Visual Material Traits
Recognizing Per-Pixel Material Context

Gabriel Schwartz
gbs25@drexel.edu

Ko Nishino
kon@drexel.edu




Materials as Visual Context

What tells us this road is unsafe?

Material Category Recognition Methods

- Give single predictions for the entire image

Images from [2].

Adelson [1]

Liu et al. [4]

Hu et al. [3]

Sharan et al. [6]

Material Category Recognition Methods

- Give single predictions for the entire image
- Require object information
  - object mask
  - bounding box

Images from [2].

Adelson [1]

Liu et al. [4]

Hu et al. [3]

Sharan et al. [6]

Material Category Recognition Methods

- Give single predictions for the entire image
- Require object information
  - object mask
  - bounding box
- Predict categories that are really object properties

Images from [2].

Adelson [1]

Liu et al. [4]

Hu et al. [3]

Sharan et al. [6]

Material Category Recognition Methods

- Give single predictions for the entire image
- Require object information
  - object mask
  - bounding box
- Predict categories that are really object properties
- Object information not always available

Adelson [1]

Liu et al. [4]

Hu et al. [3]

Sharan et al. [6]

Images from [2].

Intra-Class Appearance Variability

Images from [7].

Visual Material Traits: Characteristic Material Properties

Image from [5].

Visual Material Traits: Characteristic Material Properties

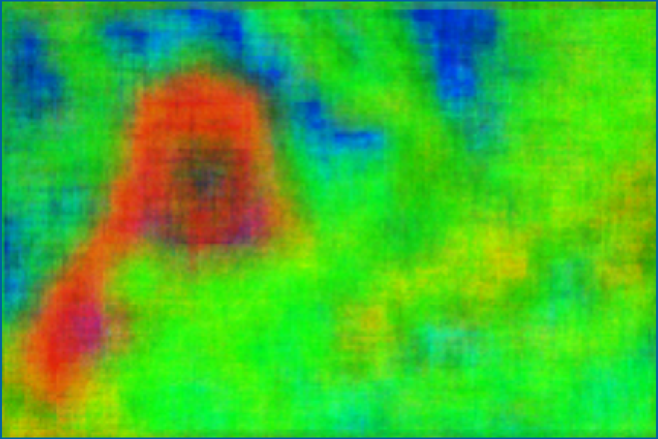

Fuzzy Organic Smooth

Material traits are locally-recognizable material properties.

Image from [5].

Visual Material Trait Appearances

- What material properties can we see locally?

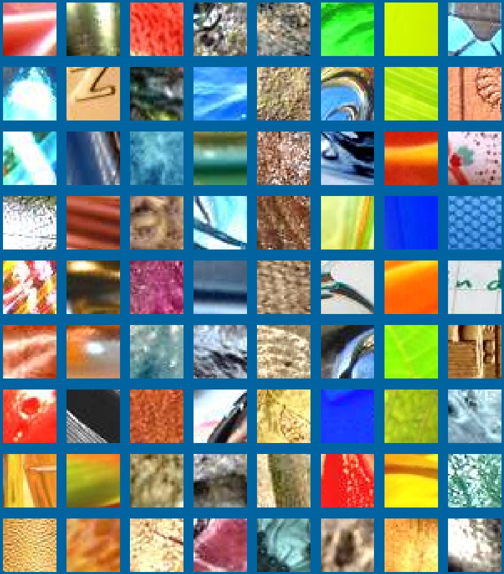

Images from [7].

Visual Material Trait Appearances

- What material properties can we see locally?
- Certain properties are easy to describe

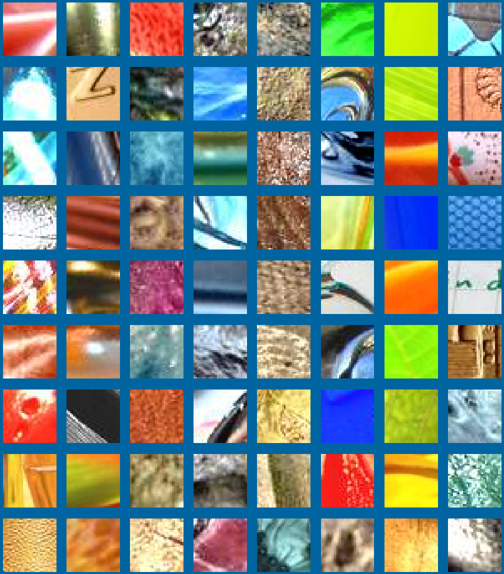

Images from [7].

Visual Material Trait Appearances

- What material properties can we see locally?
- Certain properties are easy to describe
  - Shiny

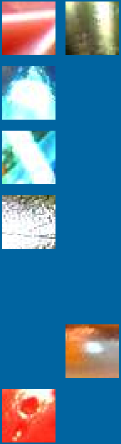

Images from [7].

Visual Material Trait Appearances

- What material properties can we see locally?
- Certain properties are easy to describe
  - Shiny
  - Smooth

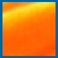

Visual Material Trait Appearances

- What material properties can we see locally?
- Certain properties are easy to describe
  - Shiny
  - Smooth
- Some are more challenging

Visual Material Trait Appearances

- What material properties can we see locally?
- Certain properties are easy to describe
  - Shiny
  - Smooth
- Some are more challenging
  - Fuzzy? Soft?

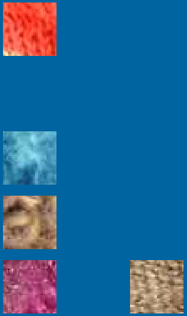

Visual Material Trait Appearances

- What material properties can we see locally?
- Certain properties are easy to describe
  - Shiny
  - Smooth
- Some are more challenging
  - Fuzzy? Soft?

How do we represent these traits?

Learning to Represent Material Traits

- Learn features that model the appearance of material traits

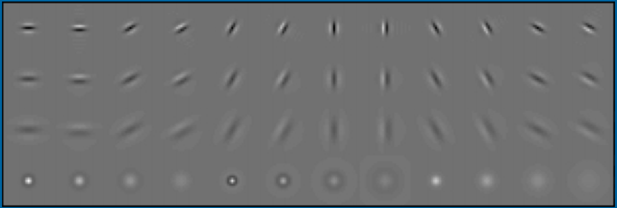

Learning to Represent Material Traits

- Learn features that model the appearance of material traits
- Features should be:

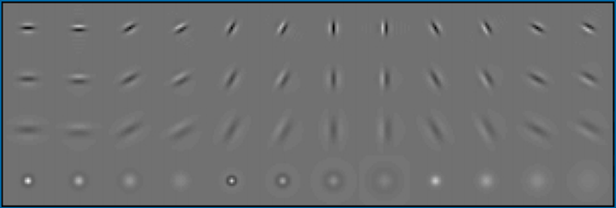

Learning to Represent Material Traits

- Learn features that model the appearance of material traits
- Features should be:
  - Fast to compute

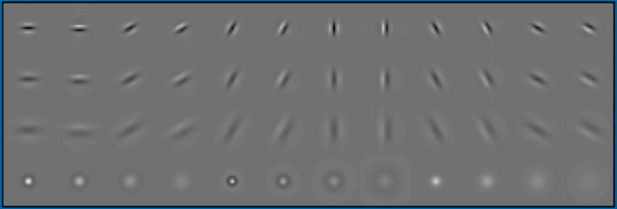

Learning to Represent Material Traits

- Learn features that model the appearance of material traits
- Features should be:
  - Fast to compute
  - Able to be extracted anywhere

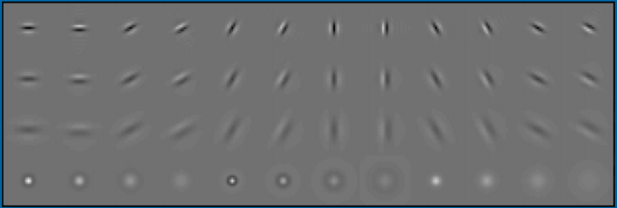

Learning to Represent Material Traits

- Learn features that model the appearance of material traits
- Features should be:
    - Fast to compute
    - Able to be extracted anywhere
    - Discriminative

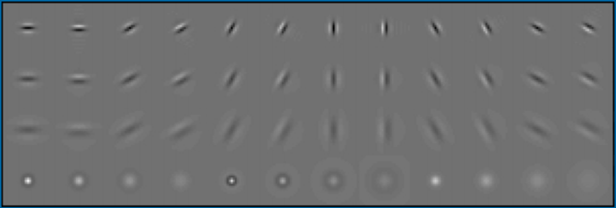

Learning to Represent Material Traits

- Learn features that model the appearance of material traits
- Features should be:
  - Fast to compute
  - Able to be extracted anywhere
  - Discriminative
- Convolution filters may satisfy all of these properties

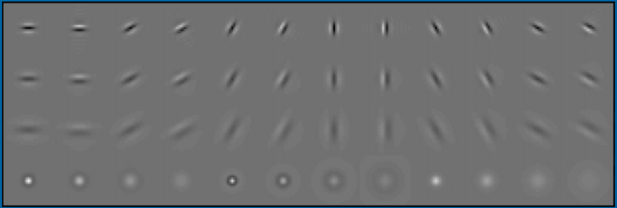

Learning to Represent Material Traits

- Learn features that model the appearance of material traits
- Features should be:
  - Fast to compute
  - Able to be extracted anywhere
  - Discriminative
- Convolution filters may satisfy all of these properties
- How do we learn them?

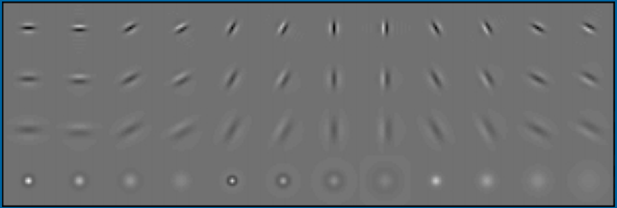

Learning Filters for Trait Representation

- Convolutional Autoencoder (CAE) model for feature learning
- Find optimal filters ($\mathbf{W}$) s.t. they:

$$\min_{\mathbf{W},\mathbf{W}'}$$

Learning Filters for Trait Representation

- Convolutional Autoencoder (CAE) model for feature learning
- Find optimal filters ($\mathbf{W}$) s.t. they:
  - Model trait patches

$$\mathbf{E}_i = h\left(\mathbf{W} * \mathbf{I}_i + b_e\right)$$

$$\mathbf{R}_i = \mathbf{W}' * \mathbf{E}_i + b_r$$

$$h(x) =$$

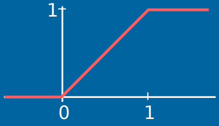

$$\min_{\mathbf{W},\mathbf{W}'} \frac{1}{N} \sum_{i=1}^{N} \|\mathbf{I}_i - \mathbf{R}_i\|_{\mathrm{F}}^2$$

Learning Filters for Trait Representation

- Convolutional Autoencoder (CAE) model for feature learning
- Find optimal filters ($\mathbf{W}$) s.t. they:
  - Model trait patches
  - Form a sparse encoding

$$
\begin{aligned}
\mathbf{E}_i &= h\left(\mathbf{W} * \mathbf{I}_i + b_e\right) \\
\mathbf{R}_i &= \mathbf{W}' * \mathbf{E}_i + b_r
\end{aligned}
$$

$h(x) =$

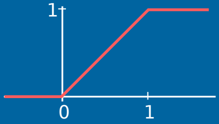

$$
\min_{\mathbf{W},\mathbf{W}'} \frac{1}{N} \sum_{i=1}^{N} \left\| \mathbf{I}_i - \mathbf{R}_i \right\|_{\mathrm{F}}^2 + \alpha \left\| p - \frac{1}{N} \sum_{i=1}^{N} \mathbf{E}_i \right\|_{\mathrm{F}}^2
$$

Learning Filters for Trait Representation

- Convolutional Autoencoder (CAE) model for feature learning
- Find optimal filters ($\mathbf{W}$) s.t. they:
  - Model trait patches
  - Form a sparse encoding
  - Have constrained magnitude

$$\begin{aligned} \mathbf{E}_i &= h\left(\mathbf{W} * \mathbf{I}_i + b_e\right) \\ \mathbf{R}_i &= \mathbf{W}' * \mathbf{E}_i + b_r \end{aligned}$$

$$h\left(x\right) =$$

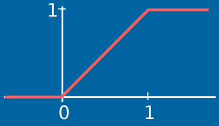

$$\min_{\mathbf{W},\mathbf{W}'} \frac{1}{N}\sum_{i=1}^{N} \|\mathbf{I}_i - \mathbf{R}_i\|_{\mathrm{F}}^2 + \alpha \left\| p - \frac{1}{N}\sum_{i=1}^{N} \mathbf{E}_i \right\|_{\mathrm{F}}^2 + \beta \left(\|\mathbf{W}\|_{\mathrm{F}}^2 + \|\mathbf{W}'\|_{\mathrm{F}}^2 \right)$$

Learning Filters for Trait Representation

- Convolutional Autoencoder (CAE) model for feature learning
- Find optimal filters ($\mathbf{W}$) s.t. they:
  - Model trait patches
  - Form a sparse encoding
  - Have constrained magnitude

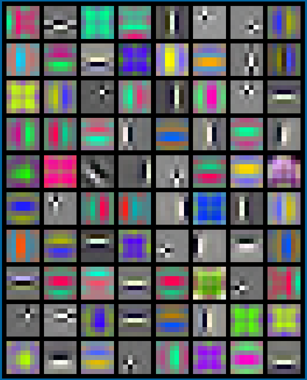

$$
\begin{aligned}
\mathbf{E}_i &= h\left(\mathbf{W} * \mathbf{I}_i + b_e\right) \\
\mathbf{R}_i &= \mathbf{W}' * \mathbf{E}_i + b_r
\end{aligned}
$$

$$h(x) =$$

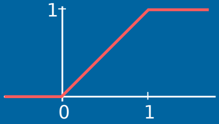

$$\min_{\mathbf{W},\mathbf{W}'} \; \frac{1}{N}\sum_{i=1}^{N}\|\mathbf{I}_i - \mathbf{R}_i\|_{\mathrm{F}}^2 + \alpha \left\| p - \frac{1}{N}\sum_{i=1}^{N}\mathbf{E}_i \right\|_{\mathrm{F}}^2 + \beta\left(\|\mathbf{W}\|_{\mathrm{F}}^2 + \|\mathbf{W}'\|_{\mathrm{F}}^2\right)$$

Learned Filters

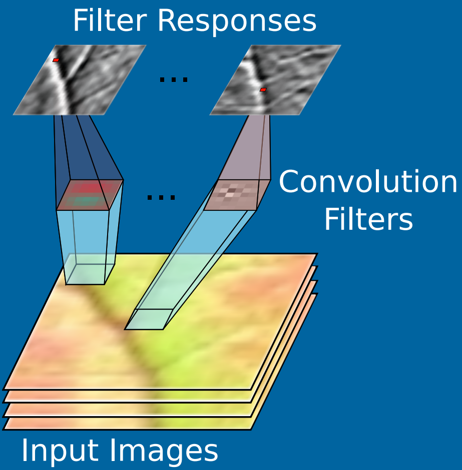

Learned Filters + Supplemental Nonlinear Features

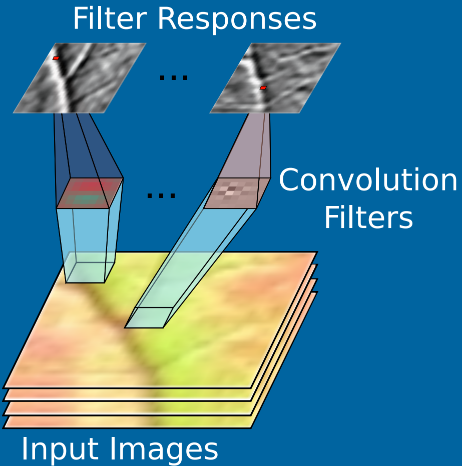

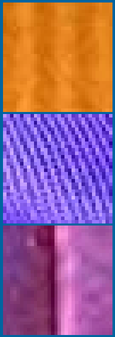

- Describe appearances CAE cannot:

Learned Filters + Supplemental Nonlinear Features

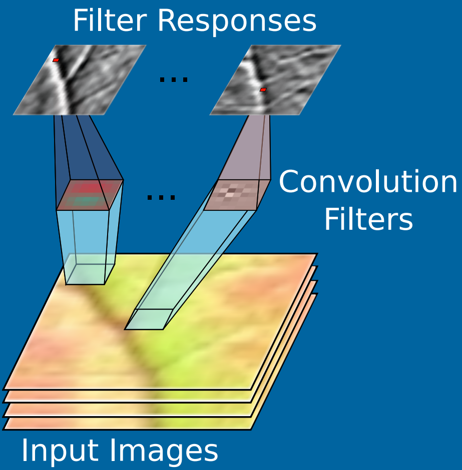

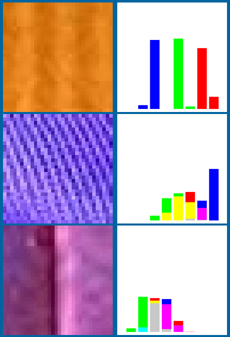

- Describe appearances CAE cannot:
  - Color Histograms

Learned Filters + Supplemental Nonlinear Features

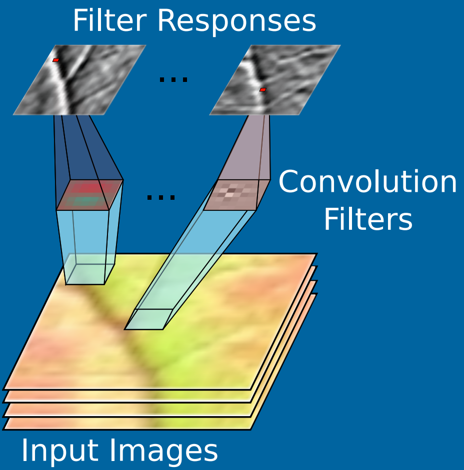

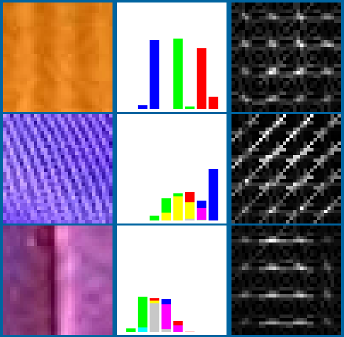

- Describe appearances CAE cannot:
  - Color Histograms
  - HOG

Learned Filters + Supplemental Nonlinear Features

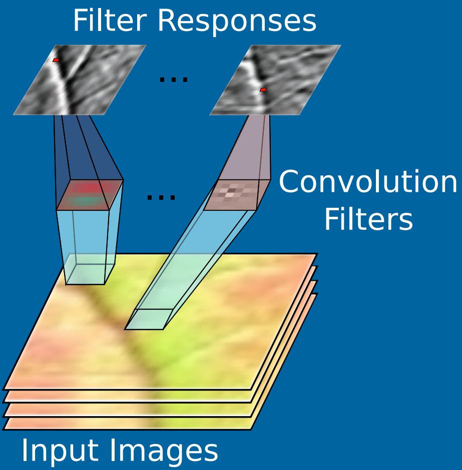

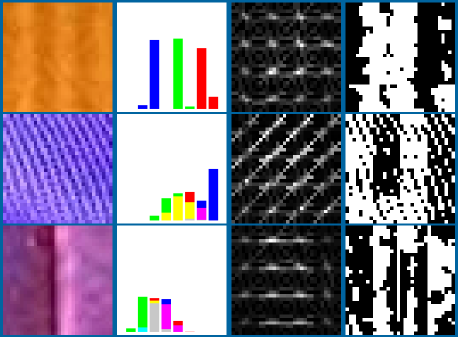

- Describe appearances CAE cannot:
  - Color Histograms
  - HOG
  - LBP

Material Trait Recognition Process

- Training data: Flickr Materials Database (FMD) [7] images with trait annotations

Learn Filters

Material Trait Recognition Process

- Training data: Flickr Materials Database (FMD) [7] images with trait annotations

Learn Filters

Select Features

Trait	CAE	Oriented	HOG	LBP	Color Histograms
Shiny	•				•
Fuzzy		•		•	
Transparent	•	•	•		
$\cdots$ (13 Material Traits)					
Total Uses	**7**	4	6	9	7

Material Trait Recognition Process

- Training data: Flickr Materials Database (FMD) [7] images with trait annotations

Learn Filters → Select Features → Train Per-Trait Classifiers

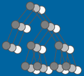

Trait	CAE	Oriented	HOG	LBP	Color Histograms
Shiny	•				•
Fuzzy		•		•	
Transparent	•	•	•		
$\cdots$ (13 Material Traits)					
Total Uses	**7**	4	6	9	7

Material Trait Recognition Process

- Training data: Flickr Materials Database (FMD) [7] images with trait annotations

Learn Filters

Select Features

Train Per-Trait Classifiers

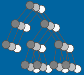

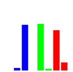

Extract Features

Trait	CAE	Oriented	HOG	LBP	Color Histograms
Shiny	•				•
Fuzzy		•		•	
Transparent	•	•	•		
· · · (13 Material Traits)					
Total Uses	**7**	4	6	9	7

Material Trait Recognition Process

- Training data: Flickr Materials Database (FMD) [7] images with trait annotations

Learn Filters

Select Features

Train Per-Trait Classifiers

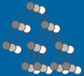

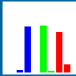

Extract Features

Recognize Traits

Trait	CAE	Oriented	HOG	LBP	Color Histograms
Shiny	•				•
Fuzzy		•		•	
Transparent	•	•	•		
· · · (13 Material Traits)					
Total Uses	**7**	4	6	9	7

Material Trait Recognition Accuracy

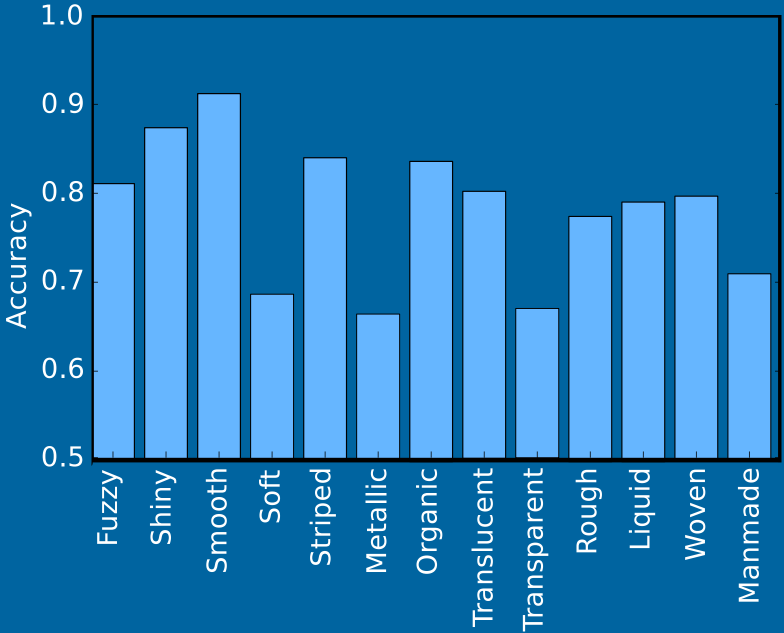

Patch Recognition Results

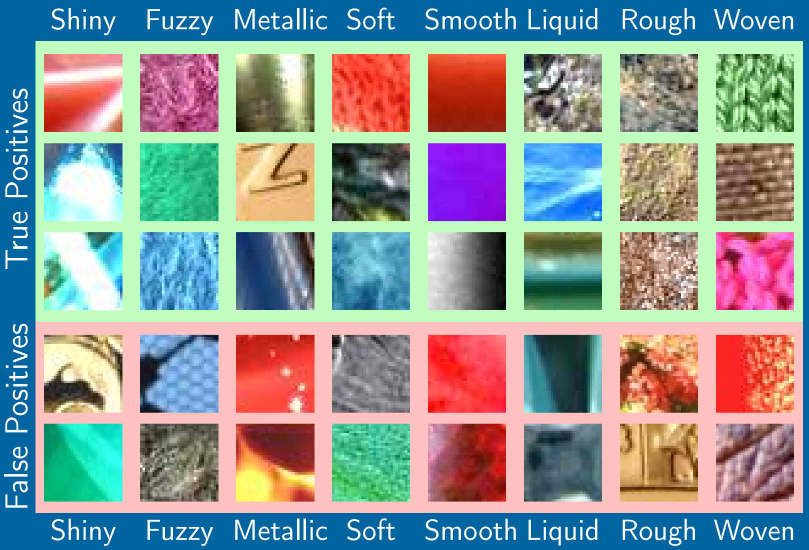

Per-Pixel Material Trait Maps

Image from [5].

Per-Pixel Material Trait Maps

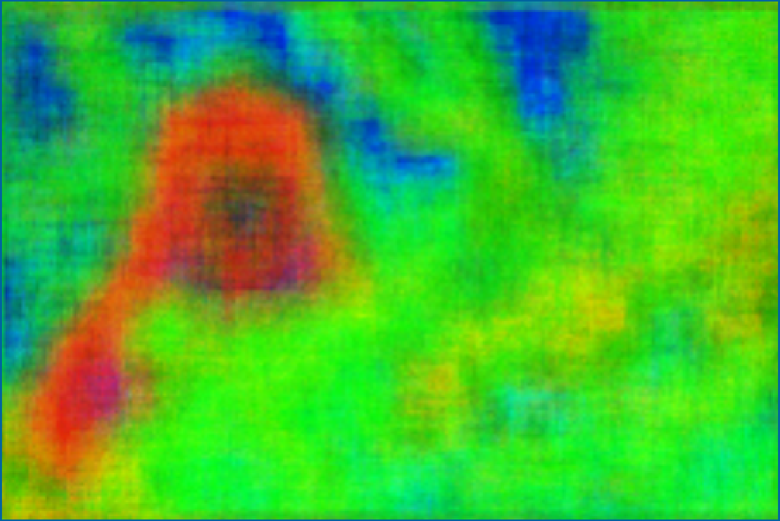

Fuzzy Organic Smooth

Image from [5].

Per-Pixel Material Trait Maps

Image from [7].

Per-Pixel Material Trait Maps

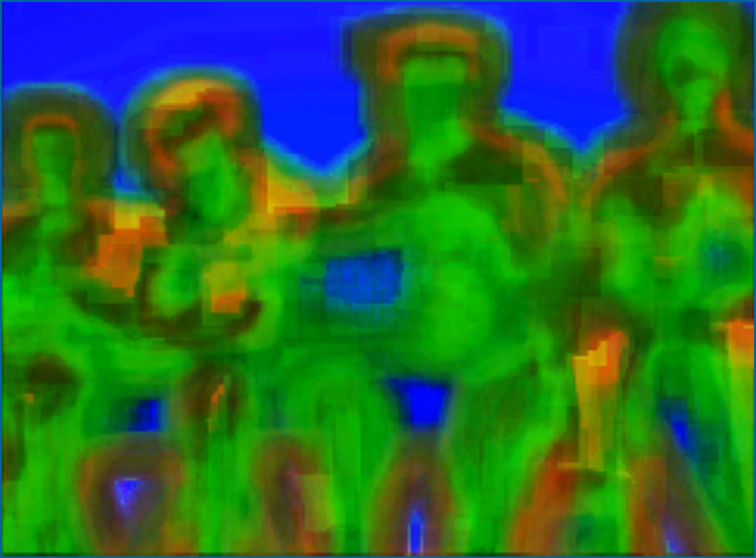

Shiny Metallic Smooth

Image from [7].

What can we do with these material traits?

Material Recognition via Trait Distributions

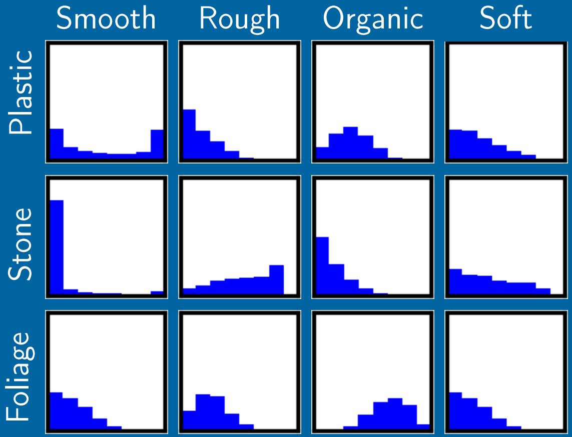

$$p\left(\text{trait}_i \middle| \text{category}_j\right)$$

Material Recognition Accuracy: Flickr

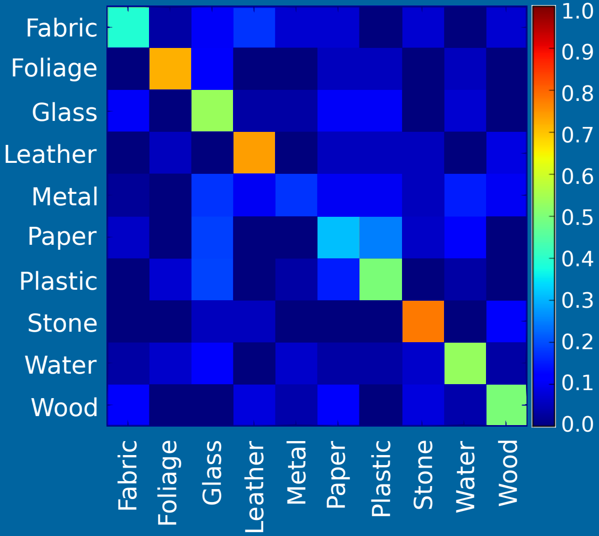

Average: 49.2%

[Liu et al.] (w/o obj): 42.6% (w/obj): 57.1%

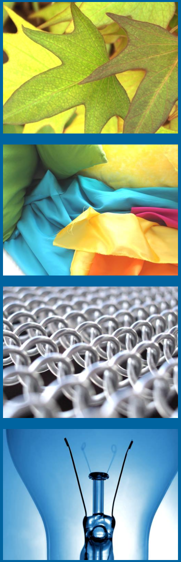

Images from [7].

Material Recognition Accuracy: ImageNet

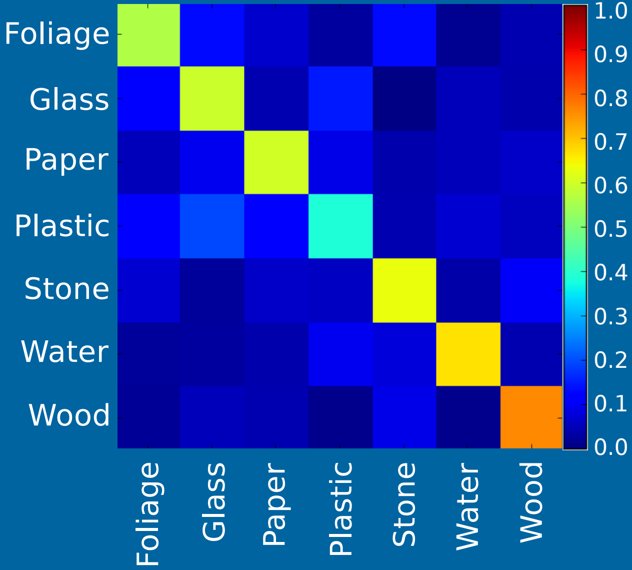

Average: 60.5%

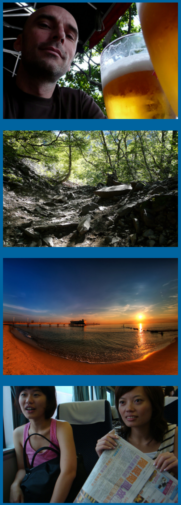

Images from [2].

Segmentation with Material Traits

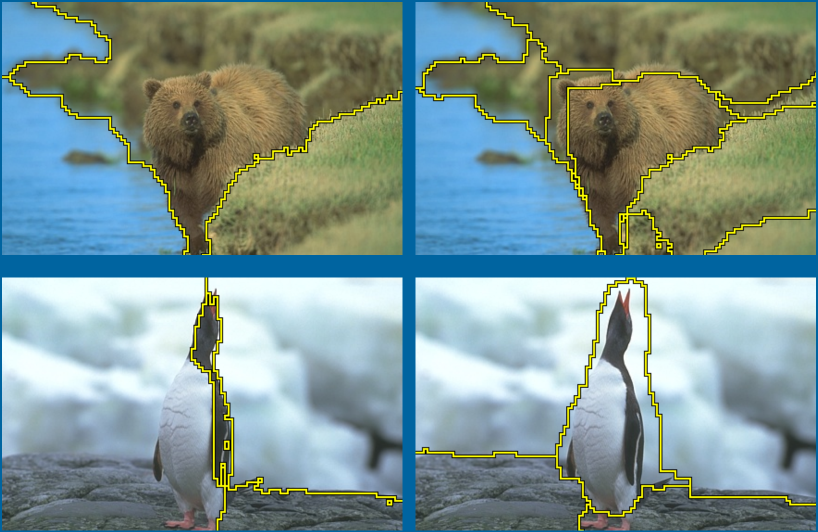

Images from [5].

Baseline NCuts | With Traits

Summary

- Material traits:
  - may be recognized locally and accurately

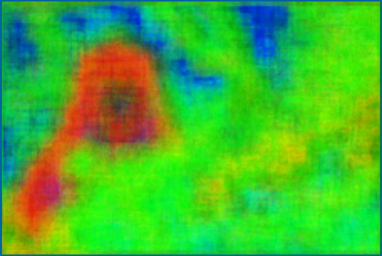

Summary

- Material traits:
  - may be recognized locally and accurately
  - have distributions that encode material categories

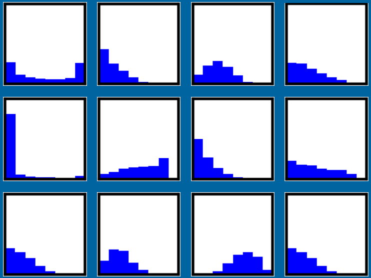

Summary

- Material traits:
  - may be recognized locally and accurately
  - have distributions that encode material categories
  - segment images into intuitively separate regions

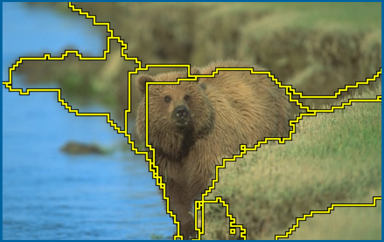

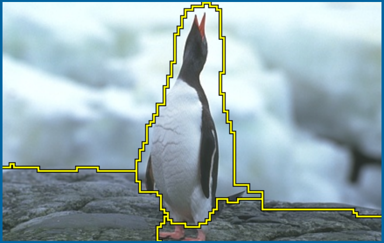

Summary

- Material traits:
  - may be recognized locally and accurately
  - have distributions that encode material categories
  - segment images into intuitively separate regions
- Future work:
  - Discover new traits
  - Improve applications

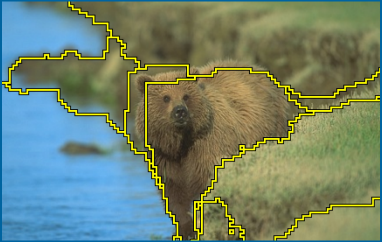

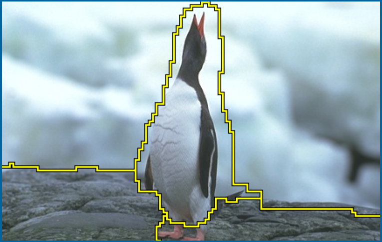

References

E. H. Adelson.
On Seeing Stuff: The Perception of Materials by Humans and Machines.
In *SPIE*, pages 1–12, 2001.

J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei.
ImageNet: A Large-Scale Hierarchical Image Database.
In *CVPR*, 2009.

D. Hu, L. Bo, and X. Ren.
Toward Robust Material Recognition for Everyday Objects.
In *BMVC*, pages 48.1–48.11, 2011.

C. Liu, L. Sharan, E. H. Adelson, and R. Rosenholtz.
Exploring Features in a Bayesian Framework for Material Recognition.
In *CVPR*, pages 239–246, 2010.

D. Martin, C. Fowlkes, D. Tal, and J. Malik.
A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics.
In *ICCV*, pages 416–423.

L. Sharan, C. Liu, R. Rosenholtz, and E. H. Adelson.
Recognizing Materials Using Perceptually Inspired Features.
International Journal of Computer Vision, 2013.

L. Sharan, R. Rosenholtz, and E. Adelson.
Material Perception: What Can You See in a Brief Glance?
Journal of Vision, 9(8):784, 2009.