

Large-Scale Concepts and Classifiers for Describing Visual Sentiment in Social Multimedia

Shih-Fu Chang

Columbia University

<http://www.ee.columbia.edu/dvmm>

June 2013



Damian Borth



Tao Chen



Rongrong Ji



Yan-Ying Chen

From Photo Shoebox to Social Multimedia



- Images
 - 300 million photos uploaded to Facebook every day.
- Videos
 - 4 billion videos watched per month on YouTube.
- Social media
 - 30 billion content shared on Facebook per month.

Social Sharing -> Opinion Expression

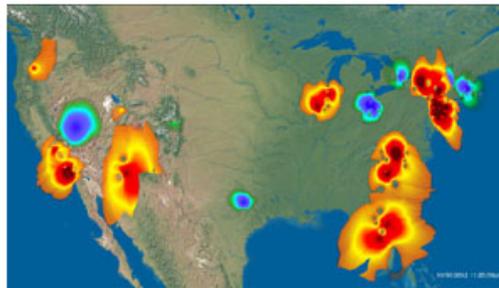


Tracking the Sentiment on Social Media

Twitter Heartbeat (SGI/Uni. Illinois)

Heat Maps of Sentiment on Twitter

Hurricane Sandy

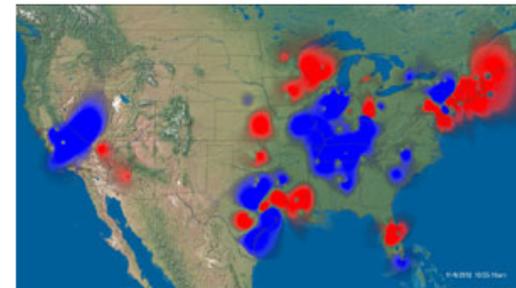


[+ view larger](#)

RED represents more **Negative** sentiment.
BLUE represents more **Positive** sentiment.

[Download](#) full resolution image (ZIP)

US Presidential Election 2012

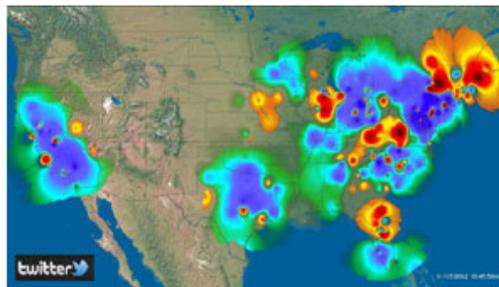


[+ view larger](#)

RED represents tweets about **Romney**.
BLUE represents tweets about **Obama**.

[Download](#) full resolution image (ZIP)

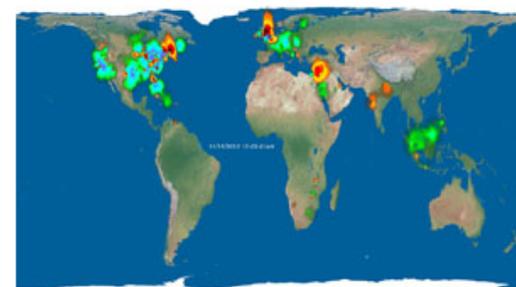
US Sentiment from Live Twitter Feed



[+ view larger](#)

RED represents more **Negative** sentiment.
BLUE represents more **Positive** sentiment.

Global Sentiment from Live Twitter Feed



[+ view larger](#)

RED represents more **Negative** sentiment.
BLUE represents more **Positive** sentiment.

People are Sharing Rich Emotions



Search for “Happiness”

-  **ARTPOP Makes Me Glow** @LadyGagaReturn 43m
Without this woman I would not find **happiness**. I would not have this wonderful family. #GagaYouAreOurHappiness
pic.twitter.com/zpNA4eOU9f
[View photo](#) [Reply](#) [Retweet](#) [Favorite](#) [More](#)
-  **Mikey™** @MikeySoAmazin 1h
You're my baby, my headache, my love, my smile, my frown, my wrong, my right, my pain, my **happiness**, my everything. You're MINE
[Expand](#)
-  **Brain Pickings** @brainpickings 1h
For Mental Health Week, the science of how our mind-wandering is robbing us of **happiness** j.mp/13QCfcz
[Expand](#)
-  **Damn Its TRUE!** @damnitstrue 1h
When you really care about someone, their **happiness** matters more than yours.
[Expand](#)
-  **Facts About Girls!** @girlnotes 1h
Your **happiness** depends on you, & only you.
[Expand](#)
-  **Nikki Woods** @nikkiwoods 1h
The only keeper of your **happiness** is YOU. Stop giving ppl power to control your smile, your worth, your attitude.
[Expand](#)
-  **Sol** @Solzilla 1h
There is no iPhone app for **happiness**... Disconnect and reunite with the real world. #Zilla
[Expand](#)
-  **Because We're Rich** @CauseWereRich 2h
Whoever said you can't buy **happiness**... Was shopping at the wrong store.
[Expand](#)

More Visual Content in Different Culture



Search for “幸福”

竹本人：为什么中国人民勤劳但不幸福
府，三公消费；2.央企垄断，与民争利；3
年收入；4.坑爹的企业污染环境，损害人民
的企业食品造 <http://t.cn/zHUUUWS>



5分钟前 来自开心网

净琳念佛-永远20岁的雨心：//@白羊座
不在于富足，而在于满足

@哲语禅思：人生的幸福，不在于富足
不在于多加燃料，而在于减少火苗；
而在于减少欲念。人的欲望是无止境的
的，以有限的生命追求无尽的欲望，又
下贪欲，追求平实简朴的生活，是获取
法。每天参一禅，让心灵变得豁达安宁

收起 | 查看大图 | 向左转 | 向右转



5月20日02:23 来自新浪微博

5分钟前 来自Android客户端

奶茶小猫V6：言言，谁要是
福！

@陈柏言：做完我都不饿了。



今天12:43 来自新浪微博

八达岭2003：希望在当下，不在昨天。优雅地转身，才能邂逅越来越
多的幸福。

@anni1986空间：放不下过去，如同开反方向的车，会离幸福
越来越远。人生如驾车，路过的，都是风景，风景再美，既已路
过，就要离开，无论舍不得，人生都要前行。放得下，过去就是
支撑未来的基石；放不下，昨天就是拖累今天的包袱。人生有希望
才有活头，而希望在当下，不在昨天。优雅地转身，才能邂逅越来越
多的幸福。

收起 | 查看大图 | 向左转 | 向右转



52分钟前 来自华为MediaPad

转发(19) | 评论(4)

3分钟前 来自iPhone客户端

转发 | 收藏 | 评论

北京巴沟韩颖10：[微喊话]韩颖于2012年3月3日交到海淀法院立案庭的
三个案子到现在已经449天了，多次催促立案，海淀法院武会珍庭长坚定
地说：立不了案！但始终不说因什么原因立不了案。这个案子有明确的
被告，有明确的诉讼请求，有基本的证据，属于你院管辖，请武庭长依
法做出立案或不立案裁定。期盼立案，渴望幸福



3分钟前 来自iPhone客户端

转发 | 收藏 | 评论

渊海紫坪：幸福就是外面哗啦啦的下着雨，我虽然早醒却依然可以再闭
上眼睡一个回笼觉。这天适合在家睡个觉，开着台灯，看本小说。科
科。



3分钟前 来自Android客户端

转发 | 收藏 | 评论

The Power of Social (Visual) Multimedia

- A picture is worth one thousand words

Example Tweets

@BarackObama: Four more years.



@Brynn4NY: Rollercoaster at sea.

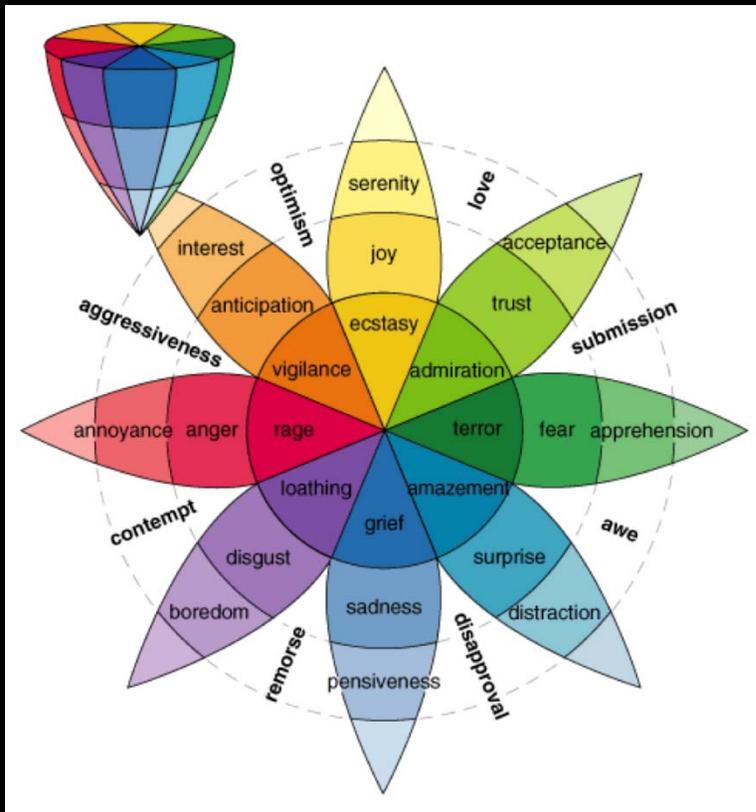


@Fang-Ru: Queen of the far far away land.



Research: Which 1000 sentimental concepts in pictures?

-- Web + big data + computer vision + psychology



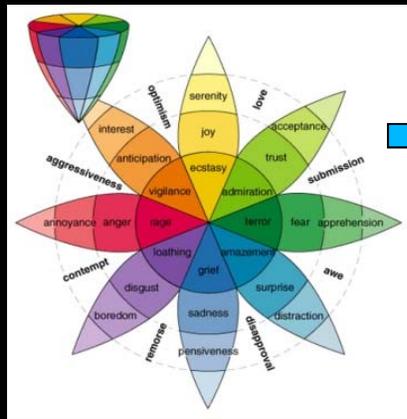
Psychology emotion wheel
(24 emotions, by Robert Plutchik)



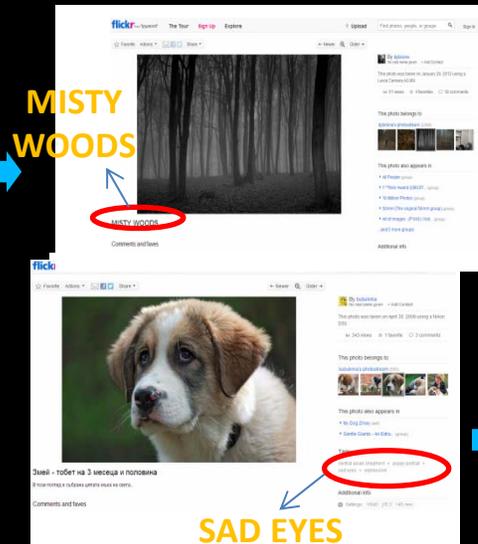
Plenty on the Web:
"For content to go viral, it needs to
be emotional," Dan Jones

Research: Which 1000 sentimental concepts?

-- data mining to discover visual sentiments in social media



Psychology emotion wheel (24 emotions)



Analyze tags with strong sentiments

Build Sentiment Ontology



Select Adj-Noun Pairs

Concurrent tags with emotions

joy	terror	amazement	disgust
joy	terror	amazing	disgusting
happy	horror	beautiful	gross
love	zombie	nature	food
smile	fear	wonder	nasty
beautiful	dark	light	sick
flowers	street	love	dirty
light	halloween	sky	dead
nature	war	eyes	face
kids	undead	clouds	blood
christmas	bomb	landscape	insect

From 6 million tags on Flickr and YouTube
Color code: text sentiment values

Not all concepts/entities are detectable!
-- which 1000 concepts to focus in pictures?

Target Concepts for CV – Adj-Noun Pair

- Adjective (268): needed for expressing emotions
 - frequent positive Adj: beautiful, amazing, cute
 - frequent negative Adj: sad, angry, dark
- Nouns (1187): feasible for computer vision
 - Noun categories: people, places, animals, food, objects, weather
- Standard steps:
 - remove named entities like “hot dog” via wikipedia
 - Choose sentiment rich ANP concepts by tools “Senti-WordNet” “SentiStrength”

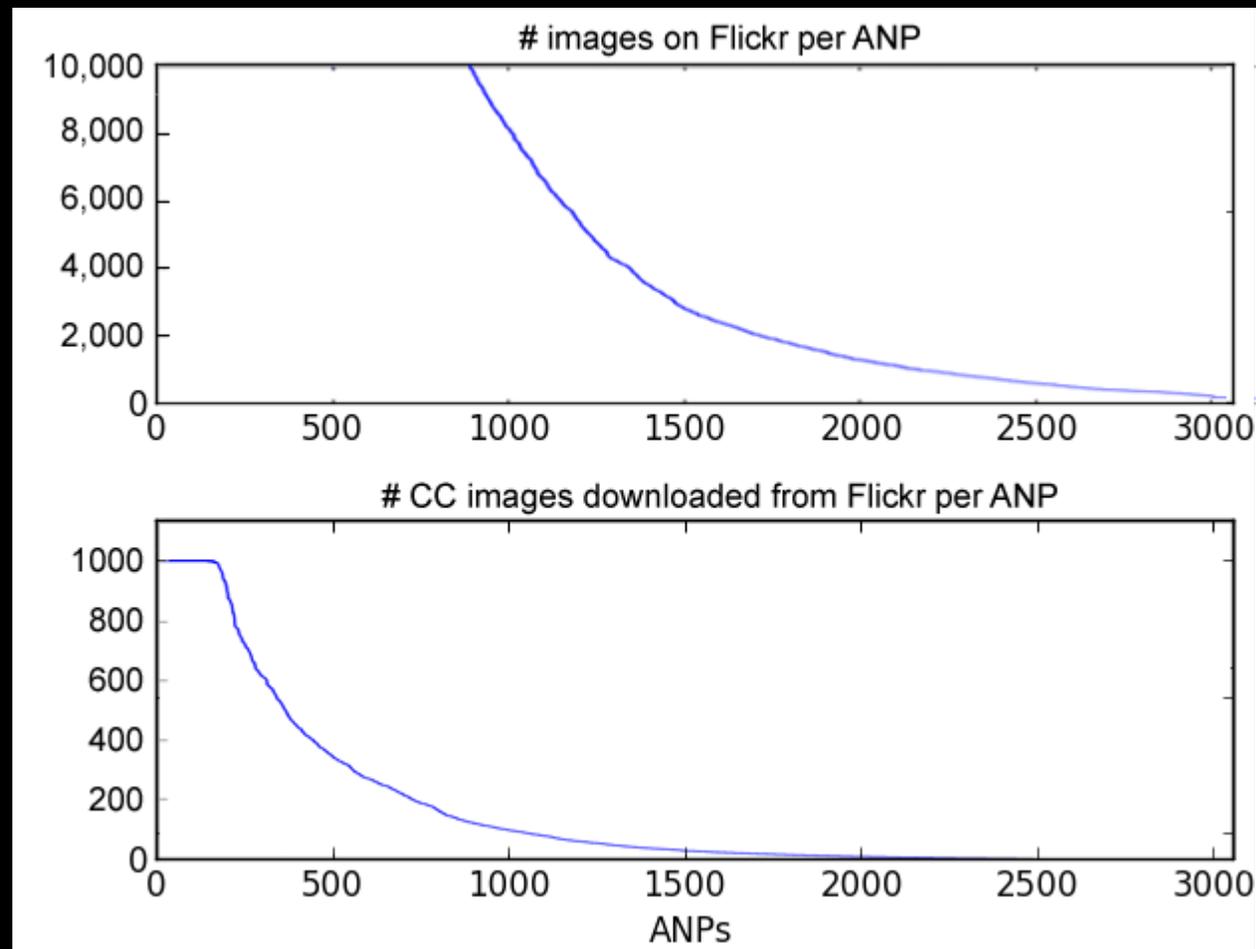
Slippery Slope
Discovered Serpentine



Currently, we have found 3000+ ANBs
Discovered Serpentine ANBs for Social Media Photos

Image Datasets

About 0.5 million images over 3000 concepts

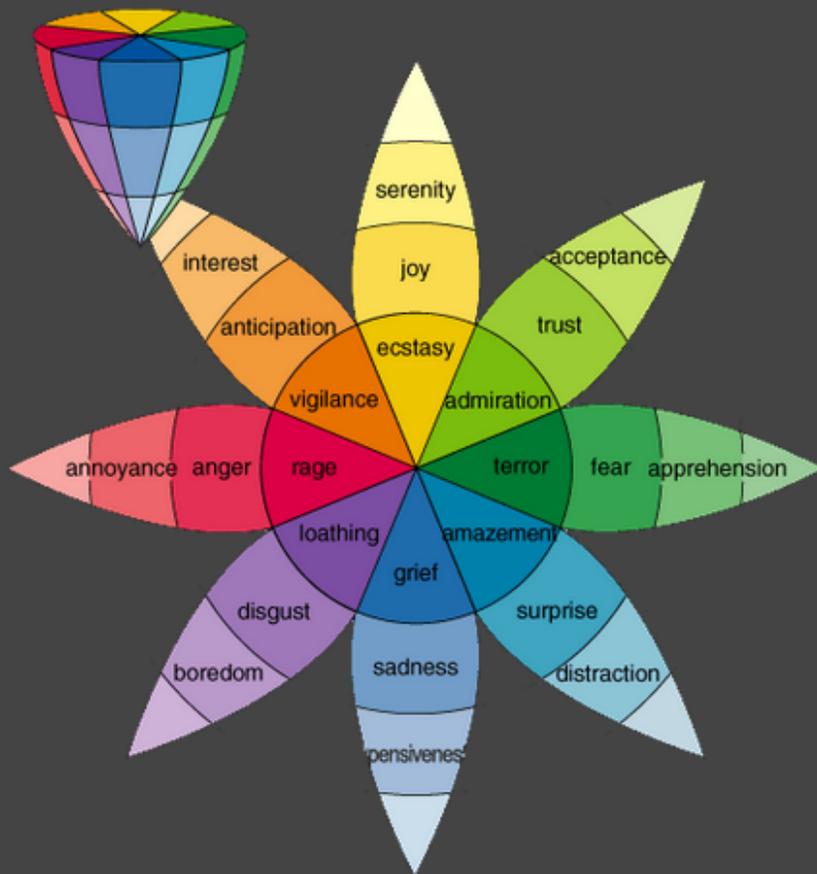


Visual Sentiment Ontology (Browser)

Home :: Ontology :: Adjective Noun Pairs :: Downloads :: About

Visual Sentiment Ontology 

Ontology :: Emotional Mapping



Selected Emotion:

◀ joy [591] ▶

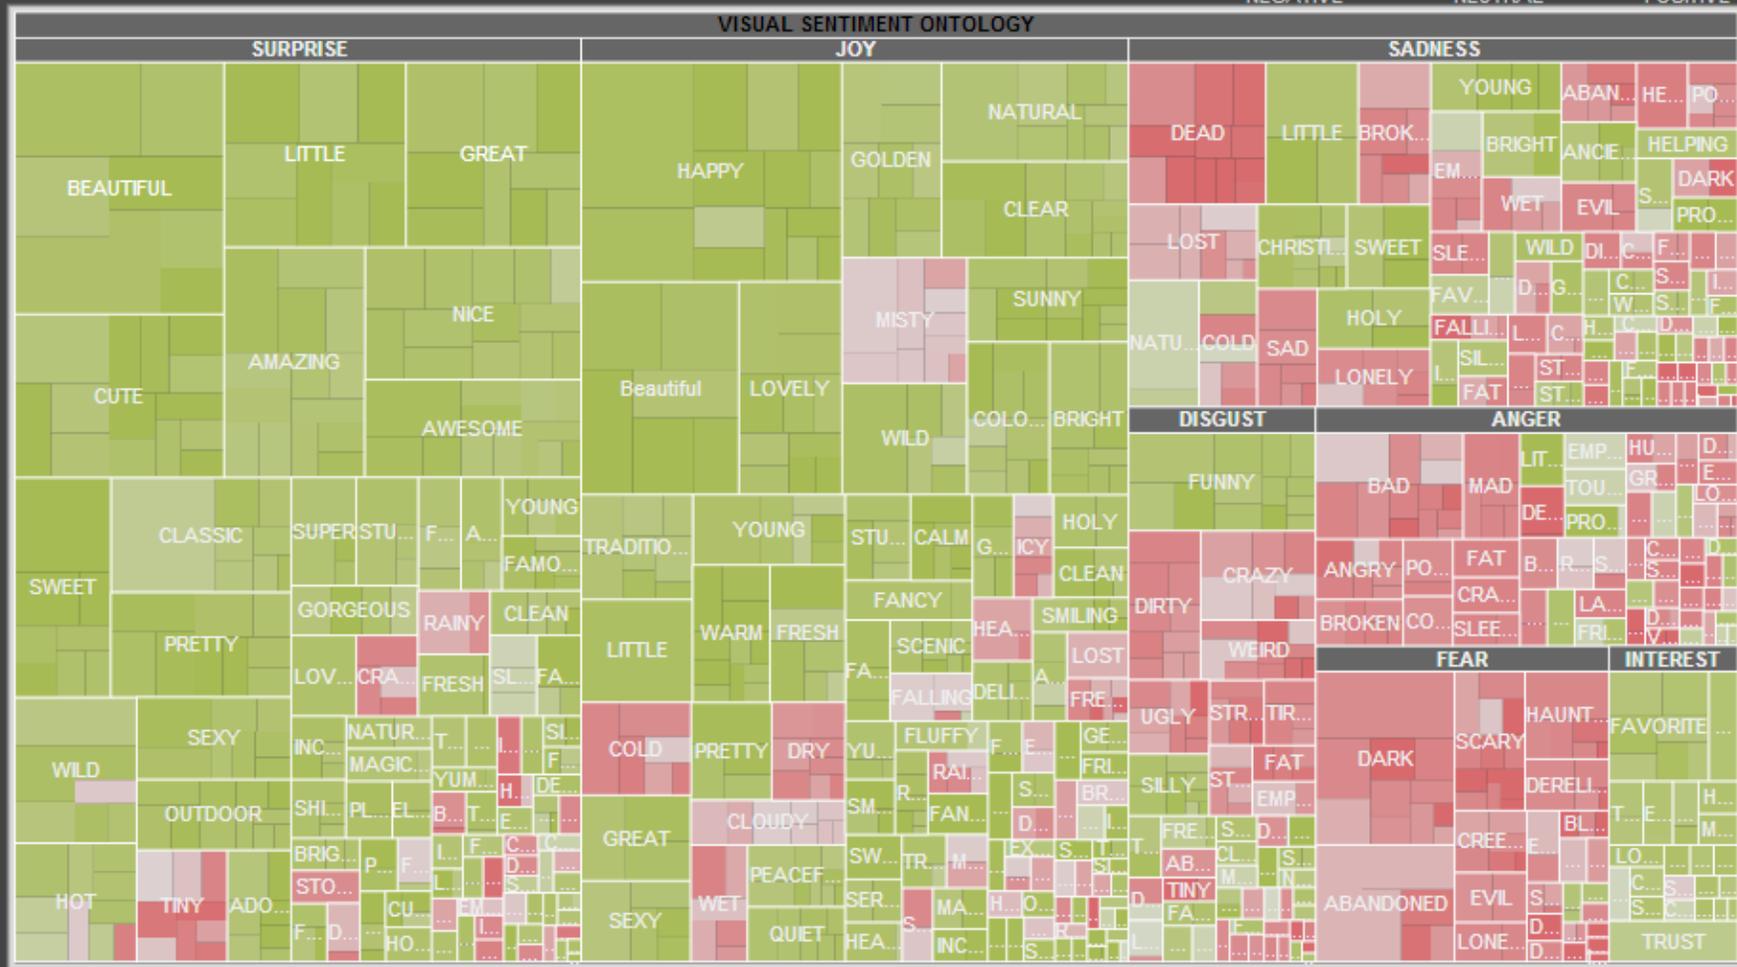
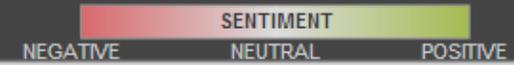
Top ANPs	sentiment	emotion
1. ▶ happy smile	1.92	0.388
2. ▶ innocent smile	1.92	0.376
3. ▶ happy christmas	2.0	0.373
4. ▶ happy father	2.0	0.358
5. ▶ happy wedding	1.72	0.348
6. ▶ friendly smile	1.92	0.346
7. ▶ delicious cupcake	1.71	0.341
8. ▶ shy smile	0.62	0.340
9. ▶ charming smile	1.92	0.339
10. ▶ happy birthday	1.79	0.337
11. ▶ warm smile	1.92	0.336
12. ▶ happy mother	2.0	0.333
13. ▶ happy halloween	1.81	0.331
14. ▶ delicious drink	1.59	0.330
15. ▶ happy heart	2.0	0.329
16. ▶ happy kids	2.0	0.323
17. ▶ healthy food	1.69	0.312
18. ▶ happy guy	1.61	0.312
19. ▶ fresh food	1.59	0.307
20. ▶ delicious pie	1.76	0.296

Visual Sentiment Ontology (Browser)

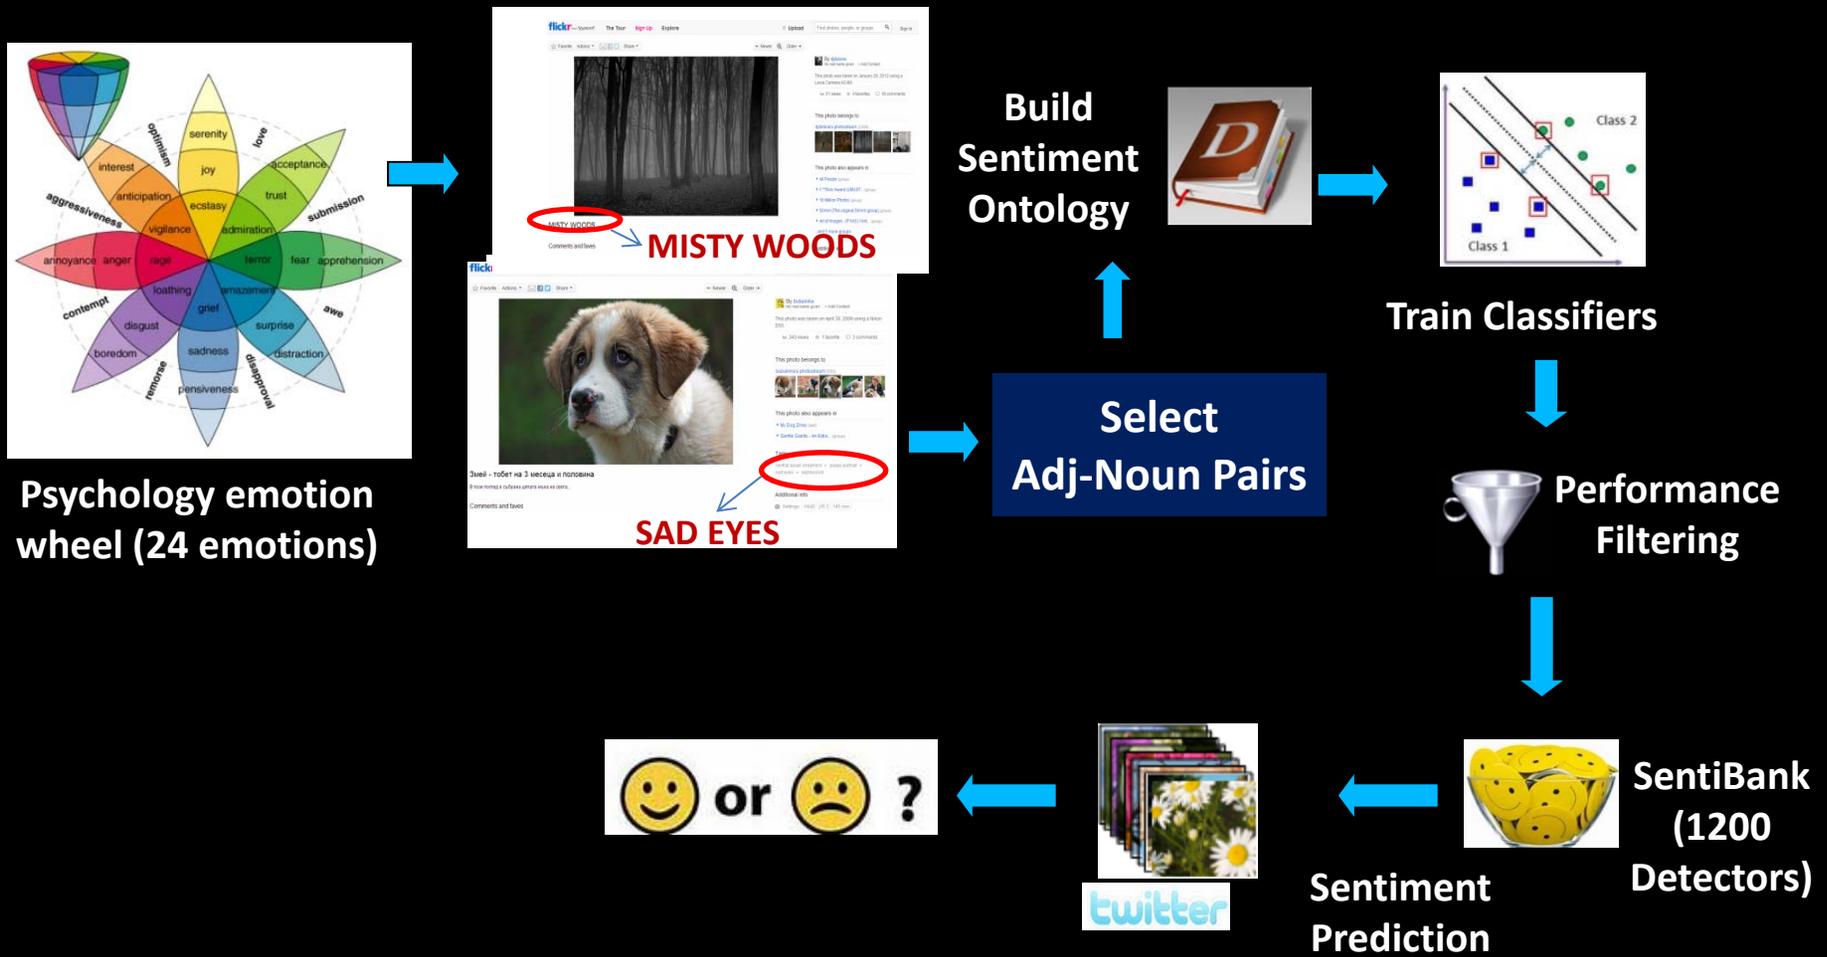
Home :: Ontology :: Adjective Noun Pairs :: Downloads :: About

Visual Sentiment Ontology 

Ontology :: Treemap Visualization



Next Step: Teach Machine to Recognize Visual Sentiments



Standard Classifier Training

- LibSVM, 5-fold cross validation
- Features
 - RGB Color Histogram (3x256 dim.)
 - GIST descriptor (512 dim.)
 - Local Binary Pattern (52 dim.)
 - SIFT Bag-of-Words (1,000 codewords
2-layer spatial pyramid, max pooling)
 - Classemes descriptor (2,659 dim.)

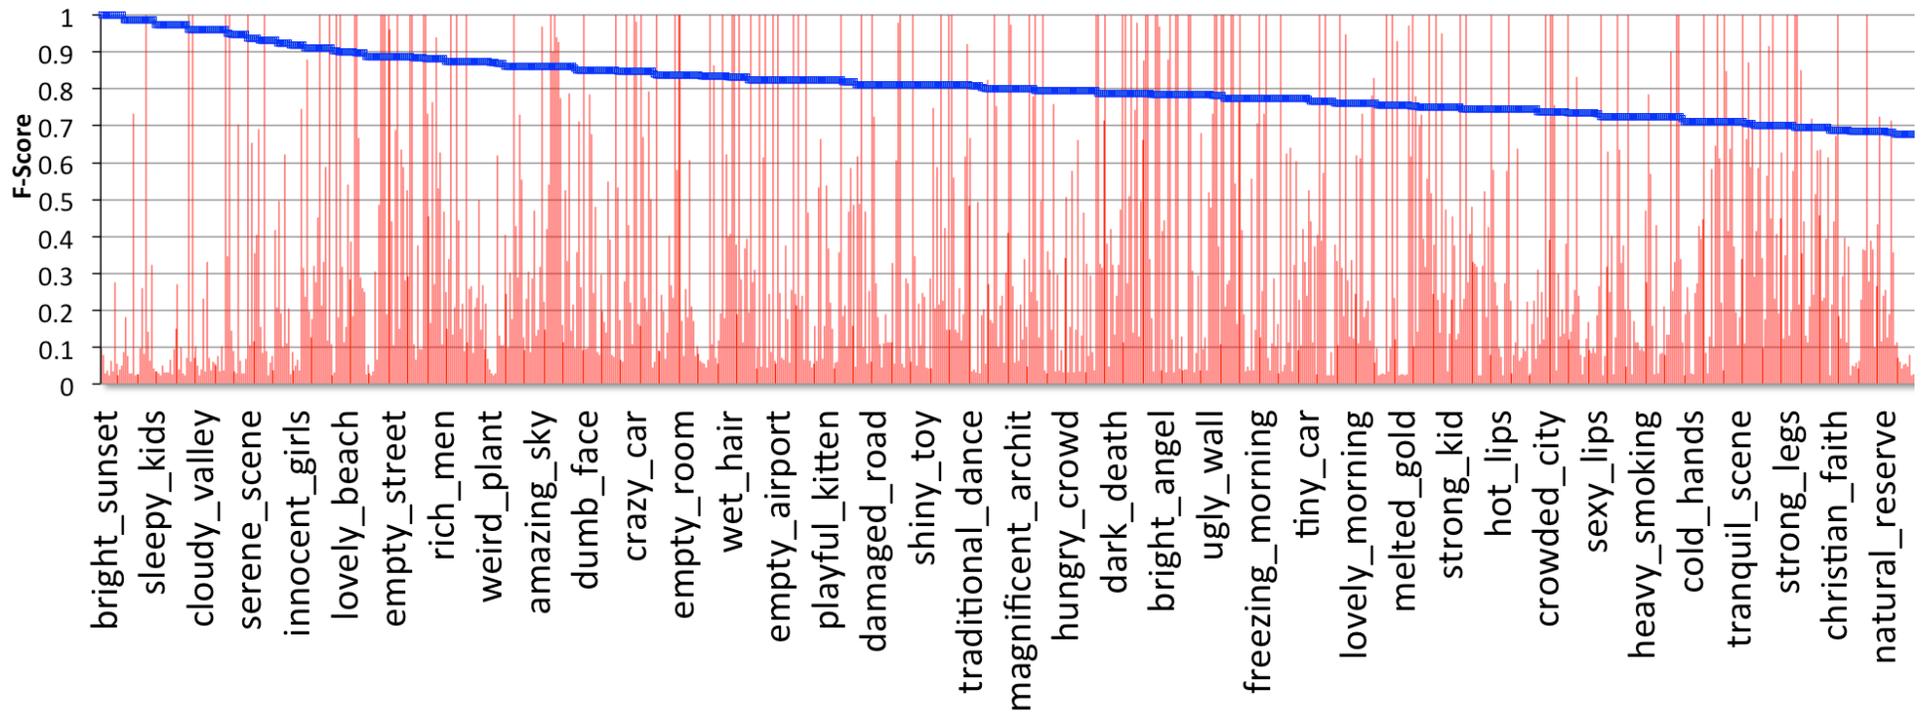
Beautiful flower
Happy face



Machine Detected Visual Sentiments
More than 600 classifiers with F score @40 > 78%

Green: Correct Red: Incorrect

Detector Accuracy vs. Frequency

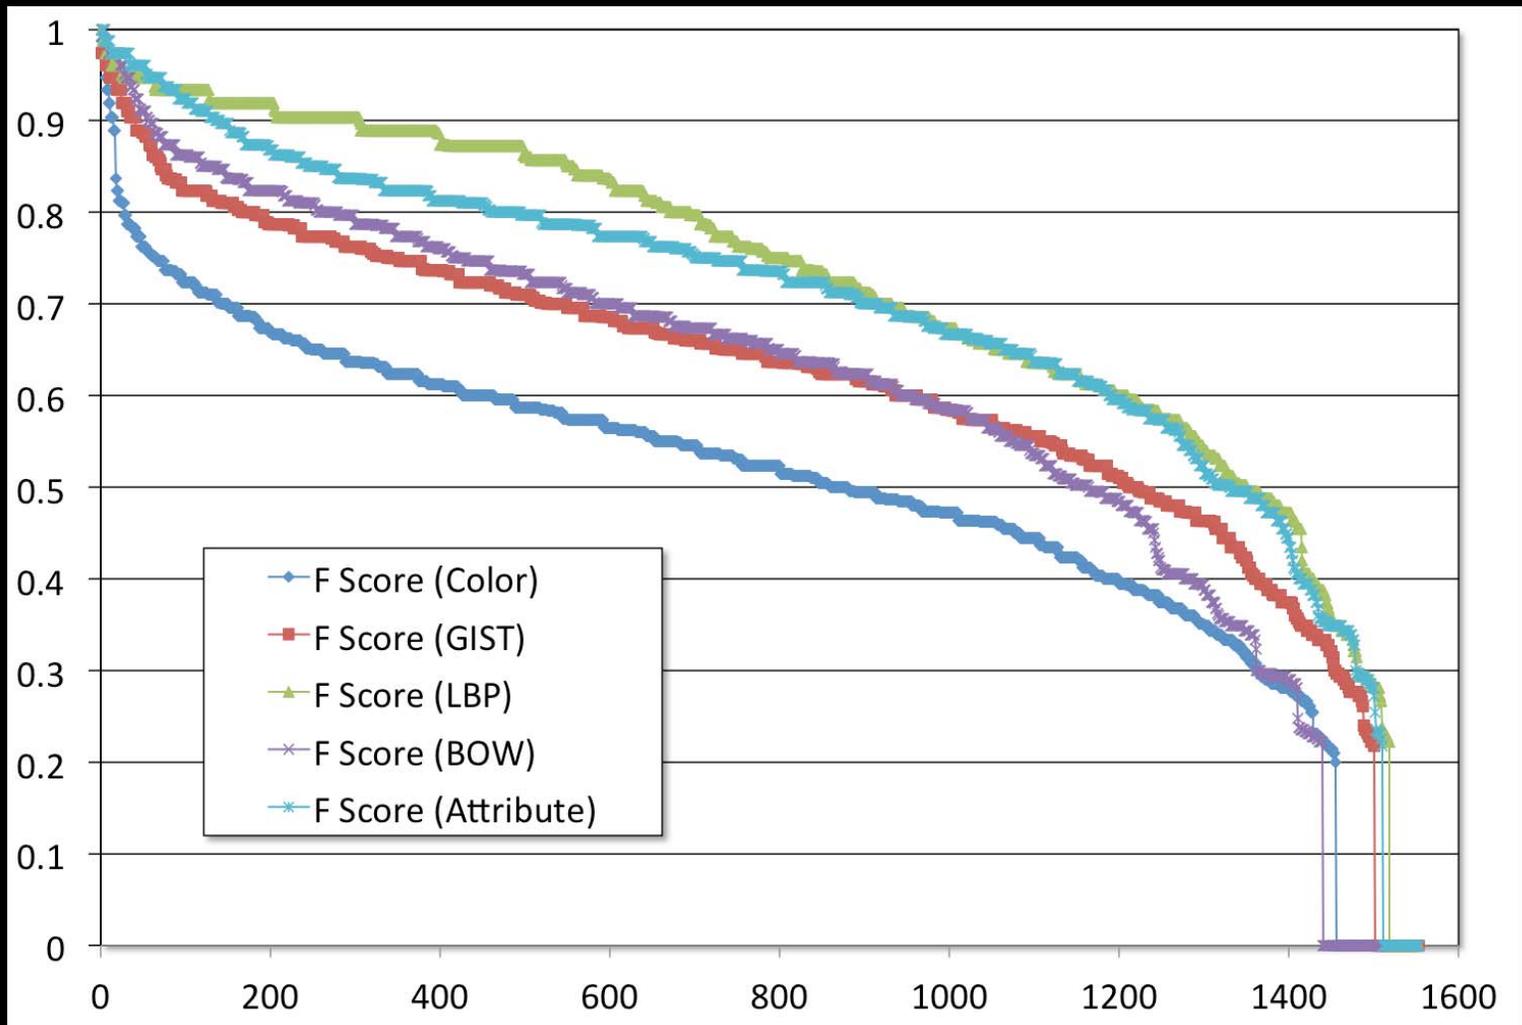


Examples

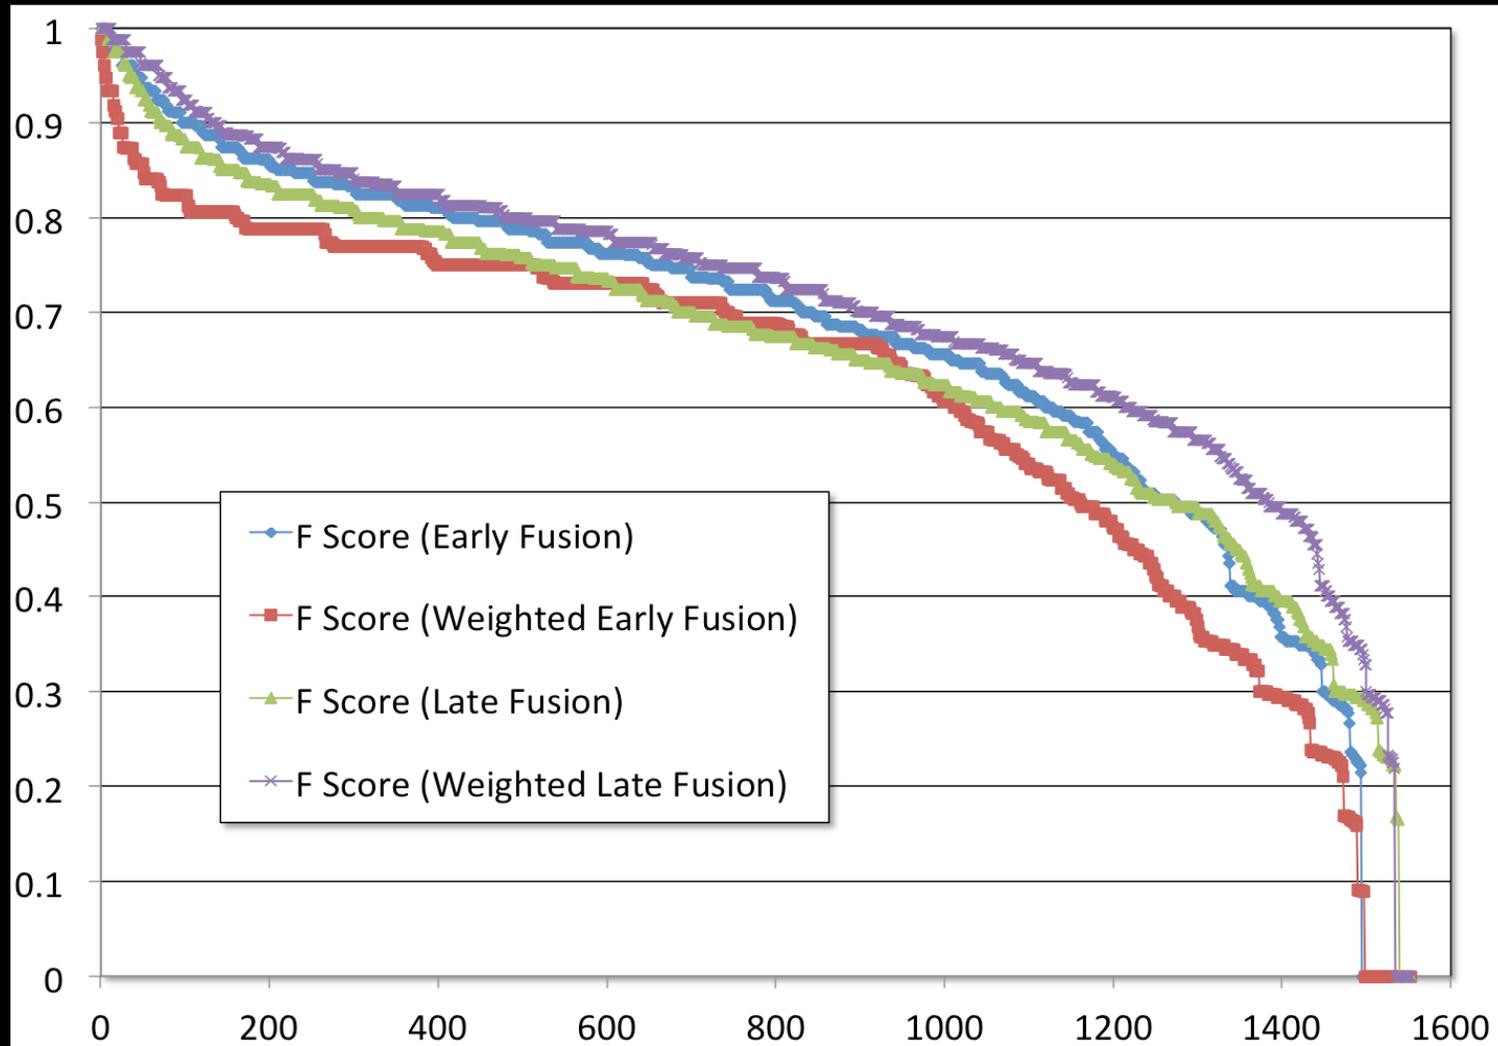
Good Results:



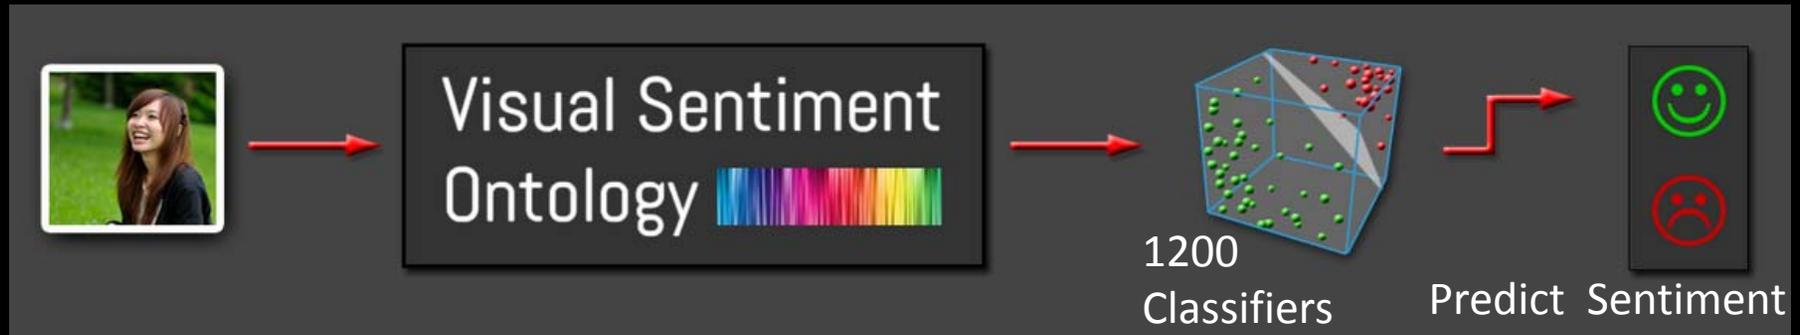
Performance vs. Features



Performance vs. Fusion



Application: Live Sentiment Prediction



PhotoTweet Stream:

The PhotoTweet stream consists of several elements:

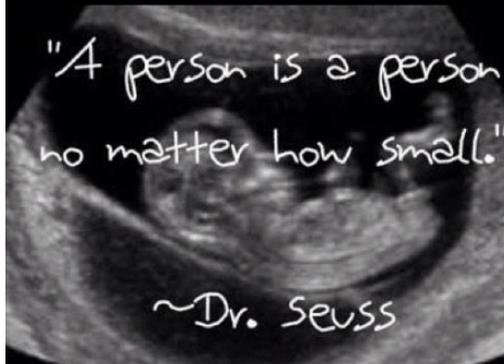
- A quote: "A person is... no matter how... ~Dr. S"
- A photo of a beach with destroyed buildings, captioned: "True stuff. I have many ladies that DO NOT get... @nickespo89 #groundzero #hurricane #newjersey @charleslawrence"
- A photo of a damaged police car, captioned: "Ouch mr police man @radiodario"

A large blue arrow points from the collage to the following text:

Positive?
Neutral?
Negative?

Viewer Response Depends ...

- Responses depend on viewer's perspective
- Multi-user sentiment AMT labeling over 2000 phototweets



True stuff. I have mad respect for all the ladies that DO NOT give in to abortion.

@nickespo89

Amazon Mechanic Turk Sentiment/Emotion Label:

(image-based labeling)

worker 1: Positive, trust:acceptance

worker 2: Neutral, interest:unlabeled,sad:pensiveness

worker 3: Positive, interest:interest

(text-based labeling)

worker 1: Positive, joy:serenity,trust:acceptance

worker 2: Positive, anger:neutral,interest:interest,joy:serenity,trust:acceptance

worker 3: Negative, sad:sadness

(text-image-based labeling)

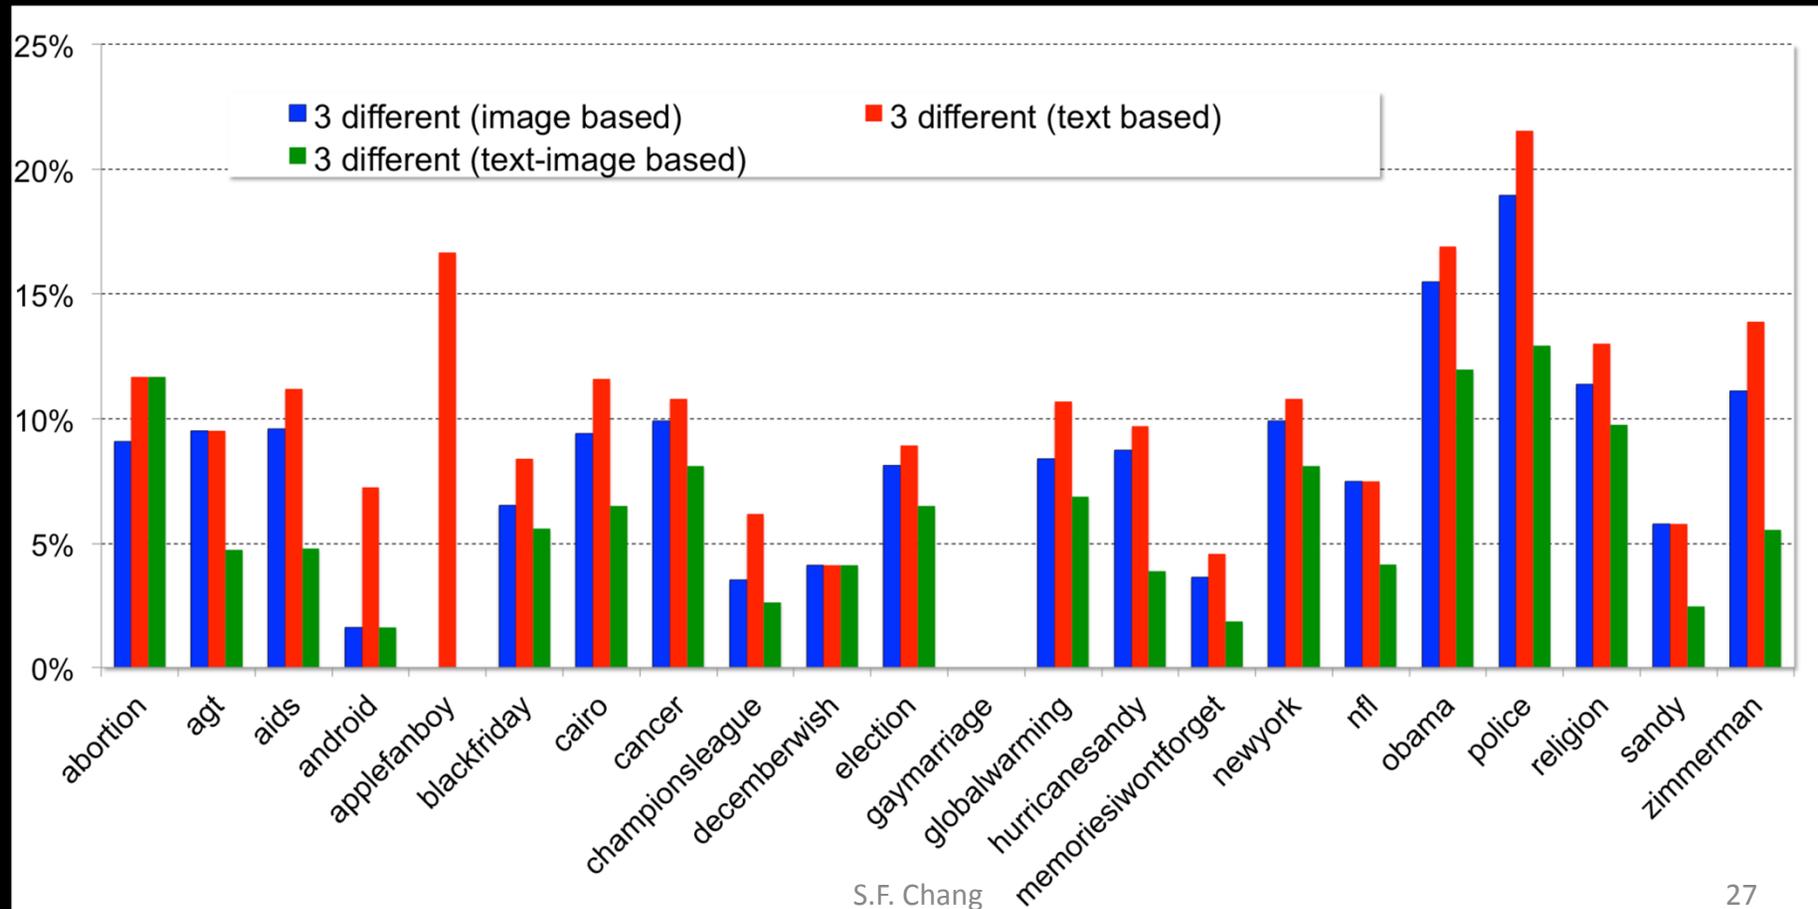
worker 1: Positive, joy:serenity,sad:neutral

worker 2: Positive, interest:interest,joy:joy,sad:neutral,surprise:distracted

worker 3: Positive, joy:serenity,surprise:neutral,trust:trust

Response also Depends on Topic

- Text more controversial than image in invoking responses
- Response inconsistency varies across topics



Sentiment Prediction Performance

Sentiment Prediction Accuracy

Text	0.61
Visual	0.65
Text-Visual (Joint)	0.74

[Demo](#)

Examples



young_teen happy_heart young_friends fat_girls happy_face
cute_girls fluffy_cat sweet_girls cute_dog friendly_smile
funny_kids



young_friends cold_feet stupid_hat heavy_winter
waiting_area crazy_hair stupid_sign fat_face harsh_winter



happy_heart sweet_girls friendly_smile traditional_wedding
grumpy_face young_teen handsome_face beautiful_flower
wedding_friends happy_wedding



violent_crime bad_guy dark_blood clean_air
ancient_sculpture funny_comic angry_men gorgeous_girls
tired_eyes tired_men dark_death dark_eyes traditional_tattoo



sweet_child great_night tired_eyes creepy_horror
dark_places dark_blood dark_woods wet_window dark_room
favorite_book young_friends dark_death weird_face
hardcore_band favorite_club hardcore_punk

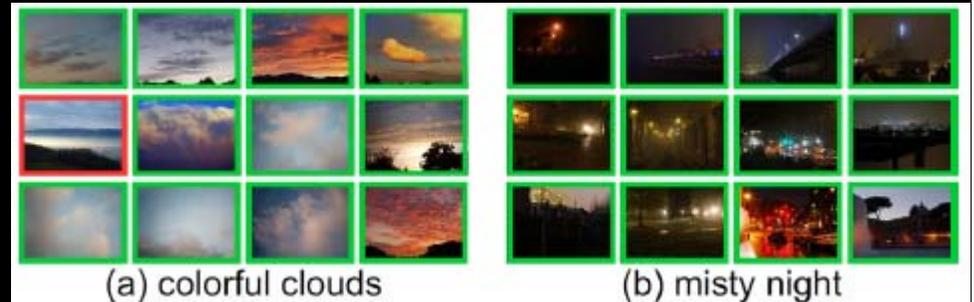
Conclusions

- Effort to build visual sentiment ontology
 - Psychology and Web folksonomy
 - Unique adjective-noun pair concepts
- Initial results in large-scale detectors
 - Ontology 3000 concepts, **SentiBank** 1200 detectors
 - Datasets (0.5 million images) and tools available
- Applications
 - Multi-modal sentiment monitoring
 - Intuitive visualization tools

Open Issues

- Improve detection of objects and sentiment attributes
 - E.g., object/scene aesthetic attributes, face emotions

Good Results:



Not so Good Results:



Open Issues

- Generalization
 - Adapt ontology and detectors to different domains, data types like video, etc.
- Relation with Audience Sentiments
- New applications – editing, recommendation