Album-Oriented Face Recognition For Online Social Networks

Zhongkai Han, Syed Zain Masood, Jason Hochreiter,
Spencer Fonte and Marshall Tappen

Presenter: Baoyuan Liu

April 23, 2013

People Like to Photograph Events

UCF

Notice: The same small set of people appear in the photos

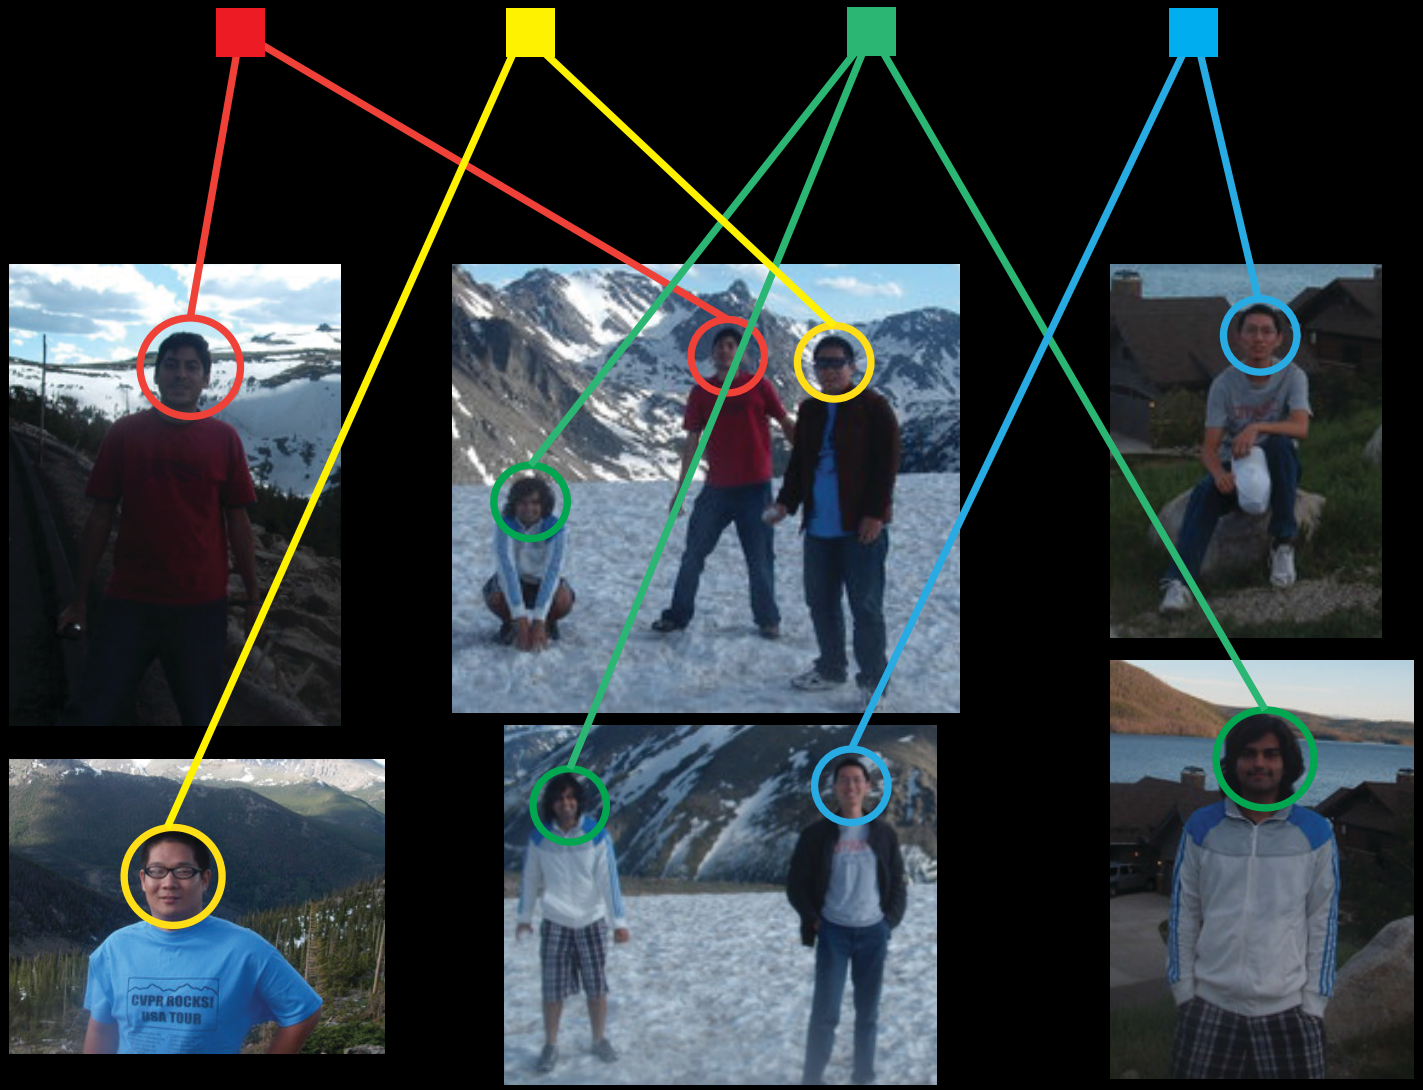

UCF

How can we use the *structural* information of the albums to improve recognition?

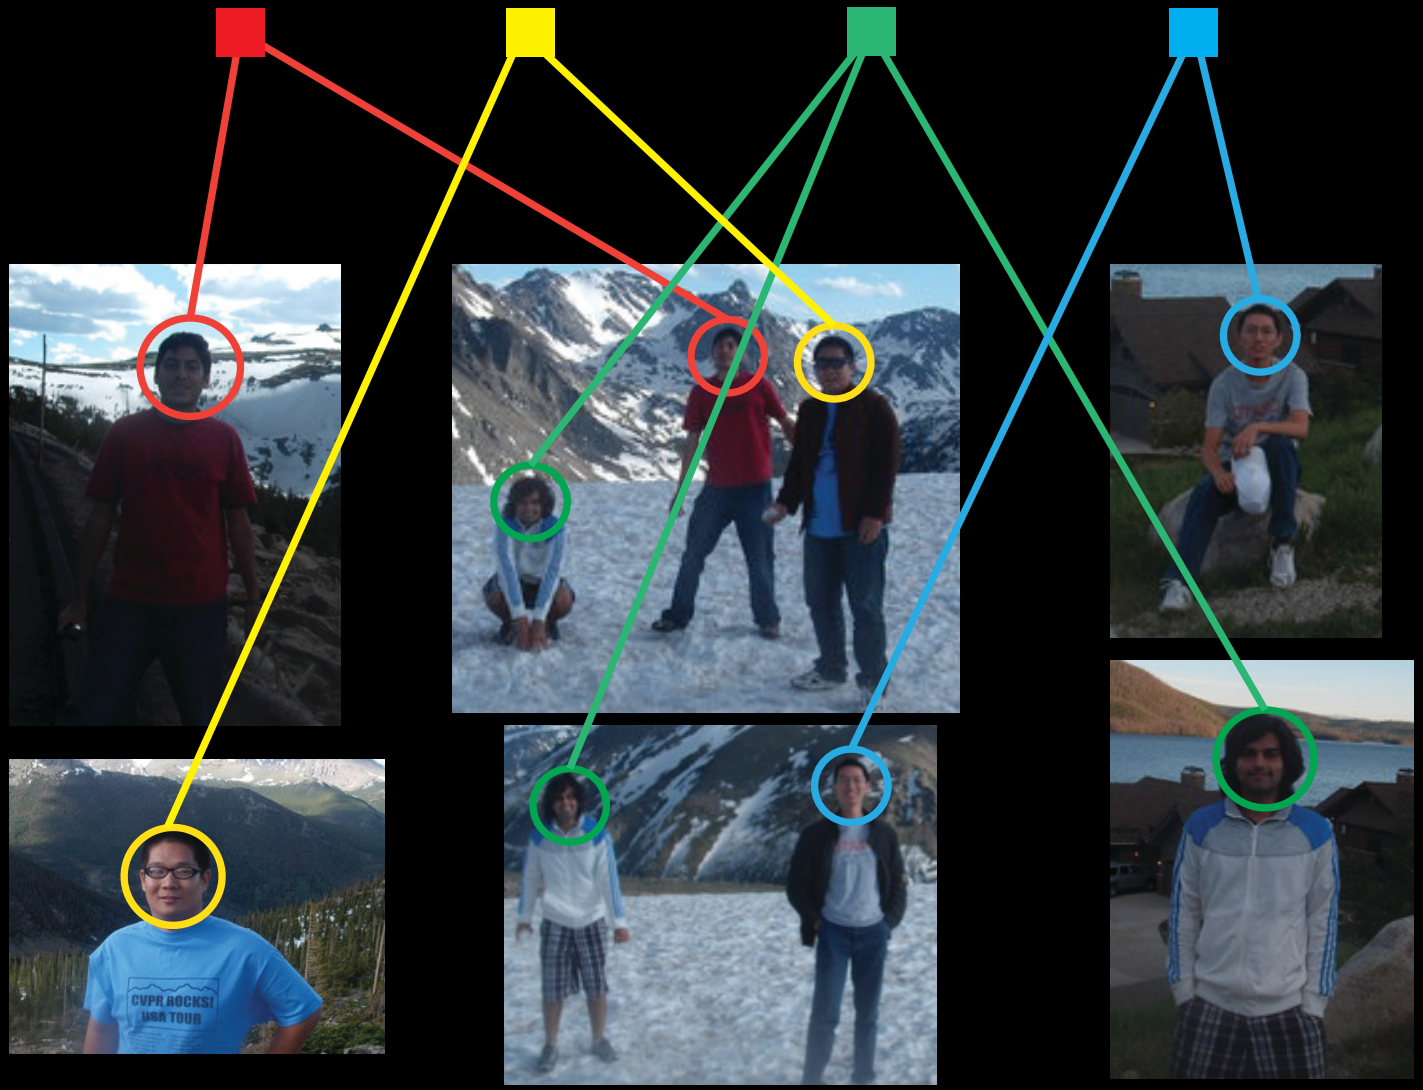

UCF

Idea of Our Paper

- Basic assumption: albums tend to contain multiple photos of a small number of people
- We propose a mathematical model that combines album information and image features to significantly improve recognition
- Extend with other social features based on facebook meta data

Previous Work on Co-occurrence

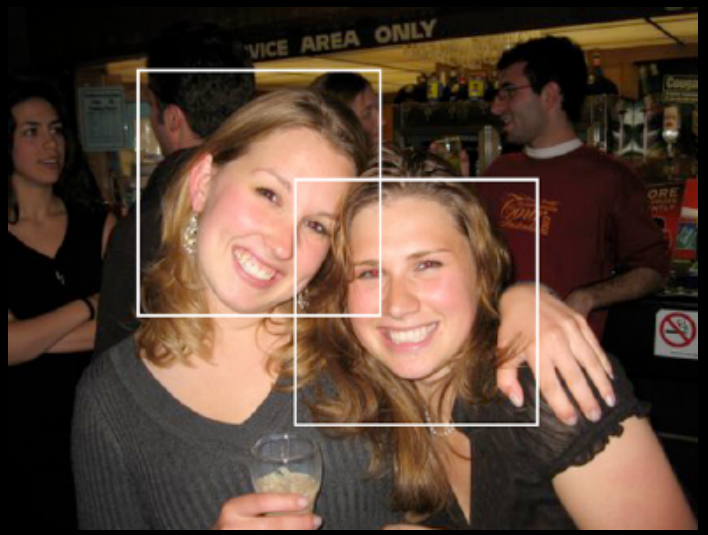

- Stone et al. proposed an image-level co-occurrence model
  - Consider the frequency of people appear together in single photos
  - Requirement: Only considered photos with ***exactly*** two faces.
  - Is this requirement too **strict** to apply to realistic facebook images?

Z. Stone, T. Zickler, and T. Darrell. Autotagging Facebook: Social network context improves photo annotation. IEEE, 2008.

UCF

Quantitative Study

- Use Facebook API to download all pictures visible to a single user's account
- Total: 8078 pictures of 2849 people
  - Considered only tagged faces
  - Validated with OpenCV face detector

Model Applicability

- 71% photographs(5735/8078) contained only one tagged face
- Image-level co-occurrence only helps when there are two or more people in the photograph
- A model based on image-level co-occurrence is only useful for 29% of photos

We treat Co-occurrence differently

- We consider co-occurrence at the album level
  - Need not be restricted by number of people appearing in each photograph
  - Albums provide additional social information that can benefit recognition

UCF

An Example

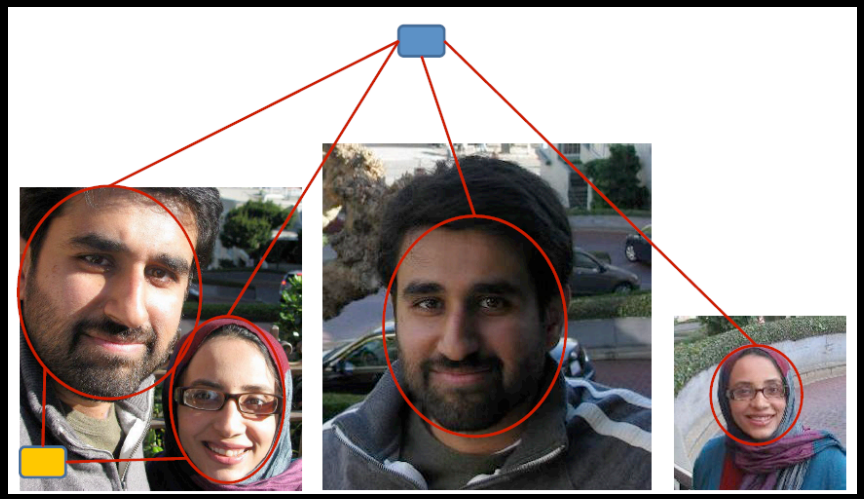

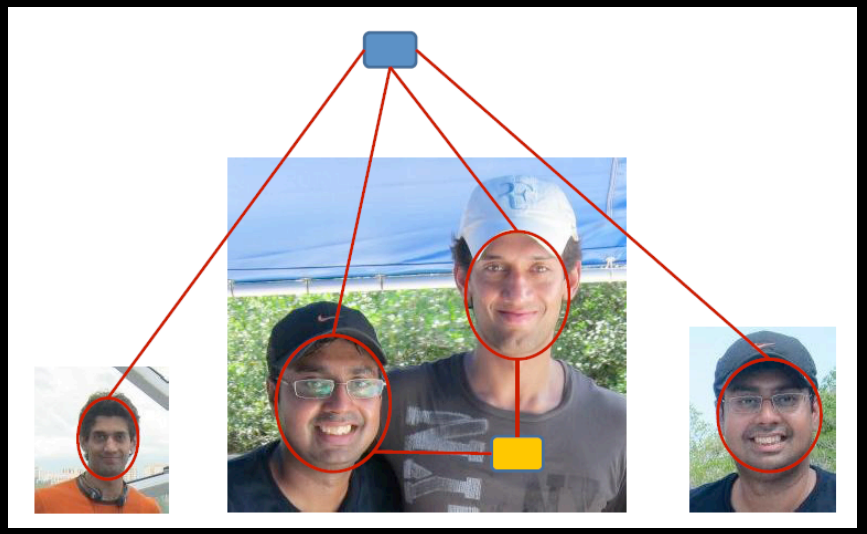

- For these two albums, an image-level co-occurrence model (yellow) can help with only 1 of the 3 photos
- A model based on album co-occurrence (blue) can help recognize faces in all photos

UCF

Can Album Level Co-Occurrence Help?

- Suppose an album prior is only useful if there are at least twice as many photos as people in an album
  - High standard intuitively, but valid for 57% of photos
  - When standard is not met, recognition degrades only 2%

UCF

Our method

- Our method – introduce a label cost
- Assign a cost to each distinct label present in an album, regardless of the number of times it appears
- Effectively limits the number of people appearing in an album
- Matches our notion of a typical album

UCF

Formulation

- A traditional face recognition system can be described by the energy function

$$E(\vec{y}) = D(y; \vec{x}) = \sum_{f \in F} D_f(y_f; \vec{x}_f)$$

All facial images in album *F*

where D_f is the data cost of assigning label y_f to face $f \in F$ and $\vec{x}_f$ is the vector of features gathered from the image corresponding to the facial image f

Data cost

- The data cost term is essentially the result of a baseline, image-only face recognition system
- We used high-dimensional V1-like features
  - Shown to achieve excellent face-recognition results
- Each feature vector is large – 86400 entries
  - Greedily selected 400 features

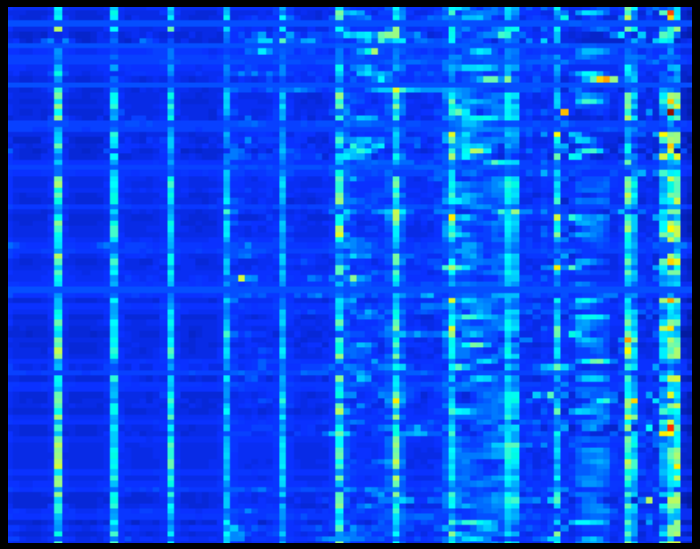

N. Pinto, J. DiCarlo, and D. Cox. How far can you get with a modern face recognition test set using only simple features? *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on,* 0:2591-2598, 2009.

UCF

Label Cost

- Our model builds on this by adding a label cost $C(\vec{y})$ which is a function of label assignment $\vec{y}$

$$E(\vec{y}) = \sum_{f \in F} D_f(y_f; \vec{x}_f) + C(\vec{y})$$

Album labeling

UCF

Overview of Label Cost

- Personal label cost: Every label should pay a cost to enter into the album. This cost restricts the number of individuals in an album

- The social label cost: Compatibility of different labels in the album based on facebook meta information

UCF

Personal Label Cost

- Expresses the idea that a limited number of people should appear in an album
  - Every label that enters the album must pay its cost

$$C_{personal}(\vec{y}) = \sum_{l \in \mathcal{L}} \lambda \mathrm{I}(l, \vec{y})$$

All possible labels

$$\mathrm{I}(l, \vec{y}) = \begin{cases} 1 & \text{if } l \in \vec{y} \\ 0 & \text{otherwise} \end{cases}$$

Social Label Cost

- The social label cost is constructed using various social metadata available via Facebook
  - Consider the interaction between a label and all other labels in the current album
    - Friendship Cost
    - Co-occurrence Cost
    - Uploader Cost

$$C_{social}(\vec{y}) = \sum_{l \in \mathcal{L}} S(l, \vec{y}) \mathrm{I}(l, \vec{y})$$

$$S(l, \vec{y}) = \sum_{j \in \mathcal{L}} (\alpha_f C_f(l, j) + \alpha_c C_{co}(l, j)) \mathrm{I}(l, \vec{y}) + \alpha_u C_u(l)$$

UCF

Social Label Cost

- **Friendship cost** – measures whether individuals co-occurring in the album are facebook friends

$$C_f(i,j) = \begin{cases} 0 & i \text{ and } j \text{ are friends} \\ 1 & \text{otherwise} \end{cases}$$

UCF

Social Label Cost

- **Co-Occurrence cost** – measures whether individuals in the album ever co-occurred in training albums

$$C_{co}(i,j) = \begin{cases} 0 & i \text{ and } j \text{ have co-occurred} \\ 1 & \text{otherwise} \end{cases}$$

Social Label Cost

- **Uploader cost** – indicates whether a user has previously appeared in photographs uploaded by a given uploader

$$C_u(i) = \begin{cases} 0 & \text{if } i \text{ has appeared in images} \\ & \text{uploaded by the owner of } F \\ 1 & \text{otherwise} \end{cases}$$

Inference

- Inference in the overall energy function is NP-hard
  - Our model is similar to the *uncapacitated facility location problem*, which allows for a greedy approximation
    - Add one label to the album that maximizes the energy function at each iteration
    - Stop when adding new candidate labels does not result in further improvement

A. A. Kuehn and M. J. Hamburger. A heuristic program for locating warehouses. *Management Science*, 9(4):pp. 643-666, 1963.

UCF

Learning

- We use the Structural SVM (SSVM) – it can optimize parameter values even if inference can only return approximate solutions
- Need to learn weights for the personal label cost and the individual social costs

$$[\lambda, C_f, C_{co}, C_u]$$

Datasets

- Two datasets gathered from Facebook accounts of volunteers
  - First dataset: gathered from all albums visible to one user
  - Second dataset: gathered from a larger set of volunteers
- Accessed all photos in all albums available, and store relevant social network information

Datasets

- We collected photographs of hundreds of individuals, many of whom had very few photographs
- Pruned dataset to include only individuals with a large number of photos
  - First dataset: 1951 images of 25 people
  - Second dataset: 1994 images of 15 people

UCF

Datasets – Privacy

- This allows us to validate our model on real-world datasets
- Difficult to share data due to privacy concerns, potential for de-anonymization
- We will email the code if anyone asks

Partitioning

- Partitioned each dataset into three parts:
  1. A set used to train the weights for the data cost (image-only training)
  2. A set used to train the weights for the personal label and social costs
  3. A testing set (454 and 487 images for the two datasets)
- Albums were partitioned in the order of time

Experimental Results

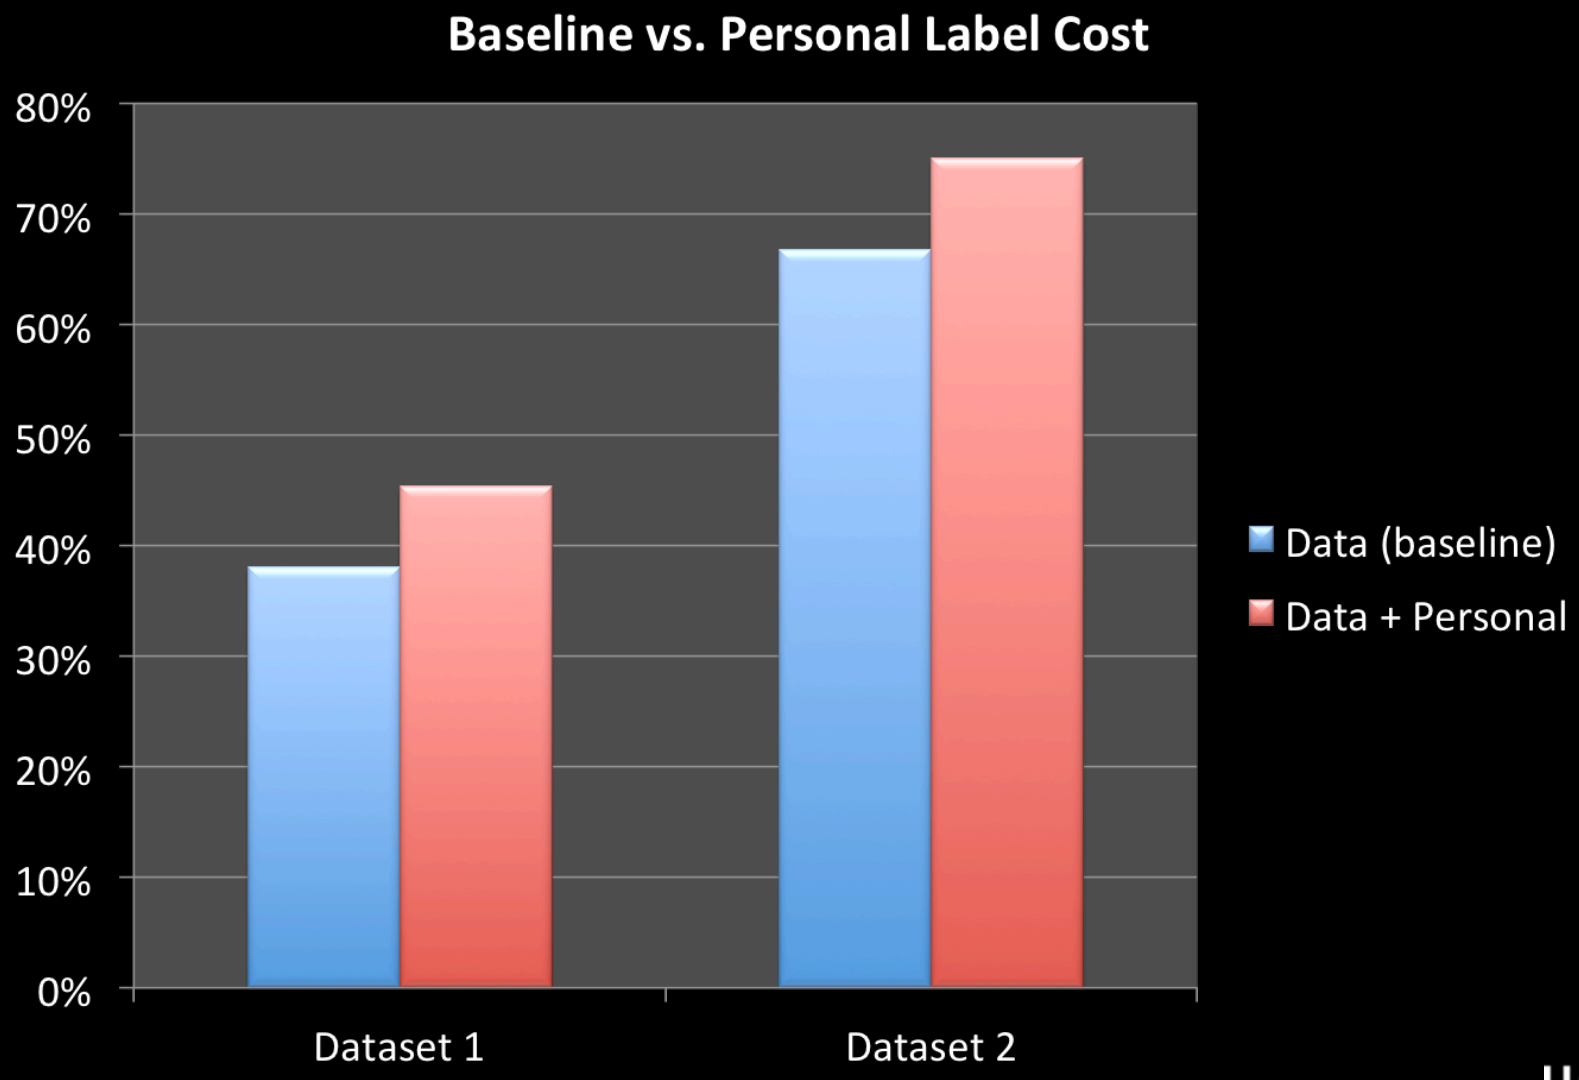

UCF

Experimental Results

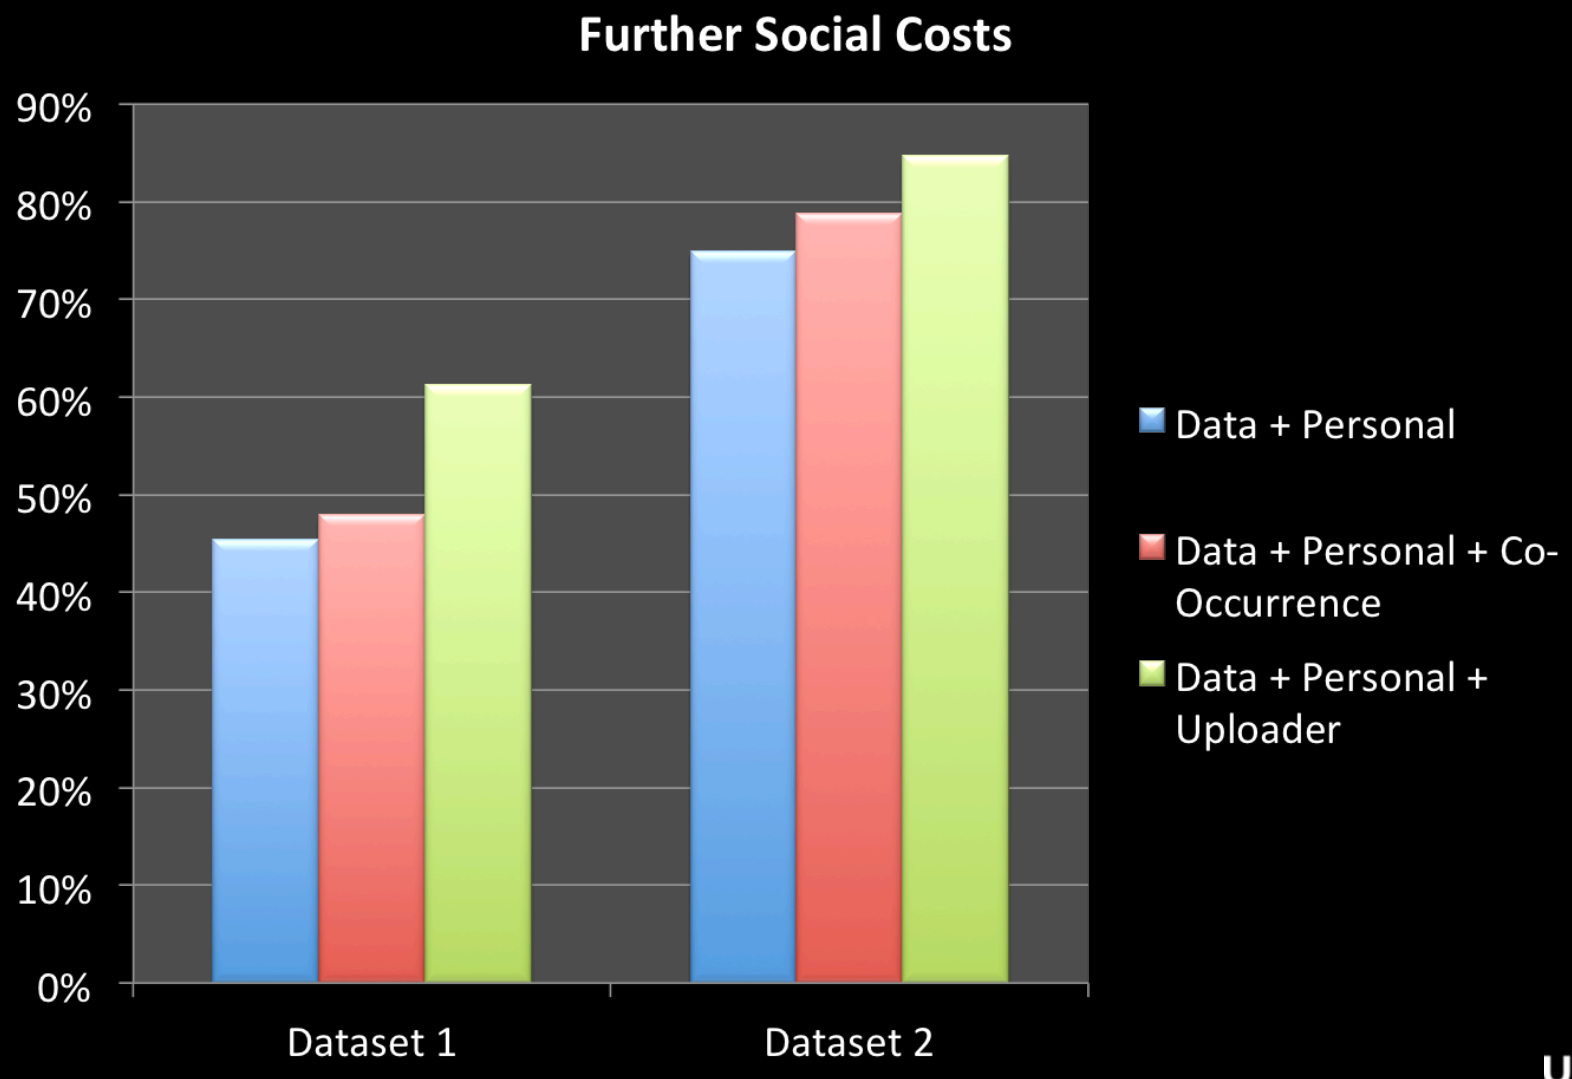

UCF

Experimental Results

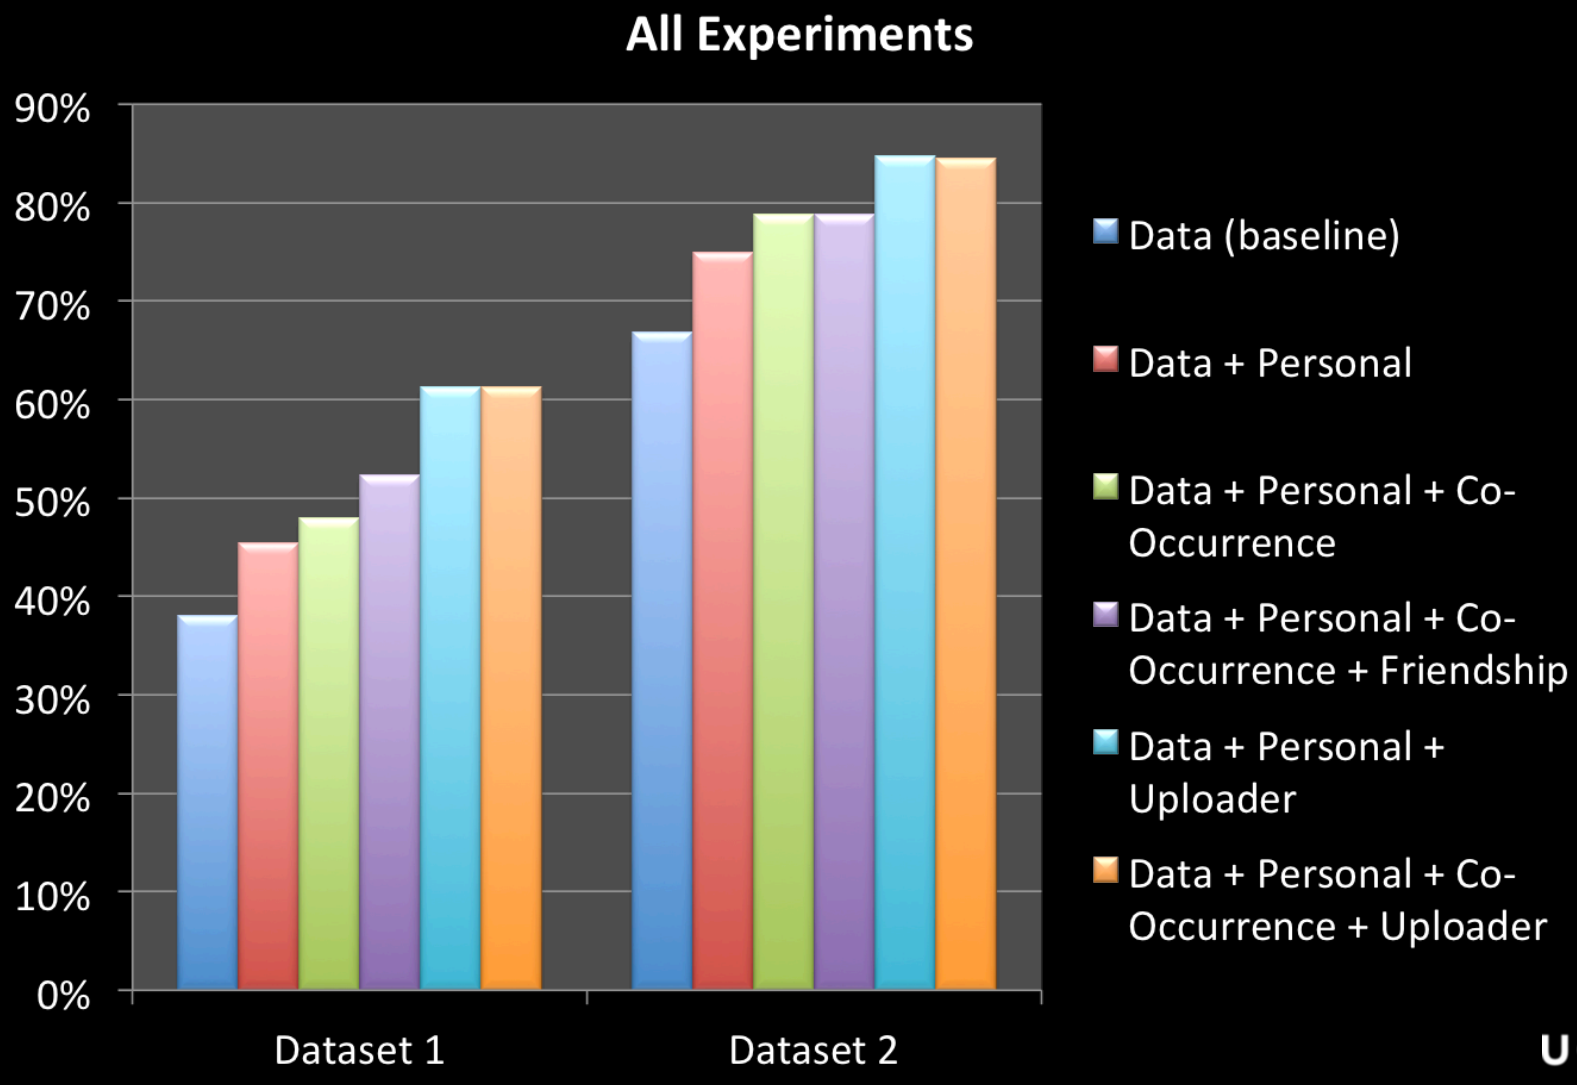

UCF

Conclusion

- We propose using album-level co-occurrence information to improve face recognition
- We build a mathematical model that incorporates traditional image features and label information
- Our experiments show that our method significantly improves the face recognition performance

UCF

Acknowledgement

- This work was funded by NSF grants IIS-0905387 and IIS-0916868.
- We'd like acknowledge our funders of this work

UCF

Thank you

UCF