

Perceptual visualization: experiment, measure and apply

R. Jøsok¹ ¹University of Bergen, Norway

Abstract

Effectiveness, ease-of-use and accuracy are only some of the metrics by which a perceptual technique can be evaluated. In this paper we will look at applications of perceptual visualization and the visual image properties they seek to enhance. We will look at how researchers measure aspects of the human perceptual system and how new visualization techniques are evaluated today. Data from user studies have been used to define intrinsic image quality metrics that can be computationally solved. In the field of human shape perception a computable intrinsic quality metric has yet to be defined. We will examine the field of shape enhancement techniques in detail and how we measure human shape perception.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: perceptual visualization, quality metric, intrinsic image quality, shape perception

1. Introduction

Chaomei Chen rated the top ten unsolved problems in information visualization in 2005. [Che05] Our understanding of elementary cognitive-perceptual tasks, the effect of aesthetics in information visualization and the lack of intrinsic quality measures are some of the issues he felt should be solved for the field to progress. In this paper we will examine the fields history and progress since his publication.

Perceptual visualization techniques are tightly bound to cognitive and visual sciences and evolve based on new knowledge about our elementary cognitive-perceptual abilities. New perceptual visualization techniques are evaluated on a case-to-case basis even when they have a common quality metric. We will look at some approaches to perceptual visualization techniques, their goals and their quality metrics. We will look at how user studies and research have been used to automate the process of quality evaluation. We will examine areas where the statistics derived from user studies have been used to automate the evaluation of perceptual visualization models.

We will look at a specific quality metric definition in detail, namely how to measure human shape estimation. We will examine issues surrounding current shape estimation evaluation tools, and discuss the existence of intrinsic image qualities that affect human shape-from-shading. Can an automated shape perception accuracy quality metric be defined? Let us start by looking at the field of perceptual vi-

sualization in general. Where do we get data to support or discard the quality of a visualization?

2. Expanding the user study

User studies can give valuable data in visualization fields. The issue of gathering data from a large and diverse population is common. The use of Mechanical Turk to assess visualization design has been explored by Heer and Bostock (2010). [HB10] Mechanical Turk allows a researcher to collect data from a multitude participants while keeping the cost low. Issues surrounding the collection of data were noted, such as forming a definition of an arbitrary task to a participant. Since researchers never get to assess the users that participate in the user study, they have to filter out corrupted data. A level of error tolerance needs to be defined. These issues were also noted by Cole et al. (2009) when they used Mechanical Turk in their user studies. [CSD*09]

3. Quality metrics

A quality metric is a measurement of quality. Perceptual quality is usually measured during user studies. A perceptual quality metric is the measurement of a visualization technique in relation to some task. Measurements can be precision, speed and user preference. Techniques can be the use of color, shapes, lines and textures in a visualization. Perceptual tasks include visual search and interpretation of visualized data.

User studies can be expensive and time-consuming and are not a viable option for many researchers. User studies need to be carefully executed to retain the integrity of the data. Building on data from user studies it is possible to derive a collection of visual qualities that are inherent in the image, an intrinsic quality. By defining an intrinsic image quality metric it is possible to compute the quality of a visualization. Researchers can focus on developing new visualization techniques and estimate their quality without the use of time-consuming and expensive user studies. There are many presentational techniques in visualization and they all have their strengths and weaknesses. Is it possible to automatically select and update a visualization based on user need?

4. Presentational quality

Conati et al. (2011) discuss the possible automation of visual presentation. A visualization that automatically updates the information on the screen according to detected user needs. They consider the possible use of eye-tracking data as a metric to determine a visualizations effect on visual attention, cognitive processing and presentational effectiveness. The presentation can then be refined by determining the personal traits, cognitive ability, preferences and the domain knowledge and goals of the user. [CCH^{*}11] The possible use of eye-tracking data as a measurement of visual and cognitive processes have been explored by Gotz and Wen (2009), and by Brusilowsky et al. (2006).

5. Artistic quality

Nonphotorealistic techniques have seen an increased use in visualization design. Visualizations that approach artistic techniques have been measured by their artistic value. In a user study by Healey and Enns (2002) the artistic value of their abstract visualization technique was measured. [HE02] They used painting techniques from the impressionist movement to render nonphotorealistic images of abstract weather data. Stroke direction, brush texture, brush size, color and intensity where mapped to aspects of the multivariate data collected from weather stations. It was an effort to preserve the integrity of multi-dimensional data while improving visual interaction. [Hea92] In a user study participants were asked to rate the artistic quality of a set of images that included human made art and painterly computer visualizations.¹ Their findings suggested that a scientific visualization had the potential of being considered a work of art.

Any visualization technique that approach an artistic style can be measured by its artistic "correctness". If the quality of an artistic technique is well defined it is possible to automate the evaluation of artistic visualizations. In a study by Cole et al. (2008) they determined where artists draw lines. They presented artists with shaded 3D models on a computer screen and asked them to recreate the "best" representation of the model with a line-drawing. They could

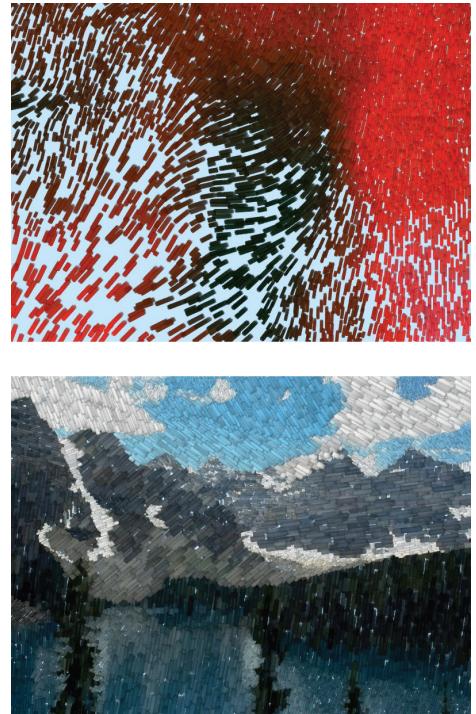
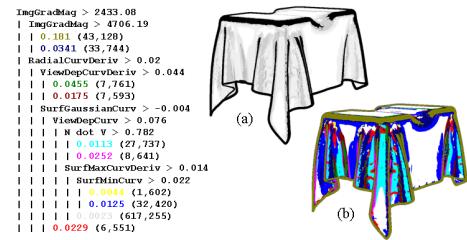


Figure 1: Example images used in the artistic value study by Healey et al. [HE02] Top: impressionist weather data. Bottom: a nonphotorealistic rendering of a photograph of Lake Moraine in Banff, Canada. The artistic value of a piece was evaluated in relation to the other artworks.

then compare the placement of lines of current computational line-drawing techniques to artistic line-drawing techniques. [CGL^{*}08] They examined the possibility of artistic line placement prediction (2). They also used their statistics to determine if multiple surface properties were evaluated by artist when determining line placement. Current computational line-drawing techniques only evaluated local properties.



Another example of an artistic quality metric can be found in the paper by A. Gooch et al. (1998). Inspired by existing technical illustrations, they use introduce a phong-shading model that use colored shading. The correct placement of lines and the possibility use of color shades was determined by an examination of existing, human-made technical illustrations. [GGSC98]

6. Accentuation of relevant data: The salience metric

A subset of perceptual visualization techniques attempt to guide human visual attention to areas of importance. The salience quality metric is a measurement of a visualizations success at achieving this goal. A high quality salient image will reduce the time it takes a participant to locate task-relevant data. Eye-tracking techniques have also been used to determine areas of high visual attention in images and their correlating visual properties.

A salience metric is the measurement of visual properties of an image that affect visual attention. An automated salience quality metric use a data relevance mask and a visual salience map, as demonstrated by Janicke et al. (2010). [JC10] The data relevance mask is a mapping of data importance to the pixels in an image. A salience map is constructed to check if the data relevance mask correlate to the visual salience of an image. The salience map shows areas in the image that attract viewer attention. A salience map can be calculated from the weighted influence of intensity, color and orientation of pixels as shown in the formula below.¹ There are many visual features that can affect image salience. This have lead to a multitude of possible parameters and parameter influence weights in salience mappings.

$$S = \lambda_1 N(\overline{F_I}) + \lambda_2 N(\overline{F_C}) + \lambda_3 N(\overline{F_O}) \quad (1)$$

$N(F)$: normalized mapping

F_I : map of pixel intensity influence on salience

F_C : map of pixel color influence on salience

F_O : map of pixel orientation influence on salience

$$\lambda_1 + \lambda_2 + \lambda_3 = 1$$

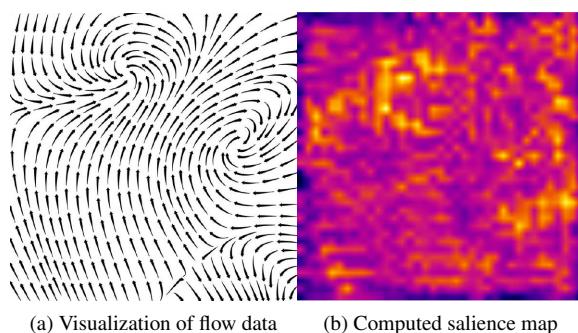


Figure 3: Salience mapping of flow data. [JC10]

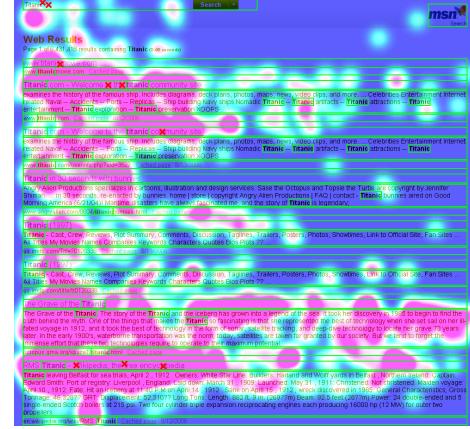


Figure 4: Heat map visualization of the number of fixations across 3 users on a page of search results for an informational task with long contextual snippets. Boxes indicate defined areas of interest. [Cut07] Notice the similarities to the salience map generated in figure 3b. Boxes of interest would be equivalent to the data relevance mask used in salience metrics.

Corcoran et al. (2010) used this quality metric to evaluate the effect of the saliency reduction filters used in their volumetric renderer. They combined several non-photorealistic shading techniques in their renders to assist users in locating important features of the data while keeping extraneous information on the screen for reference. [CRD10] Two saliency reductions appear in their work. The Kuwahara filter produces a painterly effect by blurring internal parts while keeping edges sharp, and saturation abstraction was used to desaturate unimportant visual structures.

7. Visibility of relevant data

Occlusion and shadows are often added to a visualization to enhance depth perception in computer graphics. Shadows add depth information by reducing attenuation (pixel intensities) in the image.

Soltészová et al. (2011) [SPV11] used chromatic shadows to reduce over-occlusion in volume visualizations. They added an attribute of "shadowiness" to voxels in the volume and implemented a shadow transfer function that mapped "shadowiness" value to colors in RGB space. They received input from a medical professional and an artist to determine the color space to test in three user studies. In the first study they measured participant surface orientation perception with the gauge figure task. The visual stimuli were phong-shaded volume renders with chromatic shadows from black to lighter shades of blue. The measured data suggested a trend towards surface perception improvement with the use of chromatic shadows. The largest increase was found from

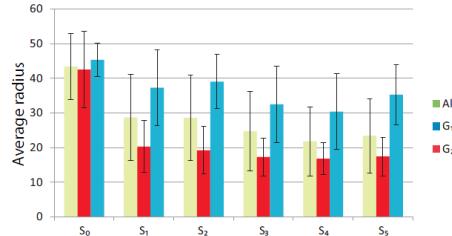


Figure 5: Average radius of the area with respect to shadow colors $S_0 \dots S_5$ compared to calculated standard deviation. G_1 have strong features in the neighbourhood while G_2 have weak features. [SPV11] As seen in the model, the chromatic shading techniques improves the visibility of data in edical volumetric rendering techniques

the transfer from black shadows to the darkest shade of blue shadows.

The second user study measured contrast perception. They determined the smallest circular area in which a user could notice a change in surface detail. They detected an increase in contrast perception with models using light chromatic shadows. They also noted user confusion when presented with large regions of light chromatic shadows. Participants could not determine if the apparent difference in contrast came from shading or from illusory visual artifacts. In the third user study the effect of chromatic shadows on depth perception was examined.

Another novel approach to enhance shape perception in visualizations is presented in a paper by Vergne et al. [VPB*09] Their method can increase visibility of important surface features with arbitrary illumination setups and materials, as seen in the figure ???. The algorithm exploit the inability of the human visual systems to locate local inconsistencies in light direction on local surfaces. [FTA04,?] In their method they locally stretch the sphere of potential illumination directions to enhance or attenuate the depiction of surfaces, and reformulate the reflectance formula of surfaces to take warped lighting into account in their renders.

The method first extracts the view-centric visual features of a model. Since the method is view-centric it scales relevant to the object distance from the viewpoint. In comparison, an object-centric model would have to simplify the object manually to scale well in complex scenes. The view-centric algorithm determine local curvature properties (silhouette, crease and inflection) by computing the first and second order derivatives of the relevant depth gradient map across the viewing plane, and determines a curvature tensor field H as seen in the formula below (2). K_u and K_v are the principal curvatures, and u and v correspond to the principal directions. They generate a salience map by mapping mean curvature

$$H = (K_u + K_v)/2$$

to the color scale seen in the figure (??). They can then exaggerate the deformation of reflection patterns that are reflected in curved surfaces by scaling the curvature of the mapping. The intensity of the scaling need to conform to the principal ratio of curvature. [?]

$$H = (u \ v) \begin{pmatrix} K_u & 0 \\ 0 & K_v \end{pmatrix} (u \ v)^T \quad (2)$$



Figure 6: The color scale used to visualize surface curvature properties. [VPB*09] Warm hues are concave while cold hues are convex.

R. Shacked (2001) used a perceptual quality metric to position and orient lighting in a phong-shaded 3D scene. [SL01] He defined the perceptual quality metric as the linear combination of six predefined target terms by the formula

$$f_Q = f_{grad} + f_{edge} + f_{var} + f_{mean} + f_{hist} + f_{dir} \quad (3)$$

The algorithm takes as input a model M , view parameters and produces a rendered version of the model. Then it performs an image analysis on f_Q to determine the current perceptual quality of the render. It reduced the evaluation and improvement of render quality to an optimization task. Perceptual quality improves as f_Q approach zero. While the techniques are quiet old, the definition of a computable perceptual quality metric is novel.

8. Human shape perception

8.1. Intrinsic image qualities that affect human shape-from-shading

Visual properties in images that affect shape and depth perception include perspective, parallax motion, occlusion, shadows, shading, colors and texture. [DAG02] Human shape-from-shading perception display remarkable adaptability when global shading cues are removed. To define a shape perception metric it is necessary to determine how the

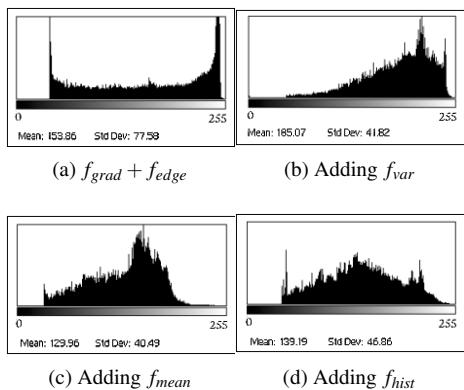


Figure 7: Evaluation of the quality function defined by Shacked. [SL01] Cumulative effect of the selected target terms on the histogram. The data was retrieved by image analysis of 3D renders.

human visual system process images, and to determine image properties that affect human shape perception.

Possible sources of human shape understanding have been explored: depth [BM88], orientation [Koe84] and curvature [TM83]. Another measurements is shape comparison or recognition.

Visual ambiguities appear as a consequence of the transformation from multi-dimensional data to a two-dimensional representation. The four dimensions of the world (3D space and time) are not represented in static 2D images. With the loss of parallax motion and stereoscopic depth cues, the human visual system displays some ambiguities.

The convex-concave ambiguity is born from assumptions the human visual systems have about the real world. Objects in the world are usually globally convex and light will generally come from above, so it is a naturally evolved efficiency mechanism. Koenderink et al. felt that the focus on this ambiguity was unnecessary as it is visually corrected when the visual system is presented with global shading cues. [Koe03]

O'Shea et al. (2008) used the gauge figure task to evaluate human accuracy at determining the light direction of a locally shaded lambertian surface. [OBA08] They confirmed that our visual system is assumes a light direction from above when viewing a shaded surface. They determined from the data that human evaluation of shape orientation is peaks when the angle between light direction and viewing direction is 20-30 degrees above the viewport. This gives us a measurement to evaluate two shape representation quality metrics; light orientation quality and light position quality. Fleming et al. (2004) have proven that the use of specular reflection in surface shading increase human shape estimation accuracy. [FTA04]

The human visual system also loose the ability to dis-

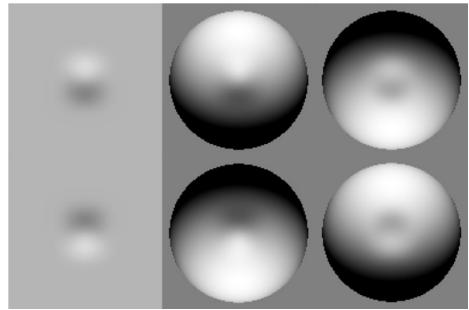


Figure 8: A visual example by Koenderink et al. [Koe03] where the convex-concave ambiguity is visually solved when presented with global shading cues.

tinguish between surface orientation variation and surface albedo change. Change in attenuation can be the result of change in surface orientation in addition to a change in the reflective properties of the surface. Another notable ambiguity is the bas-relief ambiguity presented by Belhumeur et al. (1999). [BKaY99] Human shape from shading can only be computed up to an affine transformation of the surfaces.

8.2. Measuring human shape perception

In the field of computer vision an algorithm is evaluated on how well it can reproduce the geometry depicted in an image. When the "real" geometry is known, the generated geometry can be compared to check for accuracy. The efficiency of an algorithm is usually also measured to be compared to the relative running time of other shape-from-shading algorithms. [CS99] The human analogy of this evaluation would be to have people sculpt what they see on the screen and compare accuracy and speed. While it is possible to take this approach, it is extremely inefficient and riddled with possible sources of erroneous stimuli.

Some of the earliest shape estimation evaluation tools used real-world comparison [TM83] and stereoscopic depth cues [BM88]. Todd et al. used cylinders with different curvature intensities and asked participants to rate the relative intensity of a computer generated surface. Bulthoff et al. asked participants to position depth probes at the perceived surface in 3D space. While the researchers managed to gather data about human depth perception using these tools, the methods suffered from a lack of an intuitive input mechanism and sufficiently complex objects to analyze. Another metric is human shape recognition.

Koenderink et al. introduced a new input mechanism with their gauge figure task (1992). [KV92] The task was fully implemented on the computer and used input from the mouse. The task is to orient a gauge figure (9) so that the base is aligned with the tangent plane and the pin with the surface normal. The perceived surface orientation can then be compared with the actual surface normal of the geometry

and degree of error can be computed. The gauge figure has seen extensive use since its introduction and is the main SEE tool in use today. [OBA08, CSD*09, CRD10, SPV11]

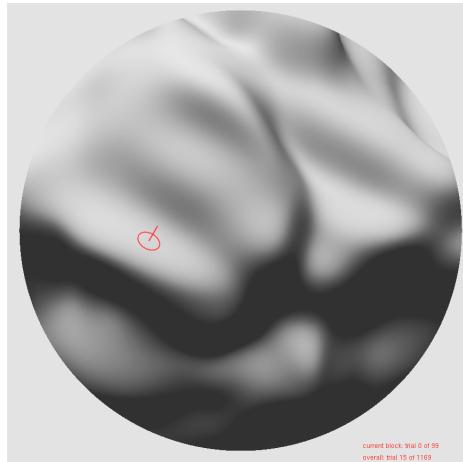


Figure 9: An implementation of the gauge figure task by O’Shea et al. [OBA08] Participants were asked to orient the gauge (red pin) to the incoming light direction.

Georgieva et al. (2008) performed a medical study of human shape understanding using a multiple shape estimation evaluation tools and a multitude of shape representations. [GTPO08] In the study they measure human brain activity during both passive and active human shape-from-shading activity. They performed control experiments to determine cognitive activity as a result of tool interaction. The main focus of the study was to determine brain activity during shape-from-shading tasks and the effect of different shading techniques. The impact of tool usage was not clearly documented.

8.3. Dissecting the gauge figure

There is a certain ambiguity about the gauge figure task itself. A user is asked to evaluate the orientation of a computer generated surface with a computer generated gauge. When O’Shea et al. (2008) used the gauge figure task to experiment on human light direction bias, they noted that participants in the experiment might have misperceived the slant of the gauge figure. [OBA08] They conclude that further research into the gauge figure task is necessary to accurately measure human shape estimation abilities.

Cole et al. (2009) use the gauge figure task to determine how adapt line drawing are at depicting shape. [CSD*09] Test subjects were asked to estimate the orientation of shaded surfaces, line drawings made by artists and computer generated line drawings. Cole et al. also noted that estimation errors could originate from user misinterpretation of the gauge. The gauge task also reward shape estimation on

renders that match the contour of the surfaces. Artist line-drawings were not accurate enough and the data collected from these studies had to be adjusted to account for the inaccuracy.

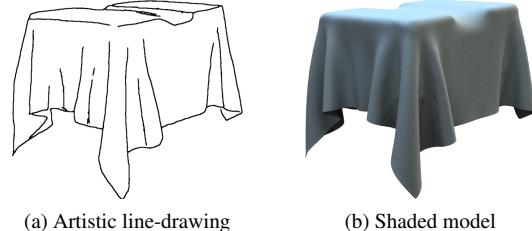


Figure 10: Example stimuli used in the user study by Cole et al. (2009) [CSD*09] They reduced structural noise and aliasing of all the line-drawings to normalize the aesthetic quality of the images.

8.4. Dissecting the collected data

Cole et al. performed a careful study of the data from their experiment on line-drawings. [CSD*09] They evaluated the introduction of erroneous data at multiple stages before, during and after data collection. They used Mechanical Turk [HB10] to gather data from a diverse population. The workers of Mechanical Turk are more likely to do tasks that take 5-10 minutes. To accommodate for this they had to restrict the number of gauge placements during a session. To evaluate worker accuracy they had the workers place two gauges at each control point during a session. If duplicated gauges vary by a large degree it can indicate a guess of surface orientation. To avoid training effects they asserted that no worker would see the same control points or model versions twice. This reduced the number of times a worker could perform the task to 52 sessions. They gathered a total of 275K gauge placements during the experiment.

Further processing on the gauge data was performed. Cases of erroneous gauge orientations due to workers misunderstanding the task appeared. The data had to be filtered. Two rules were set to automatically filter out bad data. If fewer than 70% of the duplicated gauges were placed within 30 degrees of each other, the session was ruled to have been rushed. If all gauges in a session were placed within 5 degrees of each other, the session was ruled to have been misunderstood.

Two compensational transformations were then performed on the data. Some of the stimuli in the study used perspective view, and these had to be translated to orthographic view to correlate to the orthographic gauge figure. Secondly, the bas-relief ambiguity was taken into account. [BKaY99] They define an affine transformation (4) based on the work of Konderink (2001).

$$(x, y, f(x,y)) \rightarrow (x, y, f(x,y) + \lambda x + \mu y) \quad (4)$$

The variable λ affect scaling and μ and ν affect shearing. They use a non-linear optimization procedure to find values for λ , μ and ν that minimize the angular differences between the (normalized) transformed normals and the ground truth. In the paper they only used the bas-relief transformation in global studies of the data, but they noted that it might be possible to calculate the transformation variables for local surfaces as well. At this stage the gauge data is estimated to be filtered of the most common sources of error, and the data is ready to be used in further statistical analysis.

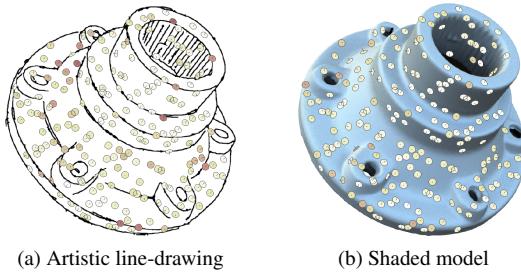


Figure 11: Visualization of the gauge figure orientations of participants in the user study by Cole et al. (2009) [CSD*09]. The gauges are colored darker when the orientation of the gauge deviates from the ground truth.

9. Conclusion

The field of visualization has progressed since Chen wrote his paper in 2005. Novel approaches to visual design and further statistical analysis of visual and perceptual qualities have lead to the implementation of several exciting visualization techniques. An example would be saliency mapping used with accentuation techniques. [JC10, CRD10] The computation of a salience map have even seen implementation in other quality metrics, such as the surface curvature salience map proposed by Vergne et al. [VPB*09]

Further research in conflicting quality metrics would be of use. The identification of quality metrics that define intrinsic qualities in an image can assist the visualization field in combining multiple qualities in a visualization. The use of chromatic shadows [SPV11], (chromatic shadows: adding depth cues with shadows while preserving visibility qualities)

Shape presentation quality to assist human shape-from-shading has seen little change since the early 90s. [KVK92] The gauge figure and user studies are still the most applied techniques today. Some new shape perception metrics have been reported, such as the effect of light direction and orientation on human shape from shading accuracy. [OBA08]

An extensive examination of the shape perception data from Cole et al. [CSD*09] could be used to determine trends in human shape-from-shading evaluation. The further identification of shape presentation quality metrics will surely help the field to progress.

References

- [BKAY99] BELHUMEUR P., KRIEGMAN D., A.L. YUILLE: The bas-relief ambiguity. *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition* 35, 1 (1999), 1060–1066. 5, 6
- [BM88] BÜLTHOFF H. H., MALLOT H. A.: Integration of depth modules: stereo and shading. *Journal of the Optical Society of America A Optics and image science* 5, 10 (1988), 1749–1758. 5
- [CCH*11] CONATI C., CARENINI G., HARATI M., TOCKER D., FITZGERALD N., FLAGG A.: User-Adaptive Visualizations: Can Gaze Data Tell Us When a User Needs Them? In *Workshops at the Twenty-Fifth AAAI Conference on Artificial Intelligence* (2011). 2
- [CGL*08] COLE F., GOLOVINSKIY A., LIMPAECHER A., BARROS H. S., FINKELSTEIN A., FUNKHOUSER T., RUSINKIEWICZ S.: Where do people draw lines? *ACM Transactions on Graphics* 27, 3 (Aug. 2008), 1. 2
- [Che05] CHEN C.: Top 10 unsolved information visualization problems. *IEEE computer graphics and applications* 25, 4 (2005), 12–6. 1
- [CRD10] CORCORAN A., REDMOND N., DINGLIANA J.: Perceptual Enhancement of Two-level Volume Rendering. *Computers & Graphics* 34, 4 (Aug. 2010), 388–397. 3, 6, 7
- [CS99] CRYER J., SHAH M.: Shape-from-shading: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 21, 8 (1999), 690–706. 5
- [CSD*09] COLE F., SANIK K., DECARLO D., FINKELSTEIN A., FUNKHOUSER T., RUSINKIEWICZ S., SINGH M.: How well do line drawings depict shape? *ACM Transactions on Graphics* 28, 3 (2009), 1. 1, 6, 7
- [Cut07] CUTRELL E.: What are you looking for?: an eye-tracking study of information usage in web search. *Proceedings of the SIGCHI conference on* (2007). 3
- [DAG02] DURAND F., AGRAWALA M., GOOCH B.: Perceptual and artistic principles for effective computer depiction. *2002 Course# 13* (2002). 4
- [FTA04] FLEMING R. W., TORRALBA A., ADELSON E. H.: Specular reflections and the perception of shape. *Most* (2004), 798–820. 4, 5
- [GGSC98] GOOCH A., GOOCH B., SHIRLEY P., COHEN E.: A non-photorealistic lighting model for automatic technical illustration. *Proceedings of the 25th annual conference on Computer graphics and interactive techniques - SIGGRAPH '98* (1998), 447–452. 3

[GTPO08] GEORGIEVA S. S., TODD J. T., PEETERS R., ORBAN G. A.: The Extraction of 3D Shape from Texture and Shading in the Human Brain. *Cerebral Cortex New York NY 18*, 10 (2008), 2416–2438. 6

[HB10] HEER J., BOSTOCK M.: Crowdsourcing graphical perception: using mechanical turk to assess visualization design. In *Proceedings of the 28th international conference on Human factors in computing systems* (2010), ACM, pp. 203–212. 1, 6

[HE02] HEALEY C., ENNS J.: Perception and painting: a search for effective, engaging visualizations. *IEEE Computer Graphics and Applications 22*, 2 (2002), 10–15. 2

[Hea92] HEALEY C.: Visualization of multivariate data using preattentive processing. 2

[JC10] JÄNICKE H., CHEN M.: A Salience-based Quality Metric for Visualization. In *Computer Graphics Forum* (2010), vol. 29, Wiley Online Library, pp. 1183–1192. 3, 7

[Koe84] KOENDERINK J.: What does the occluding contour tell us about solid shape. *Perception 13*, 3 (1984), 321–330. 5

[Koe03] KOENDERINK J.: Shape and shading. *The visual neurosciences 4* (2003). 5

[KV92] KOENDERINK J. J., VAN DOORN A. J.: Surface shape and curvature scales. *Image and Vision Computing 10*, 8 (1992), 557–564. 5

[KVK92] KOENDERINK J., VAN DOORN A., KAPPERS A.: Surface perception in pictures. *Attention, Perception, & Psychophysics 52*, 5 (1992), 487–496. 7

[OBA08] O’SHEA J. P., BANKS M. S., AGRAWALA M.: The assumed light direction for perceiving shape from shading. *Proceedings of the 5th symposium on Applied perception in graphics and visualization - APGV ’08* (2008), 135. 5, 6, 7

[SL01] SHACKED R., LISCHINSKI D.: Automatic lighting design using a perceptual quality metric. In *Computer graphics forum* (2001), vol. 20, Wiley Online Library, pp. 215–227. 4, 5

[SPV11] SOLTÉSZOVÁ V., PATEL D., VIOLA I.: Chromatic shadows for improved perception. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Non-Photorealistic Animation and Rendering* (2011), no. c, ACM, pp. 105–116. 3, 4, 6, 7

[TM83] TODD J. T., MINGOLLA E.: Perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology: Human Perception and Performance 9*, 4 (1983), 583–595. 5

[VPB*09] VERGNE R., PACANOWSKI R., BARLA P., GRANIER X., SCHLICK C.: Light warping for enhanced surface depiction. *ACM Transactions on Graphics (TOG) 28*, 3 (July 2009), 1–8. 4, 7