Scene Analysis From Range Data

Unsupervised Discovery of Categories

Felix Endres

Master's Thesis

Albert-Ludwigs-Universität Freiburg
Department of Computer Science
Autonomous Intelligent Systems Laboratory

Outline

- Classification
- Features
- Learning
- Experiments
- Summary

Classification

- Goal: Categorize objects in 3D range data
- Requirement: No information about classes

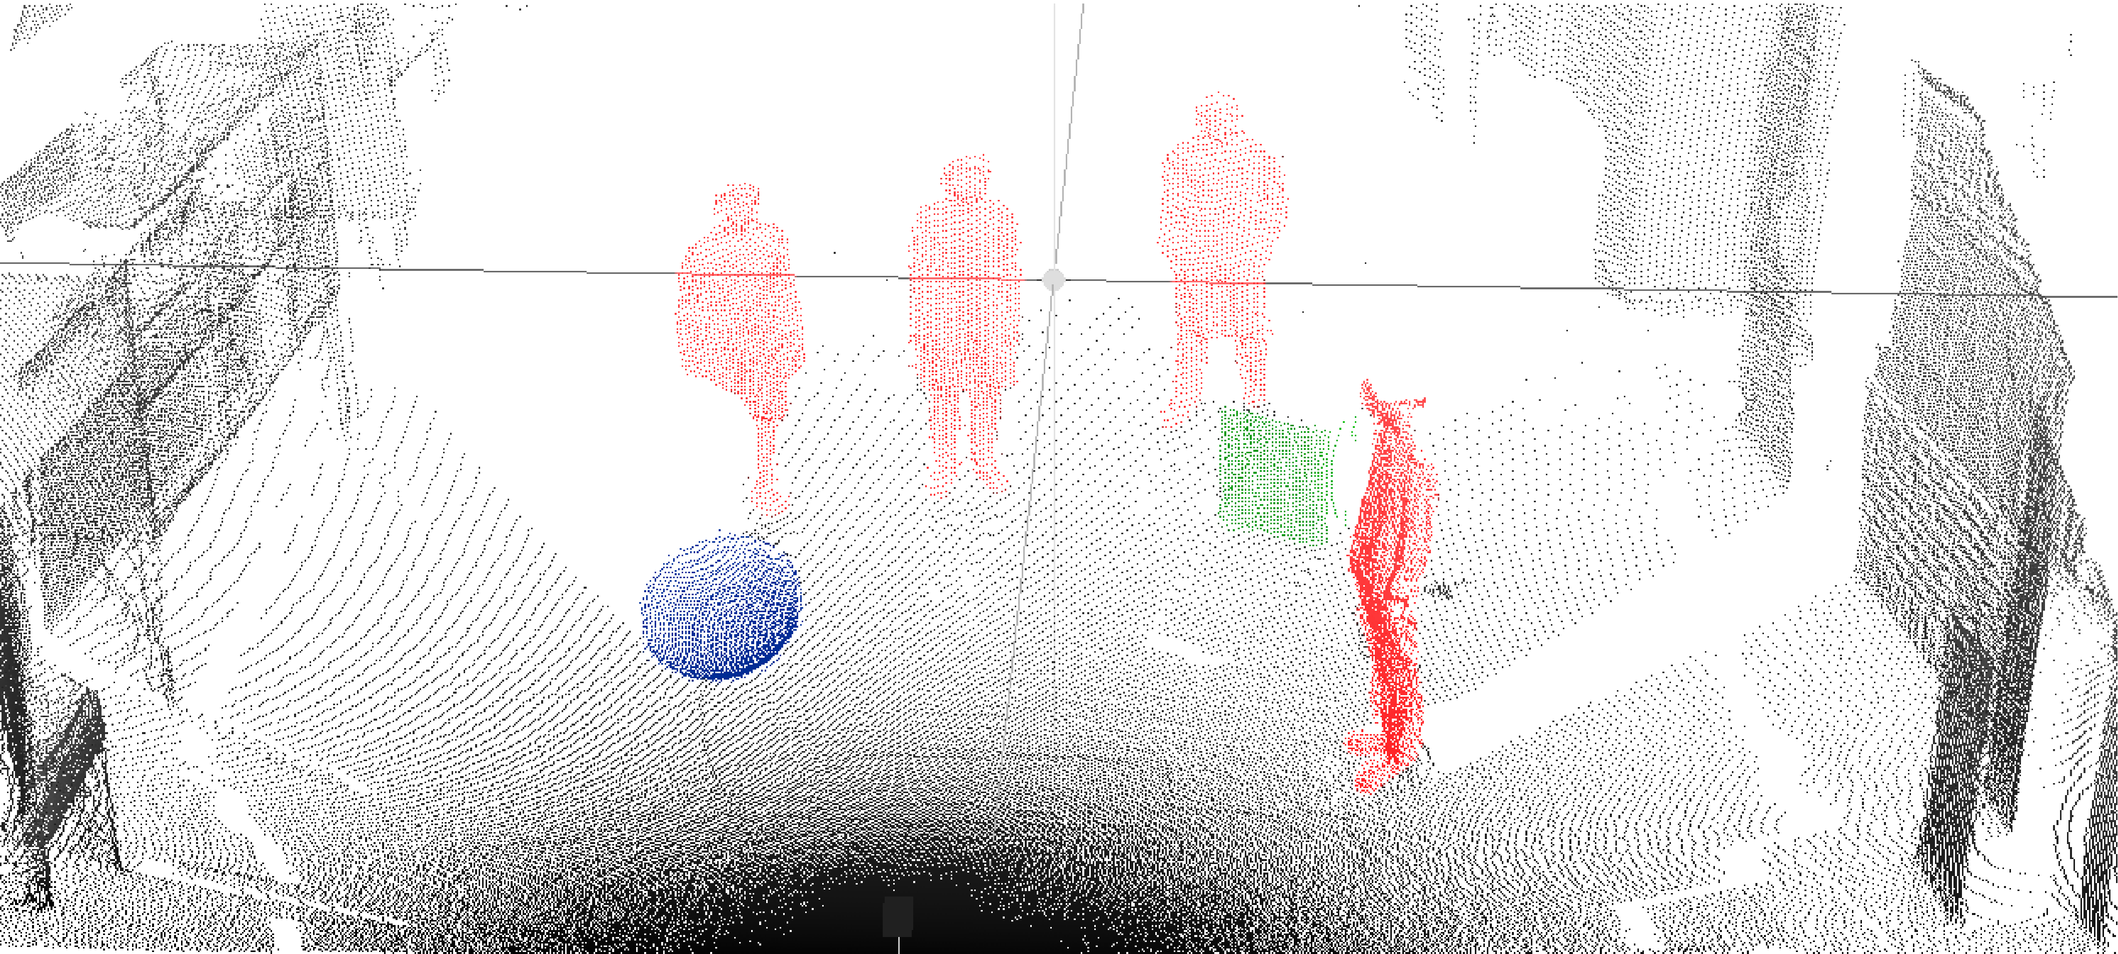

Classification

- Interpret object classes as feature distributions
- Infer the object's class

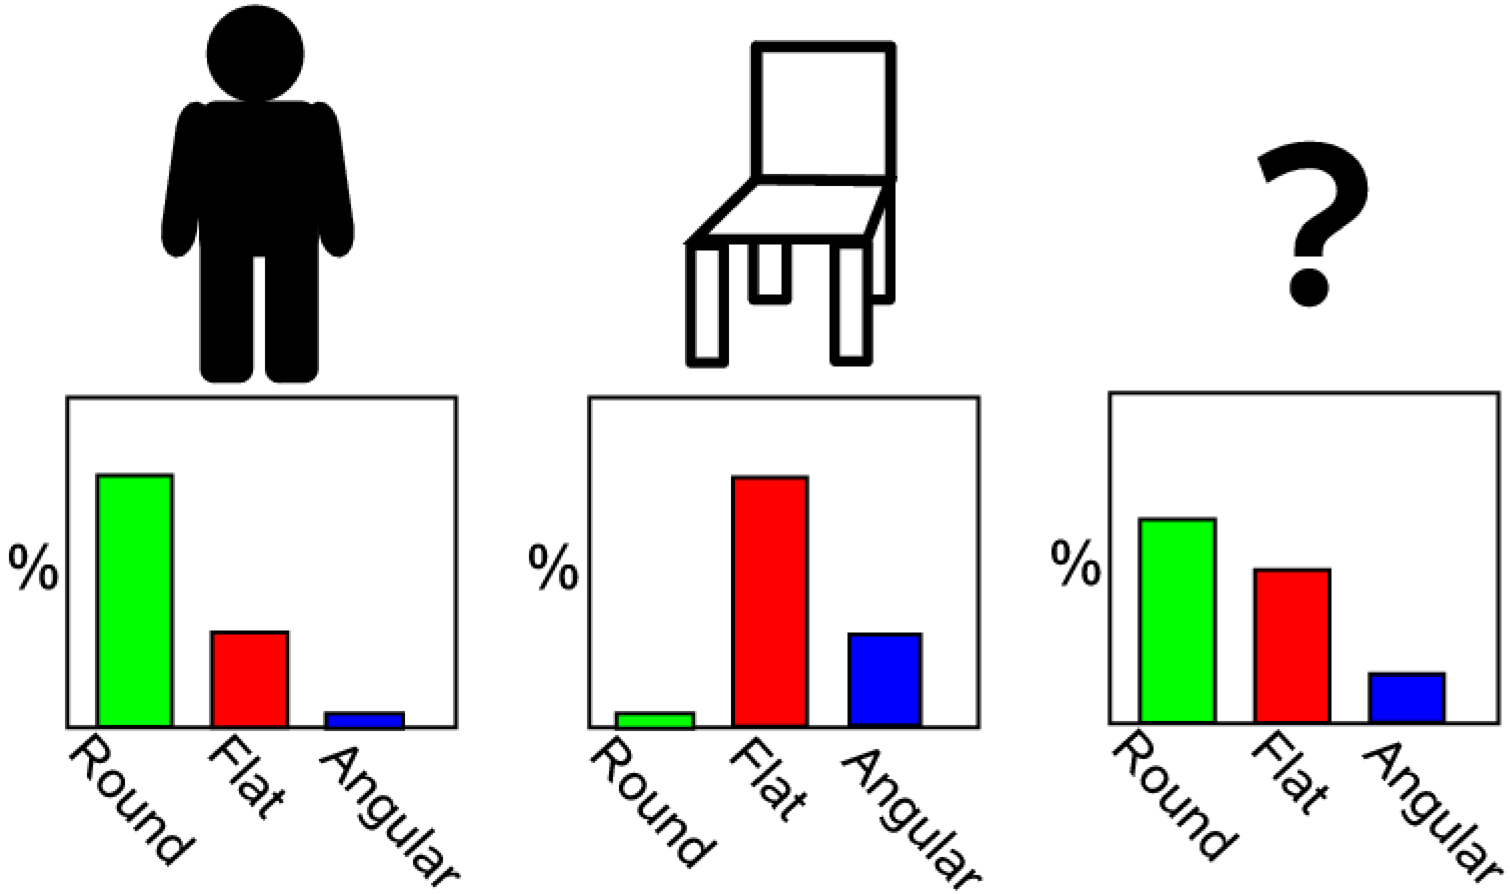

Classification

Given a number of data segments, can we find the category distributions without further information?

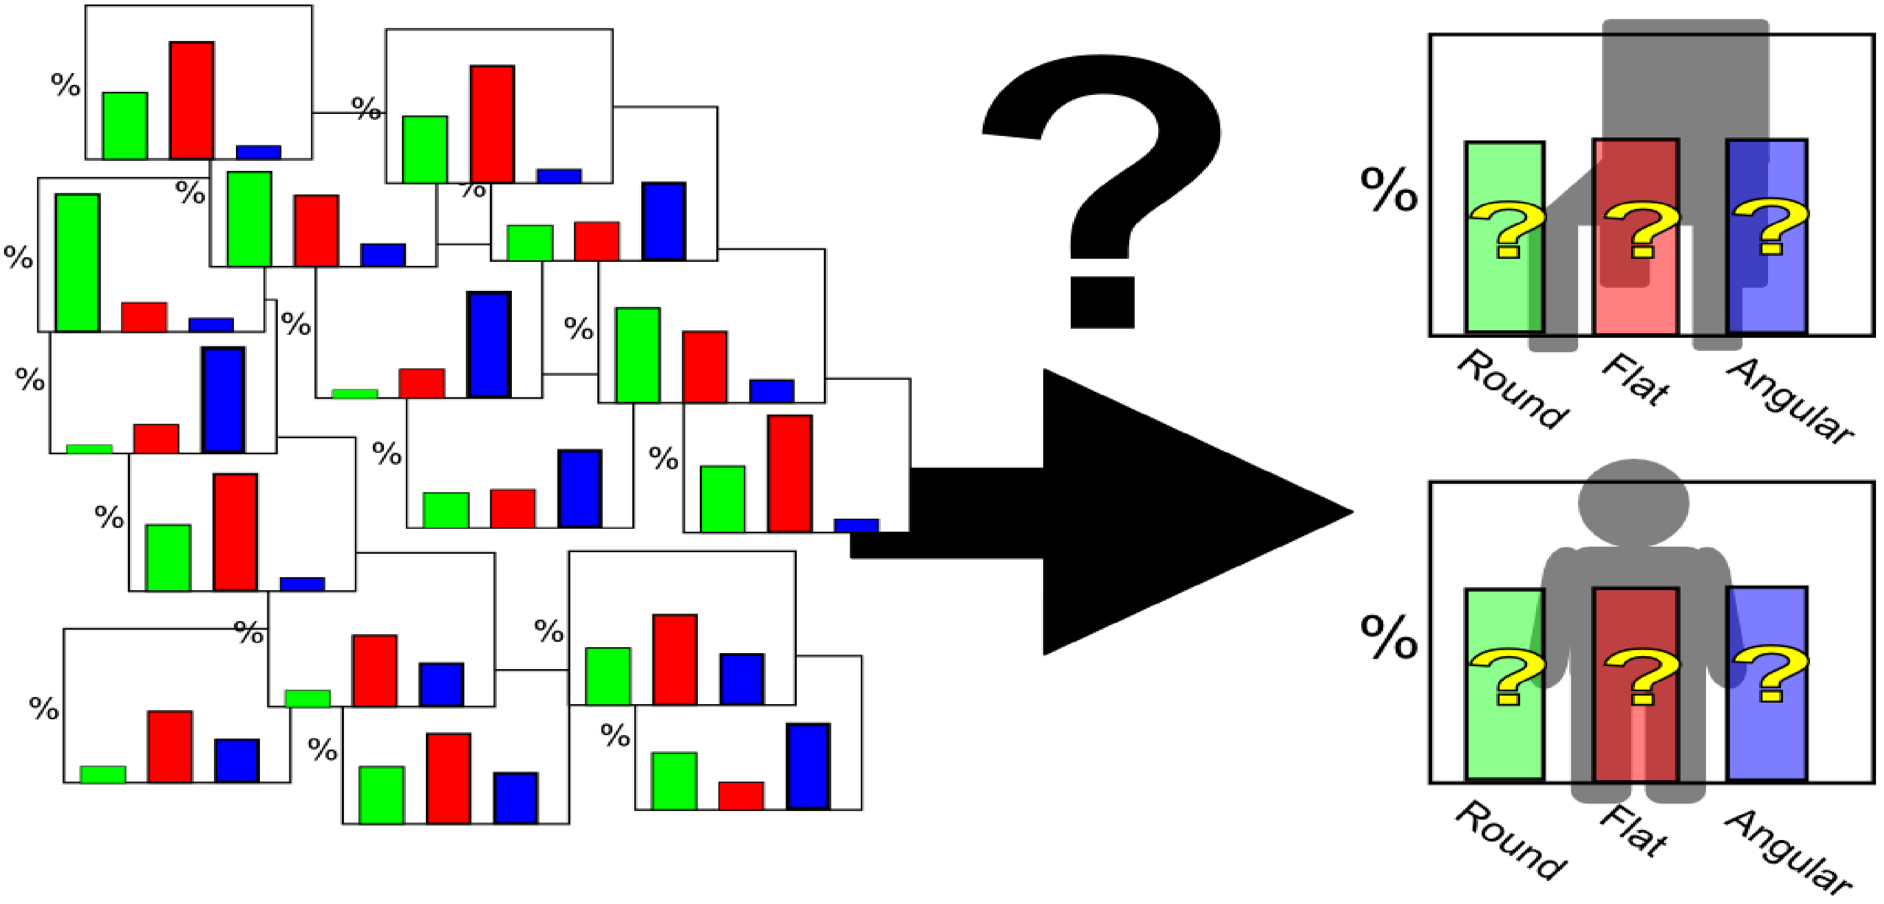

Outline

- Classification
- **Features**
- Learning
- Experiments
- Summary

Features - Spin Images

- Introduced by Andrew Johnson
- Description of the local environment of a point
- Coordinate system aligned to surface normal

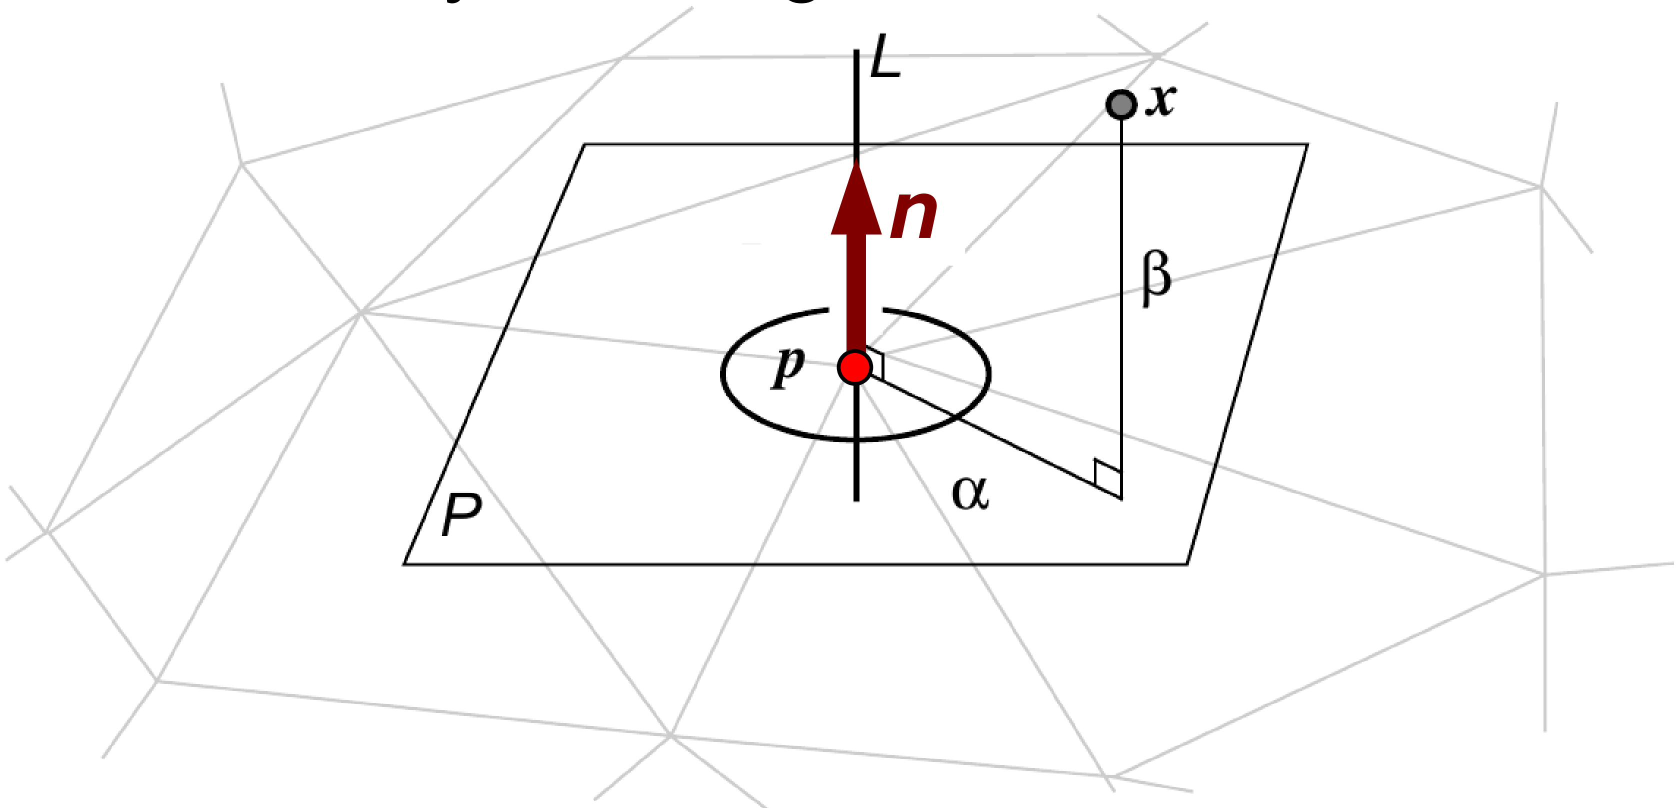

Features - Spin Images

- Robust to
  - Shape variation
  - Occlusion
  - Clutter

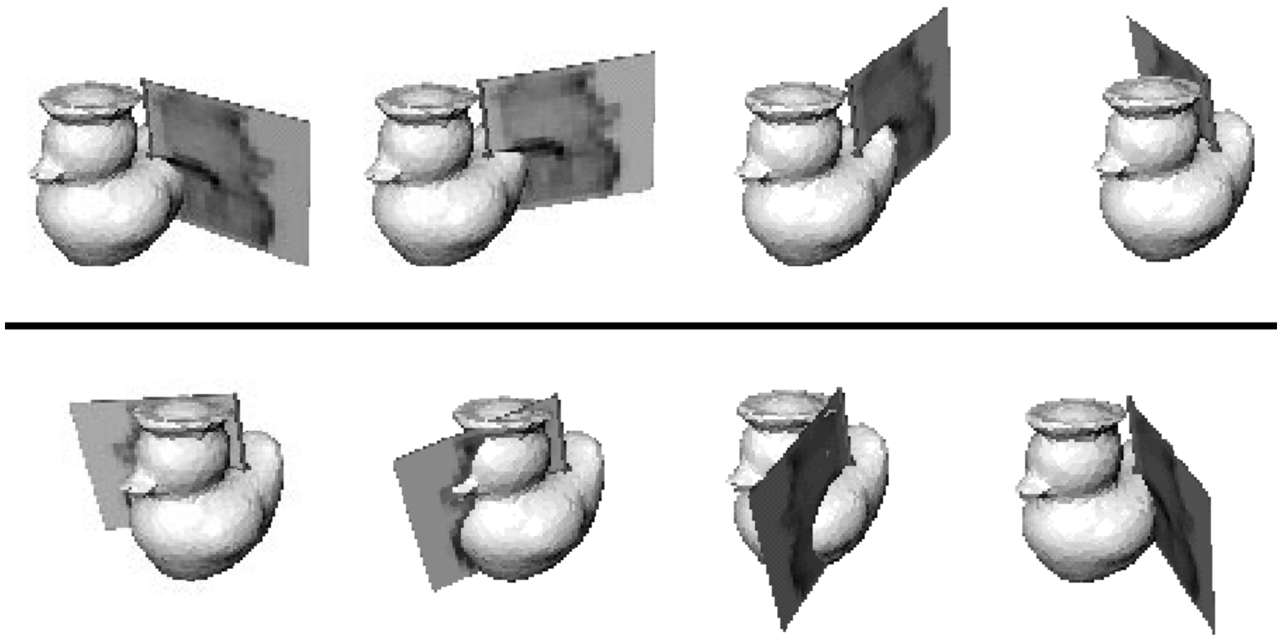

Features - Spin Images

- Robust to
  - Shape variation
  - Occlusion
  - Clutter

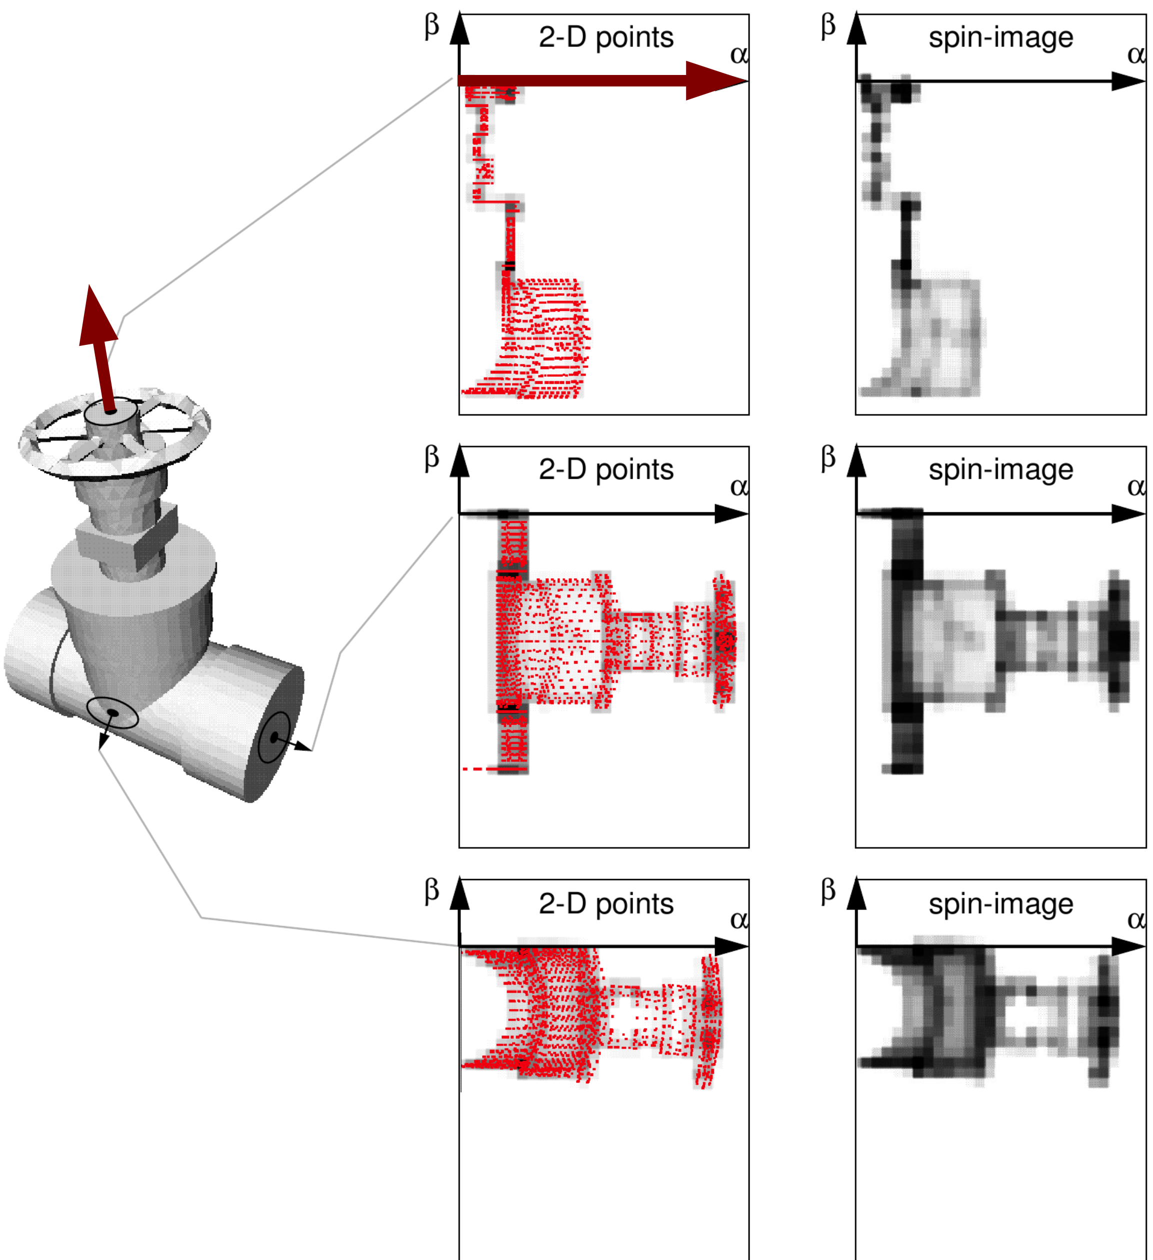

Features - Spin Images

- Parameters
  - **Support Distance** (top)
  - **Raster Resolution** (middle)
  - **Discretization Resolution** (bottom)

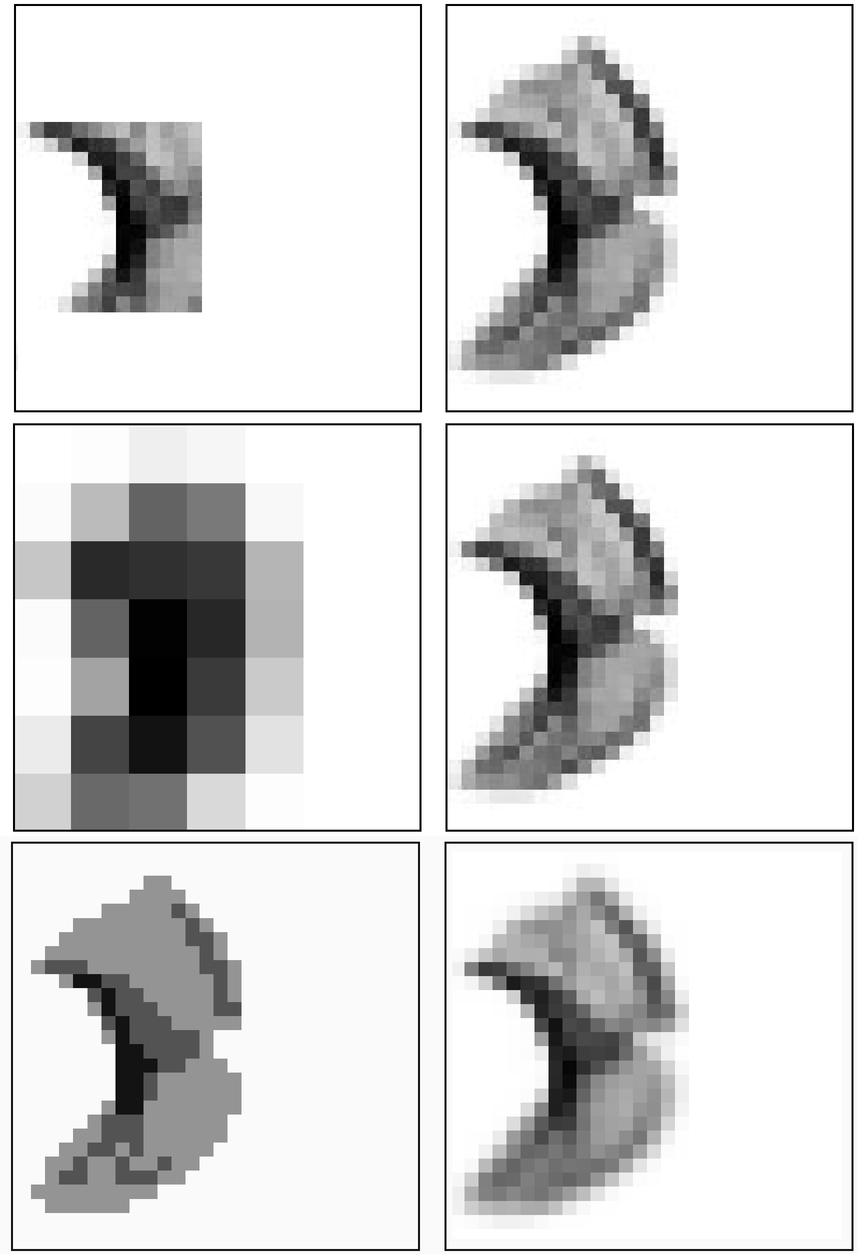

Features – Enhanced Spin Images

- We now consider the surface normals of all points
- Instead of point counts, we us the **average angle**
- Same parameters as for standard spin images

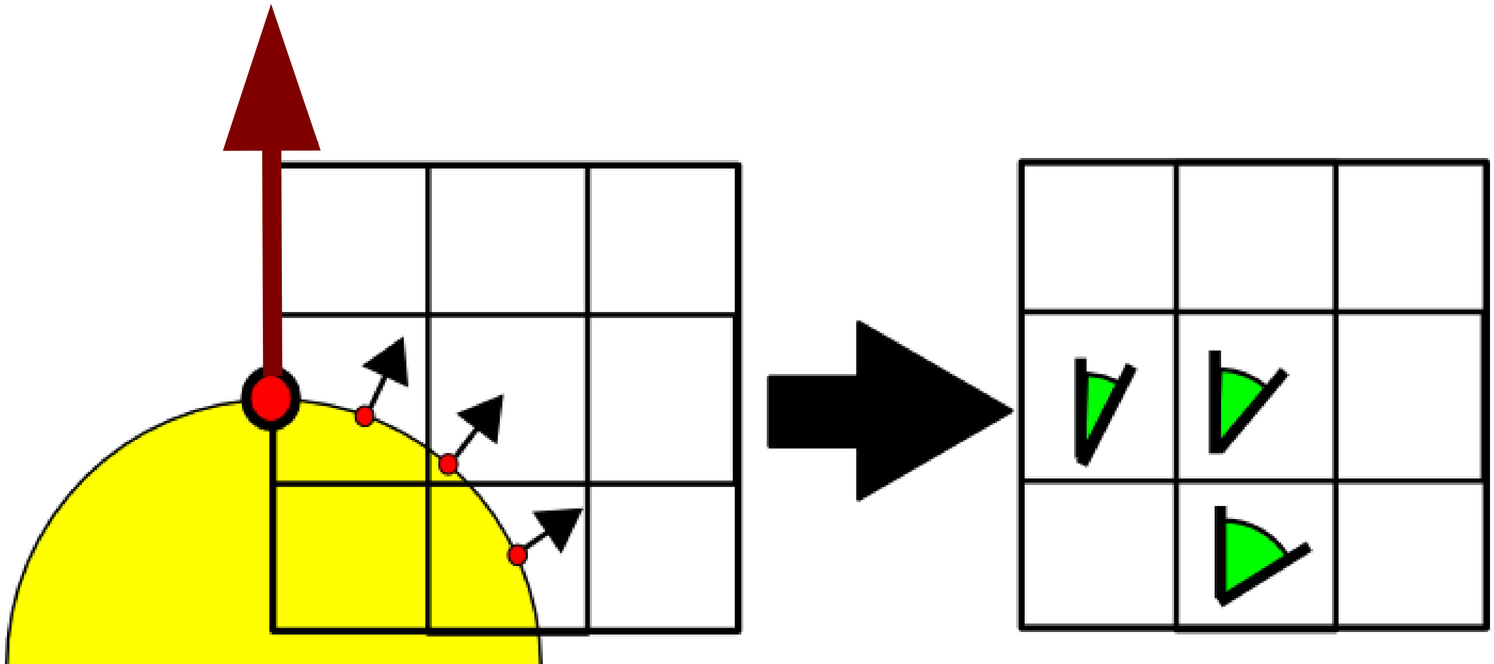

Outline

- Classification
- Features
- **<u>Learning</u>**
- Experiments
- Summary

Learning

- Probabilistic framework: latent Dirichlet allocation
- Introduced by David M. Blei in 2003
- Unsupervised: No class information necessary
- No explicit distance measurement necessary
- Category discovery from feature co-occurrence

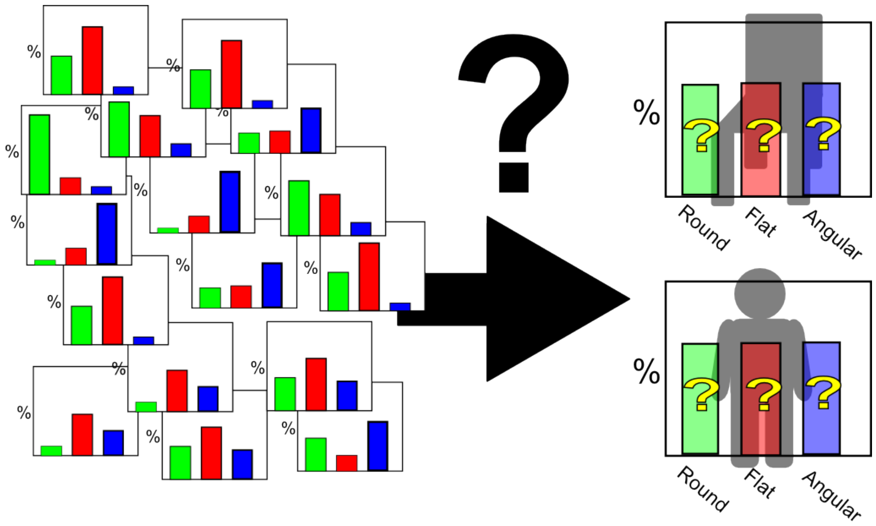

Latent Dirichlet Allocation

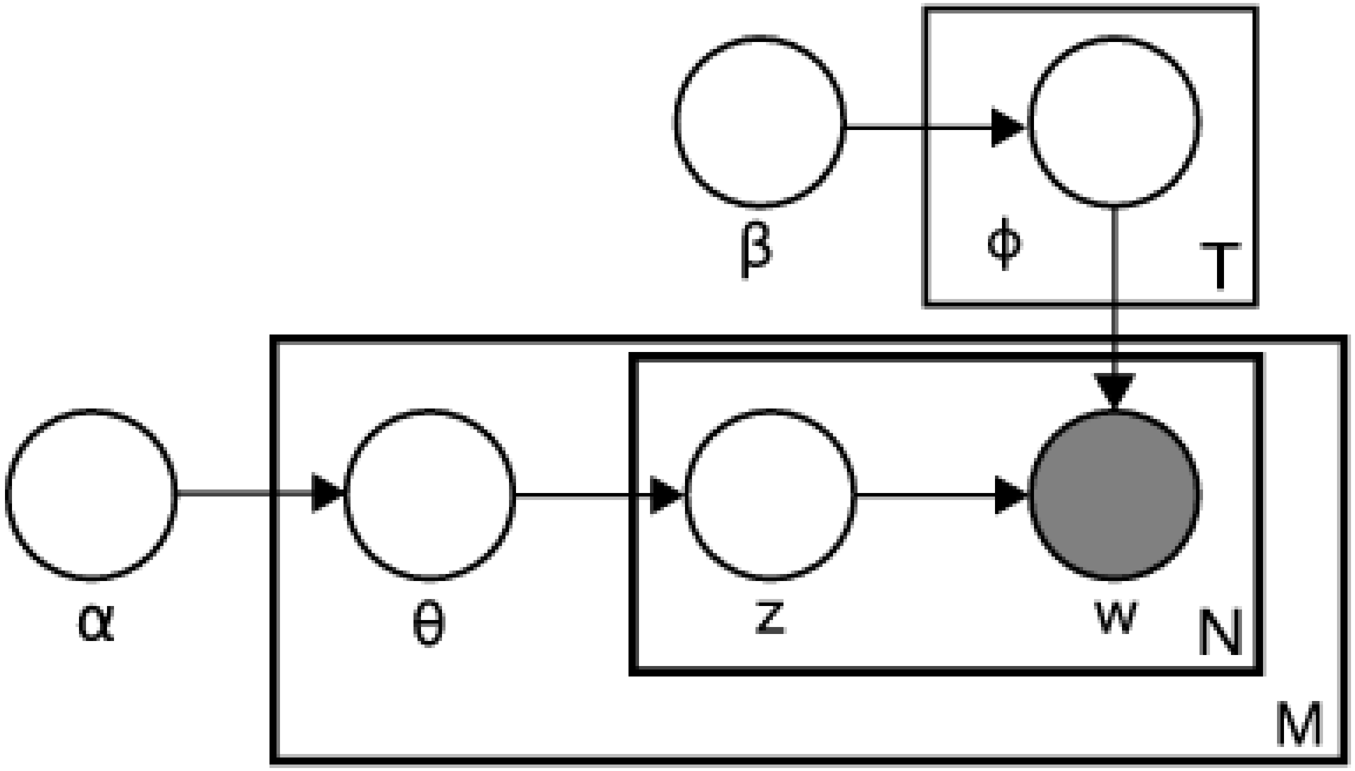

Plate Notation:

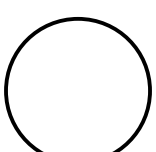
Unobserved Variable

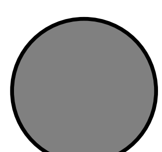
Observed Variable

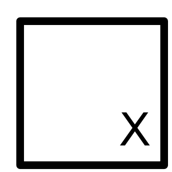

Sample Contained Variables X Times

Latent Dirichlet Allocation

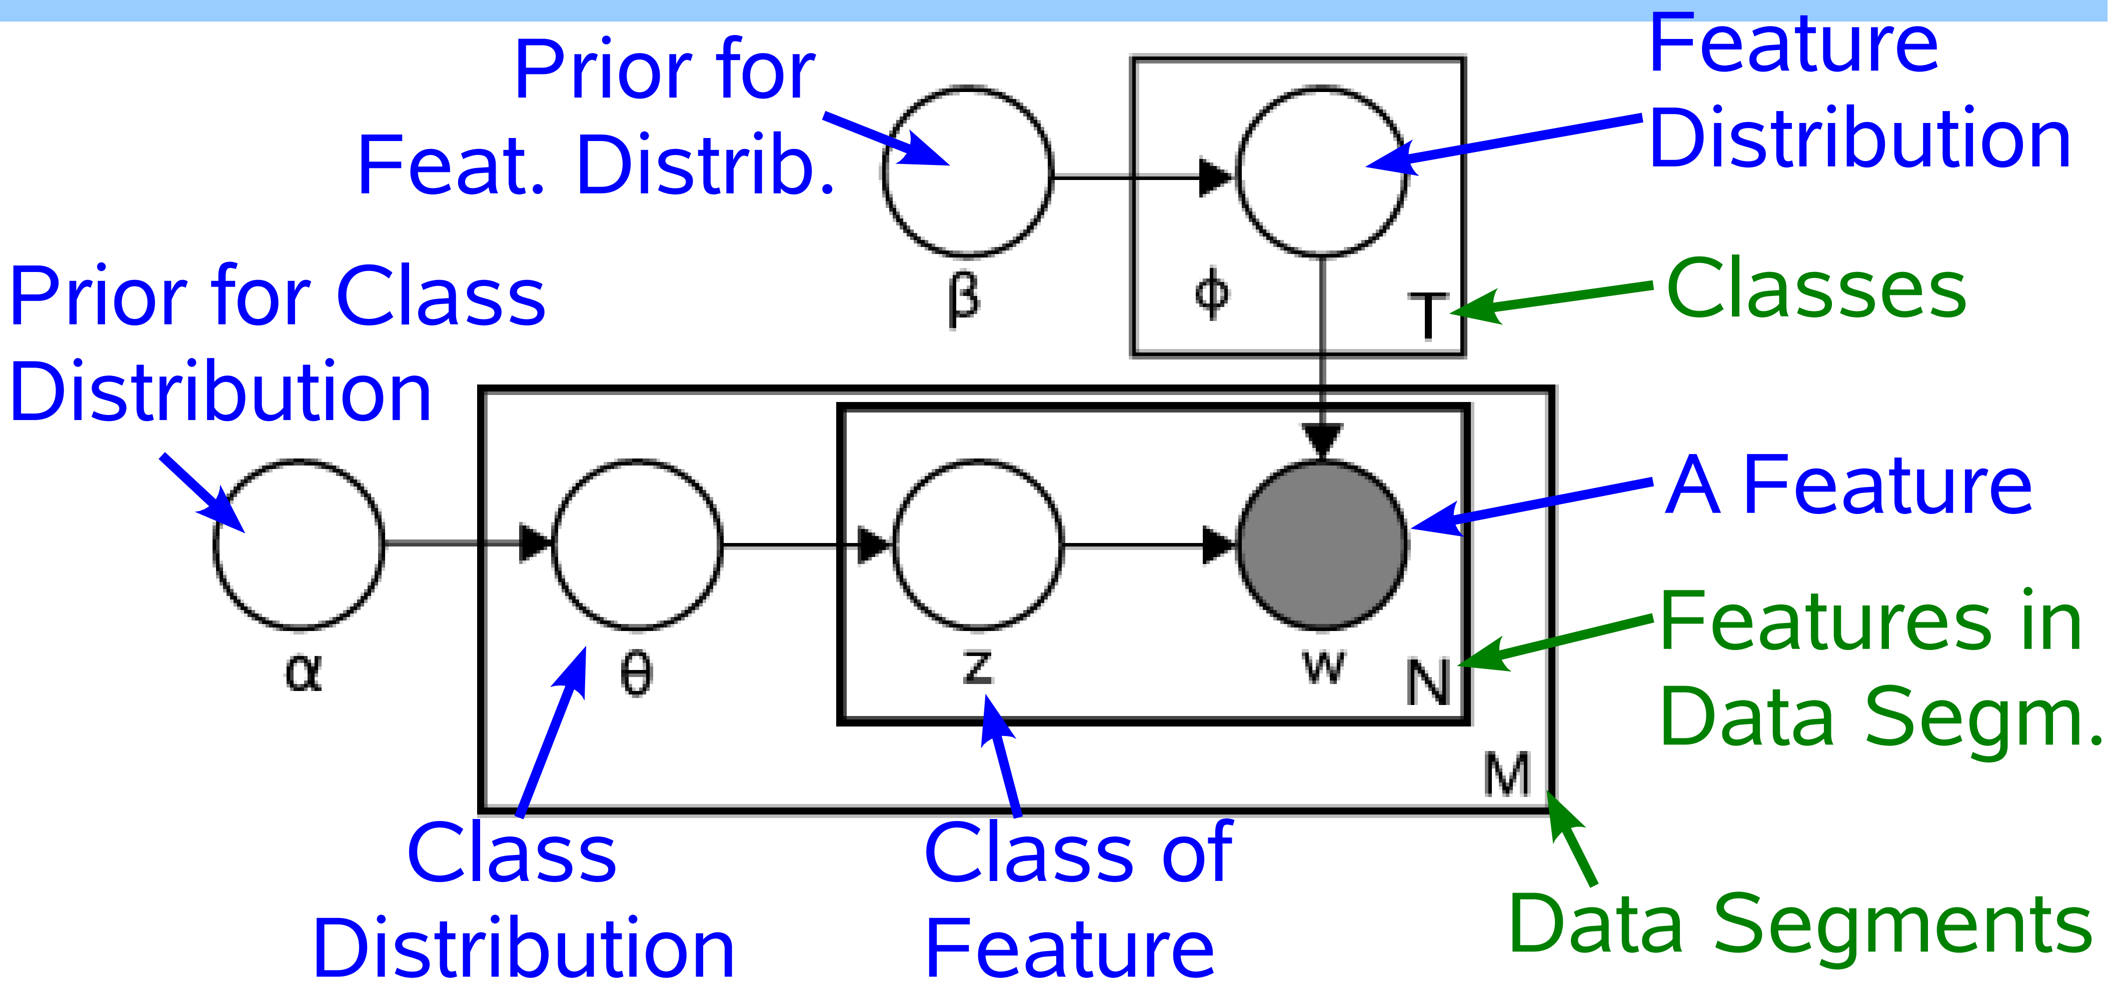

Plate Notation:

Unobserved Variable

Observed Variable

X

Sample Contained Variables X Times

Latent Dirichlet Allocation - Example

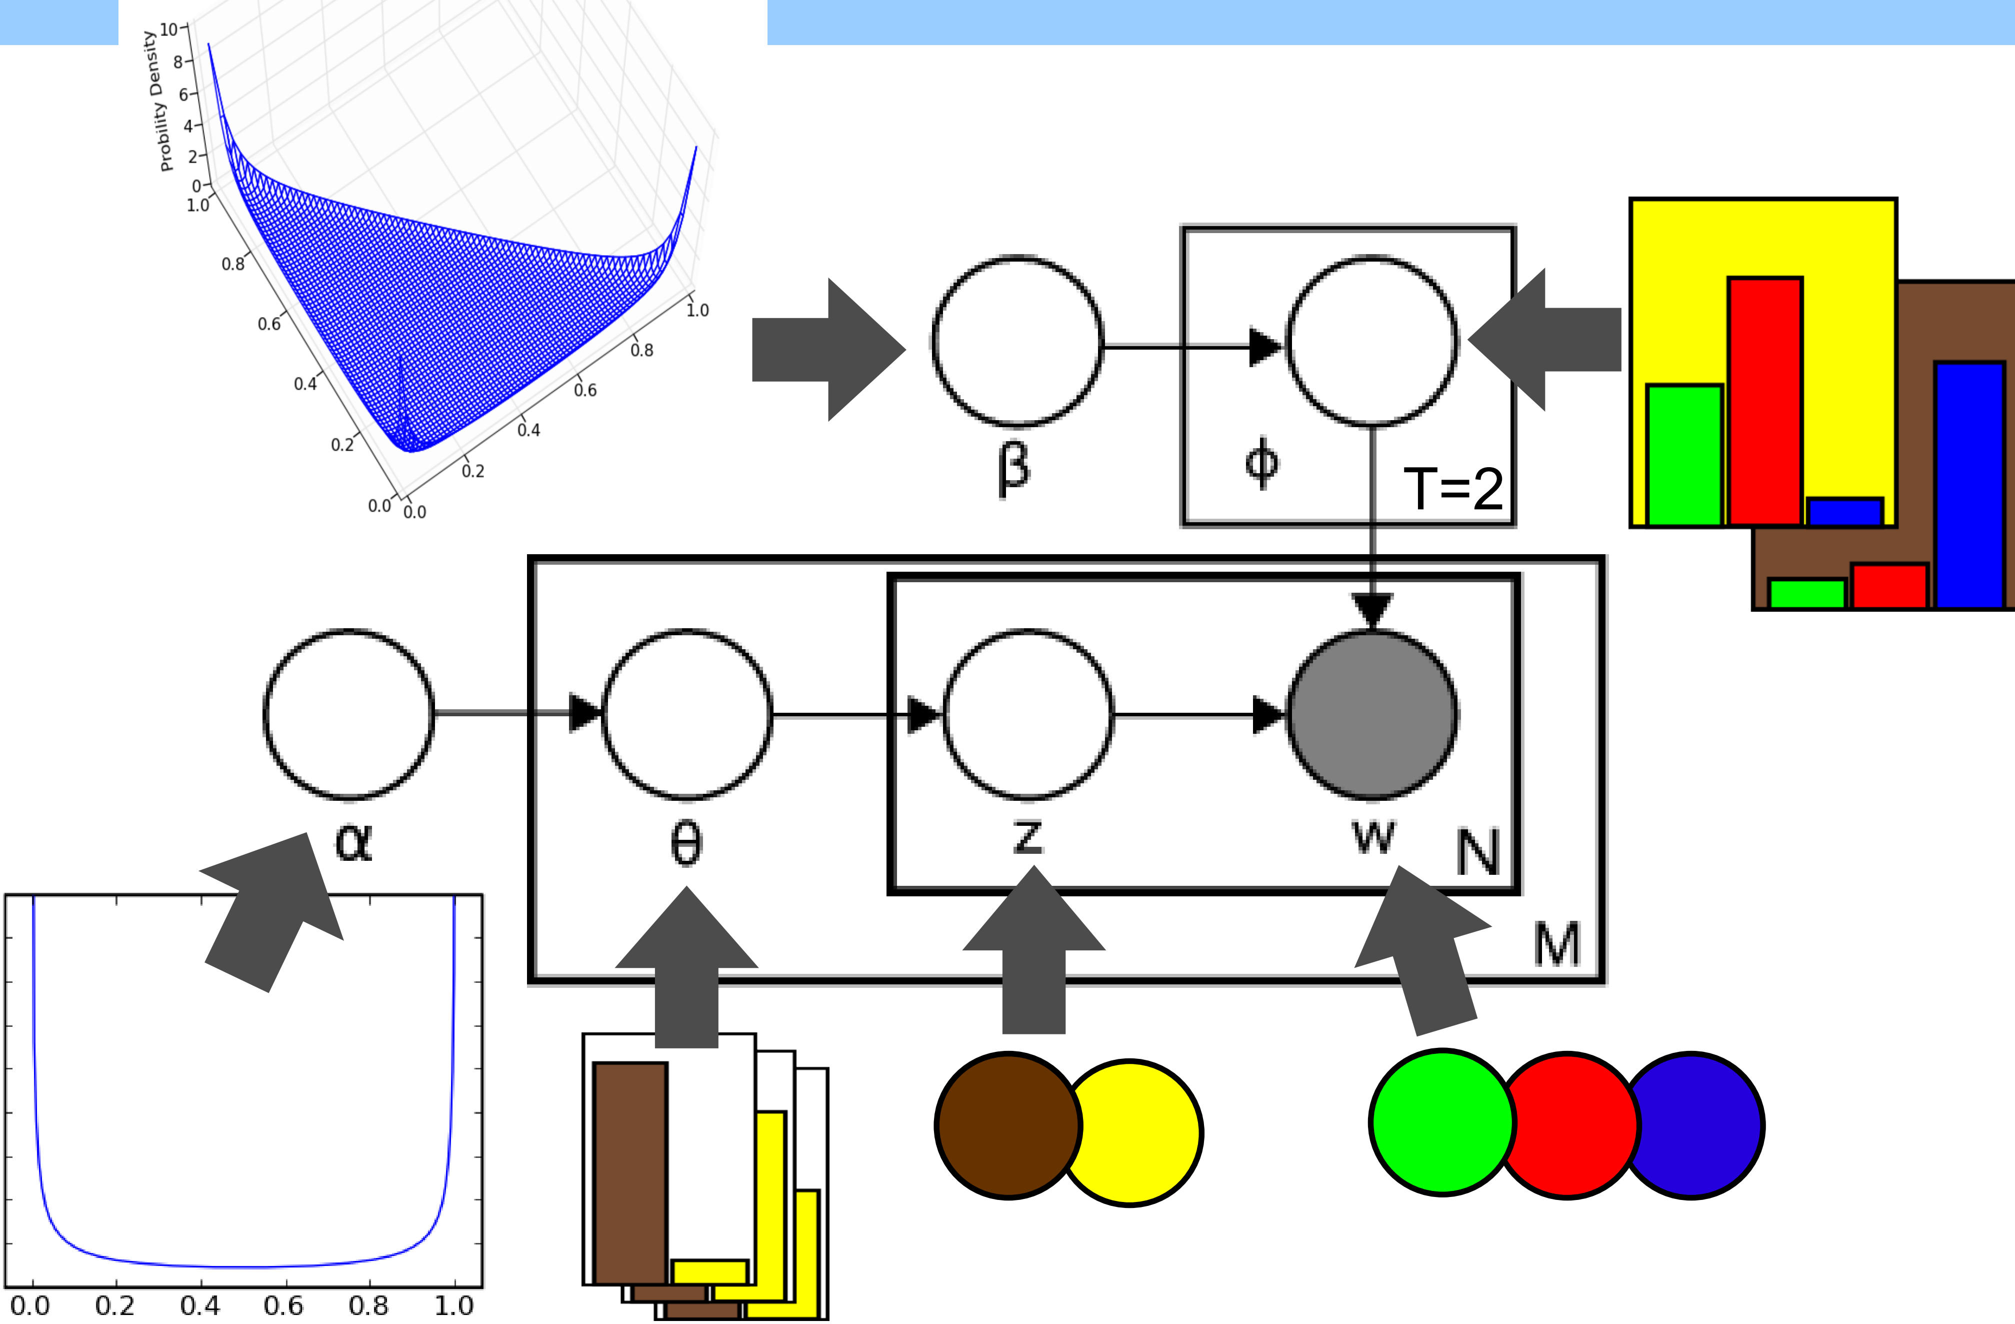

Goal

- **First Goal**: Infer most probable topic assignments for all features in a scene from co-occurrence

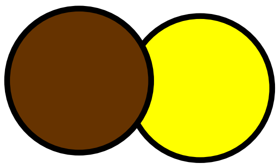

- **Second Goal**: Infer the feature distributions of the categories (to apply to unseen data)

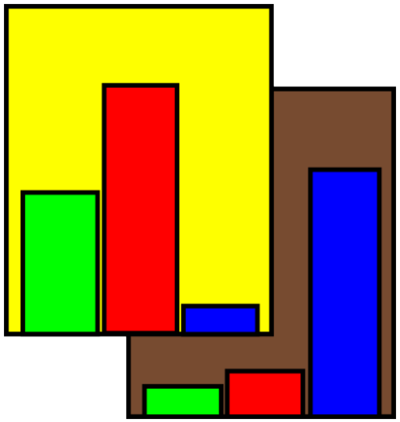

Problem

Using Bayes' rule to determine the probability of classification:

$$P(\mathbf{z} \mid \mathbf{w}) = P(\mathbf{w} \mid \mathbf{z})\, P(\mathbf{z}) \,/\, P(\mathbf{w})$$

All class assignments

All feature occurrences with assignment to data sets

Intractable, involves T^N terms

Solution

- Use Markov chain Monte Carlo for approximation of P(**z** | **w**)

1. Initialization: Assign random classes to all feature occurences
2. Gibbs Sampling: Sample the class of each feature occurrence i from $P(z_i \mid z_{\setminus i}, \mathbf{w})$
3. Repeat 2. until convergence
4. Use further samples to approximate P(**z** | **w**)
5. Use the sample statistics to approximate φ

High Level Algorithm

1. Scan scenes
2. Extract background
3. Spatial segmentation into scan segments
4. Learning of class assignments and definitions

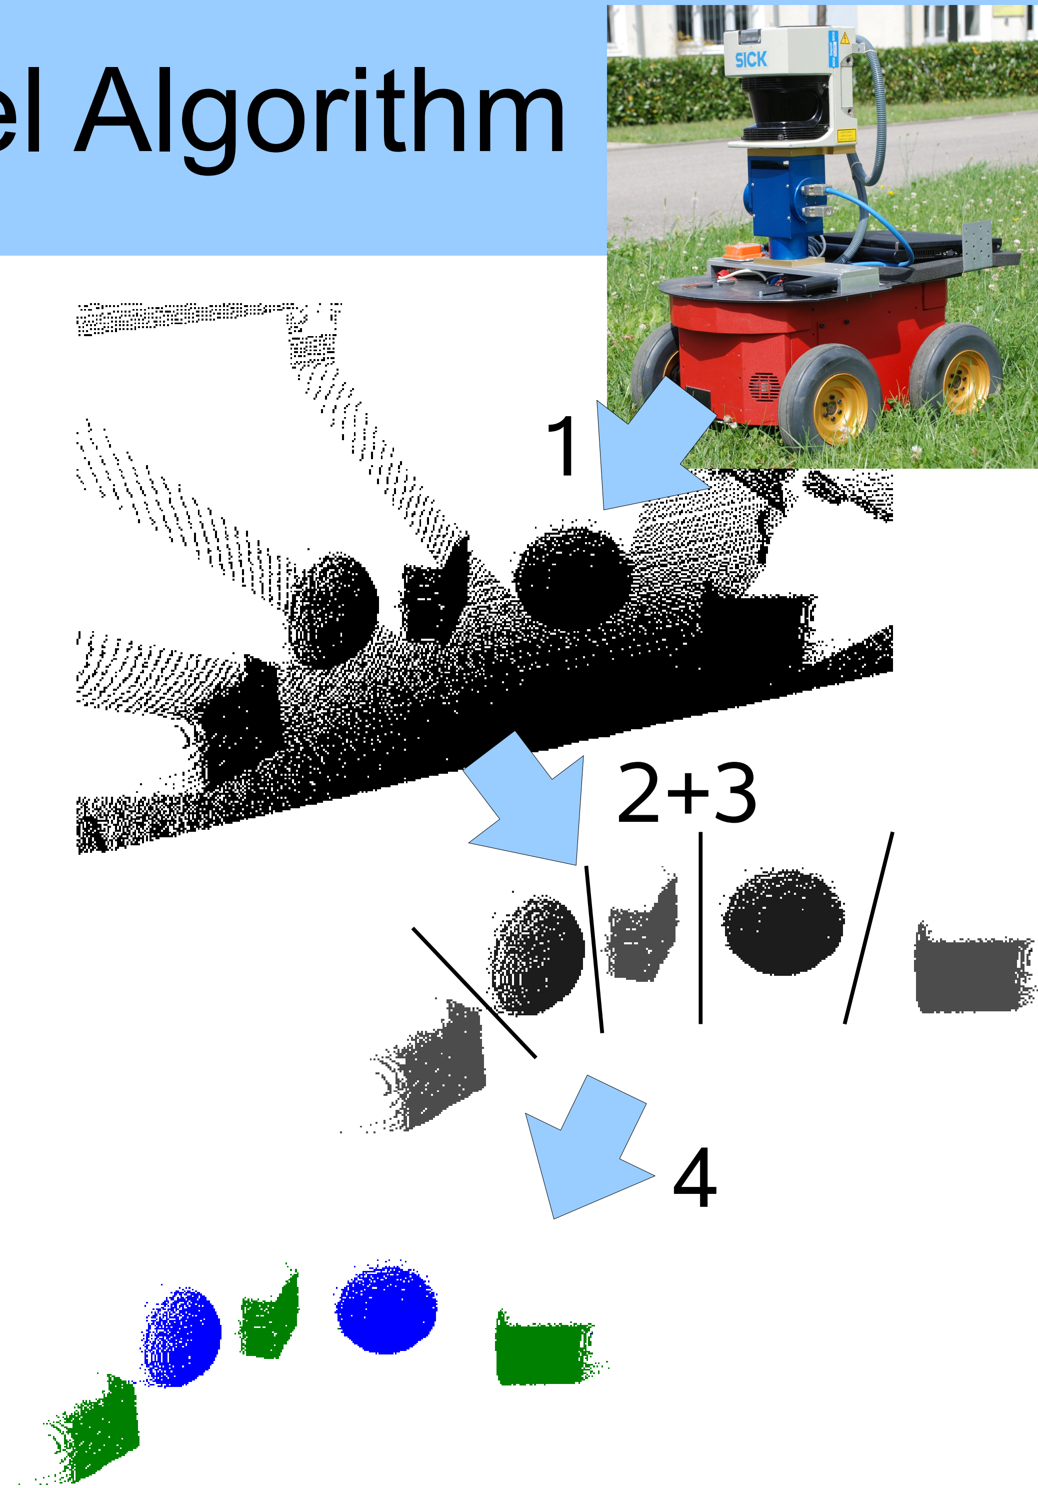

Outline

- Classification
- Features
- Learning
- **Experiments**
- Summary

Experiments

- Two corpora A/B of differing complexity
- Two/five object classes
- 12/39 scenes
- 31/82 scan segments
- Four parameters in feature feneration
- Two parameters for latent Dirichlet allocation

Experiments – Enhanced Spin Images

- Improved differentiation between object classes
- Increased similarity within object classes

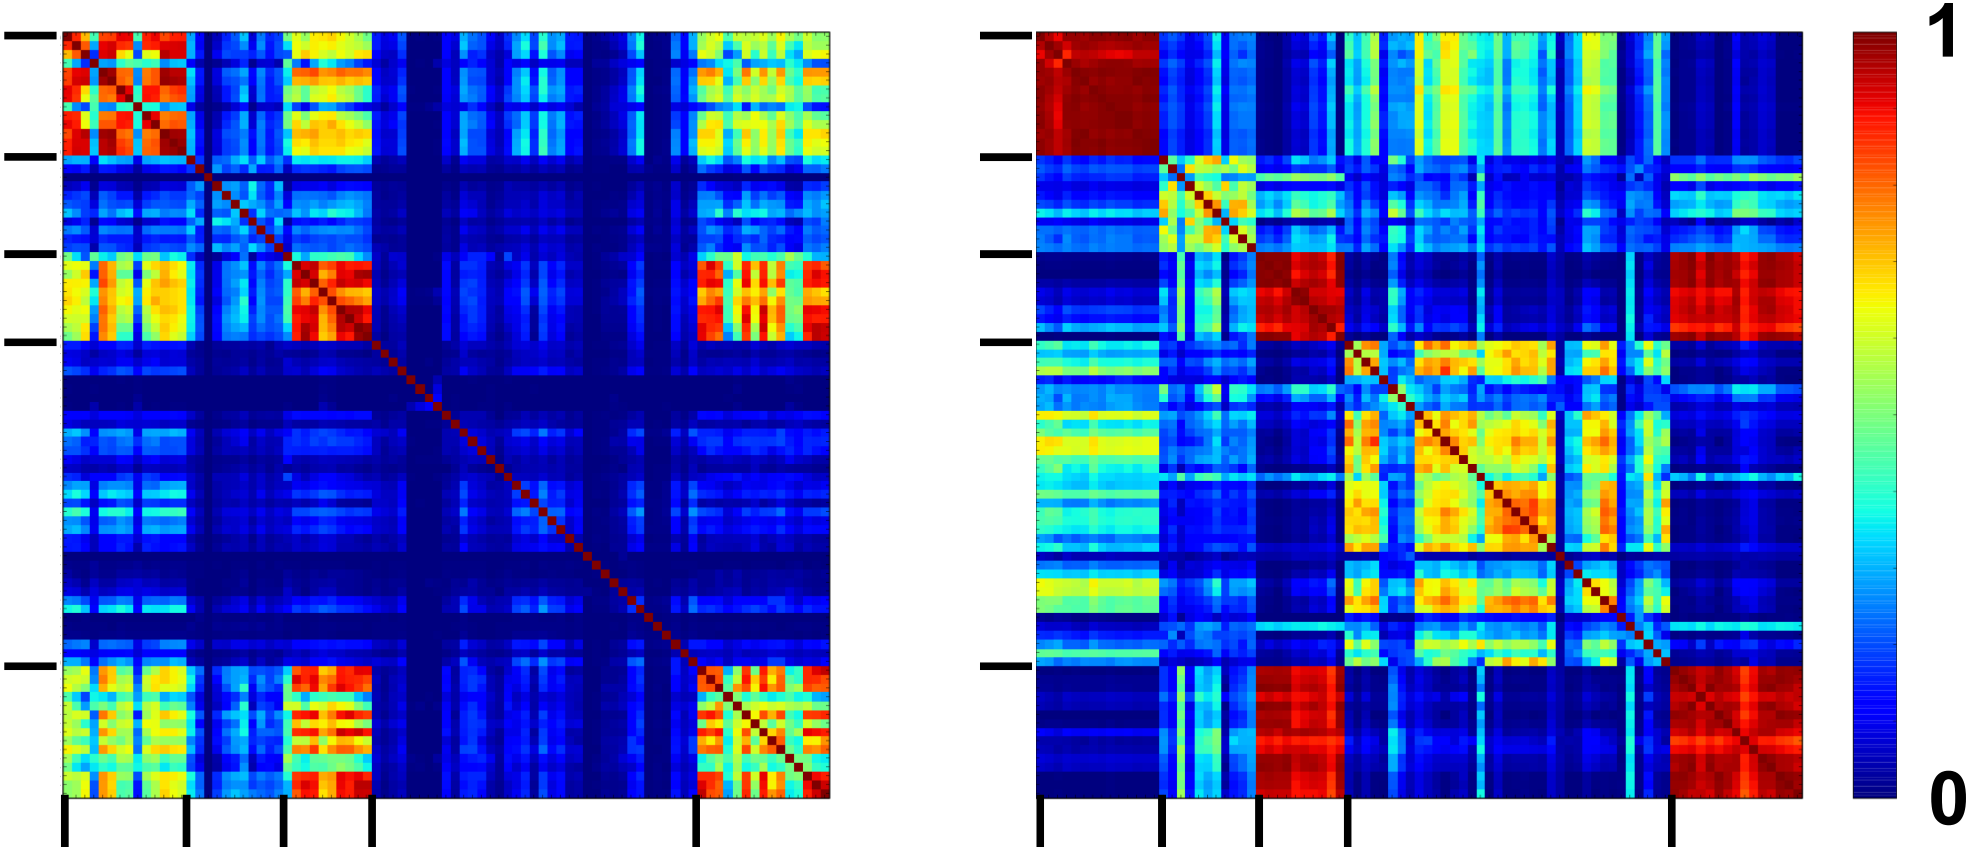

Standard Spin Images   Enhanced Spin Images

Experiments – Enhanced Spin Images

- Improved differentiation between object classes
- Increased similarity within object classes

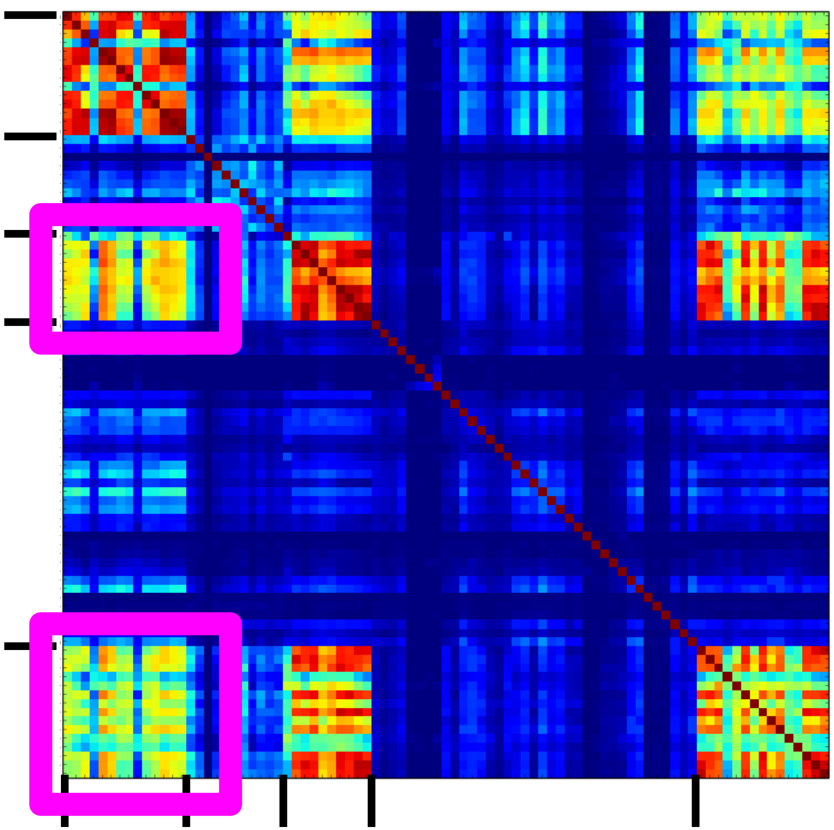

Standard Spin Images

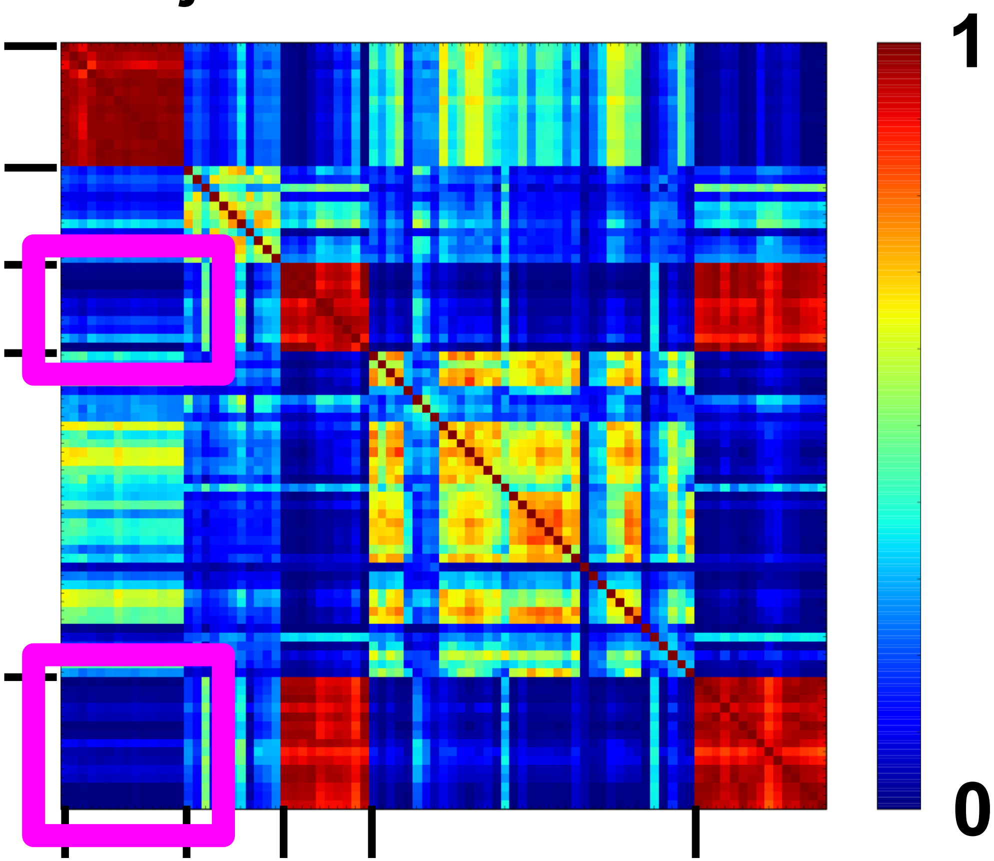

Enhanced Spin Images

Experiments – Enhanced Spin Images

- Improved differentiation between object classes
- Increased similarity within object classes

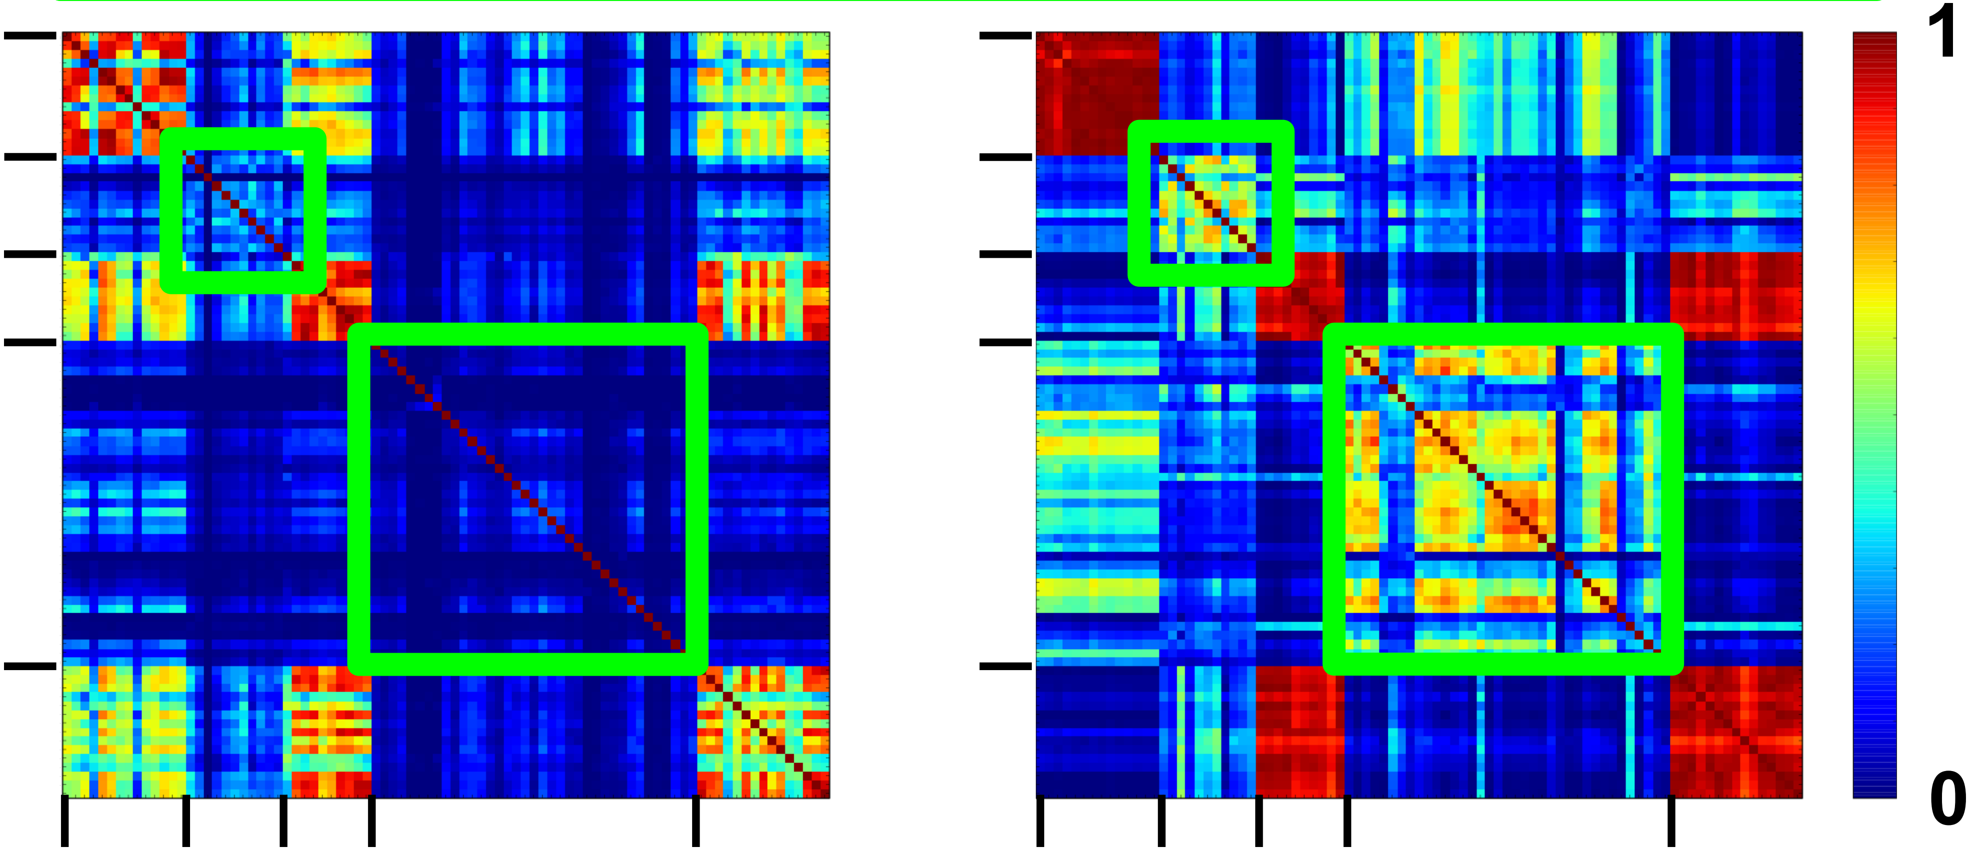

Standard Spin Images   Enhanced Spin Images

Experiments – Corpus A

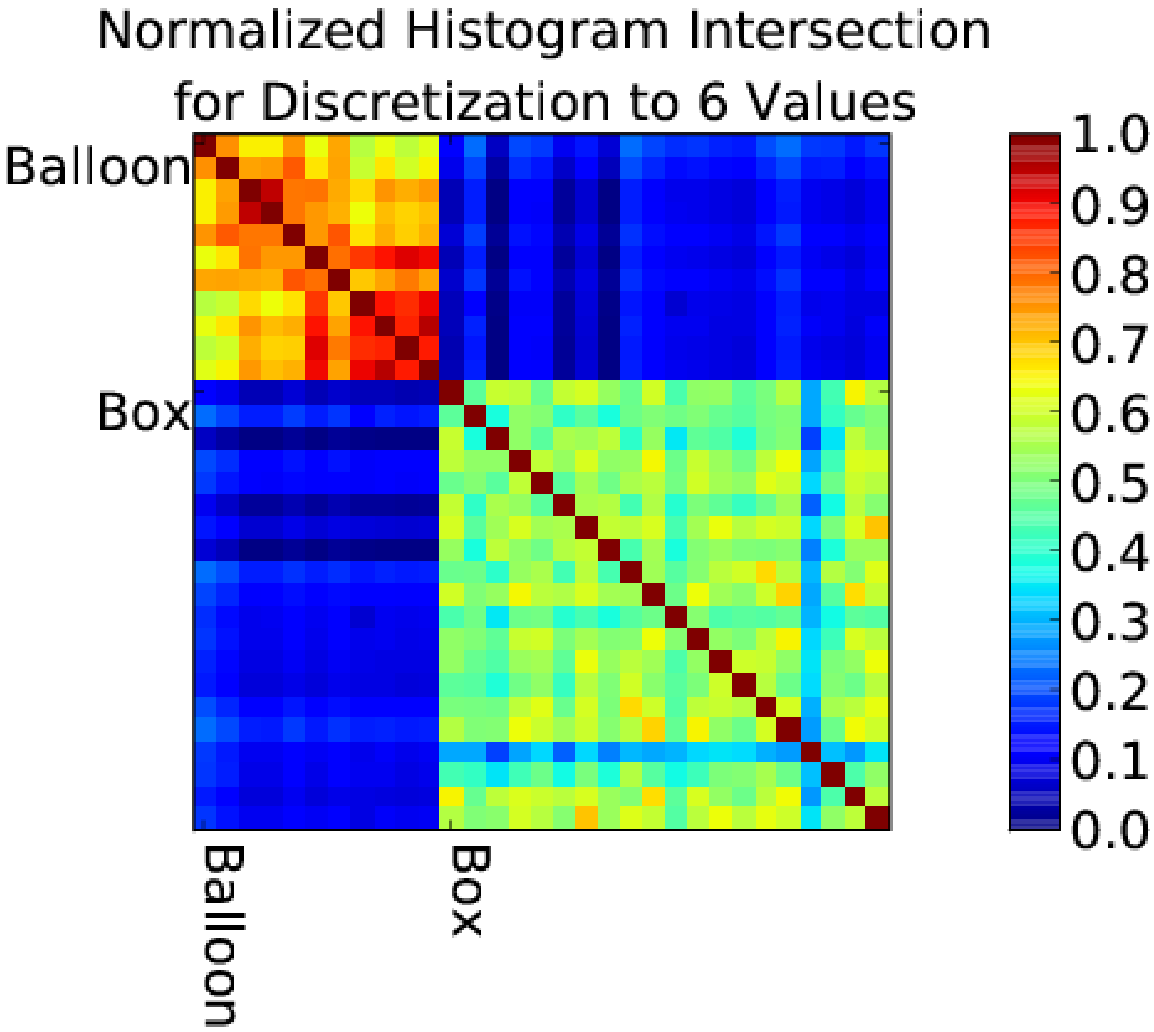

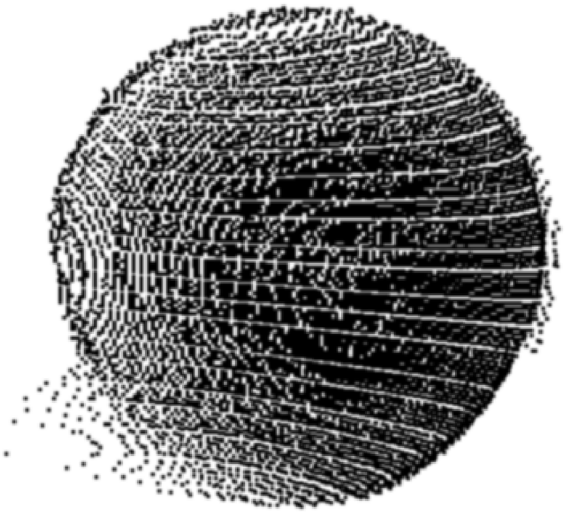

Experiments – Corpus A

LDA
Latent Dirichlet Allocation

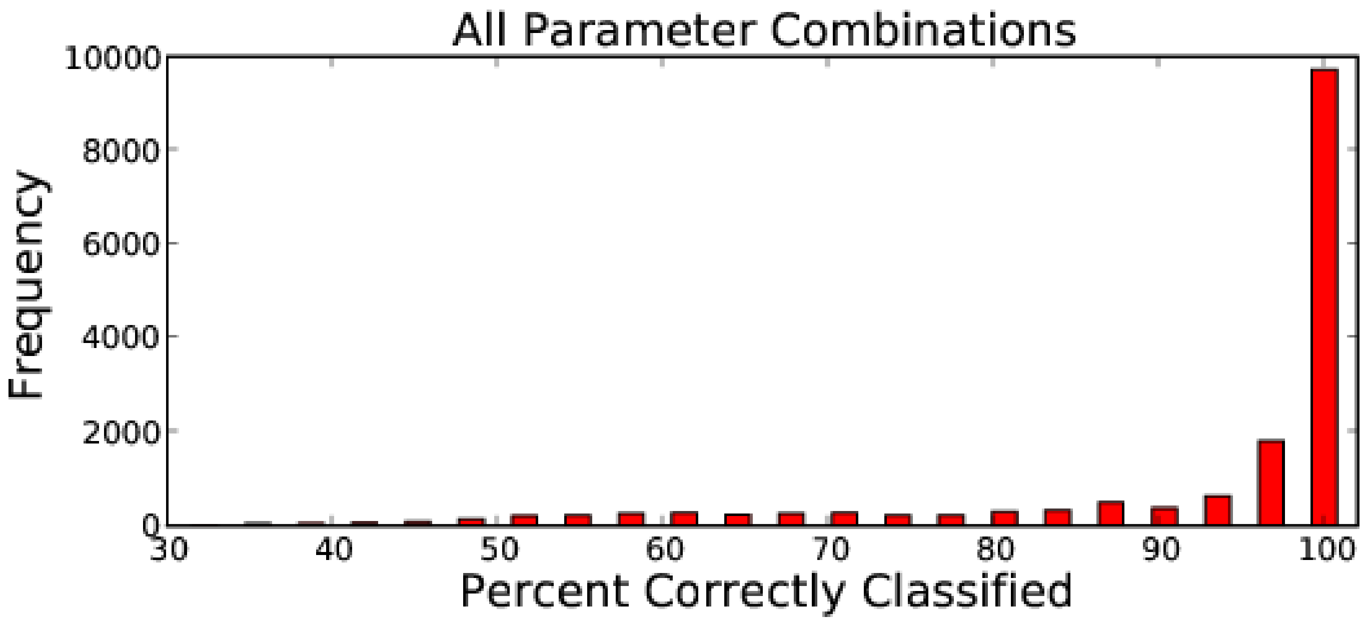

HC
Hierarchical Clustering

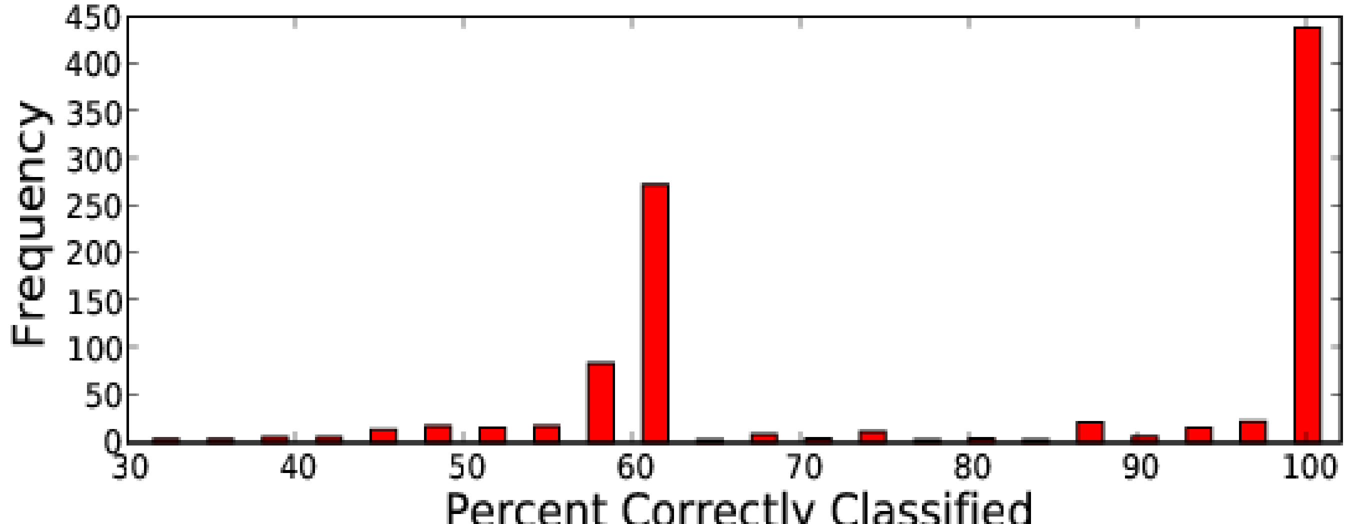

Experiments – Corpus B

- More Difficult Data:
  - Additional Object Classes
  - Similar Classes
  - Variation within Classes

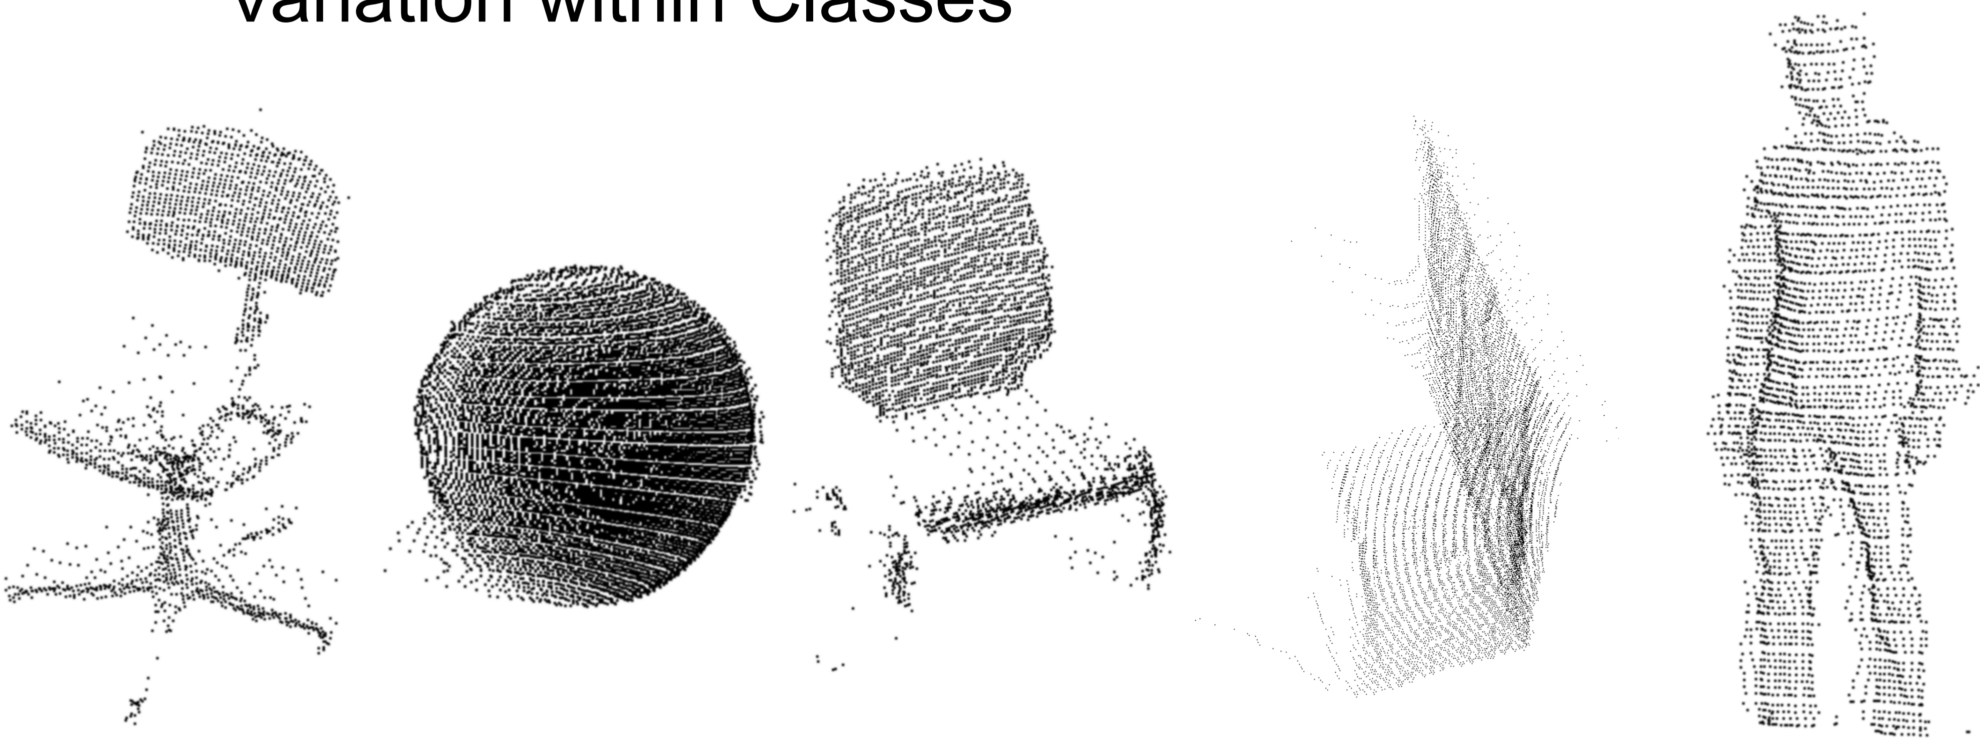

Experiments – Corpus B

- Class differentiation is more difficult
- Feature parameters become very important
- Reliable parameter settings needed

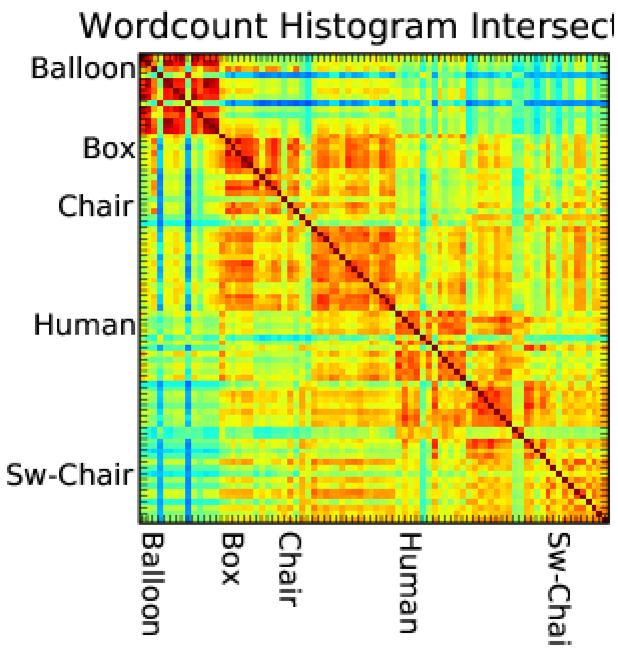

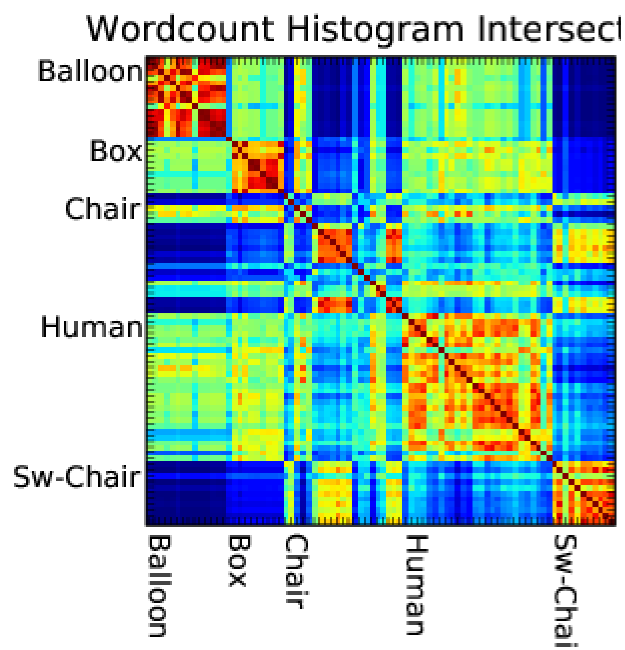

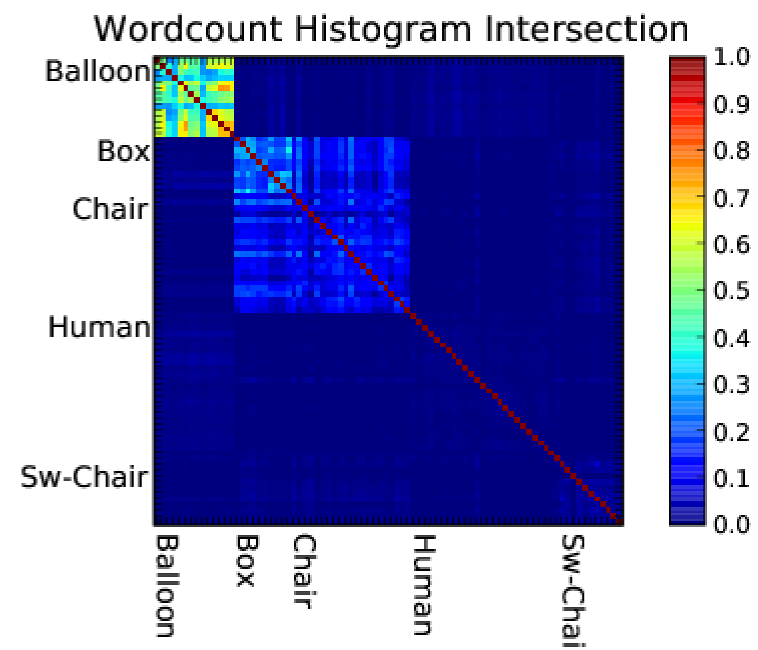

Experiments – Corpus B

Feature Parameter Selection Results

- Feature Type:
  Enhanced Spin Imgs
- Support Distance:
  Small (10cm)
- Raster Resolution:
  Low (3x3 to 5x5)
- Discretization:
  5 to 27 Values

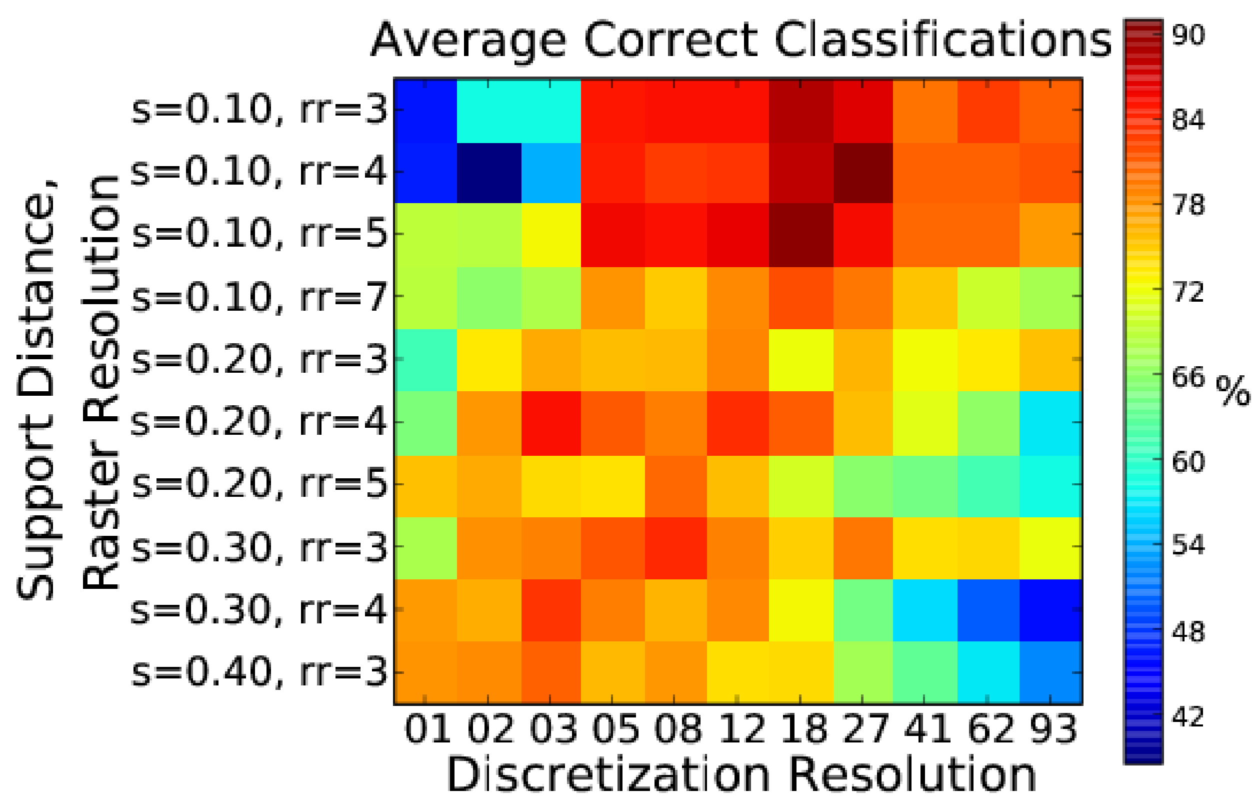

Experiments – Corpus B

LDA Hyperparameter Selection Results

- Alpha < 1.0
- Beta < 0.4

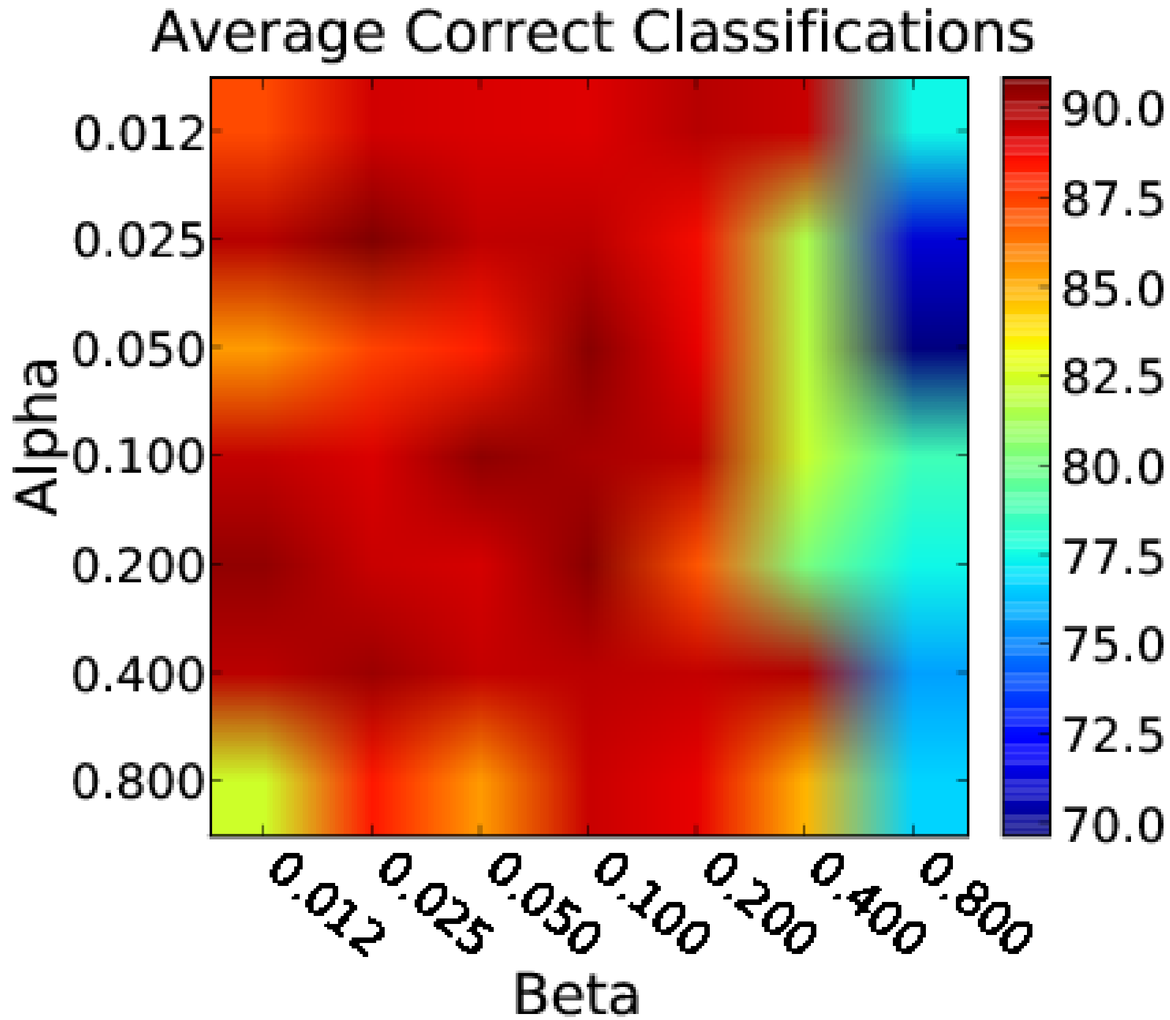

Experiments – Corpus B

HC Parameter Selection Results

- Four parameters for feature generation
- One clustering parameter (linkage type)
- No robust parameter settings found

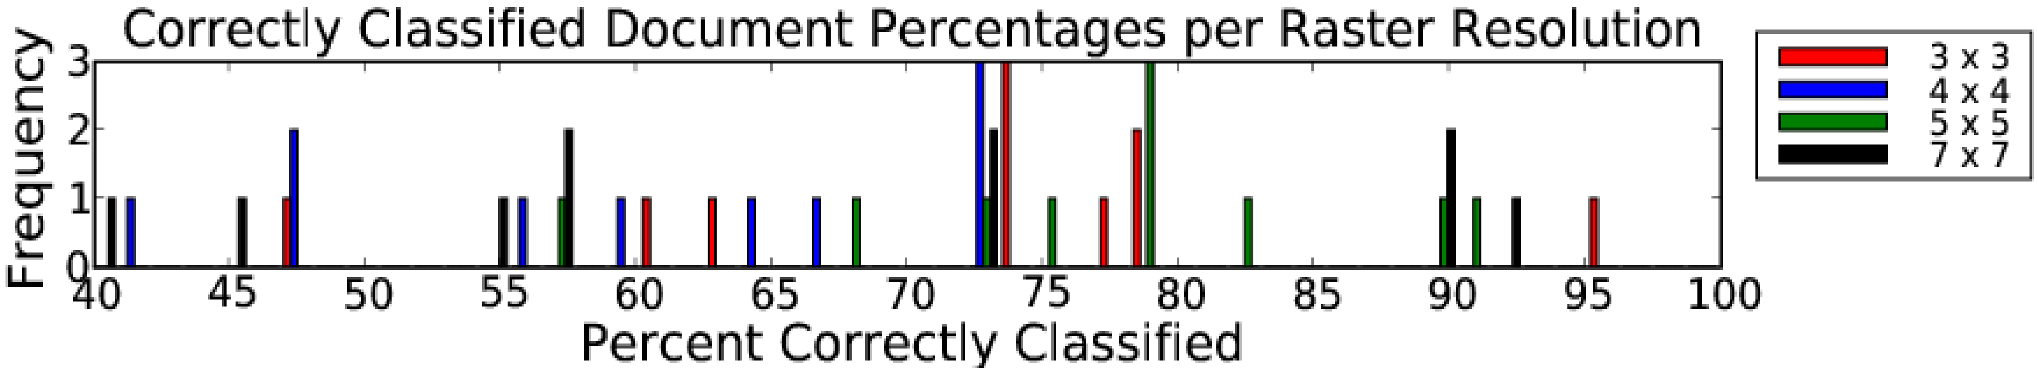

Experiments - Results

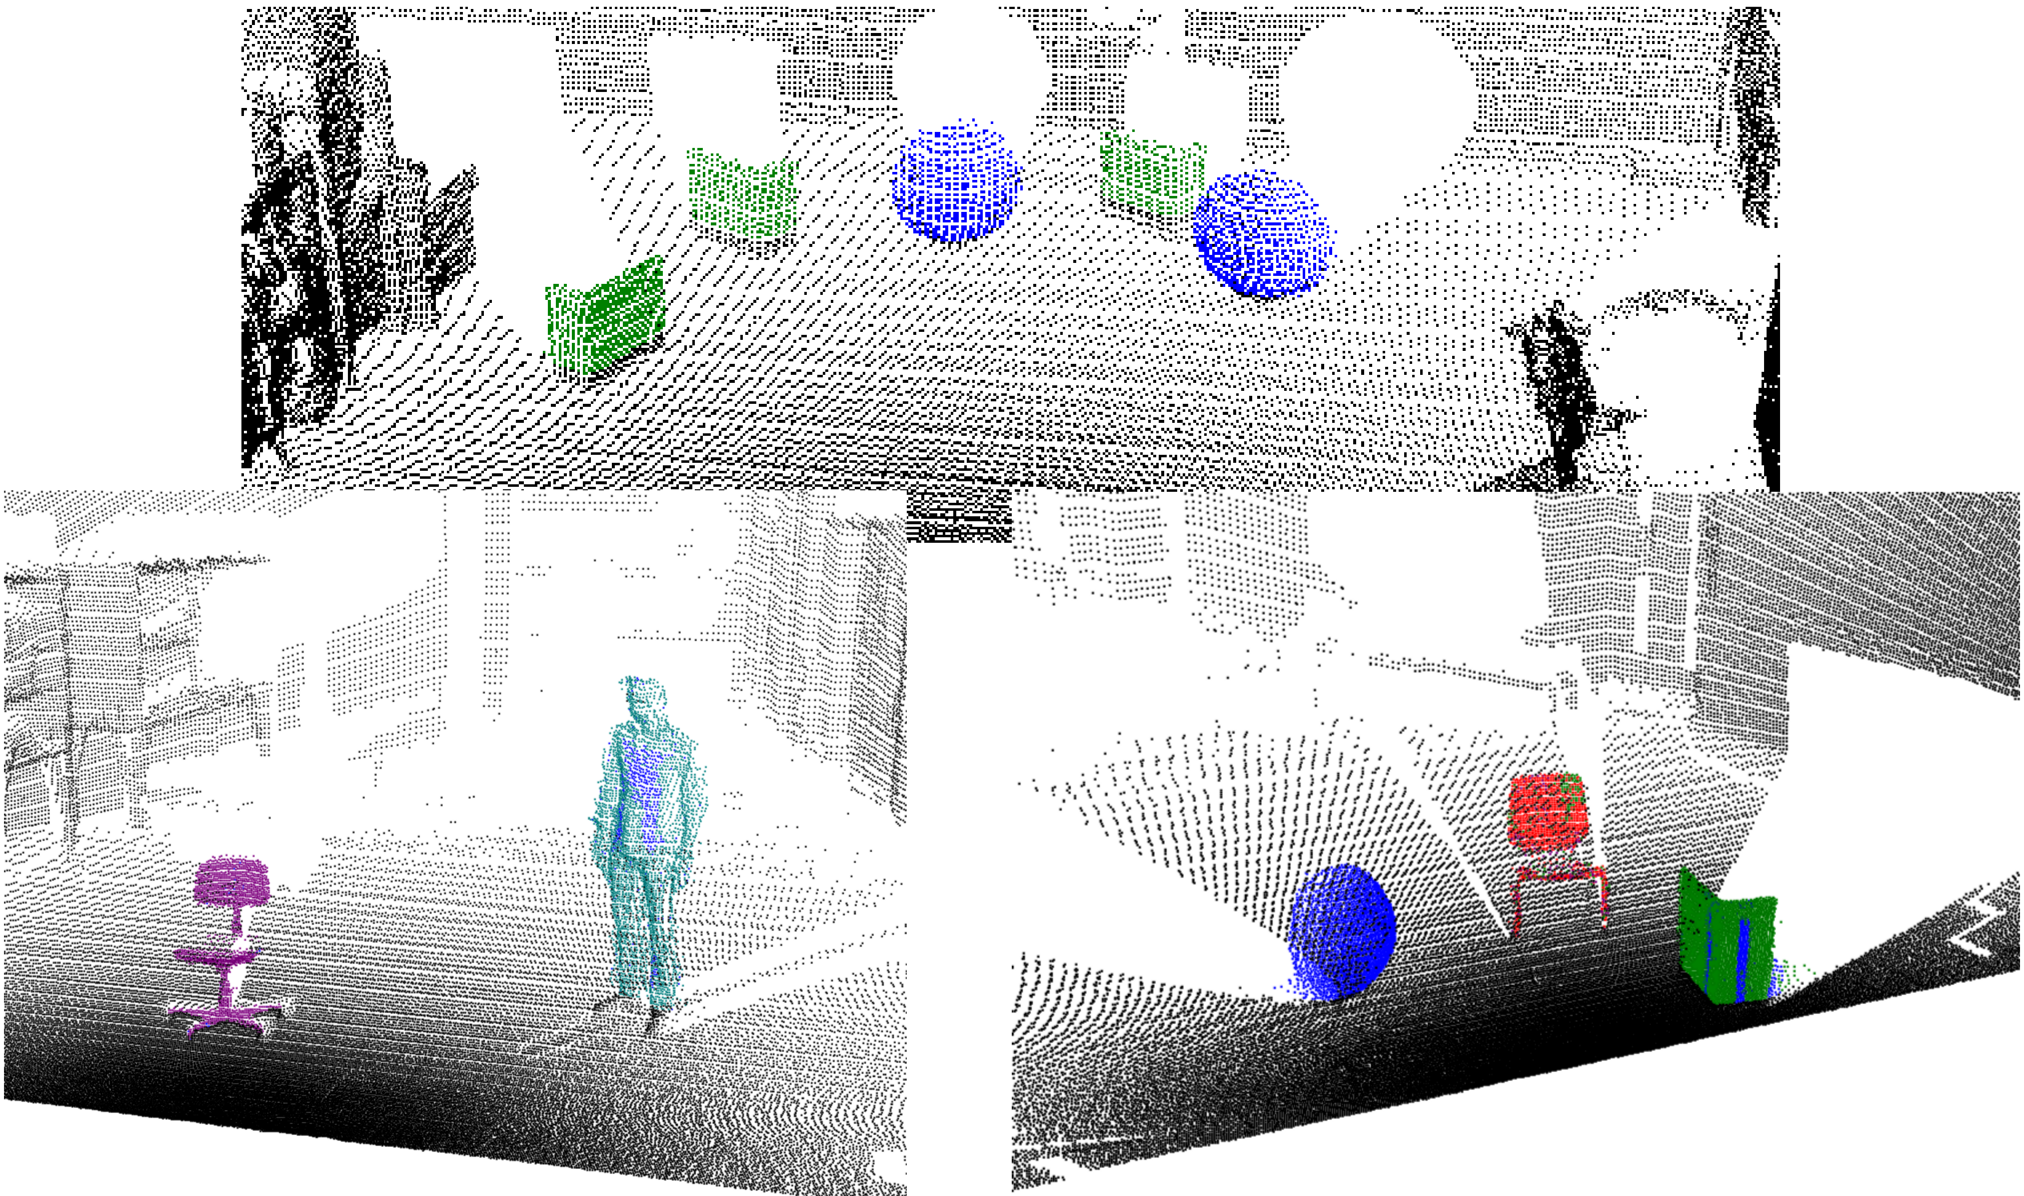

Summary

- Shape-based discovery of object classes
- Clustering of feature distributions
- Spin image enhancements improve differentiation
- LDA greatly outperforms hierarchical clustering
- Highly satisfactory classification performance

References

- Johnson, A. *“Spin-Images: A Representation for 3-D Surface Matching”*, 1997
- Blei, D. M. *“Latent dirichlet allocation”*, 2003
- Griffiths, T. L. and Steyvers, M. *“Finding scientific topics”*, 2004

Thanks for Listening

Questions?