J. Coat. Technol. Res.
DOI 10.1007/s11998-011-9378-2

Computer rendering and visual detection of orange peel

Jonathan Konieczny, Gary Meyer

Abstract The computer graphic simulation of a common spray painting artifact, called orange peel, is discussed. Orange peel distorts surface reflections and is commonplace in product design applications. The orange peel measurements from a standard industrial instrument are used to construct a height field, and this surface is rendered using traditional normal mapping techniques. Comparisons are made between real panels with orange peel and simulations of those panels. A simple visual model for detecting the presence of orange peel is also presented and evaluated. User testing of the model confirms that orange peel is more visible on dark paint colors than on light paint colors. The latter outcome suggests that to minimize application time, but still keep orange peel below visual threshold, paint application systems should be designed to take paint color into account.

Keywords Orange peel, Computer graphics, Simulation, Measurement

Introduction

The simulation of appearance defects is an important part of achieving complete realism in computer graphic pictures. Although early realistic images featured flawless surfaces with shading and color produced by evaluating idealized surface reflection models, recent study has attempted to include imperfections such as the effect of aging and weathering on the appearance of objects exposed to their surrounding environments.[1] The development of interesting patinas that add character to an object's appearance has been studied,[2,3] as has the weathering of statues constructed from stone.[4,5] Gouges and dents due to impacts and collisions have also been accounted for in the creation of realistic pictures.[6,7,8] Even the cracking of paint with age[9,10] and the accumulation of dust on a surface[11] have been included as part of the image synthesis process.

Orange peel is an important practical surface defect that occurs as a result of using a spray gun to apply paint to a surface. Orange peel is an irregularity in the top surface coating that causes distortions in the manner in which a surface reflects its surrounding environment. The complete modeling of cosmetic surfaces in a manufactured environment requires that this imperfection be included as part of the simulation. Fortunately, there are instruments for measurement that can be used to characterize the nature of orange peel and that can be employed to help generate pictures of this flaw. Adding orange peel to the catalog of successfully simulated surface blemishes is one contribution of this article.

Detecting the existence of orange peel and minimizing its effect is an important problem in the surface coatings industry. One of the differentiating factors between a luxury automobile and an inexpensive sedan is the absence or the presence of orange peel on each vehicle. In this article, we use our orange peel simulation to develop and test a simple metric that determines whether orange peel is visible to a human observer. If the relationship between spray gun parameters and orange peel creation is known, then this metric can be used to guide the development of optimal spray paint application systems. This demonstration, that computer

Electronic supplementary material The online version of this article (doi:10.1007/s11998-011-9378-2) contains supplementary material, which is available to authorized users.

J. Konieczny, G. Meyer (✉)
Department of Computer Science, Digital Technology Center, University of Minnesota, Minneapolis, MN, USA
e-mail: meyer@cs.umn.edu

J. Konieczny
e-mail: jonathankon2632@yahoo.com

Published online: 29 November 2011

graphic simulation can be used to develop quality control tools for paint application, is another contribution of this article.

Background

Quality control of surface finishes requires measurements to be made with special purpose instruments. The primary purpose of these measurements is to determine whether or not a particular sample is ready for the marketplace, or if modifications need to be made to the surface finish or how it is applied. These instruments are therefore designed to provide an intuitive and accurate method for making this determination.

As these measurement instruments become more sophisticated, they begin to not only measure a material's general acceptability, but also give a very accurate characterization of the surface finish's color, gloss, and even small scale geometry. This can aid in making specific corrections to the manufacturing or application process. It also provides an opportunity to the computer graphics community: the accuracy and usability of these instruments can be leveraged to make accurate renderings of simulated surfaces.

While many computer graphics techniques for measuring geometry and color require elaborate rigs to take pictures of the material from all possible angles, these industrial instruments measure only the portions of the sample that professionals have found are relevant to the final appearance of that material. This has three important advantages over completely characterizing the surface reflection of the sample. First, the measurement instrument itself can be much less complicated, allowing for a very rapid and simple measurement process that anyone can do with little or no training. Second, by studying the instrument that industry professionals use, a computer graphics simulation can be set up, which mimics the measurements made by that instrument. This allows a designer that is not trained in the computer graphics field to make use of a familiar tool to design and create new prototype colors and materials on a computer. Finally, these instruments are commercially available: they can be purchased "off the shelf" already set up and properly calibrated. A user does not need to worry about constructing an elaborate data acquisition gantry or calibrating the instrument against known measurements: this is already done for them.

This research leverages the above advantages of a commercially available instrument to create graphical simulations from the Byk-Gardner Wave-Scan™. This instrument measures orange peel: a rippling across the surface of the paint so named because of its visual similarity to the skin of an orange (although paint orange peel usually has a smaller amplitude, and therefore more subtle look than an actual orange). At close viewing distances, these ripples can be directly seen by the eye. At longer distances, the surface changes integrate into each other, causing the reflection on the surface of an object with orange peel to appear more blurry than on an object without the defect. Figures 1 and 2 illustrate the effects of orange peel. While orange peel is seen in the automotive industry as a defect, in some industries, such as furniture and upholstery, an orange peel like roughness is purposely added to diffuse light or hide defects such as scratches. Thus, orange peel and similar surface roughness can be designed to both give defect tolerances as well as be purposely added into furniture, automotive interiors and other materials.

The study described here also utilizes sophisticated computer graphic-rendering techniques and demonstrates how they can be employed in the surface coatings industry. Image synthesis methods have improved to such a stage that it is now possible to mathematically define a scene and create pictures that are so realistic that there is little difference between them and photographs of an actual physical environment. Computer graphics hardware has also made significant progress and certain types of photorealistic

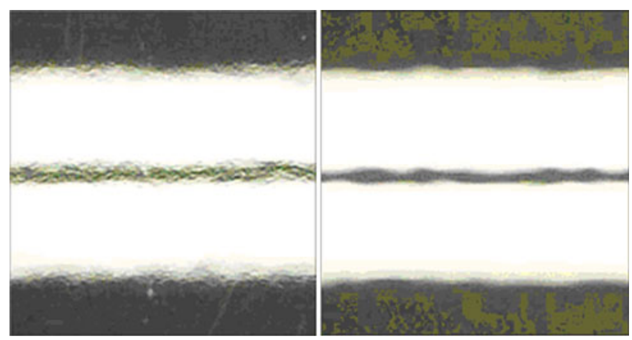

Fig. 1: A picture of the effect orange peel has on the reflection of a light source. The image on the left is a gloss coat painted over rough steel, and therefore has worse orange peel than the image on the right, which is a gloss coat painted over smoothed steel. Taken from reference (12)

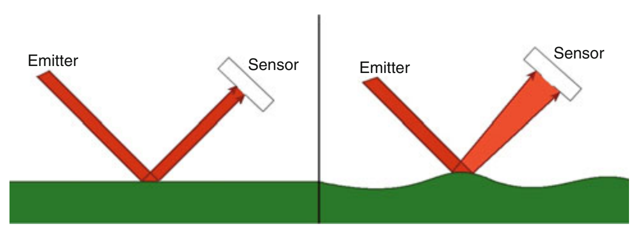

Fig. 2: An illustration of how the Wave-Scan instrument takes measurements of the surface profile. The light from a laser strikes the surface, and reflects back into a photosensor. Distortions in this reflection are measured to characterize the orange peel. The perfect mirror surface (without orange peel) reflects the light back into the photosensor without distortion. The orange peel on the right causes the reflection to widen, making the photosensor measure less light

pictures can now be generated in real time. Because visual appearance can be an important consideration in evaluating a surface coating, these advances in realistic image synthesis have important implications for the paint industry. Images of hypothetical new coatings and finishes can be created without actually formulating and manufacturing the material. The result is a virtual design process for coatings that is similar to the computer-aided geometric design approach that is routinely used today to create new three-dimensional objects. This article demonstrates how this new design methodology can be used to simulate and evaluate the appearance of a specific surface coating phenomena: orange peel.

The rest of this article describes research on both rendering of orange peel as well as the detection of surface roughness under varying viewing conditions and with different surface reflectance properties. The study reported here is one part of a larger research project devoted to simulating the application and appearance of automotive paint.[13] First, information is provided on how the Wave-Scan works. Second, a rendering algorithm is developed based on the measurements given by the Wave-Scan. Next, an analytic model that predicts the visibility of orange peel is developed, given both surface reflectance and human perceptual factors. Finally, this system is used to run user experiments that equate surface reflectance and geometric properties with the ability of humans to detect the roughness of the surface. This shows that the Wave-Scan instrument measurements can be combined with reflectance and lighting models to design defect tolerances for specific paint and lighting conditions.

Relevant work

Rendering of gloss and paint sparkle has been performed by a number of authors.[14,15,16] Orange peel has been less studied. Dumont-Becle et al.[16] performed a rendering of orange peel using a bump map layer in the simulation. In a similar manner, our study uses normal maps to perform the rendering. However, our normal maps are generated using actual measurements from an industry device made for the specific purpose of testing orange peel. This permits our simulation to be far more precise, allowing a more realistic final rendering as well as design and prototyping of materials that have orange peel.

The orange peel visibility predictor developed in this article is based on the visible differences predictor given by Daly.[17] In addition, it uses Ward's just noticeable differences algorithm.[18] The model introduced in this article adds a reflectance model, in this case the Phong[19] model, to create a prediction of visibility that includes not only perceptual factors, but also geometry and surface reflectance.

Our current rendering system for orange peel builds on top of a metallic paint rendering and design system created by Shimizu et al.[20] While complicated optical models have been developed by coating scientists to describe the reflectance of light from rough surfaces,[21] the approach taken in Shimizu's study is to develop a reflectance model that provides acceptable accuracy[15] and can be evaluated in real time. This system allows a user to design different metallic paint colors and view a photorealistic rendering of the designed color in real time. Our system builds on this, adding our real-time orange peel simulation to the rendered image.

Definition and measurement of orange peel

Orange peel is an extremely common defect that occurs with many sprayed paints, including metallic car paints. It is a series of bumps or ripples across the surface of the paint, called orange peel due to its visual similarity to the skin of an orange (although paint orange peel usually has a smaller amplitude and therefore a less pronounced appearance than an actual orange). At close viewing distances, these ripples can be directly seen by the eye. At longer distances, the surface changes integrate into each other, causing the reflection on the surface of an object with orange peel to appear more blurry than on an object without the defect. Figures 1 and 2 illustrate the effects of orange peel.

Orange peel is a significant enough problem in the paint industry that special devices have been created to measure it and to determine if a paint job has an "acceptable" level of orange peel. The Byk-Gardner Wave-Scan[12] evaluates orange peel by running a laser over the surface of the paint, and then measuring the reflection of the laser off the paint. Ideally, the laser should reflect in the perfect mirror direction off the paint, which causes it to be directly bounced into the center of the photometer in the instrument. However, orange peel will cause the laser to be distorted (by shorter wavelength orange peel) or shifted away from the center of the photometer (for longer wavelength orange peel). See Fig. 2 for an illustration of how the Wave-Scan makes individual measurements.

The Wave-Scan records measurements of these distortions as the instrument is run across the surface of a painted sample. Once a sufficient number of samples have been taken, the surface is characterized by the magnitude of distortions at five different magnitudes: 0.1–0.3, 0.3–1, 1–3, 3–10, and 10–30 mm (see Fig. 3). The magnitudes of these five wavelengths are then used to determine if the orange peel on the surface is acceptable or not, as well as possible reasons for what needs to be fixed if the orange peel is not acceptable (for instance, if the 3–10 mm amplitude is too high, one possible fix is to spray on a thicker gloss coat). The actual output of the Wave-Scan instrument is five numbers ranging from 1 to 100. These five values are referred to as W_a, W_b, W_c, W_d, and W_e, and they are scaled so that each step (i.e., 1–2 or 43–44) provides a perceptually uniform increase in appearance.

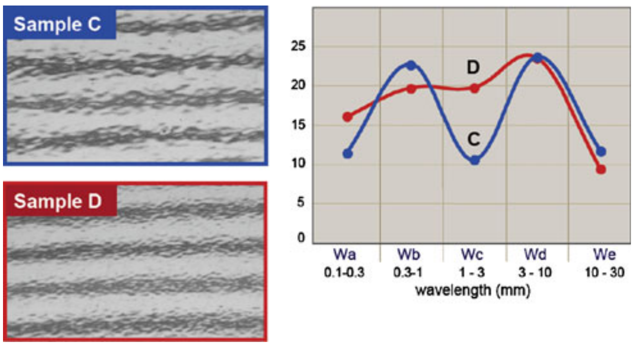

Fig. 3: A Wave-Scan graph showing the amplitude profile of the orange peel across varying wavelengths. Taken from reference (12)

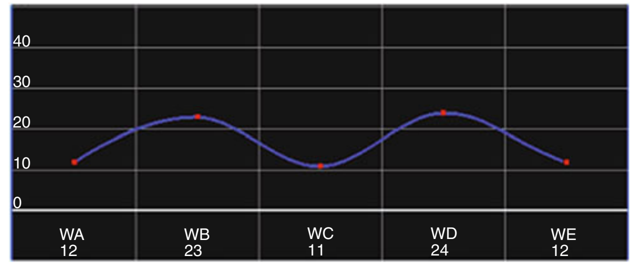

Fig. 4: An image of the user interface provided for inputting Wave-Scan values. This correlates to the graph of how the Wave-Scan gives its measurement values, shown in Fig. 3

Orange peel simulation

The purpose of our simulation is to effectively reverse-engineer what the Wave-Scan instrument does: use the five wavelength amplitudes to construct a visual simulation of what that surface looks like. This is done in three steps. The first is a relatively simple user interface that allows the user to enter the five wavelength values and to display them in a similar manner to that shown in Fig. 3. The second is to take these amplitudes and construct a virtual surface that is representative of the original measured surface. The final step is to render this surface realistically.

User interface

The user interface provides a method for the user to enter the value of the five wavelengths, and displays the result as a spline-fit curve of the surface profile. This is designed to have an appearance similar to the actual Wave-Scan graph readings shown in Fig. 3. Besides manually inputting the values, the user can directly drag the data points on the display graph. Figure 4 shows a screenshot of the provided user interface.

In addition to the Wave-Scan value input interface, the interface given by Shimizu et al.[20] is provided to the user. This allows the user to modify the color and glossiness of the paint and view how this affects the visibility of the orange peel. The "Orange peel detection" section shows how the combination of these tools can then be used to prototype defect tolerances.

Surface construction

Once the five wavelength values have been entered, those values must be converted into a virtual surface that has a similar appearance to the surface originally scanned by the device. This is an under-constrained problem: an infinite set of possible surfaces could have mapped to the same Wave-Scan values. The information that the Wave-Scan values provide is the five wavelengths: set at their original values from 0.1 to 30 mm, and their corresponding amplitudes: a number from 1 to 100. The system must fill in the missing parameters to generate a visually plausible surface.

The first step in this process is to translate the numbers between 1 and 100 into amplitudes for the surface bumps. As mentioned in the "Definition and measurement of orange peel" section, the Wave-Scan works by making a series of light flux measurements. Each of these measurements corresponds to an approximate slope at that point. The Wave-Scan stores these slopes in an array, which is then sampled at different intervals for each wavelength. The standardized variance of the slope value, $W_{va}, W_{vb}, \ldots, W_{ve}$, is then found for each of these wavelengths. Finally, each of these values is scaled by a logarithmic function[12] to place it into a perceptually uniform range from 1 to 100.

While it is impossible to obtain the original sample points from the final $W_a, W_b, \ldots, W_e$ values, it is possible to obtain $W_{va}, W_{vb}, \ldots, W_{ve}$ using the following approximation[12]:

$$W_v = 0.0015 * W^3 - 0.1288 * W^2 + 3.6117 * W - 3.4844 \quad (1)$$

where W_v is the variance, and W is the given input value for the Wave-Scan. This is performed for each $W_a, W_b, \ldots, W_e$. The square root of this is then taken to obtain the standard deviation for the slope at each wavelength. See Fig. 5 for the resulting correspondence between input W parameters and standard deviation of the surface slope. For our simulation, the average slope is assumed to be zero; the negative and positive slope values cancel each other out.

The next problem is to generate an actual surface from these parameters. Fortunately, the problem of generating visually plausible surfaces given a sparse set of parameters has already been solved in another area of computer graphics: landscape generation. In particular, the Fourier transform can be used to generate surfaces of a given wavelength and amplitude with convincing results.[22] This method takes a randomly seeded 2D height-field, and applies the Fourier transform to it to convert it to frequency space. Then, a 1/f noise filter is used to convert the random distribution into fractional Brownian motion. Brownian motion is a fractal distribution shown to mimic many natural fractal distributions.[23] After this filter has been applied, the inverse is then used to convert the data back into a spatial representation, resulting in a fractal height field.

This process gives a height field corresponding to a single wavelength value. This process is therefore repeated four more times to create a height field for each of the Wave-Scan values. These five height fields are then linearly combined to generate a final surface distribution (see Fig. 6). While this algorithm normally cannot be run in real-time, we observed that the height field of a single wavelength can be linearly scaled to create a surface of the same wavelength and bump distribution, but differing amplitude. Therefore, various surface distributions can be pre-computed, then scaled and combined with each other in real time. This allows a user to alter the Wave-Scan values and have a new surface presented to them interactively.

Rendering

The final step in the system is to display the orange peel surface to the user. Two rendering methods have

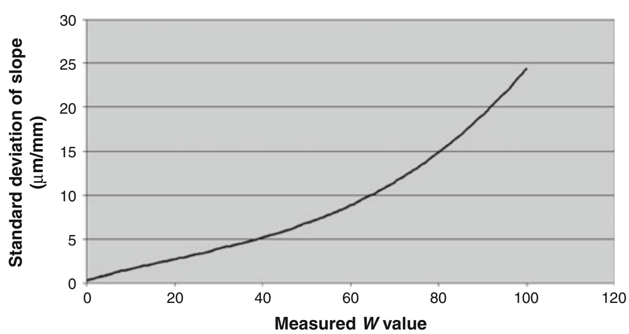

Fig. 5: The correspondence between the Wave-Scan W value and slope standard deviation. These values are used to generate a surface with the same characteristics that were measured by the instrument

been developed. The first is a straightforward rendering that shows how a uniform artificial light source reflects off the orange peel. This rendering system is meant to replicate the conditions used to take the pictures shown in Fig. 1. The second system is a more complete paint simulation system, which effectively adds orange peel effects to the metallic paint design system presented by Shimizu et al.[20]

In both systems the height map that is generated in the above section is converted into a normal map. Normal mapping is a widely accepted method for real time display of surfaces that do not require full geometric representations.[24] Usually, a normal map consists of three channels: the R 0–255 channel represents X ranging from -1 to 1, G maps to Y, and B maps to Z. However, we found that this mapping led to an unacceptable level of error due to orange peel having very slight variations in normal direction (all orange peel has its Z value at very near 1). This caused portions of the orange peel that should have had different normal directions to map to the same RGB values in the normal map. Therefore, the map was dynamically rescaled to have 0 represent the lowest possible value in the map (usually -0.1 or even less) and 255 the highest, to allow the system to always use the full 0–255 range of values. This change gave the normal map more acceptable results. The same remapping procedure can also be performed on floating point normal maps to obtain extremely high accuracy.

The first rendering system uses the normal map to make a reflection calculation involving an artificial light source, which is just a series of cylindrical white lights. There is no simulation of any of the diffuse underlying paint: it effectively just renders the gloss coat of the paint. The second rendering system adds the normal map calculation into the environment map lookups performed by Shimizu's system. Currently, only the gloss layers of the paint simulation have the orange peel added, although orange peel effects in the underlying diffuse layers could also be estimated using this method. However, this is not done because the Wave-Scan instrument does not effectively measure diffuse surfaces, as it makes the assumption that the laser light bounces off the surface as if that surface was a mirror: diffuse reflections add error to the calculation.

Fig. 6: Height fields for two separate wavelengths (left and center) and the composite surface that results from linearly combining them (right). Five different height fields (one for each wavelength) were added together to produce the final orange peel simulation

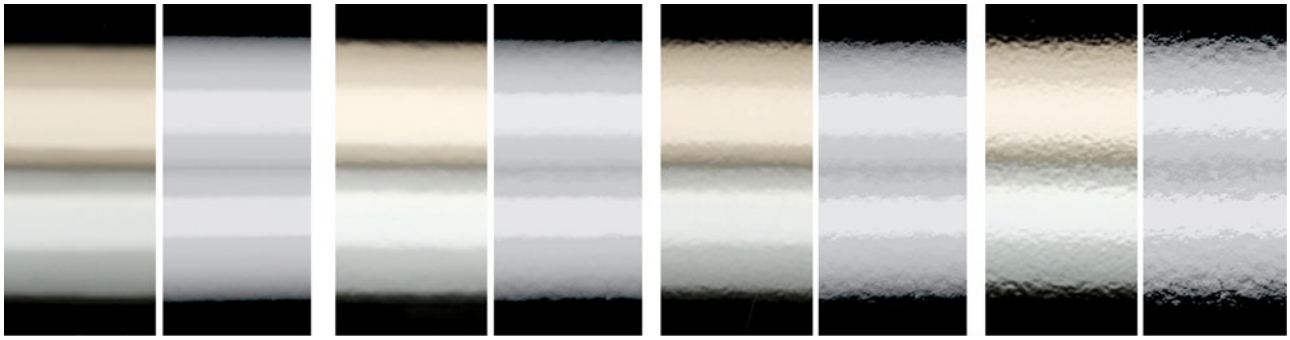

Fig. 7: Comparison between photographs of real orange peel (left in each pair) and simulations made using the same Wave-Scan values (right in each pair). Each image shows the reflection of a fluorescent tube fixture from a painted test panel

Fig. 8: Two images produced by the rendering system. On the left is an image of orange peel on metallic dark blue paint. On the right is the same orange peel on a less glossy sample of light blue paint; this makes the orange peel less noticeable

Figures 7 and 8 show the results produced by the complex rendering system. In Fig. 7a, comparison is made between pictures of real and simulated orange peel. Wave-Scan measurements were made on test panels with real orange peel, and these measurements were used to produce a simulation of that orange peel. Figure 8 shows how the orange peel looks on a three-dimensional shape with environment map based lighting.

Orange peel detection

The ability to predict if humans detect surface roughness can be beneficial in a number of industries, including the automotive, fabric, interior design, and computer graphics industries. There are two possible motivations for predicting visibility of surface roughness. The first is to design defect tolerances for phenomena such as orange peel. In the automotive area, viewers should ideally see the surface as perfectly smooth. However, in other industrial design applications, such as electronics packaging, adding roughness into a surface can prevent small scratches and normal wear and tear from being seen. Thus, effective design and control of such roughness would be a major benefit to these industries.

How we detect surface roughness is a combination of multiple factors: the lighting environment, surface micro and macro structures, and the human visual system. The software created in the above discussion combines a photorealistic lighting environment and physical surface to provide a simulation that humans should perceive in the same way as a real material sample. This section validates the system as a perceptual and design tool by both providing a mathematical framework for detection of simple, uniform surface roughness, as well as experimental validation through user testing.

An analytic model of surface roughness detection

This section presents a simple analytic model of surface roughness detection based on previous mathematical and perceptual studies. In order to simplify the factors involved, certain assumptions are made:

(1) The surface has bumps of a single, repeating amplitude, and wavelength.
(2) The lighting environment is composed of only two neutral colored lights (light and dark).
(3) The surface reflectance properties are modeled using the Phong reflection model (a single diffuse layer and a single specular layer).

This simple lighting and surface model is combined with previous research on human perception of contrast sensitivity and brightness[17] to predict if a viewer will be able to detect that a surface is or is not perfectly uniform.

The first step in this algorithm is to determine the difference in reflectance between points on the surface. In order to obtain the largest difference in perceived lightness, only two points need to be considered: the point that reflects the most light into the viewer's eye, and the point that reflects the least. Assuming that the user is viewing the surface at a 90 degree angle (perpendicular), these two points will be at the peak amplitude and halfway between that peak and the trough. These equate to the surface normal pointing directly toward the eye, and the normal oriented at the greatest angle with respect to the line of sight (see Fig. 9 for an example of such a surface). All other point pairs will have smaller perceptual differences between them.

Once the two points are determined, the reflectance, R, from each of these points is calculated:

$$R = \sum_{\text{lights}} k_{\text{d}}(\overline{L} \cdot \overline{N}) r_{\text{d}} + k_{\text{s}}(\overline{M} \cdot \overline{V})^{\alpha} r_{\text{s}} \tag{2}$$

This is the standard Phong model equation,[19] and is performed for each point. In this equation, $\overline{L}$ is the vector to the light source, $\overline{N}$ is the surface normal, $\overline{V}$ is the vector to the viewer, $\overline{M}$ is the half angle vector between $\overline{N}$ and $\overline{V}$, r_{d} is the diffuse reflectance; r_{s} is the specular reflectance, α is the specular coefficient, k_{d} is the relative amount of diffuse reflectance, and k_{s} is the relative amount of specular reflectance. For the purpose of analyzing the visibility of surface roughness, all the values except the surface normal are the same for each point. This yields the actual difference in lightness that the eye receives. However, in order to predict the detection of surface roughness, human contrast sensitivity must also be taken into account.

The two equations used in this research for contrast sensitivity are the contrast sensitivity function (CSF)[25]:

$$\text{CSF} = 2.6 * (0.0192 + 0.114 * f) * e^{-(0.114 * f)^{1.1}} \tag{3}$$

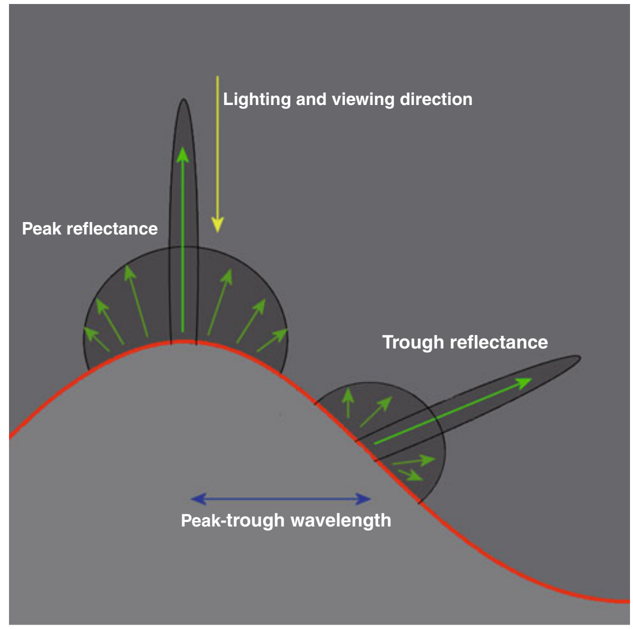

Fig. 9: An example of a single bump and the determination of the minimum and maximum reflection points along that bump. The reflectance on the left is the peak Phong reflectance: the point where the reflected light from the surface has the specular Phong lobe pointing directly toward the viewing direction. The reflectance on the right is the trough Phong reflectance: the point where the specular direction is furthest from the viewing direction and the diffuse reflectance is also at a minimum. The analytic model takes the physical lighting differences of these points into account along with human visual factors, lightness, and spatial acuity, to calculate a final perceived luminance difference between the points

where f is the frequency of the bumps; and just noticeable difference (JND)[18]:

$$\Delta L = 0.0594 * \left(1.219 + L_{\text{a}}^{0.4}\right)^{2.5} \tag{4}$$

where L_{a} is the brightness of the surround and ΔL is the least difference in luminance that would allow a human to distinguish the lighter foreground from the background. The JND formula takes both lightness adaptation and the varying response of humans to lightness levels into account. The brighter the surround, the more light that is required for a person to be able to perceive a difference between the brighter foreground and dimmer background.

The CSF is used as a normalization factor for the calculated Phong reflectances, and the result is divided by the JND formula to determine the number of luminance steps between the peak and trough values:

$$L_{\text{steps}} = ((R_{\text{p}} * \text{CSF}) - (R_{\text{t}} * \text{CSF}))/\Delta L \tag{5}$$

where R_{p} is the Phong reflectance at the peak reflectance, and R_{t} is the reflectance at the lowest reflectance.

A result greater than one means the bump can be detected, while a result of less than one means it cannot. Note that, as shown in Fig. 9, high wavelength amplitude causes the specular lobe of the trough point to completely miss the eye, causing sharp contrast between the peak and trough reflectance. Therefore, for higher levels of surface roughness, this model represents the ''worst-case'' scenario in terms of orange peel visibility: There are no obfuscating factors other than human visual acuity and the diffuse reflectance.

User evaluation

The orange peel simulation discussed in this article allows the analytic model derived above to be compared to the physical simulation of the actual bump distribution. Directly comparing the two can validate both the analytic model and the orange peel system.

In order to test the validity of these models, two experiments have been performed. The first test replicates the circumstances given in the analytic model using the orange peel system. The lighting is a checkerboard environment map with alternating light and dark checkers, and only a single bump wavelength is used. The amplitude of the Wave-Scan value determines the normal direction used in the Phong equation, and the result from that equation is used in the JND calculation. The CSF is determined by the wavelength frequency on the computer monitor combined with the viewing distance of the user. The purpose of this experiment is to verify that viewer detection of the simulated orange peel matches that predicted by the analytical algorithm derived above.

Each of the six subjects first took the Snellen visual acuity test to verify that he/she had 20/20 vision. Next, the subject was placed exactly one meter from a pre-calibrated computer monitor and adapted to the proper level of ambient illumination. The subject was then asked to view a series of 84 randomized images. Seventy five of the images had a varying wavelength and amplitude of orange peel, as well as one of three different checkerboard contrasts. Nine of the images were control images with no orange peel. Each subject was asked to say ''yes'' or ''no'' to whether or not they could see any roughness in the surface of the object displayed on the monitor. In the case of a ''yes'' answer, subjects were also asked how large the bumps were—either small, medium, or large. Subjects stated ''no'' to 51 of the 54 control images (94%).

The resulting answers were then averaged to determine at what amplitude a particular wavelength/contrast combination could be seen, and they were compared to the predicted results given by the analytical model. The results of this experiment are summarized in Fig. 10. Overall, the results match the analytical model very well. In nearly every wavelength/contrast combination, the subjects' ability to see the surface roughness averaged to within 10% of the expected orange peel wavelength value given by the analytical model.

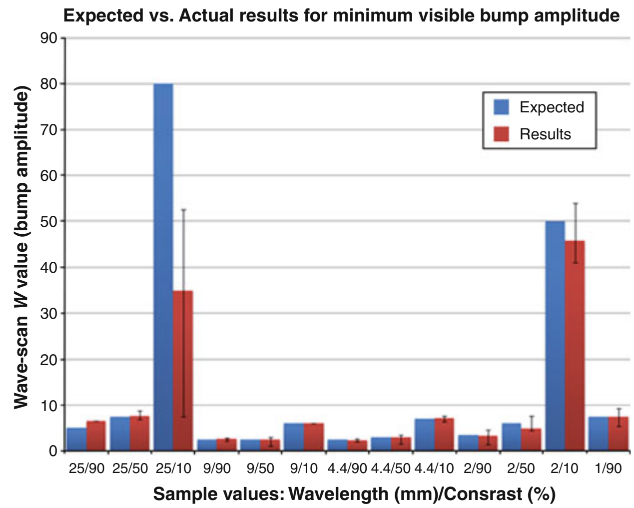

Fig. 10: A summary of the first orange peel experiment results. The blue bars (left of each pair) represent the expected value at which viewers should be able to perceive surface roughness predicted by the analytical equation. The red bar (right of each pair) represents the actual value at which the average user was able to perceive the surface roughness. The 1mm wavelength bars for 50% and 10% contrast are omitted as both the expectation and user result was that the roughness was completely imperceptible at all the tested wavelength amplitudes.

The only wavelengths that did not match within 10% were 25 mm for 90% contrast and 10% contrast. For 90%, the value was within 20% and small in terms of its absolute variance from the expected value. However, the results for 10% contrast were not within a reasonable range of the expected value, and the subjects had a very high variance in their answers. One possible explanation for this is that at such a high wavelength, other structural cues were present in the surface of the sample that were not predicted by the analytical model. Another possible explanation is that the users had such a difficult time determining this wavelength that many more samples need to be taken to gain a valid experimental result. In either case, more experiments will need to be performed to address the discrepancy.

In the second experiment, a more complicated lighting environment was used along with varying Wave-Scan parameters to show that the orange peel program can be utilized to prototype surface tolerances based on various factors such as lighting environment and paint color. Specifically, it shows that lighter paint colors obscure visibility of orange peel, and therefore less stringent surface quality requirements can potentially be set for lighter surfaces than dark surfaces.

For this test, each of eight subjects was asked to view a series of 55 images. Each image consisted of a shape half coated with one value of lightness/orange peel, and the other half coated with a different orange peel/lightness combination (see Fig. 11). All the samples had neutral coloration (gray). Each user was then asked which half of the sample was "rougher." They were not allowed to say the samples were the same.

The set of images consisted of 36 images with differing lightness and orange peel between the halves, 13 images of differing lightness but the same orange peel, and 6 images of the same lightness and different orange peel. All images that had differing orange peel and lightness had an image with the lightness–orange peel combination swapped (the same orange peel on different halves of the sample). This way, if users viewed the orange peel in the same way regardless of lightness, they should pick the lighter side 50% of the time, and the darker side 50% of the time. A result of greater than 50% selection of darker samples means that the lighter samples obfuscate the visibility of surface roughness.

The task of identifying precise orange peel roughness is fairly difficult. Users got 87.5% of the control cases correct, and 64.5% correct overall. In cases where the user guessed wrong, dark was selected 81% of the time and light 19% of the time. In cases where orange peel was the same, users selected dark 77% of the time and light 23% of the time. Figure 12 summarizes the results of this experiment. Note that the deviation between the dark and light selection does not overlap. This means that there is effectively no chance that the choice is simply a 50/50 guess ($p = 2.34 \times 10^{-5}$) confirming the original hypothesis that lighter shades of color obfuscate the visibility of surface roughness. However, in order to determine the exact deviation caused by lightness, more samples would need to be taken.

The results of these two experiments show that the orange peel program described in this article has the potential to be used to prototype surface roughness. When designing a new orange peel/color combination, a designer can use the orange peel program to study the effects of wavelength, amplitude, color shade, and lighting environment. Using the program in this manner could therefore facilitate faster and better design of automotive paints and spray systems.

Conclusions

In this article, we have added orange peel to the list of surface defects that have been successfully simulated using computer graphics. Orange peel is an important artifact that is present in most spray-painted surfaces. Faithfully reproducing orange peel adds realism to the aesthetic surfaces that are used in industrial and product design. Orange peel warps reflections from

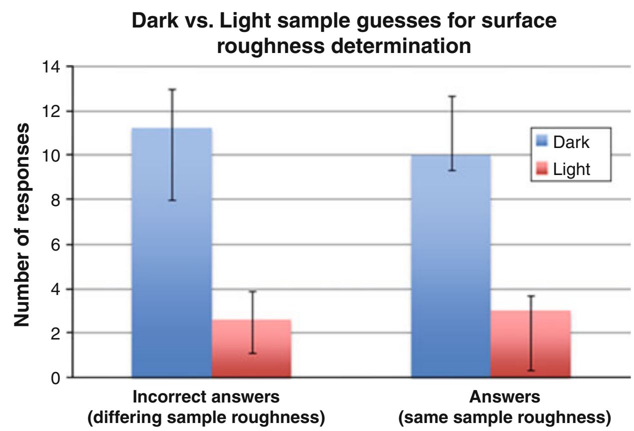

Fig. 12: The result of the second orange peel experiment. Left: The plot of average number of dark (left bar) vs light (right bar) incorrect guesses. Note that dark was chosen far more than light (approximately 81% of the time). Right: Plot of dark (left bar) vs light (right bar) guess when orange peel was actually the same. Once again, dark was seen as being rougher a significant percent of the time (77%)

Fig. 11: Two images from the second orange peel experiment. Left: An image with differing lightness and orange peel values. Right: One of the control images with varying orange peel values but the same lightness

these surfaces, and so computer graphic renderings which expertly reproduce these reflections should include this distortion. Carefully simulating orange peel also makes it possible to develop training tools and metrics that can be used to improve industrial quality control procedures.

This research has introduced another virtual measurement and rendering system based on an industry standard measurement device. Specifically, an orange peel measurement instrument has been used to generate a virtual ''landscape'' that matches its measurements. This leads to precise, photorealistic renderings based on measurements given by the device. This process has two benefits: First, it allows renderings of the effects these devices measure to be very accurate. An industry professional viewing these renderings can have confidence that the rendering corresponds to how the material being simulated actually looks. Second, it allows industry professionals (or graphic artists) to prototype materials based on these measurements. Since the professionals are already familiar with use of these devices, this forms an intuitive way for them to generate a material they would like, and view it without having to create a physical prototype. This saves time and material costs.

In addition to generating images for materials that were intended to be measured by these instruments, the use of industry devices creates some interesting possibilities for the computer graphics community. For instance, the Wave-Scan instrument is capable of measuring any small scale, isotropic bump distribution to a very high degree of accuracy. This means that it could be used to measure and render a number of real world materials, such as leather, certain other paints (paint applied with a paint roller, for instance), some rubbers and plastics.

Finally, this article has introduced an analytical formula for predicting the visibility of small scale roughness effects such as orange peel. That formula can be used along with the orange peel program to predict and design tolerances for surface defects in the automotive industry. User experiments were performed to show the viability of both the analytic model and orange peel algorithm for use as design tools.

It seems very promising that more such devices will be made available in the future by industries that wish to accurately measure their own products. This study paves the way for the computer graphics community to leverage these devices in the generation of accurate renderings of the materials which these devices measure. In addition, the concrete correlation between the industry standards and graphical models allows the graphics community to create programs that professionals can easily use to design and generate prototypes for use in their own field.

Acknowledgments The authors would like to thank Byk-Gardner for providing us with Wave-Scan data and documents. We would also like to thank Ford Motor Company for providing sample orange peel panels and photographs. This study was partially funded by NSF Grant IIP-0438693.

References

1. Dorsey, J, Hanrahan, P, ''Modeling and Rendering of Metallic Patinas.'' *SIGGRAPH: ACM SIGGRAPH 2005 Courses*, p. 2, 2005
2. Dorsey, J, Hanrahan, P, ''Modeling and Rendering of Metallic Patinas.'' *Proceedings of SIGGRAPH*, pp. 387–396, 1999
3. Chang, Y, Shih, Z, ''Physically-Based Patination for Underground Objects.'' *Comput. Graph. Forum*, **19** 109–117 (2000)
4. Dorsey, J, Edelman, A, Jensen, H, Legakis, J, Pedersen, H, ''Modeling and Rendering of Weathered Stone.'' *Proceedings of SIGGRAPH*, pp. 225–234, 1999
5. Cutler, B, Dorsey, J, McMillan, L, Muller, M, Jagnow, R, ''A Procedural Approach to Authoring Solid Models.'' *Proceedings of SIGGRAPH*, pp. 302–311, 2002
6. Paquette, E, Poulin, P, Drettakis, G, ''Surface Aging by Impacts.'' *Proceedings of Graphics Interface*, pp. 175–182, 2001
7. Merillou, S, Dischler, J, Ghazanfarpour, D, ''Surface Scratches: Measuring, Modeling and Rendering.'' *Vis. Comput.*, **17**, 30–45 (2001)
8. Bosch, C, Pueyo, X, Merillou, S, Ghasanfarpour, D, ''A Physically Based Model for Rendering Realistic Scratches.'' *Comput. Graph. Forum*, **23** 361–370 (2004)
9. Paquette, E, Poulin, P, Drettakis, G, ''The Simulation of Paint Cracking and Peeling.'' *Proceedings of Graphics Interface*, pp. 59–68, 2002
10. Iben, H, O'Brien, J, ''Generating Surface Crack Patterns.'' *Proceedings of SIGGRAPH*, pp. 177–185, 2006
11. Hsu, S, Wong, T, ''Simulating Dust Accumulation.'' *Comput. Graph. Appl.,* **15** 18–22 (1995)
12. Kigle-Bockler, G, ''Surface Quality Control on High and Medium Gloss Surfaces: Wave-Scan Dual.'' *4th Wave-Scan User Meeting*, 2006
13. Konieczny, J, ''Application, Rendering and Display of Automotive Paint.'' Ph.D. thesis, University of Minnesota, 2009
14. Ershov, S, Khodulev, A, Kochin, K, ''Simulation of Sparkles in Metallic Paints.'' *Proceeding of Graphicon*, pp. 121–128, 1999
15. Westlund, H, Meyer, G, ''Applying Appearance Standards to Light Reflection Models.'' *Proceedings of SIGGRAPH*, pp. 501–551, 2001
16. Dumont-Becle, P, Ferley, E, Kemeny, A, Michelin, S, Arqus, D, ''Multi-Texturing Approach for Paint Appearance Simulation on Virtual Vehicles.'' *Proceedings of the Driving Simulation Conference*, pp. 123–133, 2001
17. Daly, S, ''The Visible Differences Predictor.'' In: Watson, AB (ed.) *Digital Images and Human Vision*, pp. 179–206. MIT Press, Cambridge, MA, 1993
18. Ward, G, ''A Contrast-Based Scalefactor for Luminance Display.'' In: Heckbert, PS (ed.) *Graphics Gems IV*, pp. 415–421. Academic Press, Boston, MA, 1994
19. Phong, B, ''Illumination for Computer Generated Pictures.'' *Commun. ACM*, **18** 311–317 (1975)
20. Shimizu, C, Meyer, G, Wingard, J, ''Interactive Goniochromatic Color Design.'' *Color Imaging Conference*, pp. 16–22, 2003
21. Stover, J, *Optical Scattering: Measurement and Analysis*. SPIE Optical Engineering Press, 1995

22. Mastin, G, Watterberg, P, Mareda, J, ''Fourier Synthesis of Ocean Scenes.'' *IEEE Comput. Graph. Appl.*, **7** 16–23 (1987)
23. Gondek, J, ''Fractal Landscape Synthesis in Computer Graphics.'' B.S. thesis, University of Oregon, 1992
24. Blinn, J, ''Simulation of Wrinkled Surfaces.'' *Proceedings of SIGGRAPH*, pp. 286–292, 1978
25. Mannos, J, Sakrison, D, ''The Effects of a Visual Fidelity Criterion on the Encoding of Images.'' *IEEE Trans. Inform. Theory*, **20** 525–535 (1974)