Reflectance Function Approximation for Material Classification

Edward Wild

CS 766 Final Project Report

Abstract

Reflectance functions are approximated from data using kernel regression and used to classify materials. Classification algorithms are proposed to deal with unseen materials. Experimental results show that some reflectance functions can be approximated quite accurately with kernel regression, and that accurate approximations can be used to classify materials. The kernel regression techniques here use convex optimization techniques that are simpler than the nonlinear techniques often used to fit more sophisticated reflectance models to data. Preliminary results suggest that it is possible to extend classification to work with unseen materials, which has important implications towards the scalability of the method.

1 Introduction

This report summarizes the results of a project to approximate reflectance functions and classify materials based on those approximations. The report motivates the importance of studying reflectance function approximation for material classification, and states the specific conditions that will be considered. The methods used are summarized, and some experimental results are presented. Finally, some future directions of research are proposed and briefly discussed.

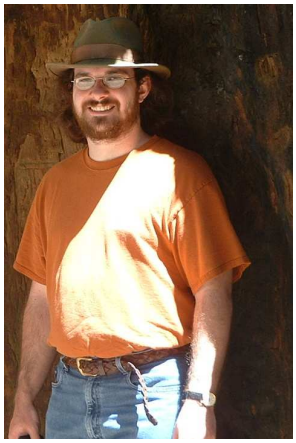

Figure 1: A scene that could be analyzed using reflectance function information.

2 Motivation

The ability to recognize materials based on their reflectance has many potential applications. For example, the ability to recognize the materials cotton, denim, and skin in figure 1 would make it easier to do segmentation, feature tracking, and object recognition. Segmentation would be easier because even though the illumination across the shirt appears to make an edge on the shirt, by using reflectance information a segmentation algorithm would realize that the edge is merely an effect of illumination, not an edge between two different regions. Feature tracking would be improved because it would be possible to enforce the constraint that features be on the same material between frames. Object recognition would be improved because the combination of the materials of denim, cotton, and skin in such a configuration is extremely likely to mean that the picture is of a person. Furthermore, algorithms that make use of reflectance information would be able to deal with illumination and pose changes.

3 Problem Statement

The reflectance functions studied will be bidirectional reflectance distribution functions. For all of the work here, measurements of the reflectance functions are considered to be exact, meaning that scene geometry and illumination are

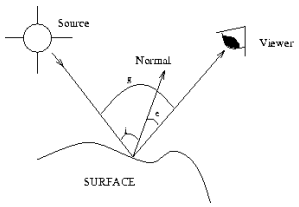

Figure 2: An illustration of reflectance by Wallace and Price. Available online, http://www.dai.ed.ac.uk/CVonline/LOCAL_COPIES/MARBLE/medium/shading/reflect.htm.

known. The reflectance functions are real valued functions of four variables. The variables are the angles that the view and illumination directions make with the surface normal. The situation is illustrated in figure 2. Let (θ_v, ϕ_v) be the angle of the view direction with the surface normal, and (θ_i, ϕ_i) be the angle of the illumination direction with the surface normal. Then the reflectance function f is a function of the variables $(\theta_v, \phi_v, \theta_i, \phi_i)$.

The CUReT database of reflectance functions contains measurements of the bidirectional reflectance distribution function of 61 materials. Each material is measured from 205 different values of $(\theta_v, \phi_v, \theta_i, \phi_i)$. The database is available at http://www1.cs.columbia.edu/CAVE/curet/.

4 Related Work

The use of reflectance functions for classification is explored by Dror et al. in [5, 4]. They find that it is possible to distinguish materials based on reflectance properties. Their motivation appears to be work by Fleming et al. that indicates that humans make use of reflectance properties to identify materials [8].

The CUReT database of reflectance functions is described by Dana et al. in [3]. The Oren-Nayar reflectance model [15] and the Koenderink et al. representation [9] are fit to the CUReT data by Dana et al. in [2].

There has been a great deal of research on reflectance models. The work that was most influential to this project was the work of Lafortune et al. [10], McCool et al. [14], and especially the work of Lensch et al. [11]. In each of these papers, reflectance functions were approximated using a combination of basis functions. However, the reflectance functions were approximated using more complicated nonlinear optimization techniques than those used here.

5 Method

There are multiple steps required to recognize materials using reflectance properties. First, an accurate approximation of measured reflectance functions from known materials must be made. Then, these accurate approximations can be used to classify materials.

5.1 Kernel approximation for reflectance functions

The CUReT database, described by Dana et al. in [3], contains BRDF measurements for 61 different materials. These measurements are of the form $f(\theta_v, \phi_v, \theta_i, \phi_i) = y$. This project requires accurate approximations of these functions. The approach described here is to use kernel regression techniques [6, 17, 7, 13] to approximate the given functions as accurately as possible.

Kernel regression is implemented by solving the optimization problem:

$$\min_{(\alpha,b)} \frac{1}{2}\|\alpha\| + \frac{\nu}{2}\|K(A, A')\alpha + be - y\| \tag{1}$$

Where each row A_i of the matrix A is the coordinate $(\theta_v, \phi_v, \theta_i, \phi_i)$ corresponding to the measured BRDF value y_i. Here, K is the commonly used Gaussian kernel [18, 1, 12] given by: $(K(A, B))_{ij} = exp(-\mu\|A'_i - B_{\cdot j}\|^2), \mu > 0$. The vector e is a vector of ones of appropriate dimension. The two parameters of this problem are μ, which controls the width of the Gaussian functions, and ν which determines the weight given the error between the measured data and the approximation. If ν is large, then the amount of error between the approximation and the measured data will be reduced, which is appropriate if the measurements are believed to be very accurate. The solution of the problem is the vector α and the scalar b which have the property that $f(\theta_v, \phi_v, \theta_i, \phi_i) = y_i \approx K((\theta_v\ \phi_v\ \theta_i\ \phi_i), A')\alpha + be$. The optimization problem turns out to be convex for fixed μ and ν. Since it is differentiable, its minimum can be found by taking its derivative and setting it equal to zero. The result is the linear system of equations:

$$\left(\nu \begin{pmatrix} K(A, A') & e \end{pmatrix}' \begin{pmatrix} K(A, A') & e \end{pmatrix} + \begin{pmatrix} I & 0 \\ 0 & 0 \end{pmatrix}\right) \begin{pmatrix} \alpha \\ b \end{pmatrix} = \nu \begin{pmatrix} K(A, A') & e \end{pmatrix}' y \tag{2}$$

The procedure used for selecting μ and ν is given as algorithm 1. Algorithm 1 uses ten fold cross validation to select the best parameters μ and ν from

the given sets M and N respectively. Ten fold cross validation is necessary to choose parameters that are likely to have good performance on unseen data. Algorithm 2 describes the cross validation procedure in detail.

Algorithm 1 Kernel Regression(A, y, M, N):(α, b)

```
e ← ∞
p ← ∞
for all μ ∈ M do
  for all ν ∈ N do
    t ← tfcve(A, y, μ, ν) {Apply algorithm 2}
    if t < e then
      p ← (μ, ν)
      e ← t
    end if
  end for
end for
(μ, ν) ← p
(α, b) ← kr-solve(A, y, μ, ν) {Apply algorithm 3}
```

5.2 Classification of reflectance functions

Once approximations of reflectance functions are available, they can be used to classify materials. Given a set of reflectance function approximations, a nearest-neighbor algorithm can be used to classify a set of new measurements. The classification is performed by finding the function that is closest to the measured values at the measured points. The procedure for classifying a substance is described in detail in algorithm 4.

A limitation of algorithm 4 is that it only works with materials whose reflectance functions have already been measured and approximated. The experimental results discussed in section 6.2 indicate that it is possible to identify materials with fewer measurements than are needed to make an accurate reflectance function approximation. Define an unseen material to be a material whose reflectance function has not been approximated. Then it is possible to determine from only a few measurements of an unseen material m that m is not one of the materials with a reflectance function that has been approximated. This procedure is described in detail by algorithm 5.

Algorithm 2 tfcve(A, y, μ, ν):e

Shuffle (A, y)
Partition (A, y) into folds $\{(A_1, y_1), (A_2, y_2), \ldots, (A_10, y_10)\}$ so that each (A_i, y_i) contains approximately $\frac{1}{10}$ of the data in (A, y) and that the (A_i, y_i) are mutually disjoint.
$e \leftarrow 0$
for $i = 1$ to 10 **do**
  $(A_{train}, y_{train}) \leftarrow \bigcup_{j \neq i} (A_j, y_j)$
  $(\alpha, b) \leftarrow$ kr-solve(A_{train}, y_{train}, μ, ν) {Apply algorithm 3}
  $e \leftarrow e + mean(|K(A_i, A'_{train}) * alpha + be - y_i|)/mean(|y_i|)$ {Add the average error for this fold}
end for
$e \leftarrow e/10$

Algorithm 3 kr-solve(A, y, μ, ν):(α, b)

Solve equation (2) with A, y, μ, and ν to obtain α and b.

Algorithm 4 classify(F, A, y):(n, e)

F is the set of known reflectance functions. For $f \in F$, $f(A) = \hat{y}$, where $\hat{y}$ is the approximate value of f at the measured points A.
$e \leftarrow \infty$
$n \leftarrow -1$
for all $f \in F$ **do**
  $\hat{y} \leftarrow f(A)$
  if $\|\hat{y} - y\| < e$ **then**
    $n \leftarrow \mathrm{id}(f)$ {id(f) is the identifier of the material corresponding to f}
    $e \leftarrow \|\hat{y} - y\|$
  end if
end for

Algorithm 5 Classify, and recognize unseen materials(F, A, y, t):(n)

Will return either the id of the material, or -1 if the material is unseen.
$(n, e) \leftarrow$ classify(F, A, y) { Apply algorithm 4}
if $e \geq t$ **then**
  $n \leftarrow -1$
end if

In practical terms, algorithm 5 is only slightly more useful than algorithm 4. The improvement in utility is small because algorithm 5 still only recognizes materials whose reflectance functions have been measured. The only improvement is that it will no longer incorrectly classify an unseen material using the material that happens to have the closest reflectance function. The new algorithm does not, for example, measure how similar two unseen materials are.

A proposed improvement to algorithm 5 is algorithm 6, which would be able not only to recognize that unseen materials, but also to classify them. That is, the new algorithm would be able to recognize whether two unseen materials m and n, which are classified at different times, are the same material. The improvement comes from a step in the algorithm that approximates the reflectance function of an unseen materials with the reflectance functions of the k closest materials. Using neighboring reflectance functions as a basis for reflectance function approximation was inspired by the work of Lensch et al. [11].

Algorithm 6 Classify with unseen materials(F, A, y, t, k):(n, F)

May augment F with the reflectance functions of new materials.
$(n, e) \leftarrow$ classify(F, A, y) {Apply algorithm 4}
if $e \geq t$ **then**
  $f \leftarrow$ approximate the reflectance function as a linear combination of its k nearest neighbors
  $\mathrm{id}(f) \leftarrow$ new unique id for this material
  $n \leftarrow \mathrm{id}(f)$
  $F \leftarrow F \cup \{f\}$
end if

6 Experimental Results

The experiments investigate the performance of the techniques described in the previous section for reflectance function approximation and material classification.

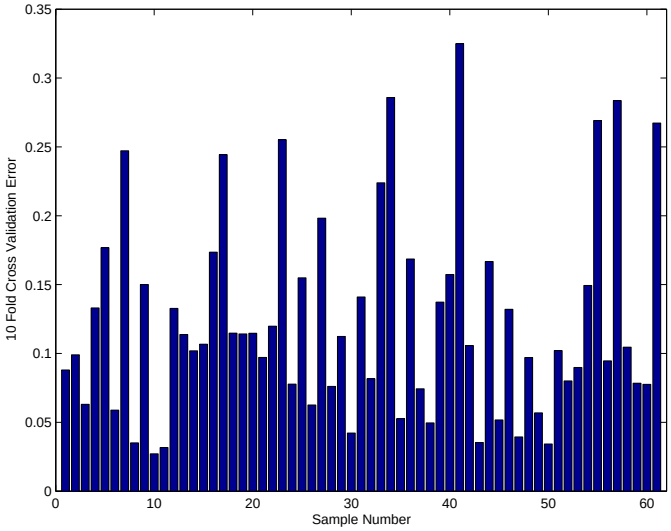

Figure 3: Approximation error for the 61 materials in the CUReT database.

6.1 Reflectance function approximation from data

This experiment demonstrates the performance of kernel regression on reflectance function approximation. Each of the 61 materials in the CUReT database was approximated from the measured data. The results are shown in figure 3. The error reported is the 10 fold cross validation error for each material. That is, the data for each material was broken into 10 training and testing sets. Algorithm 1 was run on each training set to obtain an approximation, which was then evaluated on the testing set. The error reported is the average error over all the testing sets. Note that this cross validation procedure is performed in addition to the procedure described in algorithm 2 that selects the parameters.

Figure 3 shows that many of the materials can be approximated very accurately using kernel regression. Note that 27 of the approximations have less than 10 percent average error. The parameters used in obtaining these results were $M = \{0.1, 1, 10, 15\}$ and $N = \{10, 1000, 1000000\}$. It is likely that better results could be obtained with larger M and N. The most accurate approximation was that of sample 10, plaster, which had average error 2.71 percent. The least accurate approximation was that of sample 41, brick, with average error 32.49 percent.

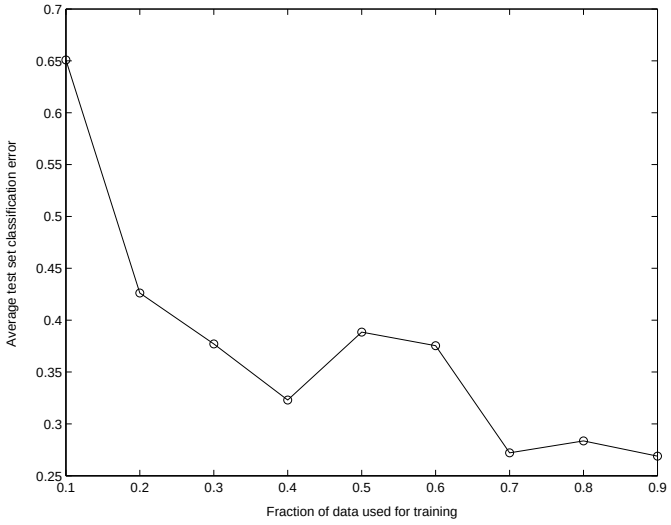

Figure 4: Classification error for the CUReT database.

6.2 Classification of materials using reflectance functions

To test classification of materials using reflectance functions, approximations were made using only some of the CUReT data. The data not used to create the approximations was then used to evaluate the accuracy of the classification system. Figure 4 shows the results of an experiment that evaluates the classification accuracy. This experiment was performed by fitting the approximations to a given fraction of the data for the substance, and then attempting to classify the substance using the remaining data and algorithm 4. This process was repeated 10 times for each fraction of the data, with a random fraction of the data selected for fitting the approximation each time.

In figure 4, note that the trend is for the classification accuracy to improve as the fraction of the data used to make the approximation increases. The approximations were made using algorithm 1, with $M = \{0.1, 1, 10, 15\}$ and $N = \{0.1, 1, 10, 15\}$. It is likely that performance would improve with larger M and N.

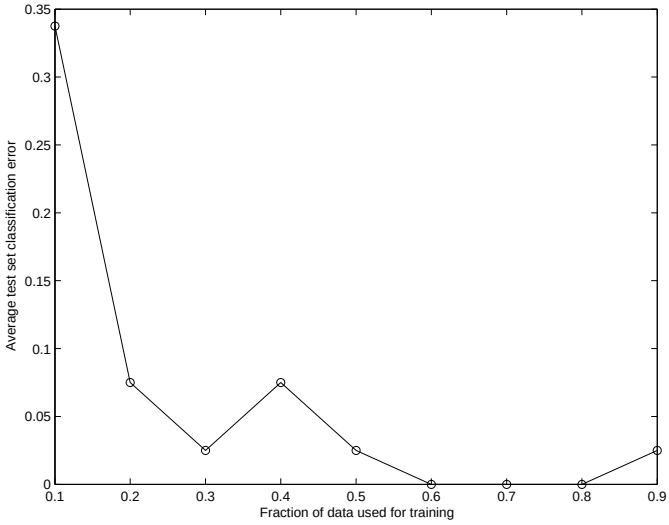

Figure 5: Classification error for the reduced database.

6.3 Results on a reduced database

In order to deal with unseen materials, very accurate approximations of the reflectance functions are needed. To achieve this accuracy, a reduced database containing only those materials whose approximation error was less than 10 percent. These substances are pebbles, three kinds of plaster, ribbed paper, salt crystals, stones, and concrete. Although one cannot make the claim that the results achieved on this reduced database will extend to the full database, this smaller database allows algorithms to be developed and tested quickly. Also, there is a hope that improving the accuracy of other reflectance function approximations will cause the results achieved on this database to apply to the full database. On the other hand, algorithms that fail on this database can be discarded. Thus, the reduced database serves as a convenient filter for algorithm development.

Figure 5 shows the results of classifying materials based on their reflectance functions. This figure was created using the same method as figure 4 in section 6.2, but only the 8 samples in the reduced database are considered. As seems natural, the average classification error is much reduced from figure 4. This reduction implies that classification error may go down with decreasing approximation error even for the full database.

The improved classification accuracy on the reduced database indicates that it may be possible to use the reduced database to test the procedures for classifying unseen materials. In preliminary experiments, it was possible to identify unseen materials correctly using algorithm 5 with little or no error, although algorithm 6 was less successful. This result may indicate that it is necessary to use a more sophisticated method of reflectance function approximation for the unseen materials, such as that of Yu et al. [19] or Ramamoorthi and Hanrahan [16].

The method used to test the performance of algorithm 6 was as follows. This procedure simulates the arrival of measurements from an unseen material at two different times. The goal is for the system to recognize that the two separate sets of measurements are of the same material. 40 percent of the measurements for each material was randomly selected for classification. This data is referred to as the classification data. The remaining data was used to obtain function approximations with algorithm 1, resulting in a set F of reflectance functions. Then, each of the materials m is given a turn at being "unseen" in the following manner. One half of the classification data for the material was used in an attempt to classify the material using the known reflectance functions $F' = \{f | id(f) \neq m\}$. If the result was that the material was unseen, then linear least squares was used to find a set of weights w with the k nearest materials, and the function $f'(\theta_v, \phi_v, \theta_i, \phi_i) = [f_k(\theta_v, \phi_v, \theta_i, \phi_i)]w$, where $[f_k(\theta_v, \phi_v, \theta_i, \phi_i)]$ is the row vector of functions that are one of the k nearest neighbors. Then the remaining half of the classification data was used in an attempt to classify the material, this time using $F'' = F \cup \{f'\}$. The result was considered correct only if the closest material had function f', and it was within the threshold used to accept or reject materials. The average error over all of the materials was reported. Using a threshold of 20 percent and $k = 7$, an error of 50 percent was obtained. This result means that 50 percent of the time, the algorithm correctly recognized that the first classification measurements were of an unseen material and that the second classification measurements were of the same unseen material. As a control, algorithm 5 was run using the same threshold, achieving an error of 0. The performance of this algorithm may be improved by a method to tune or the parameters t and k.

7 Conclusion and Future Work

This project has applied kernel regression to the problem of reflectance function approximation and presented a mechanism that allows the resulting approximations to be used for material classification. The process was extended to consider unseen materials, and some preliminary results were reported. Although the results for dealing with unseen materials are not as high as might be hoped, there is potential to improve them using more sophisticated techniques. Future testing is needed to evaluate the performance of the algorithms for dealing with unseen data on the complete dataset.

Some obvious extensions to the project include using larger numbers of possible parameters in M and N, and exploring other reflectance function models. Such improvements would hopefully improve the accuracy of the approximation. To improve the classification accuracy, classification algorithms such as support vector machines or artificial neural networks can be investigated. To make such algorithms more effective, substances with similar reflectance functions could be grouped together, possibly by using unsupervised clustering algorithms. This grouping process would also make the approach scale better for larger numbers of materials, and could remove the parameter k from the procedure for approximating reflectance functions of unseen materials in algorithm 6. This parameter would no longer be needed, since each unseen material could be approximated using the functions in its own group.

Further extensions to the project would involve making the classification process work when the geometry and illumination are unknown. Techniques mentioned in class could be used to approximate the geometry, and methods such as that of Ramamoorthi and Hanrahan [16] can be used to recover illumination under some conditions. In fact, use of these methods is likely to result in an approximation of the reflectance function. However, if it is known *a priori* that the materials in the scene have reflectance functions that are already approximated, then the material with the closest reflectance function to the approximation can be used to improve the estimate of the geometry and the illumination.

References

[1] V. Cherkassky and F. Mulier. *Learning from Data - Concepts, Theory and Methods*. John Wiley & Sons, New York, 1998.

[2] Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. Reflectance and texture of real-world surfaces. In *IEEE Conference on CVPR*, page 151, 1997. http://citeseer.nj.nec.com/article/dana97reflectance.html.

[3] Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, 1999. http://citeseer.nj.nec.com/article/dana97reflectance.html.

[4] R. Dror, E. Adelson, and A. Willsky. Surface reflectance estimation and natural illumination statistics. In *Proceedings of the 2nd IEEE Workshop on Statistical and Computational Theories of Vision*, 2001. http://citeseer.nj.nec.com/dror01surface.html.

[5] Ron O. Dror, Edward H. Adelson, and Alan S. Willsky. Estimating surface reflectance properties from images under unknown illumination. In *Proceedings of the SPIE 4299: Human Vision and Electronic Imaging IV*, San Jose, California, January 2001.

[6] H. Drucker, C. J. C. Burges, L. Kaufman, A. Smola, and V. Vapnik. Support vector regression machines. In M. C. Mozer, M. I. Jordan, and T. Petsche, editors, *Advances in Neural Information Processing Systems -9-*, pages 155–161, Cambridge, MA, 1997. MIT Press.

[7] T. Evgeniou, M. Pontil, and T. Poggio. Regularization networks and support vector machines. In A. Smola, P. Bartlett, B. Schölkopf, and D. Schuurmans, editors, *Advances in Large Margin Classifiers*, pages 171–203, Cambridge, MA, 2000. MIT Press.

[8] R. Fleming, R. Dror, and E. Adelson. How do humans determine reflectance properties under unknown illumination. In *Proceedings of CVPR Workshop on Identifying Objects Across Variations in Lighting: Psychophysics and Computation*, Colocated with CVPR 2001. Kauai, Hawaii, December 2001. http://citeseer.nj.nec.com/fleming01how.html.

[9] J. J. Koenderink, A. J. van Doorn, and M. Stavridi. Bidirectional reflection distribution function expressed in terms of surface scattering modes. In *ECCV '96. 4th European Conference on Computer Vision Proceedings*, volume 2, pages 28–39, 1996.

[10] Eric P. F. Lafortune, Sing-Choong Foo, Kenneth E. Torrance, and Donald P. Greenberg. Non-linear approximation of reflectance functions. *Computer Graphics*, 31(Annual Conference Series):117–126, 1997. http://citeseer.nj.nec.com/lafortune97nonlinear.html.

[11] Hendrik Lensch, Michael Goesele, Jan Kautz, Wolfgang Heidrich, and Hans-Peter Seidel. Image-Based reconstruction of spatially varying materials. In Steven Gortler and Karol Myszkowski, editors, *Proceedings of the 12th Eurographics Workshop on Rendering*, pages 103–114, London, Great Britain, 2001. Springer. http://citeseer.nj.nec.com/lensch01imagebased.html.

[12] O. L. Mangasarian. Generalized support vector machines. In A. Smola, P. Bartlett, B. Schölkopf, and D. Schuurmans, editors, *Advances in Large Margin Classifiers*, pages 135–146, Cambridge, MA, 2000. MIT Press. ftp://ftp.cs.wisc.edu/math-prog/tech-reports/98-14.ps.

[13] O. L. Mangasarian and D. R. Musicant. Large scale kernel regression via linear programming. *Machine Learning*, 46:255–269, 2002. ftp://ftp.cs.wisc.edu/pub/dmi/tech-reports/99-02.ps.

[14] Michael D. McCool, Jason Ang, and Anis Ahmad. Homomorphic factorization of BRDFs for high-performance rendering. In Eugene Fiume, editor, *SIGGRAPH 2001, Computer Graphics Proceedings*, pages 185–194. ACM Press / ACM SIGGRAPH, 2001. http://citeseer.nj.nec.com/mccool01homomorphic.html.

[15] M. Oren and S.K. Nayar. Generalization of the lambertian model and implications for machine vision. *IJCV*, 14(3):227–251, April 1995. http://citeseer.nj.nec.com/oren95generalization.html.

[16] Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for inverse rendering. In Eugene Fiume, editor, *SIGGRAPH 2001, Computer Graphics Proceedings*, pages 117–128. ACM Press / ACM SIGGRAPH, 2001.

[17] A. Smola and B. Schölkopf. On a kernel-based method for pattern recognition, regression, approximation and operator inversion. *Algorithmica*, 22:211–231, 1998.

[18] V. N. Vapnik. *The Nature of Statistical Learning Theory*. Springer, New York, second edition, 2000.

[19] Yizhou Yu, Paul Debevec, Jitendra Malik, and Tim Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs from. In Alyn Rockwood, editor, *Siggraph99, Annual Conference Series*, pages 215–224, Los Angeles, 1999. Addison Wesley Longman. http://citeseer.nj.nec.com/article/yu99inverse.html.