

Effect of Appearance on Perception of Deformation

Donghui Han*, John Keyser†
Texas A&M University




Figure 1: Frames taken from two animations of 64 objects. Both are rendered using the same physically based simulation of deformable cylinder models. Two appearances (Gastank and Sausage) are used for rendering. Both animations are used as stimuli in the study. High-level information in these static images hints at a rigid gastank and deformable sausage. In animation, however, low-level cues may allow deformations in the gastank to be perceived more easily. We study the potential influences using three user studies.

Abstract

We provide an experimental validation that appearance of an object can influence the perception of its level of deformation in a 3D simulation. Our study provides helpful insights in how to improve visual plausibility of deformation, which may allow artists to adjust their designs to enhance or minimize the perceived deformation in a model. We use a physically-based deformation model to simulate simple geometric shapes undergoing deformation. We apply a number of different appearance and rendering parameters to these objects, and then use two user studies to measure whether appearance used for an object can have a statistically significant effect on the perception of its deformation. In another study, we adjust the number of objects simulated and investigate how this can influence the effect of appearance.

We find that appearance can potentially influence people’s sensitivity to differences of deformation as well as subjective rating of softness in our studies. Further analysis shows that, in simple scenarios, the effect of low-level cues in appearance can be dominant, even if high-level information delivered by appearance has the opposite implication. The third study shows that as the number of objects in a scenario increases, objects are perceived to be stiffer. Also, the effect of low-level cues is weaker.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation;

Keywords: perception, physical simulation, texturing

*e-mail: donghui@cs.tamu.edu

†e-mail: keyser@cse.tamu.edu

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org. SCA '15, August 7 - 9, 2015, Los Angeles, California. © 2015 ACM. ISBN 978-1-4503-3496-9/15/08...\$15.00 DOI: <http://dx.doi.org/10.1145/2786784.2786797>

1 Introduction

3D modeling and texturing are important phases of the 3D animation production pipeline. They directly determine the appearance of an object or a character, and the information delivered by the appearance. They can potentially influence the visual plausibility of the object or character in rendered animation. In Fig. 2, the same deformed tetrahedral model of a torus is rendered with two different embedded rendering meshes and textures. By looking at the static images separately, one may have the impression of a more rigid rim, even if it looks plastic and damaged, and a softer tire. Such judgment is based on the *high-level information* delivered by the appearance of an object—the overall appearance of the model as interpreted by a viewer’s prior knowledge. Although such high-level concepts can provide much information, we focus in this study on its effect on perception of material deformability. Just as the modeling and texturing of an object is used to give a viewer the sense that they are seeing an object from real life, it is intuitive to think that such high-level information will also influence perception of the deformability of the object. However, as we will see from our studies here, this is not an accurate assumption.




Figure 2: One frame of simulation of a torus model rendered as rim (left) and tire (right)

In animation, spatial-temporal information is available. Research on *low-level visual cues*, such as high-contrast (salient) boundaries or corner points, has proved the importance of these cues in peo-

ple’s visual cognition since the work of Robson [1966]. Low-level cues of various amounts and in different forms can have potentially significant influence on people’s perception of deformation. We are faced with a question of whether high-level information can still have a dominant influence on people’s perception. In other words, does the rim still look more rigid than the tire in animation? O’Sullivan and her colleagues have investigated people’s sensitivity to distortion in physically based simulation [O’Sullivan et al. 2003; García et al. 2010]. However, there is little study on how low-level cues and high-level information influence people’s perception of deformation interactively. Our study can potentially help artists to adjust the design of appearance to balance between faithful high-level information and specifically tailored low-level cues if they want to enhance or reduce the deformation perceived by viewers.

Contributions: As a primary study, we try to investigate the effect of appearance on physically based simulation of deformation. The low-level cues are represented by intensity and color variation patterns in all rendered frames in the spatial and temporal domain. Rendering mesh, texture, even hair of a character model can all contribute to low-level cues. We ignore the source of low-level cues and only qualitatively compare appearance with scenarios having more or fewer low-level cues. We perform three psychophysics experiments to study two aspects of people’s perception of deformation: sensitivity to detectable difference of deformation and subjective rating of stiffness.

- We validate that appearance can potentially have a statistically significant effect on people’s sensitivity to difference of deformation and people’s subjective judgment of softness.
- We discover that low-level cues seem to potentially dominate the effect of appearance in simple scenarios like free-falling and colliding of a single object. This indicates that people’s judgment may be affected more by subjective perception of low-level cues rather than high-level knowledge.
- We discover that as the number of objects simulated increases, objects look stiffer to people. The effect of low-level cues is reduced. The high-level information may begin to have more effect.

As a first step in this study, we focus on answering the primary questions above. We leave the quantitative measurement of metrics on various aspects of appearance (e.g. spatial frequency of low-level intensity cues) to future work.

2 Background

We refer the reader to [Nealen et al. 2006] for review of concepts and technologies in the area of deformation modeling in computer graphics.

Perception of Physically Based Simulation

Although the influence of visual cues on perception has attracted much attention recently in computer graphics [McDonnell et al. 2008; Held et al. 2010; Bousseau et al. 2013], the influence on perception of physically based deformation is seldom studied.

Visual information in physically based simulation has been considered for improving the visual plausibility of simulation [Chenney and Forsyth 2000; Twigg and James 2007]. O’Sullivan et al. [2003] evaluates the visual plausibility of simulations where physical parameters have been distorted or degraded. They find some interesting biases and derive a set of probability functions which can be used as a metric to evaluate the visual plausibility of simulations. McDonnell et al. [2006] investigates the effect of cloth deformation on LOD representation. Reitsma et al. [2009] further examines

how a realistic environment setting can affect visual plausibility of physically based simulations. Han et al. [2013] shows that a specific simplification of large scale rigid body simulation methods does not significantly influence the visual plausibility. However, few works have considered deformable objects. Our work addresses this gap by investigating the influence of visual cues on deformation simulation.

Perception of Elastic Deformation

Many works have studied the impact of visual cues on the *haptic* perception of mechanical stiffness in virtual environments. Moody et al. [2009] demonstrates that a single physical surface can be made to “feel” both softer and harder than it is in reality by presenting visual information. These studies tend to indicate that visual feedback is profitable to haptic perception.

Only a few works have studied visual cues as a sole factor in perception of deformation. Argelaguet et al. [2013] introduces the so-called Elastic Images, a pseudo-haptic feedback technique which provides perceivable and exploitable sensation of images without the need of any haptic device. Elastic local deformation of images, procedural shadows and creases are used as visual cues. Pejisa et al. [2012] studies which properties of deforming objects people perceive, and which vocabulary terms they use to describe these properties. Garcia et al. [2010] designs a perceptual experiment to test whether the appeal of deformation simulation was improved by their method and studies local deformations effect in perception of contacts. In contrast, we consider the visual cues coming from appearance and study how such information can influence perception of deformation.

Perception of Low-level Features

In psychology, an early study by Mannos and Sakrison [1974] makes use of a Contrast Sensitivity Function (CSF) [Robson 1966] for measuring the visual fidelity of monochrome still images. CSF indicates that people are most sensitive to middle frequency gratings. Other work studies the perception of motion using visual cues [Ullman 1979]. Low-level features like specular highlights and shading have also been shown to be important in perception of 3D shape [Norman et al. 2004].

In computer graphics, Sweet and Ware [2004] investigates the relationships between view direction, texture orientation and surface orientation in surface shape perception. Winnemöller et al. [2007] presents a psychophysical experiment to determine the effectiveness of various low-level shape cues for rigidly moving objects in an interactive, highly dynamic task.

Few works have studied the function of low-level cues in perception of deformation. Also, in real applications, high/low-level information may take effect simultaneously. Our work puts more weight on realistic applications by studying how high/low-information from appearance interact to influence perception of deformation.

3 Setup

3.1 Physically Based Deformation Model Used

Among the several popular deformation models used in computer graphics, we choose the Lagrangian mesh based methods, which are well-developed and widely used. FEM based methods constitute a big part of this class of methods. They are based on a detailed discretization of traditional mechanics theory. Also, people have optimized the performance of this class of methods in different aspects. Therefore, they are more stable, reliable and accurate. Thus, they can provide realistic simulations.




Figure 3: All appearances used to render physically based elastic deformation simulations of sphere model. From left to right: Billiard, Checker-High, Checker-Low, Clay, Concrete, Refraction, Soccer.

In continuum elasticity theory, the undeformed shape of an object is usually a continuous connected domain $M \subset \mathbb{R}^3$. The coordinates of an undeformed point $\mathbf{m} \in M$ is called the *material coordinates* of that point. In the deformed configuration, the same point locates at its *world coordinates* $\mathbf{x}(\mathbf{m})$. The displacement from *material coordinates* to *world coordinates* $\mathbf{u}(\mathbf{m}) = \mathbf{x}(\mathbf{m}) - \mathbf{m}$ is a vector field defined on M . The displacement gradient is a simple measurement of the variation of the displacement field.

$$\nabla \mathbf{u} = \begin{bmatrix} u_{,x} & u_{,y} & u_{,z} \\ v_{,x} & v_{,y} & v_{,z} \\ w_{,x} & w_{,y} & w_{,z} \end{bmatrix} \quad (1)$$

Two commonly used strain tensors in mechanics and computer graphics are the Green-Saint-Venant strain tensor ε_G and Cauchy's strain tensor ε_C .

$$\varepsilon_G = \frac{1}{2}(\nabla \mathbf{u} + [\nabla \mathbf{u}]^T + [\nabla \mathbf{u}]^T \nabla \mathbf{u}) \quad (2)$$

$$\varepsilon_C = \frac{1}{2}(\nabla \mathbf{u} + [\nabla \mathbf{u}]^T) \quad (3)$$

To integrate the status of the system to the next time step, we need to know the internal force given the strain tensor in the entire domain M . The stress tensor is a fundamental force descriptor from which we can compute force density and traction. Two commonly used stress tensors are Cauchy's stress tensor σ and the Piola-Kirchhoff stress tensor P . A constitutive law builds the connection between the strain tensors and the stress tensors. Different derivations can be used for different materials. For elastic materials, Hooke's law can be used:

$$\sigma = C \cdot \varepsilon \quad (4)$$

For hyperelastic material, the derivative of deformation energy density with respect to the displacement gradient can be used. For isotropic materials, this connection is influenced by two independent values, Young's Modulus E and Poisson's Ratio ν . Between these two, Young's Modulus is a measure of material stiffness. The higher this value, the higher stress is caused for a specific strain. In our study, we adjust this parameter to achieve different stiffness.

For deformation modeling, we use the VegaFEM [Barbič et al. 2012] implementation of a corotational model and an implicit backward Euler integrator because of their stability in practice. We use the Bullet physics library for collision detection.

3.2 Animation Creation

We use Blender to create and edit models used in our study. All surface models are then passed to Tetgen [Si 2013] to generate the

tetrahedral mesh needed for FEM simulation. We use barycentric interpolation to embed the secondary rendering mesh. Simulation and rendering models are adjusted so that most triangles in rendering models are very close to elements in simulation models if they are not contained by any element. We use PovRay to render the deformed rendering mesh in each simulated frame.

To suppress the interference of other factors that may influence people's perception of deformation, we choose a very simple simulation scenario: free falling of objects. The falling objects collide with the ground and bounce up or pile up. The ground is in green and background is black for better observation. Videos are of resolution 1024×768 . They are displayed with a 23 inch screen. We designed and displayed the script for both the pilot study and formal experiment using Psychopy [Peirce 2007].

3.3 Pilot study

We first carried out a pilot study to collect primary data on people's perception of deformation. We choose to measure people's sensitivity to deformation in our animations. To do this, we use a yes-no experiment scheme. In the experiment, eight participants are asked to report whether they witness any deformation in the objects in each animation displayed. Each animation is 6 seconds long. Four simulation models are used: sphere, cylinder, cube and torus. We simulate each model with 32 different Young's Modulus levels ranging from $9e4$ to $8e7$. Thirteen appearances including seven shapes are then used to render these simulation results. We generate animations beforehand at fixed discrete Young's Modulus levels to gain better rendering quality. We then take advantage of the randomly interleaved staircase scheme [Cornsweet 1962]. To determine the next stimuli to display for each appearance, we follow the method in interleaved staircase to compute the next stimuli level, but choose the closest Young's Modulus level among the 32 instead. This scheme converges to psychometric thresholds fast. For each appearance, we use Psignifit tools to fit a logistic psychometric curve for each participant [Fründ et al. 2011]. Point of Subjective Equality (PSE), which represents stimuli parameter level corresponding to 50% probability that an object is perceived as deformable, is an important index in this result.

The result of our pilot study shows a great variance in people's PSE. However, people's answer on stimuli with Young's Modulus level near the two ends ($9e4$ and $8e7$) are very consistent. There are differences in PSE for different appearances in most people's data. Based on the pilot study, in our formal study, we choose a range for Young's Modulus that covers most people's PSE and extends to the ends where people's answer is consistent. Another observation from the study is that once they get used to the experimental scheme, people can give an answer after viewing a small fraction of video. Thus shorter stimuli are used in the formal study.

4 Effect of Appearance

In our formal study, we quantitatively measure *whether appearance can have a statistically significant effect on people’s perception of deformation in animation of simple scenarios*. There is only one object in each stimulus. We choose to measure people’s sensitivity to detectable differences of deformation. We use a fitted psychometric function and thresholds of Young’s Modulus corresponding to detectable differences in deformation as a representation of people’s perception of deformation. We intentionally include different high-level information and low-level cues in the objects’ appearances. Two studies are performed separately. We further analyze *whether high-level information or low-level cues are dominant in the potential effect*.

4.1 Study I: Simulation of Sphere Model

4.1.1 Stimuli Design

We use simulation of a sphere model in this study. To limit variables, appearance differs only in texture. Figure 4 shows the simulation and rendering meshes used in this experiment. The low resolution spherical model is converted to a tetrahedral mesh with 158 elements for simulation.




Figure 4: Spherical model used in simulation(left, 80 triangles) and the secondary model used for rendering(right, 5120 triangles).

Seven appearances used in this study are given in Figure 3. Obviously, these appearances have different amounts of low-level cues as well as completely different high-level information. We use the 2 alternative forced choice (2AFC) protocol [Blackwell 1953] which is bias free compared to a Yes-No scheme. Each stimulus is paired with an animation with the highest Young’s Modulus as reference. The two stimuli are displayed successively in a random order. Participants are told that only one of each pair has rigid objects and they need to choose the one with deformable objects. Simulations at 9 Young’s Modulus levels in the range ($9e4, 9e5$) with about even intervals are used. The number of times these animations are displayed at these levels are 2,4,6,6,6,6,8 and 10. Animations with the softest object are displayed only twice because these can be easily distinguished from the hardest one. Each stimulus is 1 second which is long enough for participants to make a judgment. Sixteen people participated in this study(11M/5F), most of which are naive to physically based elastic deformation simulation and rendering. Their ages ranged from 21 to 31. We explained each appearance to participants before the experiment. Full length example videos were then displayed to allow participants to observe each appearance. Then, the users began watching stimuli.

4.1.2 Result and Analysis

For each appearance, we fit a logistic psychometric curve for each participant as well as for overall merged data. Figure 5 shows the fitted curve of every appearance using overall merged data. In a 2AFC experiment, the lower asymptote of a fitted function is equal to 0.5, which is what one would get by guessing. So, we choose the Just Noticeable Difference (JND) threshold as a representation

of people’s perception of deformation. In this study, JND is a relative modification of Young’s Modulus with respect to the most rigid level. JND corresponds to 75% of a chance that a stimulus is recognized correctly as deformable. This threshold indicates people’s sensitivity to differences in perceived deformation. The point and interval (95% confidence interval) estimation of JND for each appearance is listed in Table 1.

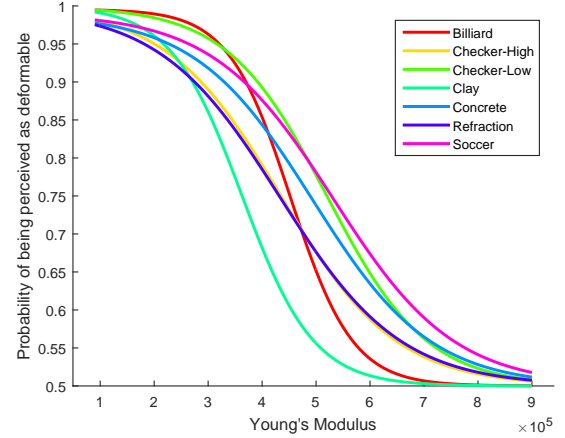


Figure 5: Psychometric curves fitted using accumulated data for every appearance in study I.

Table 1: Point estimation and interval estimation(95% confidence, relative to Point) of JND

Texture	JND	
	Point	Interval
Billiard	49.70%	-3.97% , +3.45%
CheckerBoardH	51.67%	-5.13% , +4.64%
CheckerBoardL	42.18%	-4.99% , +4.47%
Clay	59.67%	-3.67% , +3.71%
Concrete	45.14%	-5.78% , +5.51%
Refraction	51.91%	-5.14% , +4.99%
Soccer	40.81%	-2.22% , -10.67%

We further estimate the JND for each participant and each appearance. A one-way repeated measures ANOVA is performed. Appearance is found to have a significant effect on people’s JND ($F_{6,90} = 2.9341, p < 0.015$). A Dunn-Sidak post hoc analysis shows that JND of appearance Clay is significantly different from the JND of appearance corresponding to the rightmost curves in Figure 5. This indicates that appearance can have a significant effect on JND in this study. Further investigation of curves in Figure 5 and estimation of JND for appearances in Table 1 shows a consistent result. Curves for different appearances diverge from each other in a wide probability range. Point estimation of JND for different appearances spread from 40% to nearly 60%.

Although not every pair of JND for different appearances are significantly different, we can still tell the relative order of perceived stiffness of different appearances. An appearance with higher estimated JND generally corresponds to curves on the left in Figure 5. For such appearances, a larger reduction to Young’s Modulus is necessary for people to perceive detectable deformation with the same accuracy as for other appearances. On the other hand, people have lower accuracy in detecting the difference of perceived deformation when a stimulus with such an appearance is displayed, compared to a stimulus with another appearance. In this sense, ap-

pearances with higher JND and curves on the left are perceived as less deformable. In contrast, lower JND indicates more deformable.

A careful investigation of Figure 5 leads to an interesting observation. Clay is perceived as the most rigid appearance. Billiard and Refraction (Glass-like) are all perceived as more deformable than Clay. Concrete is even closer to be perceived as the most deformable appearance. This order of perceived stiffness is not consistent with the order of stiffness simply inferred from the high-level information in these appearances. On the other hand, Checker-Low is perceived as one of the most deformable. Checker-High is closer to Clay. But neither of these two has clear high-level information since a sphere with such texture is not commonly seen in real life. All these lead to a positive correlation between the amount of low-level cues and people’s perceived softness. Soccer, Checker-Low and Concrete, which are perceived as more deformable, all have lots of high-contrast local features which are easy to track and recognize. In contrast, Clay, Refraction and Checker-High, either have plain color on the surface or have high frequency intensity variation that is hard for people to track. Comparing to texture with abundant low-level cues, it can possibly be more difficult for people to recognize deformation from these textures. In this sense, the more easy-to-track low-level cues an appearance has, the more deformable an object with such appearance is perceived to be. The effect of high-level information can be opposite to that of the low-level cues in this study and may weaken the effect of each other. This can possibly explain why there is a single pair of significantly different textures. Eventhough, the order of these textures is more consistent with the effect of low-level cues. In simple scenarios, the effect of low-level cues is thus more dominant.

4.2 Study II: Simulation of Torus Model

4.2.1 Stimuli Design

To further investigate the effect of appearance we perform a second study using a torus model. In this study, we further consider visual information contributed by the rendering mesh. We use two secondary rendering meshes with similar bounding shapes but different structure and details to create four different appearances.




Figure 6: Torus model used in simulation(left, 310 triangles) and the secondary models used for rendering: tire(middle, 16416 triangles), rim(right, 11864 triangles).

Figure 6 shows the triangular models we use for simulation and rendering. We use the left model in Figure 6 for simulation. It is a low resolution approximation of both rendering meshes. This model is further tessellated into 370 tetrahedral elements for simulation.

Four appearances used in this study are shown in Figure 7. We use the same experimental design as in study I. We choose simulation at 10 Young’s Modulus levels in the range $(1e5, 3.3e6)$. The number of times we display each of the four stimuli at the ten Young’s Modulus levels are 2,4,6,6,6,6,6,8 and 10. Each stimulus is 2 seconds. We recruit another 11 volunteers (3F/8M) in this study. All of them are naive to the study. Their age ranges from 22 to 31.



Figure 7: All appearance used to render physically based elastic deformation simulation of torus model. From left to right: Rim-Bright, Rim-Dark, Tire-Black, Tire-Checker.

4.2.2 Result and Analysis

Again, we fit logistic psychometric curves for each participant as well as for overall merged data. The fitted curve for each appearance using the merged data is shown in Figure 8.

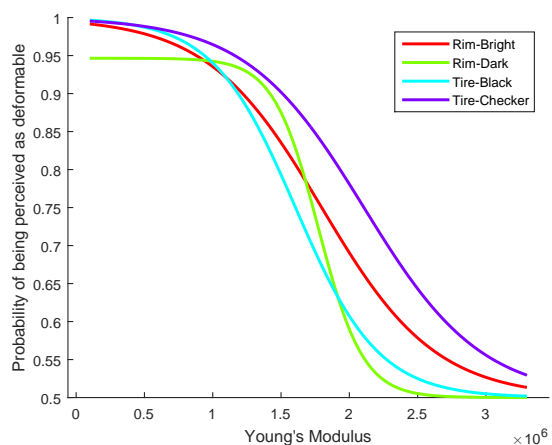


Figure 8: Psychometric curves fitted using accumulated data for all four appearances used in study II.

Table 2: Point estimation and interval estimation(95% confidence, relative to Point) of JND

Texture	JND	
	Point	Interval
Rim-Bright	45.53%	-6.59% , +5.28%
Rim-Dark	46.27%	-5.68% , +3.29%
Tire-Black	51.34%	-4.78% , +3.96%
Tire-Checker	36.16%	-7.44% , +6.44%

Estimation of JND for each appearance using overall data is reported in Table 2. One-way repeated measures ANOVA finds appearance to be a significant factor again ($F_{3,30} = 5.1647, p < 0.006$). Post hoc analysis finds the two tire appearances to have significantly different JND. The curves for the two tire appearances are on the opposite sides in Figure 8. The estimated JND for these two appearances have an approximate difference of 15%.

It is interesting to notice that Tire-Black, which should be the most deformable appearance according to the clear high-level information, is perceived as the most rigid one, even comparing to rim appearance. Tire-Checker has no clear high-level information. But it is perceived as more rigid than both rim appearance. Obviously, high-level information does not have significant influence

here. Tire-Checker has the most easy-to-track and recognize low-level cues. It is thus perceived as the most deformable one. Both rim appearances do not have colorful texture. But the complex rendering mesh causes more easy-to-track low-level cues in rendered frames than Tire-Black appearance. Again, the amount of low-level cues in appearance is positively correlated with the perceived softness. And this effect is dominant even if high-level information has opposite influence.

This study confirms the significant effect of low-level cues in appearance. It further shows that low-level cues can be from both texture and rendering mesh. Similarly, the amount of low-level cues can also be controlled in many other ways. When designing the appearance of an object which will be simulated in 3D animation, visual reality is important. But if an artist wants to achieve a special effect by enhancing or reducing the perceived deformation, more attention may need to be paid to design of low-level details to control the visual cues.

5 Study III: Influence of complexity of Scene

In the first two studies, we statistically validate the significant effect of appearance in people’s perception of deformation. Experimental data indicates that low-level cues have a dominant effect while high-level information does not. However, people’s perception of low-level cues rely heavily on tracking of visual features. There has been study on the limit of people’s ability in tracking multiple objects [Vul et al. 2009]. Complex scenes may contain too many low-level cues which are beyond people’s ability to recognize and process. We designed our third experiment to study *how complexity of a scene can influence the effect of appearance in people’s perception of deformation*. We adjust the complexity of the scene by controlling the number of objects simulated. Furthermore, we measure people’s perception of deformation by using their subjective rating of stiffness of objects. Although it may contain higher variance, it also provides more direct and abundant information.

5.1 Stimuli Design

We use the cylinder model in Figure 9 for simulation. It is tetrahedralized into 258 elements. To reduce the variance among people’s rating, we display reference animation at the beginning and middle of the experiment. We use an extra appearance (Checker) in the reference animation. Only Gastank and Sausage are used as appearance in formal stimuli.




Figure 9: From left to right: cylinder model used in simulation (224 triangles) and the secondary rendering models gastank (7102 triangles), sausage (7680 triangles), checker (1344 triangles).

We use simulation at 7 Young’s Modulus levels in the range (9e4, 1.6e6). Five reference animations with different Young’s Modulus levels evenly sampled in this range are displayed side-by-side for people to observe. All objects in these animations have a checker appearance. An integer number (1 to 5) is displayed on each animation to indicate its level of stiffness. Higher numbers correspond to stiffer objects. Reference animations are displayed in a loop until a participant is ready to proceed. For each appearance at each Young’s Modulus level, we perform simulations of 1, 4, 16,



Figure 10: All appearance used in study III. From left to right: Gastank, Sausage, Checker (only used in reference animation).

64 and 256 objects. Figure 1 shows example frames of stimuli with 64 objects. Overall, we have 70 distinct stimuli: 2 appearances \times 5 Number of Objects (NoO) \times 7 Young’s Modulus levels (YM). Each 2-second stimulus is displayed four times in total. A random order is used for each participant to display the 280 stimuli. After watching each stimulus, participants are asked to rate the stiffness of objects by choosing a number from 1 (most deformable) to 5 (most rigid). We use the average of four ratings of each distinct stimuli as the final rating. We recruit 11 new volunteers (6M/5F) for this study. Their ages range from 23 to 40. All of them are naive to this study.

5.2 Result and Analysis

We first carry out a three way repeated measures ANOVA on the final rating. The result is reported in Table 3. To help understand the result, we also fit a function that represents the relationship between rating and Young’s Modulus for each combination of appearance and NoO levels. These functions are shown in Figure 11.

One obvious and foreseeable result is the significant effect of YM in rating. We further apply the Dunn-Sidak post-hoc analysis on this main effect. We find, as we would expect, that all YM levels are significantly different from each other, and the rating mean increases as the YM level increases. In Figure 11, rating curves for all NoO and appearance combinations increase monotonically, which reflects the same effect. This indicates that people are sensitive enough to the difference between YM levels used in this study.

We do not find enough support for the significance of appearance as a main effect. However, the interaction between appearance and NoO is found to have significant effect ($F_{4,40} = 2.735, p < 0.045$). Post-hoc analysis indicates that mean ratings for Gastank and Sausage are significantly different only when NoO is 1. The curves for one Gastank object and one Sausage object in Figure 11 are diverged from each other throughout the entire YM range. The rating for Sausage is uniformly higher than Gastank when there is only one object. This result for a single object is consistent with our discovery in the previous two studies. Although the high-level information in appearance Gastank hints for more rigid material, there are more low-level cues in the form of high-contrast visual features in Gastank compared to Sausage. These features may allow the recognition of deformation to be easier. Thus Gastank can be perceived as more deformable. In simple scenarios with one object, the effect of the low-level cues obviously dominates the effect on people’s rating again, even though the high-level information has the opposite influence. This is not significant when there are more objects, however.

Another interesting observation is that NoO also has a significant effect on rating. Post-hoc analysis and fitted curves show that more objects generally cause higher rating for stiffness. In other words, as the number of objects under simulation increases, people begin to feel objects are less deformable, even though both the Young’s Modulus and appearance of objects are not changed.




Figure 11: Fitted curves that represent the relationship between people’s rating and Young’s Modulus. From left to right, Curves for objects with appearance Gastank, Curves for objects with appearance Sausage, Curves for 1 and 256 objects with appearances Gastank and Sausage.

Table 3: Three way repeated measures ANOVA on people’s rating of stiffness - 2 Textures \times 5 Number of Objects (NoO) \times 7 Young’s Modulus levels (YM)

Factor	F-value	pValue	Post-hoc
NoO	$F_{4,40} = 44.6081$	$p < 4e - 14$	Higher rating for more objects Higher rating for higher Young’s Modulus Appearance has significant effect only for single object scenario Number of Objects has more effect in Young’s Modulus range (2e5, 8e5)
YM	$F_{6,60} = 170.0229$	$p < 1e - 35$	
Appearance \times NoO	$F_{4,40} = 2.735$	$p < 0.045$	
NoO \times YM	$F_{24,240} = 8.0001$	$p < 1.2e - 19$	

The study on people’s limit in a multiple objects tracking task [Vul et al. 2009] can provide a potential explanation of this observation. As number of objects increase, the amount of low-level cues increase accordingly. Tracking these cues and recognizing relative deformation is a more difficult task than simply tracking one object. Thus, the low-level cues in a scene can simply overload people’s ability to process them. When this happens, people are more likely to be distracted by overloaded cues rather than processing them meaningfully. As a result, in complex scenes with lots of objects, people are less likely to recognize deformation and thus give higher rating. An experienced viewer may focus on a single object or local area in animation to avoid distraction. However, in entertainment scenarios where people are relaxed, they can easily be distracted by too many low-level cues when watching unfamiliar scenes for a short period. This also explains why texture does not have a significant effect when NoO is high. Low-level cues in multiple Gastank or Sausage objects stimuli can both overload people’s attention. Thus the effect of low-level cues are depressed in both cases.

With more objects, there is also no significant effect that high-level information matters. However, in Figure 11, the functions for 64 and 256 Gastank are higher than the corresponding functions of Sausage in a wide range of YM. This may be a hint that as a perceptual system is overloaded, the high-level information indeed becomes important in perceived deformation. More study would be needed to confirm or refute this, though.

It is also important to notice that the effect of NoO is weaker when YM is closer to the two ends of selected range. The effect seen in interaction between NoO and YM is partly caused by this. Intuitively, NoO has less influence on very soft and very rigid objects, which are more obviously soft or rigid.

6 Conclusion and Future Work

We performed three primary experiments to study the effect of objects’ appearance on people’s perception of deformation in com-

puter generated animation. We validate that appearance can potentially have a significant effect on people’s sensitivity to detectable differences of deformation and subjective rating of stiffness of objects. However, we find this effect to be significant only in simple scenarios with a single object of moderate stiffness. Furthermore, the effect of low-level cues in appearance is dominant comparing to high-level information in our studies. When low-level and high-level cues contradict each other, low-level cues will win out. We further investigate the influence of number of objects in the scene, controlling for the complexity of the scene. We discover that increasing the number of objects in stimuli can cause people to feel the objects are more rigid. This influence also depresses the effect of appearance on rating. Significant differences between ratings for one Gastank object and one Sausage object decrease as the number of objects in the stimuli increase. One possible explanation is that redundant low-level cues in complex scenes could simply overload people’s perceptual ability, thus causing more distraction than meaningful influence. This depresses the effect of low-level cues as well as the overall effect of appearance. However, we have seen some hints that high-level information could actually assume a dominant role when large numbers of objects are simulated.

The results provide an important guide to artists in designing objects for animation. First, our quantitative validation proves that perception of deformability of an object can be adjusted by the design of its appearance. Second, our discovery shows that simply designing the appearance of an object as it is in real life may not be an optimal choice in controlling the perceived deformability. To enhance or reduce the perceived deformability, one can balance the design to include more or fewer low-level cues in various forms (e.g. rendering mesh, texture, decoration, furs, etc.) or to achieve more realistic design. In complex scenes, however, more attention may need to be put on designing appearance that can deliver clearer high-level hints.

As a primary study, we only validate the significance of appearance, especially that of the low-level cues, in people’s perception of deformation. We qualitatively distinguish more and fewer low-level

cues in appearance while ignoring the source. Our three studies are limited in the number of models and textures tested. More specifically designed experiments are necessary to verify our discovery in more general scenarios. Our future work will propose quantitative metrics of the amount and form of low-level cues in influencing people's perception. To do that, we may study texture, shape, object moving speed, environment, etc. separately as a single source of low-level cues. We need to study typical aspects of these sources which could be quantitatively measured (e.g. spatial-temporal frequency of intensity variation in texture). By removing the effect of high-level information, we can propose quantitative metrics which could be directly used by artists. On the other hand, other quantitative aspects of people's perception of deformation (e.g. sensitivity in distinguishing different stiffness levels) could be considered as important indices. Also, we are interested in discovering exactly when high-level information could have a more significant effect.

Acknowledgements

We would like to give our thanks to Benjamin22 and bryanajones for sharing the rim model and gas tank model in Blendswap. The user studies were approved via IRB2014-0258D.

References

- ARGELAGUET, F., JÁUREGUI, D. A. G., MARCHAL, M., AND LÉCUYER, A. 2013. Elastic images: Perceiving local elasticity of images through a novel pseudo-haptic deformation effect. *ACM Trans. Appl. Percept.* 10, 3 (Aug.), 17:1–17:14.
- BARBIČ, J., SIN, F. S., AND SCHROEDER, D., 2012. Vega FEM Library. <http://www.jernejbarbic.com/vega>.
- BLACKWELL, H. R. 1953. *Psychophysical thresholds: experimental studies of methods of measurement*. Engineering research bulletin. Engineering Research Institute, Univ. of Michigan.
- BOUSSEAU, A., O SHEA, J. P., DURAND, F., RAMAMOORTHI, R., AND AGRAWALA, M. 2013. Gloss perception in painterly and cartoon rendering. *ACM Transactions on Graphics* 32, 2 (April).
- CHENNEY, S., AND FORSYTH, D. A. 2000. Sampling plausible solutions to multi-body constraint problems. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '00, 219–228.
- CORNSWEET, T. N. 1962. The staircase-method in psychophysics. *The American Journal of Psychology* 75, 3, pp. 485–491.
- FRÜND, I., HAENEL, N. V., AND WICHMANN, F. A. 2011. Inference for psychometric functions in the presence of nonstationary behavior. *Journal of Vision* 11, 6 (May).
- GARCÍA, M., OTADUY, M. A., AND O'SULLIVAN, C. 2010. Perceptually validated global-local deformations. *Comput. Animat. Virtual Worlds* 21, 34 (May), 245–254.
- HAN, D., HSU, S.-W., MCNAMARA, A., AND KEYSER, J. 2013. Believability in simplifications of large scale physically based simulation. In *Proceedings of the ACM Symposium on Applied Perception*, ACM, New York, NY, USA, SAP '13, 99–106.
- HELD, R. T., COOPER, E. A., O'BRIEN, J. F., AND BANKS, M. S. 2010. Using blur to affect perceived distance and size. *ACM Trans. Graph.* 29, 2 (Apr.), 19:1–19:16.
- MANNOS, J., AND SAKRISON, D. J. 1974. The effects of a visual fidelity criterion of the encoding of images. *Information Theory, IEEE Transactions on* 20, 4, 525–536.
- MCDONNELL, R., DOBBYN, S., COLLINS, S., AND O'SULLIVAN, C. 2006. Perceptual evaluation of lod clothing for virtual humans. In *Proceedings of the 2006 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, SCA '06, 117–126.
- MCDONNELL, R., LARKIN, M., DOBBYN, S., COLLINS, S., AND O'SULLIVAN, C. 2008. Clone attack! perception of crowd variety. *ACM Trans. Graph.* 27, 3 (Aug.), 26:1–26:8.
- MOODY, L., WATERWORTH, A., ARTHUR, J. G., MCCARTHY, A. D., HARLEY, P. J., AND SMALLWOOD, R. H. 2009. Beyond the visuals: Tactile augmentation and sensory enhancement in an arthroscopy simulator. *Virtual Real.* 13, 1 (Feb.), 59–68.
- NEALEN, A., MLLER, M., KEISER, R., BOXERMAN, E., AND CARLSON, M. 2006. Physically based deformable models in computer graphics. *Computer Graphics Forum* 25, 4, 809–836.
- NORMAN, J. F., TODD, J. T., AND ORBAN, G. A. 2004. Perception of three-dimensional shape from specular highlights, deformations of shading, and other types of visual information. *Psychological Science* 15, 8, 565–570.
- O'SULLIVAN, C., DINGLIANA, J., GIANG, T., AND KAISER, M. K. 2003. Evaluating the visual fidelity of physically based animations. In *ACM SIGGRAPH 2003 Papers*, ACM, New York, NY, USA, SIGGRAPH '03, 527–536.
- PEIRCE, J. W. 2007. Psychopypsychophysics software in python. *Journal of Neuroscience Methods* 162, 12, 8 – 13.
- PEJSA, T., 2012. Understanding perception of simulated elastic deformations. Course Project Report.
- REITSMA, P. S. A., AND O'SULLIVAN, C. 2009. Effect of scenario on perceptual sensitivity to errors in animation. *ACM Trans. Appl. Percept.* 6, 3 (Sept.), 15:1–15:16.
- ROBSON, J. G. 1966. Spatial and temporal contrast-sensitivity functions of the visual system. *J. Opt. Soc. Am.* 56, 8 (Aug), 1141–1142.
- SI, H. 2013. *Tetgen - A Quality Tetrahedral Mesh Generator and a 3D Delaunay Triangulator*. Weierstrass Institute for Applied Analysis and Stochastics (WIAS), Mohrenstr. 39, 10117 Berlin, Germany.
- SWEET, G., AND WARE, C. 2004. View direction, surface orientation and texture orientation for perception of surface shape. In *Proceedings of Graphics Interface 2004*, GI '04, 97–106.
- TWIGG, C. D., AND JAMES, D. L. 2007. Many-worlds browsing for control of multibody dynamics. In *ACM SIGGRAPH 2007 papers*, ACM, New York, NY, USA, SIGGRAPH '07.
- ULLMAN, S. 1979. *The interpretation of visual motion*. Massachusetts Inst of Technology Pr.
- VUL, E., ALVAREZ, G., TENENBAUM, J. B., AND BLACK, M. J. 2009. Explaining human multiple object tracking as resource-constrained approximate inference in a dynamic probabilistic model. In *Advances in neural information processing systems*, 1955–1963.
- WINNEMÖLLER, H., FENG, D., GOOCH, B., AND SUZUKI, S. 2007. Using npr to evaluate perceptual shape cues in dynamic environments. In *Proceedings of the 5th International Symposium on Non-photorealistic Animation and Rendering*, ACM, New York, NY, USA, NPAR '07, 85–92.