Image-Based Reconstruction of Spatially Varying Materials

Hendrik P.A. Lensch   Jan Kautz   Michael Goesele
Wolfgang Heidrich   Hans-Peter Seidel

Presentation by Tim Isganitis
CS 667 - Spring 2005
Cornell University

Goal: Quickly and accurately model the reflectance properties of real-world objects made up of multiple materials.

- Automatically acquire models (BRDFs) from images of an object.
- Account for variations of reflectance properties between different materials as well as within a given material.

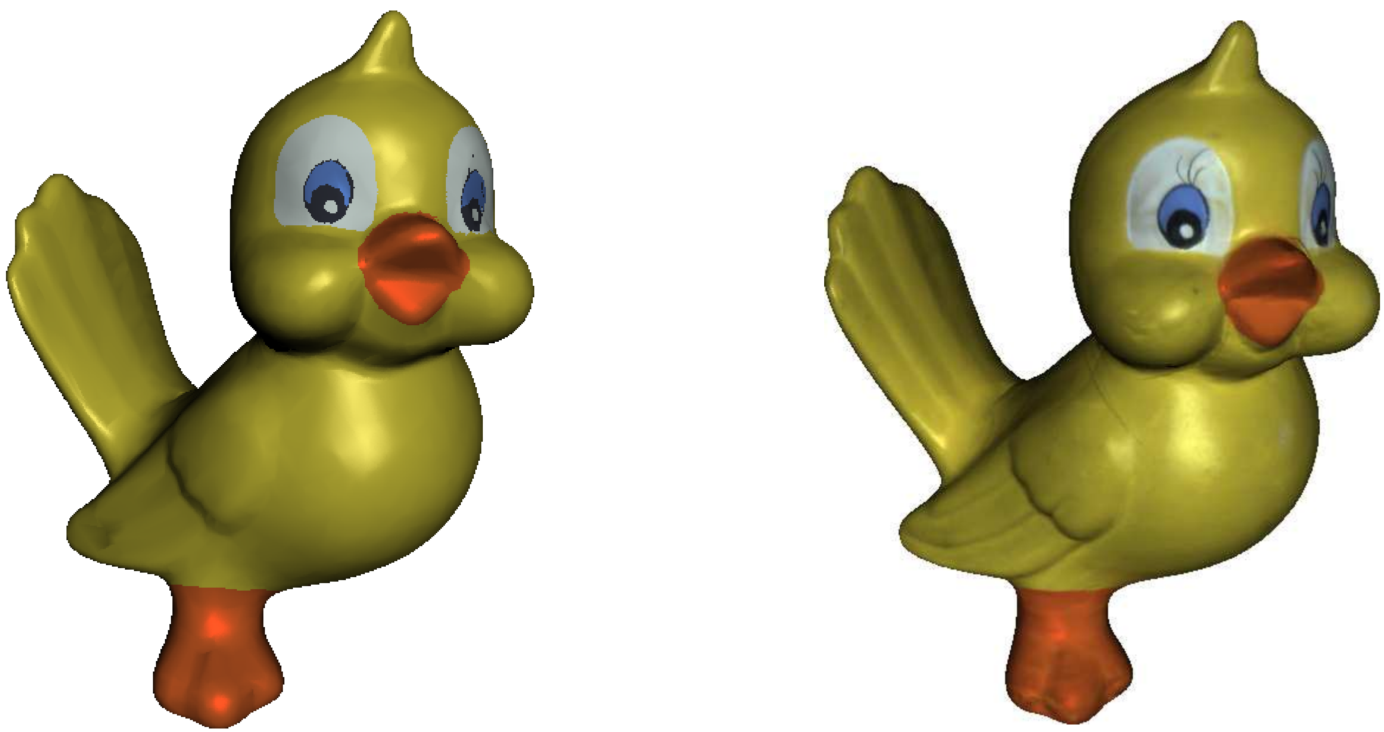

Previous Approaches

- Light Field and Image Database Solutions
- Dense sampling of radiance values for known lighting conditions.
  - Expensive to compute (images and time) and to store (tabular format).
  - Does not account for spatial variation within materials.
- Fitting a reflectance model to a small set of radiance samples.

Proposed Solution

- Use fitted reflectance models to reduce the number of samples necessary.
- Generate a set of basis BRDFs for the object using a hierarchical clustering/fitting technique.
- Represent a unique BRDF for each point as a linear combination of the basis BRDFs.

Process:

1. Acquisition
2. Resampling
3. BRDF Fitting
4. Clustering
5. Projection
6. Rendering

Acquisition

The system requires the following input:

- An accurate 3D model of the object
  - Represented in the form of a triangle mesh with known vertex normals
- A set of images of the object
  - The positions of the camera and the light source relative to the object must be known.
  - HDR

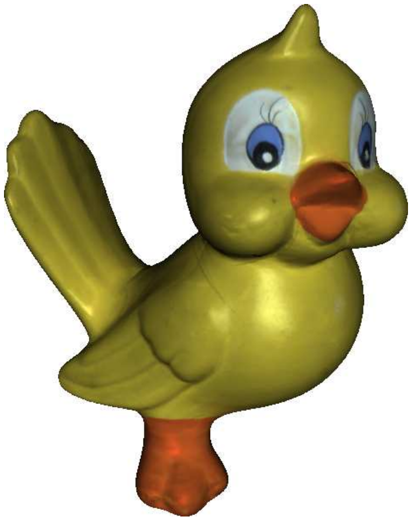

Tweety uses 14000 triangles and 25 views.

Resampling

Generate sample points (lumitexels) on the mesh for which radiance values will be stored and BRDFs computed.

- For each triangle, sample points chosen as projections of the pixel centers from the best image of that triangle.

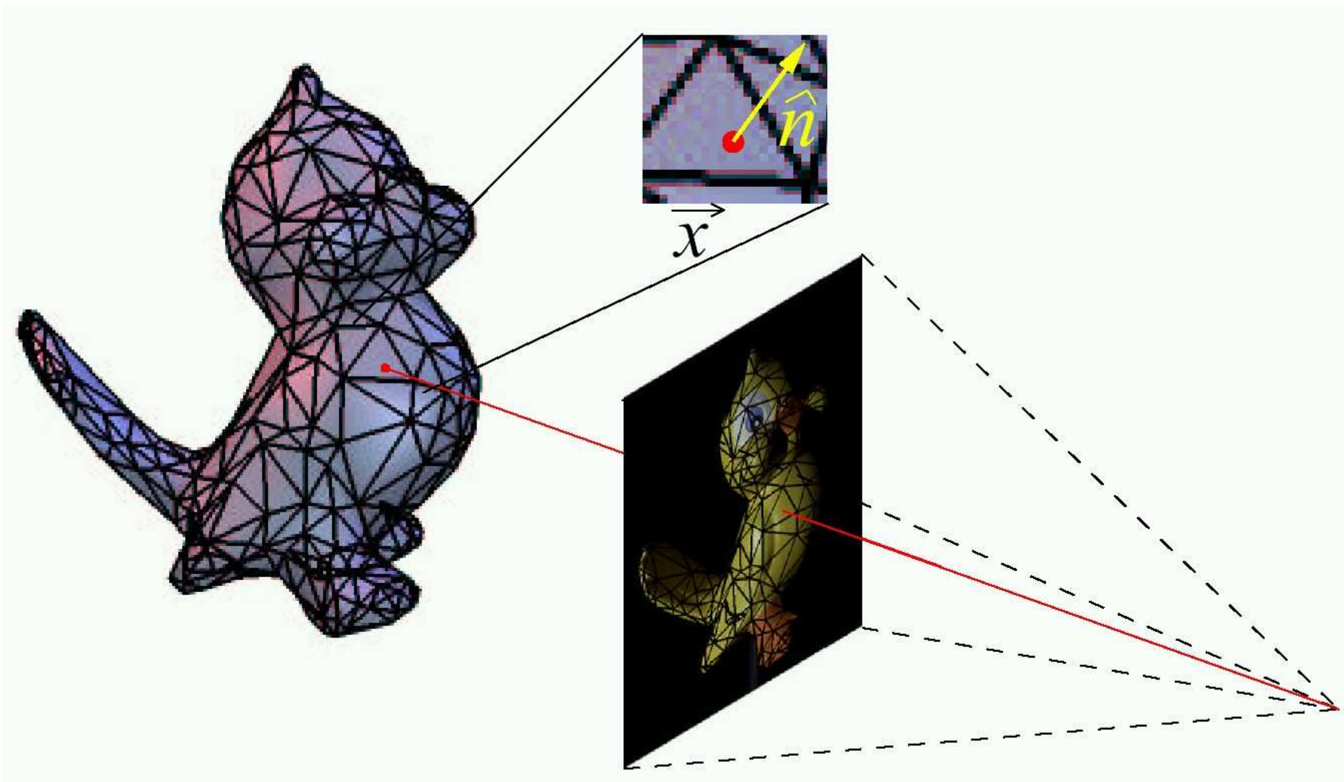

Resampling

Each lumitexel contains:

- Position on the mesh
- Normal at that point (interpolated)
- A set of radiance samples

There is one radiance sample for each image in which the lumitexel is visible, containing:

- Direction to the light source
- Direction to the camera (both relative to the surface normal)
- Outgoing radiance (image color scaled by source brightness and squared distance)

BRDF Fitting

Lafortune Model:

$$f_r(\hat{u},\hat{v}) = \rho_d + \sum_i \left[C_{x,i}(u_x v_x + u_y v_y) + C_{z,i} u_z v_z\right]^{N_i}$$

Parameters:

ρ_d – The diffuse component

$C_{x,i}$ $C_{z,i}$ – Determine off-specularity, retro/forward-reflectivity and albedo of the i^{th} lobe

N_i – The specular exponent

For more info see:

E. Lafortune, S. Foo, K. Torrance and D. Greenberg. Non-Linear Approximation of Reflectance Functions. In *Proc. SIGGRAPH*, pages 117-126, August 1997.

BRDF Fitting

Fitting the parameters of the model to a set of radiance samples is a non-linear least squares problem.

The standard approach to solving this problem for BRDF fitting is Levenberg-Marquardt optimization, which produces:

- An optimal set of parameters
- A covariance matrix of the parameters

BRDF Fitting

Given a lumitexel L and a BRDF f_r we define the error of the approximation of the reflectance properties of that point to be:

$$E_{f_r}(L) = \frac{1}{|L|} \sum_{\Re_i \in L} s * I(f_r(\hat{u}_i, \hat{v}_i) u_{i,z}, r_i) + D(f_r(\hat{u}_i, \hat{v}_i) u_{i,z}, r_i)$$

$I(r_1, r_2)$ measures intensity difference

$D(r_1, r_2)$ measures color difference

s is a weight to compensate for noisy data

Clustering

Find a set of clusters $\{C_i\}$ and associated BRDFs $\{f_i\}$ such that each f_i provides the best approximation for the set of lumitexels $\{L_k\}_i$ in C_i.

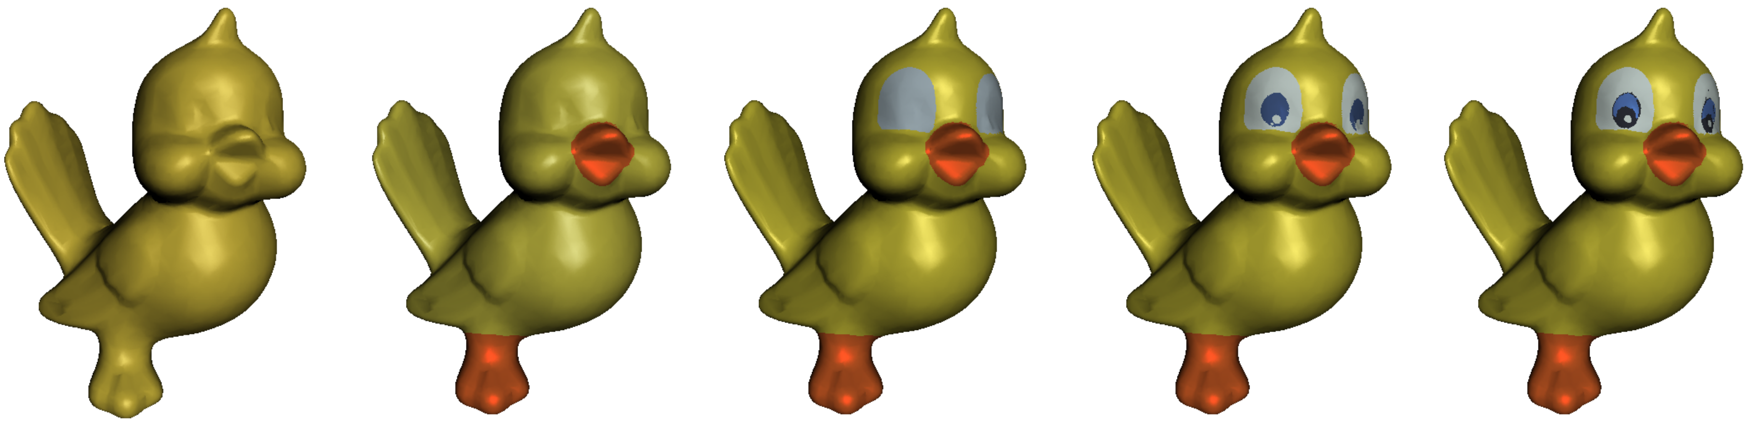

Input:

- Small proportion of the lumitexels, chosen to include samples most likely to be near specular phenomena in the image set.
- A user-specified number of materials M on the object.

Output:

- 2M-1 clusters and a BRDF representing the lumitexels in each cluster

Clustering Algorithm

Top-down, semi-hierarchical clustering:

initialize first cluster to contain all lumitexels
do
  choose the cluster C *with the largest error* $E(C_j) = \sum_{L_i \in C_j} E_{f_r}(L_i)$
  split C *into two new clusters*
  for every lumitexel L_i
    move L_i *to the cluster* C_j *st* $j = \arg\min_k E_{f_k}(L_i)$
  endfor
until 2M-1 clusters are formed

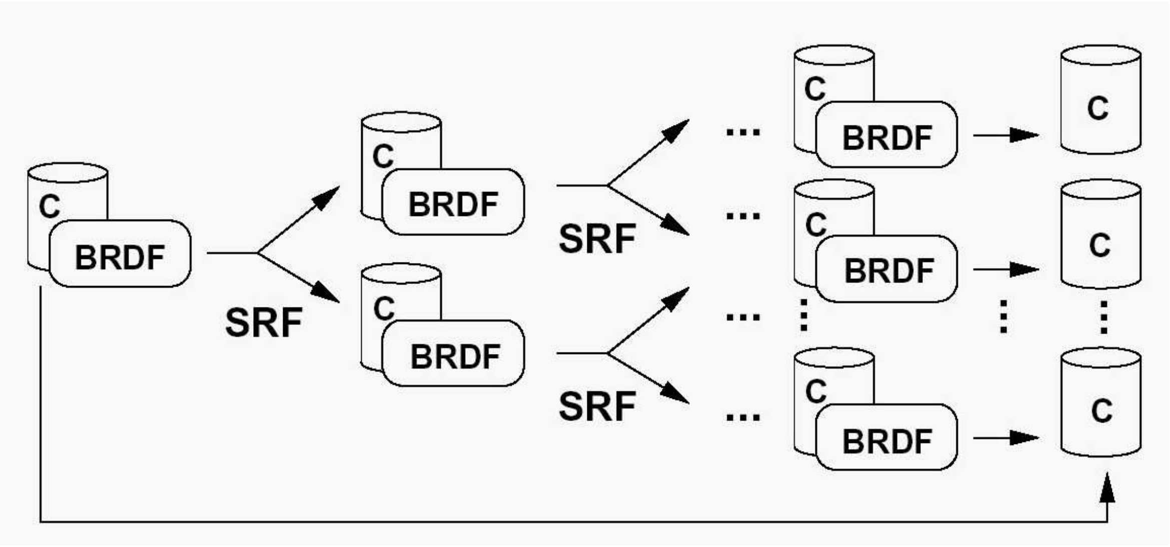

Clustering Algorithm

Splitting algorithm:

Given: *cluster C, its lumitexels $\{L_i\}$ and its BRDF f*
the largest magnitude eigenvalue λ and associated eigenvector $\vec{e}$ of the parameter covariance matrix

Create two new BRDFs:

$$f_1(\vec{a} + \tau\lambda\vec{e}; \hat{u}, \hat{v}) \qquad f_2(\vec{a} - \tau\lambda\vec{e}; \hat{u}, \hat{v})$$

do
for each lumitexel L_i
if $E_{f_1}(L_i) < E_{f_2}(L_i)$ add L_i to new cluster C_1
else add L_i to new cluster C_2
endfor
refit f_1 and f_2 to lumitexels in C_1 and C_2, respectively
until the lumitexels in each cluster do not change

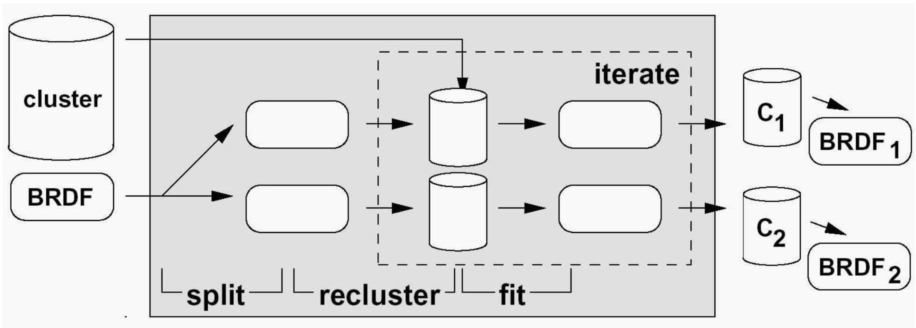

Projection

For each lumitexel L_i in cluster C_j, we find a unique BRDF:

$$f_{\pi i} = t_{i,1} f_{j,1} + t_{i,2} f_{j,2} + \ldots + t_{i,m} f_{j,m}$$

1. Choose a basis of BRDFs for each cluster.

$$\{ f_{j,1}, \ldots, f_{j,m} \}$$

2. Project each lumitexel in a cluster into its basis.
   - This is a linear least squares problem.
     - The authors solve using SVD
     - Constrain t to be non-negative

Projection

A set of basis functions for a cluster are chosen to minimize the following error function:

$$E_{\pi}(C) = \frac{1}{|C|} \sum_{L_i \in C} E_{f_{\pi i}}(L_i)$$

Authors solve using principal function analysis (PFA):

- Use following initial basis:
  - f_C, BRDF fit to the cluster C
  - Two BRDFs defived from f_C (ρ_d and N perturbed)
  - BRDFs of spatial neighbors of cluster
  - BRDFs of similar materials
- Use Levenberg-Marquardt, refitting lumitexels to projection on each iteration.
- Experiments show initial basis usually an excellent approximation of a minimum of the error function.

Rendering

What do you have?

Set of lumitexels with associated clusters and BRDFs

How can you use it?

- Authors use forward projection.
- Create texture maps for cluster membership and basis weights.

Results

model	T	V	L	R	C	B	1-RMS	C-RMS	P-RMS	F-RMS
angels	47000	27	1606223	7.6	9	6	.2953	.1163	.1113	.1111
bird	14000	25	1917043	6.3	5	4	.1513	.0627	.0387	.0387
bust	50000	16	3627404	4.2	3	4	.1025	.0839	.0583	.0581

T (triangles)
V (views)
L (lumitexels)

R (samples/lumitexel)
C (clusters)
B (BRDFs per cluster)

C-RMS (error for per cluster BRDF)
P-RMS (error after projection)
F-RMS (error after PFA)

Actual Image

Rendered Image