

# Generating image captions with external encyclopedic knowledge

Sofia Nikiforova, Tejaswini Deoskar, Denis Paperno, Yoad Winter

Utrecht University, Utrecht, the Netherlands

{s.nikiforova, t.deoskar, d.paperno, y.winter}@uu.nl

## Abstract

Accurately reporting what objects are depicted in an image is largely a solved problem in automatic caption generation. The next big challenge on the way to truly humanlike captioning is being able to incorporate the context of the image and related real world knowledge. We tackle this challenge by creating an end-to-end caption generation system that makes extensive use of image-specific encyclopedic data. Our approach includes a novel way of using image location to identify relevant open-domain facts in an external knowledge base, with their subsequent integration into the captioning pipeline at both the encoding and decoding stages. Our system is trained and tested on a new dataset with naturally produced knowledge-rich captions, and achieves significant improvements over multiple baselines. We empirically demonstrate that our approach is effective for generating contextualized captions with encyclopedic knowledge that is both factually accurate and relevant to the image.

## 1 Introduction

Most modern image captioning systems are designed to produce a straightforward description of the visual content of the image. They abstract away from the image context and have no access to information that is not directly inferred from the image. By contrast, humans can extend beyond a purely visual description and produce a caption that is influenced by the image context and informed by real world knowledge. As a result, human-generated and automatically generated captions for the same image can be drastically different, as seen in Figure 1. This contrast motivates the task of *contextualized* image captioning, which sets the goal of automatically generating captions that would account for relevant image-external knowledge.

In this paper we propose an approach to contextualizing an image captioning system and enriching



**Human:** Clock Tower, Palace of Westminster. Completed in 1859, the clock tower houses the bell known as Big Ben.

**BUTD (Anderson et al., 2018):** a tall clock tower towering over a city

**ClipCap (Mokady et al., 2021):** A clock tower is seen against a cloudy sky.

Figure 1: Human-generated caption vs. captions generated by standard automatic captioning systems.

it with a wide range of encyclopedic knowledge. Our key contribution is in constructing an image-specific knowledge context that consists of relevant facts from an open-domain knowledge base, retrieved based on image location metadata. We bridge the gap between two lines of research in contextualized image captioning: one utilizes knowledge from external data sources, establishing its relevance through image recognition (Bai et al., 2021; Huang et al., 2020), the other uses geographic metadata to identify the names of geographic entities in and around the image (Nikiforova et al., 2020). To the best of our knowledge, we are the first to combine these two approaches — we use geographic metadata to access an external knowledge base and obtain various kinds of non-geographic encyclopedic facts relevant to the image, with their subsequent integration into the captioning pipeline.

We implement the proposed approach in a knowledge-aware caption generation system with image-external knowledge incorporated at both the encoding and decoding stages. We also present a new dataset of images with naturally created knowledge-rich captions and corresponding geographic metadata. Our experiments on this dataset show the effectiveness of our approach: the system greatly outperforms multiple baselines in commonly used captioning metrics and, crucially, the

correctness of the generated facts.

## 2 Related Work

Two main aspects that define each approach to contextualized image captioning are (i) the source of external knowledge and the way to identify the data relevant to a particular image, and (ii) the method of incorporating external knowledge into caption generation.

**External knowledge source** In a popular sub-task of news image captioning, captions are generated for images that accompany news articles (Zhao et al., 2021; Hu et al., 2020; Tran et al., 2020; Jing et al., 2020; Chen and Zhuge, 2020; Biten et al., 2019). Naturally, the article texts themselves are the main source of context for captioning, supplying information about important events and entities. In a more general case images are not paired with the relevant context directly. A common way to connect images to an external source of knowledge is to use an object detection mechanism to identify objects in the image and then use their labels to query a database. In Huang et al. (2020); Zhou et al. (2019); Wu et al. (2017) detected labels are used to extract useful information about common objects featured in the images (such as “dog”, “pot”, “surfboard”) from ConceptNet (Speer et al., 2017) and DBpedia (Auer et al., 2007). Zhao et al. (2019) use Google Cloud Vision APIs to identify not only common objects but also entities (people, car brands, etc). In Bai et al. (2021) custom classifiers are trained to detect specific image attributes (e.g. the author, the artistic style), which are then used to retrieve relevant information from Wikipedia. These approaches exclusively use the visual content of the images to contextualize the captioning process. Thus, the extent of contextualization is limited by the quality of the object detection algorithms, and the potential benefit of utilizing additional data (e.g. image metadata) is left unexplored.

Certain types of image metadata can be used to build upon general object detection and identify specific entities and events in the image. Lu et al. (2018) use the time metadata of the image (the date when a given photo was taken) and its associated tags to collect similar photographs and to retrieve the names of relevant entities (e.g. people depicted in the image) from their captions. In Nikiforova et al. (2020), geographic metadata (latitude and longitude coordinates of the image location) is used to extract information about the surrounding objects

from a geographic database, which allows their system to refer to concrete locations relevant to the image in the generated captions. These papers demonstrate the effectiveness of using image metadata for contextualized captioning but are limited to establishing the names of relevant entities, and do not utilize them to get further data that could be useful for generating even more informative captions.

In contrast to the works described above, we use image metadata as a grounding “anchor”, with which we can not only identify entities relevant to the image, but also retrieve a wide range of related encyclopedic knowledge from an external database. Specifically, we use geographic metadata, which has the benefit of being easily available for many real-life photographs due to the built-in GPS in modern cameras and phones, making it easier to collect the data for training and testing the system.

**Incorporating external knowledge into caption generation** There are two dominant methods of incorporating external knowledge into the caption generation process: template-based and context-based. In template-based approaches, a caption is generated with placeholder token slots that are later filled with the most fitting named entities extracted from an external knowledge source (Bai et al., 2021; Jing et al., 2020; Hu et al., 2020; Biten et al., 2019). This is an especially common technique in news image captioning, where named entities are taken from the news article associated with the image. Still, the straightforward fill-in-the-slot method can be problematic if none of the available entities fit the already generated placeholder slot.

In context-based approaches, external knowledge informs the caption generation process along with the image features, influencing the choice of produced tokens. For example, Zhou et al. (2019) extract ConceptNet terms related to the image and use their embeddings to initialize the caption generation module. Huang et al. (2020) also use ConceptNet to identify relevant external knowledge and increase the output probabilities of the vocabulary tokens if they match the extracted entities. The downside of context-based models is their inability to generate tokens that are present in the external knowledge but happen to be out-of-vocabulary for the generator’s language model, which is common for named entities.

Our model, like several other approaches (Chen and Zhuge, 2020; Tran et al., 2020; Nikiforova

et al., 2020; Whitehead et al., 2018), combines characteristics of both template-based and context-based methods. Similarly to template-based architectures, some of the tokens produced by our caption generation model are taken directly from external knowledge sources, and as in the context-based methods, external knowledge in our model influences the generation of regular vocabulary words.

### 3 Image caption dataset

For the purposes of training a knowledge-aware image captioning system, we require a dataset of images with knowledge-rich captions. Standard image captioning datasets, such as MSCOCO (Lin et al., 2014) or Flickr30k (Young et al., 2014), contain heavily curated captions, purposefully stripped of any contextualization and references to encyclopedic knowledge. Moreover, they lack image metadata, which is critical for our approach. Therefore, we compile our own dataset of 7128 images with the associated naturally created captions and image location metadata. The source of our data is the website of the Geograph project<sup>1</sup>, which contains a large number of photographs with English-language captions and rich metadata, including the geographic coordinates of the photograph locations. The images we selected for our dataset depict geographic entities, mostly historical man-made structures and buildings, to ensure that each caption contains at least one encyclopedic fact, for example, “Theatre Royal Haymarket. Dating back to 1720”. More details of the dataset, including its split into train, validation and test sets, are given in Appendix A.

### 4 Approach

Our approach begins with identifying real world entities that are related to the image. Specifically, we use the coordinates of where the photograph was taken to retrieve a set of geographic entities around that location (Section 4.1). Then, we collect a wide range of relevant encyclopedic facts. We extract them from an open-domain knowledge base in the form of <subject, predicate, object> triples, where subjects are the previously identified geographic entities (Section 4.2). Finally, we integrate the collected knowledge into an otherwise standard encoder-decoder image captioning pipeline to produce knowledge-aware captions (Section 4.3).

<sup>1</sup><http://www.geograph.org.uk/>

### 4.1 Geographic Context

We adapt an approach from Nikiforova et al. (2020) to retrieve geographic entities relevant to the image from the OpenStreetMap<sup>2</sup> database and thus to construct an image-specific geographic context. Formally, the geographic context  $G$  of a given image is a set of  $n$  geographic entities ( $e_1 \dots e_n$ ) located within a radius  $r$  from the image location. Based on preliminary experiments, we set  $n$  at 300 and  $r$  at 1 kilometer as the hyperparameters of our system.

Each geographic entity  $e_i$  is associated with its name and a set of geographic features: distance  $d_i$  and azimuth  $a_i$  between the entity and the image location, the entity’s size  $s_i$  and type  $t_i$ . In addition to that, we introduce two new features, intended to reflect the salience of the entity through the amount of information available about it in a knowledge base: a binary indicator  $\exists f_i$  that shows whether or not the entity corresponds to any facts in the knowledge context, and the number of facts  $\#f_i$  that correspond to the entity in the knowledge context. A sample fragment of a geographic context, with the entities mapped to their names and features, is shown in Figure 2.

```
e1 - Theatre Royal - (d1=0.02km, a1= -132°, s1=0.001km2, t1=theatre, ∃f1=1, #f1=2)
e2 - Haymarket - (d2=0.05km, a2= -30°, s2=0.0004km2, t2=primary, ∃f2=0, #f2=0)
...
en - Charing Cross - (dn=0.37km, an= -81°, sn=0.0km2, tn=station, ∃fn=1, #fn=1)
```

Figure 2: A fragment of a geographic context.

The features are combined in vector representations for the entities, called “geographic embeddings”. For an entity  $e_i$  a geographic embedding is computed as follows:

$$\text{GEOEMB}(e_i) = \text{Concat}[d_i, \text{norm}(a_i), s_i, \exists f_i, \#f_i, \text{Emb}_t(t_i)] \quad (1)$$

where  $\text{norm}$  is an azimuth normalization function,  $\text{Emb}_t$  is an embedding function for the entities’ types. For implementation details, see Appendix C. Our experiments show that our straightforward concatenation strategy achieves the same performance level as the more complicated approach employed by Nikiforova et al. (2020).

### 4.2 Knowledge Context

For a given image with the geographic context  $G$ , the knowledge context  $K$  is defined as a set of  $m$  facts ( $f_1 \dots f_m$ ) about the entities in  $G$ .

<sup>2</sup><https://www.openstreetmap.org/>

We obtain the facts from the DBpedia knowledge base, where they are stored as triples of the form  $\langle \text{subject}, \text{predicate}, \text{object} \rangle$ . First, we select all the facts, in which the subject is one of the geographic entities from  $G$ , e.g.  $\langle \text{Theatre Royal}, \text{built\_in}, 1720 \rangle$ ,  $\langle \text{Theatre Royal}, \text{architect}, \text{John Nash} \rangle$ ,  $\langle \text{Theatre Royal}, \text{rebuilt}, 1879 \rangle$ , etc. Then, we train a logistic regression model to rank the facts based on how likely they are to be mentioned in the caption. The model takes into account the fact’s predicate, the ranking of the fact’s subject in the geographic context and its geographic features. The top  $m$  facts of the ranked list constitute the knowledge context of the image, with  $m$  as another hyperparameter of the system, which we set at 50.

Knowledge contexts of all the images in our dataset contain altogether facts with 1542 unique predicates, 336 of which appear at least once in the ground truth captions. The detailed statistics of the collected knowledge data (e.g. fact and predicate distribution per entity, image and caption) are provided in Appendix B.




Figure 3: A fragment of a knowledge context.

Figure 3 shows a fragment of the knowledge context corresponding to the geographic context in Figure 2. As seen in the figure, we modify the way the original DBpedia facts are represented. Instead of the  $\langle \text{subject}, \text{predicate}, \text{object} \rangle$  triples, we treat the objects as labels of the facts (similarly to the names, such as “Theatre Royal”, that act as labels of the geographic entities). In our approach, it is the fact objects that constitute the ‘vocabulary’ of fact-related tokens that appear in captions. For example, in the caption “Theatre Royal Haymarket. Dating back to 1720”, the fact  $\langle \text{Theatre Royal}, \text{built\_in}, 1720 \rangle$  is realized through the presence of the fact’s object “1720”, while “Theatre Royal” comes from the geographic context and “dating back” comes from the regular vocabulary. So, given the contexts in Figures 2 and 3, this caption will actually be processed by the system as “ $e_1 e_2$ . Dating back to  $f_1$ ”.

Thus, each fact in the knowledge context is represented through its object as a label, which can appear in a caption, and through the  $\langle \text{subject}$ ,

predicate  $\rangle$  pair, which reflects the fact’s meaning. We encode this meaning in a vector form, a “fact embedding”, which is computed as follows:

$$\text{FACTEMB}(f_i) = \text{GEOEMB}(e_i) + \text{Emb}_p(p_i) \quad (2)$$

where  $e_i$  is the subject of the fact  $f_i$  (an entity from the geographic context),  $p_i$  is its predicate and  $\text{Emb}_p$  is an embedding function for the predicates. Predicate embeddings are initialized randomly and updated in an end-to-end fashion during the training of the captioning model.

This approach to encoding facts in the knowledge context provides the captioning system with the information it needs to select an appropriate fact based on the previously generated tokens. For any given fact, the system can take into account whether or not its subject is already present in the caption and estimate whether the previous caption tokens are consistent with the fact’s predicate. One possible alternative for encoding structured external knowledge is applying graph embedding methods (such as a popular TransE algorithm (Bordes et al., 2013)) to a knowledge graph, and then utilizing the embeddings of the graph nodes (entities) and edges (relations) during text generation. Our approach is less time- and resource-consuming, since it requires no extra pre-training: predicate embeddings are trained together with the captioning model, we re-use geographic embeddings to encode fact subjects, and embeddings of the fact objects are a linear combination of the predicate and subject embeddings. Moreover, using geographic embeddings for encoding the fact subjects provides additional grounding to the knowledge context and a direct link between the knowledge context and the geographic context.

### 4.3 Model Architecture

Our caption generation system is an end-to-end trainable neural network with an encoder-decoder architecture. The overview of the system’s architecture is given in Figure 4. As shown in the figure, both encoder and decoder make use of the external knowledge to produce knowledge-rich captions. See Appendix C for concrete implementation details.

**Encoder** The encoder’s function is to convert input data into an informative representation that is subsequently used by the decoder to generate a caption. In a standard image captioning pipeline, the input data consists only of the image itself, and

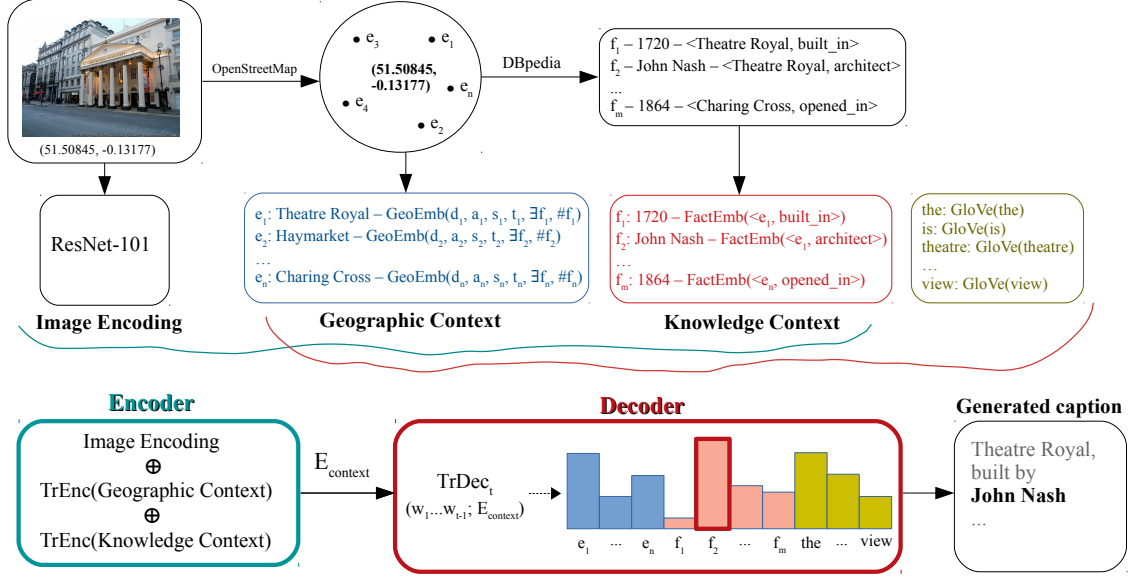


Figure 4: An overview of the system architecture (best viewed in color).

its encoding is a dense representation of its visual features. In our system, we also use geographic and knowledge contexts as the additional sources of input data.

For the image encoding  $E_{img}$ , we use a deep convolutional neural network (CNN), pre-trained on an image classification task. The particular CNN that we use is ResNet-101 (He et al., 2016), trained on the images from the ImageNet database (Russakovsky et al., 2015). In addition, we encode information contained in the geographic and knowledge contexts. First, each of their elements is embedded with the embedding functions introduced in Sections 4.1, 4.2:

$$\begin{aligned}
 EmbG &= (\text{GEOEMB}(e_1) \dots \\
 &\quad \dots \text{GEOEMB}(e_n)), e_i \in G \\
 EmbK &= (\text{FACTEMB}(f_1) \dots \\
 &\quad \dots \text{FACTEMB}(f_m)), f_i \in K
 \end{aligned} \quad (3)$$

They are subsequently encoded with two separate Transformer encoders (TrEnc), with a standard structure proposed in Vaswani et al. (2017). Finally, we concatenate the encodings of the image, the geographic context and the knowledge context:

$$\begin{aligned}
 E_{context} &= \\
 &= \text{Concat}[E_{img}, \text{TRENC}(EmbG), \text{TRENC}(EmbK)]
 \end{aligned} \quad (4)$$

The result of the concatenation is the combined representation of the visual features of the image

and the relevant information from the geographic and knowledge contexts.

**Decoder** The decoder accepts the combined context representation from the encoder and generates an output sequence — the caption. The goal of the decoder is to produce a caption that would be fitting to the image and include accurate references to the external knowledge. The decoder generates a caption token by token, at each step  $t$  taking into account the previously generated tokens  $w_1 \dots w_{t-1}$  and the context representation  $E_{context}$ . In the process, each input token is represented by a sum of its vector embedding and the encoding of its position in the sequence.

$$\text{PosEmb}(w_i) = \text{Emb}(w_i) + \text{Pos}(w_i) \quad (5)$$

We distinguish three types of tokens: regular vocabulary words, names of geographic entities (the names mapped to entities in the geographic context) and fact-related tokens (the objects of the facts in the knowledge context). We use pre-trained GloVe word embeddings (Pennington et al., 2014) for the regular vocabulary tokens. However, conventional pre-trained embeddings are not particularly informative for the geographic entity names and the fact-related tokens, as their semantics is conveyed poorly by their distribution patterns (e.g. the various contexts in which the token “1720” appears in a large scale corpus are too diverse for a precise representation of “1720” as the year when Theatre Royal was built). It is important that the

decoder can utilize information about their most meaningful characteristics: physical properties of the geographic entities and the facts’ subjects and predicates. For this reason, we use the GEOEMB and FACTEMB embedding functions to represent them.

$$Emb(w_i) = \begin{cases} \text{GEOEMB}(w_i), & \text{if } w_i \in G \\ \text{FACTEMB}(w_i), & \text{if } w_i \in K \\ \text{GLOVE}(w_i), & \text{otherwise} \end{cases} \quad (6)$$

We employ a Transformer decoder (TrDec) with a standard structure. It attends to the positional embeddings of the previously generated tokens and to the encoder’s output, the combined representation of the image contexts.

$$h_t = \text{TRDEC}(\text{PosEmb}(w_{1\dots t-1}); E_{context}) \quad (7)$$

In a standard captioning pipeline, the output of the decoder  $h_t$  is then passed to a final linear layer that acts as a classifier, estimating the probability distribution over all the tokens in the vocabulary  $V$ . The vocabulary is usually fixed and consists of the words or subwords from the training dataset. In our case, captions also include entity names and facts from the geographic and knowledge contexts, which are image-specific, and therefore, not all relevant entity names and fact-related tokens would necessarily be a part of  $V$ . So, we modify the last stage of the decoding process by computing three sets of scores: the scores for the vocabulary tokens  $v_1\dots v_k \in V$ , the scores for the geographic entity names  $e_1\dots e_n \in G$  and the scores for the facts  $f_1\dots f_m \in K$ .

$$\begin{aligned} y_{v_1\dots v_k} &= (h_t \circ p_{ind} W_{pred}) W_{vocab} \\ y_{e_1\dots e_n} &= (EmbG \text{DIAG}(h_t)) \vec{w}_{geo} \\ y_{f_1\dots f_m} &= ((EmbK \circ g_{ind}) \text{DIAG}(h_t)) \vec{w}_f \end{aligned} \quad (8)$$

where  $W_{vocab}$  and  $W_{pred}$  are trainable linear transformation matrices,  $\vec{w}_{geo}$  and  $\vec{w}_f$  are trainable linear transformation vectors,  $\text{DIAG}(h_t)$  denotes a diagonal matrix with the  $h_t$  vector in the main diagonal,  $\circ$  stands for a Hadamard product.

The predicate indicator  $p_{ind}$  and the geo entity indicator  $g_{ind}$  provide additional contextualization to the generation process<sup>3</sup>. The geo entity indicator specifies for each knowledge context fact

<sup>3</sup>We have also experimented with simpler architectures, without the predicate and geo entity indicators. The performance of these alternative systems was overall lower, especially in the accuracy of the generated encyclopedic facts. See Appendix D for more details.

whether its subject has already been mentioned in the caption at a given time step (1 for ‘has been mentioned’, 0 for ‘has not been mentioned’). Facts with the subjects that are already in the caption would generally get a higher generation probability assigned to them by the model.

The predicate indicator ensures that the external knowledge has influence over the generation of regular vocabulary words. It reflects for each of the known predicates whether the knowledge context contains a fact with this predicate *and* the subject that has already been mentioned in the caption. Introducing this indicator into the calculation of vocabulary scores reduces the probability of producing phrases that would require a specific type of fact being generated afterwards if no such facts are available in the knowledge context.

The three sets of scores are then concatenated and fed to a softmax layer, which produces an overall probability distribution over the tokens  $(v_1\dots v_k) \in V$ ,  $(e_1\dots e_n) \in G$  and  $(f_1\dots f_m) \in K$  (see the diagram in Figure 4). The token with the highest probability is generated at position  $t$ .

$$\begin{aligned} w_t &= \arg \max_{w_i} P(w_i), w_i \in V \cup G \cup K \\ \text{where } P(w_i) &= \\ &= \sigma_i(\text{Concat}[y_{v_1\dots v_k}, y_{e_1\dots e_n}, y_{f_1\dots f_m}]) \end{aligned} \quad (9)$$

## 5 Results and Discussion

We trained and tested our captioning system on the dataset described in Section 3. We evaluate its performance with standardly used automatic image captioning metrics: BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), METEOR (Denkowski and Lavie, 2014) and CIDEr (Vedantam et al., 2015). We also present *factual accuracy* as a separate focus of evaluation. Simply generating captions with verifiable statements is not enough — they have to be accurate in their description of the entities in the image. We introduce a custom metric that measures the correctness of the generated captions against the DBpedia database.

**Baselines** For the first baseline, we examine the performance of a standard caption generation system (Mokady et al., 2021) on our test set. The standard system has no contextual components and was trained on the out-of-domain images from the MSCOCO dataset.

	BLEU-1	BLEU-2	BLEU-3	BLEU-4	ROUGE	METEOR	CIDEr
Standard (Mokady et al., 2021)	0.33	0.09	0.02	0.0	5.62	2.45	0.14
No-knowledge	16.29	8.43	4.38	2.25	23.14	7.24	1.75
Knowledge-from-image	19.71	9.68	4.88	2.4	23.48	7.76	1.63
Knowledge-from-metadata	<b>26.84</b>	<b>17.98</b>	<b>12.85</b>	<b>9.5</b>	<b>33.47</b>	<b>13.52</b>	<b>32.41</b>

Table 1: Metric scores of the different systems, measured on the test set (a single run).

We create two more baseline captioning systems<sup>4</sup> and train them on our dataset. The first one is a “no-knowledge” system, in which both geographic and knowledge components have been removed. Its performance represents the level that can be achieved by a standard image captioning pipeline with no additional contextualization.

The second baseline utilizes external knowledge based solely on the image. This system, which we call “knowledge-from-image”, draws from the existing research on contextualized image captioning (Huang et al., 2020; Zhou et al., 2019). Similarly to these works, our baseline uses a pre-trained object recognition model<sup>5</sup> to detect objects and scenes in the image (e.g. “beach”) and then queries the ConceptNet DB to retrieve the terms related to these objects (e.g. “coast”, “driftwood”, “ocean”). We identify top 5 objects in every image and extract top 10 closest ConceptNet terms for each of them. We further pass their GloVe embeddings through a Transformer encoder and concatenate the result with the encoding of the image to construct a combined context representation for caption generation.

## 5.1 Caption Quality Evaluation

Table 1 shows the comparison between the metric scores of our “knowledge-from-metadata” system and the three baselines. The standard system produces captions that are very different from the ground truth ones, which is reflected in the particularly low metric scores. This is expected, since the dataset it was trained on, as well as the architecture of the system itself, did not account for the context of the images and instead focused on the

visual descriptions only. The no-knowledge and knowledge-from-image baselines achieve higher scores (the latter with a marginal improvement over the former in most of the metrics), but both are outperformed by our knowledge-from-metadata system<sup>6</sup>. All the improvements of the knowledge-from-metadata system are statistically significant (two-sample t-test,  $p < 0.001$ ). The difference in performance mainly stems from the ability of our system to produce contextually appropriate tokens from the external knowledge ‘vocabularies’; even though the same tokens are present in the vocabularies of the baselines and can be generated as regular vocabulary words, their generation in the baselines is dominated by noise and almost never matches the correct facts from the ground truth captions. The most radical improvements are in the CIDEr metric, which gives a higher weight to the words that are more informative according to the TF-IDF score; geographic names and fact-related tokens are usually rare in the corpus and highly informative, so they contribute a lot to this metric.

Table 2 shows examples<sup>7</sup> of the captions generated by the knowledge-from-metadata system and the baselines, as well as the original human-written captions. Here, the standard system from Mokady et al. (2021) produces mostly accurate descriptions of what can be seen in the images but includes no references to their context. The no-knowledge and knowledge-from-image baseline systems fail to generate correct geographic entity names and facts, as they draw all the caption tokens out of the general vocabulary. The captions generated by the knowledge-from-metadata system demonstrate that it is able to successfully use both geographic and knowledge contexts and produce accurate facts about relevant geographic entities.

<sup>4</sup>The existing contextualized captioning systems (e.g. the ones referenced in Section 2) are not suitable for direct use as baselines here. Most of these systems are specifically designed for certain types of data, which makes it unfeasible to run them on our dataset. For example, the news image captioning models require a news article text as a context for the image; in Bai et al. (2021) the system is built and trained to generate specifically the descriptions of fine-art paintings.

<sup>5</sup>We employ the VGG16 CNN model (Simonyan and Zisserman, 2014) trained on the Places365 database (Zhou et al., 2017) for scene recognition. The model’s checkpoint is provided at <https://github.com/GKalliatakis/Keras-VGG16-places365>.

<sup>6</sup>We also experimented with a “geo-from-metadata” baseline, which has access to the geographic context but not to the knowledge context. Its performance is higher than that of the other baselines, although still lower than our system’s. For more details see Appendix E.

<sup>7</sup>The images for these examples and additional examples from the test set are given in Appendix F.

---

Ground truth: Kelso Bridge. Below the confluence of the Rivers Tweed and Teviot. John Rennie engineered the bridge, which was built between 1800 and 1803.

Standard (Mokady et al., 2021): A river with a bridge and a train on it.

No-knowledge: the *river dee. farndon bridge* was opened in 1339 by monks from *farndon bridge*.

Knowledge-from-image: *chertsey bridge. chertsey bridge* dates from 1785.

Knowledge-from-metadata (ours): **kelso bridge**. the bridge over the **river tweed** was built in **1800**, and was designed by **john rennie the elder**.

---

Ground truth: St. Mary’s Lighthouse. On St Mary’s Island, just north of Whitley Bay. Completed in 1898 the lighthouse remained operational until 1984 when it was superseded by modern navigational techniques.

Standard (Mokady et al., 2021): A lighthouse with a white tower on top.

No-knowledge: the *naze tower, walton on naze*. the *naze tower* was built in 1720 as a navigational aid to 1720 as a navigational aid to aid

Knowledge-from-image: *portland bill car park. portland bill* is a popular destination for visitors to the lighthouse, built in 1906.

Knowledge-from-metadata: **st marys lighthouse**. the lighthouse was built in **1898** and is now a **grade ii listed** building.

---

Table 2: Examples of the generated captions. Correct geographic references and facts are given in **bold**; incorrect ones are given in *italics*.

	Accuracy, %
No-knowledge	0.0
Knowledge-from-image	0.0
Random fact	53.35
Knowledge-from-metadata	<b>86.08</b>

Table 3: Fact accuracy scores (a single run).

## 5.2 Generated Facts Accuracy

An important indicator of the quality of the generated image captions is their factual accuracy. We measure it as the percentage of correct facts among all generated facts. To check the correctness of a given fact, we use a rule-based metric, which first verifies that the fact’s subject entity is related to the image and then searches for common predicate-specific key phrases (e.g. “designed by”) and fact objects from the knowledge context that should co-occur with the detected phrases (e.g. the name of the actual designer), with subsequent manual verification.

In addition to the no-knowledge and knowledge-from-image baselines introduced earlier, we also create a “random fact” baseline. It takes the captions generated by our system and replaces the fact-related tokens in each caption with the tokens randomly picked from the knowledge context, preserving the token type (e.g. replaces generated years with randomly picked years, generated names with randomly picked names, etc.). This creates quite a strong baseline because the probability of any entity from the knowledge context to be relevant to the image and thus to appear in the caption

is high by design.

Table 3 shows the accuracy scores of the three baselines and our knowledge-from-metadata system. The no-knowledge and knowledge-from-image baselines did not produce any accurate facts. The strong random fact baseline’s score is much higher but is still greatly outperformed by the knowledge-from-metadata system.

## 6 Conclusions

To mimic human behavior in describing an image, it is essential that image captioning systems take into account the context of the image and related real world knowledge. In this paper, we have presented a novel method of contextualizing an image captioning pipeline with real world encyclopedic facts that are relevant to the image but are not directly inferable from it. We compiled a new image captioning dataset with naturally produced informative captions and geographic metadata, which we use to obtain diverse image-specific knowledge from external data sources. The experiments on this dataset demonstrate the effectiveness of our approach: our captioning system is able to generate captions with correct facts about entities relevant to the image. Compared to a number of baseline systems, it achieves substantial improvements in the standard metrics as well as in the accuracy of the generated facts. In future work we plan to extend the proposed approach to other domains, such as medical image captioning (Ayesha et al., 2021), which would benefit greatly from incorporating external knowledge from medical databases.



## Limitations

In our approach, we rely on image location metadata to be the grounding “anchor” for retrieving external knowledge. Thus, if there is no image metadata or similar annotations available, our approach cannot be applied without substantial modifications. A side effect of this is that it is impossible to evaluate our system on the most common captioning datasets, such as MSCOCO or Flickr30k, which do not include image metadata.

Another limitation of our work is that we propose an encoding mechanism for external knowledge that is specific to the type and format of the knowledge utilized in our system ( $\langle$ subject, predicate, object $\rangle$  triples). A different encoding strategy would have to be devised if the knowledge was supplied in a different format (e.g. free-form text from Wikipedia).

Finally, the performance of our system depends on the coverage of the external data sources. If there is no relevant knowledge in a database for a given image, the generated caption is unlikely to contain accurate encyclopedic facts.

## Ethics Statement

This paper presents a dataset of images, captions and geographic metadata collected from the Geograph project website. All the original data is licensed for reuse under the Creative Commons BY-SA 2.0 license. In compliance with the license terms, our dataset is distributed under the more recent version of the same license — Creative Commons BY-SA 4.0. All the images in this paper are credited with a stamp that indicates the license, the author (the copyright owner) and the link to the original image.

This paper also presents an image captioning system, which can be deployed and run freely. The code of the system is made available under the Creative Commons BY-SA 4.0 license. We do not expect any specific harmful consequences of using the system or any potential for misuse.

## References

Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. [Bottom-up and top-down attention for image captioning and visual question answering](#). *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 6077–6086.

Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. [DBpedia: A nucleus for a web of open data](#). *Lecture Notes in Computer Science*, 6:722–735.

Hareem Ayesha, Sajid Iqbal, Mehreen Tariq, Muhammad Abrar, Muhammad Sanaullah, Ishaq Abbas, Amjad Rehman, Muhammad Farooq Khan Niazi, and Shafiq Hussain. 2021. [Automatic medical image interpretation: State of the art and future directions](#). *Pattern Recognition*, 114:107856.

Zechen Bai, Yuta Nakashima, and Noa Garcia. 2021. [Explain me the painting: Multi-topic knowledgeable art description generation](#). In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5422–5432.

Ali Furkan Biten, Lluís Gomez, Marçal Rusiñol, and Dimosthenis Karatzas. 2019. [Good News, Everyone! Context driven entity-aware captioning for news images](#). *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12466–12475.

Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). *Advances in neural information processing systems*, 26.

Jingqiang Chen and Hai Zhuge. 2020. [A news image captioning approach based on multimodal pointer-generator network](#). *Concurrency and Computation: Practice and Experience*, page e5721.

Michael Denkowski and Alon Lavie. 2014. [Meteor universal: Language specific translation evaluation for any target language](#). *Proceedings of the ninth workshop on statistical machine translation*, pages 376–380.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.

Anwen Hu, Shizhe Chen, and Qin Jin. 2020. [ICE-CAP: Information concentrated entity-aware image captioning](#). In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4217–4225.

Feicheng Huang, Zhixin Li, Haiyang Wei, Canlong Zhang, and Huifang Ma. 2020. [Boost image captioning with knowledge reasoning](#). *Machine Learning*, 109(12):2313–2332.

Yun Jing, Xu Zhiwei, and Gao Guanglai. 2020. [Context-driven image caption with global semantic relations of the named entities](#). *IEEE Access*, 8:143584–143594.

Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. [Microsoft COCO: Common objects in context](#). In *European conference on computer vision*, pages 740–755. Springer.
- Di Lu, Spencer Whitehead, Lifu Huang, Heng Ji, and Shih-Fu Chang. 2018. [Entity-aware image caption generation](#). *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4013–4023.
- Ron Mokady, Amir Hertz, and Amit H Bermano. 2021. [ClipCap: CLIP prefix for image captioning](#). *arXiv preprint arXiv:2111.09734*.
- Sofia Nikiforova, Tejaswini Deoskar, Denis Paperno, and Yoad Winter. 2020. [Geo-aware image caption generation](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3143–3156, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [BLEU: a method for automatic evaluation of machine translation](#). *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. [PyTorch: An imperative style, high-performance deep learning library](#). *Advances in Neural Information Processing Systems*, pages 8024–8035.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. [GloVe: Global vectors for word representation](#). *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. [ImageNet large scale visual recognition challenge](#). *International journal of computer vision*, 115(3):211–252.
- Karen Simonyan and Andrew Zisserman. 2014. [Very deep convolutional networks for large-scale image recognition](#). *arXiv preprint arXiv:1409.1556*.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. [ConceptNet 5.5: An open multilingual graph of general knowledge](#). *Thirty-First AAAI Conference on Artificial Intelligence*, pages 4444–4451.
- Alasdair Tran, Alexander Mathews, and Lexing Xie. 2020. [Transform and tell: Entity-aware news image captioning](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13035–13045.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. [CIDEr: Consensus-based image description evaluation](#). *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575.
- Spencer Whitehead, Heng Ji, Mohit Bansal, Shih-Fu Chang, and Clare Voss. 2018. [Incorporating background knowledge into video description generation](#). *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3992–4001.
- Qi Wu, Chunhua Shen, Peng Wang, Anthony Dick, and Anton Van Den Hengel. 2017. [Image captioning and visual question answering based on attributes and external knowledge](#). *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1367–1381.
- Kelvin Xu, Jimmy Lei Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. 2015. [Show, attend and tell: Neural image caption generation with visual attention](#). *32nd International Conference on Machine Learning, ICML 2015*, pages 2048–2057.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. [From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions](#). *Transactions of the Association of Computational Linguistics*, 2(1):67–78.
- Sanqiang Zhao, Piyush Sharma, Tomer Levinboim, and Radu Soricut. 2019. [Informative image captioning with external sources of information](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6485–6494.
- Wentian Zhao, Yao Hu, Heda Wang, Xinxiao Wu, and Jiebo Luo. 2021. [Boosting entity-aware image captioning with multi-modal knowledge graph](#). *arXiv preprint arXiv:2107.11970*.
- Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. 2017. [Places: A 10 million image database for scene recognition](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1452–1464.
- Yimin Zhou, Yiwei Sun, and Vasant Honavar. 2019. [Improving image captioning by leveraging knowledge graphs](#). *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 283–293.

	Per image	Per caption	Per entity
Average number of facts	36.49	2.2	7.34
Average number of predicates	24.96	2.2	5.82

Table 4: Average numbers of facts/predicates.

## A Dataset

**Image selection** The main selection criterion for our image caption dataset is that each caption has to include at least one encyclopedic fact about a geographic entity relevant to the image. This ensures that the knowledge-aware captioning system has enough data to learn from. For practical reasons, we also select only captions that are no longer than 100 tokens.

**Train/validation/test split** We split our captioning dataset into train, validation and test sets that constitute, respectively, 75%, 12.5% and 12.5% of the whole dataset. In order to avoid assigning different photographs of the same geographic entities to both train and validation/test sets, we base the split on the latitude of the image location instead of splitting the dataset randomly. The photographs that were taken to the north of the 54.8975° latitude are assigned to the test set, between the 53.5706° and the 54.8975° latitude to the validation set, and the rest to the train set. With the latitude-based split, we ensure testing on previously unseen data, which helps to detect possible overfitting.

**Caption pre-processing** During pre-processing, captions are converted to lower case and split into tokens with the standard NLTK Tokenizer package. Formatting artifacts are removed and some tokens are replaced with others for unification purposes (e.g. “&” is replaced with “and”, “saint” is replaced with “st”). By simple string matching against the geographic and knowledge contexts, we identify geographic entities and facts in the captions and create a “mask” that shows the type of each token (an entity, a fact or a regular vocabulary word). The mask is supplied to the captioning system along with the pre-processed caption tokens and used for selecting the appropriate embedding function for each token (see Equation 6).

## B External Knowledge Statistics

Table 4 provides the quantitative statistics of the external knowledge data used to train and evaluate our image captioning system. The average number

per image shows how many distinct facts and predicates there are in the knowledge context for a given image. The average number per caption shows how many facts and predicates actually appear in the ground truth captions. Finally, the average number per entity shows how many unique facts and predicates a single subject (geographic entity) generally has in the knowledge context. The number of unique predicates is lower than the number of unique facts because in some cases several distinct facts may share a predicate, e.g. <Sadler’s Wells Theatre, rebuilt, 1765> and <Sadler’s Wells Theatre, rebuilt, 1998>.

Figure 5 shows the distribution of the top 50 frequent predicates to appear in the captions. Many of the predicates displayed in the figure are in fact synonymous, i.e. reflect the same relation between the subject and object entities. For example, “opened” and “openingyear” are both used to link a geographic entity (the subject) and the year when it was opened (the object). Ideally, every unique relation between entities should be represented by a single predicate. On this account, we mapped frequent synonymous predicates to a single predicate that represents their common meaning, e.g. both “opened”, “openingyear” are merged into “opened”. The distribution of the top 50 most frequent predicates after merging the synonymous ones can be seen in Figure 6.

## C Implementation

Our system is implemented in the PyTorch framework (Paszke et al., 2019), version 1.9.0, with CUDA version 11.0. The versions of the all the packages we use are provided in the requirements file that is distributed with the code. The system hyperparameters have been selected without automatic tuning; here we report the chosen values.

**Encoder** All the images in the dataset are resized to 256x256 pixels to go through the pre-trained CNN image recognition module, which we do not fine-tune. The output of the image encoder has a size 14x14 with 2048 color channels.

The Transformer Encoder used for encoding the geographic and knowledge contexts is a stan-




Figure 5: Distribution of the top 50 frequent original predicates.




Figure 6: Distribution of the top 50 frequent predicates after merging the synonymous ones.

standard PyTorch implementation, with 3 layers and 10 heads in a layer, with a dropout value of 0.5 and a feedforward network with a dimension size 512.

**Context** The size of both the geographic embedding for the entities in the geographic context and the fact embedding for the facts in the knowledge context is 300, which matches the size of the pre-trained GloVe embeddings used for the vocabulary words. In the geographic embedding, the distance and the size parameters are used without any normalization. The azimuth of the angle between the object and the image location is normalized as follows:

$$a_i = \text{Concat}[a_{north}; a_{east}]$$

where

$$a_{north} = |a_i|/180$$

$$a_{east} = \begin{cases} |90 - a_i|/180, & a_i \geq -90 \\ (90 + |a_i + 180|)/180, & \text{otherwise} \end{cases} \quad (10)$$

The normalization ensures that the azimuth values that are close to each other “on the compass”

get similar values after normalization.

The types of the geographic entities (e.g. “town”, “river”) and the fact predicates (“architect”, “opened”) are encoded with two separate embedding layers. Both are fine-tuned during the training of the image captioning system.

**Decoder** The decoder follows a standard PyTorch implementation of a Transformer Decoder. Similarly to the encoder, the decoder has 3 layers and 10 heads in a layer, the dimension of the feedforward network is 512, the dropout value is 0.5. The embeddings of the vocabulary words are initialized from the GloVe embeddings of size 300, pre-trained on the Common Crawl data and fine-tuned during training.

**Training** We train the system with the Adam optimizer with the learning rate of  $4e-4$ . During backpropagation, the decoder gradients are clipped to the absolute value of 5.0. We use cross entropy loss; the early stopping is enabled after 20 consecutive epochs without a loss decrease. The training was carried out with a 4GB NVIDIA Quadro P1000 GPU, with an estimated total of 210 GPU hours for

all the experiments, including baselines (a single training run takes approximately 3 hours, a single evaluation run — approximately 3.5 minutes). The total number of trainable parameters in the network is 16,121,781.

**Evaluation** The implementation of the standard captioning metrics (BLEU, ROUGE, METEOR, CIDEr) is adopted from <https://github.com/tylin/coco-caption>. For one of the baselines we use a pre-trained ClipCap captioning model, with the implementation and checkpoint provided at [https://github.com/rmokady/CLIP\\_prefix\\_caption](https://github.com/rmokady/CLIP_prefix_caption).

## D Alternative Architectures

**No predicate indicator** In an alternative setup without the predicate indicator  $p_{ind}$ , the scores calculation from Equation 8 is modified for the vocabulary words  $v_1 \dots v_k \in V$  as follows:

$$y_{v_1 \dots v_k} = h_t W_{vocab} \quad (11)$$

In this case, the influence of the knowledge context on vocabulary word selection is not as strong. This leads to a decrease in the generated fact accuracy, which is measured at 80.74% (as compared to the fact accuracy of 86.08% when the predicate indicator is used). A typical mistake with this setup is generating a phrase from the vocabulary that warrants a certain type of fact that is not present in the knowledge context. For example, generating “elie house, founded in 1697” when actually Elie House was *built* in 1697. Since there is no fact in the corresponding knowledge context with the subject “Elie House” and the predicate “founded”, generating “founded in” after “elie house” would very likely lead to producing an incorrect fact. Accounting for the knowledge context while generating vocabulary words is the purpose of the predicate indicator.

**No geo entity indicator** Without the geo entity indicator  $g_{ind}$ , the scores calculation from Equation 8 is modified for the facts  $f_1 \dots f_m \in K$  as follows:

$$y_{f_1 \dots f_m} = (EmbK \text{ DIAG}(h_t)) \vec{w}_f \quad (12)$$

Without this indicator, the model is more prone to generating facts, the subjects of which have not been mentioned in the caption. This often results in a mistake when the model selects a fact with an appropriate predicate but an incorrect subject, e.g.

generating “neidpath viaduct, built in 1263” with the following two facts in the knowledge context:  $\langle \text{Neidpath Viaduct, built, 1863} \rangle$  and  $\langle \text{Neidpath Castle, built, 1263} \rangle$ . The architecture without the geo entity indicator scores 76.86% in the generated fact accuracy, almost 10 percentage points lower than when the geo entity indicator is present.

## E Geographic-Knowledge-from-Metadata Baseline

In order to evaluate the significance of the knowledge context component in our image captioning system, we conduct an experiment with another baseline, called “geo-from-metadata”. This baseline has access to the geographic context but not to the knowledge context: the output of the encoder is the concatenation of the image representation and the geographic context encoding alone, and during caption generation the decoder can only pick from vocabulary tokens and geographic entity names. The difference between the performance level of the geo-from-metadata and knowledge-from-metadata systems demonstrates the impact the knowledge context has on the generated captions.

The standard captioning metrics scores of the geo-from-metadata system<sup>8</sup> show a significant improvement over the no-knowledge baseline, which does not include any contextualization at all; however, the scores remain much lower than the ones achieved by our knowledge-from-metadata system. Examples of the captions generated by the geo-from-metadata system are given in Appendix F. As is evident from the captions, the geo-from-metadata system can utilize the geographic context available to it to produce accurate geographic references but, not being able to use the knowledge context, does not produce correct encyclopedic facts. With a few coincidental right guesses, the fact accuracy on a test set is measured at 6.56%. Thus, incorporating knowledge context is shown to be crucial for generating captions with accurate encyclopedic facts about entities relevant to the image.

## F Examples of the Generated Captions

<sup>8</sup>BLEU-1 — 21.91, BLEU-2 — 13.25, BLEU-3 — 8.35, BLEU-4 — 5.32, ROUGE — 30.19, METEOR — 10.98, CIDEr — 16.59.





(a)		<p><u>Ground truth:</u> Kelso Bridge. Below the confluence of the Rivers Tweed and Teviot. John Rennie engineered the bridge, which was built between 1800 and 1803.</p> <p><u>Standard</u> (Mokady et al., 2021): A river with a bridge and a train on it.</p> <p><u>No-knowledge:</u> <i>the river dee. farndon bridge</i> was opened in 1339 by monks from <i>farndon bridge</i>.</p> <p><u>Knowledge-from-image:</u> <i>chertsey bridge. chertsey bridge</i> dates from 1785.</p> <p><u>Geo-from-metadata:</u> the <b>river tweed</b>. the bridge is a <i>grade ii listed</i> building. the bridge was built in 1826 and completed in 1805.</p> <p><u>Knowledge-from-metadata:</u> <b>kelso bridge</b>. the bridge over the <b>river tweed</b> was built in <b>1800</b>, and was designed by <b>john rennie the elder</b>.</p>
(b)		<p><u>Ground truth:</u> St. Mary's Lighthouse. On St Mary's Island, just north of Whitley Bay. Completed in 1898 the lighthouse remained operational until 1984 when it was superseded by modern navigational techniques.</p> <p><u>Standard</u> (Mokady et al., 2021): A lighthouse with a white tower on top.</p> <p><u>No-knowledge:</u> <i>the naze tower, walton on naze. the naze tower</i> was built in 1720 as a navigational aid to 1720 as a navigational aid to aid</p> <p><u>Knowledge-from-image:</u> <i>portland bill car park. portland bill</i> is a popular destination for visitors to the lighthouse, built in 1906.</p> <p><u>Geo-from-metadata:</u> <b>st marys lighthouse</b>. the lighthouse was built in 1833, and is <b>grade ii listed</b>.</p> <p><u>Knowledge-from-metadata:</u> <b>st marys lighthouse</b>. the lighthouse was built in <b>1898</b> and is now a <b>grade ii listed</b> building.</p>
(c)		<p><u>Ground truth:</u> Milngavie railway station, East Dunbartonshire. Opened in 1863 by the Glasgow &amp; Milngavie Junction Railway. View south west at forecourt.</p> <p><u>Standard</u> (Mokady et al., 2021): A car driving down a street past a building.</p> <p><u>No-knowledge:</u> <i>st johns wood church, threapwood. the church</i> was built in 1848 by the architect <i>john soane. grade ii listed</i> building (english heritage building id: link).</p> <p><u>Knowledge-from-image:</u> <i>staines station. staines station</i> is a station on the <i>waterloo to reading line, and the junction of windsor</i>. the station was opened in 1848.</p> <p><u>Geo-from-metadata:</u> <b>milngavie station</b>. the station is located on the <b>milngavie branch</b>. the station is located on the site of the station, and the station was opened in <b>1863</b>.</p> <p><u>Knowledge-from-metadata:</u> <b>milngavie station</b>. the station was opened in <b>1863</b>. it is now part of the <b>abellio scotrail</b>.</p>
(d)		<p><u>Ground truth:</u> Lenzie Old Parish Church. A Category C listed church [Link], dating from 1874.</p> <p><u>Standard</u> (Mokady et al., 2021): A church with a clock tower and a cross on top.</p> <p><u>No-knowledge:</u> <i>st peters church, colwyn bay. st asaph. st pauls church</i> was consecrated in 1887.</p> <p><u>Knowledge-from-image:</u> <i>christ church, heaton norris. christ church</i> was built in 1846. a developed <i>early english style</i>. the church was built in 1846. the church was designed by <i>sir christopher wren</i>.</p> <p><u>Geo-from-metadata:</u> <b>lenzie old parish, lenzie</b>. built in 1846, the church was built in 1846.</p> <p><u>Knowledge-from-metadata:</u> <b>lenzie old parish church , lenzie. lenzie old parish church</b> was built in <b>1874</b> and designed by <b>clarke and bell</b>. it is now a <i>grade ii listed</i>.</p>

Table 5: Examples of the generated captions. Correct geographic references and facts are given in **bold**; incorrect ones are given in *italics*.

Image references:

(a) <https://www.geograph.org.uk/photo/5710322>;

(b) <https://www.geograph.org.uk/photo/4421113>;

(c) <https://www.geograph.org.uk/photo/3278560>;

(d) <https://www.geograph.org.uk/photo/6186571>