DOI: 10.1111/j.1467-8659.2009.01632.x

EUROGRAPHICS 2010 / T. Akenine-Möller and M. Zwicker
(Guest Editors)

Volume 29 (2010), Number 2

Printed Patterns for Enhanced Shape Perception of Papercraft Models

Su Xue Xuejin Chen Julie Dorsey Holly Rushmeier

Computer Graphics Group, Yale University

Abstract

Papercraft models can serve as inexpensive prototypes in shape design applications. However, in making the models some geometric detail is necessarily lost, and artificial creases may be visible, thereby limiting the utility of these models. To compensate for these practical limitations, we introduce the use of printed patterns on papercraft models to enhance the perception of the shape they are intended to represent. We propose pattern generation schemes that modulate the sizes, directions, and densities of glyphs of patterns based on geometric attributes. We present a psychophysical experiment designed to explore the effect that printed patterns have on the perception of the papercraft model shapes. We find that models with printed patterns are perceived to represent the intended shape more accurately, and, further, that the type of printed pattern has an impact on the perceived shape.

Categories and Subject Descriptors (according to ACM CCS): I.3.1 [Computer Graphics]: Hardware Architecture—Hardcopy devices; I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations

1. Introduction

The creative design of shape is fundamental in many applications such as industrial design and architecture. Although extensive software systems have been developed to assist this process, practitioners continue to use tangible 3D models to evaluate their designs. As a result, physical rapid prototyping devices have become standard equipment in many design studios. A 3D digital model is a powerful tool that allows inspection of a design in novel digitally defined environments, with different material compositions that are not possible to observe physically. However, having a physical object in hand affords examination of an object in ways that are impossible with a computer, even using advanced haptic devices. Indeed, visual simulations in the computer and physical hardcopy outside the computer are complementary in design evaluation.

Unfortunately, many physical rapid prototyping devices based on layered manufacturing remain relatively expensive and slow. Most models built are small because of material costs and the fragile nature of the parts that are produced. Service bureaus produce models, with the introduction of additional time delays and at a cost of hundreds or thousands of dollars depending on the part and process used. For early stages in the design process, lighter weight output methods are needed. A less expensive and greener (in terms of energy consumption) potential alternative for early in the design process is the development of papercraft models from digital 3D counterparts.

Papercraft models have received increased attention over the last five years in computer graphics. These methods are capable of generating approximations of freeform objects that are used in many products as well as in architectural ornamentation. Papercraft models have the problem though that, as they are generally made simple to be quick to assemble by folding and gluing ordinary paper, much geometric detail is lost, and some artificial creases are introduced in the folded model. These shortcomings severely limit the utility of papercraft models in shape design applications.

In this work we address these practical limitations, with an eye to expanding the utility and applicability of these models. We make the following contributions:

- the idea of using printed patterns to enhance the percep-

tion of the intended shape represented by papercraft models,
- the consideration and exploration of pattern generation schemes that modulate the sizes, directions, and densities of glyphs, based on attributes of the shape,
- a psychophysical experiment that explores the role of printed patterns on the perception of physical shapes, and
- findings indicating that models with printed patterns are perceived to represent the intended shape more accurately – and moreover, that the type of pattern has an impact on the perceived shape.

2. Previous Work

The developing technologies in computer graphics facilitate the process of designing, assembling and displaying products in virtual space. On the other hand, physical prototypes are simply present and continuous in physical space, while providing a feel of real 3D products. Physical prototyping, which provides a natural shared experience, still plays a critical role in the development of consumer products [Wri05, WDB*07, CGPS08]. Rapid prototyping devices have decreased in price and have become accessible to more designers. However, inexpensive devices (e.g. the $5000 Desktop Factory®) still produce models of limited size (maximum 5in x 5in x 5in volume) (`http://www.desktopfactory.com/our_product/`) and require a significant amount of time to build up a model.

A less expensive alternative for early in the design process is the development of papercraft models. The layout for cutting and folding to make a papercraft model can be generated from a 3D mesh by methods such as unfolding polygons [JKS05], identifying triangle stripes [MS04], or by fitting conics to form ruled surfaces [MGE08]. Curved folds were incorporated to approximate almost developable surfaces by discrete developable surfaces [KFC*08]. These methods are capable of generating a wide range of freeform shapes. However, papercraft models have the drawback that for many shapes they lack full geometric detail and include artifacts such as extra creases. We propose an approach to generate patterns on papercraft models to better represent the desired shape of objects.

In computer graphics, applying patterns to virtual objects as texture maps has been used to mask surface approximation [FSPG97], to add the appearance of greater geometric detail [RRP00], and to convey shape in data visualization [IFP96]. We explore if printed patterns on papercraft models generate the same benefits on physical objects as are obtained from patterns texture mapped on virtual objects.

The effects of visual masking by patterns of similar spatial frequency are demonstrated in [FSPG97]. A typical example is a cylinder that is represented as a set of flat facets for real time rendering. Without texture, the facets are clearly visible. However, with texture of an appropriate frequency, orientation and contrast, the boundaries between the facets become less noticeable, and the collection of flat facets appears more like a smooth cylinder. We test a variety of patterns on papercraft models, with the goal of masking the artificial creases introduced into the curved surfaces of these models.

Texture mapping is a classic technique in computer graphics to add the illusion of geometric detail. Different forms of texture maps have been studied to preserve shape appearance during simplification. In addition to static color textures on surfaces, bump maps [Bli78] and normal maps [COM98, HSRG07] are used to dynamically adjust surface appearance according to changes in lighting. Rushmeier et al. [RRP00] considered the perceptual trade-off between representing geometry directly and using shaded textures to represent geometric detail. They found that for models with high spatial frequency variation, the perception of model quality could be maintained despite large simplifications in geometry by using a texture map derived from the lit full-resolution model. However, this method was not as effective for models with low spatial frequency variations. In this paper, we seek patterns that are effective for communicating all types of shapes.

Interrante et al. investigated various textures to convey surface form in the context of data visualization. An example is applying sets of strokes as texture on a transparent surface representing a data isosurface enclosing an opaque isosurface [IFP97, Int97]. Principal curvatures and principal directions are used to generate the strokes. While we also use principal curvatures and principal directions to control patterns, other geometric attributes of intended shapes and the facet sizes of papercraft models are considered to mask folds as well as convey shape details.

In the study of human vision, it is well-known that surface patterns provide cues in human shape processing. Inferring shape information according to texture cues in images is also known as the *shape-from-texture* problem in the computer vision literature [AS88, For01]. Computer graphics gives us the capability of generating synthetic images with a variety of patterns on the same shape, with or without other shape cues. This has been used to investigate how textures affect 3D surface perception in works such as [IK01] and [TOKK04]. The research of [TOKK04] in particular shows that human perception of 3D shapes based on anisotropic and inhomogeneous texture is quite robust. Our work builds on this insight.

3. Framework

Our goal is to create papercraft models that communicate the shape of high-resolution digital models. Our approach for producing models is dictated by the constraints imposed by papercraft and printing on paper. In consideration of these constraints, we propose a pipeline in which the printed patterns used to enhance shape perception are formed from

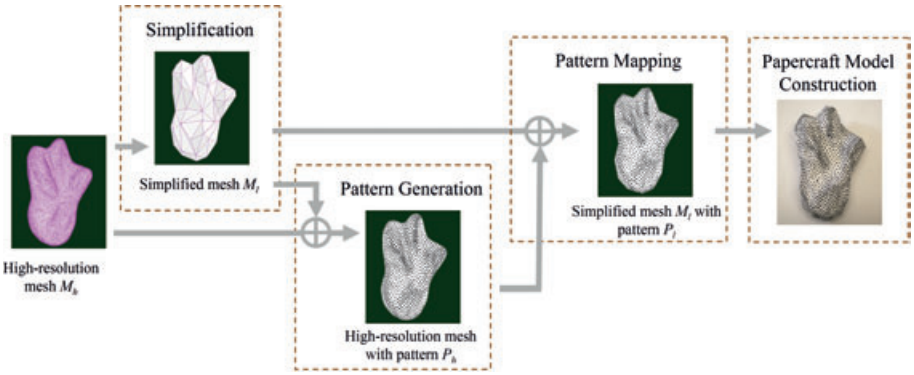

Figure 1: *The pipeline of making papercraft models with printed patterns.*

black and white glyphs. We describe here the constraints, and our basic pipeline. In Section 4 we will describe the specific methods we use to produce patterns of glyphs.

3.1. Papercraft Constraints

As noted in Section 2, numerous types of maps are used on simplified models to give the illusion of high-resolution detail. In creating papercraft models, some of these options that rely on dynamic adjustments in response to lighting changes – such as bump and normal maps (i.e. as in [Bli78, COM98]) – are obviously not possible. One option is to texture the model with detailed shading computed for the high- resolution model, as in [RRP00]. A computer display or a printed page are essentially very poor geometric representations of an object (since they are just simple planes), but they give the impression of a detailed 3D object using detailed shading.

We produced a number of preliminary test models, such as the front-lit bunny model shown in Fig. 2. Two major issues emerged from these tests: lighting conditions and the appearance of folds and seams. On a computer display, the light is controlled by virtue of being emitted by the display. On a printed page the illumination on a shaded figure is essentially uniform. Our papercraft model does not emit light, and the illumination varies across the model because of its shape. If we want to shade the model appropriately we need to "undo" the physical shading and replace it with a texture that gives the impression of light reflected from a high-resolution model. Since we have no control over the lighting conditions where the papercraft model is viewed, we can't undo the physical shading, and it competes with our computed shading. The result we found is that the texture we add to simulate the high-resolution shading often makes the papercraft look soiled.

As noted in [FSPG97], spatial artifacts are masked by features of similar spatial frequency. The artifacts we want to mask on papercraft models are high spatial frequency folds and seams. A texture that represents shading on a smooth surface, such as that shown in Fig. 2 often has relatively low

Figure 2: *A papercraft model of the simplified "Bunny" with a shading texture.*

spatial frequency. Further, slight mismatches in gluing together seams result in even greater contrast at the seams than if an untextured model were used, enhancing the visibility of artifacts we seek to hide.

Consideration of printing to emulate shading on the object leads to two observations: we need patterns with appropriate spatial frequency to mask the folds and seams with a certain spatial frequency, and further, we need high contrast patterns that do not compete with the physical light reflected from the model.

Pen-and-ink drawings, in just black and white, are frequently used to convey shape in printed documents. Other non-photorealistic rendering methods, which employ strokes and hatching, are also widely used to convey shades. Using long strokes on a papercraft object, however, can result in highly visible artifacts, particularly in places where the strokes are not perfectly aligned with the seams that must be glued together. More compact markings are desirable to avoid these artifacts.

These considerations suggest forming patterns composed of small black and white glyphs. Glyphs are used to convey orientation and other values on geometries in data visualization [War04]. As a preliminary exploration, we test three typical types of glyphs, which are *stroke*, *cross* and *dot*, as shown in Fig. 3. Stroke and cross glyphs can carry directional information, while dot glyphs provide stronger contrast. Our approach is to create patterns to print on models that are formed by sets of glyphs chosen to convey local geometric properties. We note that additional choices or

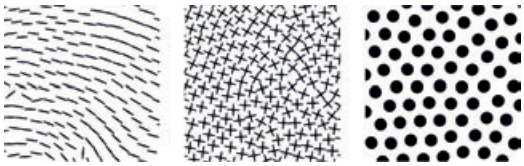

Figure 3: *Three types of glyphs used in our paper: stroke, cross and dot.*

combinations of glyphs are possible. In this initial study, our focus is on exploring the *role* of patterns, not their optimal arrangement or composition.

3.2. Basic Pipeline

Fig. 1 shows the basic pipeline of making papercraft models with printed patterns. A high-resolution mesh is treated as the starting point of our framework for creating a papercraft model, which is intended to represent the original shape. The mesh may be obtained by sampling a high-resolution representation (such as NURBs surface), or by scanning an existing physical object. Given a high-resolution mesh M_h, the whole pipeline of creating a papercraft model consists of four basic steps:

Simplification. We generate a simplified mesh M_l from M_h that contains a reasonable number (less than 200) of faces to facilitate the easy construction of papercraft models.

Pattern Generation. In this step, a pattern P_h is generated and stored in a texture atlas for the high-resolution mesh M_h. The generation scheme takes the facet size of M_l into consideration. One example is shown in Fig. 4 (a).

Pattern Mapping. The pattern P_l for the low-resolution mesh M_l is calculated by projecting the texture atlas P_h from M_h to M_l, as shown in Fig. 4 (b).

Papercraft Model Construction. Finally, a papercraft model is made from the low-resolution mesh M_l with pattern P_l.

As this paper focuses on exploring the effect patterns have on enhanced shape perception, we only consider variations in the second step "pattern generation," while holding the methods for the other steps constant. In the next section, we describe the details of our proposed pattern generation scheme.

4. Pattern Generation

To enhance the perception of the intended shape from the simplified model, we compute shape metrics based on the high-resolution mesh and then encode the shape metrics into a printed pattern. We use *principal curvatures*, *principal directions*, and the *shape index*, along with the facet size (i.e., the folding frequency) of M_l [FSPG97], as the geometric

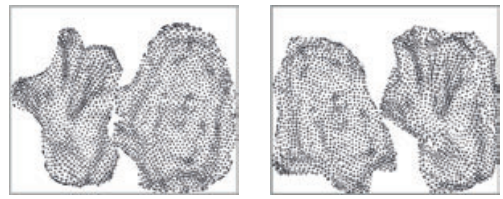

(a) (b)

Figure 4: *Pattern Mapping. (a) The pattern P_h stored in a texture atlas for the high-resolution mesh M_h. (b) The pattern P_l calculated by projecting P_h from M_h to the simplified mesh M_l.*

attributes that control the *sizes*, *directions* and *densities* of glyphs to form patterns.

As in the choice of glyphs, we do not claim that the pattern-generation schemes we have employed are in any sense optimal; clearly, there are many additional dimensions to explore.

4.1. Shape Metrics

Breton et al. [BILZ92] showed a pair of shape metrics including a curveness measure c and a shape index s,

$$c = max(|\kappa_1|, |\kappa_2|) \quad (1)$$

$$s = cos^{-1}\left(\frac{\kappa_1 + \kappa_2}{2c}\right), \quad (2)$$

where κ_1, κ_2 are the two signed principal curvatures. While the shape index reveals the category of shape, curveness serves to represent the scale of shape. Thus, this pair of shape metrics in essence provides a parametric form of shape representation by two coordinates. We adopt c and $S = cos(s)$, that is,

$$S = \frac{\kappa_1 + \kappa_2}{2c} \quad (3)$$

as the shape metrics for generating patterns based on the original shape. Additionally, we consider the maximal principal directions in generating the directional patterns of strokes and crosses.

Given a high-resolution mesh M_h of the model, we first normalize its size to be bounded in $[(-1,-1,-1),(1,1,1)]$. Principal curvatures and directions of vertices are computed using the methods in [MDSB03]. The shape metrics c and S for each vertex are then derived from Eqs. (1) and (3). For an arbitrary position p on the mesh, the shape metrics are interpolated from the metrics of three vertices of the triangle where p is located, based on the barycentric coordinates of p on this triangle.

4.2. Encoding Shape Metrics

We arrange glyphs centered at a set of sampled 3D positions $\{p_i\}$ over the surface of high-resolution mesh M_h. The critical attributes of glyphs, say, sizes, directions and densities,

are controlled according to the described shape metrics of M_h and the facet size of the simplified mesh M_l.

Glyph Size Following [Int97, IFP97, IK01], it is natural to control the size r of the glyph centered at position p according to its maximal principal curvature $c(p)$. We define

$$r(p) = a_1 c(p) + a_2, \tag{4}$$

where a_1 and a_2 are parameters determined by user adjustments. For the size-normalized M_h, r ranges from 0 to 0.2 in our system.

Glyph Direction When using maximal principal directions to align the directions of glyphs, we note that computing principal directions on meshes inherently suffers from the numerical problem due to the triangulation and ambiguities at umbilical points and flat areas, as reported in [Gol01, MDSB03, AT05]. As shown in Fig. 5 (a), where each red line indicates the principal direction of a vertex on the high-resolution mesh M_h, the intersections of principal directions interfere with the perception of the shape.

Instead, we compute a more consistent vector field over M_h for aligning the directions of glyphs. This vector field is initialized by the principal directions that simultaneously satisfy two conditions. First, their corresponding principal curvatures c satisfy

$$c > c^*_{min} + T_c(c^*_{max} - c^*_{min}),$$

where $T_c \in [0,1]$ is a scale threshold to select reliable principal curvatures. c^*_{min} and c^*_{max} are defined as

$$\begin{aligned} c^*_{min} &= c_{min} + 0.05(c_{max} - c_{min}) \\ c^*_{max} &= c_{min} + 0.95(c_{max} - c_{min}), \end{aligned}$$

where c_{min} and c_{max} are the minimal and maximal c over all vertices of M_h respectively. All raw principal curvatures of the vertices in M_h used in the pattern generation are clamped to $[c^*_{min}, c^*_{max}]$ which avoids extremum outliers for computational robustness.

Second, the principal directions for initializing the vector field are computed on close-to-parabolic areas. As suggested in [HZ00], we compute $\alpha = \frac{max\{|\kappa_1|,|\kappa_2|\}}{min\{|\kappa_1|,|\kappa_2|\}}$ for each vertex of high-resolution model M_h. The positions where

$$\alpha > \alpha^*_{min} + T_\alpha(\alpha^*_{max} - \alpha^*_{min})$$

are selected as close-to-parabolic areas. $T_\alpha \in [0,1]$ is another threshold that can be adjusted for generating appropriate vector fields. α^*_{min} and α^*_{max} are the clamped version of extremes of α over all vertices of M_h, which are computed in the same way we obtain c^*_{min} and c^*_{max} by 0.05 and 0.95 ratios respectively.

The initial vectors at the other vertices that are not regarded as reliable are set to $\mathbf{0}$. Through a series of averaging operations, that is, replacing each vector with the average of the vectors within its neighborhood, a more consistent vector field is obtained. The illustration of a final vector field with $T_\alpha = 0.2$ and $T_c = 0.5$ that is used to align the directions of glyphs is shown in Fig. 5 (b).

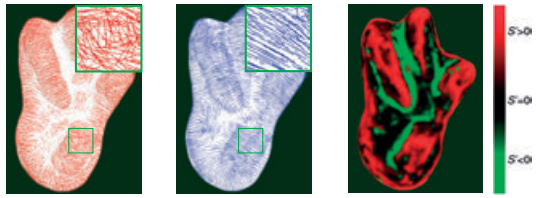

(a) (b) (c)

Figure 5: *Shape metrics used to control patterns. (a) The principal direction at each vertex on the high-resolution model; (b) The vector field generated based on reliable principal directions; (c) The distribution of S^*, the metric used to control glyph densities.*

Glyph Density The glyph's *overall density* is adjusted to better mask folds in the simplified mesh M_l, as suggested by [FSPG97]. We control the *local densities* of glyphs according to the shape index S, which reveals different categories of shapes. Further, in order to differentiate the same shapes at different scales, we use the unnormalized version of shape index, $S^* = S \times c = \frac{\kappa_1+\kappa_2}{2}$, which turns out to equal signed mean curvature. We show an example of the distribution of S^* on a high-resolution mesh in Fig. 5 (c).

To adaptively adjust the densities of glyphs according to S^*, we introduce a *repellant sphere* for every position p, which is defined as a sphere centered at p with radius

$$r_e(p) = b_1 S^*(p) + b_2, \tag{5}$$

where b_1, b_2 are model-specific parameters that are manually adjustable to achieve the best fold-masking effect according to the adequate overall folding frequency of M_l. Similar to the glyph size, the radius r_e is limited to $[0, 0.2]$ for a normalized mesh. To obtain the set of sampled positions, $\{p_i\}$, at which to place glyphs, all positions over M_h (the number of possible positions equals the resolution of P_h, which is 2048×2048 in our system) are set as available candidates, and then one position is randomly chosen as the starting point p_0. To find the next sample position, we check all available positions to find the first one that is not located within any repellant spheres of already sampled positions. All checked positions are then set as "unavailable." This sampling procedure continues until no available positions exist, and then the glyphs with modulated sizes are placed at these sample positions. With this scheme, we adaptively control the varying densities of glyphs by changing radii of the corresponding repellant spheres based on the shape indices.

Additionally, we adopt a special treatment for dot patterns to prevent the overlap of dots by adaptively shrinking dots for which an overlap is detected. The overlapping dots are shrunk to their repellent radiuses, which essentially forces dots to keep far enough away from each other.

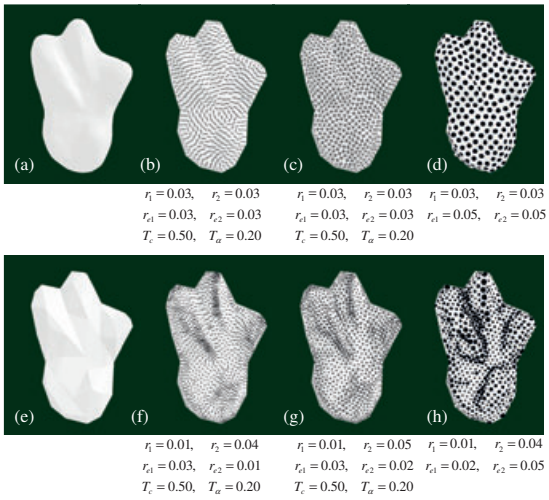

Figure 6: *A series of printed patterns on the simplified mesh generated based on the geometric attributes of the high-resolution mesh and the facet size of the simplified mesh. They are (a) the high-resolution mesh with 20,000 faces; (e) the blank simplified mesh with 150 faces; the simplified mesh with (b) uniform stroke pattern; (c) uniform cross pattern; (d) uniform dot pattern; (f) modulated stroke pattern; (g) modulated cross pattern and (h) modulated dot pattern.*

Parameter Setting To generate appropriate patterns for shape perception, the parameters $a_1, a_2, b_1, b_2, T_c, T_\alpha$ are adjusted through slider bars. The parameters can be adjusted in pairs, rather than all at once. Consider the model shown in Fig. 1 as an example. Given a normalized high-resolution mesh M_h, first T_c and T_α are adjusted to obtain a reliable vector field (Fig. 5 (b)). Instead of adjusting a_1, a_2, it is more intuitive to adjust the glyph sizes r_1, r_2 that correspond to $c = c^*_{min}$ and $c = c^*_{max}$ respectively. Given a pair of r_1, r_2, the parameters a_1, a_2 can be solved via Equation 4. Similarly, instead of adjusting b_1, b_2, we adjust the radii of repellant spheres, r_{e1}, r_{e2}, that correspond to $S = S^*_{min}$ and $S = S^*_{max}$, respectively. In order to get an appropriate overall glyph frequency for visual masking, we first adjust r_1, r_2 and r_{e1}, r_{e2} while keeping $r_1 = r_2$ and $r_{e1} = r_{e2}$. This generates a uniform pattern (Fig. 6 (b)). Finally, r_1, r_2, r_{e1}, r_{e2} are manually adjusted to generate patterns on the simplified mesh (Fig. 6 (f)). For convenience, r_1 and r_2 can be adjusted first to obtain appropriate glyph sizes, while keeping r_{e1} and r_{e2} constant; then r_{e1} and r_{e2} are slightly tuned to modulate glyph densities.

More patterns generated using our pattern generation scheme are shown in Fig. 6. Note that this pattern generation method is not limited to papercraft models; it could be also used with any of the techniques described in [JKS05, MS04, MGE08, KFC*08].

5. Psychophysical Experiment

We conducted a psychophysical experiment to examine the effectiveness of the various printed patterns described in Section 4. In the experiment, participants compared a physical shape formed from clay to various versions of papercraft models made from a digital model of the shape obtained by laser scanning. A *Two-Alternative Forced Choices* design was used, with the participant asked to choose which of two papercraft models had the shape most similar to a reference clay model. Our goal was to test the following hypotheses:

- **Hypo1**: *Patterns printed on a papercraft model can improve the perception of the model as being the shape of the high-resolution model it is intended to represent.*
- **Hypo2**: *Patterns of glyphs with size and density modulated by shape metrics are better for producing the perception of the intended shape.*
- **Hypo3**: *The enhancement of shape perception depends on the type of glyph used to form the pattern.*

5.1. Stimuli

As stimuli, we made five clay shapes of varying size, aspect ratio and geometric complexity. Two of these shapes were deliberately made very simple for use in training in the experimental tasks (Fig. 8), and three were used for data collection (Fig. 7). To focus on shape comparisons, all of the shapes were deliberately made to be abstract, rather than shapes that resemble easily named objects. Our goal was to ensure that participants would make comparisons to these reference shapes only, and not to mental models that they might previously have of named shapes.

We captured digital models of the clay shapes using a NextEngine® scanner to obtain the raw meshes (about 500K faces). We repaired and slightly smoothed the raw meshes using MeshLab v1.2.2b(*http://meshlab.sourceforge.net/*). The spatial sampling of the raw meshes was much higher than needed to represent the geometric features of the object. To facilitate the interactive adjustment of parameters described in Section 4.2, the oversampling was greatly reduced by using Garland's qslim software [GH97] to generate models of 20,000 faces, which were then used as the high-resolution mesh M_h for each shape. The qslim software was also used to produce the low-resolution meshes M_l, reducing to 100 faces for training shapes and 150 faces for the other three shapes. The meshes were parameterized and texture atlases were generated using the method described in [ZSGS04].

Given the corresponding high-resolution and low-resolution meshes, we followed the pipeline in Section 3 to generate the designed pattern maps and applied them to the low-resolution meshes. The high-resolution meshes were loaded into the interactive system described in Section 4.2 to specify the parameters for generating uniform and modulated versions of patterns using the stroke, cross and dot

glyphs. The pattern generation required one to two minutes of processing for each mesh. The generated patterns were stored in the texture atlases for the high-resolution meshes. These patterns were then projected onto the texture atlases for the low-resolution meshes using the normal shooting technique described by Sander et al. [SGG*00]. The projection of the patterns onto the lower-resolution meshes required two to three minutes per mesh.

Pepakura Designer 3® was used to generate the flat, unfolded low-resolution meshes with patterns. The flattened meshes were printed on ordinary paper, which was cut and folded into the papercraft models.

For the three clay shapes used in data collection, seven papercraft models were prepared. The seven patterns used for comparison were: *blank* (no patterns), *uniform stroke*, *modulated stroke*, *uniform cross*, *modulated cross*, *uniform dot* and *modulated dot*. In the rest of this work, we will refer to these patterns respectively as *blank*, *uni-stroke*, *modu-stroke*, *uni-cross*, *modu-cross*, *uni-dot* and *modu-dot*. For the clay shapes used in the training task, just three simplified papercraft models were prepared as shown in Fig. 8.

The experimental setup is shown in Fig. 9. The center "Reference" is always one of the clay shapes. The shapes placed in positions A and B are papercraft representations of the center reference shape. The table surface the shapes were placed on was 28 inches from the floor, and covered with white paper. The setup was located in an office, equipped with typical indirect office lighting. The A, B and Reference shapes were set on the table with the same orientation relative to the sides of the table.

5.2. Procedure

We began the experiment with each participant by reading an instruction sheet to them explaining the comparison task. Participants were instructed to observe the three shapes, and to indicate which shape, A or B, most closely resembled the Reference shape. Participants were allowed to sit or stand, and to move to observe from different angles, as long as they stayed behind a blue line marked on the floor that was 25 inches in front of the table. They were not allowed to touch the shapes. These precautions were taken to ensure that the papercraft models were not deformed in the course of the experiment. After the participant indicated either A or B for each comparison, a new comparison set was presented.

Before gathering data, each participant went through a training experiment to ensure that they understood the task. The training experiment contained eight comparison sets in which the first six used shapes shown in Fig. 8 and the last two introduced shapes shown in Fig. 7. Participants were allowed to ask questions during the training, which served to help them understand and practice what they learned from the instructions. The answers collected from the training were recorded but not used in data analysis.

Figure 7: *Three groups of models used for data gathering. Each group contains a clay model as reference (top-left), a blank papercraft model and six papercraft models with different patterns.*

The formal data gathering experiment followed the same procedure as the training, except that only the shapes shown in Fig. 7 were used. Every possible pair of papercraft models with different patterns, accompanied by their reference model, was presented exactly once. Thus in the formal experiment a participant had $3\binom{7}{2} = 63$ comparison sets to finish. The order of tests was randomized, and was different for each participant to avoid possible order-dependent effects. The time required for each participant to complete the experiment was about 30 minutes.

There were 22 participants, all over eighteen years of age. Both male and female participants were included, and none had any computer graphics background. Participation in the experiment was voluntary and unpaid. By the restrictions of

Figure 8: *Two groups of models including clay models and papercraft models with different patterns used in the training session of our psychophysical experiment.*

Figure 9: *A comparison set is placed on the test table in a typical office environment in our psychophysical experiment. The middle is a clay model used as a reference model. The other two papercraft models with different patterns are labeled A and B, respectively.*

the human subjects review board overseeing our work, the participants were informed of their right to end their participation at any time, were allowed to remain anonymous and were not asked for any additional demographic information. We found that we needed three investigators to conduct this experiment to reduce the time required for each participant to a reasonable level. One investigator prepared the comparison sets and other two were in charge of delivering them to the test position and presenting them.

After the formal data gathering was finished, we had a voluntary informal debriefing with participants. We explained the purpose of the experiment and collected comments and suggestions from the participants.

6. Results

We analyzed the results of the psychophysical experiment to test our hypotheses.

6.1. Statistical Tests

Assuming a normal distribution, we performed a series of t-tests at level $\alpha = 0.05$ to compare each pair of patterns to discover statistical significance in the responses collected in the experiment.

The tests were performed across all participants and all models. From the results, significant advantages ($p^* < 0.05$) were observed when comparing the effect on enhancing shape perception of the following pairs of patterns:

- *uni-stroke*: >*blank*
- *modu-stroke*: >*blank*, >*uni-stroke*, >*uni-cross*, >*uni-dot*
- *uni-cross*: >*blank*
- *modu-cross*: >*blank*, >*uni-stroke*, >*uni-cross*, >*uni-dot*
- *uni-dot*: >*blank*
- *modu-dot*: >*blank*, >*uni-stroke*, >*uni-cross*, >*uni-dot*

Furthermore, to quantitatively show the performance of different patterns, we applied Thurstone's Law of Comparative Judgements, Case V, to derive interval scaling values through paired comparisons, where subject responses were assumed normally distributed. The scale values of the seven patterns over three models are shown in Fig. 10. We can see that for each model the modulated patterns outperform the unmodulated ones. The *blank* (no-pattern) has the poorest performance for the first two models, and has roughly equal performance (the poorest) with the *uni-stroke* pattern for the last model. The enhancement of shape perception also relies on the type of glyphs used to form patterns. For the first model, modulated cross and dot patterns better convey the shape of the reference model than modulated line patterns. In contrast, modulated line patterns are better than the other two modulated patterns at representing the shape of the second model. There is no significant difference between the three types for the third model.

We computed the relative scaling values of the seven patterns with respect to 22 participants across all models. The result is shown in Fig. 11. We can see variability in responses of different participants. Participant 5 responded that *blank* models had a shape more similar to the reference clay models than any of the papercraft models with patterns. For participant 22, modulated patterns better presented the shape of the reference model than the blank models, while uniform patterns were worse than blanks. We note that the scales of seven patterns for each participant are computed based on only three models. The insufficient number of comparisons per participant does not yield statistically convincing results for individuals. However, we can still observe that generally the results for *blank* are poorest, and the modulated patterns outperform the unmodulated ones for most participants.

6.2. Discussion

We consider our three hypotheses in turn. With respect to **Hypo1**, we did find that considering the results across all models and all participants, the models with patterns (both the uniform and modulated) outperformed the blank models. Considering the results for the individual clay shapes, only the patterned models for the shape shown farthest to the left on Fig. 10 always clearly outperformed the blank model. For the other two shapes, the uniformly patterned are

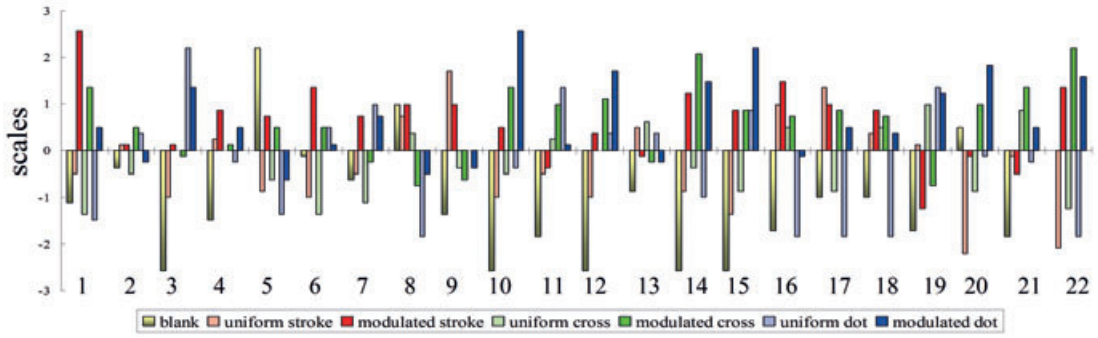

Figure 11: *Paired comparison scaling values of the seven patterns over 22 participants.*

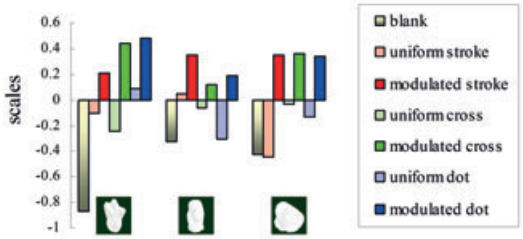

Figure 10: *Paired comparison scaling values of the seven patterns over three models.*

not clearly superior to the blank model. Therefore, while the experimental results indicate that printed patterns can make a difference, we can not conclude that simply any type of pattern will enhance the perception of the paper model.

With respect to **Hypo2**, we found that considering the results across all models and all participants, the models with patterns modulated by shape metrics outperformed models with spatially uniform patterns. This appears to be true for each individual model as well. While we can not claim that our modulated patterns are in any way optimal, this indicates that future research should continue the consideration of modulated patterns.

For **Hypo3**, the experimental results were the most surprising. We did not find significant differences between the results for different glyphs. This was true both for the spatially uniform case, and for the modulated case. We had anticipated that the stroke and cross glyphs would be clearly superior because they are able to carry more information about the high-resolution shape. This could be due to factors such as the way the various glyphs were printed. We can observe, though, that simply indicating directional information about curvature using a glyph does not necessarily result in superior perception of the shape we intended to represent with the paper model.

Considering the results for **Hypo3** in more detail, we observe that the success of enhanced shape perception by different types of glyphs varied with the reference clay shapes, Fig. 10. This raises the possibility that different aspects of shape perception are being balanced. Some participants mentioned in the debriefing that it was usually hard to compare the modulated (or uniform) patterns formed by different glyphs because the dot patterns convey a nice smooth perception, while directional patterns (stroke and cross) work better in conveying details and depth information. We speculate that the dot pattern conveys smoothness because the strong contrast between black dots and white background distracts the observers' attention from noticing the contrast between model boundaries and the background, and from noticing the creases between faces. We conjecture that the best pattern should be determined by balancing the requirements of rendering *smoothness*, *depth* and *details* for individual shapes by differing combinations of glyphs. The possibility that high-contrast dot patterns are effective for masking and directional patterns represent details well is a useful guide for us in designing even more effective patterns.

During our experiments, participants also provided many other valuable comments. Two of the participants pointed out that small imperfections in a model had some influence on their perception about the overall shape. Although we endeavored to make the physical models to be of the same quality, not all papercraft models of the same shape are absolutely identical. Slight differences still occur at some places, such as the seams where two faces are glued together.

From the analysis of the psychophysical experiment, our work opens up many interesting research problems in addition to exploring new patterns. We would like to investigate whether similar results would be obtained if the experiment were performed on synthesized images rather than real physical objects. Additionally, besides the smooth shapes used in our experiment, we would also like to investigate the effect of patterns on perception enhancement for shapes with finer details and sharp edges. And we would like to study the effectiveness of patterned papercraft models for some specific design tasks.

7. Conclusion

In this paper, we introduced the idea of using printed patterns on simplified papercraft models to enhance the in-

tended shape perception. In order to compensate for the loss of geometric details and to mask visible artificial creases caused by simplification, patterns are generated according to the geometric attributes of high-resolution models and printed on simplified papercraft models. With the proposed pattern generation scheme, the sizes, directions, and densities of pattern glyphs are modulated by principal curvatures, principal directions and shape indices respectively, with the facet size of the simplified model taken into consideration. We conducted a psychophysical experiment to explore the effect that printed patterns have on the enhanced shape perception for simplified papercraft models. The experiment results demonstrate that the models with printed patterns are perceived to represent the intended shape more accurately. Furthermore, the type of printed pattern (i.e., modulated or uniform) has an impact on the accuracy of shape perception. Our hope is that this work offers insight into human vision and will ultimately extend the usefulness of papercraft models in shape design.

Acknowledgements

The authors thank the anonymous reviewers for their comments, Steven Zucker for numerous stimulating discussions, Michelle Modest for help designing and conducting the experiment, Patrick Paczkowski for editorial suggestions, and the entire Yale Graphics Group for their support.

References

[AS88] ALOIMONOS Y., SWAIN M.: Shape from texture. *Biological Cybernetics 58*, 5 (1988), 345–360.

[AT05] AGAM G., TANG X.: Accurate principal directions estimation in discrete surfaces. In *Proceedings of the Fifth International Conference on 3-D Digital Imaging and Modeling* (2005), IEEE Computer Society, pp. 293–300.

[BILZ92] BRETON P., IVERSON L. A., LANGER M. S., ZUCKER S. W.: Shading flows and scene bundles: A new approach to shape from shading. In *ECCV* (1992), vol. 588, pp. 135–150.

[Bli78] BLINN J. F.: Simulation of wrinkled surfaces. *Proc. SIGGRAPH* (1978), 286–292.

[CGPS08] CIGNONI P., GOBBETTI E., PINTUS R., SCOPIGNO R.: Color enhancement for rapid prototyping. In *The 9th International Symposium on VAST International Symposium on Virtual Reality, Archaeology and Cultural Heritage* (2008), pp. 9–16.

[COM98] COHEN J., OLANO M., MANOCHA D.: Appearance-preserving simplification. In *Proc. SIGGRAPH* (1998), pp. 115–122.

[For01] FORSYTH D.: Shape from texture and integrability. In *Eighth IEEE International Conference on Computer Vision* (2001), vol. 2, pp. 447–452.

[FSPG97] FERWERDA J. A., SHIRLEY P., PATTANAIK S. N., GREENBERG D. P.: A model of visual masking for computer graphics. In *Proc. SIGGRAPH* (1997), pp. 143–152.

[GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. In *Proc. SIGGRAPH* (1997), pp. 209–216.

[Gol01] GOLDFEATHER J.: Understanding errors in approximating principal direction vectors, 2001. Technical Report 01-006, University of Minnesota, Computer Science and Engineering.

[HSRG07] HAN C., SUN B., RAMAMOORTHI R., GRINSPUN E.: Frequency domain normal map filtering. *ACM Trans. Graph. 26*, 3 (2007), 28.

[HZ00] HERTZMANN A., ZORIN D.: Illustrating smooth surfaces. In *Proc. SIGGRAPH* (2000), pp. 517–526.

[IFP96] INTERRANTE V., FUCHS H., PIZER S.: Illustrating transparent surfaces with curvature-directed strokes. In *Proceedings of the 7th conference on Visualization* (1996), pp. 211–218.

[IFP97] INTERRANTE V., FUCHS H., PIZER S. M.: Conveying the 3D shape of smoothly curving transparent surfaces via texture. *IEEE Transactions on Visualization and Computer Graphics 3* (1997), 98–117.

[IK01] INTERRANTE V., KIM S.: Investigating the effect of texture orientation on the perception of 3D shape. In *Proceedings of Human vision and electronic imaging* (2001), vol. 4299, pp. 330–339.

[Int97] INTERRANTE V.: Illustrating surface shape in volume data via principal direction-driven 3D line integral convolution. In *Proc. SIGGRAPH* (1997), pp. 109–116.

[JKS05] JULIUS D., KRAEVOY V., SHEFFER A.: D-charts: Quasi-developable mesh segmentation. *Computer Graphics Forum, Proceedings of Eurographics 24*, 3 (2005), 581–590.

[KFC*08] KILIAN M., FLÖRY S., CHEN Z., MITRA N. J., SHEFFER A., POTTMANN H.: Curved folding. In *ACM Trans. Graph.* (2008), pp. 1–9.

[MDSB03] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete differential-geometry operators for triangulated 2-manifolds. In *Visualization and Mathematics III* (2003), pp. 35–57.

[MGE08] MASSARWI F., GOTSMAN C., ELBER G.: Paper-craft from 3D polygonal models using generalized cylinders. *Comput. Aided Geom. Des. 25*, 8 (2008), 576–591.

[MS04] MITANI J., SUZUKI H.: Making papercraft toys from meshes using strip-based approximate unfolding. *ACM Trans. Graph. 23*, 3 (2004), 259–263.

[RRP00] RUSHMEIER H. E., ROGOWITZ B. E., PIATKO C.: Perceptual issues in substituting texture for geometry. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (june 2000), vol. 3959, pp. 372–383.

[SGG*00] SANDER P. V., GU X., GORTLER S. J., HOPPE H., SNYDER J.: Silhouette clipping. In *Proc. SIGGRAPH* (2000), pp. 327–334.

[TOKK04] TODD J. T., OOMES A. H., KOENDERINK J. J., KAPPERS A. M.: The perception of doubly curved surfaces from anisotropic textures. *Psychological Science 15*, 1 (2004), 40–46.

[War04] WARE C.: *Information Visualization: Perception for design*. Morgan Kaufmann, San Francisco, CA, USA, 2004.

[WDB*07] WEYRICH T., DENG J., BARNES C., RUSINKIEWICZ S., FINKELSTEIN A.: Digital bas-relief from 3d scenes. *ACM Trans. Graph. 26*, 3 (2007), 32.

[Wri05] WRIGHT P. K.: Rapid prototyping in consumer product design. *Commun. ACM 48*, 6 (2005), 36–41.

[ZSGS04] ZHOU K., SNYDER J., GUO B., SHUM H.-Y.: Isocharts: stretch-driven mesh parameterization using spectral analysis. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing* (2004), pp. 45–54.