You are what you tweet…pic! Gender prediction based on semantic analysis of social media images

Michele Merler, Liangliang Cao and John R. Smith

IBM T.J. Watson Research Center, USA

mimerler@us.ibm.com

Outline

- Motivation & Research Questions
- Proposed Approach
  - Visual and Textual Analytics
  - Fusion Strategies
  - Experimental Results
- Conclusions & Future Directions

Social Media is a Goldmine of for Multimedia Research

Event Discovery and Summarization

Attributes Discovery

Gender, age, location, education, political preferences, job, etc.

Training Data for Visual Classification

Sentiment Analysis

Previous Work

- Text-based Gender estimation from social media , *language dependent, performance ceiling*
  - Tweets [Burger et al. EMNLP11] [Pennacchiotti et al. ICWSM11]
  - First name [Liu and Ruths AAAI13]
  - Hashtags [Totems14]
  - Psycho-linguistic features [System U] [Kokkos et al. FM14]
  - Topic modeling on boards [Chang et al. CSCW14]
- Non-text based Gender estimation from social media
  - Collaborative Filtering (who you are friends with, who you follow) [Ito et al. ASONAM13] [Ludu CORR14] *limited performance*
  - Profile Picture face analysis, *not always available/reliable*
  - Page Colors [Alowibdil et al. CASNAM13], *limited performance, not always available*
  - Whole Feed Images [Ma et al. IWCMASM14], *small set of ad hoc classifiers, no use of profile pictures, extremely limited generalization power*
- Combinations Gender estimation from social media
  - Text + Images [Sakaki et al. ICCL14]
    small preliminary study, over-simplicistic fusion method: $\alpha p(text) + (1- \alpha)p(visual)$ *, limited performance*
- Multimodal Fusion (not for gender prediction)
  - Extensive literature, early fusion, late fusion, *general fusion strategies vs proposed specific filtered fusion (see experimental results)*

A need remains for a system that derives user gender using an effective multimodal combination of visual and non-visual cues

Research Questions

- Is there a ***correlation*** between ***gender*** and the content of the ***images*** that people post on social media?

  yes

- If so, can we ***predict*** a social media user's gender based on a semantic analysis of those images?

  yes

- Does the visual insight provide ***complementary*** information with respect to others (text)?

  yes

Summary of Invention –Extracted Information

Multimodal Cues

- Textual
  - Name
  - Description
  - Tweets (text)
- Visual
  - Profile picture
  - Header picture
  - Profile colors
  - Feed images/videos

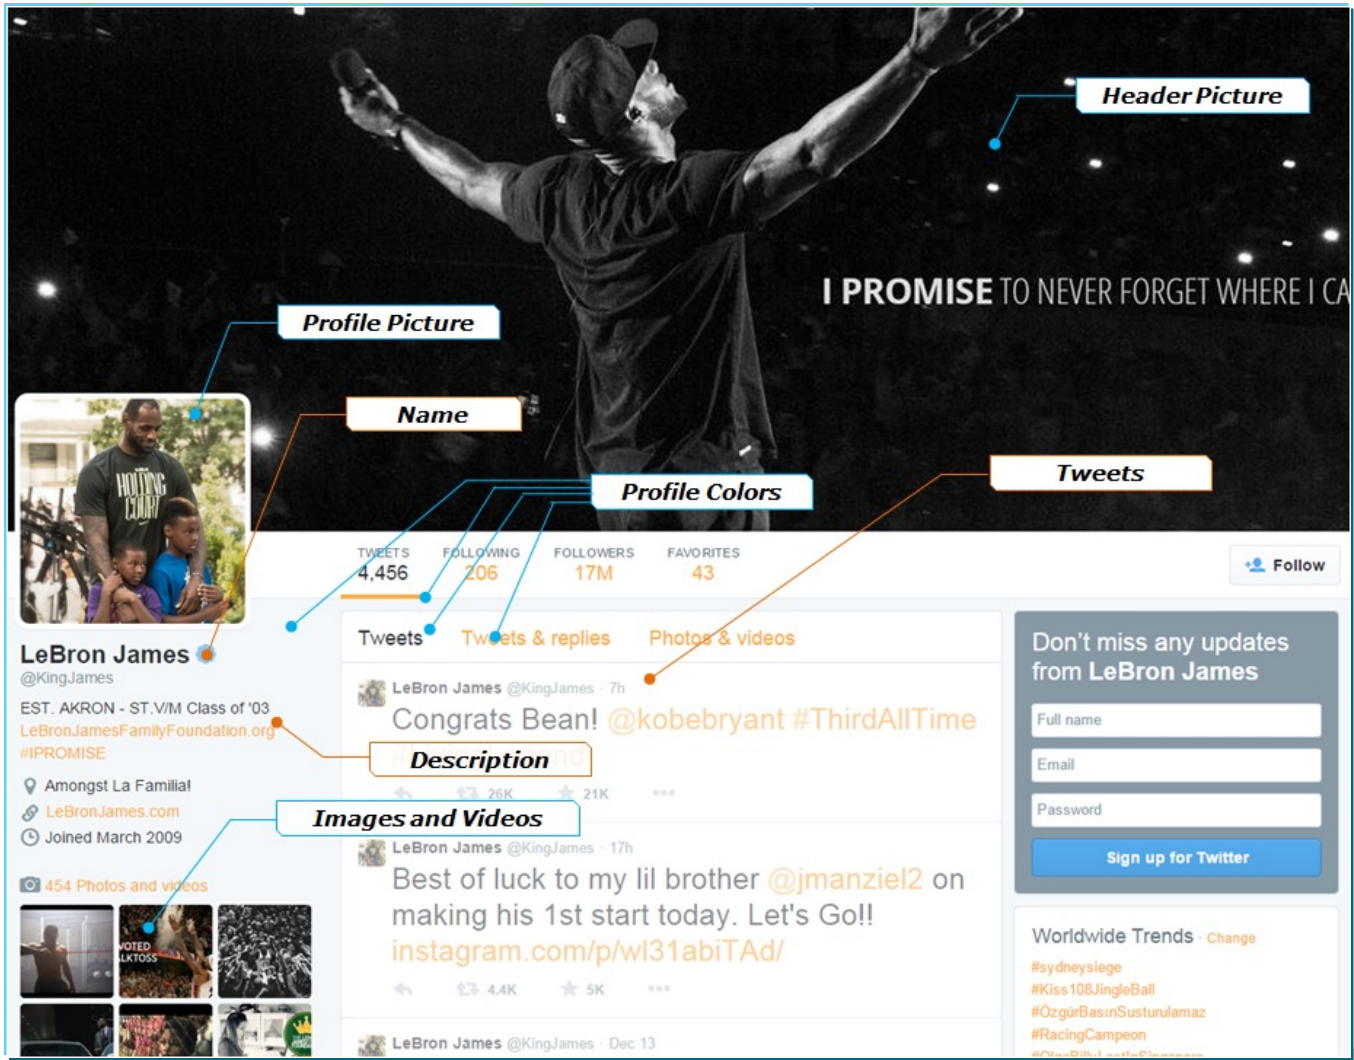

Summary of Invention

- **V1 Classifier** Based on generic visual classifiers
- **V2 Classifier** Based on face gender analysis
- **V3 Classifier** Based on low level descriptors
- **T1 Classifier** Based on name/gender frequency analysis
- **T2 Classifier** Based on standard text BOW, n-gram classifier

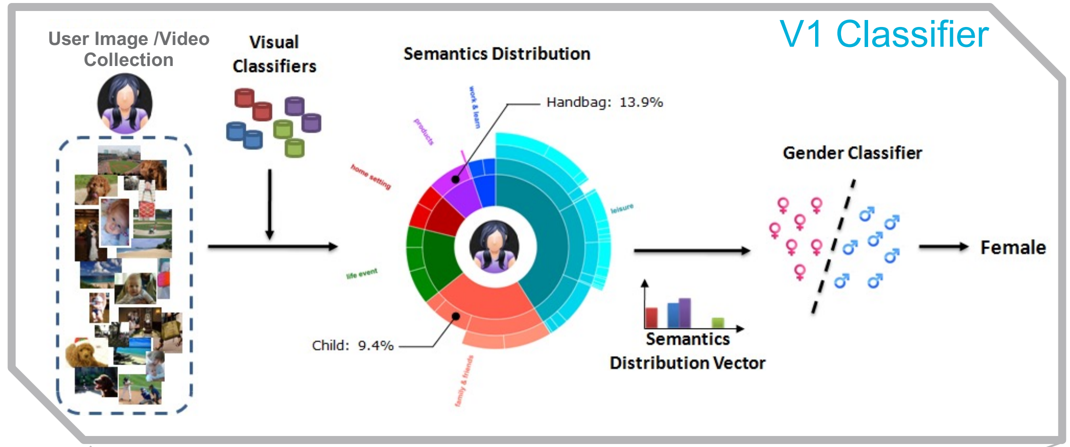

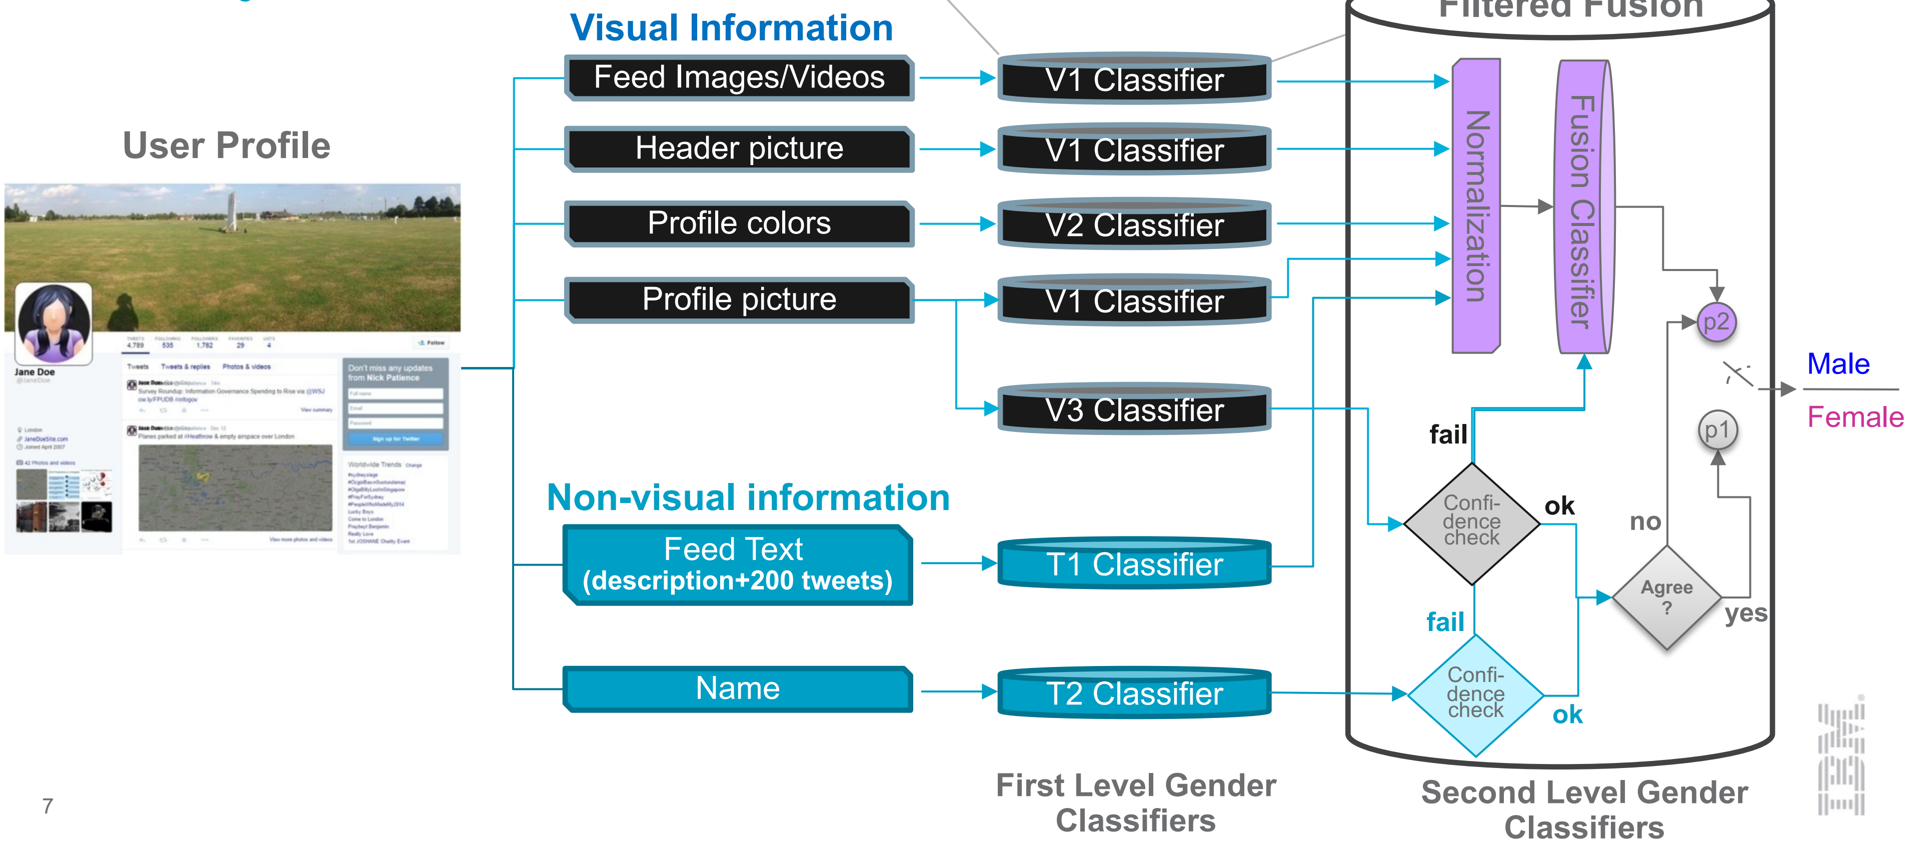

Textual Analytics

- Profile Name
- Text from tweets

Visual Analytics

- Profile Picture
- Color
- Analysis of Collection of Posted Images

What’s in a Twitter Profile Picture?

What’s in a Twitter Profile Picture?

the good (looking)

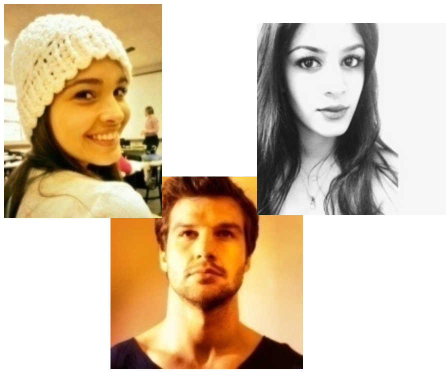

the bad

the cute

…and the weird?

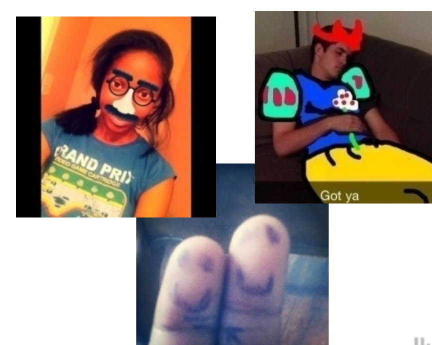

Limitations of Profile Picture Face Analysis

Misleading Clothing

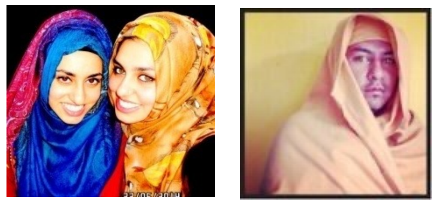

Occlusion

Cloning

Interesting angles

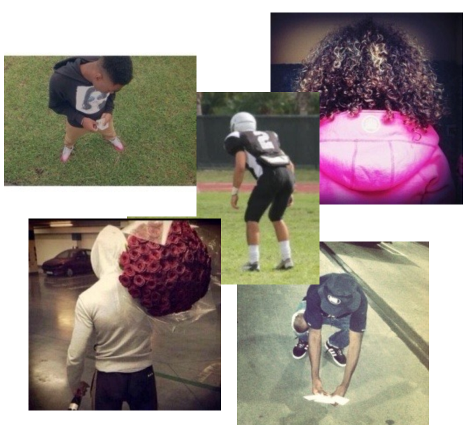

Time travel

celebrity swap

Multiple people

Non-human pictures

Limitations of Profile Picture Face Analysis

Race, glass, smiling

Age, gender

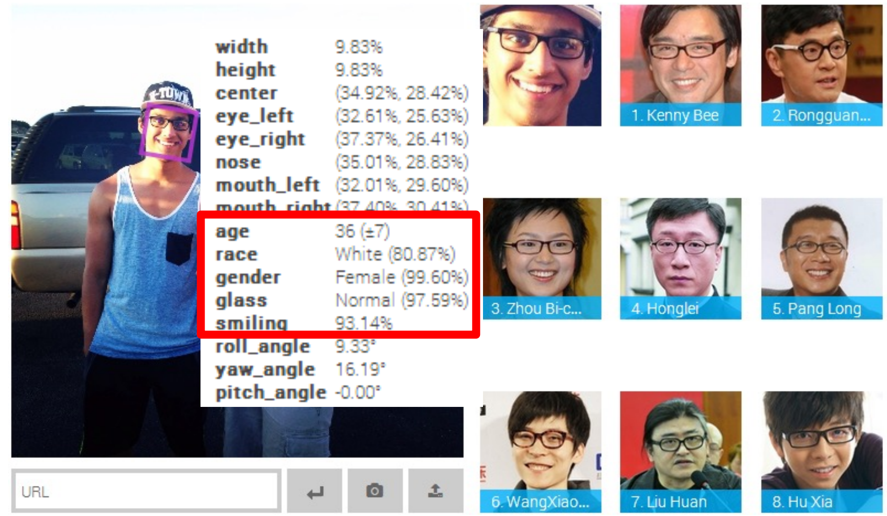

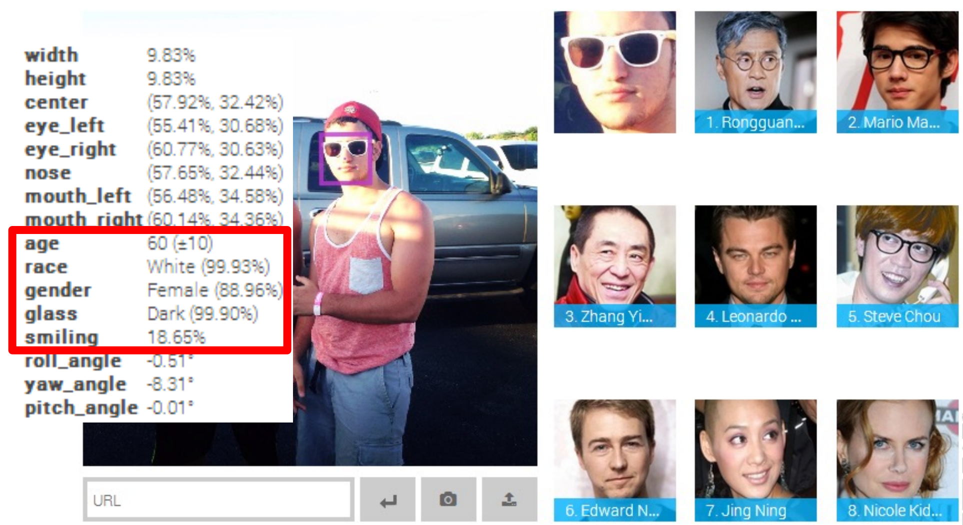

Source: http://www.faceplusplus.com/demo-search/

Profile Picture : proposed approach

- Face++ detector

- Concept Detectors for 25 categories
  - *Adult, Animal, Baby, Beach, Boy, Brand Logo, Building, CGI, Car, Cat, Child, Dog, Elderly Man, Elderly Person, Elderly Woman, Female Adult, Girl, Horse, Human Portrait View, Human, Icon, Male Adult, Motorcycle, Nature, Two People*

  - *Train SVM on top of Semantic Model Vector of concept detectors*

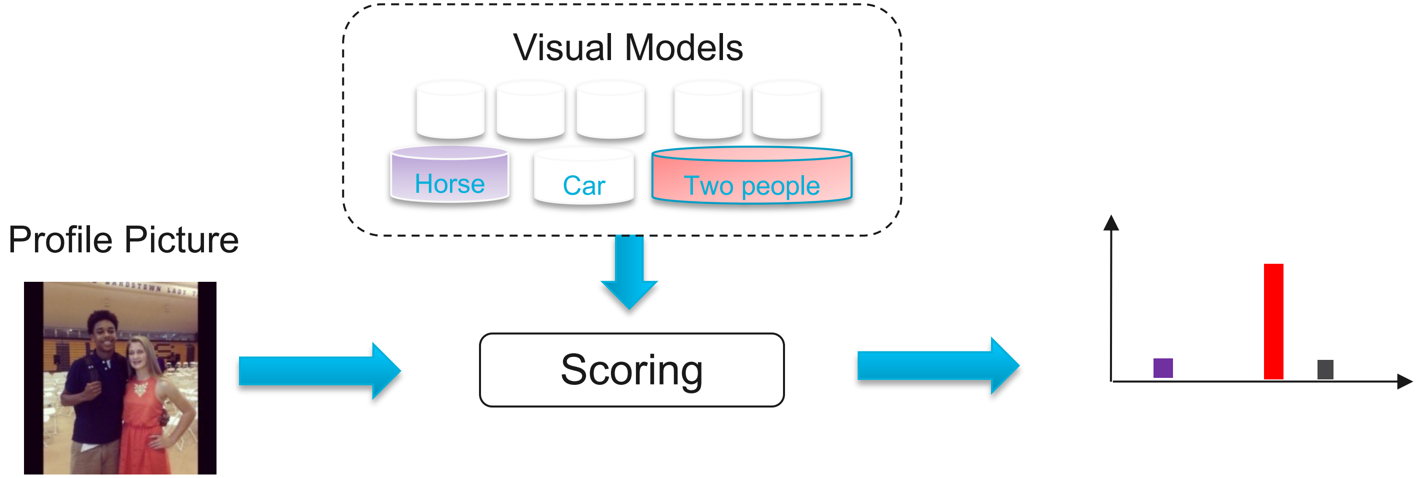

Linear Profile SVM weights by category

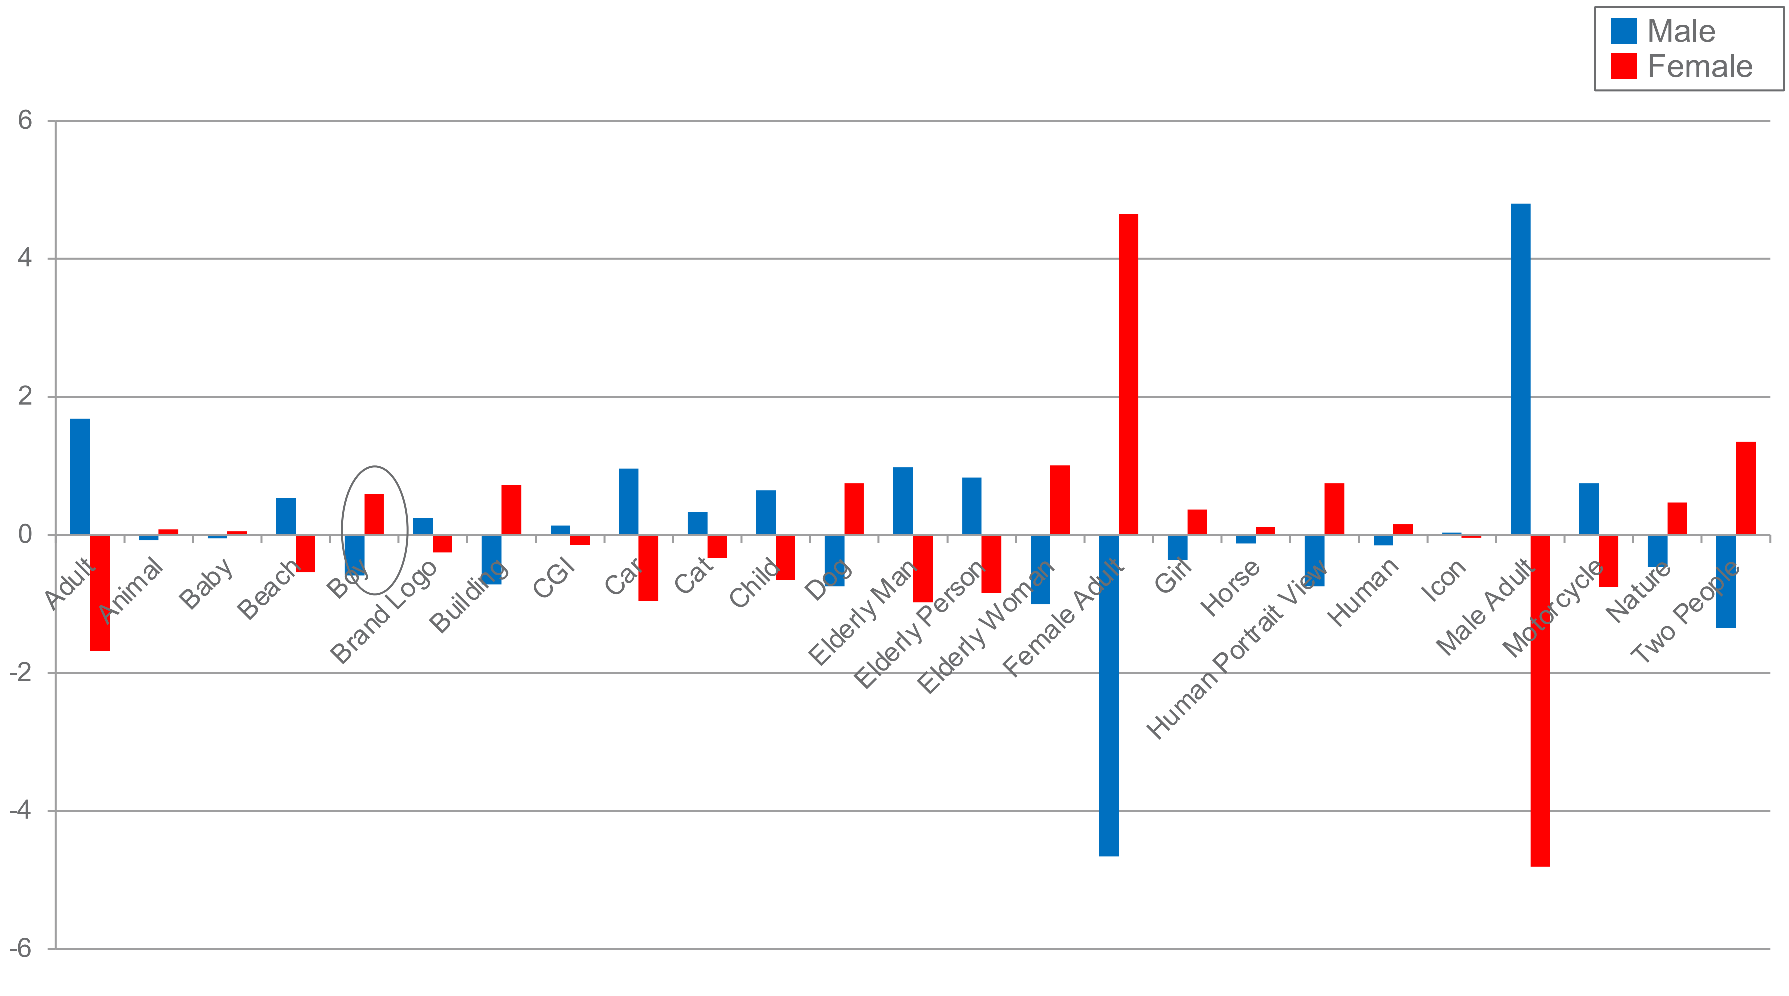

Images From the Entire Feed

- Same Semantic Model Vector Approach
  - SMV 51 (subset of IMARS Taxonomy)
  - SMV 717 (subset of IMARS Taxonomy)
  - SMV Deep (from Caffe, 1K ImageNet categories)

- Aggregation Strategies
  - Model on Images directly
  - Simple Prediction Scores Aggregations (avg, max)
  - Statistical Count scores (threshold + count)

Twitter Gender Dataset Examples

Male

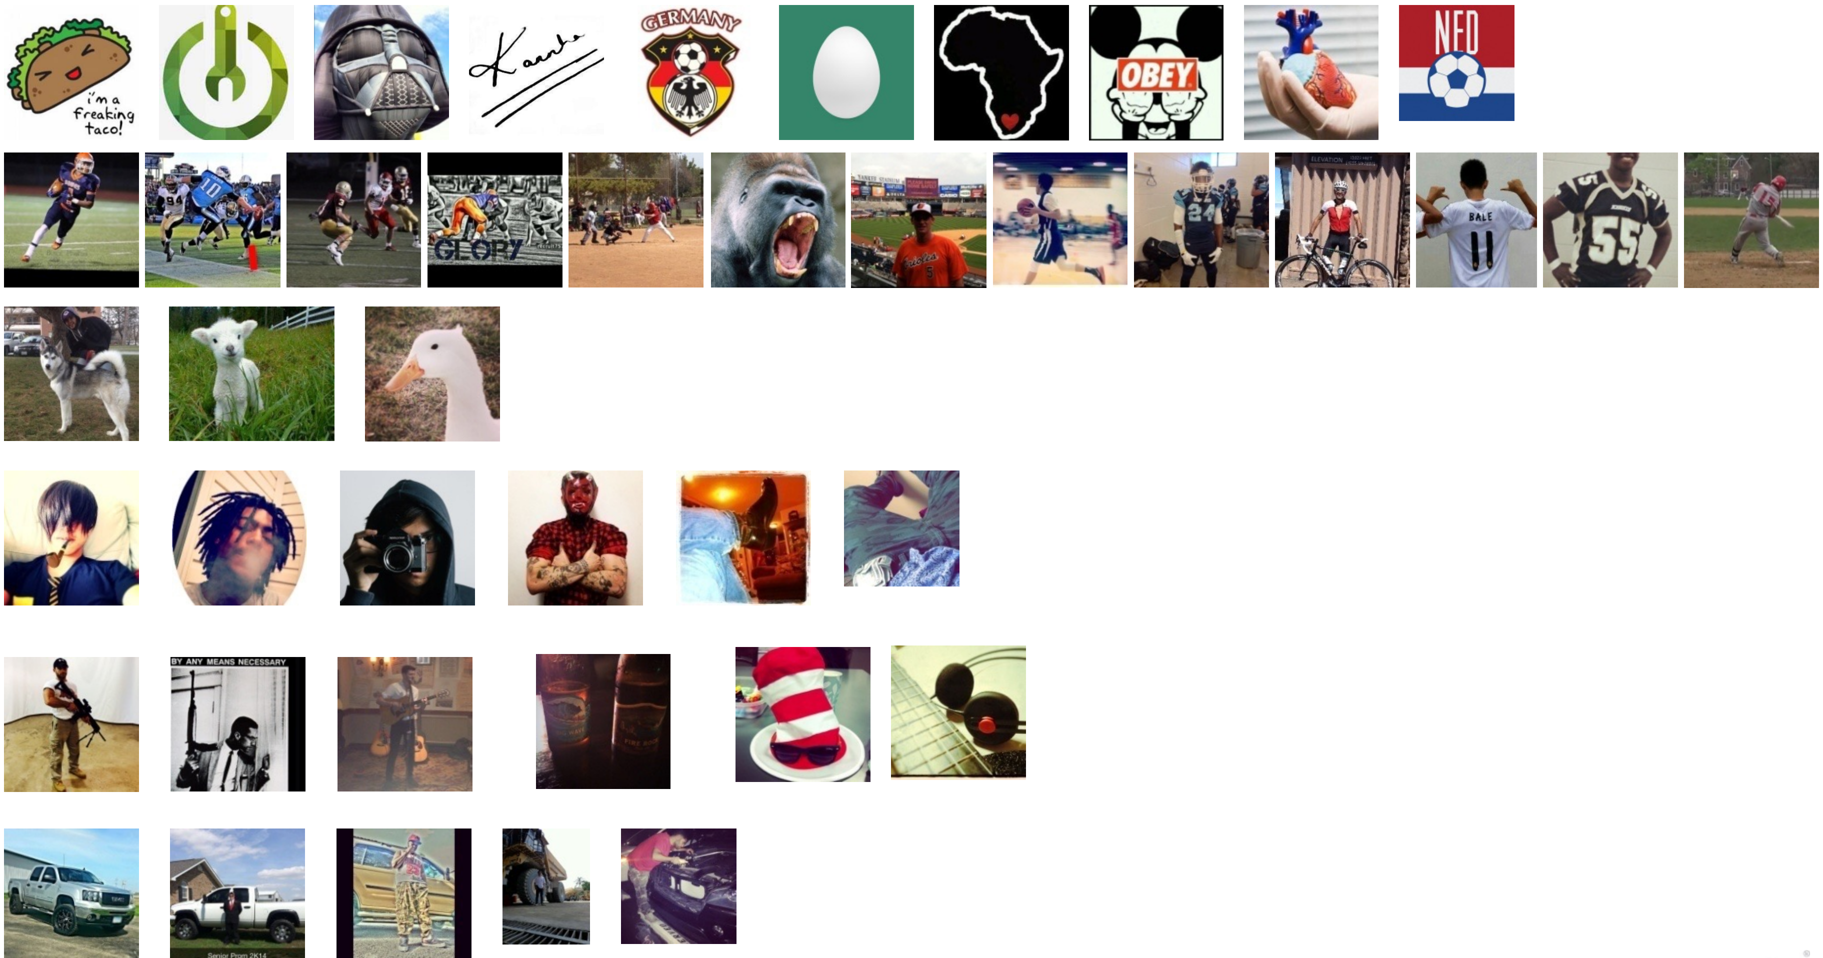

Twitter Gender Dataset Examples

Female

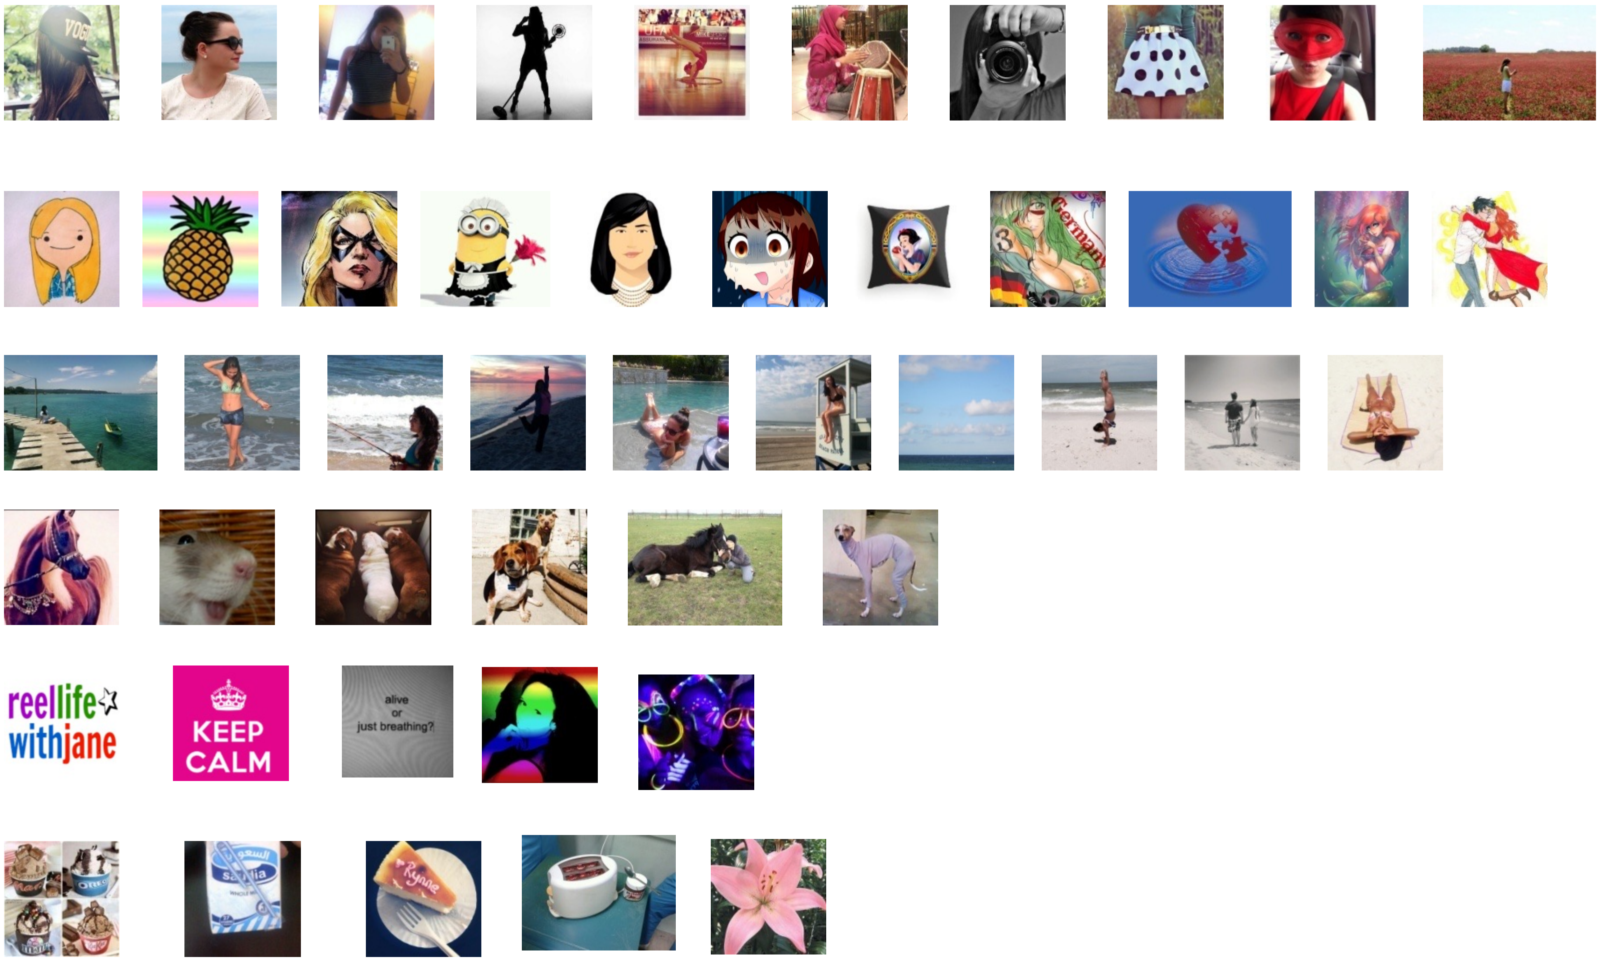

Linear SVM weights by category (SMV 51)

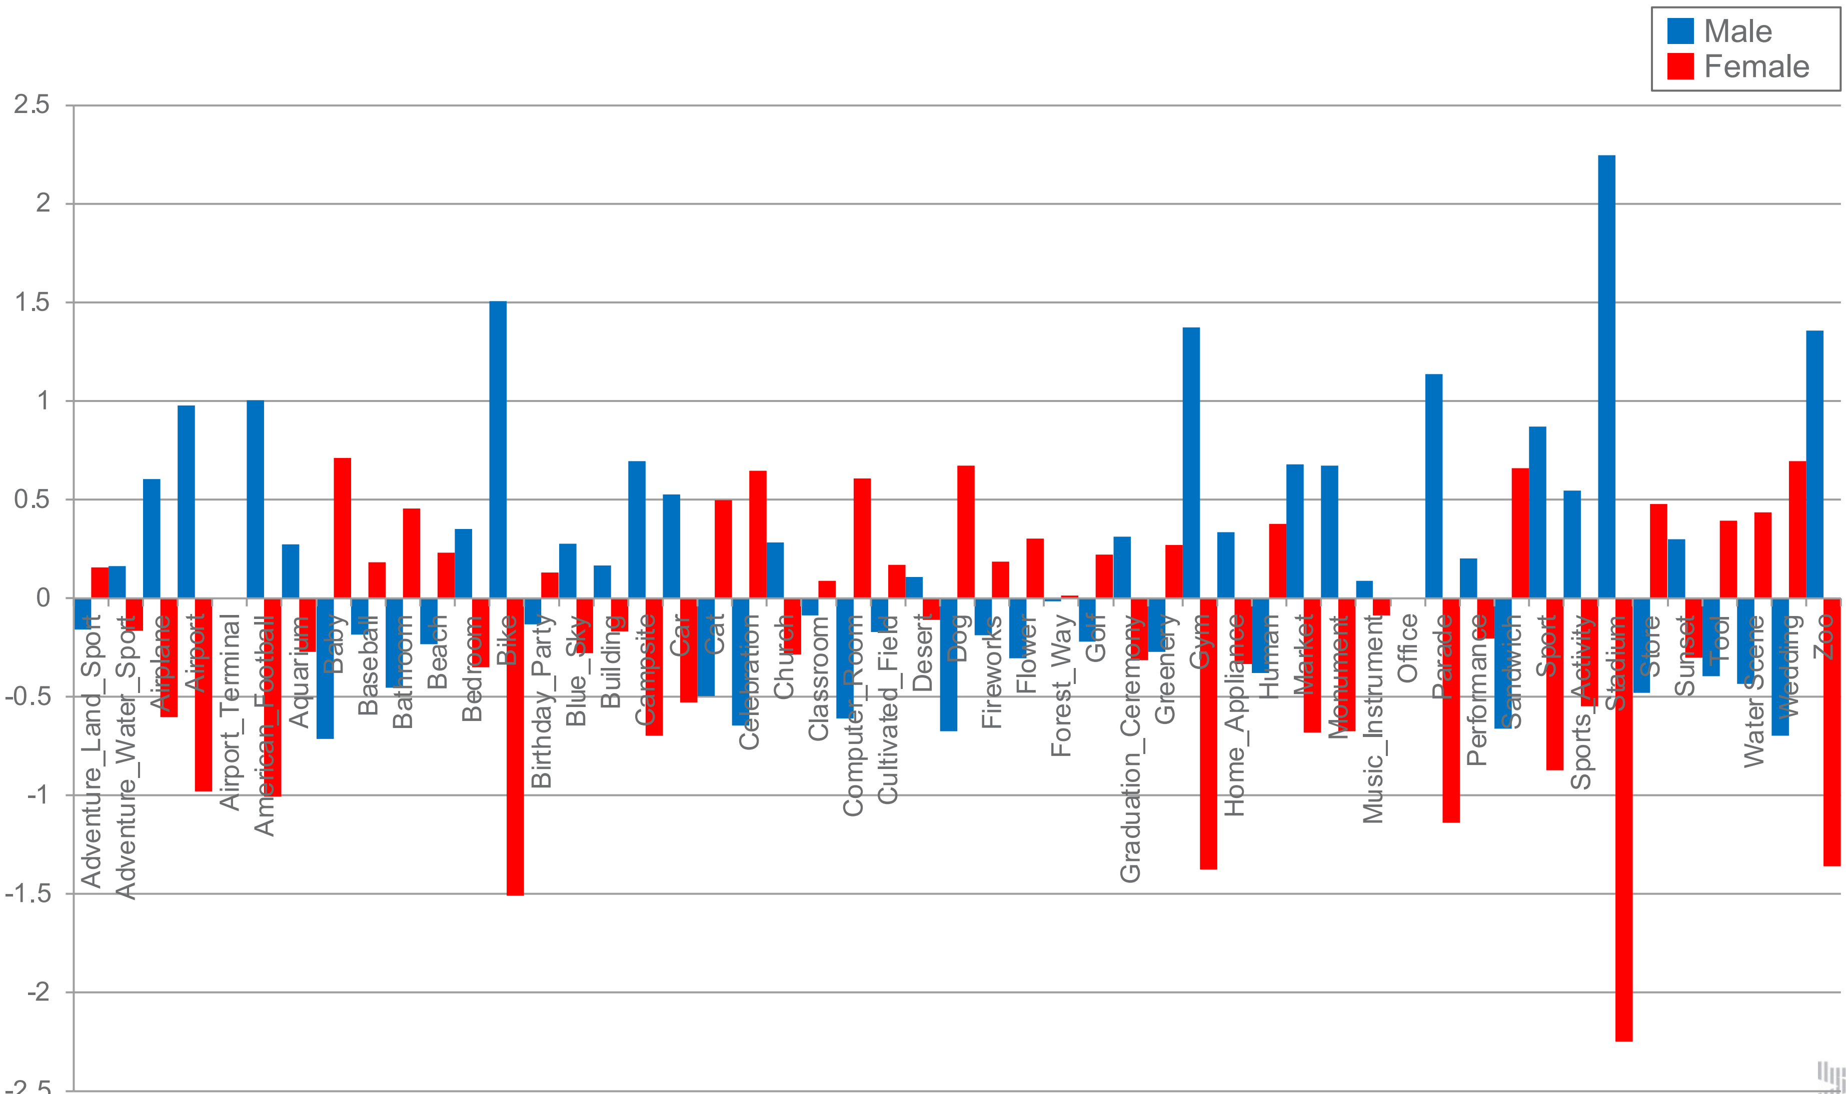

Top 14 weighted categories by gender (SMV51)

Male


Stadium


Bike

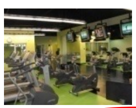
Gym


Zoo


Parade

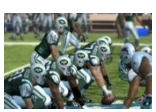
American Football


Airport


Sport


Campsite

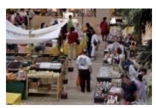
Market

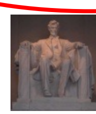
Monument


Airplane


Sports Activity


Car

Female


Baby


Wedding

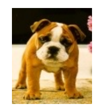
Dog

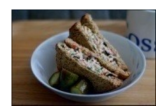
Sandwich/Food


Celebration

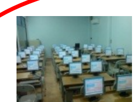
Computer Room

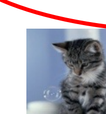
Cat


Store

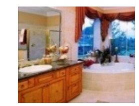
Bathroom


Water Scene


Tool


Human

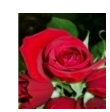
Flower


Greenery

Top 20 weighted categories by gender (Deep ImageNet)

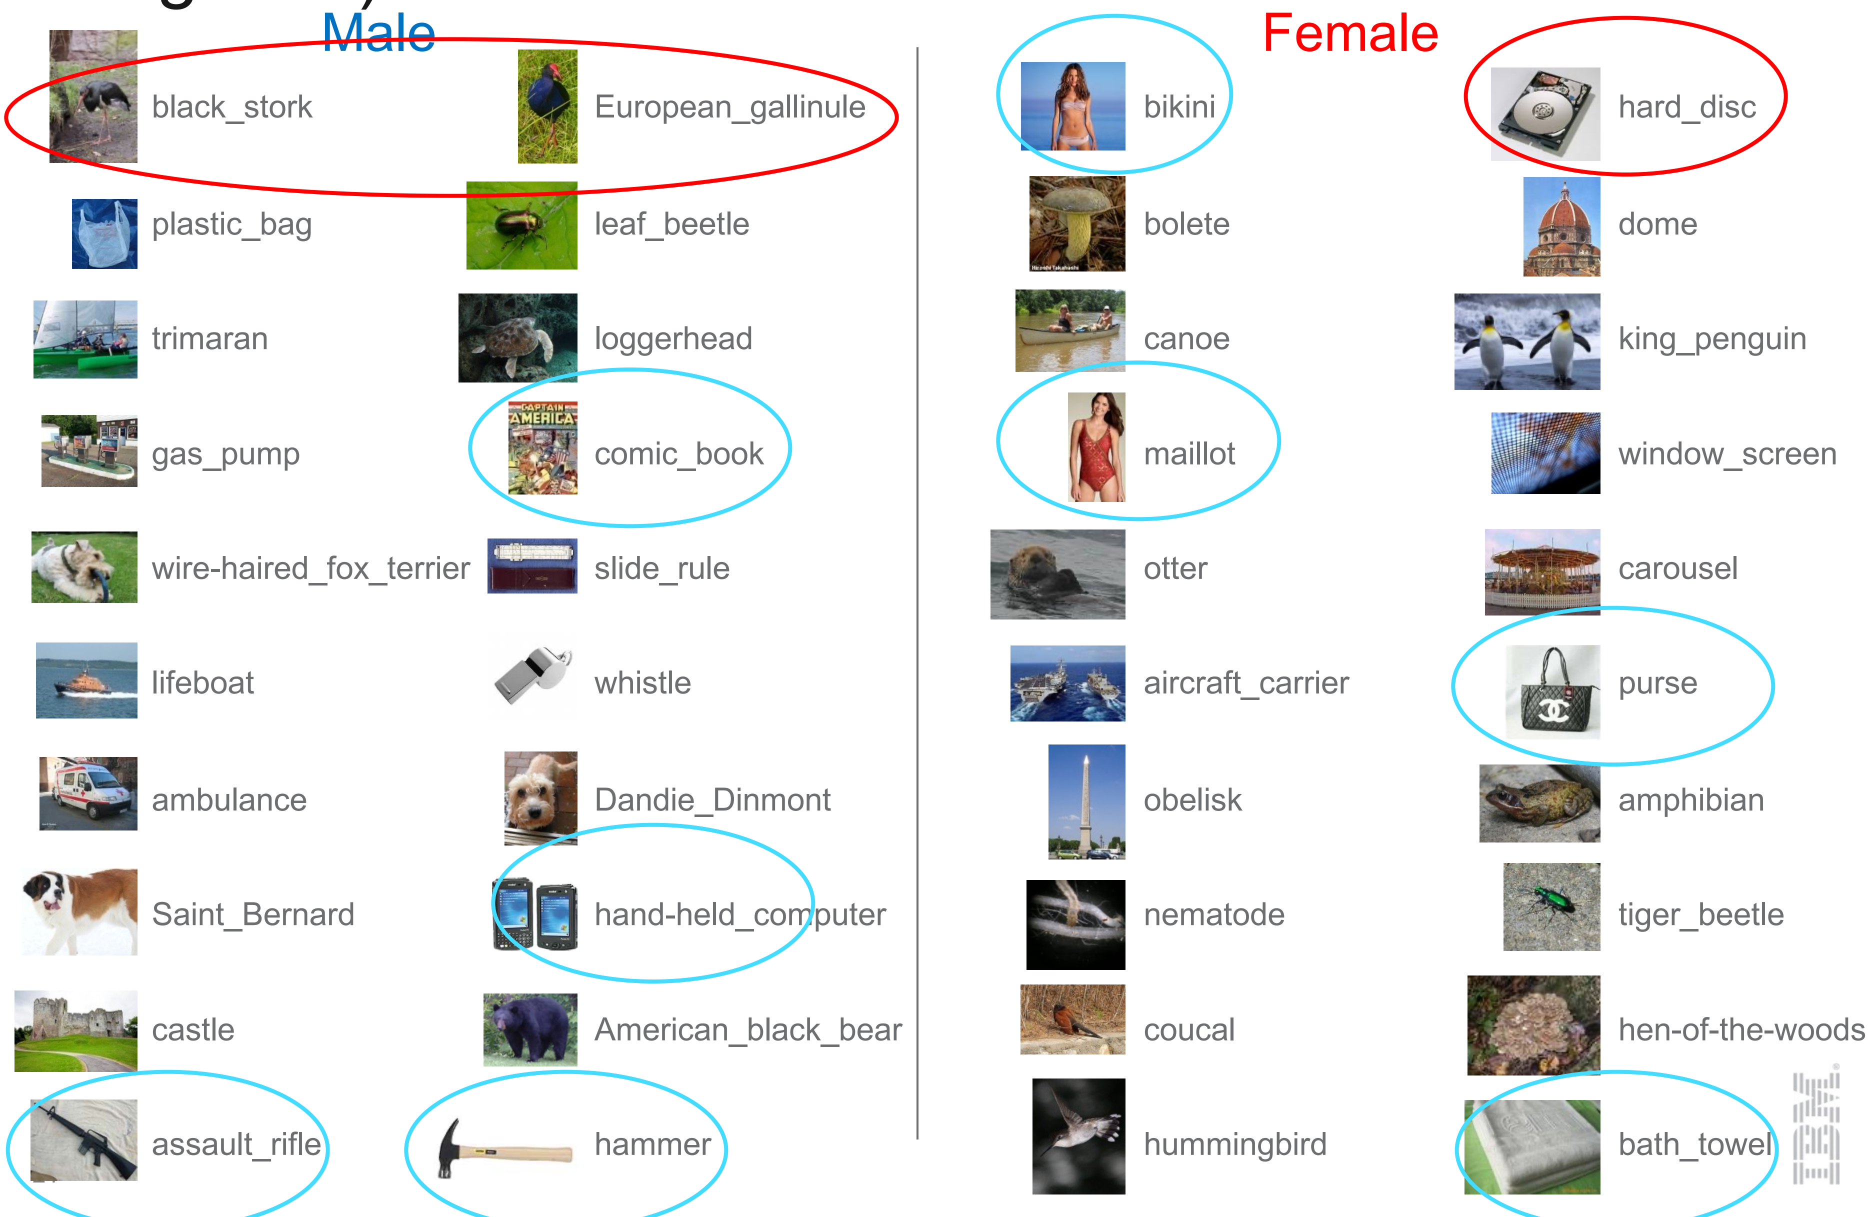

Use Profile Page Color Info

Jalal S. Alowibdi, Ugo A. Buy and Philip Yu, Language Independent Gender Classification on Twitter, 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining

- Background color
- Text color
- Link color
- Sidebar Fill color
- Sidebar Border color

Figure 4. Spectrum of popular colors for female users (left-hand side) and male users (right-hand side).

512 Quantized colors (RGB with 3 bits each)

71% Accuracy when using all 5

http://www.twitteraccountsdetails.com/

Visit Full Profile

Tweet 0

User	Twitter ID	30208909
	Username	____SachaNicole
	Full Name	So
	Description	wassup you fake ass hoe
	Location	❄
	URL	*Not provided by the user*
	Tweets	56,407
	Followers	790
	Following	649
	Favorited	782 tweets
	Listed	in 4 lists
Account	Age	5 years, 6 months and 18 days
	Language	en
	Is Verified?	No
	Is Protected?	No
	Time Zone	Quito
	UTC Offset	UTC-5h
	Geotagging	Enabled
	Is Translator?	No
	Contributors	Disabled
Design	Default Profile Theme?	No
	Default Profile Image?	No
	Profile Image Link	https://pbs.twimg.com/profile_images/...
	Header Image Link	https://pbs.twimg.com/profile_banners...
	Image Background?	Yes
	Background Image Link	https://pbs.twimg.com/profile_backgro...
	Tile Background Image	Disabled
	Background Color	#000000
	Text Color	#E600E6
	Link Color	#FAB812
	Sidebar Fill Color	#FFFFFF
	Sidebar Border Color	#FFFFFF

Data source: twitter.com/____SachaNicole

Use Profile Page Color Info

512 bins

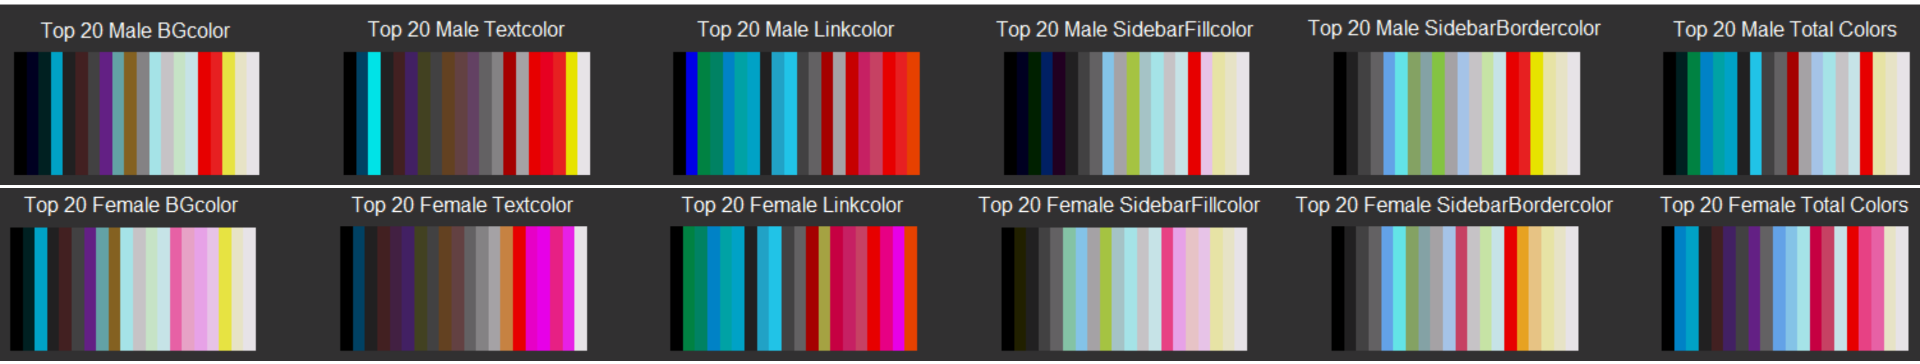

729 bins

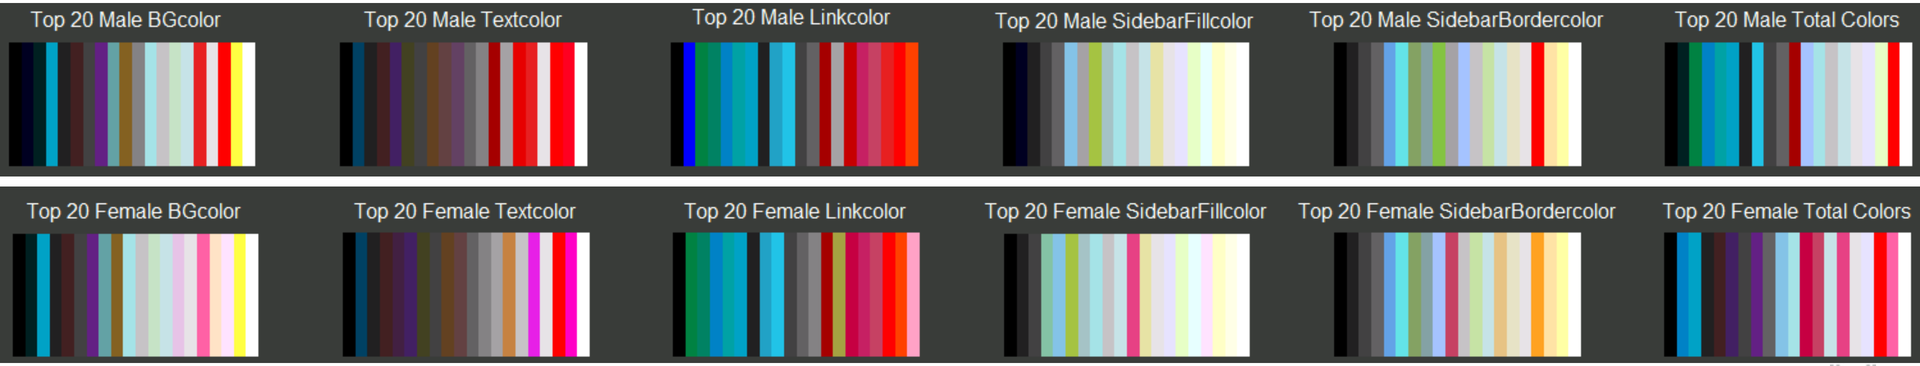

Experimental Setup and Proposed Approach

- Public Annotated Dataset of 10K Twitter users[1]

- 10 Training/Test random splits, each test split with 400 male and 400 female users

Proposed Approach : *Fusion* of 4 Multimodal Cues

  - First Name (when available): associated with frequency in Male/Female populations[2]
  - Text : standard bag of words from ~200 tweets, analyzed with LibShortText[3] library
  - Profile Picture: Semantic Model Vector (25 concepts) and gender inference from Face++
  - *Stream Pictures* : Semantic Model Vector (51, 717, 1K) aggregated over user pictures
  - Page colors
  - Background and Header picture

- Gender modeling was conducted using SVM with RBF kernel, with kernel parameters estimated via grid search

1. http://www.networkdynamics.org/static/datasets/LiuRuthsMicrotext.zip
   From paper *Liu and Ruths, What's in a Name? Using First Names as Features for Gender Inference in Twitter, AAAI 2013*
2. http://www.census.gov/genealogy/www/data/1990surnames/names files.html
3. http://www.csie.ntu.edu.tw/~cjlin/libshorttext/

Experimental Results

- Public Annotated Dataset of 10K Twitter users[1]
- 10 Training/Test random splits, each test split with 400 male and 400 female users, rest used for training
- Gender modeling was conducted using SVM with RBF kernel, kernel parameters estimated via grid search

Image/video collection (visual feed) different pooling strategies →

Method	Accuracy
Max Pooling	67.53
Avg Pooling	69.43
Avg Top-Quarter Pooling	69.56
Threshold-count	70.82
Avg Prediction Pooling	**71.38**

Individual Performance of Different Approaches

	Method	Accuracy
Visual	Background SMV717	60.11
	Header SMV717	64.41
	color	66.18
	Visual Feed SMV51	66.67
	Visual Feed SMV717	71.38
	Visual Feed SVMDeep1000	**75.40**
	Profile SMV25	69.11
	Profile Face++	74.90
Text	First Name	71.22
	First Name Frequency	69.58
	LibText 200 Tweets	**83.37**

Fusion Strategies

Method	Accuracy
Visual Feed Early Fusion	75.58
Visual Feed Late (avg) Fusion	74.34
Visual Feed Late (SVM) Fusion	**75.6**
Profile Late (avg) Fusion	77.85
Profile Late (SVM) Fusion	78.63
Profile Filtered Fusion	**79.05**
All Visual Late(SVM) Fusion	80.08
All Visual Late(SVM) Filtered Fusion	**83.36**
Textual Early Fusion	84.08
Textual Feed Late (avg) Fusion	84.53
Textual Late (SVM) Fusion	84.67
Textual Filtered Fusion	**85.72**
Visual+Text Early Fusion	84.07
Visual+Text Late (SVM) Fusion	85.97
Visual+Text Late (SVM) Filtered Fusion	**88.01**

1. http://www.networkdynamics.org/static/datasets/LiuRuthsMicrotext.zip
From paper *Liu and Ruths, What's in a Name? Using First Names as Features for Gender Inference in Twitter, AAAI 2013*

Experimental Results

- Public Annotated Dataset of 10K Twitter users[1]
- 10 Training/Test random splits, each test split with 400 male and 400 female users, rest used for training
- Gender modeling was conducted using SVM with RBF kernel, kernel parameters estimated via grid search

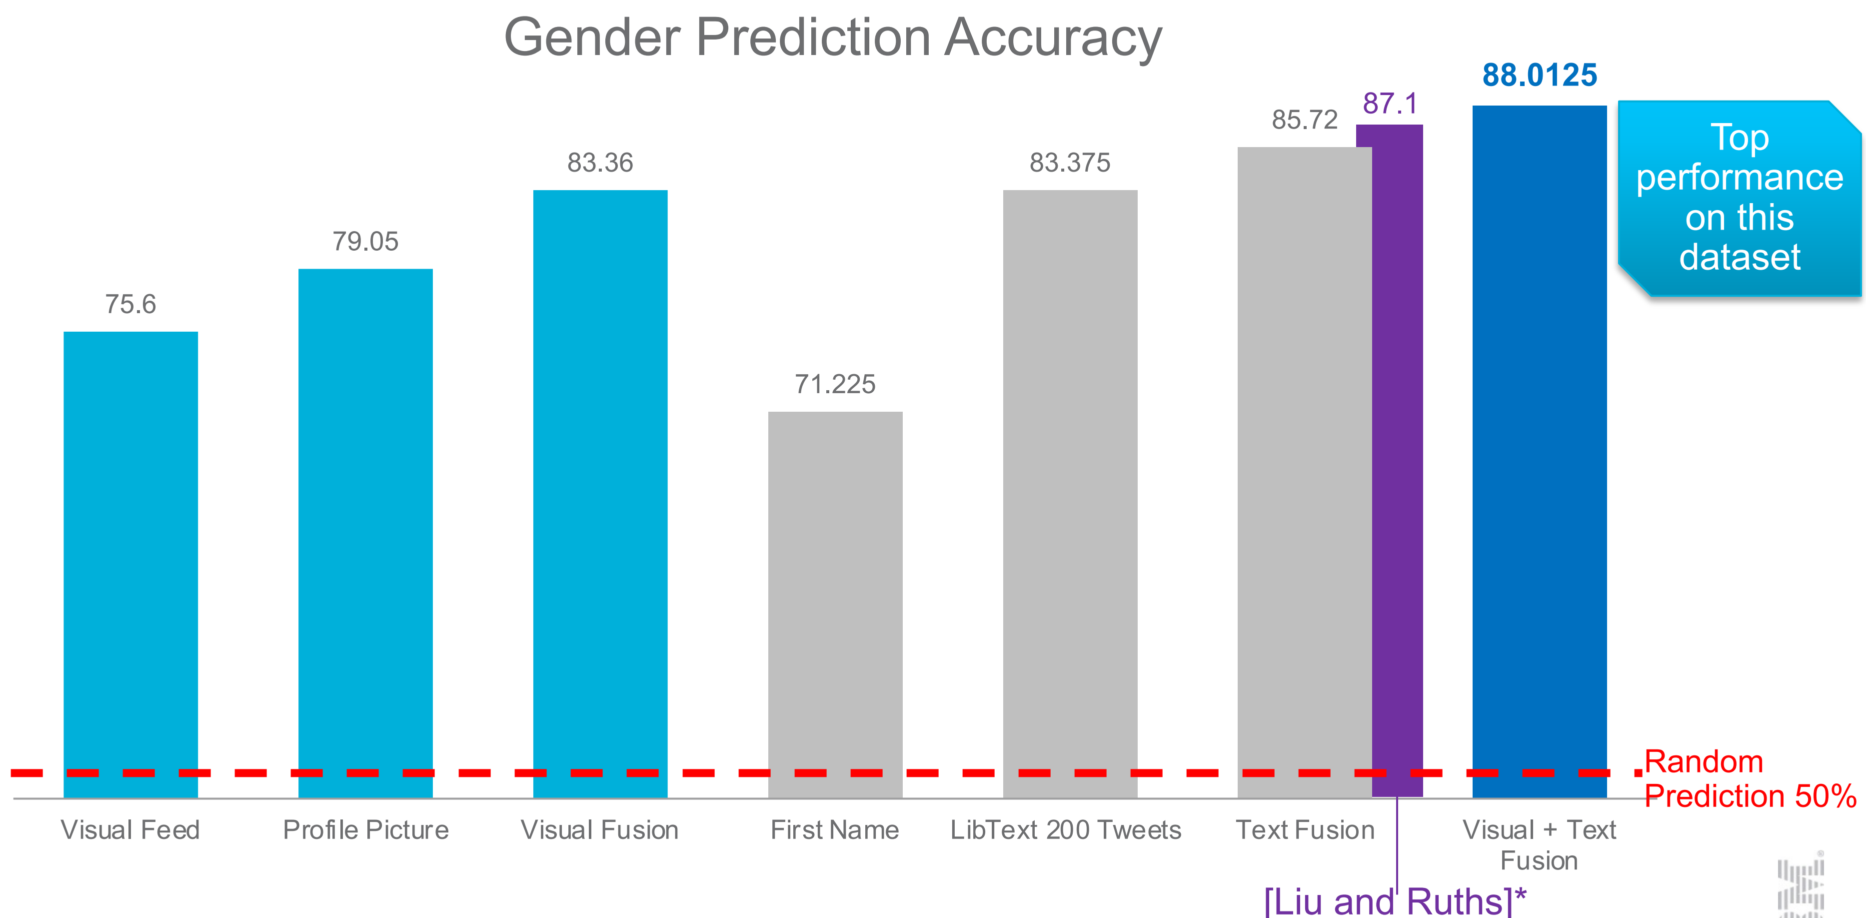

*Different splits

1. http://www.networkdynamics.org/static/datasets/LiuRuthsMicrotext.zip
From paper *Liu and Ruths, What's in a Name? Using First Names as Features for Gender Inference in Twitter, AAAI 2013*

Conclusions and Future Directions

Conclusions

- There is a correlation between the content of images posted on social media and the users' gender, which can be exploited for gender prediction
- Visual and textual information can be combined to boost gender prediction performance
- Not all sources of information are equal. Filtered fusion provides best results.

Future Directions

- Improve/tailor visual classifiers
- Cross media comparisons (Twitter vs Instagram vs Facebook….)
- Use framework to predict other attributes (age, hobbies, life events…)

Questions?

#UnaGeek
@UnaGeek

Follow

Ah, so that's the age of Ultron.

References

- [BurgerEMNLP11] J. D. Burger, J. Henderson, G. Kim, and G. Zarrella. *Discriminating gender on Twitter.* In Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP 2011
- [PennacchiottiICWSM 11] M. Pennacchiotti and A.-M. Popescu. *A machine learning approach to twitter user classication*. In ICWSM. The AAAI Press, 2011
- [LiuAAAI13] Liu and Ruths*. What's in a Name? Using First Names as Features for Gender Inference in Twitter.* AAAI 2013
- [AlowibdiICASNAM13] Jalal S. Alowibdi, Ugo A. Buy and Philip Yu. *Language Independent Gender Classification on Twitter.* 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining
- [ChangCSCW14] S. Chang, V. Kumar, E. Gilbert, and L. Terveen. *Specialization, homophily, and gender in a social curation site: Findings from pinterest.* In CSCW, 2014
- [ItoASONAM13] J. Ito, T. Hoshide, H. Toda, T. Uchiyama, and K. Nishida. *What is he/she like?: Estimating twitter user attributes from contents and social neighbors*. In Advances in Social Networks Analysis and Mining (ASONAM), 2013
- [LuduCORR14] P. S. Ludu. *Inferring gender of a twitter user using celebrities it follows*. CoRR, 2014.
- [KokkosFM14] A. Kokkos and T. Tzouramanis. *A robust gender inference model for online social networks and its application to linkedin and twitter*. First Monday, 19(9), 2014
- [NguyenCOLING14] D. Nguyen, D. Trieschnigg, A. Dogru□oz, R. Gravel, M. Theune, T. Meder, and F. de Jong. *Why gender and age prediction from tweets is hard: Lessons from a crowdsourcing experiment*. In Proceedings of COLING, 2014.
- [MaIWCMASM14] Xiaojun Ma, Yukihiro Tsuboshita, Noriji Kato. *Gender Estimation for SNS User Profiling Automatic Image Annotation.* 1st International Workshop on Cross-media Analysis for Social Multimedia, 2014
- [SakakiICCL14] S. Sakaki, Y. Miura, X. Ma, K. Hattori, and T. Ohkuma. *Twitter user gender inference using combined analysis of text and image processing*. In International Conference on Computational Linguistics, 2014.
- [Totems14] http://totems.co/blog/machine-learning-nodejs-gender-instagram/
- [FarseevICMR15] Aleksandr Farseev, Liqiang Nie, Mohammad Akbari and Tat-Seng Chua, *Harvesting Multiple Sources for User Profile Learning: a Big Data Study.* In ICMR, 2015