CSE 255 – Lecture 4

Data Mining and Predictive Analytics

Predicting image content using social metadata

Images on the web

Can we make better **predictions** about image content by understanding the **social network?**

1. Predicting content on image sharing networks

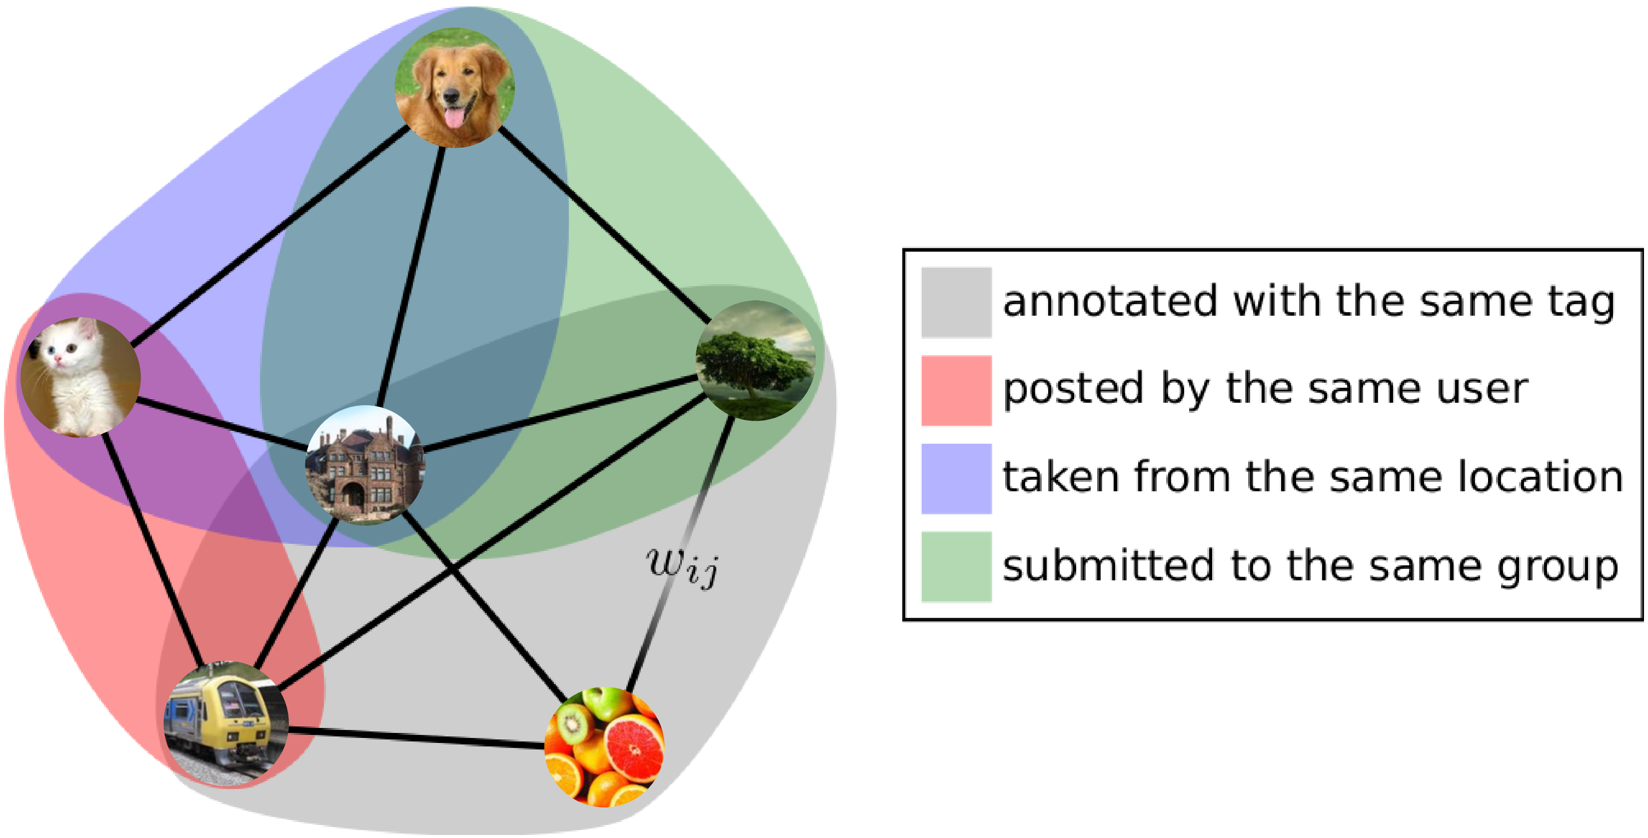

"**What** are the objects that appear in the image and how are images **related**"

ECCV 2012 (w/ Leskovec)

Dataset description

We collected Flickr metadata from four popular classification benchmarks:

	CLEF	PASCAL	MIR	NUS	Total
#photos	4546	10189	14460	244762	**268587**
#users	2663	8698	5661	48870	**58522**
#groups	10575	6951	21894	95358	**98659**
#comments	77837	16669	248803	9837732	**10071439**
#locations	1007	1222	2755	22106	**23745**
#labels	99	20	14	81	**214**

(Everingham et al. 2010; Huiskes et al. 2008; Nowak et al. 2010; Chua et al. 2009)

Does shared **social data** imply shared **labels**?

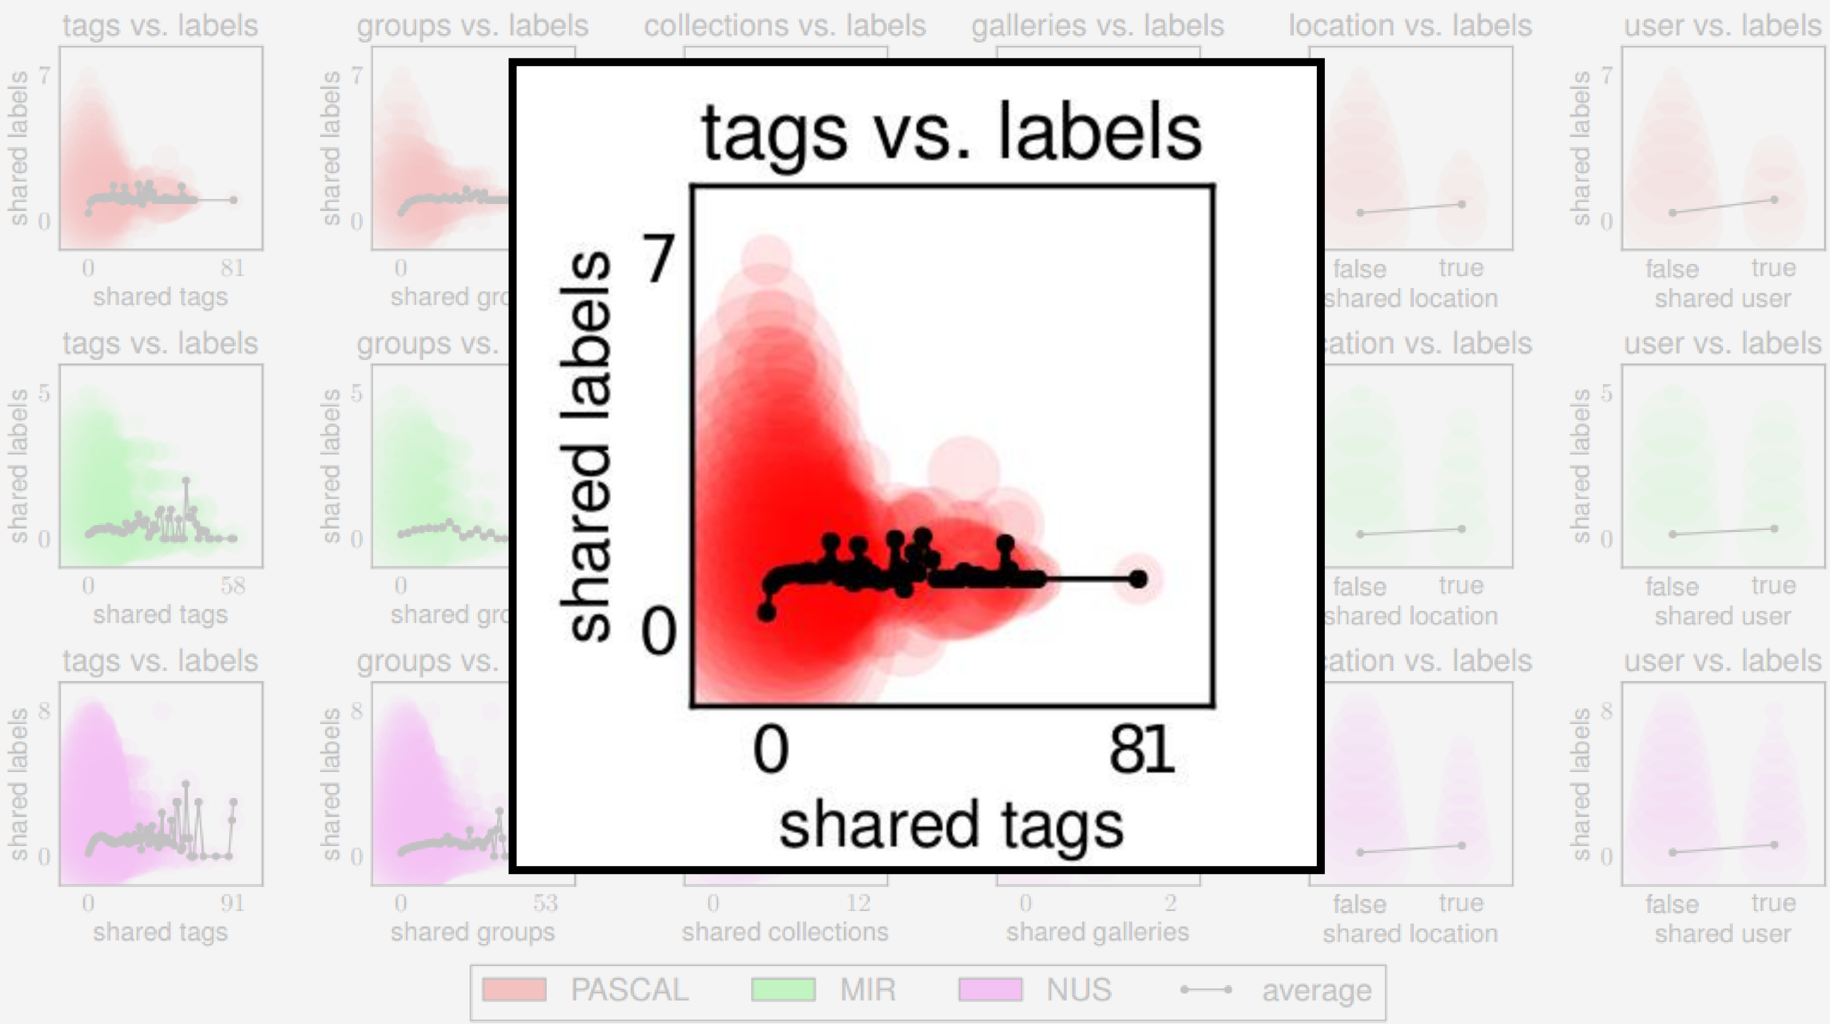

Model

Images that have **common social data** should have a higher probability of having the same labels

1. Build a **graph** of related images

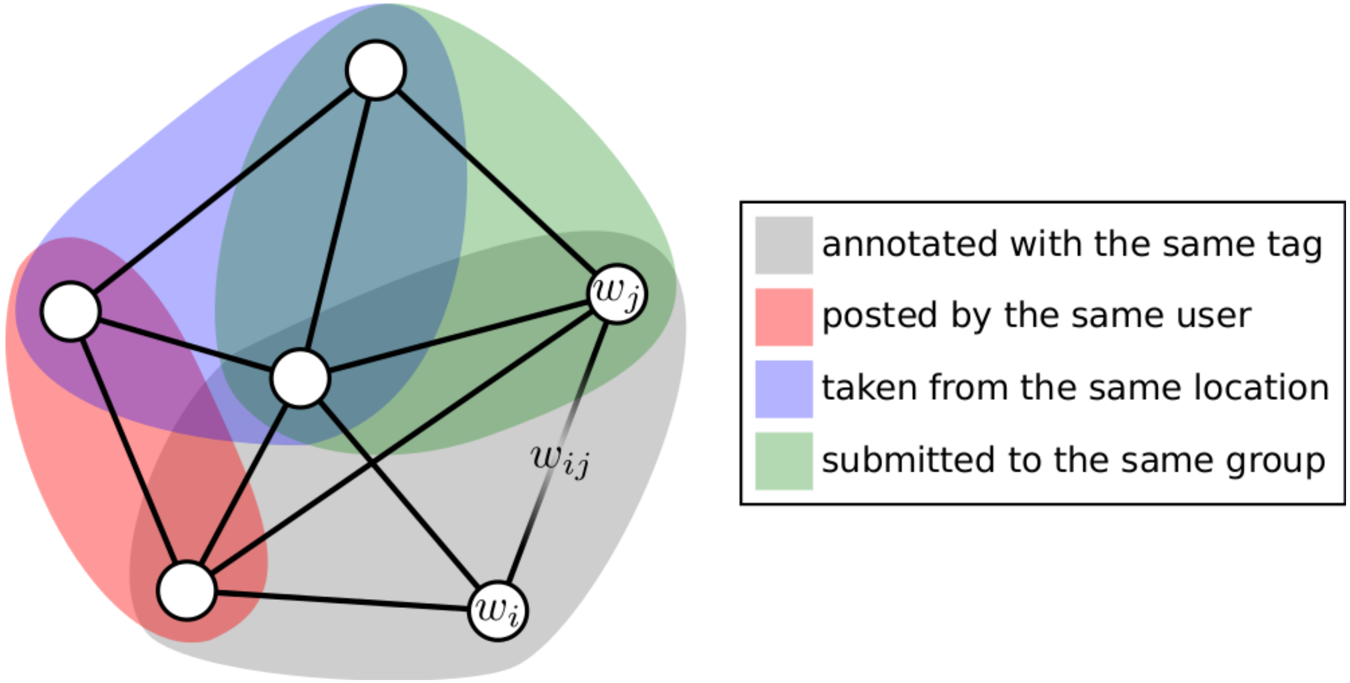

Model

Images that have **common social data** should have a higher probability of having the same labels

2. Build a **model** of image labels

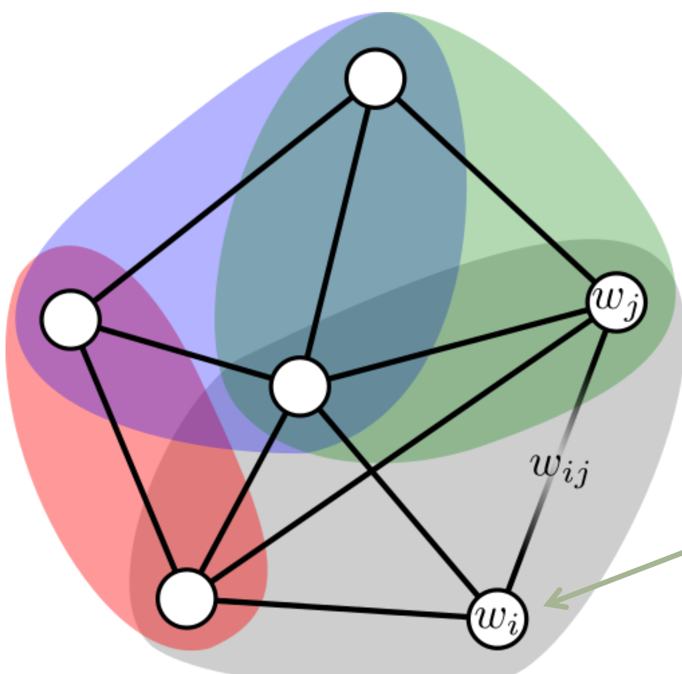

Image features

Weights for class *c*

$$\underbrace{\langle \phi_c(x_i), \theta_c^{\text{node}} \rangle}_{w_i}$$

Weight (or cost) of a positive label

Model

Images that have **common social data** should have a higher probability of having the same labels

2. Build a **model** of image labels

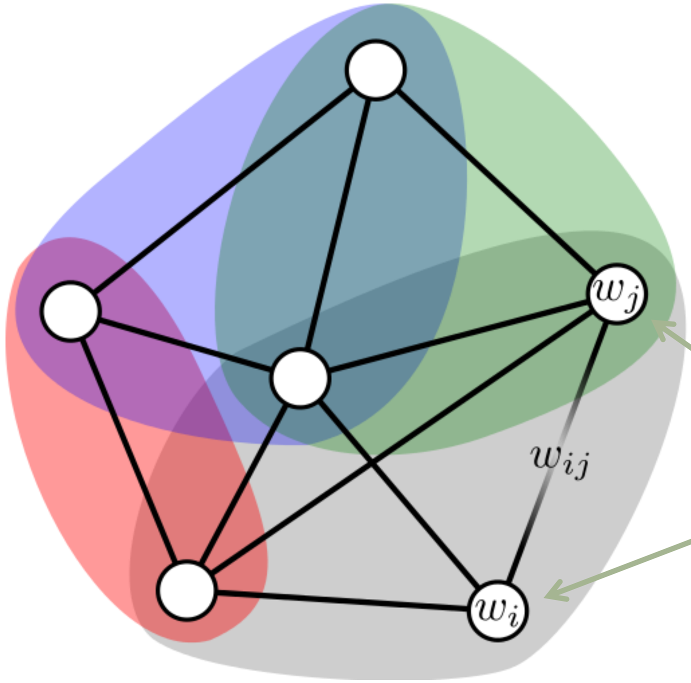

Social network features

$$\underbrace{\langle \phi_c(x_i, x_j), \theta_c^{\text{edge}} \rangle}_{w_{ij}}$$

Weight (or cost) of label disagreement

Model

Images that have **common social data** should have a higher probability of having the same labels

3. **Inference**: minimize cost for all labels

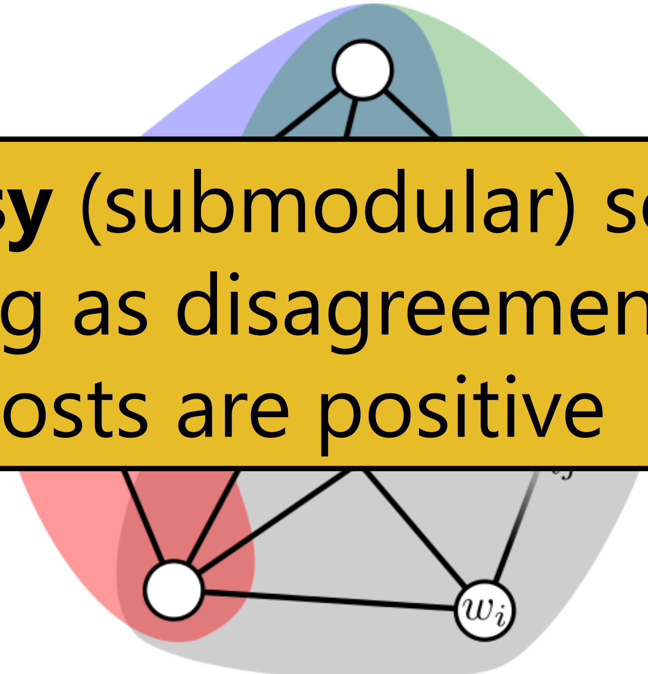

Easy (submodular) so long as disagreement costs are positive

Local cost

$$y_i \cdot \underbrace{\langle \phi_c(x_i), \theta_c^{\text{node}} \rangle}_{w_i}$$

$$+$$

$$\delta(y_i = y_j)\, \underbrace{\langle \phi_c(x_i, x_j), \theta_c^{\text{edge}} \rangle}_{w_{ij}},$$

Disagreement cost

Model

Images that have **common social data** should have a higher probability of having the same labels

4. **Learning**: choose parameters that minimize training error

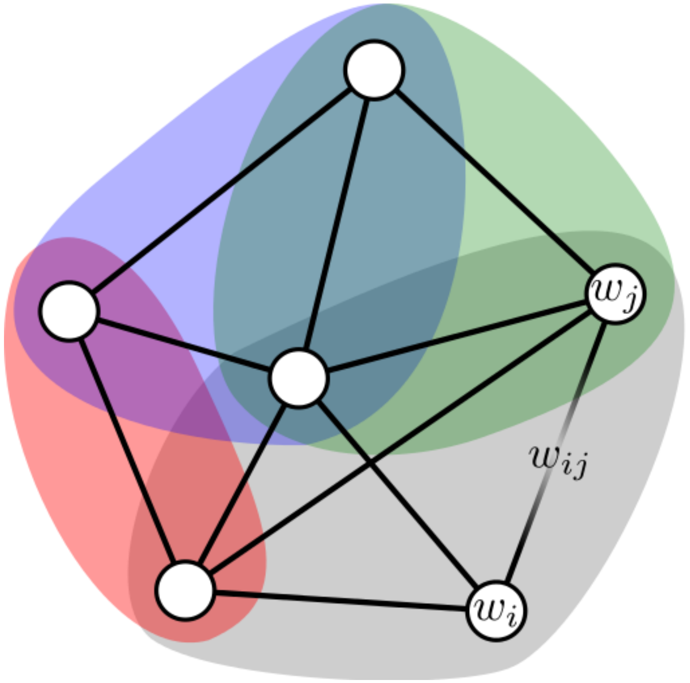

$$\underset{\Theta}{\operatorname{argmin}} \left[\underbrace{\Delta(\bar{Y}(\mathcal{X};\Theta), Y_c)}_{\text{empirical risk}} + \underbrace{\frac{\lambda}{2} \|\Theta\|^2}_{\text{regularizer}} \right]$$

We learn parameters through **structured prediction** (Tsochantaridis et al. 2005)

Performance: image classification

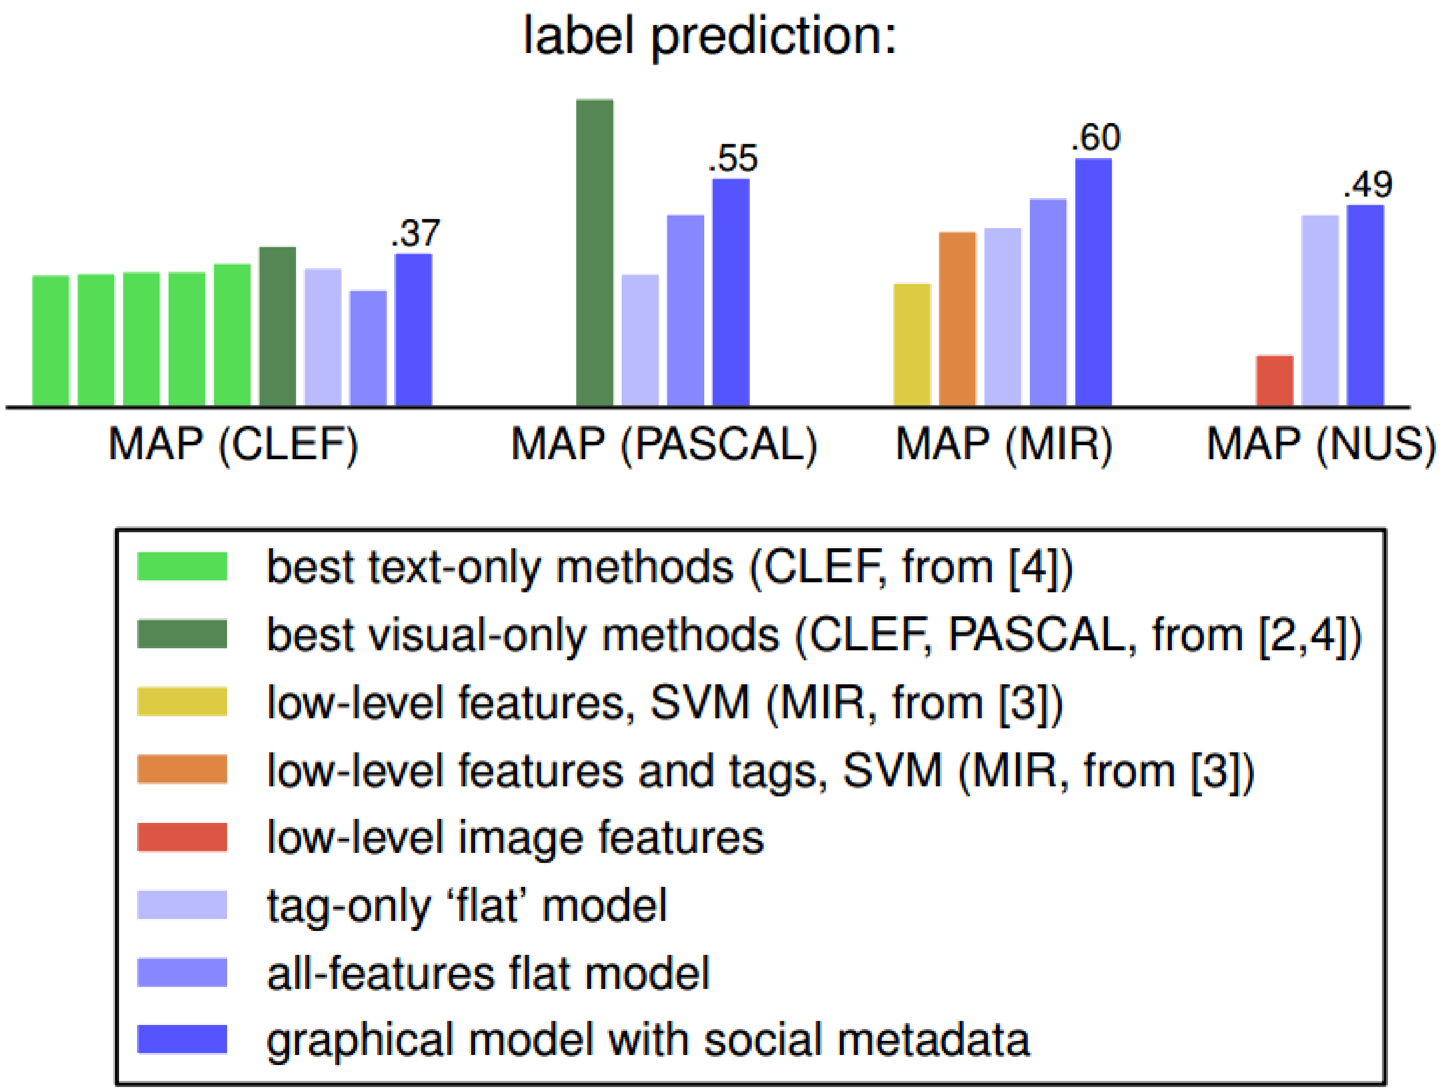

Performance: tag and group prediction

tag recommendation:

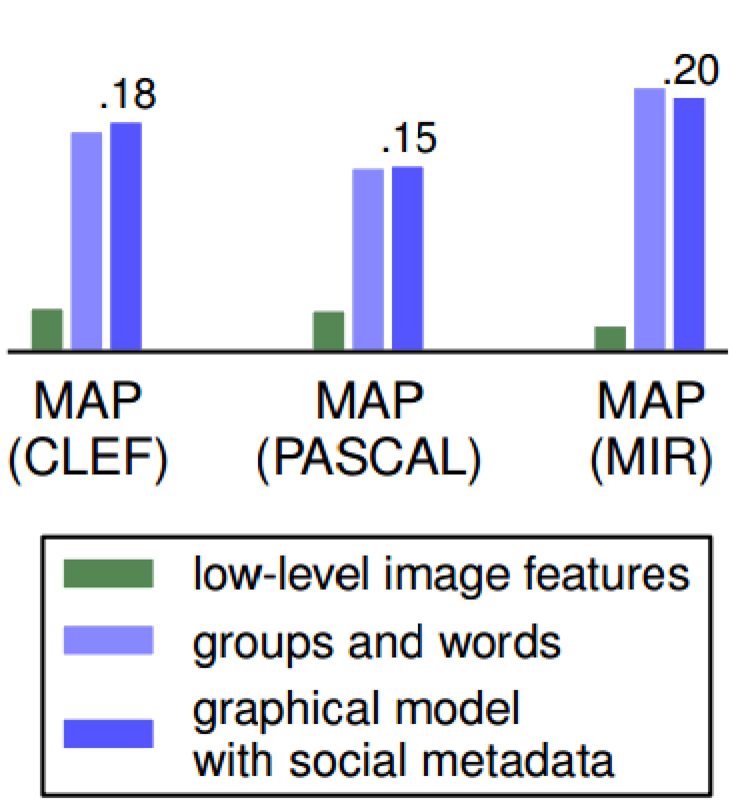

group recommendation:

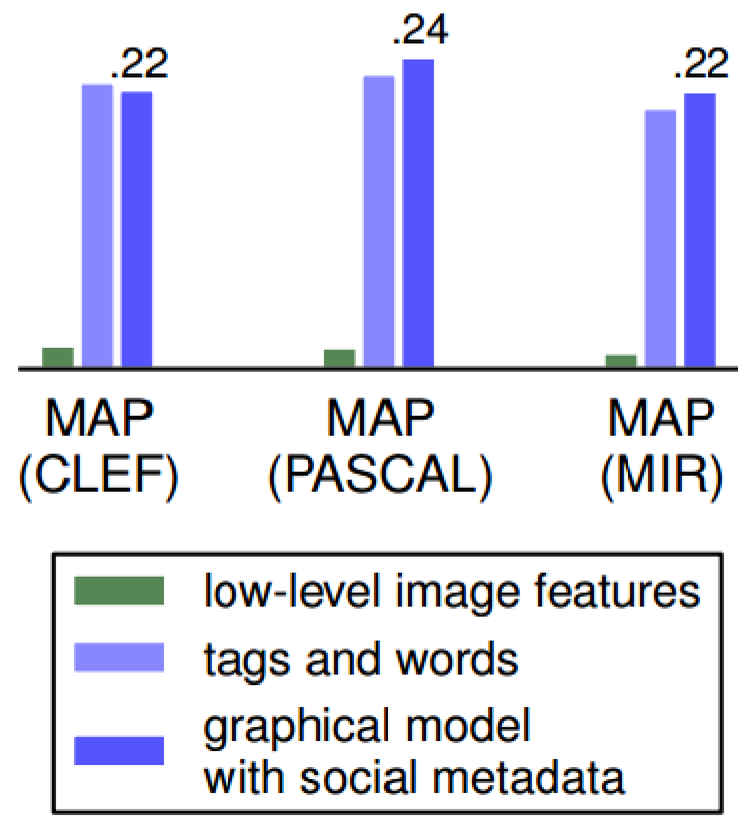

Which features are predictive?

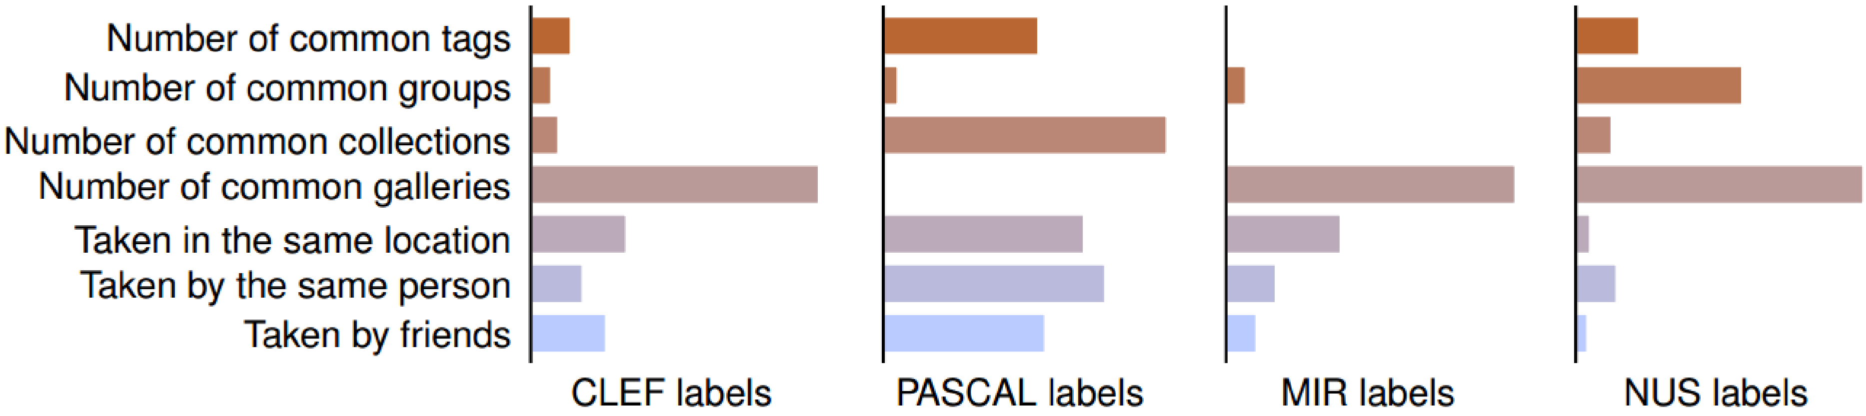

- Which features are informative depends on the dataset
- **Galleries** tend to be the most predictive, as they tend to be organized around a common concept
- The **identity of the user** and their **location** are somewhat important on most datasets

Questions?