose the ability to discount shadows. At the 25% key level, the effects of KTRF are perceived almost entirely as changes in scene brightness (M, N, O, P). This is consistent with Peli et al. [1991]’s finding that contrast perception changes as a function of illumination.

Note that the experiment was performed using anchored scales. While this is standard in psychophysics, the exact choice of the anchors might cause biases and/or noise. As a control, we re-ran the experiment without anchors (with 16 new volunteers; 6 females, 10 males, aged 18–37, avg. age 24). We added a training session in which participants were shown all possible stimuli pairs in a random order without being asked to rate them. During this main part of the experiment, participants were told that 0 meant that the stimuli are identical and 16 meant that the stimuli are extremely different. To additionally test for accuracy, each pair was viewed 4 times in this experiment, with each pair being seen once before any pair was seen a second time (blockwise randomization) resulting in a 50–60 minute experiment. The resulting MDS space was nearly identical to the anchored MDS space (Procrustes distance $d = 0.0130$ [Dryden and Mardia 1998]). We did not find any effect of anchoring or any significant improvement in the result with more repetitions per stimuli comparisons.

5.1 Parametric Model

Logvinenko and Maloney [2006] noticed that the MDS reconstruction of their data showed a “fan-like” structure. They thus parametrically modeled the perceived dissimilarity (d) of two achromatic Munsell chips, one with albedo A_i and illuminant L_k , and the other with albedo A_j and illuminant L_m , with the weight w_a and w_l as:

$$d = [w_a(\log A_i - \log A_j)^2 + w_l(\log L_k - \log L_m)^2]^{1.1} \quad (1)$$

Since our data show a similar fan-like structure, we propose using a similar parametric model, describing a circular pattern. Specifically, we propose using the following parametric model:

$$[x, y] = [R \cdot \cos(T) + c_x, R \cdot \sin(T) + c_y] \quad (2)$$

where:

$$R = w_{br} \cdot \log_2\left(\frac{100}{B}\right) + w_{kr} \cdot \log_2(K) + c_r$$

$$T = w_{bt} \cdot \log_2\left(\frac{100}{B}\right) + w_{kt} \cdot \log_2(K) + c_t$$

and:

$$B = \text{key intensity percentage (100, 75, 50, 25, etc.)}$$

$$K = \text{key-to-fill ratio (1/1, 2/1, 3/1, 4/1, etc.)}$$

We used a least-squares curve fitting tool in Matlab (*lsqcurvefit*()) to fit the model to the MDS plot coordinates (Appendix, Table 1) which yielded the 4 unknown weights (w_{br} , w_{kr} , w_{bt} and w_{kt}) and 4 unknown constants (c_x , c_y , c_r and c_t), with the residual norm of 0.08, indicating a good fit. The weight and constant values of the fit can be found in Appendix, Table 3. The yellow lines in Figure 2 depicts the proposed log-polar model approximating the MDS experiment result.

Since *Dimension 1* seems to be perceived scene brightness, we can use the model to create iso-brightness curves (red lines in Figure 2, Bottom). That is, we can produce multiple KTRF levels that maintain the same level of perceived scene brightness (see Appendix, Figure 4).

5.2 Perceptual Validation of the Model

To validate that this model can be used to produce stimuli with different KTRF levels while maintaining similar overall brightness, we ran an online “matching to sample” experiment, where 30 people (14M, 16F, age 18–52, avg. age 28) took part. The conditions for this experiment were less controlled than the previous (monitor settings, brightness of the room, etc.), so that we could ensure our results were valid in a non-laboratory setting.

For each perceived iso-brightness level (red lines in Figure 2, Bottom), we generated the adjusted stimuli (A^* - P^*) from the model (see Appendix, Figure 4). In each trial, one key intensity (100%, 75%, 50%, or 25%) was randomly chosen. The 1:1 KTRF for the chosen key level was displayed at the top of the screen as the “sample” (i.e., A, E, I or M). One of the remaining KTRF values (2:1, 3:1, 4:1) was chosen for the two “match” stimuli (original vs. adjusted), which were displayed at the bottom. The participant had to choose which of the matches they believed to be more similar to the sample, in terms of overall brightness. Each combination of key intensity and KTRF was repeated 4 times in a random order and with a counterbalanced presentation of the left and right matches, yielding 48 trials in total.

On average, participants chose our adjusted stimuli 3 out of 4 times (75.35%), which is above chance level. The percentage of the time participants chose the adjusted stimulus were submitted to a two-way, repeated-measures ANOVA with brightness (4) and KTFR (3) as within-participants factors. No main effects or interactions were found, which implies that no systematic differences across brightness or KTFR were found in the percentage of the times the adjusted stimuli were chosen. Note that some stimuli were altered more in Euclidean space than others (e.g., $M^* - P^*$ were further from the original $M - P$ than $A^* - D^*$ were from $A - D$). Thus, our log-polar model can create different KTFR levels with more consistent perceived brightness, regardless of the actual intensity.

6 DISCUSSION

In this paper, we investigated the perception of overall brightness levels and shadows strength in a carefully controlled, computer-generated scene. Our results extend previous knowledge from psychophysical experiments to a more complex three-point lighting setting with a cartoon-style virtual character.

Our key findings about human visual perception are: We have shown that *the ability to distinguish between shadows and overall illumination depends on the level of overall illumination*. When the overall illumination is high, people can reliably discount the effect of shadows to estimate the true light intensity. As the overall light intensity decreases, the effect of the shadows is increasingly mis-allocated to be a change in overall illumination. We have also shown that *the overall pattern of discrimination between brightness and shadow strength is very regular and can be captured with a log-polar function*. We demonstrated that the model can be used to create stimuli with different KTFR but perceptually equal overall brightness. In other words, key intensity and KTFR can now be more effectively and systematically sampled for future experiments, such as studying the effect of lighting on the perception of higher level factors such as emotion and appeal (e.g., [Wisessing et al. 2016]).

Who might benefit from this information? Our approach can be used for any CG character or scene after collecting a new dissimilarity matrix (which takes only 15 minutes per participant) and calculating the weights and constants. This model could also be applied to control illumination in studio photography, using the same approach. Of course, with enough data, it should be possible in the future to train the mathematical model so that the proximity structure of perceived brightness and shadow for other characters could be estimated without requiring dissimilarity ratings.

We believe our parametric model is a first step towards creating a perceptually-guided lighting tool for CG characters. The tool would allow a user to change shadow intensities without affecting the scene brightness (by automatically adjusting the lighting intensity to match the parametric model). Such a tool would have benefits in ensuring lighting consistency across shots and for maintaining appealing scene brightness [Wisessing et al. 2016] while allowing shadows to be added for mood or dramatic effect. This could also speed-up the shot brightness matching in the post-processing stage of the CG pipeline.

Our study was limited by assessing only the typical portrait lighting for cartoon characters with a middle-gray background.

Additionally, we only focus on artistic lighting which has relevance in CG and photography, as opposed to unmanipulated natural or actual light, which would capture the complex optical interactions between object appearance and context [Xia et al. 2017a,b]. Future work will investigate if our results generalize for other character models and other controlled illumination setups.

ACKNOWLEDGMENTS

This research was funded by Science Foundation Ireland under the ADAPT Centre for Digital Content Technology (Grant 13/RC/2016) and the Game Face (Grant 13/CDA/2135) project. We thank the Mery Project for the character model and rig.

REFERENCES

- Ansel Adams. 1948. *The Negative: Exposure and Development Basic Photo 2*. Morgan and Lester 98 (1948).
- Edward H Adelson. 2000. Lightness Perception and Lightness Illusions. *New Cogn. Neurosci* 339 (2000).
- Edward H. Adelson. 2011. Checkershadow Illusion. <http://persci.mit.edu/gallery/checkershadow>. (2011). Retrieved July 8, 2019.
- Wenyang Bi and Bei Xiao. 2016. Perceptual Constancy of Mechanical Properties of Cloth Under Variation of External Forces. In *Proceedings of the ACM Symposium on Applied Perception (SAP '16)*. ACM, New York, NY, USA, 19–23.
- Ingwer Borg and James Lingoes. 2012. *Multidimensional similarity structure analysis*. Springer Science & Business Media.
- David H Brainard. 2003. Color appearance and color difference specification. *The science of color 2* (2003), 191–216.
- Alfred M Bruckstein. 1988. On shape from shading. *Computer Vision, Graphics, and Image Processing* 44, 2 (1988), 139–154.
- Trevor F. Cox and Michael A. A. Cox. 2001. *Multidimensional scaling, 2nd Ed*. Boca Raton, FL: Chapman & Hall/CRC, 2001.
- Douglas W. Cunningham and Christian Wallraven. 2011. *Experimental design: From user studies to psychophysics*. AK Peters/CRC Press.
- Douglas W Cunningham, Christian Wallraven, Roland W Fleming, and Wolfgang Straßer. 2007. Perceptual reparameterization of material properties. In *Proceedings of the Third Eurographics conference on Computational Aesthetics in Graphics, Visualization and Imaging*. Eurographics Association, 89–96.
- Katja Doerschner, Roland W Fleming, Ozgur Yilmaz, Paul R Schrater, Bruce Hartung, and Daniel Kersten. 2011. Visual motion and the perception of surface material. *Current Biology* 21, 23 (2011), 2010–2016.
- Ian L Dryden and Kanti V Mardia. 1998. *Statistical shape analysis*. Vol. 4. Wiley Chichester.
- GT Fechner. 1966. Elements of Psychophysics. Adler HE, translator; Howes DH, Boring EG, editors. (1966).
- David H. Foster. 2011. Color constancy. *Vision Research* 51, 7 (2011), 674 – 700. Vision Research 50th Anniversary Issue: Part 1.
- Alan Gilchrist, Christos Kossyfidis, Frederick Bonato, Tiziano Agostini, Joseph Cataliotti, Xiaojun Li, Branka Spehar, Vidal Annan, and Elias Economou. 1999. An anchoring theory of lightness perception. *Psychological review* 106, 4 (1999), 795.
- Alan L Gilchrist. 2013. *Lightness, brightness and transparency*. Psychology Press.
- Harold Hill and Vicki Bruce. 1996. The effects of lighting on the perception of facial surfaces. *Journal of Experimental Psychology: Human Perception and Performance* 22, 4 (1996), 986–1004.
- Alan Johnston, Harold Hill, and Nicole Carman. 1992. Recognising Faces: Effects of Lighting Direction, Inversion, and Brightness Reversal. *Perception* 42, 11 (1992), 1227–1237.
- Ludwig Kardos. 1934. Ding und Schatten. Eine experimentelle Untersuchung über die Grundlagen des Farbensehens. *Zeitschrift für Psychologie und Physiologie der Sinnesorgane*. Abt. 1. *Zeitschrift für Psychologie* (1934).
- Frederick AA Kingdom. 2011. Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research* 51, 7 (2011), 652–673.
- Manuel Lagunas, Sandra Malpica, Ana Serrano, Elena Garces, Diego Gutierrez, and Belen Masia. 2019. A Similarity Measure for Material Appearance. *ACM Transactions on Graphics (SIGGRAPH 2019)* 38, 4 (2019).
- Alexander D Logvinenko and Laurence T Maloney. 2006. The proximity structure of achromatic surface colors and the impossibility of asymmetric lightness matching. *Perception & Psychophysics* 68, 1 (2006), 76–83.
- Gerald Millerson. 1991. *The Technique of Lighting for Television and Film*. Focus Press.
- Ennio Mingolla. 1983. Perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology* 9, 4 (1983), 583–595.

- Fred E Nicodemus, Joseph C Richmond, Jack J Hsia, IW Ginsberg, and T Limperis. 1992. Geometrical considerations and nomenclature for reflectance. *NBS monograph* 160 (1992), 4.
- Zahide Pamir and Huseyin Boyaci. 2016. Context-dependent lightness affects perceived contrast. *Vision research* 124 (2016), 24–33.
- Eli Peli, Jian Yang, Robert Goldstein, and Adam Reeves. 1991. Effect of luminance on suprathreshold contrast perception. *JOSA A* 8, 8 (1991), 1352–1359.
- Marion Webster Richardson. 1938. Multidimensional psychophysics. *Psychological Bulletin* 35 (1938), 659–660.
- Szymon Rusinkiewicz, Michael Burns, and Doug DeCarlo. 2006. Exaggerated Shading for Depicting Shape and Detail. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25, 3 (July 2006).
- Stanley S Stevens. 1957. On the psychophysical law. *Psychological review* 64, 3 (1957), 153.
- Maureen C Stone. 2003. A field guide to digital color. (2003).
- Warren S Torgerson. 1952. Multidimensional scaling: I. Theory and method. *Psychometrika* 17, 4 (1952), 401–419.
- Jan Jaap R. van Assen and Roland W. Fleming. 2016. Influence of optical material properties on the perception of liquids. *Journal of Vision* 16, 15 (12 2016), 12–12.
- Floris L Van Nes and Maarten A Bouman. 1967. Spatial modulation transfer in the human eye. *JOSA* 57, 3 (1967), 401–406.
- Peter Vangorp, Jurgen Laurijssen, and Philip Dutré. 2007. The Influence of Shape on the Perception of Material Reflectance. *ACM Trans. Graph.* 26, 3, Article 77 (July 2007).
- Peter Vangorp, Karol Myszkowski, Erich W. Graf, and Rafal K. Mantiuk. 2015. A Model of Local Adaptation. *ACM Transactions on Graphics* 34, 6, Article 166 (Oct. 2015), 13 pages.
- Veronika Šoltészová, Daniel Patel, and Ivan Viola. 2011. Chromatic Shadows for Improved Perception. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Non-Photorealistic Animation and Rendering (NPAR '11)*. ACM, New York, NY, USA, 105–116.
- Annette Werner. 2007. Color constancy improves, when an object moves: High-level motion influences color perception. *Journal of Vision* 7, 14 (2007), 19–19.
- Pisut Wisessing, John Dingliana, and Rachel McDonnell. 2016. Perception of Lighting and Shading for Animated Virtual Characters. In *ACM Symposium on Applied Perception*. ACM, 25–29.
- Ling Xia, Sylvia C Pont, and Ingrid Heynderickx. 2017a. Light diffuseness metric part 1: Theory. *Lighting Research & Technology* 49, 4 (2017), 411–427.
- Ling Xia, Sylvia C Pont, and Ingrid Heynderickx. 2017b. Light diffuseness metric, Part 2: Describing, measuring and visualising the light flow and diffuseness in three-dimensional spaces. *Lighting Research & Technology* 49, 4 (2017), 428–445.

APPENDIX

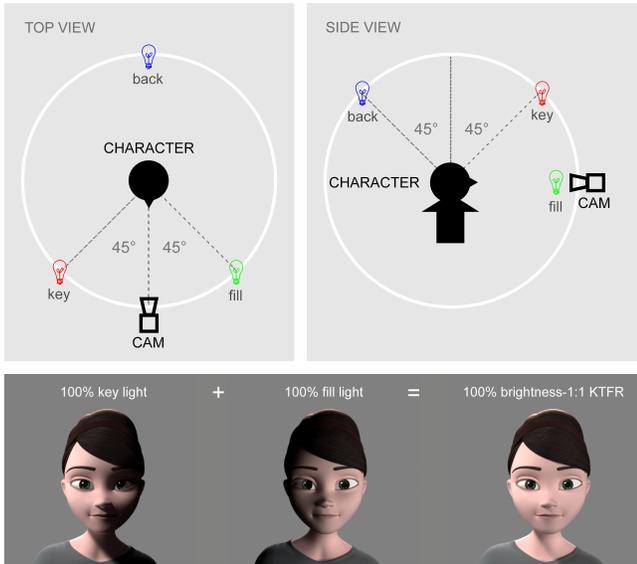


Figure 3: Top: Three-point lighting setup. Bottom: The contribution of key and fill lights in the 100% brightness-1:1 KTRF condition.

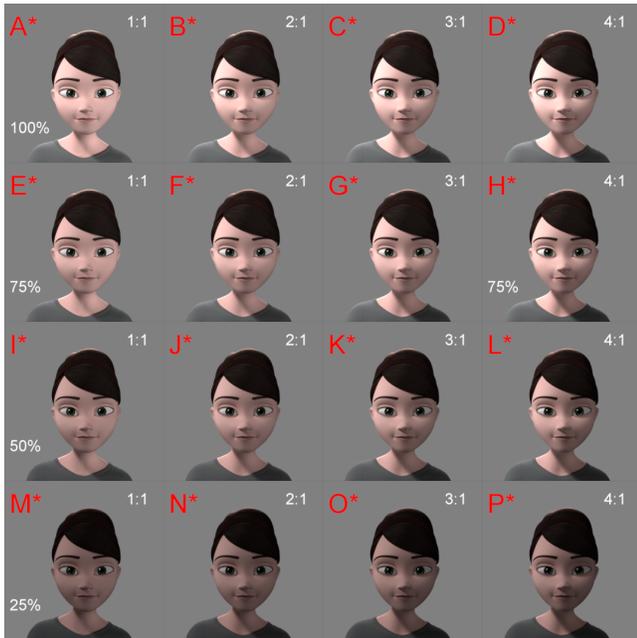


Figure 4: Our perceptually ‘adjusted’ stimuli set. Notice how the overall brightness is more consistent across the rows than in Figure 1.

Table 1: The coordinates of the blue data-points in Figure 2 (top), representing the locations of the stimuli the MDS perceptual space.

stimuli	x	y
A	-1.2052	0.5052
B	-1.0251	0.1656
C	-0.9276	-0.0089
D	-0.8683	-0.2590
E	-0.8970	0.3587
F	-0.6266	-0.0398
G	-0.5536	-0.3015
H	-0.3408	-0.4670
I	-0.3362	0.2109
J	0.0560	-0.2170
K	0.4404	-0.1682
L	0.6054	-0.4148
M	0.9824	0.1077
N	1.3832	0.1514
O	1.5901	0.2979
P	1.7230	0.0787

Table 2: The coordinates of the red data-points in Figure 2 (top), showing the location in the perceptual space of the stimuli with “adjusted” brightness.

stimuli	x	y
A*	-1.2075	0.5277
B*	-1.2075	0.2270
C*	-1.2075	0.0593
D*	-1.2075	-0.0586
E*	-0.8715	0.3099
F*	-0.8715	0.0411
G*	-0.8715	-0.1133
H*	-0.8715	-0.2230
I*	-0.2861	0.1212
J*	-0.2861	-0.1269
K*	-0.2861	-0.2725
L*	-0.2861	-0.3764
M*	0.8961	0.2324
N*	0.8961	-0.0267
O*	0.8961	-0.1800
P*	0.8961	-0.2888

Table 3: Weight and constant values of the parametric model (Equation 2) that best fits the MDS perceptual space—the yellow lines in Figure 2.

variables	values
c_x	-0.8844
c_y	2.3602
w_{br}	0.4569
w_{kr}	0.3687
c_r	1.8608
w_{bt}	5.6586
w_{kt}	24.9604
c_t	-100.00