



Using Global Illumination in Volume Visualization of Rheumatoid Arthritis CT Data

Lin Zheng, Abhijit J. Chaudhari, Ramsey D. Badawi, and Kwan-Liu Ma

University of California at Davis

Proper lighting in rendering is essential for visualizing 3D objects, but most visualization software tools still employ simple lighting models. Over the past few years, advances in graphics hardware and the introduction of global-illumination approximations for volume visualization¹⁻³ have begun to make advanced illumination feasible for visualization. Here, we show how we used global illumination in volume visualizations of x-ray computed tomography (CT) wrist data for evaluating responses to treatment of rheumatoid arthritis (RA).⁴

In the task we're studying, physicians or clinical researchers identify and examine bone erosions in RA patients' wrists and hands. To study the disease's progression, multiple CT scans are made over the treatment period. Even though volume rendering has been around for many years, the common practice is still to study individual 2D slices of the data, which leads to two problems. First, in 2D, it's hard to determine the area and extent of erosion, which is a 3D feature. Second, in 2D, it's also hard to identify and compare the same bone erosions from different scans.

Although 3D visualization clearly should be used for such tasks, current volume visualization tools raise several issues. For example, local illumination might erroneously increase the prominence of certain surface features, which affects interpretation and diagnosis. To address this issue, we conducted a study with experts who were familiar with the disease process and medical-image interpretation. In particular, we studied tasks for which we believe proper illumination can enhance user performance. Our results indicate that interactive visualization with global illumination helped the experts derive more accurate interpretations of the image data. With clinical needs and the recent advancement of volume visualization technology,

this study is timely and points the way for further research.

Visualizing Rheumatoid Arthritis

Although decreased bone density (osteolysis) as measured in CT images has many causes, such as malignancies, metabolic processes, and chronic inflammation, the most common cause is RA. The disease targets both trabecular and cortical bone. Healthy bone has a balance between bone formation and bone resorption. When the balance is broken and bone resorption by osteoclasts dominates the process, osteolysis occurs. Areas with lower density values on a CT image of the bone structures can be visualized as erosions in a 3D view. Bone erosion is an early signal of the course of RA.

Erosion and Joint-Space-Narrowing Acquisition

A team designed high-resolution multimodality medical-imaging techniques for a clinical study about monitoring response to RA therapy. Molecular imaging of RA uses the wrist as an archetype for understanding the pathophysiology of the disease group and as a bellwether for diagnosis, monitoring, and therapeutic response. A study with molecular and multimodal imaging about bone resorption and formation will provide complete information about the disease's progress. Researchers have developed advanced therapeutic strategies to promote erosion repair, and sensitive radiographic techniques are used to capture changes in bone erosion.

Optimization of volume visualization techniques for examining RA patients' wrist and finger bones could greatly affect clinical imaging. The small bones of the wrist and fingers are typically affected by the disease early and give a good picture of disease status and progression. The ma-

for anatomical imaging markers associated with RA are bone erosion and joint space narrowing (JSN), which is associated with the destruction of cartilage at joints.⁵ Robust tools that can detect, visualize, and prospectively quantify erosion or JSN in 3D are limited in the clinical field, and our enhanced visualization technique could be an important diagnostic tool. New aggressive biologic treatments for RA have become available, but early differentiation of those who respond to the drugs from those who don't is critical to spare patients from the resulting disability, side effects, and cost. Advanced methods for change analyses, such as what we intend to support, could improve detectability and consistency compared to conventional methods.

The Datasets

The University of California, Davis Institutional Review Board (IRB) approved the clinical study. IRB guidelines for obtaining informed consent were used before patients were included in the study.

For the baseline, technicians scanned symptomatic wrists of RA patients who were candidates for receiving etanercept, a biopharmaceutical. (The patients' median age was 55.) This process employed a custom-built high-resolution PET/CT (positron emission tomography and x-ray computed tomography) scanner.⁴ The patients didn't move between the PET and CT scans, so this process produced coregistered imaging data. Technicians rescanned the patients at approximately one month (30 ± 5 days) or three months (90 ± 10 days) after the treatment started.

We used the datasets of three patients, whom we'll call Patients 1, 2, and 3. Each dataset consisted of the baseline scans and the scans after 1 or 3 months of therapy. Each resulting volume was $512 \times 512 \times 512$ voxels. We used only the CT data. Figure 1 displays CT data from one of the test datasets.

Study Methods and Setting

For our study, we had to choose an adequate renderer and properly set the light sources.

The Rendering System

We created an interactive volume renderer enhanced with GPU-accelerated global-illumination calculations.³ These calculations were based on highly parallelized, grid-based numerical schemes to solve a light propagation model including both absorption and scattering. This testbed system, implemented in C++ and GLSL (Open GL Shading Language), can produce realistic, high-quality

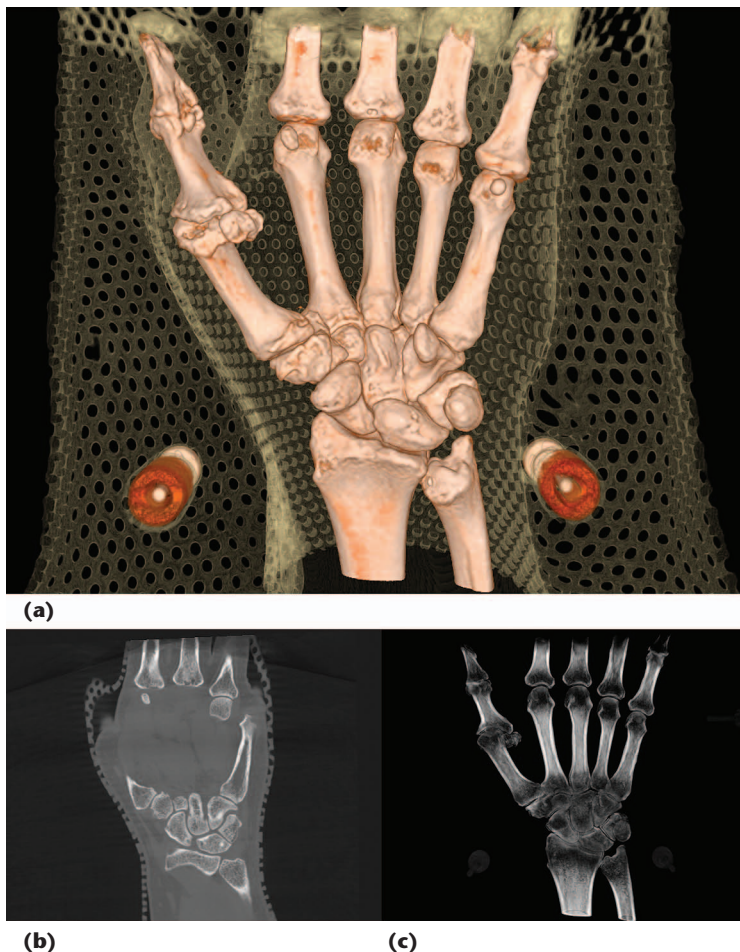


Figure 1. Images from a study of the wrist and finger bones of rheumatoid arthritis (RA) patients. (a) A direct volume rendering of a computed-tomography (CT) wrist dataset. Technicians used the mesh and plastic material to immobilize the hand during scanning. (b) A single slice and (c) an x-ray radiography-like rendering of the volume data.

shading, including shadow and multiple scattering inside translucent and heterogeneous media. Our system achieves interactive rendering rates with these limitations:

- It supports only an isotropic phase function for scattering.
- It doesn't support arbitrary complex environment lighting.
- It doesn't support diffuse or specular interreflection (multiple light bounces).

We believed these limitations wouldn't change our study's findings. However, for existing local and global-illumination methods, oversimplified shading models could lead to wrong interpretations of certain structures. For example, without proper occlusion, certain concave structures might be overilluminated, and viewers might not detect the concaveness. Another example would be some



Figure 2. Three-point lighting using the data from Patient 3. (a) With only a key light, part of the bone is unlit and not visible. (b) Using both a key light and a fill light provides some improvement, but part of the bone still isn't visible. (c) A visualization that also incorporates a back light provides the best clarity.

spatial structures with varying thickness. Identifying the local structures' thickness is easy if light scattering is present where light penetrates thin areas. As we found in our study, global-illumination methods aren't as helpful if the light sources aren't properly placed.

Multiple Lights and Their Settings

Using multiple light sources might have a complex impact on the bone structures' shape, owing to shading. Furthermore, the surface details revealed depend largely on each light source's position. You can affect perception of the shape by adjusting the light source's direction, which can change the shading pattern.⁶

Our system lets us easily change each light source's position and direction. To assist users, it supports automatic lighting setup based on the camera position and regions of interest. We adopt the three-point lighting commonly used in photography and movies. Three directional light sources illuminate the subject and generate the desired shading:

- A *key light* provides the main illumination of the features of interest in the 3D domain. It's the brightest of the three lights and is on the upper left of the main features of interest. This light by itself casts strong shadows.
- A *fill light* is a secondary light source. It has lower intensity than the key light and is on the front right of the main features of interest. It makes more of the features visible and helps soften the illumination.
- A *back light* is behind the main features of interest. It helps enhance the object's silhouette. In our setting, it also penetrates the bone's eroded parts to convey the bone's thickness.

Figure 2 shows the effect of three-point lighting, using the Patient 3 dataset. The back light is blue; the other two are white. With only a key light,

part of the bone isn't shown (see Figure 2a). Using both a key light and a fill light highlights more features of the bone surface (see Figure 2b). However, neither the key light nor fill light provides a clear shape of the bone. The bone contours are partially ambiguous. Observers have difficulty distinguishing the forefinger bone from the background. As Figure 2c shows, the back light defines the bones' boundary. Such an arrangement effectively distinguishes the object from the background and perceptually enhances the 3D shape by changing the shading.⁷ Finally, although shadows might be helpful for certain visualization tasks, illumination that can provide both depth information and a full view of the features of interest is more desirable for medical-imaging applications.

Evaluation

We evaluated visualizations incorporating global illumination to determine to what extent advanced illumination helped users in their tasks. This evaluation included

- an informal assessment of global illumination,
- validation with perceptual-quality measures,
- consideration of light sources, and
- a formal user study.

The participants were expert users interested in examining bone erosions typical of RA (manifested as breaks or holes on the bone surface) and whether the erosions' status changed during treatment. They expected that erosions in non-responders would progress over time, whereas in responders, erosive changes would be arrested and potentially reversed.⁸ The experts also wanted to examine how JSN changed during treatment.

Local versus Global Illumination

We first asked the participants to compare images using local illumination versus those using global illumination. They previously hadn't used volume




Figure 3. Comparing local and global illumination for images of Patient 1. (a) The baseline images. (b) Images taken three months after treatment began. The left images show local illumination; the right images show global illumination. For the baseline images, with global illumination, the participants concluded they could better identify the true erosions and their extent. After three months of therapy, the changes in the second finger’s metacarpophalangeal joint are clear from the 3D rendering.

visualization with global illumination. For Patient 1’s baseline dataset with local illumination (see Figure 3a, left), lighting depended only on the surface normal (the angle of the surface to the light). Part of the rendered surface was ambiguous to the experts. They could only identify rapid changes in surface topology but couldn’t gauge whether these corresponded to true erosions. The lighting effect didn’t always convey the appropriate depth relationship, which made interpretation suboptimal. The participants found the excessive surface details due to local illumination distracting. (These details might also have been a by-product of device noise.)

With global illumination (see Figure 3a, right), the participants concluded they could better identify the true erosions and their extent. This more realistic portrayal of the erosions seemed to provide an adequate level of detail and let the participants focus on the bone’s relevant parts.

We also considered complementarity between these two rendering effects. In visualization of high-dimensional medical-imaging data, a trade-

off exists between sensitivity and specificity. A diagnostic test with high sensitivity has a very high probability of picking up disease but can be severely riddled with false positives and therefore have a low specificity. On the other hand, a highly specific test can accurately identify disease-free patients but might have poor sensitivity.

In our context, local illumination yielded high sensitivity but resulted in a high false-positive rate. Global illumination resulted in high specificity but might have missed small changes in morphology that might have proved important. These results are consistent with the results of our quantitative user study. The participants therefore unanimously agreed that it was advantageous to have both locally and globally illuminated renditions. In that way, they could better disambiguate the relationship between the features of interest and therefore improve interpretation.

Figure 3b shows Patient 1’s condition after three months of therapy. With local illumination, finding the differences between the baseline images

Table 1. Quantitative perceptual quality. A smaller number means better quality.

Term	Local	Global illumination
Edge	0.50000	0.45000
Mean	0.08200	0.08400
Variance	0.15000	0.08100
Histogram	0.00016	0.00015
Total	0.73216	0.61515

and those taken after three months of therapy is difficult. However, with global illumination, the changes in the second finger’s metacarpophalangeal joint are clear. Although bone regrowth can happen during treatment, on the basis of these data, the participants couldn’t conclude whether the differences were due to CT image noise or true erosion filling. To further determine the differences’ cause, they would have had to analyze additional images.

Validation with Perceptual Measures

A perceptual-based metric can quantify an image’s perceptual quality.⁹ Images convey visual information such as shape, surface properties, and realism, which can be described by features such as edges, shading gradients, shadows, highlights, and the light source. We employed such a metric to measure the effectiveness of the volume-rendered images. Because we used different light settings in rendering, we didn’t consider the light source when measuring each image’s quality.

The quality function comprises four terms for corresponding features in the image:

$$f_Q = f_{edge} + f_{var} + f_{mean} + f_{hist},$$

where

- f_Q is the overall perceptual quality,
- f_{edge} is the detected edges’ prominence,

- f_{var} is the luminance variance’s distance to a target value (42 on a 0–255 scale),
- f_{mean} is the mean’s distance to a target mean (124 on a 0–255 scale), and
- f_{hist} measures the luminance histogram’s distance from an ideal equalized histogram.

A smaller f_Q means better perceptual quality. For details about these measures, see “Automatic Lighting Design Using a Perceptual Quality Metric.”¹⁰

We needed to compute an image map to mark different regions in an image: the edges, background pixels, and surface pixels. On the basis of what we learned from the expert users about which regions are more important and which regions are insignificant during observations, we weighted each region accordingly to quantify the image quality. Table 1 shows that with this perceptual metric, global illumination was indeed better than local illumination for the images in Figure 3, which is consistent with expert users’ opinions.

Light Settings

Because we adopted three-point lighting for global illumination, we experimentally determined optimal settings for the visualization tasks. As Figure 4a shows, although we can still see the details over the bone structures, the back light made the parts with erosions look thinner and more transparent, which might provide additional information to users. The participants thought that this enhancement was helpful and helped them determine the changes’ degree and extent. As Figure 4b shows, decreasing the overall opacity revealed more differences. During the evaluation, we let the participants move the viewpoints and lights to see the differences they could make. It was clear to them that using additional lights in key locations helped.




Figure 4. Comparing light settings for images of Patient 1. (a) The baseline image, using global illumination with a blue back light and an orange key light. (b) The baseline image, using the same light sources as in Figure 4a but with decreased opacity. (c) The image after three months of treatment, using the same setting as in Figure 4b. Comparison of Figures 4b and 4c shows that after the treatment, no significant bone reconstruction occurred.




Figure 5. Comparing light settings for images of Patient 2. (a) The baseline image, using local illumination with a distant light. (b) The baseline image, using global illumination with a blue back light and an orange key light. A green circle marks the eroded area on the second finger’s phalanges. (c) An image after one month of treatment, using global illumination with a blue back light and an orange key light. Comparison of the purple circles in Figures 5b and 5c indicates some bone reconstruction after the treatment.

One confounding factor in our study was due to the inherent noise in the CT images. This was because of the chosen CT technique, which was optimized to minimize the radiation dose and truncation artifacts. So, expert users can’t always tell whether a small feature is due to image noise or is an actual, albeit small, change.

As Figures 5a and 5b show, there’s an eroded area at the second finger’s phalanges in the baseline case; this area appears reduced in the follow-up scan (see Figure 5c). With the 3D renditions of the data using global illumination and the ability to interactively manipulate lights and the viewpoint, the participants were more confident about differentiating image noise or truncation from real change. With local illumination (see Figure 5a), the perceived erosion was harder to distinguish from other high-frequency surface structures. With global illumination, erosive changes were more prominent than other features, drawing the participants’ attention for further investigation.

The Formal User Study

The six participants in this study were radiology and medical-imaging researchers. They all had background knowledge and experience in wrist RA diagnosis with 2D x-ray images.

The participants first took a tutorial in which they viewed two datasets with global and local illumination and learned how to pick erosion holes. While examining, they could rotate the wrist. After they submitted all the erosions they thought they found, the system showed the correct ones.

Next, the participants took the real test. We used four wrist datasets from different patients. To avoid a learning effect, each participant saw each dataset in either global illumination or local illumination, not both. We assigned the datasets evenly



Figure 6. The erosions that a participant picked. We evaluated the number of correctly identified erosions, the number of incorrect identifications, and the number of missed erosions.

to the participants so that each one saw two sets with global illumination and two sets with local illumination. That is, we randomly exposed each participant to a particular combination of the four datasets.

For each case, we noted

- the visualized dataset,
- whether the dataset used global or local illumination,
- the erosions the participant picked (see Figure 6), and
- the time (in seconds) the participant took.

This study aimed to quantitatively validate the accuracy and efficiency of global illumination for wrist RA diagnosis. We had two hypotheses. First, global illumination helps users find the correct

Table 2. The study results.

Participant	Global illumination (GI)				Local illumination (LI)				GI time (s)	LI time (s)
	No. of erosions missed	No. of wrong erosions selected	Total no. of errors	Total no. of erosions	No. of erosions missed	No. of wrong erosions selected	Total no. of errors	Total no. of erosions		
1	7	1	8	20	7	15	22	28	155.35	499.53
2	2	4	6	28	5	21	26	20	225.05	403.85
3	10	0	10	27	11	6	17	21	80.28	136.16
4	5	4	9	31	8	4	12	17	207.76	531.63
5	0	4	4	17	20	10	30	31	563.29	930.18
6	9	0	9	20	9	8	17	28	184.09	246.01

erosion holes. Second, with global illumination, users spend less time finding the erosions.

Test Results

Table 2 summarizes the results. We grouped trials by whether they used global or local illumination. The total number of errors combines the number of real erosions the participants missed and the wrong erosions they selected. We counterbalanced the assignment of datasets to the participants to minimize the dataset's effect on the overall results. There were significantly fewer errors with global illumination than with local illumination. (The mean number of global-illumination errors was 7.67, the mean number of local-illumination errors was 20.67, and the two-sided p value by paired t -test was 0.0147.) If we consider the effect of the total number of correct erosions for the different groups of datasets and count the percentage of errors the participants made in each group, the two-sided p value by paired t -test was 0.0100, which also meets $p < 0.05$. Our hypothesis about accuracy is thus accepted at a significance level of 95 percent.

The results show that different rendering methods had different effects on the participants. Some participants had relatively higher accuracy with global illumination. Some weren't affected much by the rendering method. Some participants tended to select more erosions and made more mistakes. Some participants tended to select fewer erosions and missed the correct ones. However, no matter how large the effect was, global illumination generally lowered the number of errors.

Table 2 also shows the time difference between the two groups. With global illumination, participants spent an average of 235.97 s finding the erosion holes. For local illumination, the mean time was 457.89 s. With a paired t -test, the two-sided p value was $0.0124 < 0.05$. Our hypothesis about time is thus accepted at a significance level of 95 percent.

During the study, some participants were relatively faster (for example, Participant 3) and some

were slower (for example, Participant 5). The time difference didn't significantly affect their correctness. So, global illumination appeared to help the participants become more time-efficient when finding bone erosions, without sacrificing accuracy.

It's also helpful to consider the reproducibility of each user on the enhanced visualization. In future research, we'll ask the expert users to participate in a similar test after a period of time, such as three months. We'll then compare the results from the same participant to study the intra-observer variability, which will help confirm and improve our study's accuracy.

Our research effectively demonstrated advanced illumination's value for clinical radiology. Our findings suggest the proper use of multiple light sources and global illumination and will help us and others design more usable visualization systems. Global illumination is more than adding shadows. In many cases, strong shadows aren't desirable. To enable correct visual examination, we showed that certain critical aspects of the data can be perceived only with proper light settings and global illumination. In addition, the Phong illumination model (used in local illumination) can't produce correct absorption and scattering effects that reflect, for example, the eroded bone's thickness and concaveness. However, our expert users pointed out that visualizations created with only local illumination and with only global illumination are complementary.

Visualization tools such as ours might allow rapid interpretation of otherwise complex images and could improve current diagnostic tests' sensitivity and specificity. The clinical need to produce accurate change analyses in patients undergoing treatment for RA is unmet, and advanced image-processing methods such as the one we evaluated here have merits in this regard. We plan to continue validating our methods with a larger number of patients and further development of robust visualization.

One area of further research is to design appropriate user interfaces for applying advanced rendering methods in medical and nonmedical visualization tasks. We need to determine what rendering parameters should be hidden or controllable by the users. What's the most intuitive way for users to place light sources? How do we direct users to properly use lights when rendering their data? We've done some preliminary research to answer some of these questions.⁹ Finally, for multimodality data, there should be other aspects of the data, such as positron emission tomography or magnetic resonance imaging, captured by particular image modalities for setting up rendering and visualization parameters. ❖

Acknowledgments

This project was supported partly by the US National Institutes of Health through grants UL1 TR000002, K12 HD051958, and R01 CA129561 and by the US National Science Foundation through grants IIS-1320229 and CCF-1025269.

References

1. T. Kroes, F.H. Post, and C.P. Botha, "Exposure Render: An Interactive Photo-realistic Volume Rendering Framework," *PLoS ONE*, vol. 7, no. 7, 2012, article e38586.
2. Y. Zhang, Z. Dong, and K.-L. Ma, "Real-Time Volume Rendering in Dynamic Lighting Environments Using Precomputed Photon Mapping," *IEEE Trans. Visualization and Computer Graphics*, vol. 19, no. 8, 2013, pp. 1317-1330.
3. Y. Zhang and K.-L. Ma, "Fast Global Illumination for Interactive Volume Visualization," *Proc. ACM Symp. Interactive 3D Graphics and Games (I3D 13)*, 2013, pp. 55-62.
4. A.J. Chaudhari et al., "High-Resolution ¹⁸F-FDG PET with MRI for Monitoring Response to Treatment in Rheumatoid Arthritis," *European J. Nuclear Medicine and Molecular Imaging*, vol. 37, no. 5, 2010, p. 1047.
5. D. Van Der Heijde, "How to Read Radiographs According to the Sharp/van der Heijde Method: Part 3: Rheumatoid Arthritis: Response," *J. Rheumatology*, vol. 26, no. 3, 1999, pp. 743-745.
6. K. Berbaum, T. Bever, and C.S. Chung, "Light Source Position in the Perception of Object Shape," *Perception*, vol. 12, no. 4, 1983, pp. 411-416.
7. V.S. Ramachandran, "Perception of Shape from Shading," *Nature*, vol. 331, no. 6152, 1988, pp. 163-166.
8. E. Gravallesse and N. Walsh, "Rheumatoid Arthritis: Repair of Erosion in RA—Shifting the Balance to Formation," *Nature Reviews Rheumatology*, vol. 7, no. 11, 2011, p. 626.
9. Y. Zhang and K.-L. Ma, "Lighting Design for Globally Illuminated Volume Rendering," *IEEE Trans. Visualization and Computer Graphics*, vol. 19, no. 12, 2013, pp. 2946-2955.
10. R. Shacked and D. Lischinski, "Automatic Lighting Design Using a Perceptual Quality Metric," *Computer Graphics Forum*, vol. 20, no. 3, 2001, pp. 215-227.


Lin Zheng is a PhD student in computer science at the University of California, Davis. Contact her at lzheng@ucdavis.edu.

Abhijit J. Chaudhari is an assistant professor of radiology at the University of California, Davis. Contact him at ajchaudhari@ucdavis.edu.

Ramsey D. Badawi is an associate professor of radiology and biomedical engineering at the University of California, Davis. Contact him at ramsey.badawi@ucdmc.ucdavis.edu.

Kwan-Liu Ma is a professor of computer science at the University of California, Davis. He's a member of the CG&A editorial board. Contact him at ma@cs.ucdavis.edu.

Contact department editor Mike Potel at potel@wildcrest.com.

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.

