

Distortion in 3D shape estimation with changes in illumination

Franck Caniard*

Max Planck Institute for Biological Cybernetics

Roland W. Fleming†

Max Planck Institute for Biological Cybernetics

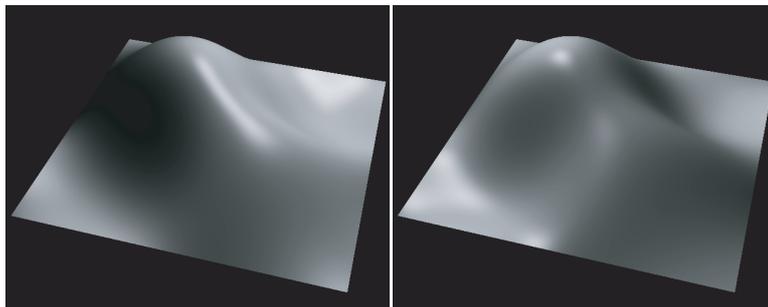


Figure 1: Both images are renderings of the same shape, using Phong shading. The elevation of the light sources is the same but the azimuth differs by 90° . The ridge of the large bump seems flat in the left image, but the right image shows that it is partially inclined.

Abstract

In many domains it is very important that observers form an accurate percept of 3-dimensional structure from 2-dimensional images of scenes or objects. This is particularly relevant for designers who need to make decisions concerning the refinement of novel objects that haven't been physically built yet. This study presents the results of two experiments whose goal was to test the effect of lighting direction on the shape perception of smooth surfaces using shading and lighting techniques commonly used in modeling and design software.

The first experiment consisted of a 2 alternate forced choice task which compared the effect of the amount of shape difference between smooth surfaces lit by a single point light with whether the position of the light sources were the same or different for each surface. Results show that, as the difference between the shapes decreased, participants were more and more biased towards choosing the match shape lit by the same source as the test shape.

In the second experiment, participants had to report the orientation at equivalent probe locations on pairs of smooth surfaces presented simultaneously, using gauge figures. The surfaces could either be the same or slightly different and the light source of each shape could either be at the same relative location or offset by 90° horizontally. Participants reported large differences in surface orientation when the lighting condition was different, even when the shapes were the same, confirming the first results.

Our findings show that lighting conditions can have a strong effect on 3-dimensional perception, and suggest that great care should be taken when projection systems are used for 3D visualisation where an accurate representation is required, either by carefully choosing lighting conditions or by using more realistic rendering techniques.

*e-mail:franck.caniard@tuebingen.mpg.de

†e-mail:roland.fleming@tuebingen.mpg.de

CR Categories: J.6 [Computer Applications]: Computer-aided engineering—Computer-aided design (CAD); I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture;

Keywords: Phong, Shape perception, Illumination, Shape from shading

1 Introduction

Computer graphics (CG) are almost always used in product design, where they are required to convey precise impressions of the 3D shape of a novel objects so that informed decisions can be made about how the model should be modified. However, designers often feel uncomfortable relying on this medium. For example, in the car industry final decisions are always made using costly full-scale physical models, as designers do not feel that the percept obtained from even high-end visualisation systems is accurate enough.

It is well known that many different cues contribute to the mental reconstruction of 3D objects from 2D images, including shape from shading, oclusions and silhouette [Todd and Reichel 1990; Koenderink and Andrea 1982; Koenderink 1984; Richards et al. 1987], textures [Kim et al. 2004] or specular reflections [Fleming et al. 2004]. Designers pay exquisite attention to producing the model itself, and in some applications, the paint characteristics are also simulated with painstaking accuracy, as both of these are features of the final product itself. By contrast, lighting is simply required for the visualization and rarely receives such detailed attention. For example, in most 3D modeling software, a simple point light source and Phong shading are used, although it is known that Phong is physically inaccurate and does not produce realistic shading [Koenderink and van Doorn 2003]. The light source position is either part of the software's default settings, or attached to the camera and never takes shape properties into account.

It is also well known that there are ambiguities caused by the interaction of lighting and shading, such as the bas-relief ambiguity [Belhumeur et al. 1997]. This work aims at studying the influence of lighting on perceived shape. Previous studies have shown contradictory results in terms of the effect of lighting and shading on shape perception. When a rich set of additional cues are present, including texture and binocular stereopsis, the direction of illumination has a negligible effect on the perception of shape [Todd

et al. 1997]. By contrast, [Koenderink et al. 1996], found systematic distortions of 3D shape when subjects judged the shape of real physical objects under different lighting conditions. Ferwarda and colleagues [Ferwerda et al. 2004] found that subjects were better able to identify subtle differences in the designs of car-like objects when global illumination methods were used to generate the images than with standard real-time shading.

This study aims at finding out whether different lighting conditions can yield different percepts of the same shape in CG. In two consecutive experiments, participants were presented sets of similar irregularly-shaped smooth surfaces lit by point light sources situated at different locations. In the first experiment, using a 2AFC design, they had to find which of two match shapes was identical to a test shape. In the next experiment, they had to report surface orientations using gauge figures [Koenderink et al. 1992] on pairs of shapes that could be different or the same, under same or different lighting conditions.

2 Experiment 1

2.1 Methods

The goal of this first experiment was to test whether lighting has an effect on shape perception.

We used a two alternate forced choice (2AFC) design. Participants were first presented an unknown shape for 3 seconds (test). It then disappeared and two match shapes appeared simultaneously. Participants had to report which of those shapes was the same as the first one. Only one shape was identical to the test shape (target). The other differed by various amounts (distractor). Participants were instructed to strictly pay attention to the shape, and time was not limited.

2.1.1 Random shape generation

The shapes were NURBS surfaces (Non Uniform Rational B-Splines). They were generated by first producing a 5×5 control points regular flat grid. Control points situated on the edges were repeated 4 times to keep the surface anchored to the edges of a square frame. The other control points were offset vertically to produce smooth bumps and valleys in the surface. The resulting NURBS was then triangulated, using a 80×80 regular grid, each square of the grid being split in two triangles, producing a very fine approximation of the NURBS surface made out of 12.800 triangles. The surface normals were accurately calculated at the vertex positions, then interpolated across the triangles on the graphics card to obtain a normal per pixel. The result is a smooth, deformed surface on which polygons are not apparent.

This study is only concerned with testing the influence of illumination on shape-from-shading. It has been shown that other cues have a strong effect on shape perception, notably the effect of occlusion and silhouette. For this reason, during the generation process only shapes that did not cause any occlusions, either with themselves or with the background. We let the generation program run until a total of 1000 surfaces were retained.

As a measure of shape differences we used the root mean square (RMS) value of the height of the surface at 100×100 regularly spaced locations, spreading over the whole object. The height was 0 at the edges. The absolute RMS value for each surface (i.e. difference between the surface and its projection on a flat plane) was the same for all generated shapes.

2.1.2 Setup

Graphics were produced in real time, using an Nvidia 7800GTX graphics card and a Sony GDM-F500R 21 inch CRT monitor, on a standard PC running Linux, using a resolution of 1280×1024 pixels, at a refresh rate of 85Hz.

The reference shape was drawn inside a centred viewport using the whole height of the screen, but only half of the width. When presenting the two possible answers, the screen was split into two vertical viewports of equal width. For a given trial, each of those three viewports was using its own virtual camera, with identical properties and situated at the same position relative to the corresponding object, so as to avoid effects of perspective and viewpoint. All backgrounds were black so that the viewports' boundaries were invisible. The NURBS surfaces were lit using a standard per-pixel Phong shader with a single white point light source. The material used was also white and cast specular reflections (plastic-like material). This was calculated in real time on the graphics card's processor.

Participants answered the task by pressing a key to the left end of the keyboard to indicate that the shape on the left was the same as the first one, and to the right end to indicate the shape on the right. Their face was situated 56cm away from the monitor and supported by a chin rest. The field of view of the virtual cameras was adjusted accordingly (29.75° vertically).

2.1.3 Conditions

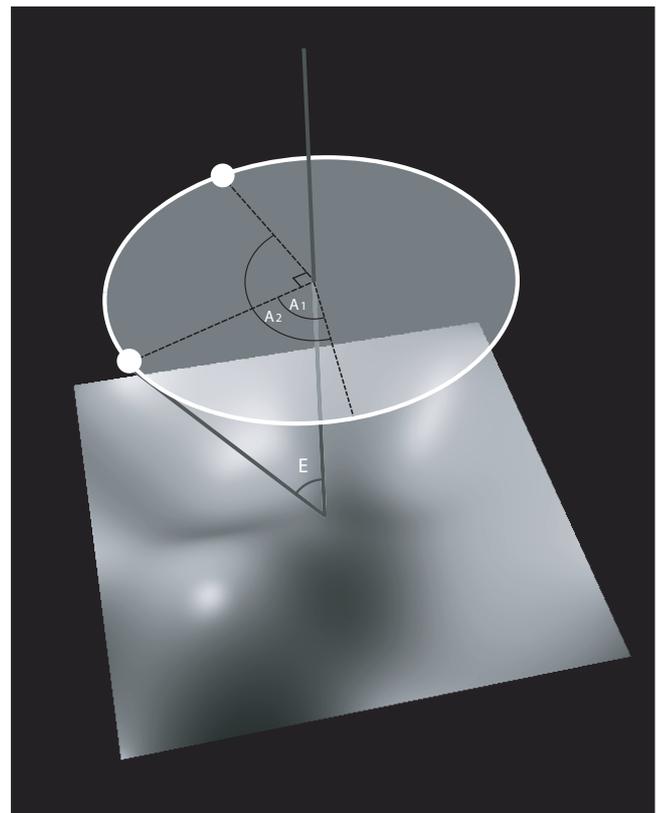


Figure 2: This shows how light sources were positioned. The two large white dots represent light sources used in a trial where the light positions were different for each shape. In both cases, the elevation (E) is the same, but the azimuths angles ($A1$ and $A2$) differ by 90° . The source lighting the above shape is the one that corresponds to an azimuth angle of $A1$.

The test shape was randomly picked from our set of 1000 NURBS. The same shape was never presented more than once. There was always one and only one right answer. The distractor was created by morphing the test shape with another randomly picked shape. The morphing was performed by linearly interpolating the position of the corresponding control points.

The main variable was the position of the light source. For each given trial, the elevation was the same for all three shapes (75.0° from the horizontal plane). The azimuth of the test shape's light source was randomly chosen in the range $[0^\circ 360^\circ]$. Then, one of the two matches' light source was attributed the same position as the test, and the other one was either the same, or offset by $\pm 90^\circ$ on the horizontal plane (see figure 2). The direction of the offset was randomly picked.

This yields three main conditions:

- The light is the same for all three shapes. Any difference in the images will only come from the shapes themselves. Participants' Value added performance should only depend on the amount of difference between the two surfaces.
- The target shape has the same lighting as the test shape ("light goes with the shape"), making the distractor appear more different than it actually is. Even in the case where the difference is very small, participants should be highly biased towards the target, leading to high performance in all cases.
- The distractor has the same lighting as the test ("light goes against the shape"). If lighting has an effect at all on shape perception, participants should be biased towards choosing the distractor, especially as the difference in shape becomes smaller.

In the first and second cases, the expected high performance should be reinforced by the fact that the rendering of the target is pixel for pixel identical to the one of the test shape.

The second important variable was the relative difference between the two matches. We quantified the difference by using the relative RMS value, calculated as described in the shape generation section. We used 7 evenly spaced RMS values. The range was arbitrarily chosen so that with the smallest value the difference was very hard to detect, and was fairly easy to see with the largest value (see figure 3).

The third parameter that we varied was the elevation of the camera relative to the horizontal plane. It could either be 90° (object seen from the top), or 60° , and was the same for all three objects in a given trial.

All parameters were combined in a factorial design, and each condition was repeated 18 times. This yields a total of 756 trials. Five evenly spaced short optional breaks were offered, during which the participant was shown their overall score, in an attempt to keep them concentrated and motivated. The score was their percentage of correct answers during the previous blocks. Each block consisted of 126 trials, and the score was calculated across conditions, so that they did not receive condition specific feedback. A complete session lasted for about one hour.

2.2 Results

We collected data from 15 naïve subjects (7 females, 8 males), all in their twenties, with normal or corrected to normal vision. They were all paid for their participation.

Figure 4 shows example data from two subjects. Namely, they exhibit the best and worst performance in the difficult condition where

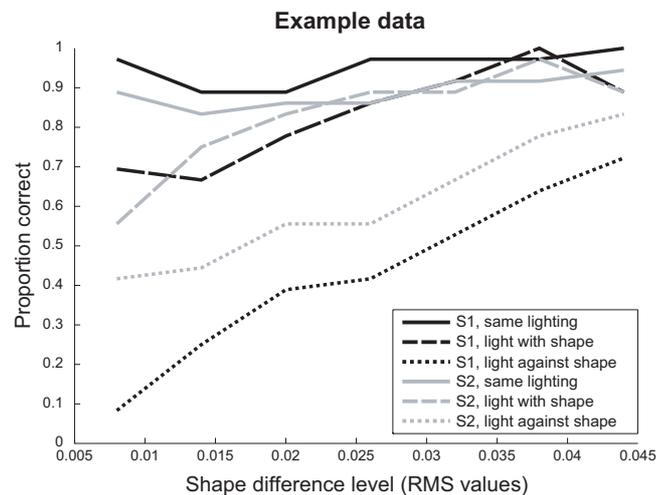


Figure 4: Performance of two participants, per condition. S1 exhibited the worst performance when the light went against the shape, and S2 the best. Both of them are well under chance level in that condition while performing well in the two others.

the light goes against the shape. They perform equally well in the two first conditions, but in the third one even the best performing subject chose the incorrect shape 58% of the time. This increased to 92% for the other subject.

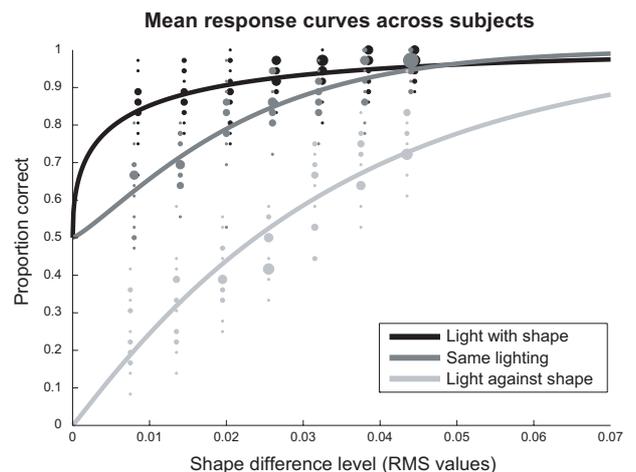


Figure 5: Global performance of all 15 subjects, per lighting condition. The diameter of the data points represents the number of samples. The data was fitted using a Weibull distribution. Note that in this graph, the columns of points have been slightly shifted for readability. However, in the experiment the same RMS values were used in each condition.

Figure 5 shows the performance across participants, split by lighting condition. The curves were fitted by a Weibull cumulative distribution function, using the Psignifit toolbox [Wichmann and Hill 2001b; Wichmann and Hill 2001a]. The horizontal axis represents the levels of difference between the presented shapes, from small to large, and the vertical axis the mean performance across repetitions.

These graphs show that participants show a marked effect of lighting conditions on perceived shape. The top most curve is the case where the light goes with the shape, biasing the answers towards

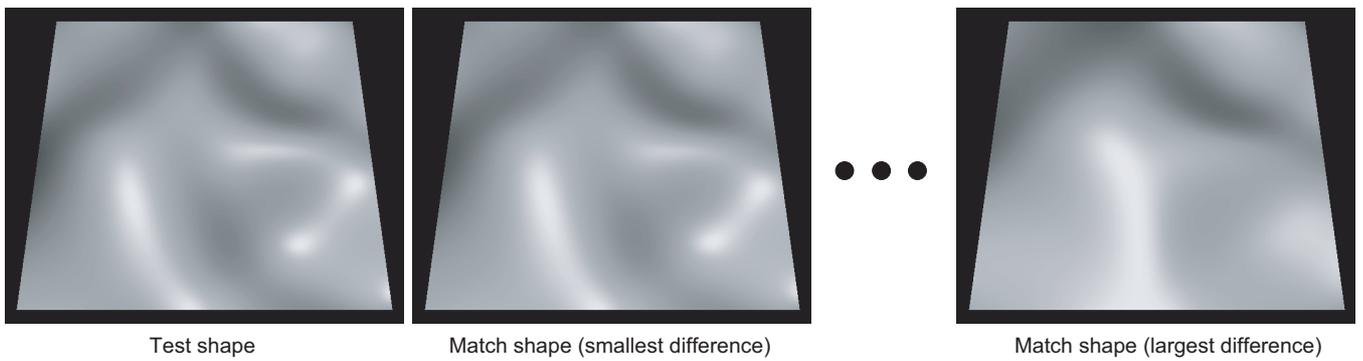


Figure 3: Example of a randomly generated NURBS surface, and variations based on it. The leftmost shape is the test. The two other shapes are morphs between the test shape and another random NURBS, showing the smallest and highest difference levels. A total of 7 evenly spaced levels, as established by our RMS metric, were used during the first experiment.

the target, even when the differences are hardly visible.

In the case where the lighting is the same for all three shapes, thus not having any influence on the relative shape perception, participants performed very well when the differences were large, but are close to chance level when the matches are almost identical.

In the third case, results show a very high bias towards choosing the distractor. When the differences were large, they tended to choose the target in the majority of cases, but when they were difficult to detect, they almost always choose the distractor (about 80% of the trials at the most difficult level).

No effect of the camera elevation angle was found (F-Test: $p=0.68$).

2.3 Conclusion

This first experiment shows that lighting condition seems to have a strong effect on shape perception. After being presented an unknown object, then having to find which of two presented objects had exactly the same shape, all of our 15 participants show a high bias towards choosing the match that was lit under the same condition as the target, even when the actual shape was different. This effect increases as the difference between shapes diminishes.

However, there are a number of limitations with the 2AFC method. First, it only allows us to measure the subject's ability to discriminate between stimuli, it does not provide any quantitative data about what shapes they perceived, or which features they relied on. Second, the task requires subjects to choose one response. It is possible that in some trials they saw neither or both shapes as matching the test object, and subjects would be unable to report such a percept. Finally, because the task emphasizes simple discrimination between the stimuli, there is no guarantee that subjects attended exclusively to the perceived 3D shape, although they were instructed to do so. It has to be noted that, in cases where the light went with the shape, the correct answer was pixel for pixel identical to the reference shape (condition 1 and 2), and it could be that subjects defaulted to a 2D pattern matching strategy, which would yield a similar pattern of results. To address these issues, we performed a second experiment with similar stimuli, in which subjects were required to report the absolute estimates of the 3D surface normals at a set of probe locations on the surface.

3 Experiment 2

The goal of this second experiment was to collect more qualitative data about the shape that the participants perceive under different

lighting conditions. We chose a task that could only be performed by forming a 3 dimensional mental representation of the objects and not by simply using 2D cues. We designed a gauge figure task, combined with a yes/no task, in which participants had to report whether they perceived two simultaneously presented smooth surfaces as being the same or different, then report locations where those differences were best perceived, and report the surface orientation of each surface at these locations. As a control condition, if they reported that the objects were the same they were required to report the surface orientations at location randomly chosen by the software.

3.1 Methods

The hardware setup and the algorithm used to generate the NURBS were identical to the first experiment. A set of 100 surfaces was created. A quick preselection was performed, during which 6 participants were asked to report, using a 3 point scale, whether the presented shape's 3D structure was easily perceived, hard to perceive or impossible to perceive. Each shape was presented 3 times, and the presentation order was randomised. Only the 20 shapes that obtained the highest score were retained for the experiment. This allowed to filter out surfaces whose shape was most difficult to comprehend. NURBS that were morphed with the base shapes to generate different targets were randomly picked from the remainder of the set.

3.1.1 Task

The task consisted in two distinct parts. During each trial, the participants first had to report whether they thought that the presented surfaces were the same or different objects. Before starting, they were instructed to exclusively pay attention to the shape, and were warned that the differences were likely to be very small. They could report their answer by clicking on either the "same" or "different" button situated at the bottom of the screen with the mouse cursor. The presentation time was not limited, the experimental software simply waited for the participant's choice. They then had to adjust 5 gauge figures on each surface, placed at the same relative probe locations. The pairs of figures were not shown simultaneously, but sequentially, so that participants were not able to adjust them relative to each other. Here again, there was no hard time limit but participants were advised not to spend more than about 5 seconds on each gauge figure in order to finish the session within a reasonable amount of time (less than 1.5 hour). If during the yes/no part their answer was that both surfaces were the same, the 5 locations were chosen by the software. But if they thought that the surfaces

were different, they were instructed to choose 5 pairs of locations where the orientation of the surfaces appeared to be different. The location could be selected by using two mouse cursors moved in parallel on top each shape. The adjustment of the gauge figures was also done using a standard computer mouse. None of the participants showed problems with the handling of the experimental software.

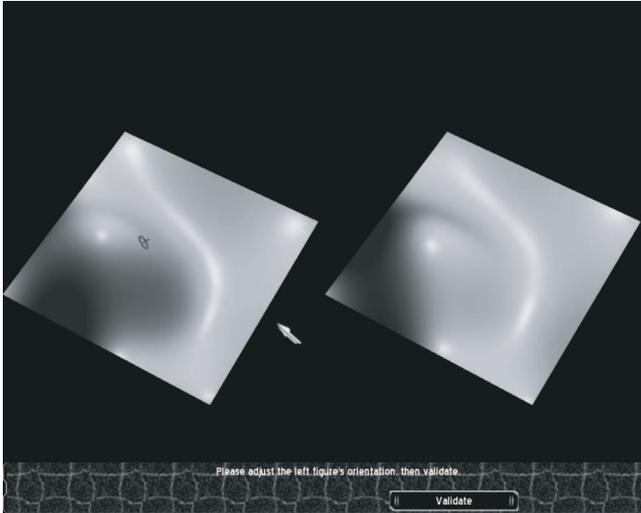


Figure 6: Screen capture of the experimental software while adjusting a gauge figure. Once the participant was happy with the adjustment, they had to click on the “validate” button, making the current gauge figure disappear and the next one appear at the same relative location on the surface to the right. Note that on this graph the thickness of the gauge figure’s lines have been increased for visibility. During the experiment they were exactly one pixel wide.

Two parameters were varied across trials. The first one is the light position which, as in the first experiment, could either be the same for both shapes, or offset by 90° on the horizontal plane (see figure 2). The direction of the offset was chosen randomly. The second parameter was whether the shapes were the same or differed by a small amount. Here we only used one level of difference (RMS = 0.024). The camera position was constant across trials. A 3/4 view was used (67° elevation, 30° azimuth, see figure 6) as we felt that it gave a better impression of depth to the scene. The parameters were combined in a factorial design, and each condition was repeated five times. The shapes were different for every trial to prevent participants from learning them. This yields a total of 20 trials, and 100 pairs of figures to adjust. All participants saw the exact same trials, but in randomised order.

Before hand, a separate program trained the participants to the handling of gauge figures and the task of reporting surface orientations. They had to adjust gauge figures on the same type of shapes, and received feedback after each adjustment. Once they felt comfortable with this part, they were taken through another step where they were shown the pairs of shapes used during the experiment, under the same lighting conditions. They simply had to report whether they saw them as same or different objects. Each pair was presented once and they received no feedback. The goal was to familiarise them with the fact that when they were present, differences were subtle.

Perceived angle difference between left and right surface

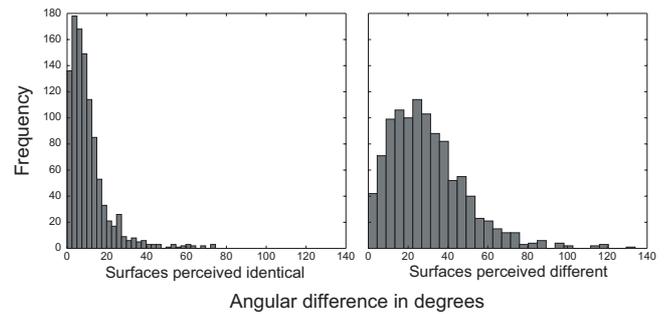


Figure 7: Histogram showing the angular difference between the gauge figures adjusted on the left NURBS and the gauge figures adjusted at the same location on the right NURBS, when the participants thought that the two objects were the same and when they thought that they were different.

3.2 Results

We collected data from 21 naïve participants (13 females, 8 males), all in their 20s, with normal or corrected to normal vision. They were all paid for their participation. They needed between 45 and 90 minutes to complete a full session.

Figure 7 shows the distribution of the angular difference between the reported orientation of corresponding gauge figures. When the shapes are seen as the same, angle values are small, producing a high peak near the origin. When the shapes are seen different, values are spread over a much wider range. The distributions are very skewed, mostly in the case where the shapes seem to be identical. This is due to the fact that the reported values were relatively small (the peaks are at about 3° on the left graph, and at about 22° on the right graph), but were bounded between 0° (perfect match) and 180° (largest possible difference that could be reported).

As the distributions are skewed, statistical tests were performed using the Friedman’s non parametric test for repeated measures, as it does not assume a Gaussian distribution. They showed a significant effect of the shape condition (shape is the same against shape is different), with $\chi^2=6.1$, $df=1$ and $p=0.0136$, which is not surprising, considering that this is what the participants were specifically asked to compare. More interestingly, they also revealed a very strong effect of lighting, with $\chi^2=6.1$, $df=1$ and $p=4.53 \times 10^{-9}$. This is clearly visible in figure 8, where the angular differences when the lighting condition is the same for both shapes is much smaller than when the light positions are offset. The small angles observed when the shapes are different and the lighting is the same can be explained by the fact that those differences were mostly very subtle. On the other hand, the 90° offset of the position of the light sources produced very large differences in terms of shading patterns on the surfaces, considerably influencing the mental reconstruction of the 3D structure of the object using only shading information.

3.3 Conclusion

The goal of this second experiment was to test whether the effect of lighting on shape perception seen in the first one was still present when participants were forced to build a 3D mental representation of the shape. This was done by asking them to report local surface orientations on pairs of similar or identical irregularly-shaped smooth surfaces under either identical or different lighting conditions. We could obtain a more quantitative measure of the strength of the effect as we could compare reported angles, although fully re-

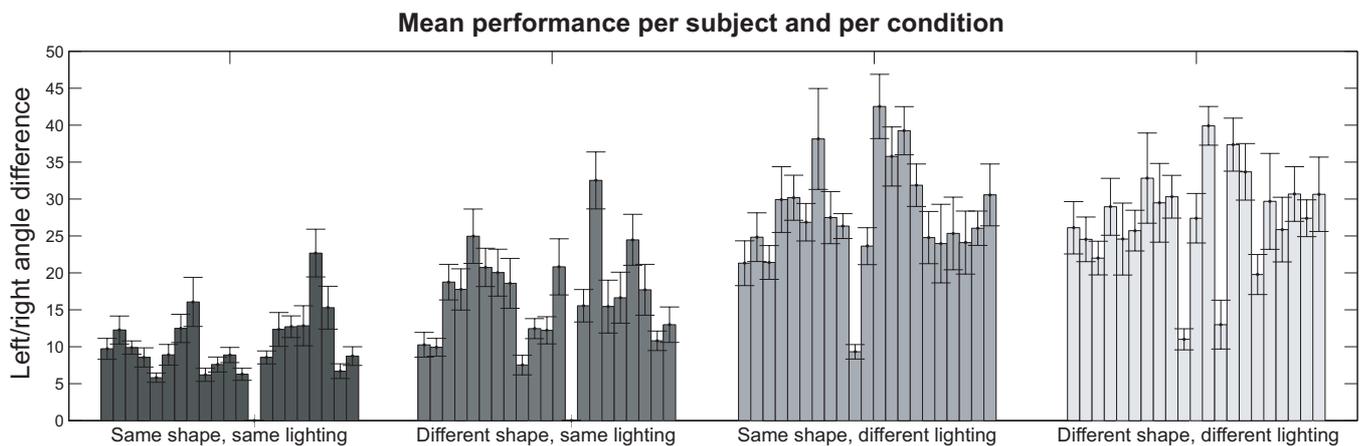


Figure 8: Mean angular difference between equivalent left and right gauge figures per condition and per participant. This shows that the effect of lighting is always present and strong among all subjects. When the light positions are the same, the effect of the shape condition is also clearly visible, but to a lesser degree.

constructing the perceived shape would require a much denser sampling of local orientations. Results clearly show a significant effect of lighting conditions. One possibility that we have not tested is that differences in perceived shape were caused by misestimation of the light source. If this were true it suggests that the apprehension of shape could be improved by embedding the object in a context that facilitates the accurate estimation of the light source.

4 Discussion

Together, our findings provide evidence for a substantial effect of lighting on the perception of 3D shape in simple computer graphics renditions. When smoothly curved objects are shaded with the Phong model under a single point source, the perceived shape is highly mutable depending on the lighting, and can differ considerably from the intended shape. This has important implications for the visualization of 3D data in industrial design, whenever precise perception of 3D shape is required. Surprisingly, Phong illumination is still probably the most commonly used shading model, both for modelling in CAD software, and for design evaluation. Our results may provide some insight into the problems experienced by designers using digital models of their designs.

How might 3D shape visualization be improved? This is a topic of active research. One obvious possibility would be to try to improve the physical realism of the renderings, for example, by employing global illumination rendering methods. Other authors [Ferwerda et al. 2004] have shown that subjects are better at perceiving subtle shape differences in car designs under global illumination than under standard local shading models. Realistic illumination has other benefits, for example, it is known to improve the perception of surface reflectance [Fleming et al. 2004; te Pas and Pont 2005].

However, there are a number of caveats to this approach. The first is that designers routinely require real-time rendering capabilities, and global illumination is still not possible at useful framerates for the mega-polygon models used in automobile design. The second caveat is that there is evidence that the dependency of perceived shape on the illumination conditions extends to real (as opposed to CG) objects [Koenderink et al. 1996]. Indeed, whether illumination is realistic or not, if the material has a substantial diffuse component, the resulting pattern of isophotes (and consequently the field of image orientations that could be measured by the front end of the visual system) varies with the illumination.

Together these suggest a couple of alternatives. The first is to identify a set of standard viewing conditions that are simultaneously (i) as realistic as possible and (ii) as unbiased as possible in terms of the afforded shape percept. In ongoing experiments, we are comparing shape estimation performance across a range of state-of-the-art realtime shaders, and across parametric variations in the illumination field. Our goal is to identify simple, canonical shading conditions that promote realistic appearance and stable shape perception. The second, perhaps more adventurous possibility is to seek non-photorealistic shaders that improve shape apprehension. The advantage of non-photorealistic methods is that they can be completely illumination free, and can be tailored to contain only features that are relevant to a particular judgment. This has been explored extensively in the field of abstract data visualization, most notably by Interrante and colleagues [Interrante et al. 1995; Interrante 1997; Interrante and Grosch 1997] who have shown several methods for depicting 3D data using textures. However, in product design, some sense of photographic realism is also desirable. This suggests the possibility of using sparse non-photo-realistic overlays to augment standard shaded images. Our future studies will also evaluate non-photorealistic shaders for 3D shape apprehension.

Acknowledgements

We are grateful to Dr Jörg Hahn, from Daimler-Chrysler in Stuttgart, for sharing his great knowledge about the process and issues involved in car design.

References

- BELHUMEUR, P., KRIEGMAN, D., AND YUILLE, A., 1997. The bas-relief ambiguity.
- FERWERDA, J. A., WESTIN, S. H., SMITH, R. C., AND PAWLICKI, R. 2004. Effects of rendering on shape perception in automobile design. In *APGV '04: Proceedings of the 1st Symposium on Applied perception in graphics and visualization*, ACM Press, New York, NY, USA, 107–114.
- FLEMING, R. W., TORRALBA, A., AND ADELSON, E. H. 2004. Specular reflections and the perception of shape. *Journal of Vision* 4(9), 798–820.

- GOOCH, A., GOOCH, B., SHIRLEY, P., AND COHEN, E. 1998. A non-photorealistic lighting model for automatic technical illustration. In *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, ACM Press, New York, NY, USA, 447–452.
- INTERRANTE, V., AND GROSCH, C. 1997. Strategies for effectively visualizing 3d flow with volume lic. In *VIS '97: Proceedings of the 8th conference on Visualization '97*, IEEE Computer Society Press, Los Alamitos, CA, USA, 421–ff.
- INTERRANTE, V., FUCHS, H., AND PIZER, S. 1995. Enhancing transparent skin surfaces with ridge and valley lines. In *VIS '95: Proceedings of the 6th conference on Visualization '95*, IEEE Computer Society, Washington, DC, USA, 52.
- INTERRANTE, V. 1997. Illustrating surface shape in volume data via principal direction-driven 3d line integral convolution. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 109–116.
- KIM, S., HAGH-SHENAS, H., AND INTERRANTE, V. 2004. Conveying three-dimensional shape with texture. In *APGV '04: Proceedings of the 1st Symposium on Applied perception in graphics and visualization*, ACM Press, New York, NY, USA, 119–122.
- KOENDERINK, J. J., AND ANDREA, J. D. 1982. The shape of smooth objects and the way contours end. *Perception 11*(2), 129–137.
- KOENDERINK, J. J., AND VAN DOORN, A. J. 2003. *The visual neurosciences*. MIT Press, Cambridge, ch. 72, 1090–1105.
- KOENDERINK, J. J., VAN DOORN, A. J., AND KAPPERS, A. M. 1992. Surface perception in pictures. *Percept Psychophys. 52*(5), 487–96.
- KOENDERINK, J. J., VAN DOORN, A. J., AND C. CHRISTOU, L. J. 1996. Perturbation study of shading in pictures. *Perception 25*(9), 1009–1026.
- KOENDERINK, J. J. 1984. What does the occluding contour tell us about solid shape? *Perception 13*, 321–330.
- RICHARDS, W. A., KOENDERINK, J. J., AND HOFFMAN, D. D. 1987. Inferring three-dimensional shapes from two-dimensional silhouettes. *Journal of the Optical Society of America A 4* (July), 1168–1175.
- TE PAS, S. F., AND PONT, S. C. 2005. A comparison of material and illumination discrimination performance for real rough, real smooth and computer generated smooth spheres. In *APGV '05: Proceedings of the 2nd symposium on Applied perception in graphics and visualization*, ACM Press, New York, NY, USA, 75–81.
- TODD, T. J., AND REICHEL, F. D. 1990. Visual perception of smoothly curved surfaces from double-projected contour patterns. *Journal of Experimental Psychology 16*(3), 665–674.
- TODD, J. T., NORMAN, J. F., KOENDERINK, J. J., AND KAPPERS, A. M. 1997. Effects of texture, illumination, and surface reflectance on stereoscopic shape perception. *Perception 26*(7), 807–22.
- WICHMANN, F. A., AND HILL, N. J. 2001. The psychometric function: I. fitting, sampling and goodness-of-fit. *Perception and Psychophysics 63*(8), 1293–1313.
- WICHMANN, F. A., AND HILL, N. J. 2001. The psychometric function: II. bootstrap-based confidence intervals and sampling. *Perception and Psychophysics 63*(8), 1314–1329.

