GeoWINE: Geolocation based Wiki, Image, News and Event Retrieval

Golsa Tahmasebzadeh TIB – Leibniz Information Centre for Science and Technology Hannover, Germany golsa.tahmasebzadeh@tib.eu

Sherzod Hakimov TIB – Leibniz Information Centre for Science and Technology Hannover, Germany sherzod.hakimov@tib.eu Endri Kacupaj University of Bonn Bonn, Germany kacupaj@cs.uni-bonn.de

Jens Lehmann University of Bonn Fraunhofer IAIS Dresden Bonn, Germany jens.lehmann@cs.uni-bonn.de

Eric Müller-Budack TIB – Leibniz Information Centre for Science and Technology Hannover, Germany eric.mueller@tib.eu

Ralph Ewerth TIB Hannover & L3S Research Center, Leibniz University Hannover Hannover, Germany ralph.ewerth@tib.eu

ABSTRACT

In the context of social media, geolocation inference on news or events has become a very important task. In this paper, we present the GeoWINE (Geolocation-based Wiki-Image-News-Event retrieval) demonstrator, an effective modular system for multimodal retrieval which expects only a single image as input. The GeoWINE system consists of five modules in order to retrieve related information from various sources. The first module is a state-of-the-art model for geolocation estimation of images. The second module performs a geospatial-based query for entity retrieval using the Wikidata knowledge graph. The third module exploits four different image embedding representations, which are used to retrieve most similar entities compared to the input image. The last two modules perform news and event retrieval from EventRegistry and the Open Event Knowledge Graph (OEKG). GeoWINE provides an intuitive interface for end-users and is insightful for experts for reconfiguration to individual setups. The GeoWINE achieves promising results in entity label prediction for images on Google Landmarks dataset. The demonstrator is publicly available at http://cleopatra.ijs.si/geowine/.

CCS CONCEPTS

• Computing methodologies \rightarrow Neural networks; • Information systems \rightarrow Information retrieval.

KEYWORDS

Geolocation Estimation; Computer Vision; Knowledge Graph

ACM Reference Format:

Golsa Tahmasebzadeh, Endri Kacupaj, Eric Müller-Budack, Sherzod Hakimov, Jens Lehmann, and Ralph Ewerth. 2021. GeoWINE: Geolocation based

SIGIR '21, July 11–15, 2021, Virtual Event, Canada

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8037-9/21/07...\$15.00 https://doi.org/10.1145/3404835.3462786 **1** INTRODUCTION

Motivation. With the rapid growth of technology, billions of images are appearing on the Internet every day. While surfing the Web one may encounter an image (e.g., depicting a place) and would like to know where it belongs to and what are similar places. Nonetheless, it would be hard to find keywords to query a search engine. With this in mind, reverse-image search is a unique way to browse the Web. Some of existing reverse-image search implementations are Google Images, eBay, Bing and Pinterest most of which are mainly based on extracting visual features from the query image or meta information to perform the retrieval task [1]. In this regard, content-based location estimation of images would enhance the retrieval task. There is an increasing research interest in exploiting geographical content of images for various tasks like image classification [11], information retrieval [2], and image verification [3]. On the other hand, information retrieval, text and Web mining, as well as knowledge graphs are increasingly intertwined with the purpose of yielding opportunities for enhanced exploration facilities.

Wiki, Image, News and Event Retrieval. In Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information

Retrieval (SIGIR '21), July 11-15, 2021, Virtual Event, Canada. ACM, New

York, NY, USA, 5 pages. https://doi.org/10.1145/3404835.3462786

In this paper, we propose GeoWINE (Geolocation Wiki-Image-News-Events) to close the gaps between geolocation estimation, information representation in knowledge graphs, and information retrieval. We utilize geolocation estimation as a core task, but also go beyond by including contextual information in order to enrich the information retrieval task. In this regard, the proposed system is an extension of geolocation estimation to highlight its potential by linking the input image to other knowledge sources such as Wikidata [13] and Open Event Knowledge Graph (OEKG) [5]. To the best of our knowledge, this is the first attempt based on image-based geolocation estimation for event and news retrieval in the context of more downstream tasks such as image verification, recommending similar places, or fact-checking.

Approach and Contribution. The proposed system is a geolocation based multimodal retrieval system that comprises five

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.



Figure 1: Overview of the GeoWINE architecture.

modules. Given an image as input, GeoWINE applies a state-ofthe-art geolocation estimation model as a starting point to retrieve data from Wikidata [13], EventRegistry¹, and OEKG [5]. The geolocation estimation model is responsible for the prediction of the input image coordinates. The second module performs a geospatial query on Wikidata to retrieve all entities of specific types located no farther than a specified radius from the predicted coordinates. Here, the entity types and the radius are given as input to our system. The third module leverages three different image embedding representations that are derived from the tasks of geolocation estimation, place recognition, and an ImageNet model for image classification. These embeddings are used to rank the most similar entities compared to the input image. Finally, the last two modules retrieve similar news and events from EventRegistry and OEKG.

We evaluate GeoWINE on Google Landmarks dataset [14], where it achieves promising performance in predicting entity labels of query images. GeoWINE enables users to retrieve entities, news, and events related to an image, through a clean and intuitive User Interface (UI), with interactive response times. To the best of our knowledge, this is the first public and open-source demo for geolocationbased multimodal retrieval through various sources. To facilitate reproducibility and reuse, all material is publicly available².

Use cases. The proposed system provides various search scenarios each of which could be addressed as a use case. One of the use cases is to assist journalists or social scientists in fact-checking or fake news detection [9, 10]. Given a query image as an input, the system could be used to contextualize the image by retrieving relevant documents. First, the geolocation is predicted and then the predicted coordinates are used to query Wikidata [13] to retrieve entities within the selected radius and entity types by the user. In this way, the user can select an entity on the map to get back the relevant news and events in order to check whether the input image is relevant to the retrieved documents or their corresponding images. Another useful application is the verification of images based on

¹http://eventregistry.org/ ²https://github.com/cleopatra-itn/GeoWINE the retrieved entities. For instance, given an image of a scenery, the system recommends a ranked list of visually similar places within a selected radius in order to verify the label of the input image. This application might be interesting for tourists. From another perspective the system could also be useful in meta information verification of news. For instance to see if the location claimed of a news matches the location of its image.

2 METHOD

In this section, we describe the five modules of the GeoWINE system. The basic functionality is shown in Figure 1.

2.1 Geolocation Estimation Module

The goal of geolocation estimation is to predict latitude and longitude coordinates where photo was captured, using only the visual content. We apply the model³ [8] based on ResNet101 [7] pretrained on a subset of the Yahoo Flickr Creative Commons 100 Million dataset (YFCC100M) [12] with around five million geotagged images. In this model, geolocation estimation is treated as a classification problem where the earth is subdivided into geographical cells. This approach exploits the hierarchical representation of multiple partitionings by taking into account the photo's scene category, i.e., whether it depicts indoor, natural, or urban setting.

2.2 Geospatial-based Entity Retrieval Module

The primary knowledge source for our system is the Wikidata knowledge graph [13] which is a free and open knowledge base that can be read and edited by both humans and machines. Since Wikidata is a rich source for data with semantic relations between entities and gives continuous access using public SPARQL (Simple Protocol and RDF Query Language) endpoint, it meets our requirements for entity retrieval.

2.2.1 *Determining Common Entity Types.* In order to determine entity types for the retrieval task, we use an index set from the

³https://github.com/TIBHannover/GeoEstimation/tree/pytorch

Google Landmarks dataset [14]. Each image in the index set is linked with the corresponding Wikimedia Commons category, which is also linked with a Wikidata entity. We query the types of Wikidata entities using the *instance of* (P31) of each image and count their occurrences. The most frequent 35 entity types are selected and clustered in 12 groups where each group contains similar entity types. For instance, *Religious building* contains entity types such as *church building, cathedral, Buddhist temple, shrine*, etc. The full list of entity groups is given in Table 1.

2.2.2 *Querying Entities from Wikidata.* The geolocation estimation module (Section 2.1) outputs the geocoordinates for an input image. In this step, we query Wikidata to retrieve entities of the specified type and located in the selected radius (in kilometers) of the predicted geocoordinates. The user specifies the radius and selects the entity types of interest from the available 12 groups.

2.2.3 *Retrieval of Metadata and Images.* Eventually, corresponding text descriptions and images are downloaded for the retrieved entities. We extract Wikimedia Commons links of the images and the Wikipedia link for each queried entity using MediaWiki and Wikidata query APIs (Application Programming Interface)⁴.

2.3 Visual Search Embeddings Module

To generate the visual representations, state-of-the-art approaches in computer vision are applied. The *Geolocation* embedding is taken from the same model used for geolocation estimation [8] (Section 2.1). The *Place* embedding is another representation based on a ResNet model [7] to recognize 365 distinct places like *beach* or *street*. Another embedding applied here is *Object* embedding based on a ResNet model pre-trained on the ImageNet dataset for image classification [6]. Each of the embeddings is encoded in a vector of size 2048. Eventually, the *Combined* is resulted from concatenation of all three to get a more general visual representation. These representations are employed to rank the retrieved entities (Section 2.2.2) based on similarity to the input query image. The similarity between the entity image and the query image is computed using the cosine similarity of their corresponding image embeddings.

2.4 News and Event Retrieval Modules

Given the query image, the user can select the corresponding news articles or events for the retrieved entities within the defined radius.

We use EventRegistry to retrieve news for each entity, which is a repository of news in different domains, locations and languages. To get relevant news for each selected entity, its Wikidata label is used as a keyword to query the relevant news.

Regarding the events, we use OEKG [5], which is a multilingual, event-centric, temporal knowledge graph composed of different datasets linked to Event-centric Knowledge Graph (EventKG) [4]. These datasets come from multiple application domains such as question answering, entity recommendation, and named entity recognition. We use the OEKG endpoint⁵ to query the relevant events from the knowledge graph for the selected Wikidata entity.

Set
S

Figure 2: Snapshots of the demonstration on a sample image.

3 SYSTEM OVERVIEW

An overview of the system is shown in Figure 1. The demonstrator consists of a frontend and a Python Flask server as the backend.

3.1 Frontend

The frontend was built using React JS (JavaScript), an open-source JavaScript library for user interface components. The page comprises four main panels: 1) the panel to provide the query image (either by selecting an existing one or uploading new images) alongside the entity type and radius, 2) the map that displays the *entity results*, 3) the panel that provides entity details, news and events, and 4) the ordered entities based on the similarity to the input image. Once the user presses the search button, the provided image, the selected entity type, and the defined radius are sent to the backend.

3.2 Backend

In the backend, the request is sent to a Python Flask application in JSON (JavaScript Object Notation) format. Then, the application forwards the request to an initialized instance of GeoWINE, which computes the results using the pipeline components specified in Section 2. Finally, the Flask application sends the computed results back to the frontend in JSON format.

4 DEMO WALKTHROUGH

In this section, we provide an example query in Figure 2 to guide the reader and user through the demo⁶.

Users can upload an image or select from the provided list of images, which are a subset of evaluation test set (Section 5). Then they specify a radius in kilometers, and select entity types (explained in Section 2.2.1) from the drop-down list and then get the results as illustrated in the top left box. Given these inputs, the geocoordinates are predicted and pinpointed on the map (middle box). In addition to the predicted location (the green marker), the corresponding

⁴https://www.mediawiki.org/wiki/API:Main_page ⁵http://oekg.l3s.uni-hannover.de/sparql

⁶http://cleopatra.ijs.si/geowine/

Table 1: The accuracy scores at *k* ranked results (in [%]) and the number of instances for different input entity groups: Bridge (Br), Historic Site (Hi), Square (Sq), Castle (Ca), Monument (Mo), Museum (Mu), Building (Bu), Religious Building (Rb), Tower (To), Tourist Attraction (Ta), Waterfall (Wa), Skyscraper (Sk).

in [%]	Br	Hi	Sq	Ca	Мо	Mu	Bu	Rb	То	Та	Wa	Sk	overall
Top 10	90	79	73	86	72	67	86	50	77	81	90	65	84
Top 5	74	70	64	77	68	60	86	50	77	81	85	55	76
Top 1	53	52	27	50	56	40	57	13	46	68	75	35	52
# instances per entity group	19	33	33	22	25	30	14	8	13	31	20	40	275

Wikidata entities are represented on the map within the selected radius (blue markers). For the pre-selected images the ground truth geocoordinates are available so that the user can compare the performance of the system. In this example, the location is predicted correctly, and the label is *Notre-Dame Paris* (top right).

In addition to the aforementioned outputs, a list of images is presented at the bottom. These images correspond to the retrieved entities on the map and they are ranked based on similarity to the input image. There are four different options here: similar geolocations/entities, similar places, overall similarity by ImageNet embeddings, and combined search. These four different ranked lists correspond to the four different visual representations: *Geolocation, Place, Object* and *Combined* respectively (Section 2.3). As shown in Figure 2 bottom box, the top-1 similar geolocation is *Notre-Dame Paris.* Eventually, the user can select one of the retrieved entities on the map and see the relevant news and events (top right box).

5 EVALUATION

In this section, we present an evaluation of the presented demonstrator on a subset of the Google Landmarks dataset.

As mentioned in Section 2.2.1, we identified the most frequent entity types and formed 12 distinct groups by clustering the types of similar imagery. Then, we select one type from each group and download the images, which yields 13280 images in total. Given the Wikimedia Commons image as the query, the goal is to see if the system is able to retrieve the corresponding Wikidata [13] entity. Initially, we run the geolocation estimation model on the query images and discard query images whose estimated geocoordinate errors are above city level (25 kilometres) based on Great Circle Distance (GCD) in comparison with the ground truth geocoordinates. From the remaining 1285 images, we extract geolocation embeddings and rank the retrieved entities based on similarity to the input image. For the purpose of a fair evaluation, we discarded all the queries where the image was the same as the Wikidata image by comparing hash codes of the query and the retrieved images.

The remaining 275 query images are our test set, for which the top-*k* accuracy is given in Table 1 for different input entity groups. As the Table shows, 84%, 76%, and 52% of the test queries are ranked among the top 10, top 5, and top 1, respectively. Among the top-10 results, *Waterfall* and *Bridge*, among the top-5 *Building* and among top-1 *Waterfall* entity groups have the highest accuracy. In all the top-k results, *Religious building* achieves the lowest accuracy. Overall, the results show the feasibility of the whole system for different types of landmarks.

6 CONCLUSIONS

In this paper, we have introduced the GeoWINE demonstrator, a geolocation-based multimodal retrieval system aimed at supporting different applications such as fact-checking, fake news detection, or image verification. The system consists of five modules: geolocation estimation, geospatial-based entity retrieval, visual search embeddings, news and event retrieval. The user can provide an image, radius, and types of results (e.g. tourist attraction) as input and receive the image geocoordinates on the map, surrounding entities from Wikidata, and visually similar entities. Furthermore, the user can find some relevant news and events based on a selected entity on the map. In order to evaluate the system, we have used a small subset of the Google Landmark dataset for labeling images. The experiments have shown promising results affirming the potential of the overall system for the applications mentioned above.

In future work, we plan to use additional resources for news and event retrieval. Also, we intend to include further entity types and input options for the user, e.g., textual descriptions or keywords, to improve the retrieval task.

ACKNOWLEDGMENTS

This work has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 812997 (CLEOPATRA ITN).

REFERENCES

- Hanaa Al-Lohibi, Tahani Alkhamisi, Maha Assagran, Amal Aljohani, and Asia Othaman Aljahdali. 2020. Awjedni: A Reverse-Image-Search Application. ADCAIJ: Advances in Distributed Computing and Artificial Intelligence Journal 9, 3 (2020), 49–68.
- [2] Jason Armitage, Endri Kacupaj, Golsa Tahmasebzadeh, Swati, Maria Maleshkova, Ralph Ewerth, and Jens Lehmann. 2020. MLM: A Benchmark Dataset for Multitask Learning with Multiple Languages and Modalities. In CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020. ACM, 2967–2974. https://doi.org/10.1145/ 3340531.3412783
- [3] Jiaxin Cheng, Yue Wu, Wael AbdAlmageed, and Premkumar Natarajan. 2019. QATM: Quality-Aware Template Matching for Deep Learning. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019. Computer Vision Foundation / IEEE, 11553-11562. https: //doi.org/10.1109/CVPR.2019.01182
- [4] Simon Gottschalk and Elena Demidova. 2019. EventKG the hub of event knowledge on the web - and biographical timeline generation. Semantic Web 10, 6 (2019), 1039–1070. https://doi.org/10.3233/SW-190355
- [5] Simon Gottschalk, Endri Kacupaj, Šara Abdollahi, Diego Alves, Gabriel Amaral, Elisavet Koutsiana, Tin Kuculo, Daniela Major, Caio Mello, Gullal S. Cheema, Abdul Sittar, Swati, Golsa Tahmasebzadeh, and Gaurish Thakkar. 2021. OEKG: The Open Event Knowledge Graph. In Proceedings of the 2nd International Workshop on Cross-lingual Event-centric Open Analytics co-located with the 30th The Web Conference (WWW 2021), Ljubljana, Slovenia, April 12, 2021 (online event due to COVID-19 outbreak) (CEUR Workshop Proceedings, Vol. 2829). CEUR-WS.org, 61–75. http://ceur-ws.org/Vol-2829/paper5.pdf

- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016. IEEE Computer Society, 770–778. https://doi.org/10.1109/CVPR.2016.90
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Identity Mappings in Deep Residual Networks. In Computer Vision ECCV 2016 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part IV (Lecture Notes in Computer Science, Vol. 9908). Springer, 630–645. https://doi.org/10.1007/978-3-319-46493-0_38
- [8] Eric Müller-Budack, Kader Pustu-Iren, and Ralph Ewerth. 2018. Geolocation Estimation of Photos Using a Hierarchical Model and Scene Classification. In Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XII (Lecture Notes in Computer Science, Vol. 11216). Springer, 575-592. https://doi.org/10.1007/978-3-030-01258-2_35
- [9] Eric Müller-Budack, Jonas Theiner, Sebastian Diering, Maximilian Idahl, and Ralph Ewerth. 2020. Multimodal Analytics for Real-world News using Measures of Cross-modal Entity Consistency. In Proceedings of the 2020 on International Conference on Multimedia Retrieval, ICMR 2020, Dublin, Ireland, June 8-11, 2020. ACM, 16-25. https://doi.org/10.1145/3372278.3390670

- [10] Eric Müller-Budack, Jonas Theiner, Sebastian Diering, Maximilian Idahl, Sherzod Hakimov, and Ralph Ewerth. 2021. Multimodal News Analytics using Measures of Cross-modal Entity and Context Consistency. Int. J. Multim. Inf. Retr. (2021). https://doi.org/10.1007/s13735-021-00207-4
- [11] Kevin D. Tang, Manohar Paluri, Fei-Fei Li, Robert Fergus, and Lubomir D. Bourdev. 2015. Improving Image Classification with Location Context. In 2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015. IEEE Computer Society, 1008–1016. https://doi.org/10.1109/ICCV.2015.121
- [12] Bart Thomee, David A. Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. 2016. YFCC100M: the new data in multimedia research. *Commun. ACM* 59, 2 (2016), 64–73. https://doi.org/10. 1145/2812802
- [13] Denny Vrandecic and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. Commun. ACM 57, 10 (2014), 78–85. https://doi.org/10.1145/ 2629489
- [14] Tobias Weyand, André Araujo, Bingyi Cao, and Jack Sim. 2020. Google Landmarks Dataset v2 - A Large-Scale Benchmark for Instance-Level Recognition and Retrieval. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020. IEEE, 2572–2581. https://doi.org/10.1109/CVPR42600.2020.00265