

MLM: A Benchmark Dataset for Multitask Learning with Multiple Languages and Modalities

JASON ARMITAGE*

University of Bonn
Germany

ENDRI KACUPAJ*

University of Bonn
Germany

GOLSA TAHMASEBZADEH*

TIB – Leibniz Information
Center for Science and
Technology
Germany

SWATI

Jožef Stefan Institute
Slovenia

MARIA MALESHKOVA

University of Bonn
Germany

RALPH EWERTH

TIB – Leibniz Information
Center for Science and
Technology
Germany

JENS LEHMANN

University of Bonn
Fraunhofer IAIS
Germany

ABSTRACT

In this paper, we introduce the MLM (Multiple Languages and Modalities) dataset - a new resource to train and evaluate multitask systems on samples in multiple modalities and three languages. The generation process and inclusion of semantic data provide a resource that further tests the ability for multitask systems to learn relationships between entities. The dataset is designed for researchers and developers who build applications that perform multiple tasks on data encountered on the web and in digital archives. A second version of MLM provides a geo-representative subset of the data with weighted samples for countries of the European Union. We demonstrate the value of the resource in developing novel applications in the digital humanities with a motivating use case and specify a benchmark set of tasks to retrieve modalities and locate entities in the dataset. Evaluation of baseline multitask and single task systems on the full and geo-representative versions of MLM demonstrate the challenges of generalising on diverse data. In addition to the digital humanities, we expect the resource to contribute to research in multimodal representation learning, location estimation, and scene understanding.

1 INTRODUCTION

At the core of recent research in machine learning is the proposal that systems trained on large sets of data will display forms of the human ability to generalise. Learning algorithms will enable systems to infer abstract representations from raw data and tackle a wide range of tasks [7]. The web provides an abundance of data for training models that enable functions in application pipelines. Data availability then is no impediment to the aims of machine learning in relation to these applications - but there remain two important ways in which systems fail to generalise.

The first of these is the inability for systems to cope with diversity in real-world data [6, 53]. Even small changes in the

distributions of data lead to steep declines in the performance of machine learning models. Information on the web and maintained in archives is present in a diverse range of modalities including language, images, and numerical data. Variety is also observed in each modality taken in isolation. Consider text in documents stored and studied by cultural and academic institutions. A system that is designed to predict contextual information - in this case, locations referenced in the documents [1] - may encounter text written in several languages and writing systems. Additional inputs with relevant information can be stored in databases or as structured text in a knowledge base. Machine learning systems that are robust to disparities between and within modalities can make use of a higher proportion of the information present across a range of materials [5].

The second way in which current systems fail to generalise is in using representations learned during training to perform new tasks. Analysis of machine learning in application workflows reveals a host of disparate models - each trained on homogeneous data to conduct specialised tasks [41]. In this approach, learning is siloed with the result that computation and data required for training and inference multiply - and there is limited sharing of parameters and representations across pipelines. Systems that conduct multiple tasks hold the promise of simplifying application architectures and improving efficiency in the use of resources. In the case of researchers conducting cultural and historical analysis, this also presents the advantages of studying or accessing materials in multiple formats [40] such as textual documents and images that reference the same subject.

The Multiple Languages and Modalities (MLM) dataset consists of text in three languages (English, French, and German), images, location data, and knowledge graph triple classes (ie component entities of triples for the sample entity that are stored in Wikidata). The resource is designed to evaluate multitask systems in relation to single task alternatives when trained simultaneously on sequences of tasks - in this case, a system that performs both

*Denotes equal contribution to this research

cross-modal retrieval and location estimation. Multitask learning aims to benefit from commonalities in input data [13] and MLM is derived from knowledge graph entities that connect all of the samples in the data. Relationships between entities are further manifested in triple classes.

1.1 Contributions

We identify four contributions resulting from this research.

- The first contribution is the generation process, which provides a consistent and reproducible approach to build a resource of diverse data with semantic relationships. As detailed below, this process will accelerate the release of future versions of MLM.
- The second contribution is the release of a novel dataset with a range of modalities and languages that - to our knowledge - is unique in the research literature. The resource tests the ability for multitask systems to exploit both formal and indirect relationships between samples.
- The third contribution is the specification of a benchmark series of tasks to evaluate multitask systems that characterise applications used and built by researchers and developers in the digital humanities. The resource presents a unique option to this community to develop novel applications that make use of multimodal frameworks for the assessment of cultural and historical phenomena [40].
- The fourth contribution is building a multitask system that combines several methods to perform in concurrence cross-modal retrieval and location estimation on highly diverse data.

1.2 Use Case

We present a use case that applies to the benchmark evaluation specified in this research and characterises the novel applications enabled by a multitask approach to diverse data in a real-world scenario. The MLM resource and benchmark tasks also extends to use cases in several other areas highlighted in Section 2.

Relations between media and geospatial location underpin applications and projects developed for cultural and historical analysis. The Edinburgh Geoparser [1] and Perdido [37] are recent instances of applications that retrieve locations from multilingual literary texts sourced in books and journals. Identifying locations depicted in visual materials is also a core function in applications that georeference historical images on landscapes and urban areas [9, 20]. In the multitask system framework presented below (see Figure 3), we enable geoparsing and georeferencing in two novel ways. First the researcher has the option of starting a query either with a text or an image. During the retrieval task (see Figure 1), related materials in the missing modality are returned. This creates the opportunity to conduct a multimodal analysis on the entity of interest [40]. Second a multitask framework returns a predicted location for both inputs.

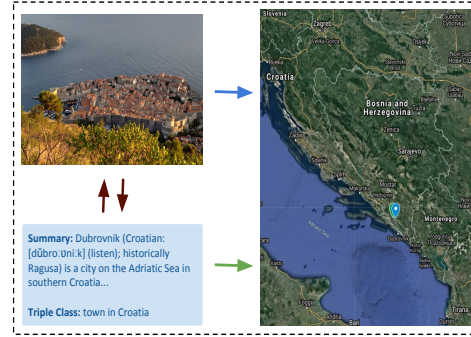


Figure 1: Cross-modal retrieval and location estimation performed by the Multitask IR+LE framework.

2 IMPACT

In this section, we highlight research areas where the MLM resource enables innovation and assess the benefits of the dataset in relation to existing resources for evaluating multitask systems.

We have outlined above new approaches in the digital humanities that are enabled by the dataset and benchmark evaluation tasks. Multitask learning systems that learn on multimodal data are also an active area of research in relation to multimodal representation learning, location estimation, and scene understanding [5, 28]. MLM is further designed to evaluate the ability for multitask systems to leverage relationships between constituent entities in data and knowledge graph properties used in the generation process. Multitask learning systems that exploit these relationships as a signal are positioned to deliver additional benefits to applications that rely on semantic data and knowledge graphs, which include recommender systems, mobile information retrieval, and bioinformatics [2].

Resources for training and evaluating multitask learning systems are characterised by homogeneity both in modalities and languages. The Office-Home and Computer Survey are examples of resources to train text-only systems [55]. Similarly multitask systems for computer vision turn to benchmark datasets where images are the primary modality [12, 32]. Multimodal datasets employed for multitask learning consist predominantly of two modalities [30, 16] and a single language. Although ImageCLEF runs challenges (see Table 1) with textual samples in multiple languages, these instances also focus on bimodal learning [49]. The Multiple Language and Modalities dataset, which we propose, approximates the diversity of online data and those stored in digital archives or databases in presenting samples in three modalities with multilingual text and semantic data. To our knowledge, this resource is unique in presenting an opportunity to train multitask systems on visual, textual, spatial and semantic web inputs for multimodal, and multilingual learning.

Geographical location is central to both the resource and task presented in this research. We are motivated by on-going investigations on representation bias in machine learning datasets to present a second geo-representative version of our resource. Representation bias has been found to penalise minority classes

Table 1: Multimodal and Multilingual Datasets

Resource	Data Types	Language	Samples
MLM	Text, Images,	DE, EN, FR	236k
	Geocoordinates		
ImageCLEF	Text, Images	DE, EN, FR	237k
IKEA	Text, Images	DE, EN, FR	4k

in classification tasks [45]. MLM-irle-gr is a subset of the full data with balanced sample sizes of entities for the 28 countries in the European Union. This enables governmental and commercial organisations that provide applications and services in this region to minimise the impact of representation bias resulting from training data when evaluating multitask learning systems.

3 MLM - MULTIPLE LANGUAGES AND MODALITIES

In this research we propose a benchmark to evaluate the strengths of multitask systems to generalise on diverse data. This section starts by defining the tasks constituting the evaluation, continues with an introduction of the dataset, and concludes with details on the generation process.

3.1 Benchmark Evaluation

The benchmark evaluation is composed of two tasks that are performed in conjunction to retrieve and locate entities represented by multimodal and multilingual inputs. The first task consists of cross-modal retrieval where a representation in an input modality x_i returns the target y_i in a corresponding modality. In this case, the inputs are one from a pair $P = \{(x_i^u, x_i^t)\}_{i=1}^n$, where x_i^u is an image and x_i^t is a combination of structured x_i^r and unstructured x_i^s textual inputs. The objective is to return the corresponding target y_i at the lowest position in the ranking of returned predictions. The objective of the second task is to use the same pair of inputs $P = \{(x_i^u, x_i^t)\}_{i=1}^n$ to perform a classification on a set of geo-cells $y = (y_1, y_2, \dots, y_n)$ - each of which groups a set of entities based on location given as geocoordinates. The final objective of the evaluation is to return results that surpass the baseline system presented below on both tasks. Additional details of the constituent tasks are included in Section 5.1.

3.2 Dataset

The Multiple Languages and Modalities comprises data points on 236k human settlements for evaluating and optimising multitask learning systems. MLM presents a dataset with a high level of diversity in terms of modality and language. For each entity, we have extracted text summaries, images, coordinates, and their respective triple classes. Text summaries are available in three languages (English, French, and German) with each entity having between one and three language entries.

Human settlements from all continents are provided in the full dataset with 72% located in Europe (see Figure 2). MLM-irle is

Table 2: MLM - Dataset Details

Num. of	MLM	MLM-irle	MLM-irle-gr
Entities	236496	218681	22501
Images	412422	314533	31621
Summaries	497899	462328	47508
Triple classes	1685	1655	452

a version of the dataset generated for the benchmark evaluation tasks and features geo-cells required for location estimation. To serve organisations that focus on the European Union, we have also created a version of the dataset - MLM-irle-gr (ie geo-representative) - that provides a geographically balanced coverage of human settlements in this region. MLM-irle-gr contains data on 24k human settlements across the EU weighted in relation to the population count¹ for each of the 28 countries.




Figure 2: MLM: Number of Entities by Continent (000).

3.3 Generation Process

In this first release, a core aim of the research was to develop a generation process that enables fast and consistent scaling for future versions. The process outlined also results from a core requirement for the resource: to evaluate the ability of multitask systems to incorporate information on entities with semantic relations. In this section, we detail stages in the framework for dataset generation. Our selection of the Wikidata Knowledge Graph² as the primary source was conditioned on the requirement for data with semantic relations between entities and the ability enabled by SPARQL to target data that met our research aims.

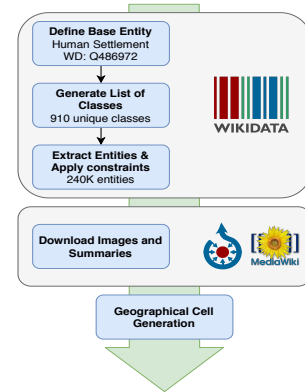


Figure 3: MLM dataset generation process.

¹ <https://population.un.org/wpp/Download/Standard/Population/>

² <https://www.wikidata.org/>

3.3.1 Define base entity type. The first step consisted of defining the base entity type that best supported developing applications that make use of geographical location. We selected the Wikidata entity human settlement (ID Q486972)³ based on the number of entities returned relative to other candidate base entity types. In the second step, instances of human settlement base type with targeted modalities were extracted resulting in over half a million entities. We applied constraints to extract entities with a Wikipedia article in at least one target language (EN, DE, FR), a minimum of one image, and a set of geocoordinates resulting in 40k human settlements.

3.3.2 Generate a list of classes. We increased the scope and granularity of human settlements by extracting entity types that are subclass of the the base type. Entities are returned for triples upto three hops from the starting entity. An example of additions is the commune of Italy (ID Q747074) - a subclass of the entity municipality (ID Q15284), which is a subclass of the human settlement base type. This stage resulted in a set of 910 unique classes where all entities derive from the base type by either one, two, or three hops.

3.3.3 Extract entities. In the extraction stage, we run queries on all entities that are an instance of at least one of the 910 unique classes in the generated list. This returned a set of 2.5M entities. A final application of the constraints related to targeted data types generated a collection of data that constituted 240k entities.

3.3.4 Download images and summaries. The final stage consists of downloading images and text summaries for the extracted entities. While querying Wikidata we were able to also extract Wikimedia Commons⁴ links of images and the corresponding Wikipedia link for each entity. Text summaries were downloaded using the MediaWiki API⁵, while images using the Wikimedia link. The process was accelerated by chunking the data and running multiple sessions. Download errors led to a small loss of entities and the final raw dataset is in Table 2.

3.3.5 Geographical cell generation. In this stage, entities along with their images and summaries are assigned to geographical cells. To this end, we first exclude all SVG images since they mostly contain maps and signs which are non-descriptive for our tasks. Next, we divide the earth into non-overlapping geographical cells C using the S2 Geometry Library⁶. In this library, initial cells are generated by projecting the earth’s surface on a cube with six sides. Based on geocoordinates of the entities, a hierarchical subdivision is applied [52], where each cell is the node of a quad-tree. Starting at the root nodes, the respective quad-tree is subdivided recursively until all cells contain a maximum of $T_{max} = 500$ entities. Finally, in order to have more uniform cells relative to the entities, all generated cells containing less than

$T_{min} = 20$ entities are removed.

4 AVAILABILITY

Versions of MLM listed in Table 2 are available for direct download and use. To support findability and sustainability, the MLM dataset is published as an on-line resource at <https://doi.org/10.5281/zenodo.3885753>. A separate page with detailed explanations and illustrations is available at <http://cleopatra.ijs.si/goal-mlm/> to promote ease-of-use. The project GitHub repository contains the complete source code for the system and generation script is available at <https://github.com/GOALCLEOPATRA/MLM>. Documentation adheres to the standards of *FAIR Data principles*⁷ with all relevant metadata specified to the research community and users. It is freely accessible under the *Creative Commons Attribution 4.0 International license*, making it reusable for almost any purpose.

4.1 Updating and Reusability

MLM is supported by a team of researchers from the University of Bonn, the Leibniz Information Center for Science and Technology, and Jožef Stefan Institute. The resource is already in use for individual projects and as a contribution to the project deliverables of the Marie Skłodowska-Curie CLEOPATRA Innovative Training Network. In addition to the steps above that make the resource available to the wider community, usage of MLM will be promoted to the network of researchers in this project. Awareness among researchers and practitioners in digital humanities will be promoted by demonstrations and presentations at domain-related events. The range of modalities and languages present in the dataset also extends its application to research on multimodal representation learning, multilingual machine learning, information retrieval, location estimation, and the Semantic Web. MLM will be supported and maintained for three years in the first instance. A second release of the dataset is already scheduled and the generation process outlined above is designed to enable rapid scaling.

5 MULTITASK LEARNING FRAMEWORK

In this section, we specify a multitask learning framework that characterises the use case above by performing cross-modal retrieval and location estimation tasks. The section starts with formulations of the tasks evaluated on MLM and concludes with details of the system and methods.

5.1 Tasks

5.1.1 Cross-modal Retrieval. We build on the approach proposed by Marin et al [35] to perform cross-modal information retrieval on visual u and textual t inputs. Embeddings v^u and v^t are learned as described below. A mapping from the visual and textual embeddings to a common space is learned using a cosine embedding loss, which minimises the distance between v^u and v^t

³ IDs are Wikidata identifiers <https://www.wikidata.org/wiki/Wikidata:Identifiers>

⁴ https://commons.wikimedia.org/wiki/Main_Page

⁵ https://www.mediawiki.org/wiki/API:Main_page

⁶ <https://code.google.com/archive/p/s2-geometry-library/>

⁷ <http://www.nature.com/articles/sdata201618/>




Figure 4: Multitask IR+LE framework.

based on cosine similarity and labels $s = \{1, -1\}$

$$L_{cos}(v^u, v^t, l) = \begin{cases} 1 - \cos(v_1^u, v_1^t), & \text{if } s = 1 \\ \max(0, \cos(v_1^u, v_1^t) - \text{margin}), & \text{if } s = -1 \end{cases} \quad (1)$$

where v^u and v^t are representations of the visual and textual inputs respectively. At inference time, an image or textual sample is presented with the aim of returning samples for the corresponding modality.

5.1.2 Location estimation. In this architecture, location estimation is regarded as a classification problem. To this end, geo-cells (described in Section 3.3.5) are regarded as classes. To formulate the location estimation task, we have q geo-cells $y = (y_1, y_2, \dots, y_n)$ for each of which there are n number of entities represented by images, summaries and triples. Each entity is labeled with its corresponding GPS coordinates in (latitude, longitude) pairs. The task is to predict y given v^u and v^t based on a cross-entropy loss:

$$L(v, y) = -v[y] + \log \left(\sum_j \exp(v[j]) \right) \quad (2)$$

5.1.3 Multitask. The multitask learning framework (see Figure 4) trains a combined pipeline to perform cross-modal retrieval and location estimation in conjunction. The system learns over the cross-entropy and cosine embedding losses with a weighted average. In order to account for the difference in magnitude between losses, log variance is calculated over the tasks and converted to log standard deviation $\log(\sigma)$. This weight is applied to the losses before computing the mean μ over the products.

5.2 System

We apply state-of-the-art embedding methods in computer vision and NLP to generate representations of the inputs. Two types of descriptors are utilized to better represent location information of the images when generating Visual embeddings v^u . One of the descriptors is taken from a ResNet [21] model pre-trained on the Places365 dataset [57], where the task is to recognize 365 distinct places - examples include beach, stadium, street. The second descriptor is taken from the model [38] based on ResNet101 [21] aimed at predicting the geolocation information of an image. This model is pretrained on a subset of the Yahoo Flickr Creative Commons 100 Million dataset (YFCC100M) [47]. The subset includes around five million geo-tagged images with geolocation labels, and was introduced for the MediaEval Placing Task 2016

(MP-16) [29]. Each descriptor results in a feature vector of size 2048, so eventually we concatenate both vectors and get a vector of size 4096 for each image. To generate textual embeddings v^t for summaries and triples we use BERT [19] to extract word vectors for all sentences in text. Since text content of news articles is long, we select maximum 500 characters at a time and extract word vectors from BERT, and average the vectors obtained from each token of the sentence.

An integrated architecture (see Figure 3) takes the visual and text embeddings to perform the two tasks in concurrence. Visual embeddings v^u are fed into a 1D convolutional layer (conv1d) with ReLU activation $\max(0, x)$ - and a linear layer with hyperbolic tangent $\tanh(u_i^l)$. An LSTM layer replaces the conv1d for text summary v^s and triple class v^r embeddings. Regularisation consists of dropout for textual embeddings and norm with a clamp $\min = 1 * 10^{-12}$ for all modalities. Linear layers with dropout receive the outputs of the norm in location estimation. In both tasks, text summary and triple classes are fused with concatenation $v^t = [v^s, v^r]$.

6 EXPERIMENTS

Experiments assess the performance of the multitask learning system in relation to single task pipelines. We align with benchmark cross-modal retrieval tasks [35] in selecting median rank (medR) as the primary metric. Median rank also takes into account both the performance of the retrieval and the position of the retrieved instance. To compute medR, a subset of 500 instances are randomly sampled from the test set. Then the experiments are repeated 10 times and report the mean results. Next, for each input image or text, items in the corresponding modality are ranked. The objective is to retrieve the sample for the target modality at the highest rank. The retrieved lists are further assessed using recall percentage at top K (R@K). This measures the percentage of queries where the matching item is ranked among the top K returned samples. The formulation of location estimation as a multiclass classification task motivates the selection of Precision (P) (Eq.3a), Recall (R) (Eq.3b) and weighted F1 score (Eq.3c) as measures of performance for the second task.

$$\text{Precision} = \frac{TP}{TP + FP} \quad (3a)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (3b)$$

$$F_1^{\text{weighted}} = \frac{1}{Q^2} \sum_{j=1}^Q Q_j \frac{2 \times P_j \times R_j}{P_j + R_j} \quad (3c)$$

In the above definitions, TP , FP and FN stand for True Positive, False Positive and False Negative respectively and Q is the number of geo-cells, and j is the j th instance of geo-cells.

6.1 System Configuration and Training

Pipelines for single task and multitask learning are trained for 100 epochs on a single GeForce GTX 1080 Ti GPU, using training/validation sets of 175k/22k and 18k/2k for MLM-irle and MLM-irle-gr respectively. Full details on the dataset splits are available at <http://cleopatra.ijs.si/goal-mlm/>. Although we trained for a number of epochs adopted for comparable systems [35], we note that validation loss was still declining for both multitask and single task systems at the end of training. For all tasks, consistent hyperparameter settings are maintained to enable comparative analysis of the results. The batch size is set at 128 with a learning rate of 0.0001 and a probability of 0.1 in dropout layers.

6.2 Results

In this section evaluation results are discussed for multitask and single task pipelines on the MLM-irle and MLM-irle-gr (ie geo-representative) datasets containing 22k and 2k test instances respectively.

6.2.1 Cross-modal Retrieval In Table 3, we report medR and Recall@K for image to text (first four columns) and text to image retrieval (the remaining columns). The multitask system attained a strong improvement on the single task equivalent in the first of these with medR at 19.9. Qualitative checks on samples consisted of extracting a random sample of 500 entities from the test set. Two examples of triple classes where the multitask system scored strongly relative to the single equivalent are “town of the United States” and “Ortsteil”. This corresponds to increases of 13% on R@5 and 20% points on R@10. Although R@1 also showed a marked increase, scores for both multitask and single task are at low levels. The multitask system also outperformed strongly on the geo-representative dataset with a 26-point improvement on medR relative to the single task pipeline. Multitask and single task systems attained stronger scores on text to image retrieval for both datasets with double the R@K values and medR at 11.9 - underling the superiority of the multitask approach in cross-modal retrieval.

6.2.2 Location Estimation A summary of performance for multitask and single task systems on location estimation is presented in Table. 4. Results are broken down by performance on predicting geo-cells using images and text consecutively. As with cross-modal retrieval, both systems return stronger scores when inputs are textual relative to visual. A notable difference with the first task is that performance improves on the geo-representative dataset. This difference is most marked in image-based estimation: the F1 score for the multitask system is 10 points higher on MLM-irle-gr relative to the larger dataset. Scores for the multitask system on both image- and text-based location estimation are

closely aligned with the single task system - indicating there are only modest gains from using the former approach on the second task.

7 RELATED RESEARCH

In the first instance, researchers designing multitask systems that learn from multimodal data turn to existing multimodal resources for benchmarking. A reliance on image captioning datasets has led to multitask learning proposals limited to paired image and text samples [8, 34]. Datasets in the medical and bioinformatics areas that have inspired multitask approaches also use these two modalities [39, 27]. Several proposals pair natural language and knowledge graph embeddings for training multitask systems [33]. In contrast to our research, tasks in this area are concentrated on applications related to NLP and knowledge graphs - and use text-only data. Multitask learning on multimodal data is an area of research in robotics [11, 44] - although the objective of training autonomous agents is distinct from the use cases for MLM. In all of the cited research, text inputs are English-only. Multilingual resources and benchmarks for training multitask architectures are motivated primarily to solve NLP tasks and are correspondingly unimodal [22, 31]. A multitask solution was proposed for the multilingual IKEA dataset with a specific objective in machine translation [58]. As with previously noted cases, this resource is bimodal and so applies to a smaller subset of constituent tasks relative to MLM.

Research on information retrieval for multimodal data has an extensive history and we focus here on resources and benchmark tasks where samples are present in multiple languages. The Cross-Language Evaluation Forum (CLEF) evolved into ImageCLEF (from 2003), GeoCLEF (from 2005) and VideoCLEF (from 2008) tracks. Resources in the LifeCLEF tracks of ImageCLEF have included geocoordinates - but the focus is primarily on images and is aimed at researchers in disciplines related to biology [23]. Information retrieval on video content has extended to multilingual scenarios where either queries, content, or both are in several languages [3]. Datasets in this area are mostly composed of comparatively small numbers of samples. Constituent modalities (ie video, audio) vary from those in MLM and place an emphasis on the spatiotemporal context of the inputs [51]. Video retrieval is also differentiated by an emphasis on domain-specific sub-tasks including shot segmentation [18] and detection of objects or scenes [51]. All of these resources exclude semantic Web elements.

Location estimation methods related to this research breakdown into three main groups [10]: natural [4, 56], city-scale [24, 14] and global [26, 48]. Several methods in the published research rely on a single modality [52, 38, 36] and we focus here on multimodal approaches. MediaEval benchmark placing datasets [29, 15] include more than five million instances with images, videos and metadata used to estimate capturing locations represented in multimedia. Kordopatis-Zilos et al. [25] build on this dataset by combining visual features with a language model based on word frequency in the text for geolocation estimation. In [50] a dataset

Table 3: Cross-modal retrieval results for MLM

		Image to Text				Text to Image			
		medR	R@1	R@5	R@10	medR	R@1	R@5	R@10
MLM-irle	Single task	43.1	0.02	0.07	0.14	32.5	0.04	0.14	0.24
	Multitask	19.9	0.05	0.20	0.34	11.9	0.11	0.34	0.47
MLM-irle-gr	Single task	53.6	0.01	0.06	0.13	43.2	0.03	0.10	0.18
	Multitask	27.9	0.03	0.14	0.25	21.8	0.07	0.23	0.35

Table 4: Location estimation results for MLM

		Image-based			Text-based		
		F1	P	R	F1	P	R
MLM-irle	Single task	0.16	0.17	0.17	0.59	0.61	0.60
	Multitask	0.15	0.17	0.16	0.58	0.61	0.59
MLM-irle-gr	Single task	0.28	0.30	0.28	0.67	0.70	0.68
	Multitask	0.25	0.27	0.26	0.67	0.69	0.68

containing image and text is collected from Wikipedia combined with around 890K satellite images and is used for interpreting satellite images. BreakingNews [42] is a resource with 100K instances in the domains of sports, politics, arts, healthcare or local news. The researchers apply a baseline that combines a Convolutional Neural Network and Recurrent Neural Networks for multimodal geolocation estimation. In addition there are other approaches that use geolocation information for image recognition [46, 17] and image understanding [54]. In [54] the authors leverage season and location context with a probabilistic framework to help improve region recognition in images. In [17] the effectiveness of using geolocation on fine-grained recognition is examined. All the aforementioned methods are based on datasets in English only. All these resources are intended for single task scenarios for the corresponding task. In contrast, this paper provides a multilingual dataset with a great number of modalities. The objective also differs in the focus on systems that perform geolocation estimation as one of a sequence of tasks.

Research on different forms of bias in data - and their impact on machine learning systems - underlines the value of resources that provide balanced coverage of their domains. Representation bias in relation to geographical location is a form of bias that has high relevance to MLM and has been identified in open source data sets used widely to train machine learning systems [43]. Representation bias results in poor classification performance on samples for classes that are under-sampled in the data. In another part of the machine learning pipeline, evaluation bias results when the test data fails to match the population of interest [45]. We are motivated by this research to address representation and evaluation bias in this resource.

8 CONCLUSION AND FUTURE WORK

In this research, we introduced the MLM resource for evaluating multitask systems on diverse data. In addition, we specified a set of benchmark tasks that characterise the specific use case of geoparsing and georeferencing information encoded in multiple modalities and languages. A baseline multitask system that makes

use of both cross-modal and knowledge fusion was specified and evaluated against single task systems. The paper details a process that generates a resource comprising linked entities. The generation process also enables consistent and rapid scaling of the resource for future releases. Finally we also ensured that the dataset meets FAIR principles and released a geo-representative version of the dataset to serve cultural and academic institutions with a focus on the European Union. We believe that the resource will improve generalisation in multitask learning systems and so streamline the pipelines of applications that learn from materials stored in digital archives. In upcoming research, we will release future versions with data in additional languages, add entries from other data sources, design systems for specific functions, and develop an application that further exploits contextual data.

9 ACKNOWLEDGMENTS

The project leading to this publication has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 812997.

REFERENCES

- [1] Beatrice Alex, Kate Byrne, Claire Grover, and Richard Tobin. Adapting the Edinburgh geoparser for historical georeferencing. *International Journal of Humanities and Arts Computing*, 9(1):15–35, 2015.
- [2] Mehdi Ali, Hajira Jabeen, Charles Tapley Hoyt, and Jens Lehmann. The KEEN universe. In *International Semantic Web Conference*, pages 3–18. Springer, 2019.
- [3] Yusuf Aytaç, Mubarak Shah, and Jiebo Luo. Utilizing semantic word similarity measures for video retrieval. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.
- [4] Georges Baatz, Olivier Saurer, Kevin Köser, and Marc Pollefeys. Large scale visual geo-localization of images in mountainous terrain. In *European conference on computer vision*, pages 517–530. Springer, 2012.
- [5] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multi-modal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2):423–443, 2018.
- [6] Tilman Becker, Edward Curry, Anja Jentzsch, and Walter Palmetschhofer. New Horizons for a Data-Driven Economy: Roadmaps and Action Plans for Technology, Businesses, Policy, and Society. In José María Cavanillas, Edward Curry, and Wolfgang Wahlster, editors, *New Horizons for a Data-Driven Economy: A Roadmap for Usage and Exploitation of Big Data in Europe*, pages 277–291. Springer International Publishing, Cham, 2016.
- [7] Yoshua Bengio. Learning Deep Architectures for AI. *Foundations and Trends in Machine Learning*, 2:71, 2009.
- [8] Alexander Binder, Wojciech Samek, Klaus-Robert Müller, and Motoaki Kawanabe. Enhanced representation and multi-task learning for image annotation. *Computer Vision and Image Understanding*, 117(5):466–478, 2013.

- [9] Nicolas Blanc, Timothée Produit, and Jens Ingensand. A semi-automatic tool to georeference historical landscape images. Technical report, PeerJ Preprints, 2018.
- [10] Jan Brejcha and Martin Čadík. State-of-the-art in visual geo-localization. *Pattern Analysis and Applications*, 20(3):613–637, 2017.
- [11] Simon Brodeur, Ethan Perez, Ankesh Anand, Florian Golemo, Luca Celotti, Florian Strub, Jean Rouat, Hugo Larochelle, and Aaron C. Courville. Home: a household multimodal environment. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Workshop Track Proceedings*. OpenReview.net, 2018.
- [12] Barbara Caputo, Henning Müller, Jesus Martinez-Gomez, Mauricio Villegas, Burak Acar, Novi Patricia, Neda Marvasti, Suzan Üsküdarlı, Roberto Paredes, Miguel Cazorla, et al. Imageclef 2014: Overview and analysis of the results. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 192–211. Springer, 2014.
- [13] Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- [14] David M Chen, Georges Baatz, Kevin Köser, Sam S Tsai, Ramakrishna Vedantham, Timo Pylvänäinen, Kimmo Roimela, Xin Chen, Jeff Bach, Marc Pollefeys, et al. City-scale landmark identification on mobile devices. In *CVPR 2011*, pages 737–744. IEEE, 2011.
- [15] Jaeyoung Choi, Claudia Hauff, Olivier Van Laere, and Bart Thomee. The placing task at MediaEval 2015. In *MediaEval 2015, Wurzen, Germany, 14-15 September 2015; Ceur Workshop Proceedings 1436, 2015*. CEUR, 2015.
- [16] Jaeyoung Choi, Martha Larson, Gerald Friedland, and Alan Hanjalic. From Intra-Modal to Inter-Modal Space: Multi-task Learning of Shared Representations for Cross-Modal Retrieval. In *2019 IEEE Fifth International Conference on Multimedia Big Data (BigMM)*, pages 1–10, Singapore, Singapore, September 2019. IEEE.
- [17] Grace Chu, Brian Potetz, Weijun Wang, Andrew Howard, Yang Song, Fernando Brucher, Thomas Leung, and Hartwig Adam. Geo-aware networks for fine-grained recognition. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [18] Bertrand Delezoide and Hervé Le Borgne. Semanticvox: A multilingual video search engine. In *Proceedings of the 6th ACM international conference on Image and video retrieval*, pages 81–84, 2007.
- [19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- [20] Mouna Harrach, Alexandre Devaux, and Mathieu Brédif. Interactive image geolocalization in an immersive web application. *International Archives of the Photogrammetry, Remote Sensing & Spatial Information Sciences*, 2019.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- [22] Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization. *arXiv preprint arXiv:2003.11080*, 2020.
- [23] Alexis Joly, Hervé Goëau, Hervé Glotin, Concetto Spampinato, Pierre Bonnet, Willem-Pier Vellinga, Julien Champ, Robert Planqué, Simone Palazzo, and Henning Müller. LifeCLEF 2016: multimedia life species identification challenges. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 286–310. Springer, 2016.
- [24] Hyo Jin Kim, Enrique Dunn, and Jan-Michael Frahm. Learned contextual feature reweighting for image geo-localization. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3251–3260. IEEE, 2017.
- [25] Giorgos Kordopatis-Zilos, Symeon Papadopoulos, and Ioannis Kompatsiaris. Geotagging text content with language models and feature mining. *Proceedings of the IEEE*, 105(10):1971–1986, 2017.
- [26] Giorgos Kordopatis-Zilos, Symeon Papadopoulos, and Yiannis Kompatsiaris. In-depth exploration of geotagging performance using sampling strategies on yfcc100m. In *Proceedings of the 2016 ACM Workshop on Multimedia COMMONS*, pages 3–10, 2016.
- [27] Tomasz Kornuta, Deepta Rajan, Chaitanya Shivade, Alexis Asseman, and Ahmet S. Ozcan. Leveraging medical visual question answering with supporting facts. In Linda Cappellato, Nicola Ferro, David E. Losada, and Henning Müller, editors, *Working Notes of CLEF 2019 - Conference and Labs of the Evaluation Forum, Lugano, Switzerland, September 9-12, 2019*, volume 2380 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2019.
- [28] Ryohei Kuga, Asako Kanezaki, Masaki Samejima, Yusuke Sugano, and Yasuyuki Matsushita. Multi-task learning using multi-modal encoder-decoder networks with shared skip connections. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 403–411, 2017.
- [29] Martha Larson, Mohammad Soleymani, Guillaume Gravier, Bogdan Ionescu, and Gareth JF Jones. The benchmarking initiative for multimedia evaluation: Mediaeval 2016. *IEEE MultiMedia*, 24(1):93–96, 2017.
- [30] Qing Li, Qingyi Tao, Shafiq Joty, Jianfei Cai, and Jiebo Luo. Vqa-e: Explaining, elaborating, and enhancing your answers for visual questions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 552–567, 2018.
- [31] Ying Lin, Shengqi Yang, Veselin Stoyanov, and Heng Ji. A multi-lingual multi-task architecture for low-resource sequence labeling. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 799–809, 2018.
- [32] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and S Yu Philip. Learning multiple tasks with multilinear relationship networks. In *Advances in neural information processing systems*, pages 1594–1603, 2017.
- [33] Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3219–3232, Brussels, Belgium, 2018. Association for Computational Linguistics.
- [34] Junyu Luo, Ying Shen, Xiang Ao, Zhou Zhao, and Min Yang. Cross-modal image-text retrieval with multitask learning. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 2309–2312, 2019.
- [35] Javier Marin, Aritro Biswas, Ferda Ofli, Nicholas Hynes, Amaia Salvador, Yusuf Aytar, Ingmar Weber, and Antonio Torralba. Recipe1m+: A dataset for learning cross-modal embeddings for cooking recipes and food images. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [36] Stuart E Middleton, Giorgos Kordopatis-Zilos, Symeon Papadopoulos, and Yiannis Kompatsiaris. Location extraction from social media: Geoparsing, location disambiguation, and geotagging. *ACM Transactions on Information Systems (TOIS)*, 36(4):1–27, 2018.
- [37] Ludovic Moncla, Mauro Gaio, Thierry Joliveau, and Yves-François Le Lay. Automated geoparsing of Paris street names in 19th century novels. In *Proceedings of the 1st ACM SIGSPATIAL Workshop on Geospatial Humanities*, pages 1–8, 2017.
- [38] Eric Muller-Budack, Kader Pustu-Iren, and Ralph Ewerth. Geolocation estimation of photos using a hierarchical model and scene classification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 563–579, 2018.
- [39] Binh D Nguyen, Thanh-Toan Do, Binh X Nguyen, Tuong Do, Erman Tjiputra, and Quang D Tran. Overcoming data limitation in medical visual question answering. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 522–530. Springer, 2019.
- [40] Luc Pauwels. A multimodal framework for analyzing websites as cultural expressions. *Journal of Computer-Mediated Communication*, 17(3):247–265, 2012.
- [41] Miguel De Prado, Jing Su, Rabia Saeed, Lorenzo Keller, Noelia Vallez, Andrew Anderson, David Gregg, Luca Benini, Tim Llewellynn, Nabil Ouerhani, Rozenn Dahyot, and Nuria Pazos. Bonseyes ai pipeline—bringing ai to you: End-to-end integration of data, algorithms, and deployment tools. *ACM Trans. Internet Things*, 1(4), August 2020.

- [42] Arnau Ramisa, Fei Yan, Francesc Moreno-Noguer, and Krystian Mikolajczyk. Breakingnews: Article annotation by image and text processing. *IEEE transactions on pattern analysis and machine intelligence*, 40(5):1072–1085, 2017.
- [43] Shreya Shankar, Yoni Halpern, Eric Breck, James Atwood, Jimbo Wilson, and D. Sculley. No classification without representation: Assessing geo-diversity issues in open data sets for the developing world. In *NIPS 2017 workshop: Machine Learning for the Developing World*, 2017.
- [44] Devendra Singh Chaplot, Lisa Lee, Ruslan Salakhutdinov, Devi Parikh, and Dhruv Batra. Embodied multimodal multitask learning. *arXiv preprint arXiv:1902.01385*, 2019.
- [45] Harini Suresh and John V Guttag. A framework for understanding unintended consequences of machine learning. *arXiv preprint arXiv:1901.10002*, 2019.
- [46] Kevin Tang, Manohar Paluri, Li Fei-Fei, Rob Fergus, and Lubomir Bourdev. Improving image classification with location context. In *Proceedings of the IEEE international conference on computer vision*, pages 1008–1016, 2015.
- [47] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016.
- [48] Michele Trevisiol, Hervé Jégou, Jonathan Delhumeau, and Guillaume Gravier. Retrieving geo-location of videos with a divide & conquer hierarchical multimodal approach. In *Proceedings of the 3rd ACM conference on International conference on multimedia retrieval*, pages 1–8, 2013.
- [49] Theodora Tsirikli, Adrian Popescu, and Jana Kludas. Overview of the wikipedia image retrieval task at imageclef 2011. In *CLEF (Notebook Papers/Labs/Workshop)*, volume 4, page 5, 2011.
- [50] Burak Uzkent, Evan Sheehan, Chenlin Meng, Zhongyi Tang, Marshall Burke, David B Lobell, and Stefano Ermon. Learning to interpret satellite images using wikipedia. In *IJCAI*, pages 3620–3626, 2019.
- [51] Xin Wang, Jiawei Wu, Junkun Chen, Lei Li, Yuan-Fang Wang, and William Yang Wang. Vatex: A large-scale, high-quality multilingual dataset for video-and-language research. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4581–4591, 2019.
- [52] Tobias Weyand, Ilya Kostrikov, and James Philbin. Planet-photo geolocation with convolutional neural networks. In *European Conference on Computer Vision*, pages 37–55. Springer, 2016.
- [53] Xue-Wen Chen and Xiaotong Lin. Big Data Deep Learning: Challenges and Perspectives. *IEEE Access*, 2:514–525, 2014.
- [54] Jie Yu and Jiebo Luo. Leveraging probabilistic season and location context models for scene understanding. In *Proceedings of the 2008 international conference on Content-based image and video retrieval*, pages 169–178, 2008.
- [55] Yu Zhang and Qiang Yang. A Survey on Multi-Task Learning. *arXiv:1707.08114 [cs]*, July 2018. arXiv: 1707.08114.
- [56] Yan-Tao Zheng, Ming Zhao, Yang Song, Hartwig Adam, Ulrich Buddemeier, Alessandro Bissacco, Fernando Brucher, Tat-Seng Chua, and Hartmut Neven. Tour the world: building a web-scale landmark recognition engine. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1085–1092. IEEE, 2009.
- [57] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464, 2017.
- [58] Mingyang Zhou, Runxiang Cheng, Yong Jae Lee, and Zhou Yu. A visual attention grounding neural model for multimodal machine translation. *arXiv preprint arXiv:1808.08266*, 2018.