

MITSUBISHI ELECTRIC RESEARCH LABORATORIES
<http://www.merl.com>

Harnessing Real-World Depth Edges with Multiflash Imaging

K-H Tan, R.S. Feris, M. Turk, J. Kobler, J. Yu, R. Raskar

TR2005-067 December 2005

Abstract

A method for capturing geometric features of real-world scenes relies on a simple capture setup modification. The system might conceivably be packed into a portable, self-contained device.

IEEE Computer Graphics and Applications

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Copyright © Mitsubishi Electric Research Laboratories, Inc., 2005
201 Broadway, Cambridge, Massachusetts 02139

Harnessing Real-World Depth Edges with Multiflash Imaging

Ramesh Raskar
Mitsubishi Electric Research Labs

Kar-Han Tan
Epson Palo Alto Lab

Rogerio S. Feris and Matthew Turk
University of California, Santa Barbara

James Kobler
Harvard Medical School

Jingyi Yu
Massachusetts Institute of Technology

A method for capturing geometric features of real-world scenes relies on a simple capture setup modification. The system might conceivably be packaged into a portable, self-contained device.

Sharp discontinuities in a depth map, or depth edges, are extremely useful 2.5D entities. The ability to localize and highlight depth discontinuities makes it possible to produce stylized photography and videos, which are useful in technical applications such as medical imaging¹ and nonphoto-realistic rendering (NPR).² In addition, depth edges are important low-level features for many computer vision tasks, such as visual recognition.³

Figure 1 shows an example of automatic depth-edge-based stylization, where a car engine is imaged under diffused lighting and shown in a stylized form, accentuating the shape-conveying depth edges even when they are low-contrast edges, while deemphasizing visual clutter, like those introduced


by rust. The resulting image resembles drawings from car repair manuals with an added degree of realism and authenticity as it was captured from a real engine.

When a rich 3D model of a scene is available, identifying and localizing depth discontinuities is a relatively well-understood task. Extending this approach to real scenes by first capturing 3D models, however, remains difficult. Our multiflash imaging method bypasses 3D geometry acquisition and directly acquires depth edges from images. In the place of expensive, elaborate equipment for geometry acquisition,⁴ we use a camera with multiple strategically positioned flashes. Instead of having to estimate the full 3D coordinates of points in the scene (using, for example, 3D cameras) and then look for depth discontinuities, our technique reduces the general 3D problem of depth edge recovery to one of 2D intensity edge detection. Our method could, in fact, help improve current 3D cameras, which tend to produce incorrect results near depth discontinuities.

Exploiting the imaging geometry for rendering provides a simple and inexpensive solution for creating stylized images from real scenes. We believe that our camera will be a useful tool for professional artists and photographers, and we expect that it will also let the average user easily create stylized imagery.

Creating stylized imagery from photographs, rather than 3D geometric models, has recently received a great deal of attention. The majority of the available techniques for image stylization involves processing a single image as the input, using image processing and computer vision techniques like morphological operations, image segmentation, and edge detection. Some methods aim for stylized depiction, while others enhance legibility. Animators have used interactive techniques for stylized rendering—such as rotoscoping with real video footage—to create animation like the groundbreaking *Waking Life* (<http://www.wakinglifemovie.com/>) and *Avenue Amy* (http://www.curiouspictures.com/shows/clips/ave_amy.html). Multiflash imaging has the potential to automate tasks where meticulous manual operation was previously required and to make it easier for filmmakers to blur the line between photographic and stylized material.

1 Car engine (a) imaged under diffused lighting, (b) stylized using depth edges computed with our technique, (c) with increased brightness, (d) and with histogram equalized.



Basics

Our multiframe imaging method is motivated by the observation that when a flashbulb (close to the center of projection of the camera) illuminates a scene during image capture, thin slivers of cast shadow are created at depth discontinuities. Moreover, the shadows' positions are determined by the relative position of the camera and the flashbulb: When the flashbulb is on the right, shadows are created on the left, and so on. Thus, if we can shoot a sequence of images in which different light sources illuminate the subject from various positions, we can use the shadows in each image to assemble a depth edge map.

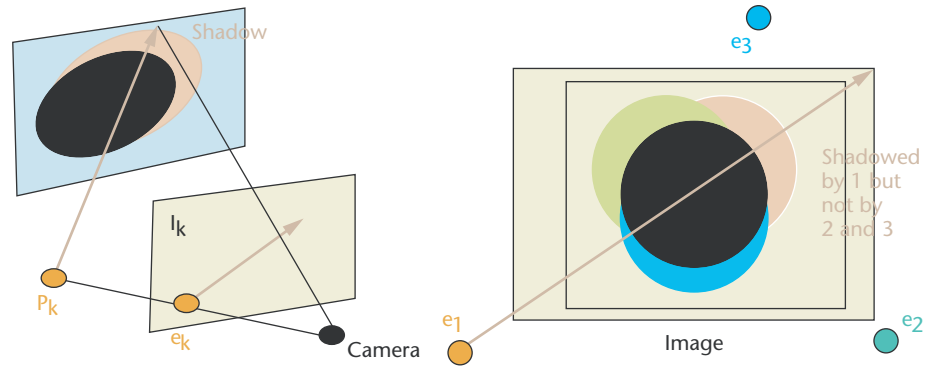
Imaging geometry

To capture the intuitive notion of how the positions of the cast shadows are dependent on the relative position of the camera and light source, we examine the imaging geometry, illustrated in Figure 2. Adopting a pinhole camera model, the projection of the point light source at P_k is at pixel e_k on the imaging sensor. We call this image of the light source the light epipole. The images of (the infinite set of) light rays originating at P_k are in turn called the epipolar rays originating at e_k . We then define the term *depth edges* as the 2D images of depth discontinuities.


Removing and detecting shadows

Our approach to reliably remove and detect shadows in the images is to strategically position lights so that every point in the scene that is shadowed in some image is also imaged without being shadowed in at least one other image. We can achieve this by placing lights so that for every light there is another light on the camera's opposite side so that all depth edges are illuminated from two sides. Also, by placing the lights close to the camera, we minimize changes across images due to effects other than shadows.


To detect shadows in each image, we first compute a shadow-free image, which can be approximated with the maximum composite image (MAX image), which is an image assembled by choosing at each pixel the maximum intensity value from among the image set. We then compare the shadow-free image with the individual shadowed images. In particular, for each shadowed image we compute the ratio image by performing a pixel-wise division of the intensity by that of the MAX image. The ratio image is close to 1 at pixels that are not shadowed, and close to 0 at pixels that are shadowed. This accentuates the shadows and also removes intensity transitions due to surface material changes.



(a)



(b)



(c)

2 Building a multiframe camera. (a) Imaging geometry. (b) Hardware schematic. (c) Prototype based on a Canon G2 camera.

Algorithm

Codifying these ideas, we arrive at the following algorithm:

Given n light sources positioned at $P_1, P_2 \dots P_n$,

- Capture n pictures $I_k, k = 1 \dots n$ with a light source at P_k
- For all pixels $x, I_{\max}(x) = \max_k(I_k(x)), k = 1 \dots n$
- For each image k , create ratio image R_k , where $R_k(x) = I_k(x)/I_{\max}(x)$
- For each image R_k , traverse each epipolar ray from epipole e_k , find pixels y with step edges with negative transition, and mark pixels y as a depth edge


Building multiframe cameras

We built a multiframe camera using a 4-megapixel Canon Powershot G2, as shown in Figure 2. A microcontroller board triggers sequentially the four flashes mounted around the camera. The board synchronizes the flashes to the image capture process by sensing the flash trigger signal from the camera hot shoe.

We have published elsewhere descriptions of more advanced prototypes and discussions on the finer points of multiframe imaging.² In this article we shall illustrate the usefulness of multiframe imaging in a number of different applications: NPR, medical imaging, biological illustrations, and visual recognition.


NPR with depth edges

Multiframe imaging can address two important issues in NPR: detecting shape contours that should be enhanced and identifying features that should be sup-



3 NPR with depth edges. (a) Bone scene, rendered with signed contour style. (b) Flower scene, rendered by removing textures and overlaying depth edges, (bottom right) with contour colors assigned with foreground object colors, (bottom left) with signed edges highlighted by modulating image intensity around depth edges.

4 Change detection (counterclockwise from top left): reference image, changed image, detected changed region, and stylized scene change depiction.



pressed.⁵⁻⁷ Depth edges correspond to physical object shape contours and silhouettes. In addition to providing the 2D location of such contours, depth edges are also signed in the sense that at a depth edge we know which side is the foreground (positive sign) and background (negative sign), since we know where the shadow appeared. This 2.5D nature of depth edges allows the design of rendering styles not possible or difficult in the absence of high-quality 3D models. Figure 3 shows examples of rendering styles.

Some static illustrations demonstrate action—for example, changing oil in a car—by making moving parts in the foreground brighter. Foreground detection via

intensity-based schemes, however, is difficult when the colors are similar and texture is lacking—for example, detecting hand gesture in front of other skin-colored parts (see Figure 4). We take two separate sets of multiflash shots, without and with the hand in front of the face, to capture the reference and changed scene. We note that any change in a scene is bounded by new depth edges introduced. Without explicitly detecting foreground, we highlight interiors of regions that contribute to new depth edges. We create a gradient field where pixels marked as depth edges in the changed scene, but not in reference, are assigned a unit magnitude gradient. The orientation matches the image space normal to the depth edge. The gradient at other pixels is zero. The reconstructed image from 2D integration is a pseudo depth map. We threshold this map at 1.0 to get the foreground mask, which is brightened.

Medical applications

In many medical applications like minimally invasive surgery with endoscopes, it's often difficult to capture images that convey the 3D shape of the organs and tissues under examination. Perhaps for the same reason, medical textbooks and articles frequently resort to hand-drawn illustrations when depicting organs and tissues. In the sections that follow we show the use of multiflash imaging to address this problem.

A multiflash imaging system captures additional shape information compared to traditional cameras and therefore has the potential to enhance visualization and documentation in surgery and pathology. The raw shadowed images can be processed to create finely detailed images comparable to medical illustrations or to enhance edge features for quantitative measurement. Alternatively, the shadowed images can be combined to generate shadow-free images, which are often desirable for documentation of specimens in the field of pathology.

Surgery

Most endoscopic procedures are now performed with the surgeon observing monitor displays rather than the actual tissue. This affords the possibility of interposing image manipulation steps, which, if they can run close to real time, can enhance the surgeon's understanding. Depth perception is an obvious deficit when using monocular endoscopes. Researchers have explored 3D imaging using stereoscopic methods explored with mixed results. A 1999 study found that stereoendoscopic viewing was actually more taxing on the surgeons than monocular viewing.⁸ Structured lighting is also under investigation as a means for calibrating endoscopic images, but this technique does not enhance 3D structures in real time.⁹

Application of enhanced shadow information to augment surgical perception has not been exploited previously. Shadows normally provide clues about shape, but the circumferential (ring light) illumination provided by most laparoscopes diminishes this information. Similarly, the intense multisource lighting used for open procedures tends to reduce strong shadow effects. Loss of shadow information might make it difficult to appreciate the shapes and boundaries of structures and thus more

difficult to estimate their extent and size. It could also make it more difficult to spot a small protrusion, such as an intestinal polyp, if no clear color differences exist. The ability to enhance the borders of lesions so that they can be measured will become more useful as endoscopes begin to incorporate calibrated sizing features.

Multiflash imaging with endoscopes

The simplest way to implement multiflash imaging in endoscopes is to use multiple instruments, where instead of inserting one endoscope, three are inserted. The middle instrument acts as the camera while the two on the side act as light sources. By synchronizing the light sources with the image-capture process for the middle endoscope, the entire setup would act as a multiflash camera. While this approach might involve inserting more imaging instruments, it's a way to systematically illuminate the subject and potentially reduce the amount of adjustments required during an operation to produce images that convey the required 3D information.


In many scenarios, it's more useful to have a single instrument capable of multiflash imaging. For example, in situations that require flexible endoscopes, it could be difficult or impossible to insert and align multiple flexible light sources with the endoscope. Fortunately, it's possible to implement multiflash imaging on a single instrument with our method because the light sources can be placed near the camera.¹ This allows for compact designs suited for use in tightly constrained spaces, unlike many traditional 3D shape recovery methods where the imaging apparatuses must be placed at large distances apart.

For proof-of-concept endoscopic imaging, we took advantage of the illumination system in the Wolf Lumina, a standard rod lens laryngeal endoscope. The illumination bundle in this endoscope bifurcates at the distal tip into two illumination ports, which are located on each side of the imaging lens. Because the illumination fibers travel largely in separate bundles, we achieved independent illumination of the two ports by selectively illuminating different halves of the bundle at their proximal end. Thus, we converted the endoscope to a multiflash system consisting of two flash sources 180 degrees apart.

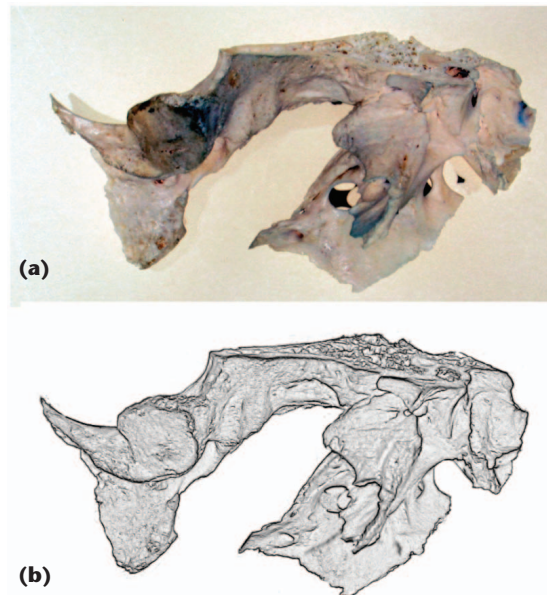
We bench-tested the multiflash endoscope with biological specimens that simulated the examination of the human larynx. We were particularly interested to see how well the system performed in detecting small surface lesions, which commonly have depth discontinuities. The processed images demonstrated the system's capability to find edges of some simulated lesions (see Figure 5). We also found that the rounded shapes and the translucency of internal surfaces can make it difficult to cast useful shadows, problems that we will address in the future by experimenting with illumination parameters.

Pathology

Pathology departments document surgical and autopsy specimens. Systems for photographing such specimens involve special methods for eliminating unwanted shadows, usually by placing the specimens on glass plates suspended over black cavities. Using the multi-



5 (a) Enhanced endoscope. (b) Input image and image with depth edges superimposed.



6 Biological illustrations with multiflash imaging. (a) Sphenoid bone. (b) Raw output from depth edge detection.


flash system and processing to obtain the MAX composite image produces a view in which almost all shadows are eliminated.

Medical and biological illustration


Often, it's desirable to generate black-and-white illustrations of medical and natural history specimens in which salient details are emphasized and unnecessary clutter is omitted.¹⁰ The most important details to emphasize are those that convey the object's shape. Typically shape is conveyed by emphasizing edges and using stippling for shading. This type of illustration is seen less frequently nowadays because of the expense involved in having artists create these graphics.

In our experiments with multiflash imaging, we have observed that the depth edge confidence maps frequently resemble hand-drawn sketches. At the same time, since they are created from photographs, the maps retain a high degree of realism. We felt that multiflash photography could make it faster, easier, and less expensive for artists and researchers to create medical illustrations.

Figure 6 shows some results generated with this camera. We showed the results to medical professionals at the Massachusetts General Hospital and received positive feedback on the usefulness of these images. In addition, they also found the MAX composite image to be



7 (a) Letter “R” in ASL alphabet. (b) Canny edges. (c) Depth edges obtained with our multiflash technique.



8 From left to right: input image, Canny edges, and depth edges.

more useful, as it’s difficult to take shadow-free images with ordinary cameras. Anatomists and medical illustrators who have seen our system are mostly interested in the depth-edge confidence image, which can be easily converted into a detailed black-and-white drawing.

Visual object recognition

The accurate detection of depth edges is also useful in many object recognition tasks and human–computer interaction applications. In this section, we show that our method can be reliably applied in automatic sign language analysis, particularly considering the finger spelling recognition problem.

Sign language is the primary communication mode used by most deaf people. It consists of two major components: word level sign vocabulary, where gestures communicate the most common words, and finger spelling, where the fingers on a single hand spell out more obscure words and proper nouns, letter by letter.

We address the finger spelling recognition component, showing the usefulness of depth edges in this task.

Although researchers have spent great effort in the past decade to develop automatic finger spelling recognition systems, most successful approaches are based on instrumented gloves, which are considered cumbersome for the user and are often expensive. In general, nonintrusive vision-based methods, while useful for recognizing a small subset of convenient hand configurations, are limited to discriminate configurations with high amounts of finger occlusions—a common scenario in most finger spelling alphabets. In such cases, traditional edge detectors or segmentation algorithms fail to detect important internal edges along the hand shape (due to the low intensity variation in skin color), while keeping edges due to nails and wrinkles, which may confound scene structure and the recognition process. Also, some signs might look similar to each other, with small differences on finger positions, thus posing a problem for appearance-based approaches.

Finger spelling recognition

We show that depth edges could be used as a signature to reliably discriminate among complex hand configurations in the American Sign Language (ASL) alphabet, which would not be possible with current glove-free vision methods.

Figure 7 shows a comparison of the Canny method, which is a standard-intensity edge detector, and depth edges extracted with our method for the letter “R” of the ASL alphabet. Important internal edges are missing in the Canny method, while unwanted edges due to wrinkles and nails are present.

We realized that depth edges are good features to discriminate among signs of finger spelling alphabets. Even when the signs look similar (for example, letters “E”, “S,” and “O” in ASL alphabet), the depth edge signature is quite discriminative (see Figure 8). This poses an advantage over vision methods that rely on appearance or edge-based representations. However, our method does not detect edges in finger boundaries with no depth discontinuity. It turns out that this is helpful to provide more unique signatures for each letter.

To quantitatively evaluate the advantages of using depth edges as features for finger spelling recognition, we considered an experiment with the complete ASL alphabet—except for letters “J” and “Z,” which require motion analysis to be discriminated. We collected a small set of 72 images using our multiflash camera (three 640×480 -resolution images per letter, taken at different times). The images showed variations in scale, translation, and rotation. The background was plain, with no clutter, since we wanted to show the importance of obtaining clean edges in the interior of the hand. Textured but flat and smooth backgrounds did not affect our method, but did make an edge detection approach (used for comparison) much more difficult.

For object classification, we used a depth edge shape descriptor similar to shape context matching, which is invariant to object translation and scaling.¹¹ For comparison, we also considered shape descriptors based on Canny edges. We obtained the recognition rate using a

leave-one-out scheme in the collected data set. Our approach achieved 96 percent correct matches, compared with 88 percent when using Canny edges.

Rebollar mentioned in his work that letters “R,” “U,” and “V” represented the worst cases, as their class distributions overlap significantly.¹² Figure 9 shows these letters and their corresponding depth edge signatures. They are easily discriminated with our technique. In the experiment described previously, the method based on Canny edges fails to discriminate them.

We collected all the images in our experiment from the same person. We plan to build a more complete database with different signers. We believe that our method will better scale in this case, due to the fact that texture edges (for example, wrinkles, freckles, and veins) vary from person to person but are eliminated in our approach. Also, shape context descriptors have proven useful for handling hand shape variation from different people.¹¹

For cluttered scenes, our method also eliminates all texture edges, thus considerably reducing clutter (see Figure 10).

For segmented hand images with resolution 96×180 , the computational time required to detect depth edges is 4 ms on a 3-GHz Pentium IV. The shape descriptor computation requires on average 16 ms. Thus, our method is suitable for real-time processing. For improving hand segmentation, depth edges could be computed in the entire image. In this case, the processing time for 640×480 images is 77 ms.


We are currently exploiting a frequency division multiplexing scheme, where flashes with different colors (wavelength) are triggered simultaneously. We hope this will allow for efficient online tracking of depth edges in sign language analysis.

Additional finger-spelling issues

What if there are no cast shadows due to lack of background? In these cases only the outermost depth edge, the edge shared by the foreground and distant background, is missed in our method. This could be detected with a foreground and background estimation technique. Let I_0 be the image of the scene taken just with ambient illumination. Then the ratio I_0/I_{\max} (image acquired with no flash over MAX composite of flash images), is near 1 in the background and close to zero in the foreground interior. This is because the faraway background is not affected by the flash, but the object will be brighter.

Another solution is to use our method to detect internal edges in the hand, while using traditional techniques (such as skin color segmentation or background subtraction) to obtain the external hand silhouette.

We noticed that depth edges might appear or disap-



9 Letters “R,” “U,” and “V.” The use of a depth edge signature can easily discriminate them.



10 Cluttered scene using (a) Canny edges and (b) depth edges.

pear with small changes in viewpoint (rotations in depth). We believe this might be a valuable cue for hand pose estimation.

A common thread in recent research on pose estimation involves using a 3D model to create a large set of exemplars undergoing variation in pose, as training data. Pose estimation is an image retrieval problem in this data set. We could use a similar approach to handle out-of-plane hand rotations. In this case, a 3D hand model would store a large set of depth edge signatures of hand configurations under different views.

We have not seen any previous technique that can precisely acquire depth discontinuities in complex hand configurations. In fact, stereo methods for 3D reconstruction would fail in such scenarios, due to the textureless skin color regions as well as low intensity variation along occluding edges.

Word level sign language recognition could also benefit from our technique, due to the high amounts of occlusions involved. Flashes in our setup could be replaced by infrared lighting for user-interactive applications. We are currently evaluating our method in a large database with different signers. We also plan to address the problem of continuous signing in dynamic scenes, using colored flashes.

Future work

Many hardware improvements are possible. The depth-edge extraction scheme could be used for spectrums other than visible light that create shadows, for example, in infrared, sonar, x-ray, and radar imaging. We described a video-rate camera for detecting depth edges in dynamic scenes described elsewhere, and we plan to build prototypes with infrared light sources invisible to humans so the resulting flashes are not distracting.² We could use a

frequency division multiplexing scheme to create a single shot multiflash photograph. ■

References

1. K.-H. Tan et al., "Shape-Enhanced Surgical Visualizations and Medical Illustrations with Multi-Flash Imaging," *Proc. Int'l Conf. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Springer-Verlag, 2004, pp. 438-445.
2. R. Raskar et al., "Non-photorealistic Camera: Automatic Stylization with Multi-Flash Imaging," *Proc. Siggraph*, ACM Press, 2004, pp. 679-688.
3. R. Feris et al., "Exploiting Depth Discontinuities for Vision-Based Fingerspelling Recognition," *Proc. IEEE Workshop Real-Time Vision for Human-Computer Interaction*, vol. 10, IEEE CS Press, 2004.
4. D. Scharstein and R. Szeliski, "High-Accuracy Stereo Depth Maps Using Structured Light," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, vol. 1, IEEE CS Press 2003, pp. 195-202.
5. D. DeCarlo and A. Santella, "Stylization and Abstraction of Photographs," *Proc. Siggraph*, ACM Press, 2002, pp. 769-776.
6. J. Wang et al., "Video Tooning," *Proc. Siggraph*, ACM Press, 2004, pp. 574-583.
7. T. Saito and T. Takahashi, "Comprehensible Rendering of 3D-Shapes," *Proc. Siggraph*, ACM Press, 1990, pp. 197-206.
8. M. Mueller et al., "Three-Dimensional Laparoscopy. Gadget or Progress? A Randomized Trial on the Efficacy of Three-Dimensional Laparoscopy," *Surgical Endoscopy*, vol. 13, Springer-Verlag, 1999, pp. 1432-2218.
9. D. Rosen et al., "Calibrated Sizing System for Flexible Laryngeal Endoscopy," *Proc. 6th Int'l Workshop: Advances in Quantitative Laryngology Advances in Quantitative Laryngology, Voice and Speech Research*, G. Schade et al., eds., Verlag, 2003.
10. E.R.S. Hodges, *The Guild Handbook of Natural History Illustration*, John Wiley & Sons, 2003.
11. S. Belongie, J. Malik, and J. Puzicha, "Shape Matching and Object Recognition Using Shape Contexts," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 4, 2002, pp. 509-522.
12. J. Rebollar, R. Lindeman, and N. Kyriakopoulos, "A Multi-Class Pattern Recognition System for Practical Fingerspelling Translation," *Proc. Int'l Conf. Multimodal Interfaces*, IEEE CS Press, 2002, p. 185.



Ramesh Raskar is a research scientist at Mitsubishi Electric Research Labs. His research interests include computer vision and graphics, projective geometry, nonphotorealistic rendering, and intelligent user interfaces. Raskar received a PhD from the University of North Carolina at Chapel Hill in computer science. He is a member of the IEEE, Society for Information Display, and ACM. Contact him at raskar@merl.com.



Kar-Han Tan is a research scientist at the Epson Palo Alto Lab. His research interests include human and computer vision, graphics, advanced sensors and displays, and intelligent human-machine interaction. Tan has a PhD from the Univ. of Illinois at Urbana-Champaign in computer science. He is a member of the IEEE and ACM. Email him at tankh@vision.ai.uiuc.edu.




Rogerio S. Feris is a PhD candidate in computer science at the University of California, Santa Barbara. His research interests include computer vision and graphics for human-computer interaction. Feris received a BS in computer engineering from the Federal University of Rio Grande, Brazil, and an MS in computer science from the University of Sao Paulo, Brazil. Contact him at rferis@cs.ucsb.edu.



Matthew Turk is an associate professor in the Computer Science Department and the Media Arts and Technology Program at the University of California, Santa Barbara. He codirects the Four Eyes Lab. Turk received a BS from Virginia Tech in electrical engineering, an MS from Carnegie Mellon University in electrical and computer engineering, and a PhD from the MIT Media Lab. Contact him at mturk@cs.ucsb.edu.



James Kobler is director of the H.P. Mosher Laryngological Research Laboratory at the Massachusetts Eye and Ear Infirmary, a faculty member of the Dept. of Otolaryngology at Harvard Medical School, and adjunct faculty in the Division of Health Sciences and Technology at MIT. His research interests include endoscopic imaging, voice neural prostheses, and vocal fold regeneration. Kobler received a PhD in neurobiology from the University of North Carolina at Chapel Hill. Contact him at james_kobler@meei.harvard.edu.



Jingyi Yu is a PhD candidate in electrical engineering and computer science at MIT and is finishing his dissertation at the University of North Carolina at Chapel Hill. His research interests include image-based modeling and rendering, image processing, and computational camera models. Yu received a BS in electrical and applied science and applied mathematics from the California Institute of Technology and an MS in electric engineering and computer science from MIT. Contact him at Jingyi@graphics.csail.mit.edu.