Placing Flickr Photos on a Map

Fateme Shirazi

Saarland University

Outline

- Introduction
- Representing Locations on the Map
- Modeling Locations
- Experimental Setup
- Results
- Error analysis
- Conclusion

Outline

- Introduction
- Representing Locations on the Map
- Modeling Locations
- Experimental Setup
- Results
- Error analysis
- Conclusion


flickr from YAHOO!
Signed in as FS24
(1 new)
Help
Sign Out
Home
You
Organize & Create
Contacts
Groups
Explore
Find a place
Search
Link to this map
Map
Hybrid
Satellite
Find my location
DSC01118 by FS24
Your stuff (9).
Sort by: Interesting • Recent
Search the map

flickr from YAHOO!
Signed in as FS24
(1 new)
Help
Sign Out
Home
You
Organize & Create
Contacts
Groups
Explore
Find a place
Search
Link to this map
Map
Hybrid
Satellite
Find my location
White Heaven by Katarina 2353
91,714,661 geotagged items
Sort by: Interesting • Recent
Search the map
1000km
750mi
Data ©2009 NAVTEQ

Introduction

- GPS-enabled cameras and mobile phones
  - latitude and longitude
- Geo-referenced photos
  - organized in a browsable taxonomy
  - or pin-pointed on a map

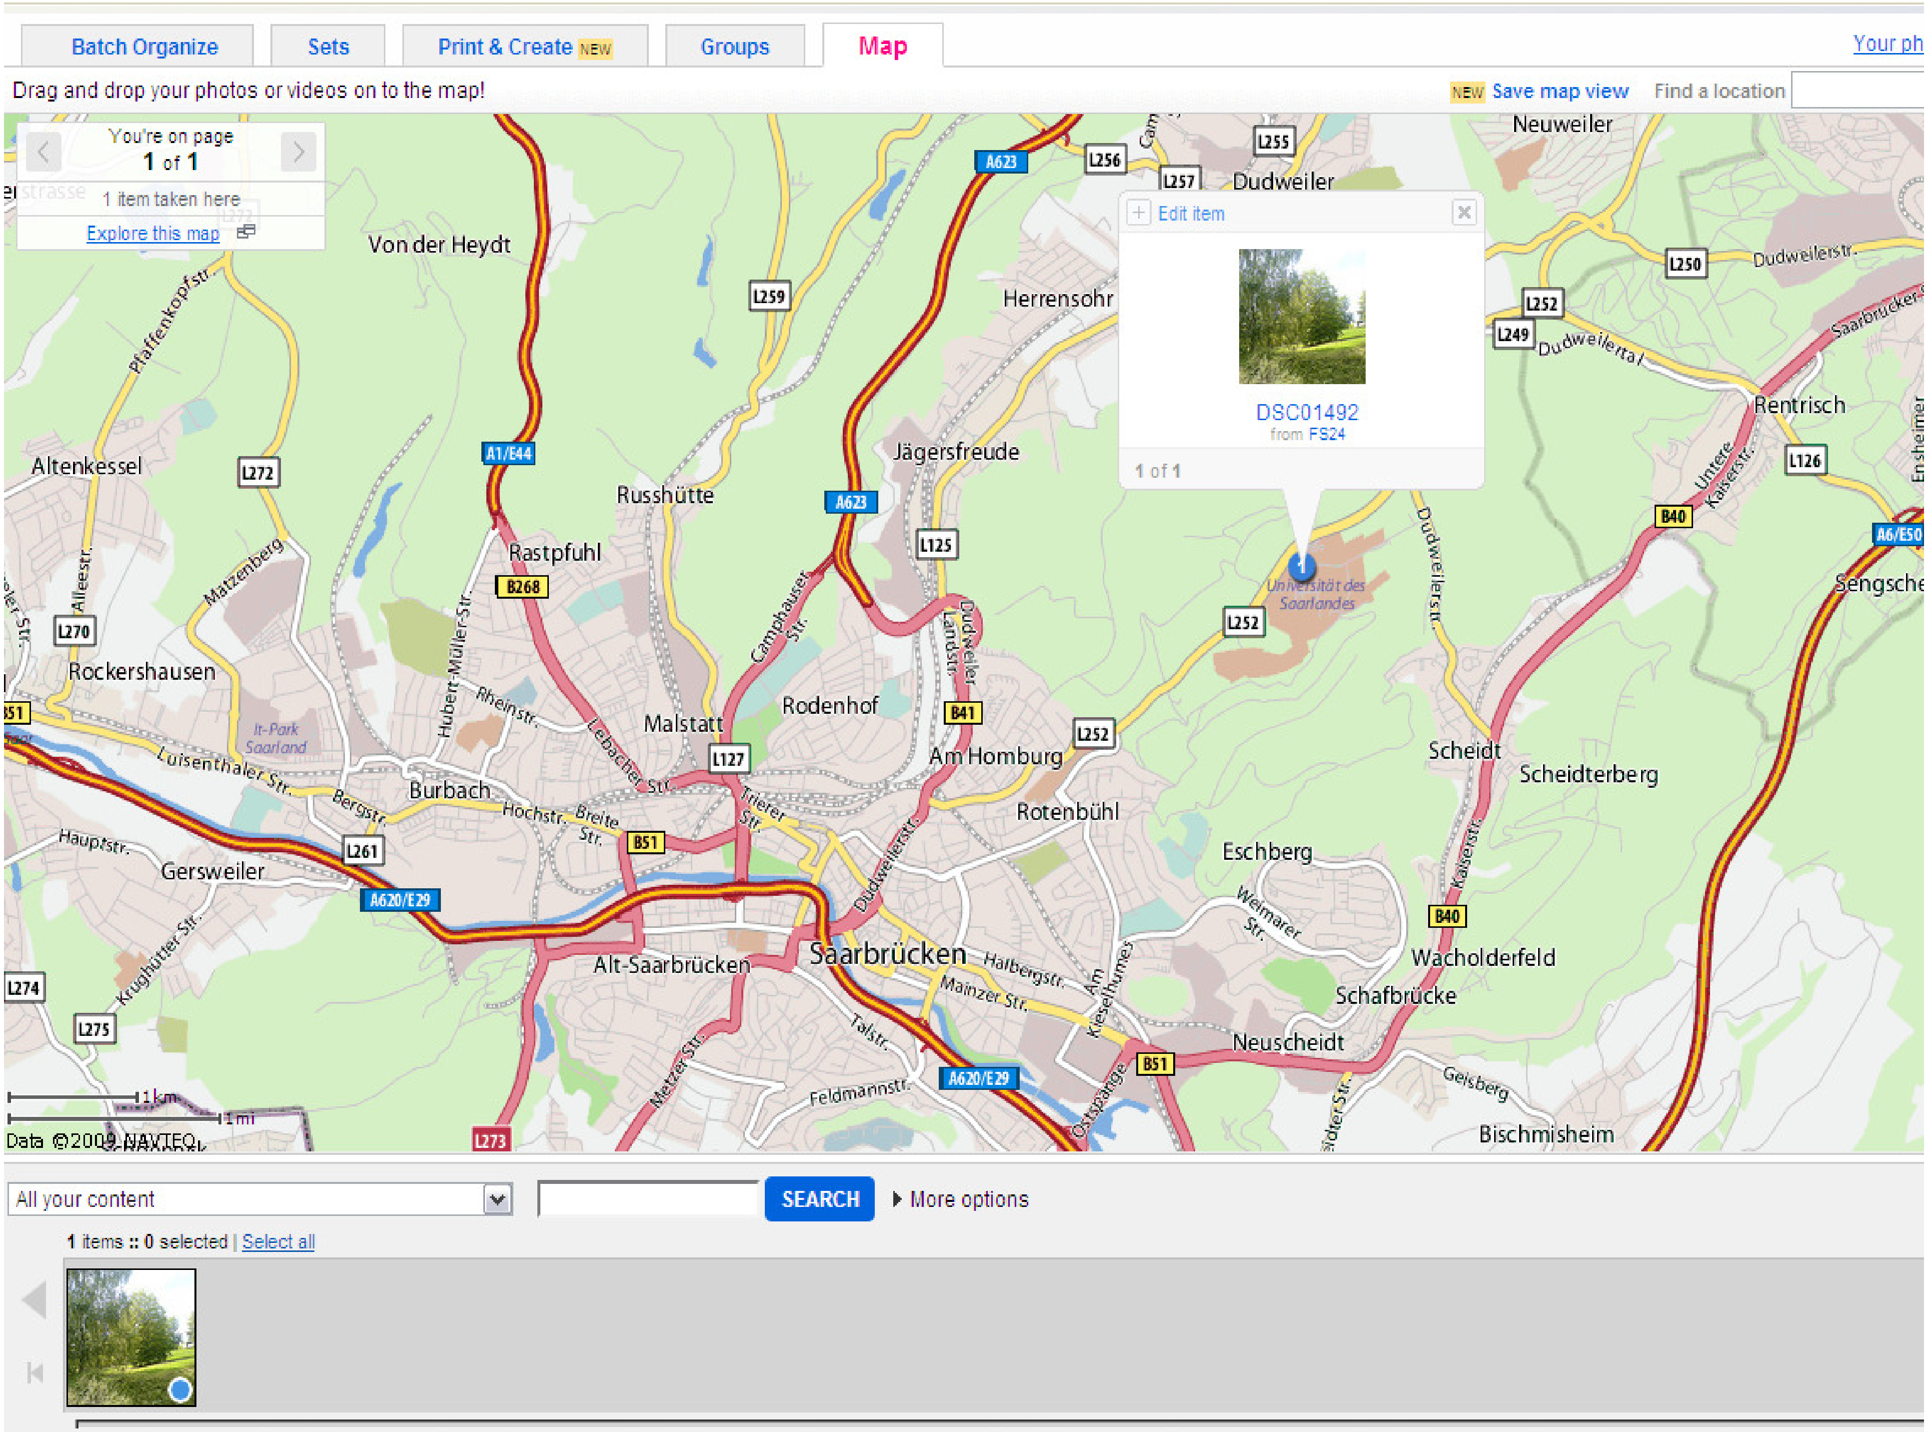

Batch Organize
Sets
Print & Create NEW
Groups
Map
Your ph
Drag and drop your photos or videos on to the map!
NEW Save map view
Find a location
You're on page
1 of 1
1 item taken here
Explore this map
Edit item
DSC01492
from FS24
1 of 1
Universität des Saarlandes
Saarbrücken
Alt-Saarbrücken
Dudweiler
Data ©2009 NAVTEQ
All your content
SEARCH
More options
1 items :: 0 selected | Select all

Introduction

- Objective of the paper:
  - Automatically **placing photos** uploaded in Flickr on the world map, **using tags** provided by the user.

Frame work of the paper:

- Constructing an $m \times n$ grid based on the longitude/latitude.
- Using a set of images whose locations are known.
- Placing images in its corresponding grid cell.
- Using the description of the images given by the Flickr users.
- Finally, investigating how to incorporate external resources into the model

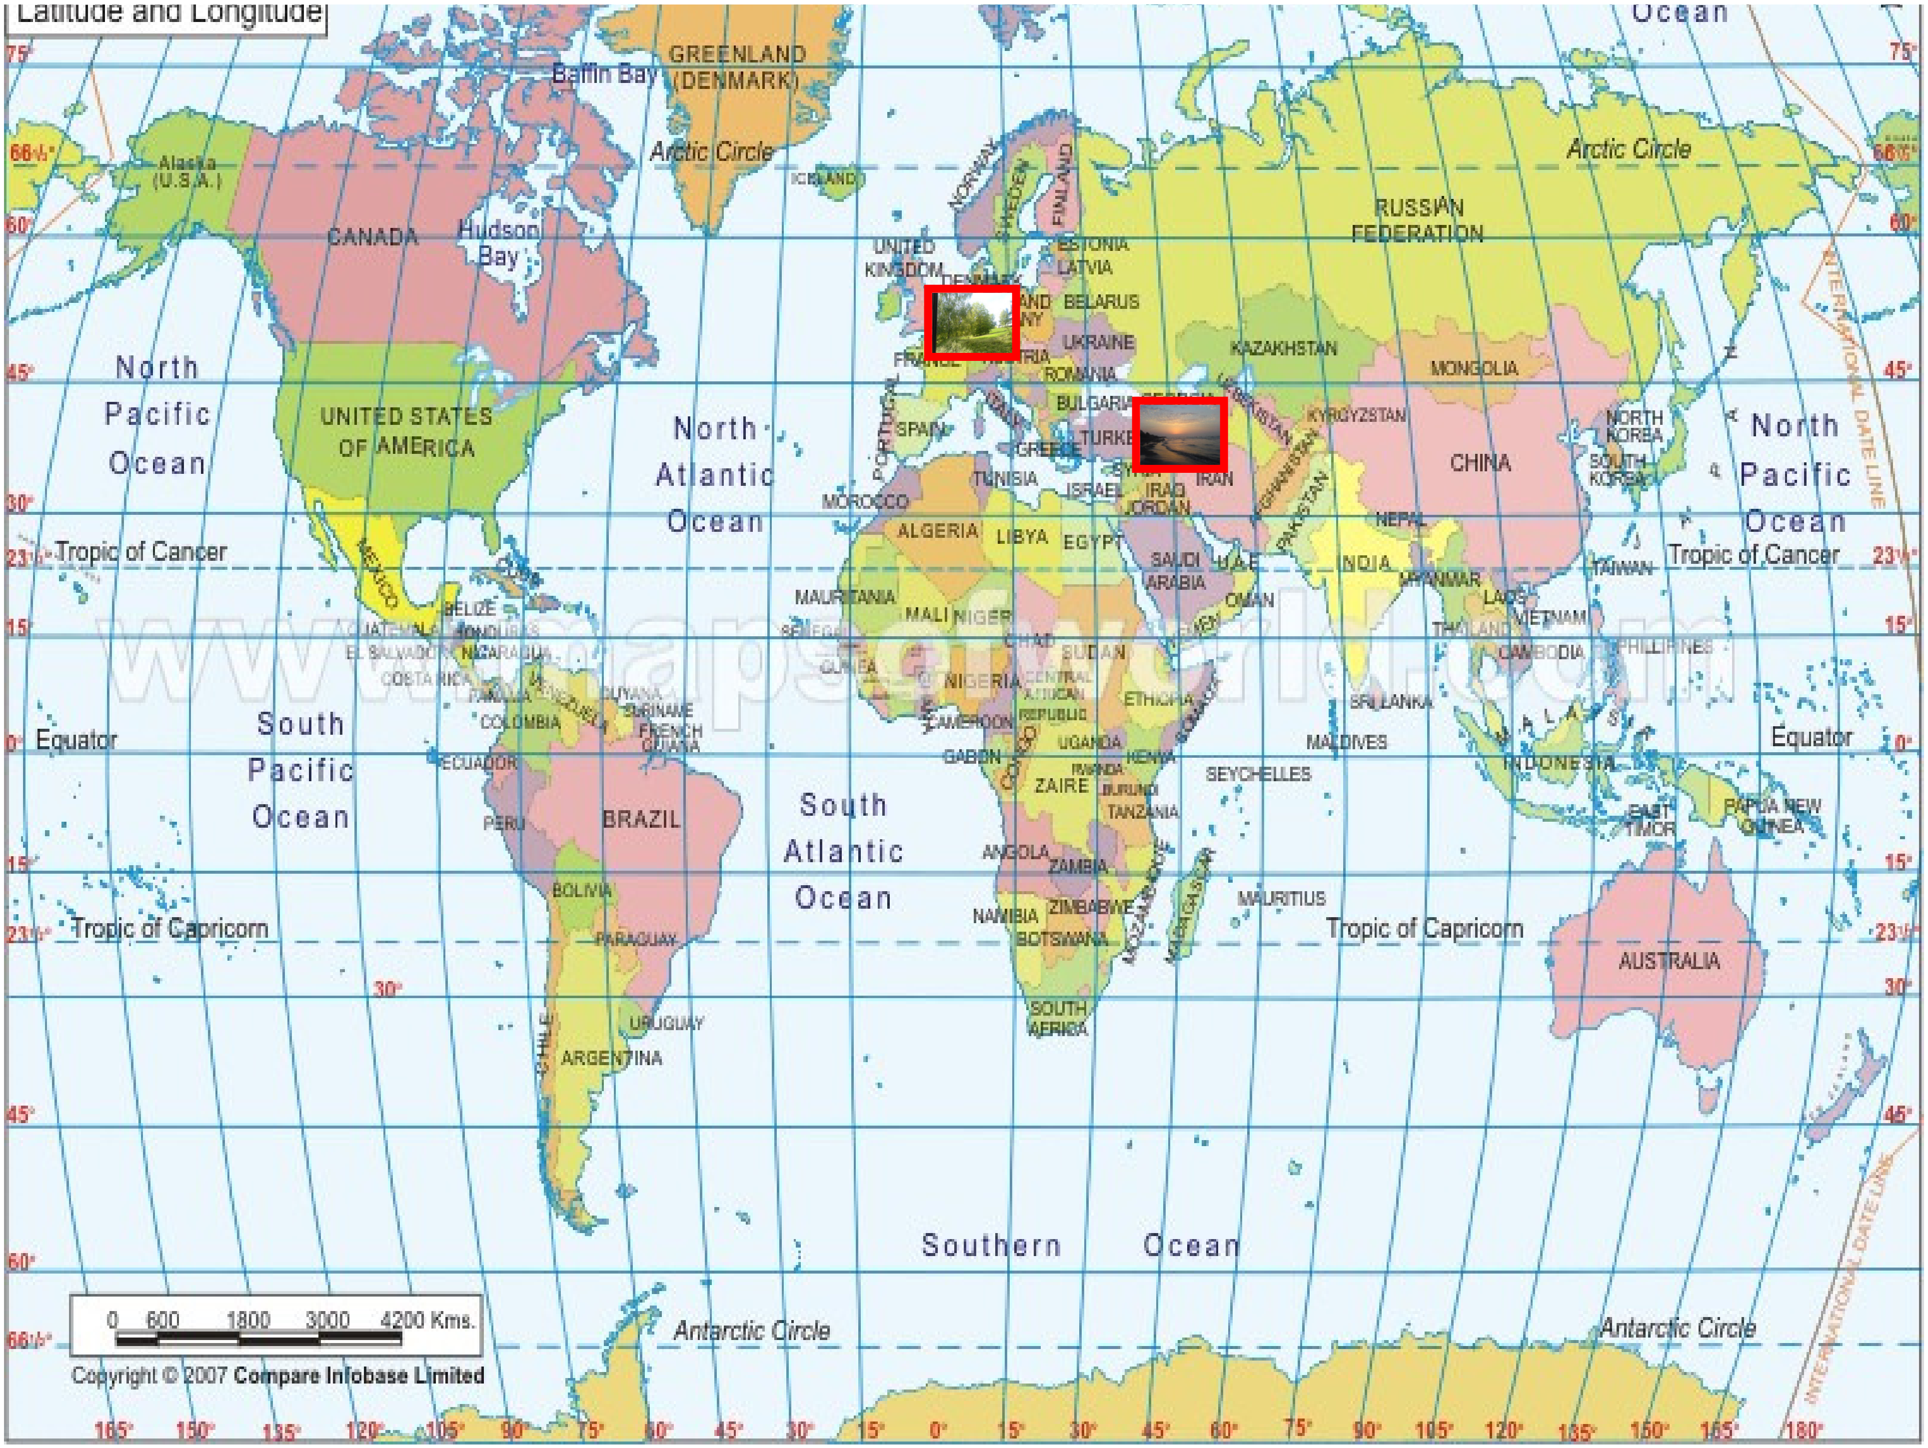

Latitude and Longitude
GREENLAND (DENMARK)
Baffin Bay
Arctic Circle
ICELAND
Alaska (U.S.A.)
CANADA
Hudson Bay
RUSSIAN FEDERATION
UNITED KINGDOM
NORWAY
SWEDEN
FINLAND
ESTONIA
LATVIA
BELARUS
UKRAINE
KAZAKHSTAN
MONGOLIA
ROMANIA
BULGARIA
FRANCE
SPAIN
PORTUGAL
GREECE
TURKEY
North Pacific Ocean
UNITED STATES OF AMERICA
North Atlantic Ocean
MOROCCO
TUNISIA
ISRAEL
IRAQ
JORDAN
IRAN
KYRGYZSTAN
CHINA
NORTH KOREA
SOUTH KOREA
AFGHANISTAN
PAKISTAN
NEPAL
Tropic of Cancer
MEXICO
ALGERIA
LIBYA
EGYPT
SAUDI ARABIA
U.A.E.
OMAN
INDIA
MYANMAR
TAIWAN
LAOS
BELIZE
MAURITANIA
MALI
NIGER
CHAD
SUDAN
THAILAND
VIETNAM
CAMBODIA
PHILIPPINES
NICARAGUA
COSTA RICA
EL SALVADOR
NIGERIA
CENTRAL AFRICAN REPUBLIC
ETHIOPIA
SRI LANKA
COLOMBIA
GUYANA
SURINAME
FRENCH GUIANA
CAMEROON
SOMALIA
Equator
South Pacific Ocean
ECUADOR
GABON
UGANDA
KENYA
MALDIVES
INDONESIA
PERU
BRAZIL
South Atlantic Ocean
CONGO
ZAIRE
RWANDA
BURUNDI
TANZANIA
SEYCHELLES
EAST TIMOR
PAPUA NEW GUINEA
ANGOLA
ZAMBIA
BOLIVIA
NAMIBIA
ZIMBABWE
MOZAMBIQUE
MADAGASCAR
MAURITIUS
Tropic of Capricorn
PARAGUAY
BOTSWANA
AUSTRALIA
SOUTH AFRICA
URUGUAY
CHILE
ARGENTINA
Southern Ocean
Antarctic Circle
INTERNATIONAL DATE LINE
0 600 1800 3000 4200 Kms.
Copyright © 2007 Compare Infobase Limited
www.mapsofworld.com

Outline

- Introduction
- Representing Locations on the Map
- Modeling Locations
- Experimental Setup
- Results
- Error analysis
- Conclusion

Representing Locations on the Map

- Locations as bags-of-tags
  - Each photo has a FlickrID, a geographic co-ordinate, and a set of tags.
  - Map each geographic co-ordinate on the map.
  - Derive a language model from tags of photos at this location.

Representing Locations on the Map

- Locations as a graph
  - Spatial relationship to build an undirected graph.
    - Where the link exists only if situated close enough on the grid.
  - Locations found within a predefined distance are linked and considered to be neighbors.

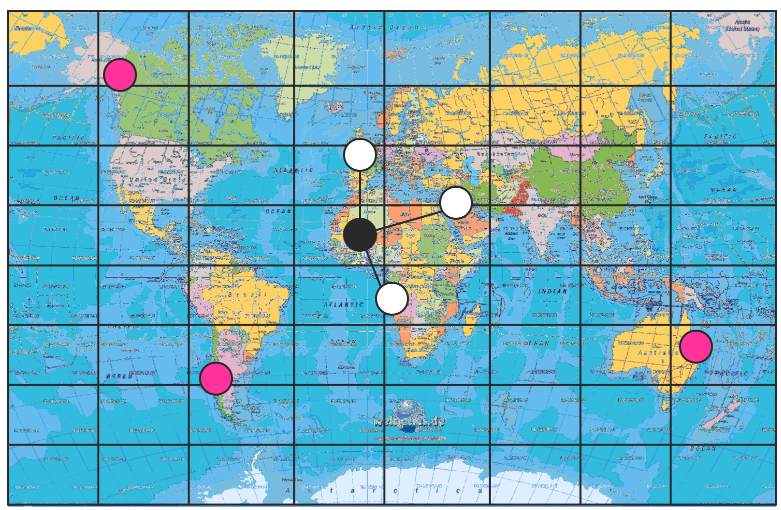

Representing Locations on the Map

- Locations as **pseudo-documents .**
- Linked locations more likely have *similar* tags .
- Locations relevant to an image being close in the graph.

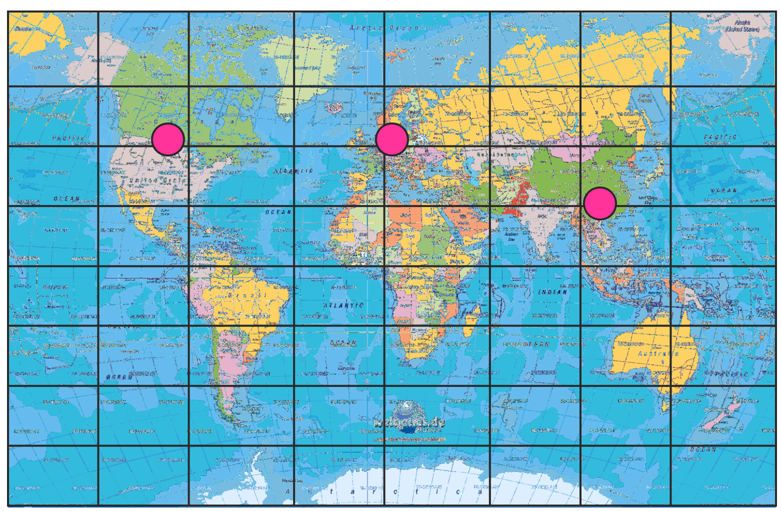

Outline

- Introduction
- Related work
- Representing Locations on the Map
- Modeling Locations
- Experimental Setup
- Results
- Error analysis
- Conclusion

Modeling Locations

- Ranking list of locations L, for a given tag set T belonging to an image taken within the bounds of L:

$$P(L|T) = \frac{P(T|L)P(L)}{P(T)}$$

- P(T|L) : the probability to generate the tag set of the image.

$$P(T|L) = \prod_{i=1}^{|T|} P(t_i|L)$$

Modeling Locations

$$P(t|L) = \frac{|L|}{|L| + \lambda} P(t|L)_{ML} + \frac{\lambda}{|L| + \lambda} P(t|G)_{ML}$$

- |L| : the size of the location L in tags.
- P(t|L) : maximum likelihood estimates of tag generation probabilities for the location's language model.
- P(t|G): maximum likelihood estimates of tag generation probabilities for the global language model.

Smoothing

- 1.Tag-based smoothing with neighbors
- 2.Smoothing cell relevance probabilities
- 3.Boosting geo-related tags
- 4.Spatial ambiguity-aware smoothing

1.Tag-based smoothing with neighbors

- To overcome data sparseness.
- Tags indicate an area exceeding the bounds of a location.

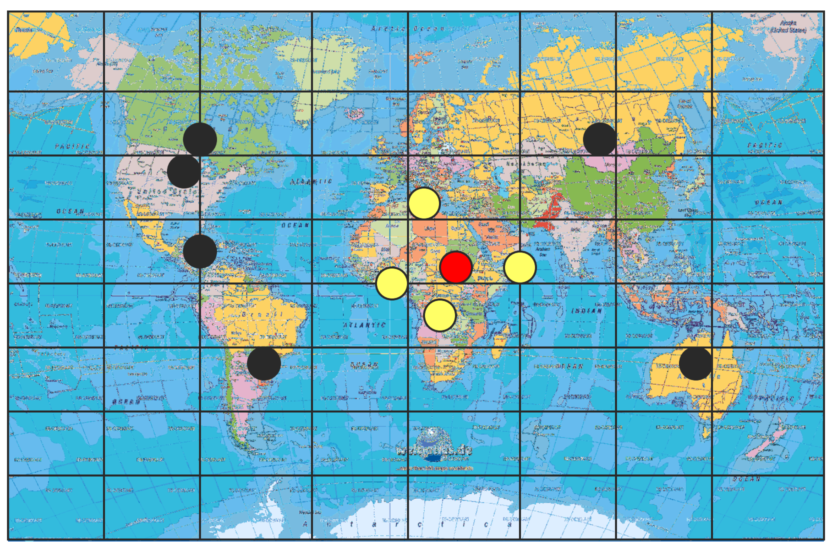

1.Tag-based smoothing with neighbors (con't)

- Considering that each tag is generated by
  - location's language model,
  - or by language models of neighboring locations:
- NB(L) :all locations included in the neighborhood of location L.

$$P(t|L) = \mu \frac{|L| \cdot P(t|L)_{ML}}{|L| + \lambda} + (1-\mu) P(t|NB(L)) + \frac{\lambda \cdot P(t|G)_{ML}}{|L| + \lambda}$$

$$P(t|NB(L)) = \sum_{L' \in NB(L)} \frac{|L'|}{|L'| + \lambda} \frac{P(t|L')_{ML}}{(2d+1)^2 - 1}$$

2.Smoothing cell relevance probabilities

- “Good” locations come from “good” neighborhoods.
- Some relevance should be propagated through the links between close locations.
- The tag set probability augmented with the probabilities of neighboring locations:

$$P(T|L) = \alpha P(T|L) + (1-\alpha) \sum_{L' \in NB(L)} \frac{P(T|L')}{(2d+1)^2 - 1}$$

Arctic Ocean
PACIFIC OCEAN
ATLANTIC OCEAN
INDIAN OCEAN
Canada
United States
Mexico
Brazil
Argentina
Russia
Kazakhstan
Mongolia
China
India
Japan
Australia
Antarctica
Greenland (Dk.)
Alaska (United States)
welt-atlas.de
...more than 850 maps worldwide

Propagation in direction of higher relevance

- Propagating relevance from locations having lower scores than the location to be smoothed.

- Support documents with enough probability to be relevant, rather than relevant documents support poor ones.

- Edges between cells become directed (from lower scored to higher scored cells)

3.Boosting geo-related tags

- Some tags are more popular near certain locations even without analysis of their spatial distribution.

$$P(t|L)_{ML} = P(t|L)_{ML}(1 + \beta P(Loc|t))/Z$$

- β: boosting coefficient
- P(Loc|t) : probability of the tag t to be location-specific.
- Z : normalisation coefficient.

4.Spatial ambiguity-aware smoothing

- Some tags are specific for more than one location:
  - Their scope exceeds the bounds of a single cell.
  - or due to their ambiguity (for example bath and Bath, UK).
  - or because they are instances that are typically spotted at a few specific locations, such as Elephants .

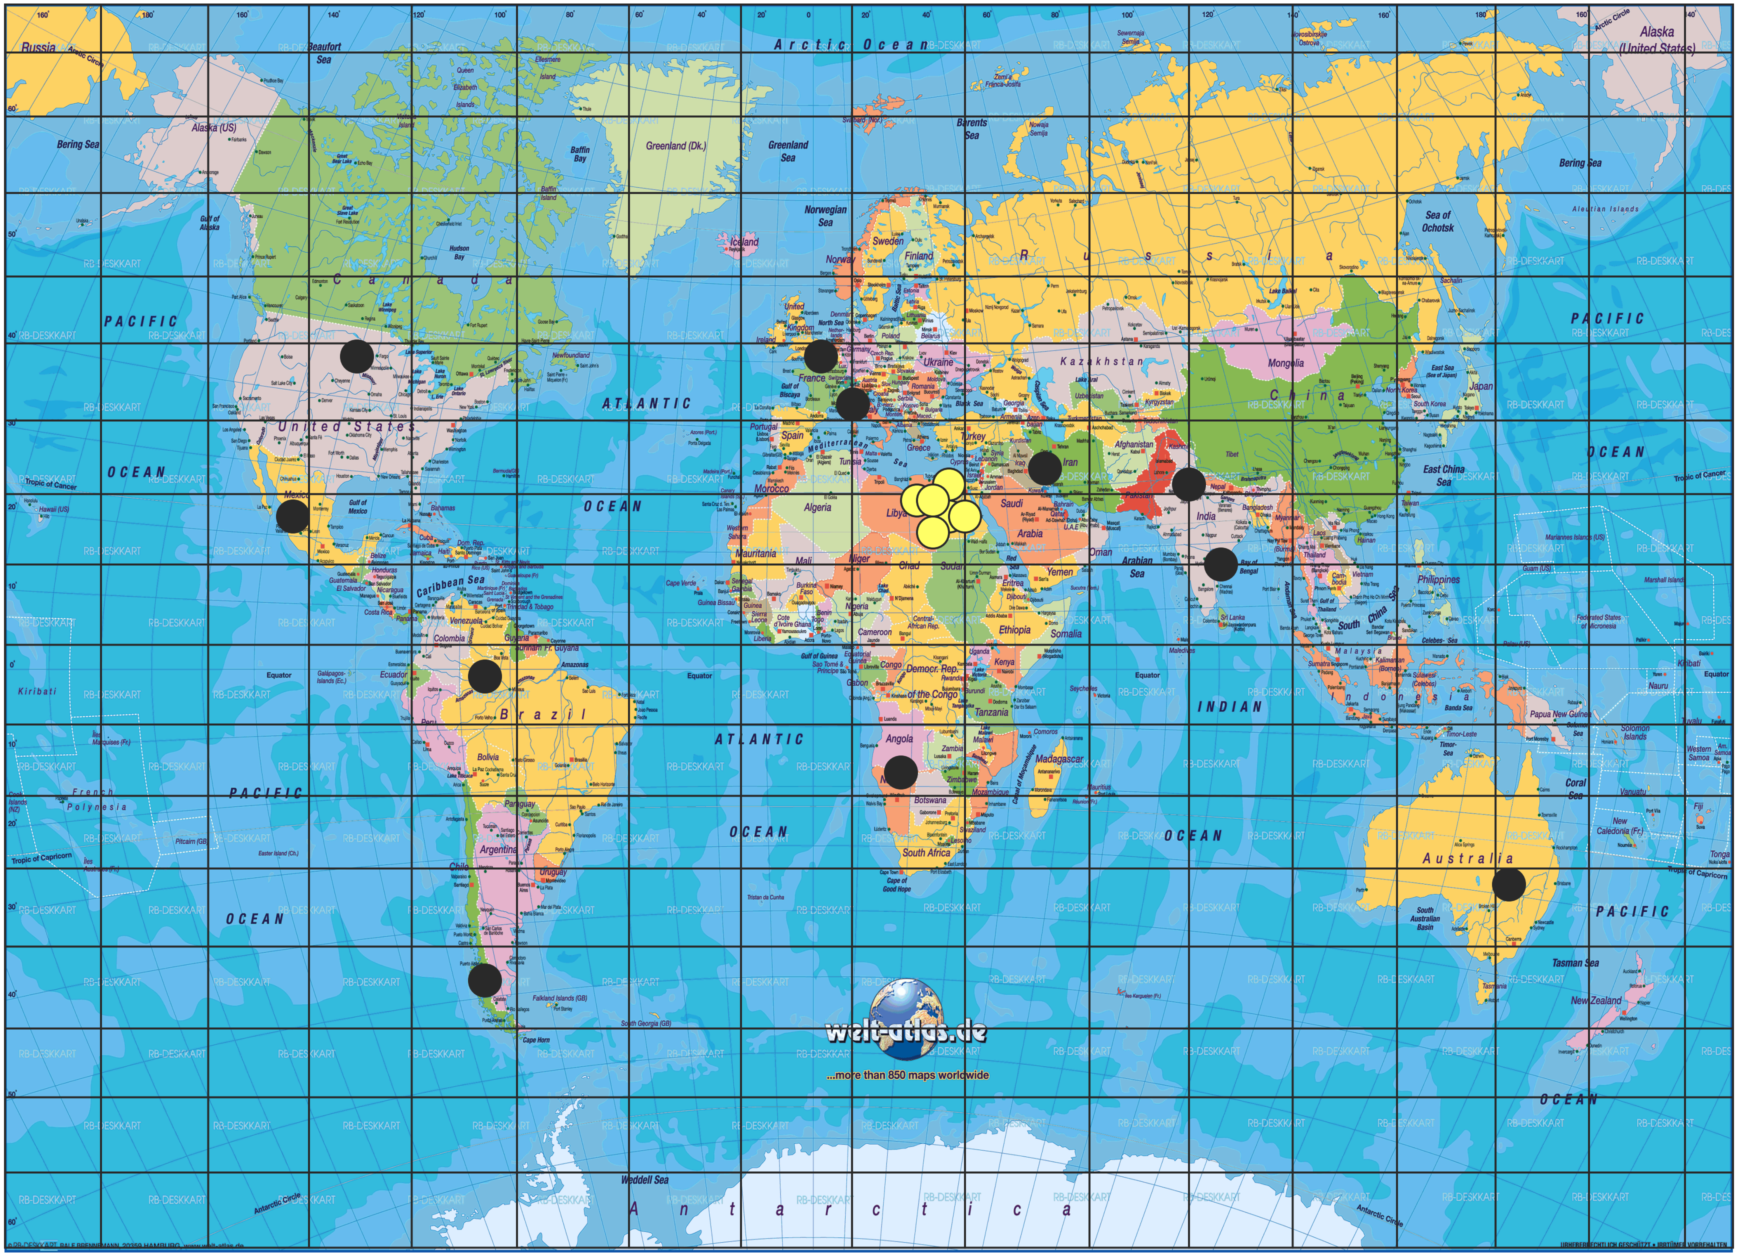
welt-atlas.de
...more than 850 maps worldwide
Arctic Ocean
PACIFIC OCEAN
ATLANTIC OCEAN
INDIAN OCEAN
Antarctica

4.Spatial ambiguity-aware smoothing (con't)

- The smoothing coefficient in Equation 3 will be tag-specific and proportional to the ambiguity of a tag:

$$\lambda(t) = \lambda + \gamma(\sigma_{lat}(t) + \sigma_{lon}(t))$$

- γ: weight coefficient to control the influence of ambiguity level on smoothing.

- Standard deviation of its latitudes and longitudes

Outline

- Introduction
- Representing Locations on the Map
- Modeling Locations
- Experimental Setup
- Results
- Conclusion

Experimental Setup

- The Flickr data set
  - A randomly sampled set of 397,000 Flickr images with the associated tags.
  - A filter applied to remove the effect of bulk uploads by users.
    - This reduces the set of photos to 140,000.
  - For better understanding the data, all images are geo-referenced with the GeoNames gazetteer.
  - Overall, the collection contains photos from about 180 different countries .

Experimental Setup

- Evaluation measures
  - (**Acc**) the percentage of correct predictions over all test examples.
  - (**MRR)** the ability of the system to find the actual location of a photo among its top recommendations.
    - $1/N \sum 1 / r$
      - r : first rank with correct position.
      - N : number of locations.

Experimental Setup

- (**Acc@K**) whether the actual location is within a K-cell distance from the predicted location.
  - Cell-based distance

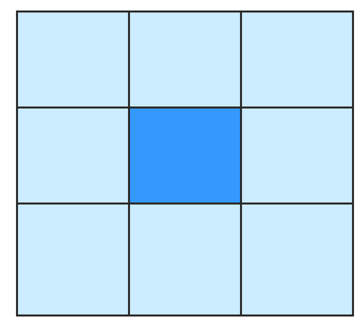

- (**PAcc**) whether the predicted location belongs to the same parent with the correct location (for instance, 100 km cells are parents for 50 km cells, 50 km cells - for 10 km cells, etc.).

Location belongs to the same parent

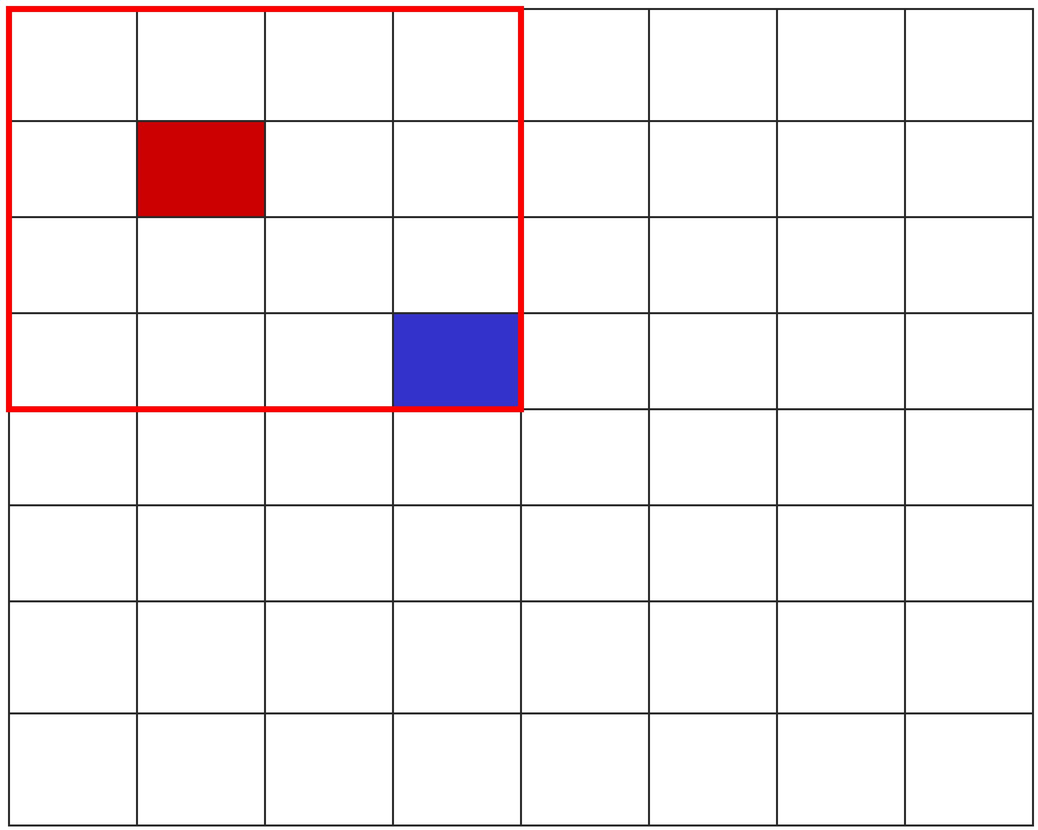

Outline

- Introduction
- Representing Locations on the Map
- Modeling Locations
- Experimental Setup
- **Results**
- Error analysis
- Conclusion

Results

- Baseline language model **LM**
- Tag-based smoothing **TS (1)**
- Cell-based smoothing **CS (2)**
- Cell based smoothing with score propagation in the direction of higher relevance **CSR (2')**
- Toponym based boosting **TB (3)**
- Ambiguity-aware tag specific smoothing **AS (4)**

Performance of the baseline LM method

Division	Acc	MRR	Acc@1	Acc@2	Acc@3	PAcc
1km	0.067	0.073	0.125	0.152	0.170	0.122
5km	0.140	0.155	0.226	0.248	0.259	0.177
10km	0.181	0.197	0.261	0.278	0.291	0.247
50km	0.256	0.277	0.332	0.354	0.378	0.289
100km	0.288	0.309	0.370	0.410	0.435	-

Results

- The accuracy increases,
  - when increasing the grid size, from 0.067 to 0.288.
- Additional performance improvement is observed
  - when analyzing the relaxed accuracy measures to include the direct neighbors of the predicted location.

Performance of neighborhood smoothing

Method	Acc	MRR	Acc@1	Acc@2	Acc@3	PAcc
	1 km					
LM	0.067	0.073	0.125	0.152	0.170	0.122
+TS	0.068	0.074	0.128	0.160	0.180	0.129
+CS	0.066	0.073	0.13	0.158	0.179	0.126
+CSR	**0.070**	**0.075**	**0.141**	**0.176**	**0.197**	**0.140**
	10 km					
LM	0.181	0.197	0.261	0.278	0.291	0.247
+TS	0.181	0.197	0.260	0.278	0.291	0.245
+CS	0.183	0.195	0.266	0.285	0.297	0.252
+CSR	**0.187**	**0.201**	**0.271**	**0.288**	**0.301**	**0.255**
	100 km					
LM	0.288	0.309	0.370	0.410	0.435	-
+TS	0.290	0.311	0.371	0.409	0.437	-
+CS	0.289	0.310	0.387	0.430	0.456	-
+CSR	**0.296**	**0.314**	**0.390**	**0.443**	**0.470**	-

Results

- Marginal improvements,
  - CSR method outperforming the other two smoothing extensions independent of grid size, 1, 10, and 100 km.
- Smoothing was only done with the immediate neighbors (d = 1)..

Performance of combinations of methods

Method	Acc	MRR	Acc@1	Acc@2	Acc@3
	10 km				
LM	0.181	0.197	0.261	0.278	0.291
+CSR	0.187	0.201	0.271	0.288	0.301
+TB	0.198	0.209	0.283	0.303	0.316
+TB+CSR	0.198	0.210	0.286	0.305	0.319
+AS	0.190	0.205	0.275	0.292	0.306
+AS+TB	0.204	0.213	0.295	0.314	0.329
+AS+CSR	0.194	0.206	0.285	0.303	0.317
+AS+TB+CSR	**0.204** (+13%)	**0.213** (+8%)	**0.297** (+14%)	**0.316** (+14%)	**0.332**(+14%)

Results

- All methods improve over baseline LM for all measures.
- The maximum performance is reached by using all three methods together.
- The results for accuracy-based measures were not tested for statistical significance.
  - They have a binary outcome (correct or incorrect).

Results

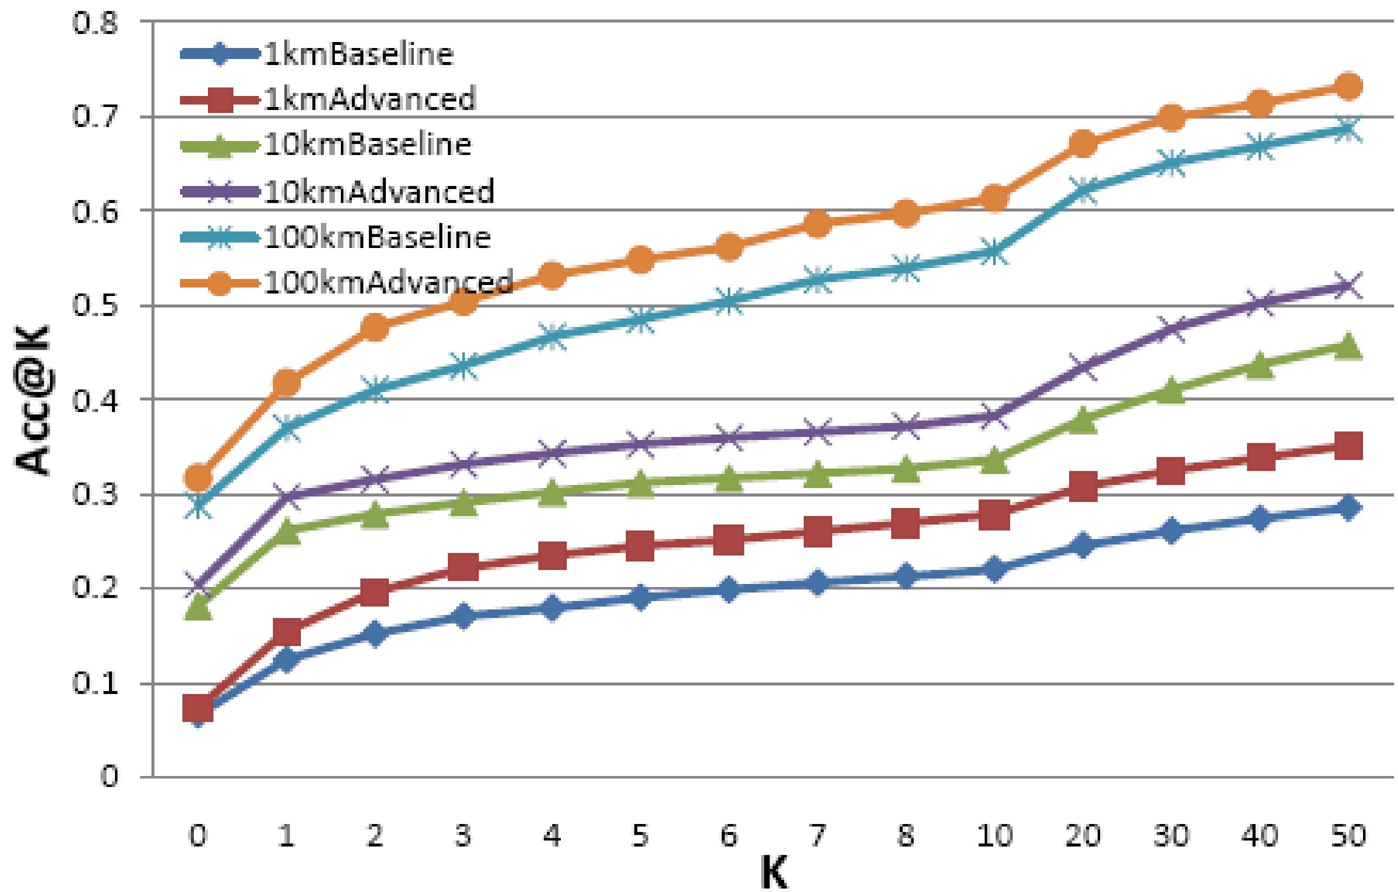

1kmBaseline
1kmAdvanced
10kmBaseline
10kmAdvanced
100kmBaseline
100kmAdvanced
Acc@K
K
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0
1
2
3
4
5
6
7
8
10
20
30
40
50

Results

- High relaxed accuracies with K up to 50.
- For regions with larger bounds the performance of the baseline method increases accordingly.
- The benefit from the advanced methods stays the same ,probably avoiding especially coarse errors.

Outline

- Introduction
- Representing Locations on the Map
- Modeling Locations
- Experimental Setup
- Results
- Error analysis
- Conclusion

Error analysis

- There are two main sources of errors:
  - caused by sparsity and noisiness
  - arising from ambiguity and incompleteness of the tag sets.
- In the first case, the right location is
  - not represented in the data,
  - or poorly represented with tags specific for this location only (e.g. containing no toponyms).

Error analysis

- For the second case, images difficult to localize :
  - (1) images with tags specific to too many locations (e.g. beach coast rocks lovers);
  - (2) images with toponyms, but with no tags disambiguating them (e.g. michigan cats dogs);
  - (3) images with a tag falsely indicating to a location (e.g. paris hilton picturing a poster in New York);
  - (4) images containing a tag specific to a region larger than a chosen grid cell size (e.g. Alaska snow for 100 km cells).

Error analysis

- The first three types of errors can be eliminated in the future by taking some contextual or user-specific evidence:
  - for instance, tags of recently uploaded images or the location of user IP.
- Highly ambiguous tag sets may be successfully mapped by relying on the history of user locations,
  - since such tag sets might be location-specific on the personal level (e.g., people celebrating their birthday in their home location).

Error analysis

- Resolved with additional image content analysis.


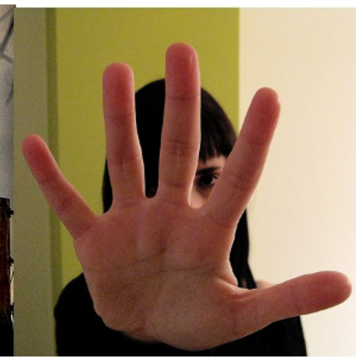

- Errors resulting from images containing tags specific to a larger region are more difficult to avoid.

Error analysis

- Photos taken in very popular tourist destinations, are represented better among correct mappings than in the entire data set,
  - tourists almost always describe their photos with location specific tags.

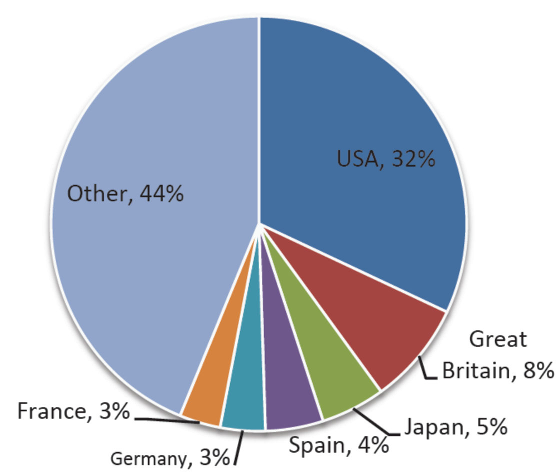

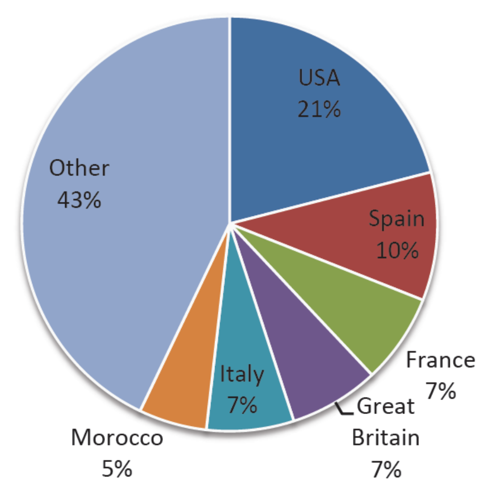

Outline

- Introduction
- Representing Locations on the Map
- Modeling Locations
- Experimental Setup
- Results
- Error analysis
- Conclusion

Conclusions

- A generic methods for automatically placing photos uploaded in Flickr on the world map.

- Estimate a language model through analysis of the terms people use to describe images taken at a particular location.

- Increasing the accuracy of the predictions
  - by incorporating ambiguity-aware smoothing,
  - cell-based smoothing with score propagation in the direction of highly relevant neighbors.

Improving suggestions

- First, to define automatically and appropriate grid division for a tag set.
- Minimize interactions between users and the system
  - by showing a map view at the optimal zoom level
- Second, study the utility of additional evidence coming from a user profile,
  - uploads history, social network or IP address.
- Finally, images used to build location models can be distinguished
  - by using common (e.g. noise ratio) or Flickr-specific (number of views, interestingness) quality measures.

Thank you!

- Pavel Serdyukov, Vanessa Murdock, and Roelof van Zwol. Placing Flickr Photos on a Map. SIGIR Conference 2009, pp. 484-491.