Bayesian Nonparametric Intrinsic Image Decomposition

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Jason Chang | Randi Cabezas | John W. Fisher III

Shading

Reflectance

[Grosse et al. ‘09]

 CSAIL

Previous Approaches

- Single Image
  - [Barron & Malik CVPR’11, CVPR’12, ECCV’12]
  - [Barrow & Tenenbaum CVS’78]
  - [Bell & Freeman ICCV’01]
  - [Gehler et al. NIPS‘11]
  - [Grosse et al. ICCV’09]
  - [Land and McCann JOSA‘71]
  - [Shen et al. CVPR’08, CVPR‘11]
  - [Tappen et al. TPAMI‘05, CVPR’06, CVPR‘07]
  - [Zhao et al. TPAMI 2012]
- Multiple Image
  - [Weiss ICCV’01]
- Object Specific
  - [Li et al. ECCV ‘14]
- Different Media
  - [Jeon et al. ECCV’14]
  - [Kong et al. ECCV’14]
- Many more…

CSAIL

Previous Approach: SIRFS

Models 3D shape, normals, and lighting

(a) Input Image (b) Modified shape (c) Modified reflectance (d) Modified light (e) Modified orientation

Is 3D necessary?

[Barron & Malik CVPR’11, CVPR’12, ECCV’12]

Previous Approach: Retinex

1. Image gradients should match reflectance gradients at edges

[Land and McCann '71]

Previous Approach: Retinex

1. Image gradients should match reflectance gradients at edges

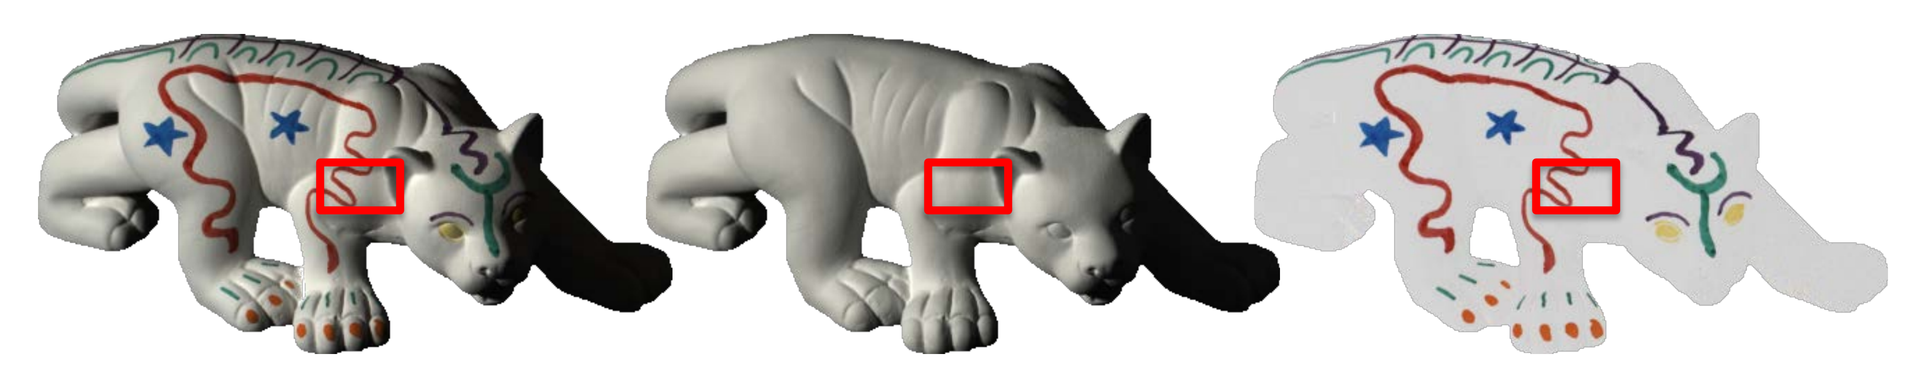

[Land and McCann '71]

Previous Approach: Retinex

1. Image gradients should match reflectance gradients at edges

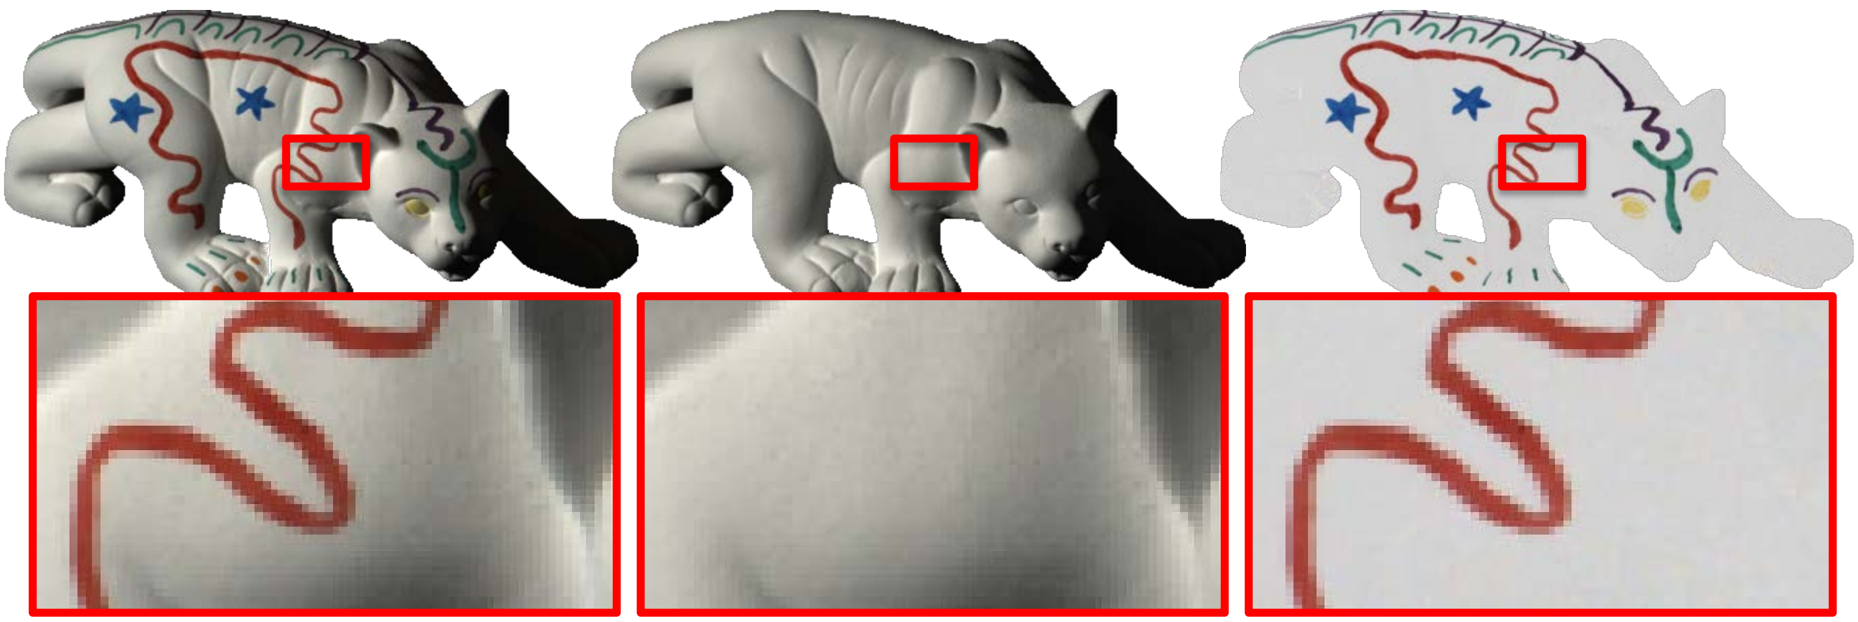

[Land and McCann ‘71]

Previous Approach: Retinex

1. Image gradients should match reflectance gradients at edges

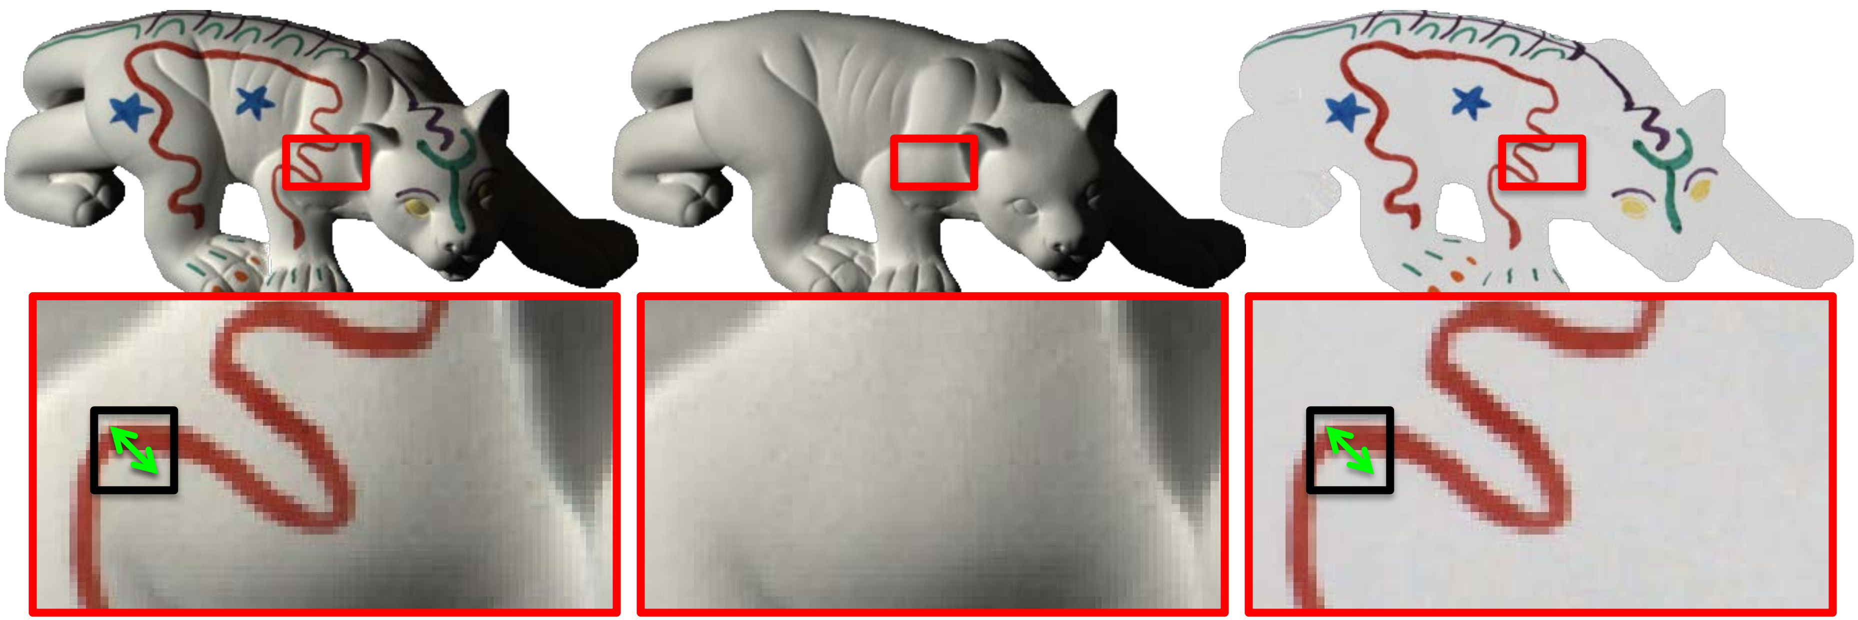

[Land and McCann '71]

Previous Approach: Retinex

1. Image gradients should match reflectance gradients at edges
2. Shading should be smooth

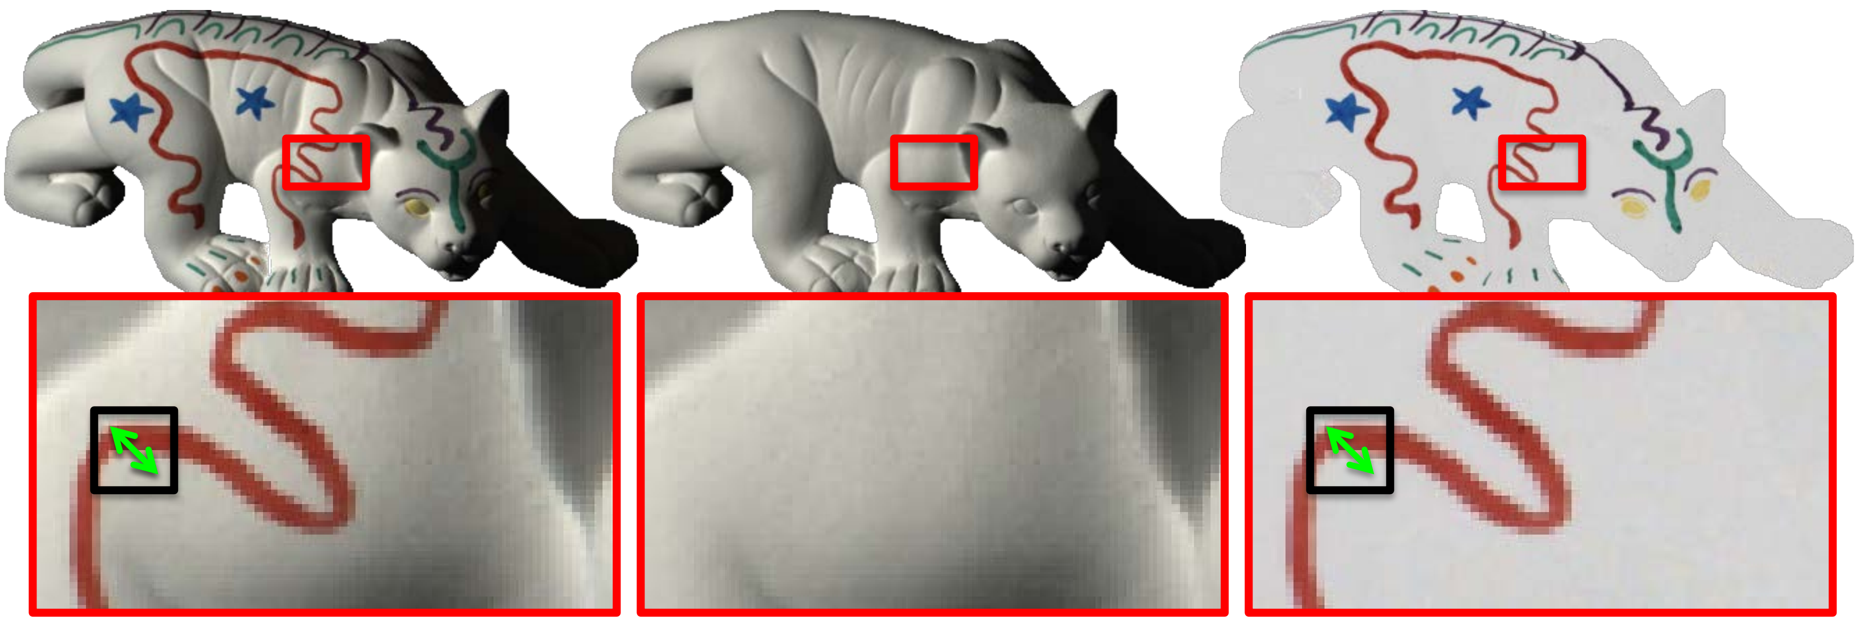

[Land and McCann ‘71]

Previous Approach: Retinex

1. Image gradients should match reflectance gradients at edges
2. Shading should be smooth

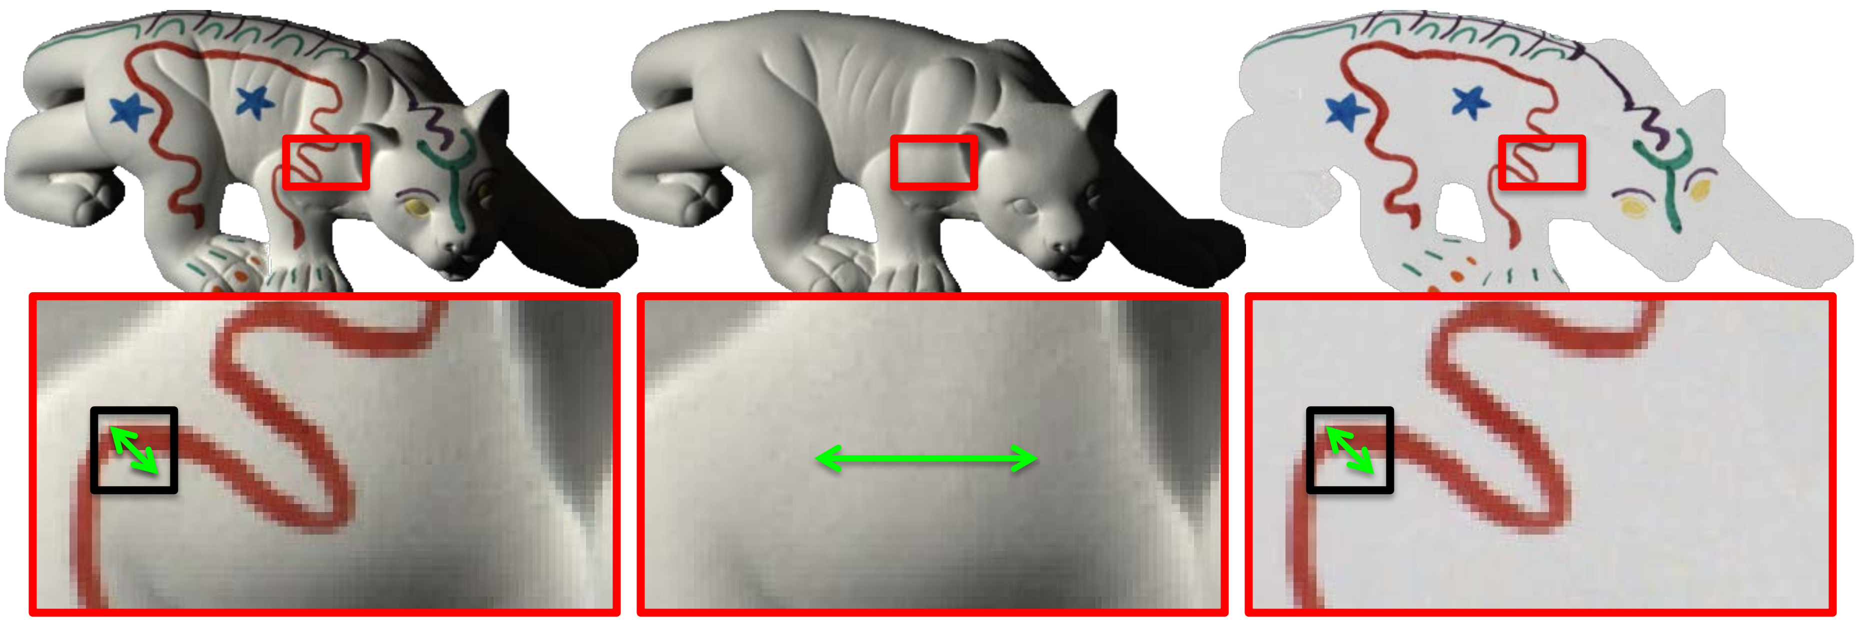

[Land and McCann '71]

Advancement: Reflectance Sparsity

1. Image gradients should match reflectance gradients at edges
2. Shading should be smooth
3. Reflectance should be sparse
   - [Barron & Malik CVPR’11, CVPR’12, ECCV’12]
   - [Gehler et al. NIPS’11]
   - [Shen et al. CVPR’ 08], [Shen & Yeo CVPR’11]
   - [Zhao et al. TPAMI’12]

Previous Approach: Gehler et al. NIPS'11

1. Image gradients should match reflectance gradients at edges
2. Shading should be smooth
3. Reflectance values come from a sparse set of colors

Is gradient matching necessary?

Previous Approach: Gehler et al. NIPS’11

1. K-Means clustering
2. Find cluster means
3. Optimize for shading (4-connected GMRF)
4. Repeat until convergence

How to set K?

Smoothness assumption?

Bayesian Nonparametrics

Observation


Shading
x
Reflectance
=
Observation

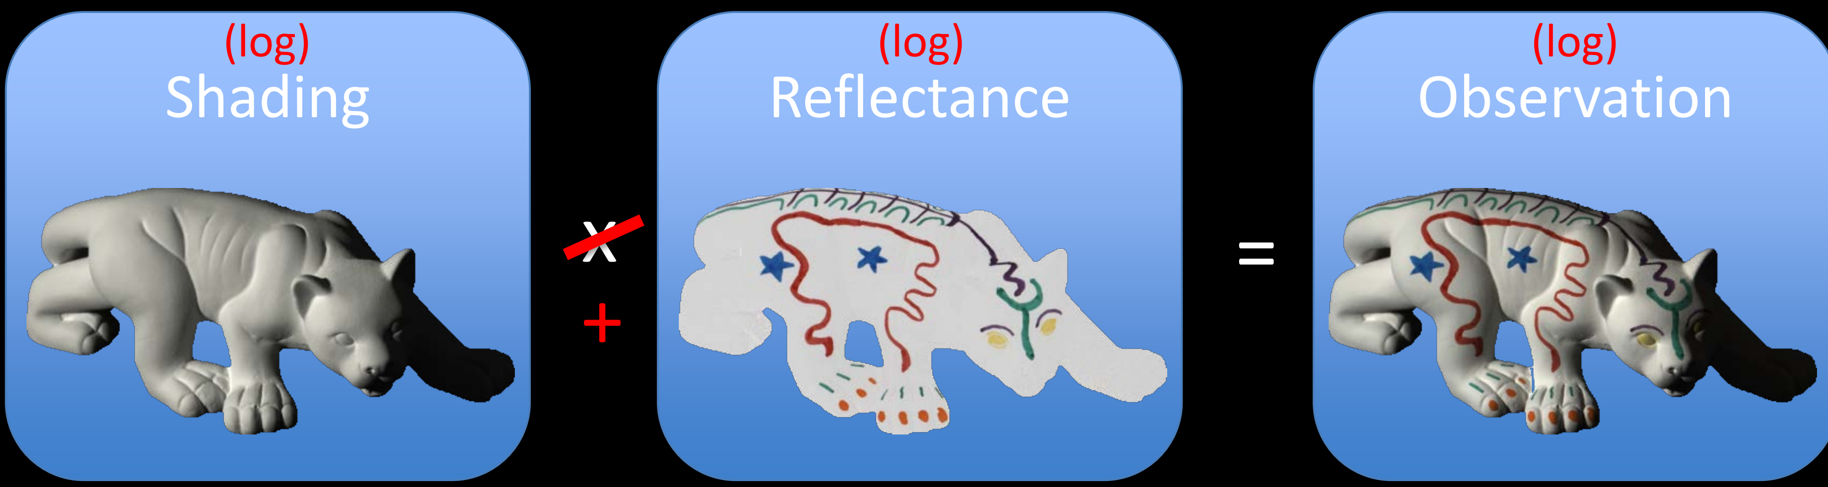
(log)
Shading
~~x~~
+
(log)
Reflectance
=
(log)
Observation

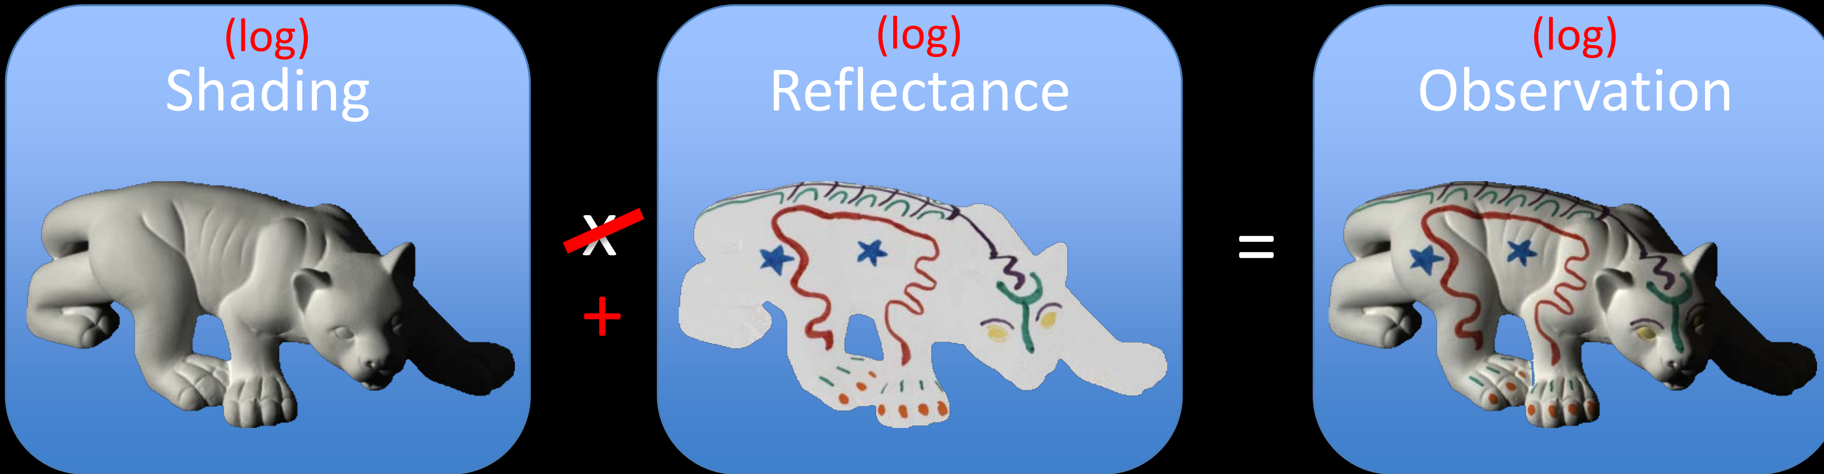

<u>Our Idea</u>

Gaussian process shading image

+

Dirichlet process Gaussian Mixture Model reflectance image

~~x~~ +

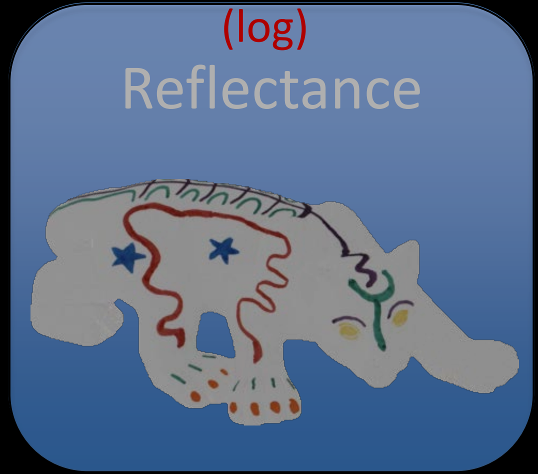

=

~~GMRF~~ Gaussian Process

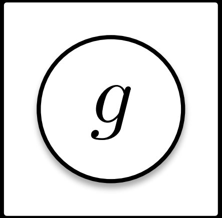

Mean

$$g \sim \mathrm{GP}(0, \kappa) \rightarrow$$ Controls smoothness

Covariance kernel

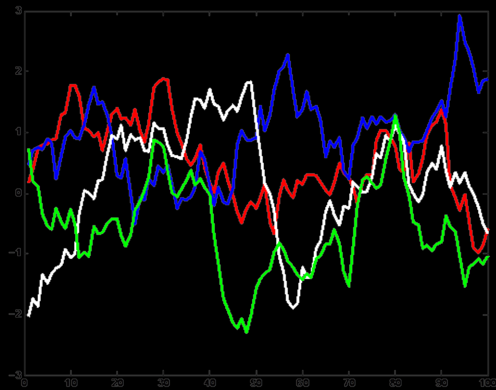

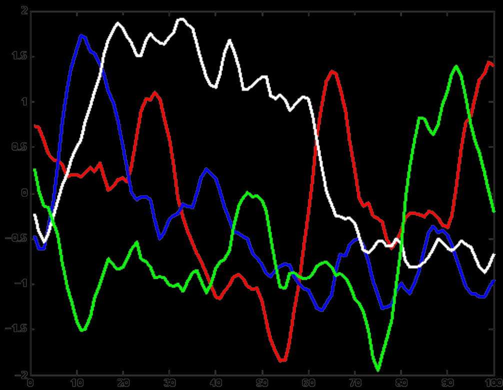

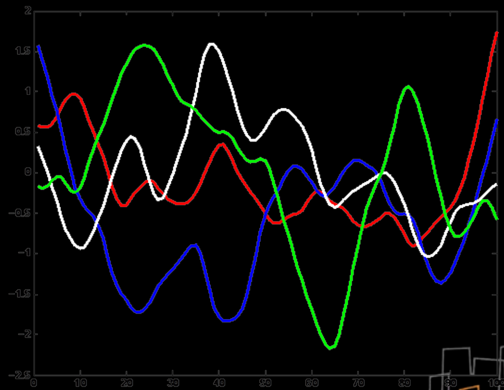

← Smoothness vs. κ →

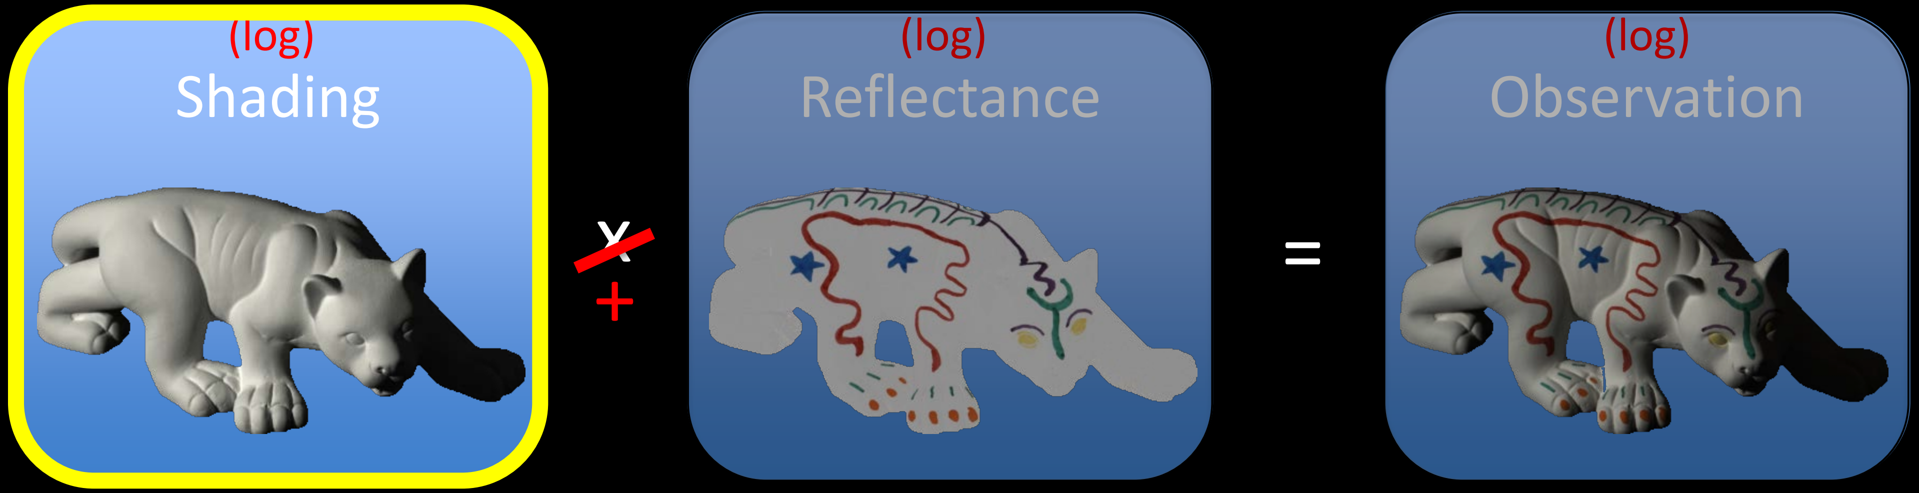

Gaussian Process 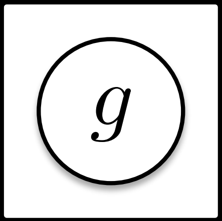

$$g \sim \mathrm{GP}(0, \kappa)$$

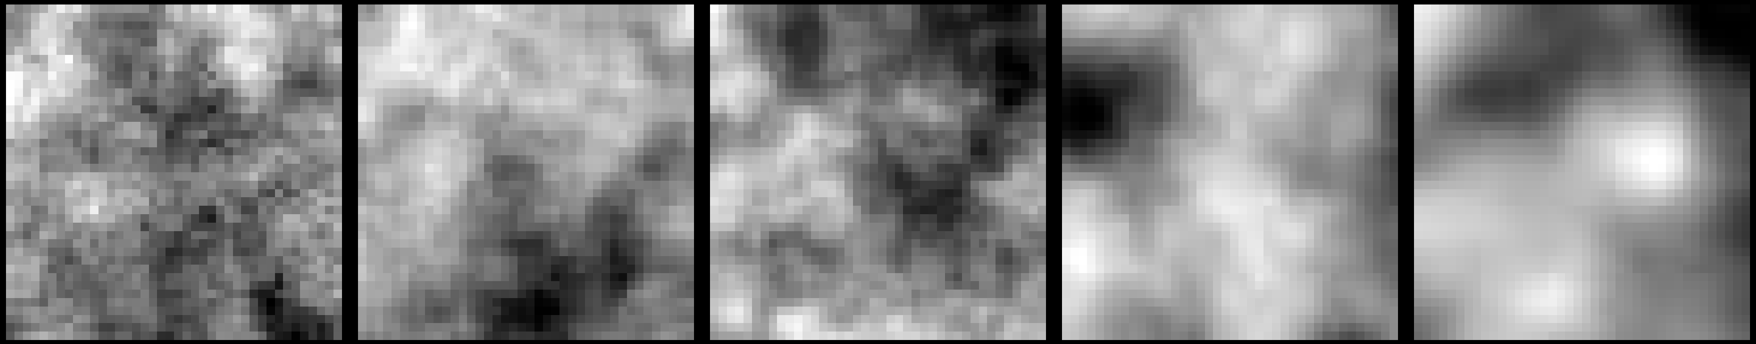

← Smoothness vs. κ →

Dirichlet Process Gaussian Mixture Model (DPGMM)

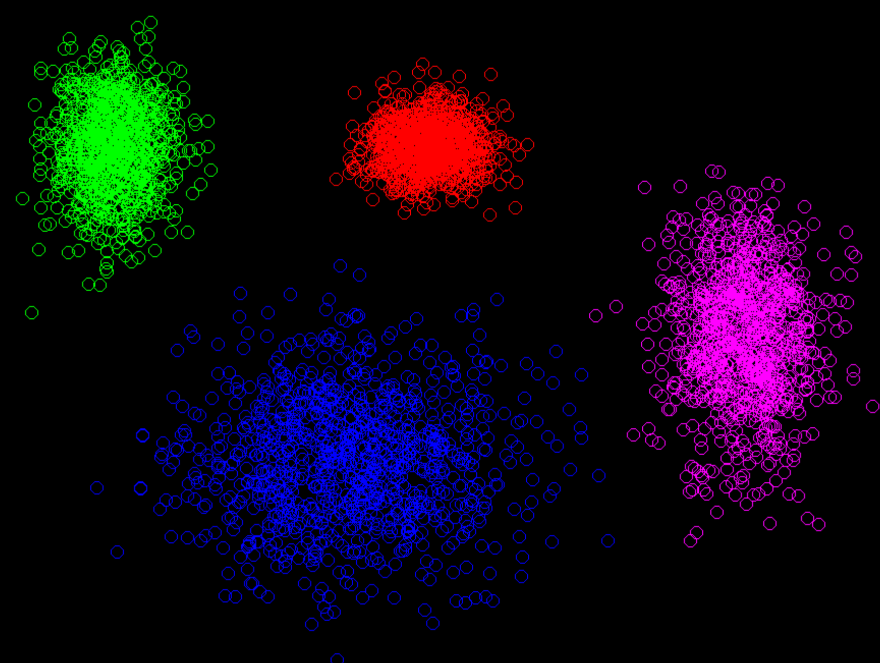

Dirichlet Process Gaussian Mixture Model (DPGMM)

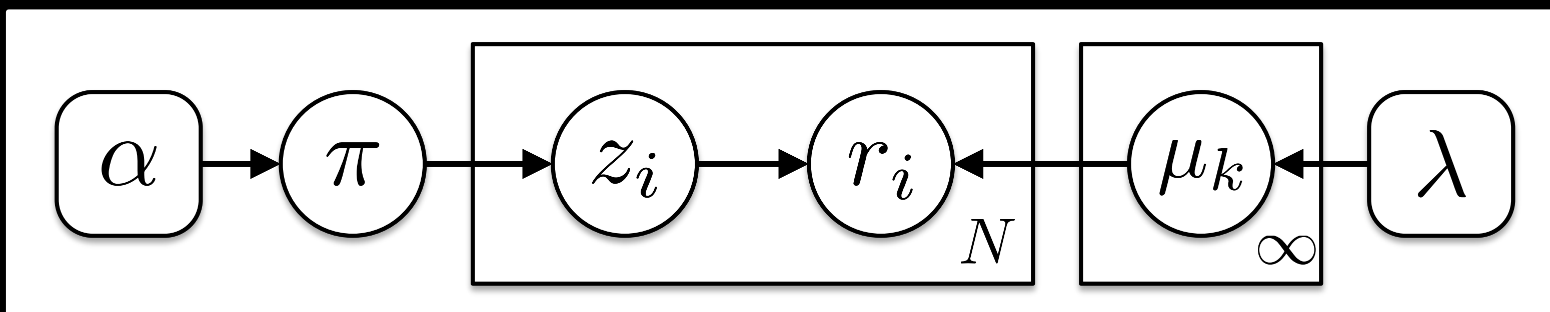

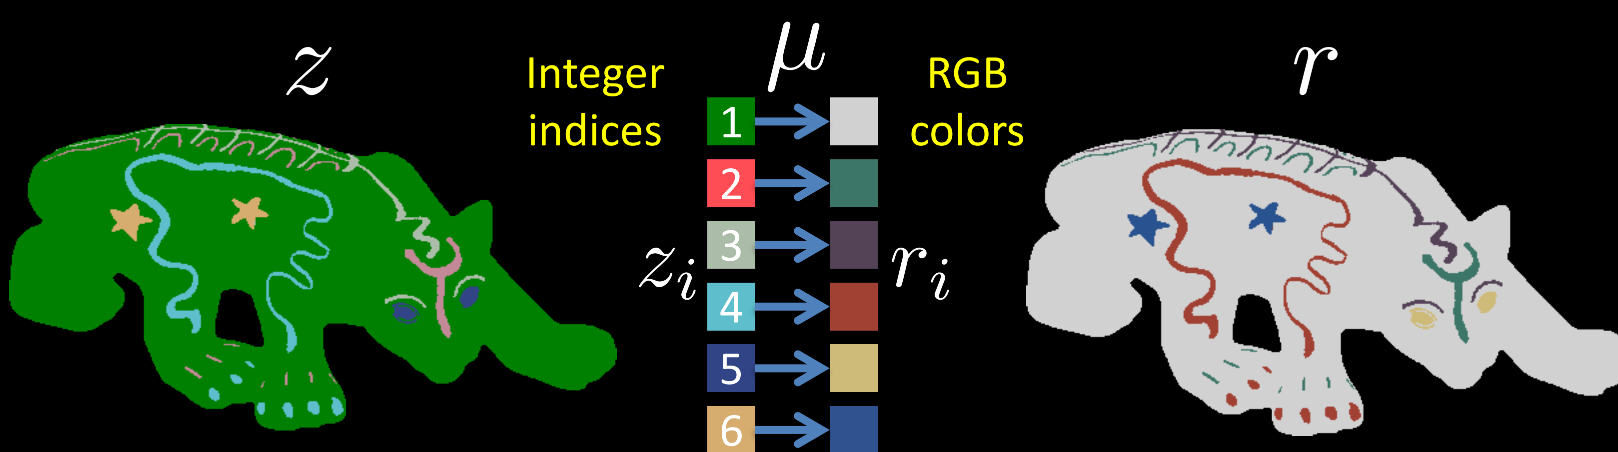

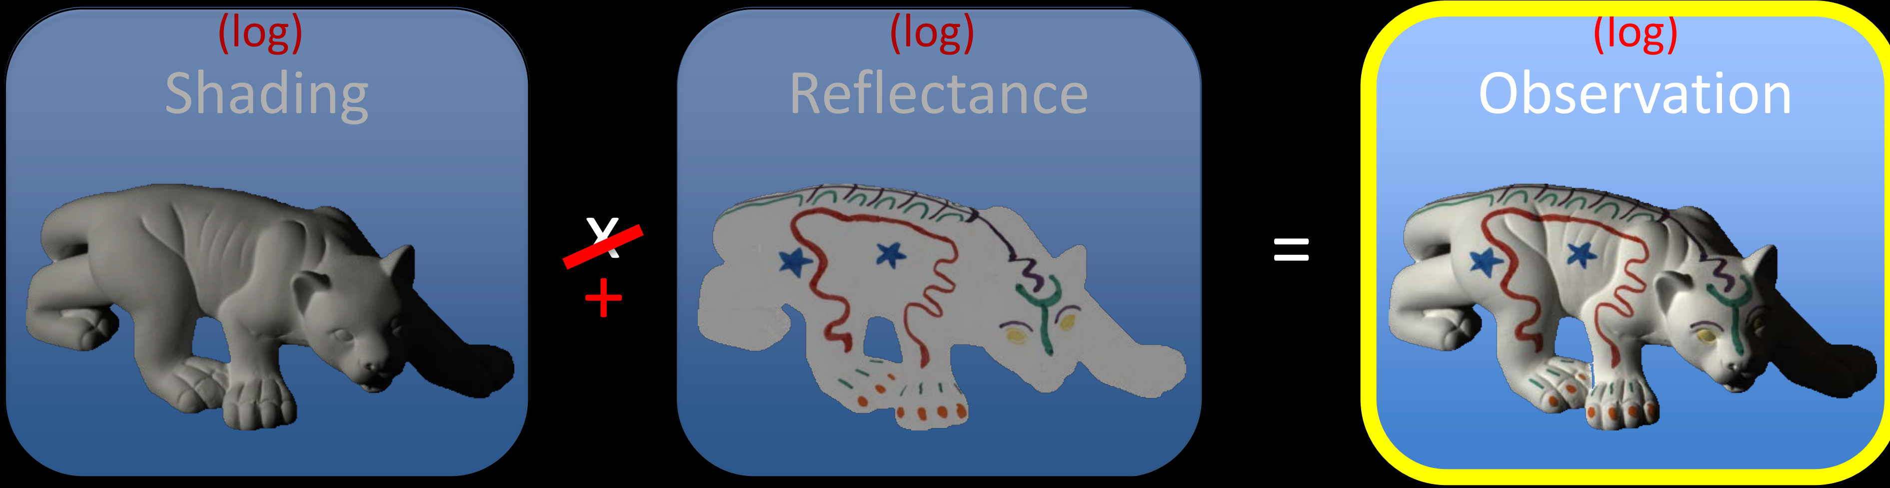

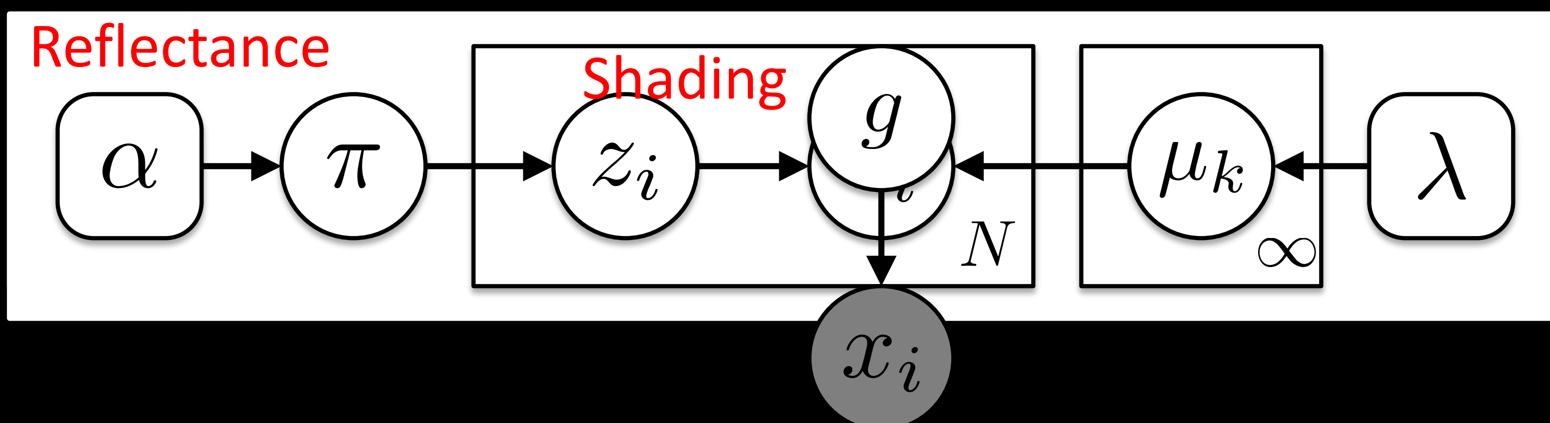

Observation

$$x_i \sim \mathcal{N}(x_i\,;\, \underbrace{g}_{\text{GP}} + \underbrace{\mu_{z_i}}_{\text{DPGMM}}\,;\, \Sigma)$$

$$x_i \sim \mathcal{N}(x_i\,;\,g + \mu_{z_i}\,;\,\Sigma)$$

$$x_i \sim \mathcal{N}(x_i \,;\, g \overset{\rightleftarrows}{+} \mu_{z_i} \,;\, \Sigma)$$

Unknowns: g, μ, z

g – GP Shading Image
μ – Reflectance colors
z – Cluster assignment
x – Observed image

Iterative Inference:
(used by [Gehler et al. 2011])

$$\begin{cases} g \mid \mu, z, x & \text{GP Regression} \\ \mu, z \mid g, x & \text{DPGMM} \end{cases}$$

-3

+3

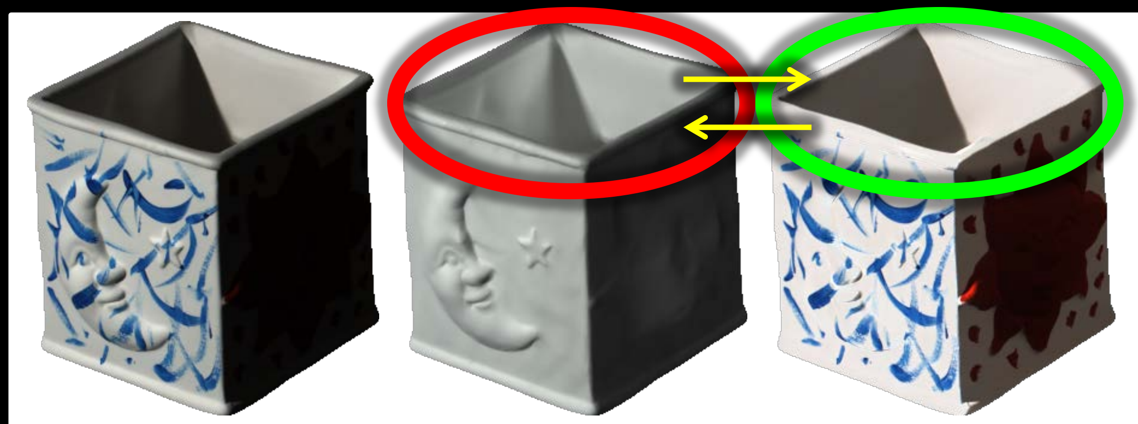

$$x_i \sim \mathcal{N}(x_i\,;\, \overset{-3 \rightleftarrows +3}{g} + \mu_{z_i}\,;\, \Sigma)$$

Unknowns: g, μ, z

g – GP Shading Image
μ – Reflectance colors
z – Cluster assignment
x – Observed image

Iterative Inference:
(used by [Gehler et al. 2011])

$$\begin{cases} g \mid \mu, z, x & \text{GP Regression} \\ \mu, z \mid g, x & \text{DPGMM} \end{cases}$$

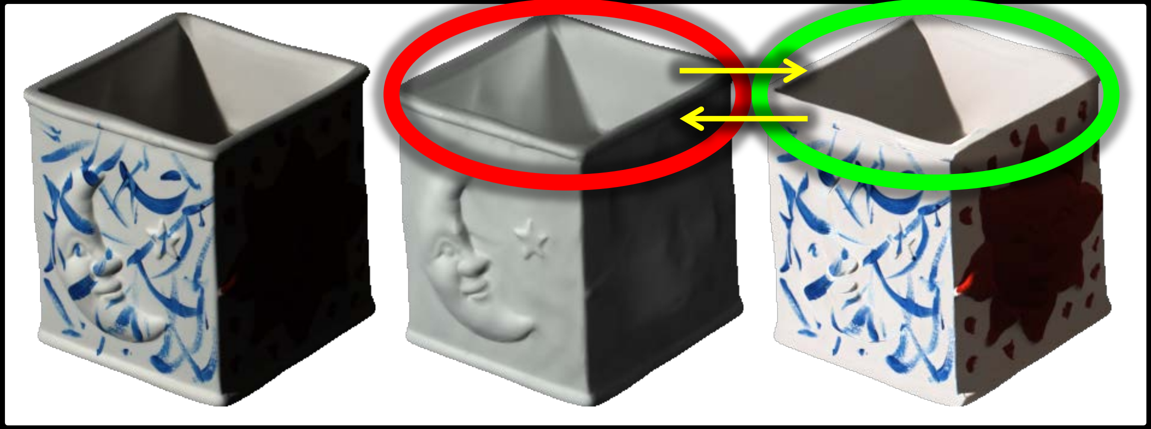

$$x_i \sim \mathcal{N}(x_i\,;\, g + \mu_{z_i}\,;\, \Sigma)$$

g – GP Shading Image
μ – Reflectance colors
z – Cluster assignment
x – Observed image

Unknowns: g, μ, z

~~**Iterative Inference:**~~
~~(used by [Gehler et al. 2011])~~

$$\cancel{\begin{cases} g \mid \mu, z, x & \text{GP Regression} \\ \mu, z \mid g, x & \text{DPGMM} \end{cases}}$$

Marginalized MCMC:

1. Marginalize $g, \mu,$ infer $z|x$
2. Large split/merge moves

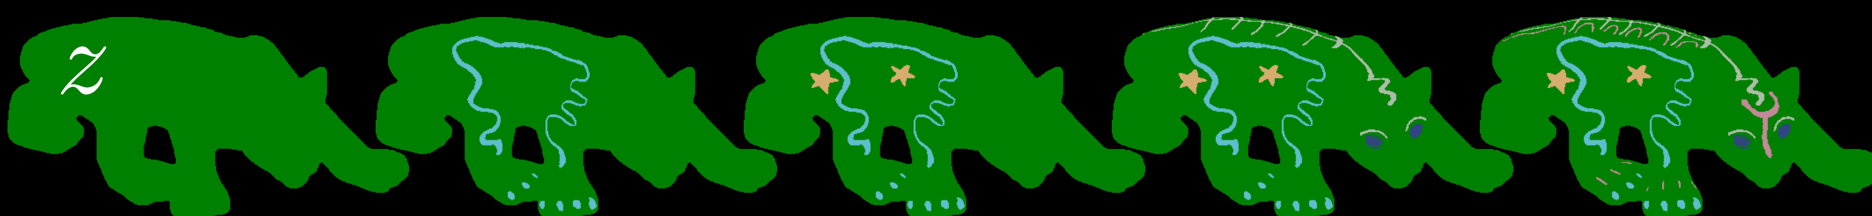

Itreation

Details at poster and in paper

Improved Results

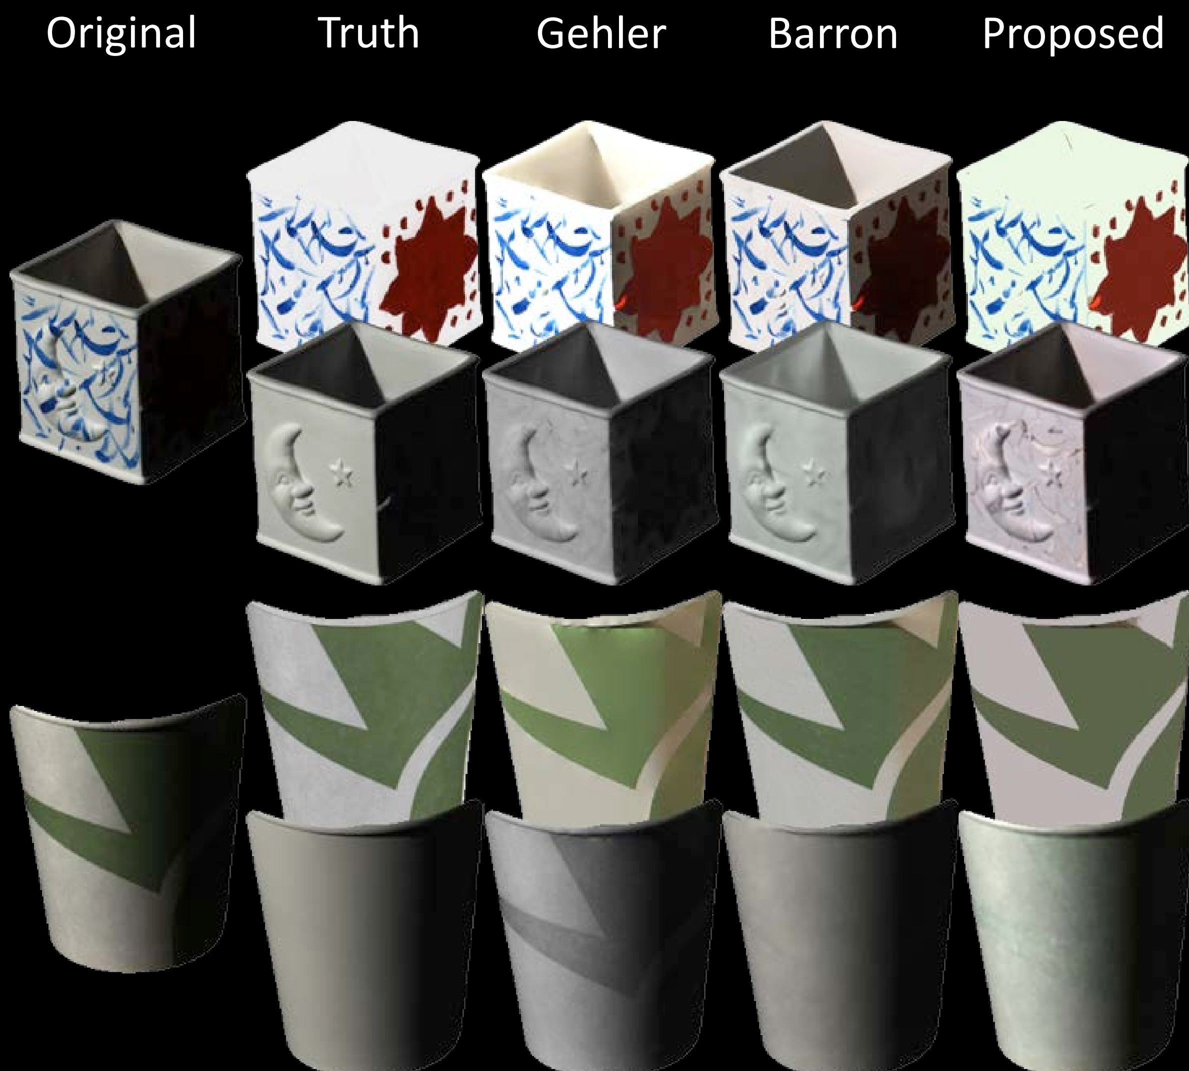

[Grosse et al. ‘09], [Barron & Malik ‘13], [Gehler et al. ‘11]

Improved Results

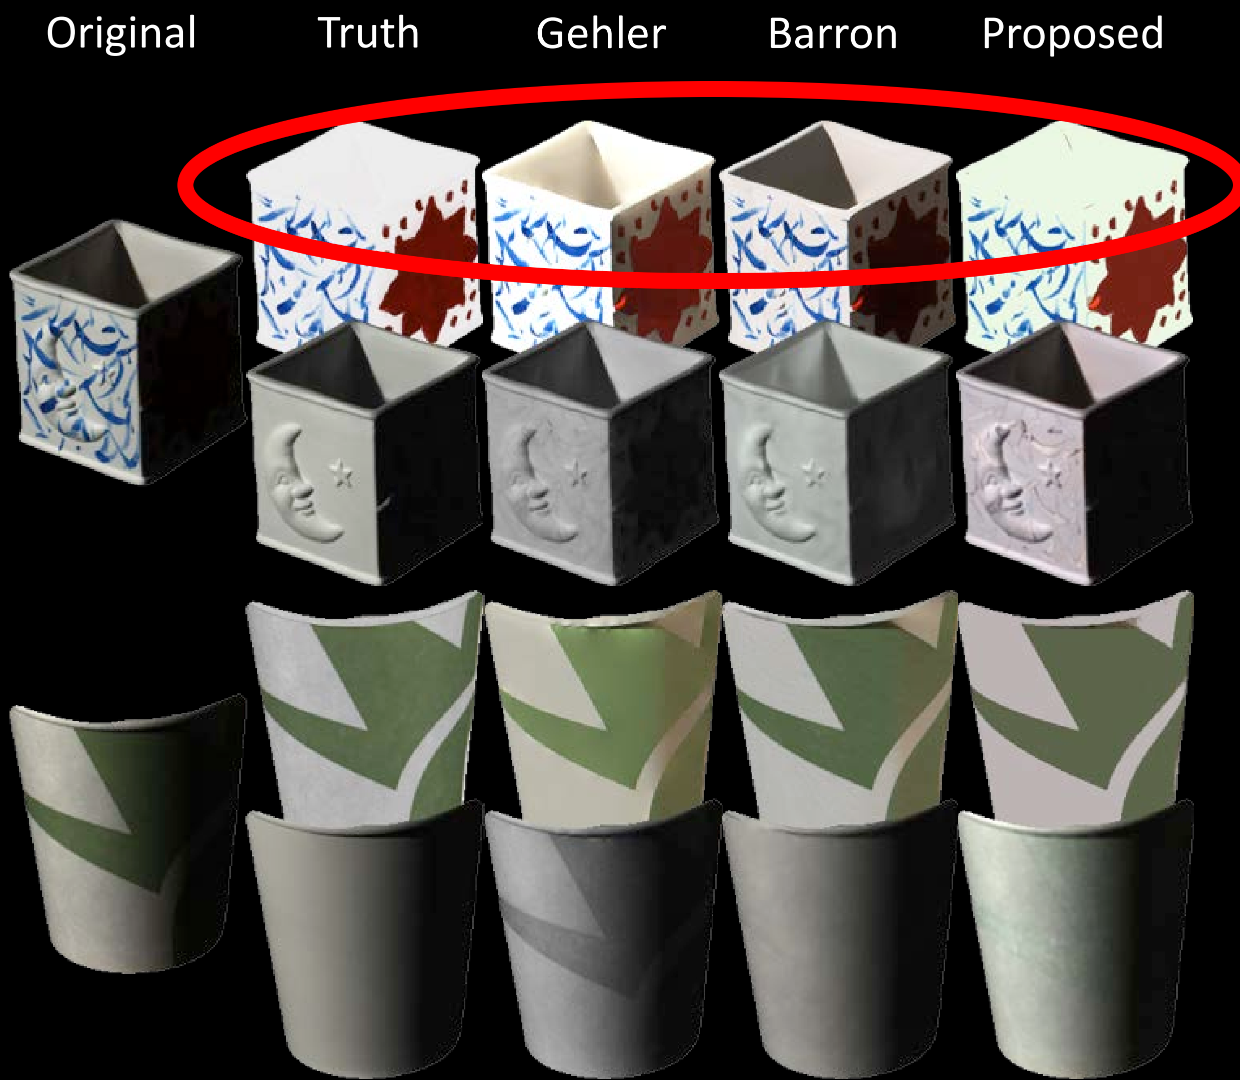

[Grosse et al. ‘09], [Barron & Malik ‘13], [Gehler et al. ‘11]

CSAIL

Improved Results

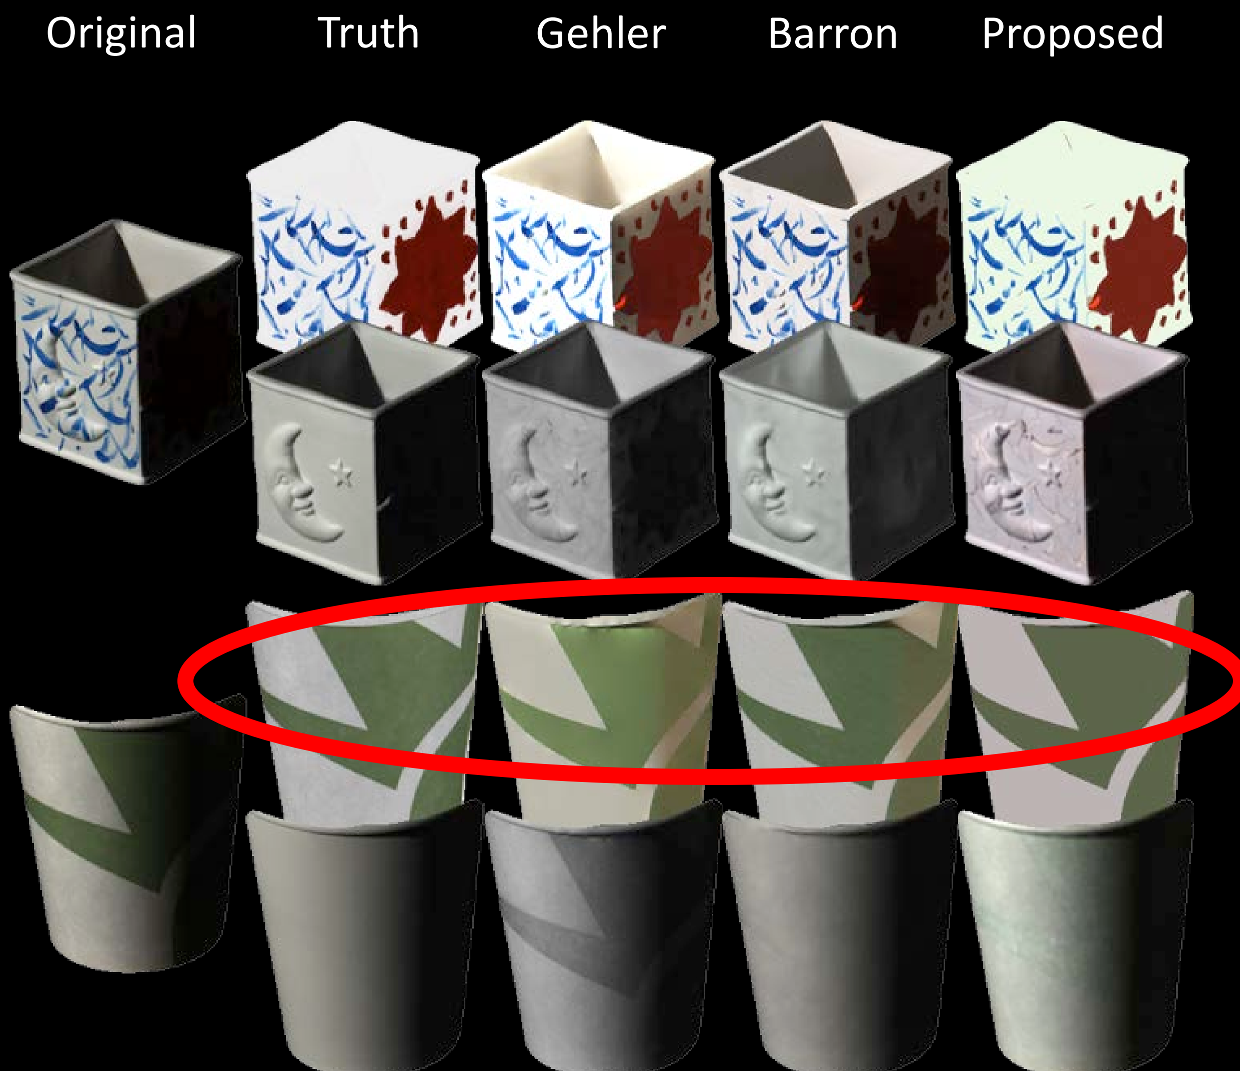

[Grosse et al. '09], [Barron & Malik '13], [Gehler et al. '11]

Improved Results

Original Truth Gehler Barron Proposed

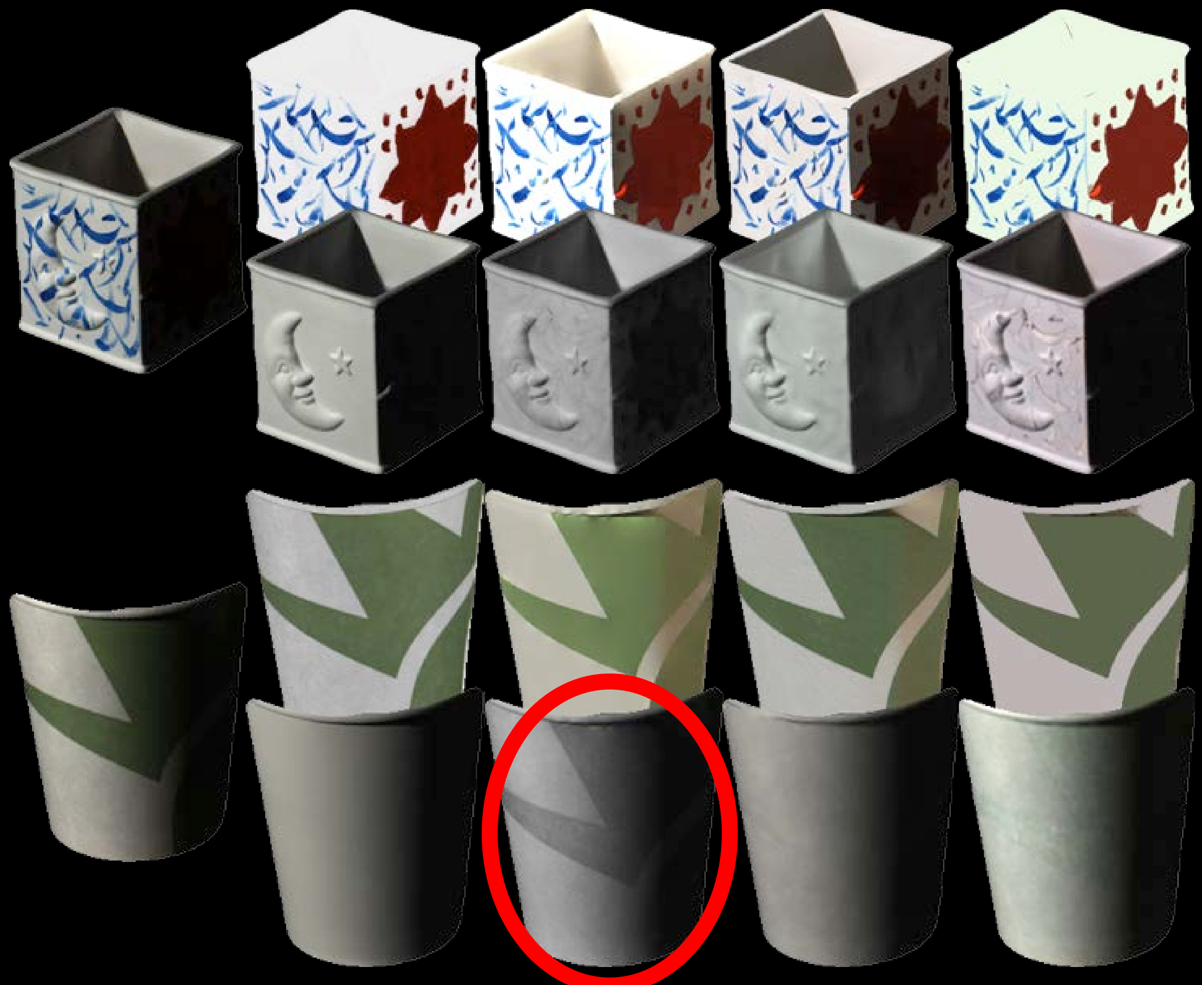

[Grosse et al. ‘09], [Barron & Malik ‘13], [Gehler et al. ‘11]

Failure Cases

[Grosse et al. ‘09], [Barron & Malik ‘13], [Gehler et al. ‘11]

Failure Cases

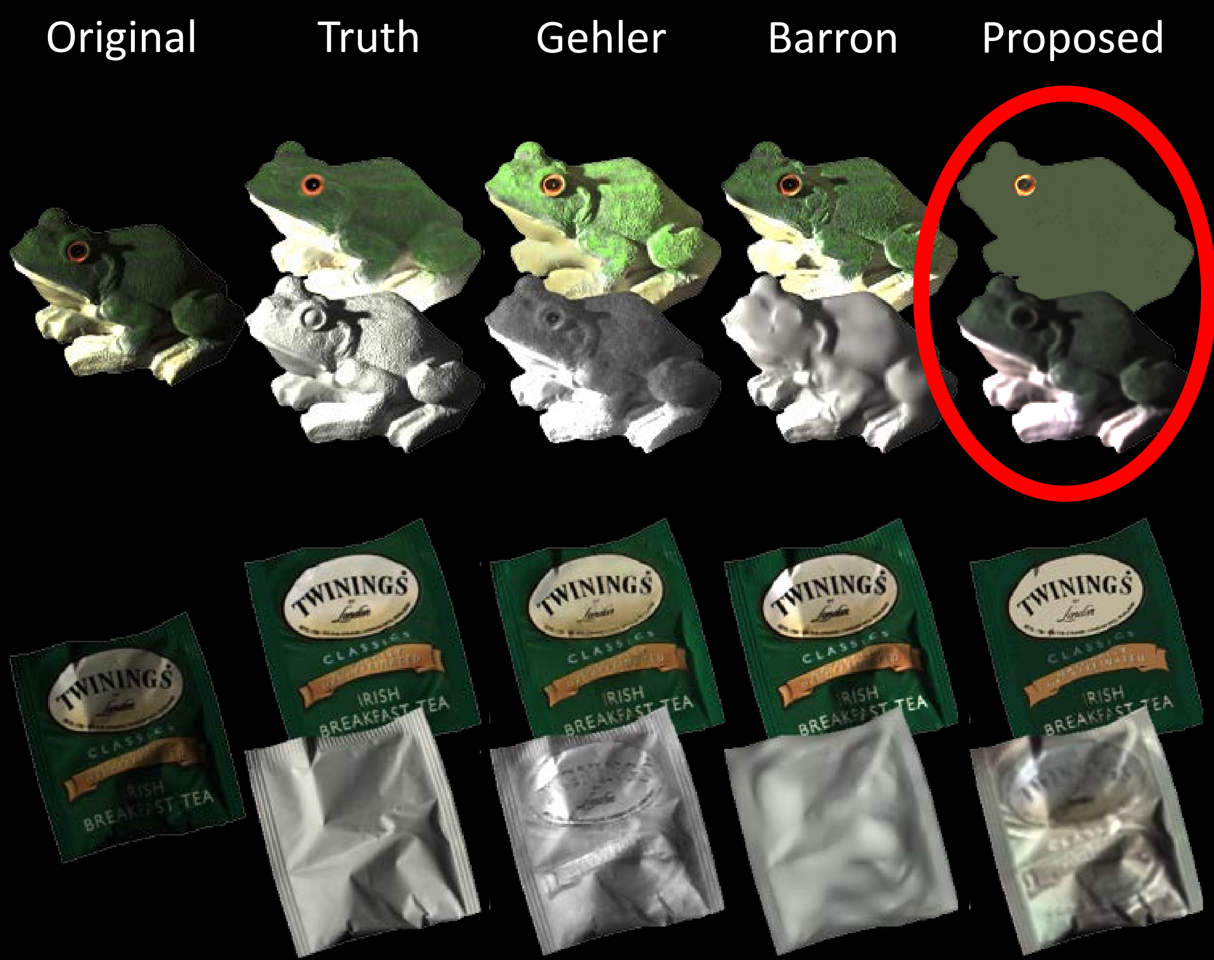

[Grosse et al. ‘09], [Barron & Malik ‘13], [Gehler et al. ‘11]

Failure Cases

Original | Truth | Gehler | Barron | Proposed

[Grosse et al. ‘09], [Barron & Malik ‘13], [Gehler et al. ‘11]

Results

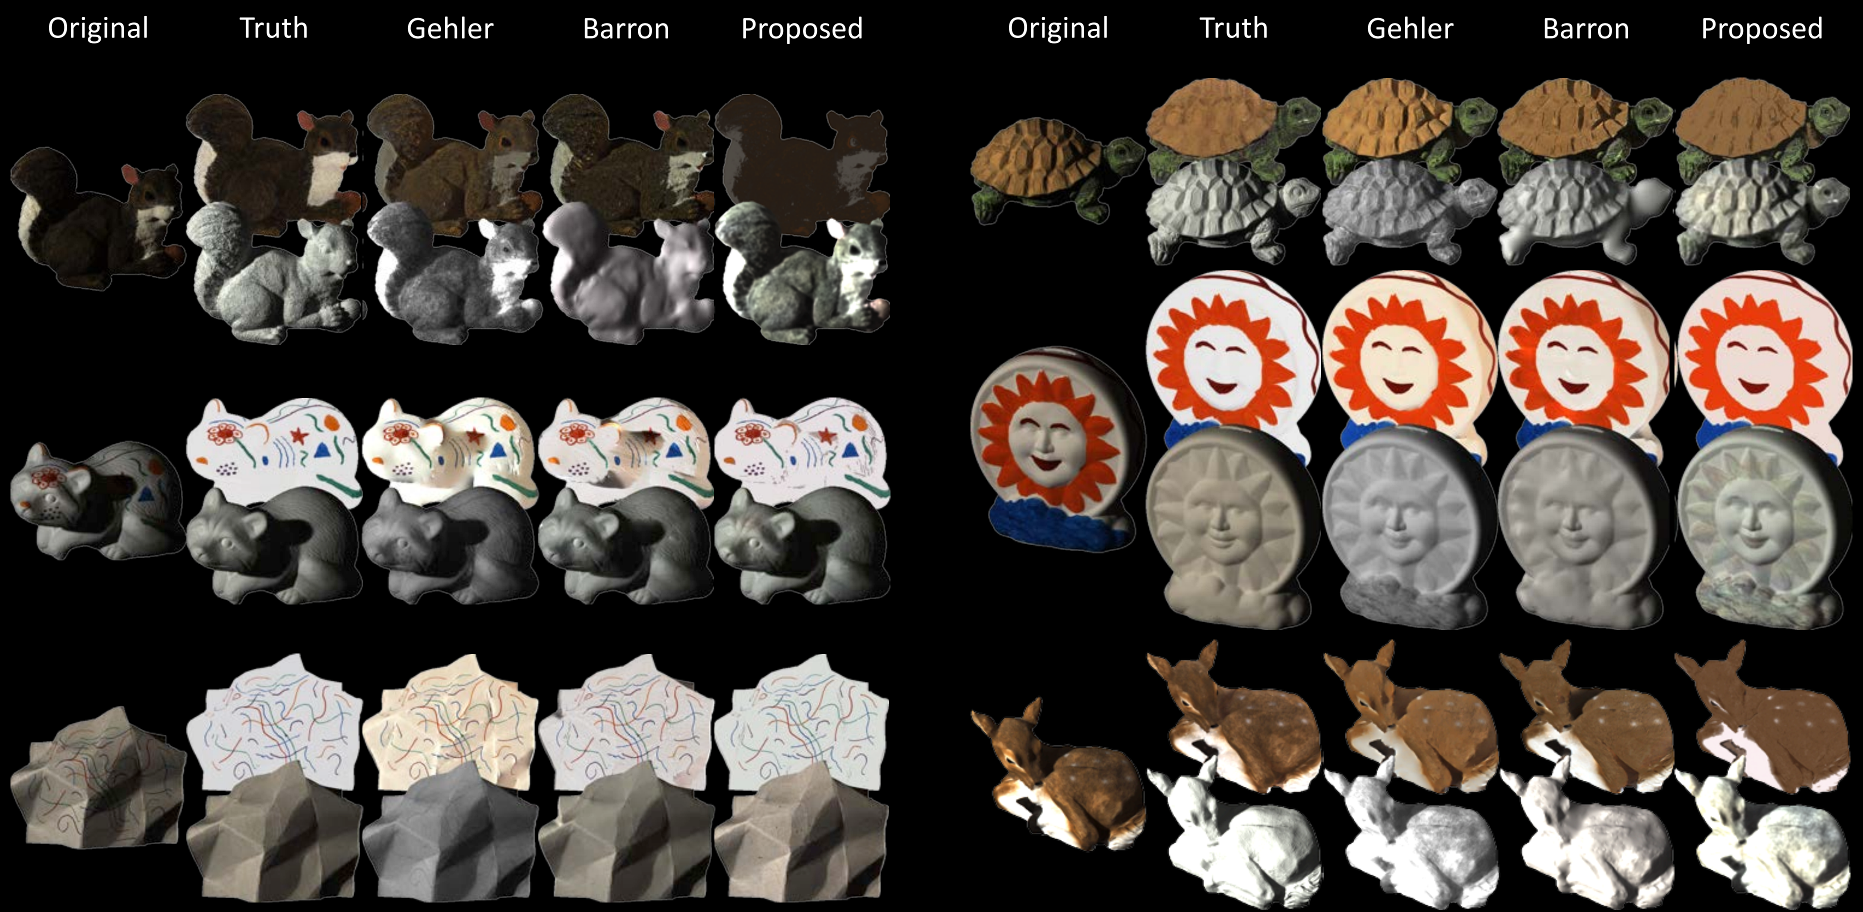

[Grosse et al. ‘09], [Barron & Malik ‘13], [Gehler et al. ‘11]

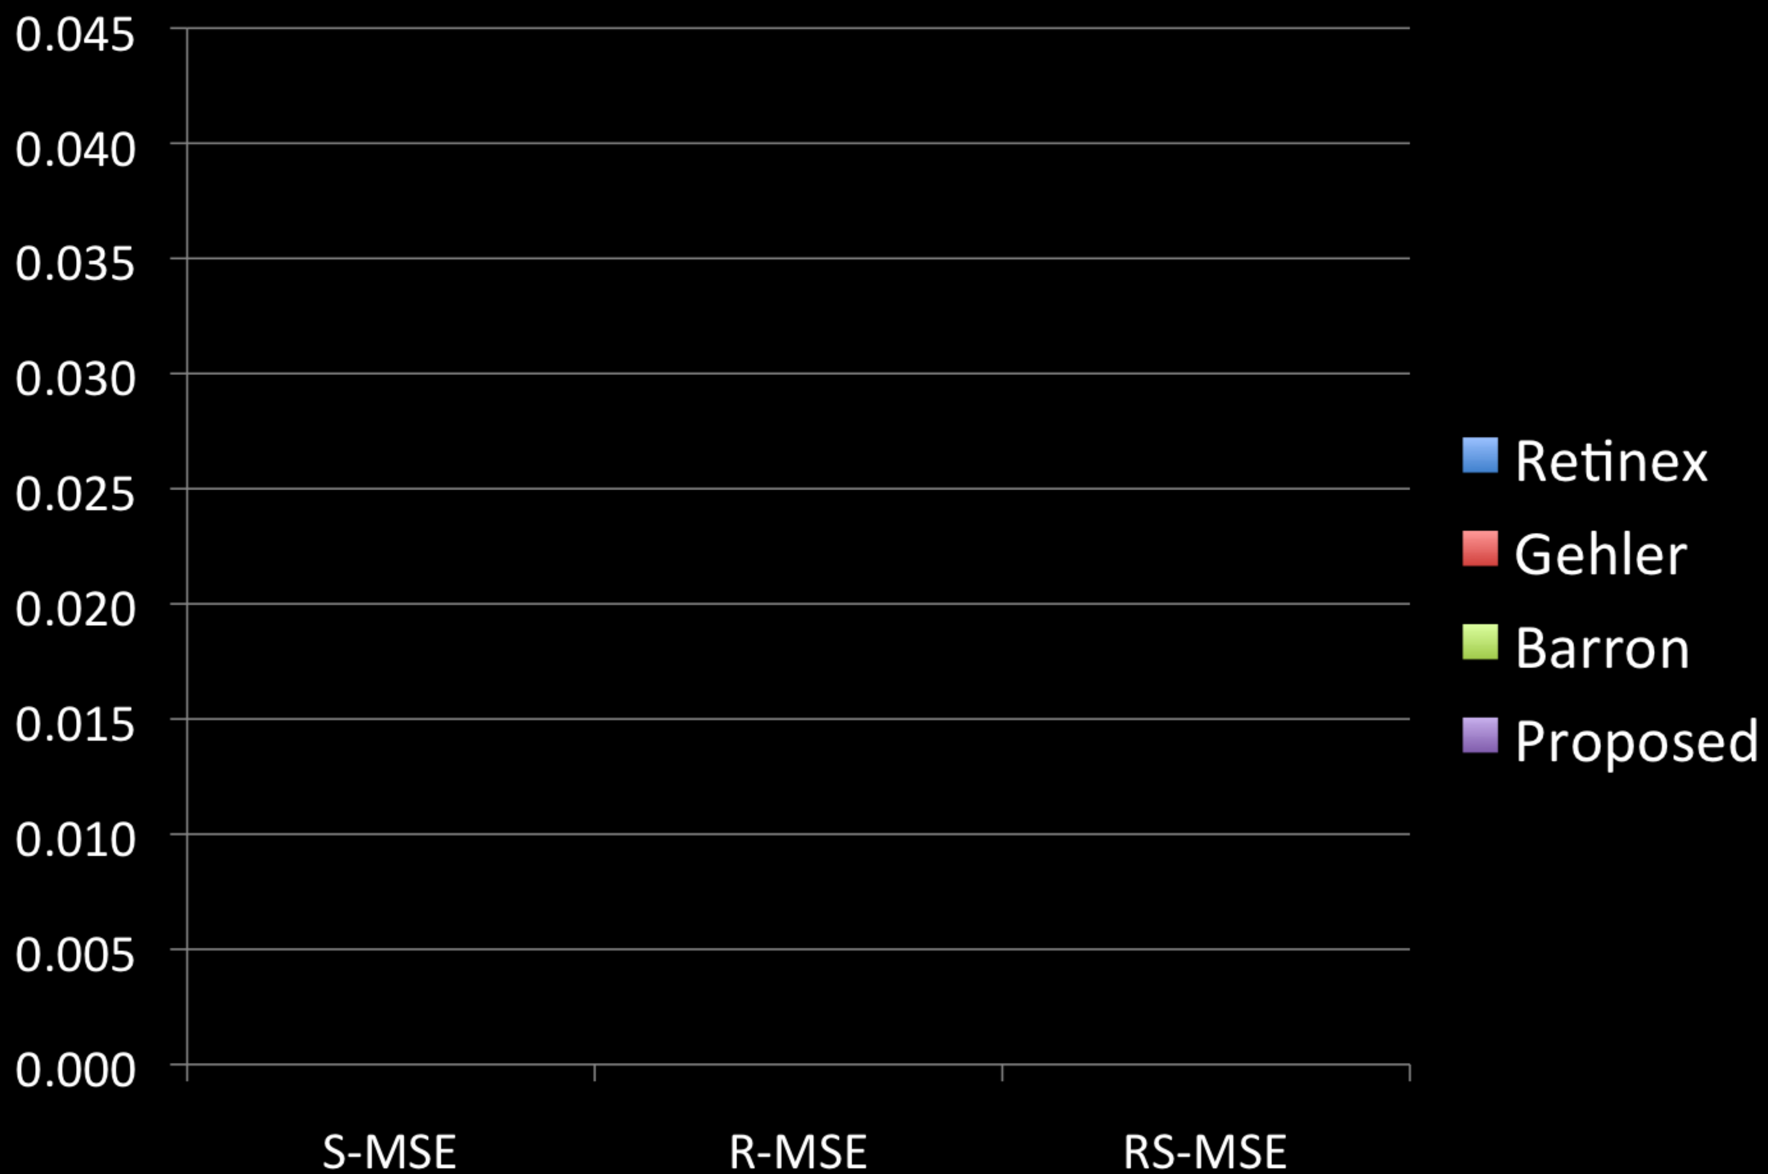

[Grosse et al. ‘09], [Land & McCann ‘71], [Barron & Malik ‘13], [Gehler et al. ’11]

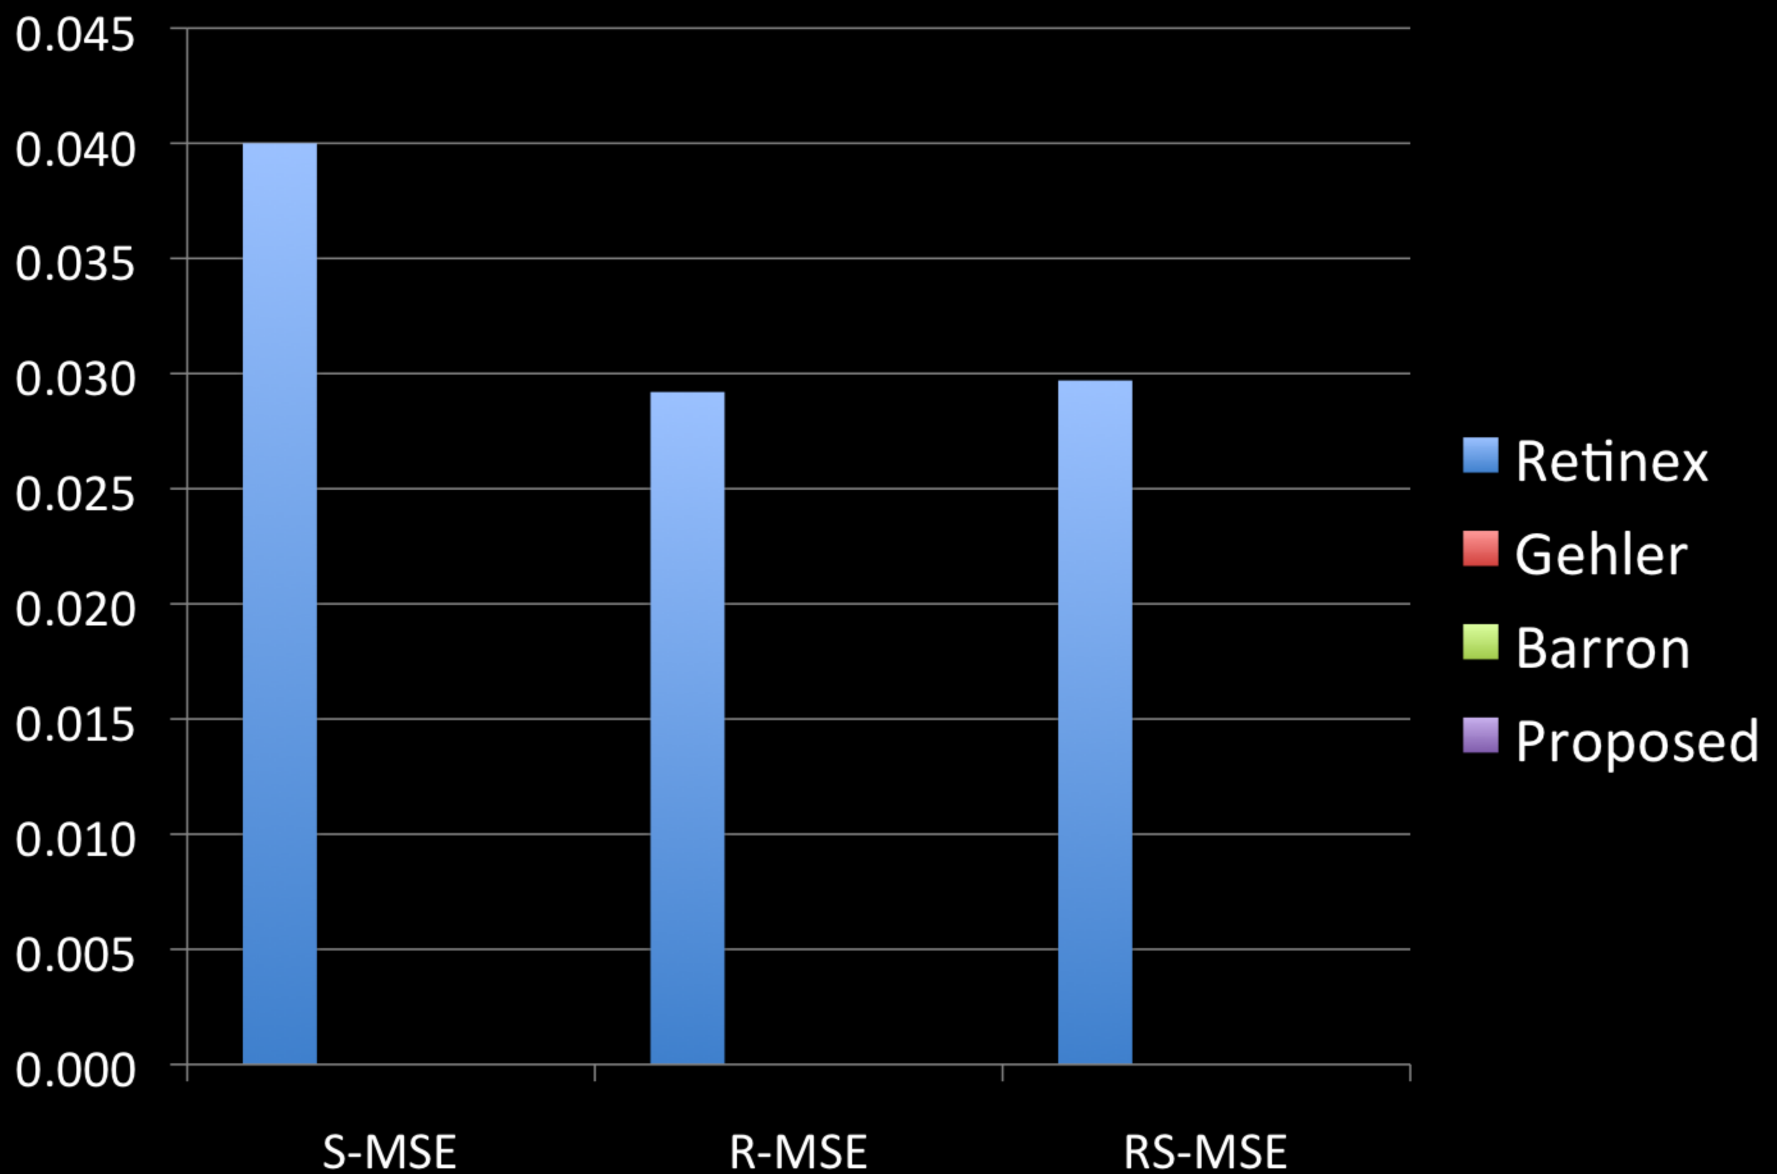

[Grosse et al. ‘09], [Land & McCann ‘71], [Barron & Malik ‘13], [Gehler et al. ’11]

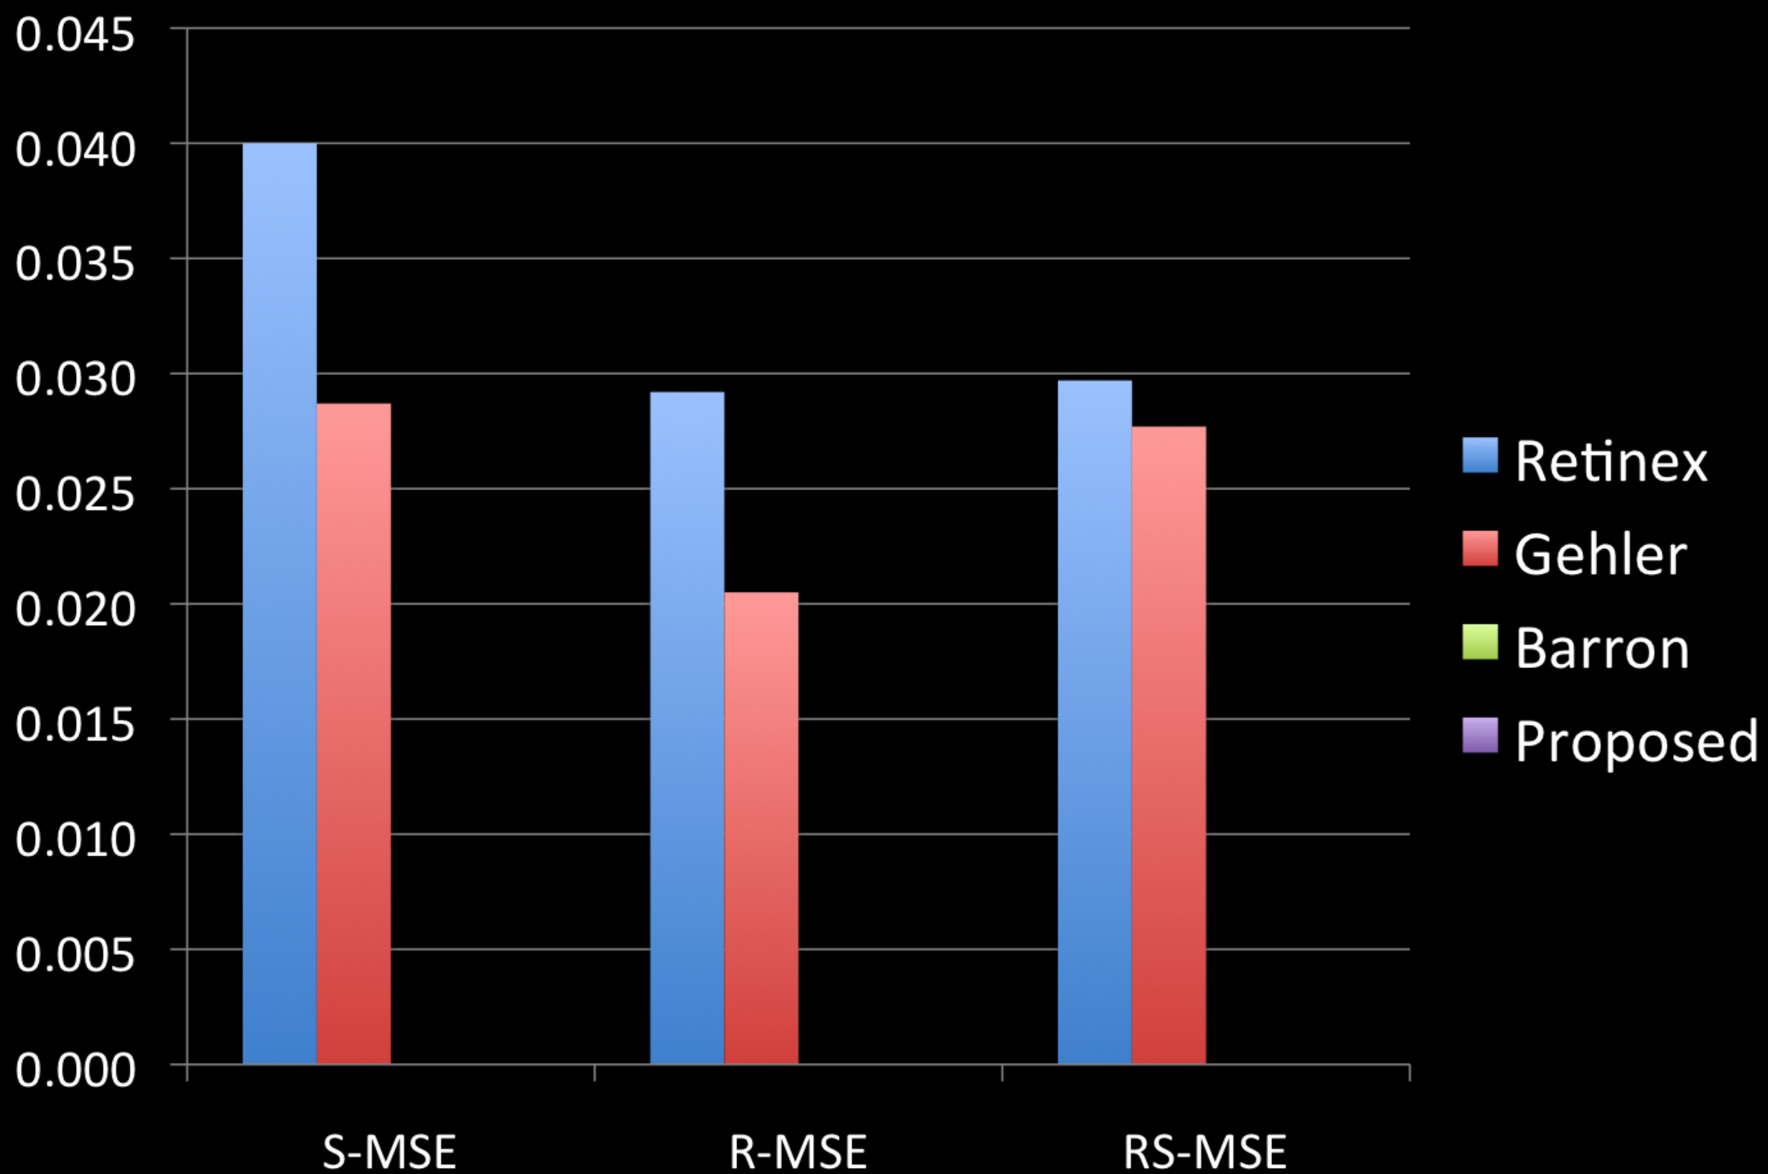

[Grosse et al. ‘09], [Land & McCann ‘71], [Barron & Malik ‘13], [Gehler et al. ’11]

CSAIL

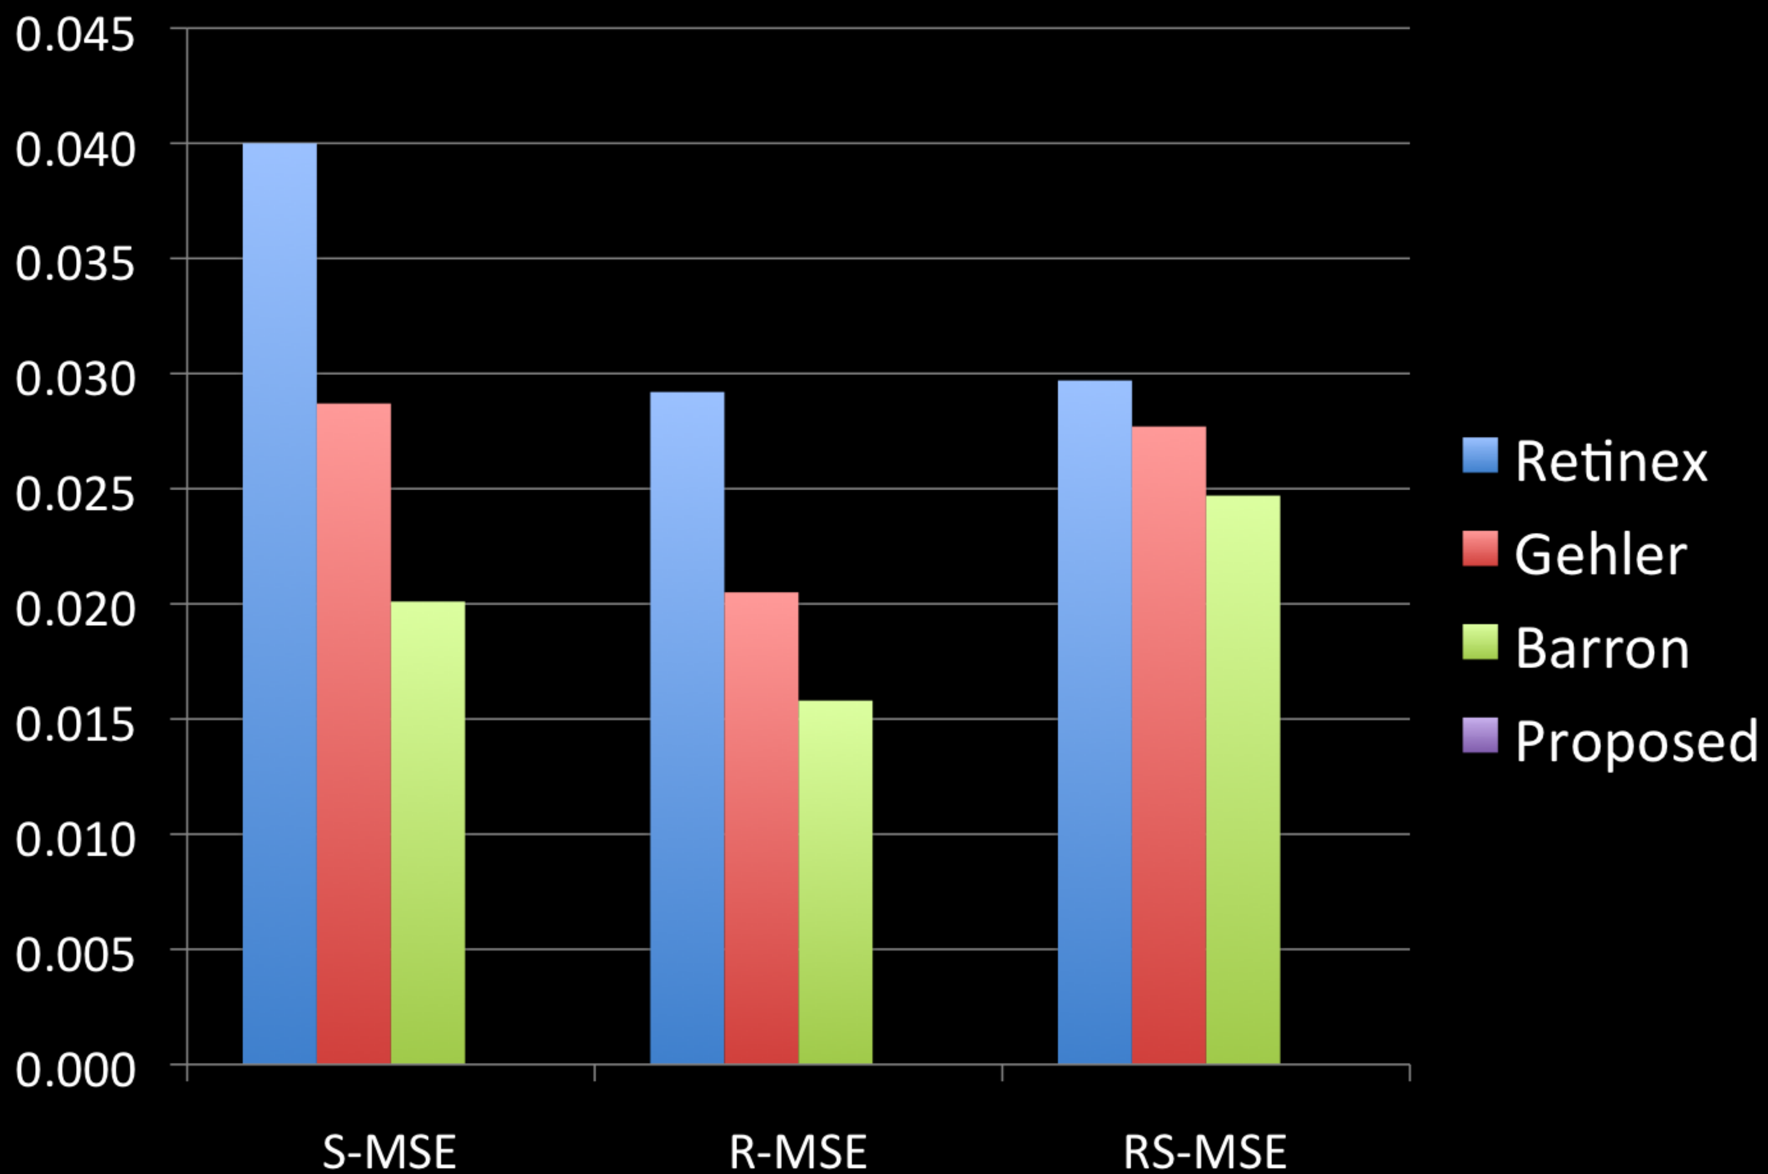

[Grosse et al. '09], [Land & McCann '71], [Barron & Malik '13], [Gehler et al. '11]

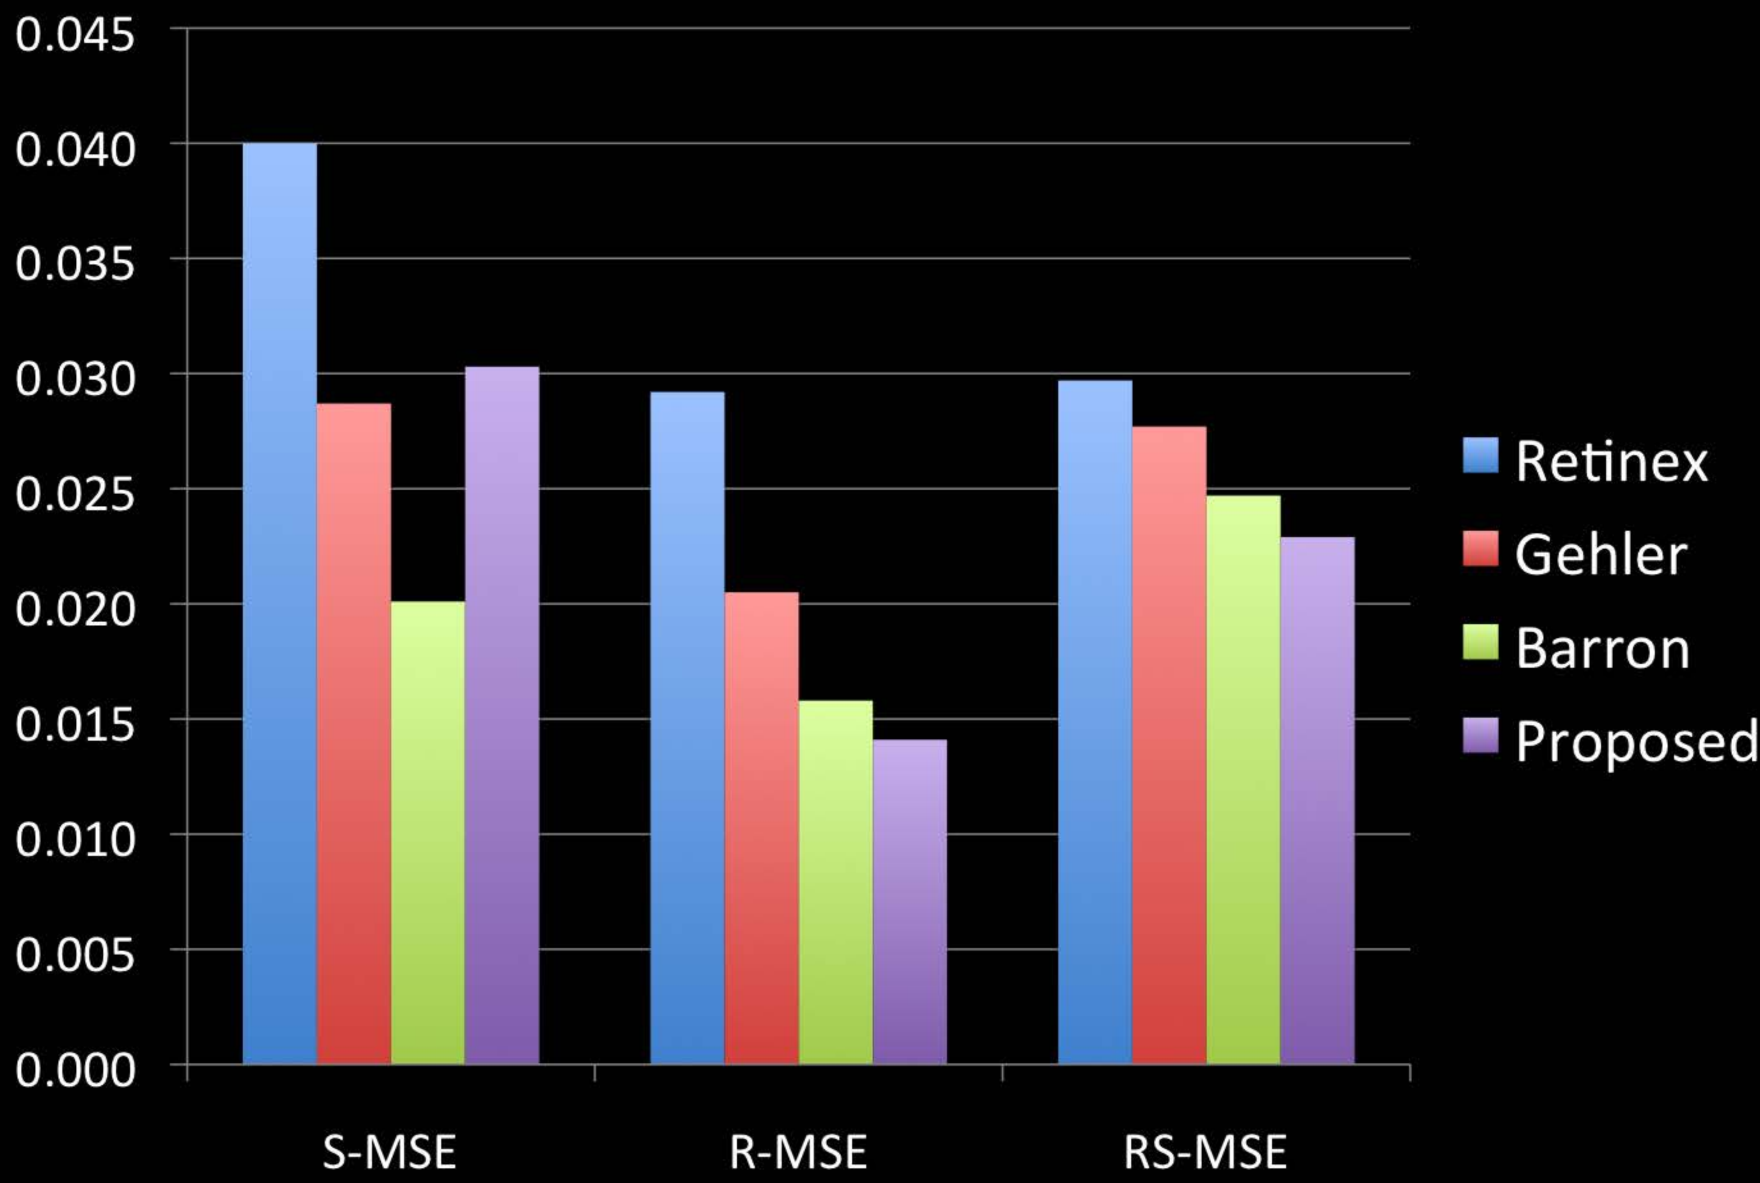

[Grosse et al. ‘09], [Land & McCann ‘71], [Barron & Malik ‘13], [Gehler et al. ’11]

Summary

– Bayesian nonparametric extension of [Gehler et al. ‘11]
– No 3D modeling or Retinex gradient terms
– On par with state-of-the-art performance

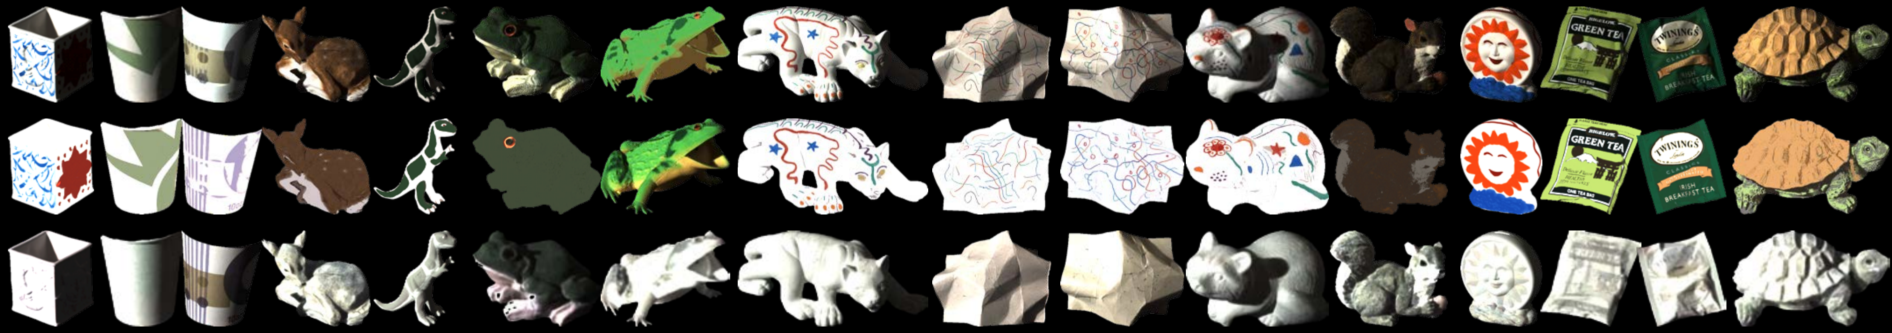

Code will be available at: http://people.csail.mit.edu/jchang7/

CSAIL