

Cross-Lingual and Cross-Cultural Variation in Image Descriptions

Uri Berger

Hebrew University of Jerusalem
University of Melbourne
uri.berger2@mail.huji.ac.il

Edoardo M. Ponti

University of Edinburgh
University of Cambridge
eponti@ed.ac.uk

Abstract

Do speakers of different languages talk differently about what they see? Behavioural and cognitive studies report cultural effects on perception; however, these are mostly limited in scope and hard to replicate. In this work, we conduct the first large-scale empirical study of cross-lingual variation in image descriptions. Using a multimodal dataset with 31 languages and images from diverse locations, we develop a method to accurately identify entities mentioned in captions and present in the images, then measure how they vary across languages. Our analysis reveals that pairs of languages that are geographically or genetically closer tend to mention the same entities more frequently. We also identify entity categories whose saliency is universally high (such as animate beings), low (clothing accessories) or displaying high variance across languages (landscape). In a case study, we measure the differences in a specific language pair (e.g., Japanese mentions clothing far more frequently than English). Furthermore, our method corroborates previous small-scale studies, including 1) Rosch et al. (1976)'s theory of basic-level categories, demonstrating a preference for entities that are neither too generic nor too specific, and 2) Miyamoto et al. (2006)'s hypothesis that environments afford patterns of perception, such as entity counts. Overall, our work reveals the presence of both universal and culture-specific patterns in entity mentions.¹

1 Introduction

Do speakers of different languages talk differently about what they see? This question is important for both machine learning and cognitive science. From a computational perspective, the answer can guide the creation of more diverse vision-and-language datasets (Ye et al., 2023). From a cognitive science

¹We provide our code at github.com/uriberger/cross_lingual_diff_in_descs and a web interface for our dataset at tinyurl.com/5ekabrb5.



	Translated captions	Saliency	
		PERSON.N.01	DANCER.N.01
id	A dancer in a traditional Balinese costume...	$\frac{2}{2} = 1$	$\frac{2}{2} = 1$
	Two dancers dressed in traditional Balinese...		
nl	Asian traditional costume during a performance	$\frac{2}{3}$	$\frac{1}{3}$
	Two people in costume for a performance		
	Traditional Balinese dancers in costumes		

Figure 1: A photo taken in an Indonesian-speaking area, corresponding captions in Indonesian (id) and Dutch (nl) translated to English, and saliency for the PERSON.N.01 and DANCER.N.01 synsets. Saliency is measured as the proportion of captions referring to the synset or any of its descendants (e.g., DANCER.N.01 is a descendant of PERSON.N.01).

perspective, it can help us study cultural effects on perception and language.

To investigate their cross-lingual and cross-cultural variation, we can compare the descriptions of the same images by different communities of speakers. In this setting, differences in semantic content must stem from linguistic rather than perceptual differences.² Crucially, describing the same image takes into account cultural effects (e.g., the saliency of an entity) while excluding environmental ones (e.g., the presence of an entity) (Liu et al., 2021; Hershovich et al., 2022).

²It must be noted that Mitterer et al. (2009) provide evidence of environmental effects on visual perception; however, these effects are limited to low-level details, such as similar colour discrimination.

Previous work investigated which cultures draw more attention to the foreground or the context of a visual scene (Miyamoto et al., 2006), and how often each culture uses negation or certain bigrams at the start of the descriptions (van Miltenburg et al., 2017). Nevertheless, these studies were conducted on a small scale, either with artificial images or focused on just a few related languages.

To induce more solid conclusions, a large dataset of image descriptions in a variety of typologically and areally diverse languages is needed (Ponti et al., 2019). We use Crossmodal3600 (XM3600, Thapliyal et al., 2022), which is both relatively large in scale (3600 images) compared with previous controlled studies and more linguistically diverse (36 languages) compared with any previous image-description dataset.

We develop automatic tools for accurately identifying phrases mentioning entities in image descriptions and map these phrases into WordNet (Miller, 1992) synsets, which function as a proxy for entity categories. We further annotate images with the categories they contain, independent of the descriptions. Finally, we study how entity mentions vary across different languages (see Figure 1 for an example).

Using this annotation as quantitative evidence, we are able to address the following questions:

1. Do **geographically, genetically, or typologically similar** languages share the same saliency of entities?
2. Which entities are **universally** (non-)salient and which ones display the **largest variance** across languages?
3. What is the semantic **granularity** of entities mentioned in different languages?
4. Is the **number** of entities mentioned affected by the familiarity or the complexity of the images?

Our method can be used both as an exploration and a verification tool for cultural effects on object saliency. For the former, we present new findings (Sections 4.2, 4.3, and 4.4). Future work may propose small-scale, controlled studies to verify these findings.

For the latter, we demonstrate how our method can validate previous findings from cognitive studies. First, annotators typically select object cate-

gories that fall in the middle of the hierarchy, with more general (lower) and more specific (higher) categories being less frequent, supporting Rosch et al. (1976)’s hypothesis of basic-level categories (Section 4.5). Second, we put to the test Miyamoto et al. (2006)’s claim that images taken in East Asian locations tend to include more entities compared to Western locations (Section 4.6).

To summarise, our contributions are as follows. First, we present new insights into cross-lingual variation in entity saliency and confirm previous findings on a wider and more diverse set of languages. Second, we release our entity mention dataset to support future research on cross-lingual and cross-cultural variation of image descriptions.

2 Background and Related Work

Our study falls under the umbrella of works that study cultural effects on perception and language. Early studies relied on small-scale, controlled experiments, while the recent rise in the amount of publicly available visual data with human annotations has prompted large-scale and real-world studies harnessing computational methods to investigate these datasets.

2.1 Small-Scale Controlled Studies

Numerous studies detail experiments involving human participants, who are grouped according to their cultural backgrounds and tasked with completing a visual assignment. Subsequently, the outcomes are examined and compared across different groups. Relevant visual tasks include, among many others, colour discrimination (Winawer et al., 2007; Mitterer et al., 2009), identification of a change in an observed scene (Miyamoto et al., 2006), visual illusions (Segall et al., 1963; Caparos et al., 2012; Linnell et al., 2018), and reading facial expressions (Jack et al., 2009).

Most relevant to the present work, Masuda and Nisbett (2001) study differences in image descriptions across cultures. They display animated clips of underwater scenes to Japanese and Americans and subsequently collect their description of the scenes. The Japanese participants provided more statements about contextual information and relationships, rather than the foreground. Senzaki et al. (2014) replicate this result for the Canadian and Japanese, while Senzaki et al. (2016) demonstrate that the effect is more pronounced for adult participants compared to children.

Miyamoto et al. (2006) put forth the hypothesis that these differences are driven by the typical environments encountered by speakers of different cultures. In particular, they show that photos taken in Japan contain more entities according to both human annotators and automated counts, and that American participants attend more on contextual information when primed with Japanese photos than with American photos.

2.2 Large-Scale Studies

Several previous studies have recognised the advantage of utilising existing public image–caption datasets as a test bed for large-scale experiments of cultural effects on perception and language.

Van Miltenburg et al. (2017) compare descriptions generated by English, German, and Dutch annotators on images taken from Flickr30k (Young et al., 2014). They study lexical differences such as the use of negation and the degree of specificity in annotations across languages. Their approach diverges from ours in several key respects. First, we focus on XM3600, which contains culturally diverse images and covers a wider range of languages. Crucially, we systematically study patterns of entity saliency by developing automated tools for entity mention identification; their method instead involves manually inspecting specific images.

Ye et al. (2023) focus on cross-lingual differences in descriptions of 7 languages from XM3600. Their motivation is computational, aiming to demonstrate that annotators from different languages highlight distinct entities, leading to a more comprehensive coverage of entities in multilingual image-caption datasets compared to monolingual ones. This, in turn, offers improved training data for vision-and-language foundation models. This echoes similar considerations from Liu et al. (2021), who created the MaRVL dataset by sourcing images from 5 diverse cultures representing culturally salient entities. In contrast, our motivation is cognitive; we recognise the existence of such cross-lingual variation in image annotation and leverage it as a framework to investigate questions about cross-cultural variation.

2.3 WordNet

WordNet (Miller, 1992) is a large database of synsets arranged into a tree structure of hyponyms and hypernyms (among other lexical relations). Each synset contains all the English words that are synonyms according to one of their senses.

For instance, the synset BANK.N.01 represents the *river bank* meaning of the word *bank* and is the direct child of the SLOPE.N.01 synset, while the synset BANK.N.02 represents the *financial institution bank* meaning and is the direct child of the FINANCIAL_INSTITUTION.N.01 synset.

3 Methods

To identify the entities mentioned in a caption, we use WordNet’s synsets as a proxy for entity categories. We translate all captions to English (Section 3.1), pre-define a target list of synsets (Section 3.2), extract noun phrases that correspond to one of these synsets (Section 3.3) and filter the resulting synset lists to remove synsets corresponding to entities that are not in the image (Section 3.4).

3.1 Captions Translation

We first translate captions from all languages into English. Translation is a necessity if we are to consider a wide variety of languages since the tools required to automatically process the captions (e.g., WordNet and part-of-speech taggers) are only available in a few languages. Nevertheless, the translation process presents certain benefits, notably the elimination of metaphorical expressions containing entities in non-English languages, and thus a possible confounding factor.

Translation quality plays a critical role in our experiments, as incorrect translations can impact our conclusions. We therefore use the Google Translation API³, which demonstrates robust performance in 31 out of the 36 languages in XM3600 according to several studies (Jiao et al., 2023; Enis and Hopkins, 2024).⁴ We exclude the remaining 5 low-resource languages⁵ from our experiments, as the Google Translation API is rated poorly by native speakers of these languages (Benjamin, 2019).

3.2 Synset Selection

We select a subset of synsets to be identified as entities based on their use in XM3600 and their relevance to our study. We start by mapping all noun

³<https://cloud.google.com/translate>

⁴Arabic (ar), Chinese-Simplified (zh), Croatian (hr), Czech (cs), Danish (da), Dutch (nl), English (en), Filipino (fil), Finnish (fi), French (fr), German (de), Greek (el), Hebrew (he), Hindi (hi), Hungarian (hu), Indonesian (id), Italian (it), Japanese (ja), Korean (ko), Norwegian (no), Persian (fa), Polish (pl), Portuguese (pt), Romanian (ro), Russian (ru), Spanish (es), Swedish (sv), Thai (th), Turkish (tr), Ukrainian (uk), Vietnamese (vi).

⁵Bengali (bn), Maori (mi), Cusco Quechua (quz), Swahili (sw), Telugu (te).

phrases in (translated) XM3600 captions to WordNet synsets. Afterwards, we select *root synsets*, i.e., synsets that (a) can be visibly identified as present or absent in an image (unlike, for example, CITY.N.01), (b) have no ancestors in the WordNet hierarchy that meet condition a, and (c) have at least 100 instantiations of the synset or its descendants in the captions, to provide sufficient statistical strength (the list of root synsets is provided in Appendix A). Next, we include *explicit synsets*, i.e., all descendants of the root synsets to which at least 100 noun phrases from XM3600 captions were mapped. Finally, we include all intermediate synsets between a root synset and its descendent explicit synsets (*implicit synsets*). For instance, ANIMAL.N.01 is a root synset and SQUIRREL.N.01 is an explicit synset: we therefore add all the synsets along the path in the tree between the two: CHORDATE.N.01, VERTEBRATE.N.01, MAMMAL.N.01, PLACENTAL.N.01, RODENT.N.01. As a result, we obtain a list of 649 synsets (25 root, 391 explicit, 233 implicit).

3.3 Synset Extraction

We then extract the selected synsets from the captions as follows. We classify the words in the caption with off-the-shelf part-of-speech taggers, extract noun phrases, identify WordNet synsets (if any exist) each noun phrase should be mapped to, and resolve ambiguities when a noun phrase may refer to multiple synsets. We now elaborate on each of these steps. Additional details can be found in Appendix B.

Noun phrase extraction. We assign each word in the captions a part-of-speech tag with Stanza (Qi et al., 2020). Then, we extract noun phrases by identifying all sequences in the caption that consist of one or more consecutive nouns.

Synset identification. We then link each noun phrase to its corresponding WordNet synsets with the NLTK API.⁶ If one of these synsets, or one of its ancestor synsets, is found in our predefined selected synset list (see Section 3.2), we assign the relevant synset to this phrase. Note that, since WordNet may map phrases to multiple synsets, by the end of this phase each noun phrase may be assigned with zero, one, or more synsets.

Resolving ambiguities in synset mapping. If a noun phrase is assigned multiple synsets, we re-

solve this ambiguity through a pre-trained language model—specifically, BERT large uncased (Devlin et al., 2019). Given a caption and candidate synsets s_1, \dots, s_n where $n > 1$, we generate n new versions of the caption, where in sentence c_i we replace the phrase in question with a phrase representing synset s_i . This representative phrase could be the synset name or an example of the synset. BERT then calculates the probability of each synset phrase via masked language modelling. We select the synset with the highest probability.

3.4 Synset Filtering

The list of synsets generated by the aforementioned process may include synsets that correspond to entities not present in the associated image for several reasons. First, a synset might mistakenly appear on the list due to errors of captioning annotators, automatic translation, or our synset extraction method (Section 3.3). Second, annotators may mention entities that are not visible in the image. For example, in the caption *A woman standing with her back to the camera*, the synset CAMERA.N.01 will be extracted even though no camera is actually depicted in the image.

Since we are interested in patterns of entity mentions for entities actually present in the image, we aim to filter out such synsets. We therefore manually annotate each image in the XM3600 dataset, identifying which of the 25 root synsets have instances in the image. We then refine the synset list extracted using the process in Section 3.3 by removing any synsets whose corresponding root synset is not represented in the image.

3.5 Validation

To validate our synset extraction method, we randomly sample 5 captions per language, for a total of 155 captions, manually annotate the synsets mentioned in the image description, and compute the precision and recall of the synset list extracted using our method.

In this sample, our method extracted 247 synsets, 240 of which matched the ground truth identified during annotation, resulting in a precision of 0.97. We manually annotated 247 synsets, 240 of which were correctly predicted by our method, yielding a recall of 0.97.

4 Experiments

In this section we experiment on our entity mentions corpus, extracted from XM3600 captions.

⁶nltk.org/howto/wordnet.html

4.1 Definitions

Each example in XM3600 associates an image with multiple captions in each language $l \in \mathcal{L}$: $\mathcal{D}_l = \{i, \{c_j\}_{j=1}^n\}_{i=1}^{3600}$. For a given entity $o \in \mathcal{O}$ (i.e., WordNet synset) and caption (a sequence of words) $c \in \mathcal{V}^+$, we define the function $f(o, c)$ as returning 1 if o is in the list of entities extracted from c , and 0 otherwise. If a synset is extracted for a given caption, all its ancestors in the WordNet tree are also considered extracted.

Entity saliency. We define the saliency of entity o in image i as the fraction of captions of this image for which o is in the list of extracted entities. Formally, given captions for i in multiple languages $\mathcal{L}_1, \dots, \mathcal{L}_m$,

$$\begin{array}{cccc} c_1^{\mathcal{L}_1}, & c_2^{\mathcal{L}_1}, & \dots, & c_{n_1}^{\mathcal{L}_1} \\ c_1^{\mathcal{L}_2}, & c_2^{\mathcal{L}_2}, & \dots, & c_{n_2}^{\mathcal{L}_2} \\ \dots, & \dots, & \dots, & \dots \\ c_1^{\mathcal{L}_m}, & c_2^{\mathcal{L}_m}, & \dots, & c_{n_m}^{\mathcal{L}_m} \end{array}$$

we denote the saliency of o in i for captions in language $l \in \mathcal{L}$ as

$$\text{saliency}(l, o, i) = \frac{\sum_{j=1}^{n_l} f(o, c_j^l)}{n_l} \quad (1)$$

See Figure 1 for an example. Composing the saliency scores across entities, languages, and images, we obtain a tensor of size $M \in [0, 1]^{|L| \times |O| \times |I|}$. The cell $M_{l,o,i}$ represents the saliency of entity o in language l for image i .

4.2 Geographic and Genetic Effects

We first study whether languages that are geographically close, genetically related, or typologically similar, also share similar patterns in terms of entity saliency. To this end, we correlate two distance metrics: typological distance and saliency distance between languages.

Typological distance. We use `lang2vec` (Littell et al., 2017) to source pre-computed geographical, genetic and featural distances between languages. Featural distance is the overlap between phonemic inventories.

Saliency distance. We define the saliency distance between two languages ($\mathcal{L}_l, \mathcal{L}_k$) as follows. Given the tensor of saliency scores M obtained from Equation 1, we measure saliency distance as the Euclidean distance between language-specific scores: $\sqrt{\sum_o^{|\mathcal{O}|} \sum_i^{|\mathcal{I}|} [(M_{\mathcal{L}_l,o,i}) - (M_{\mathcal{L}_k,o,i})]^2}$

Typological Criterion	r	p	z
Geographic	0.34	0.04*	2.12
Genetic	0.26	0.01*	2.51
Featural	-0.01	0.92	-0.11

Table 1: Mantel test between the saliency distance matrix (Euclidean distance of saliency scores) and typological distance matrices. * denotes statistical significance.

Correlation significance test. For each distance metric, we construct a distance matrix where the (l, k) cell represents the distance between \mathcal{L}_l and \mathcal{L}_k according to that metric. We use the Mantel test (Mantel, 1967) to compute the correlation between the saliency distance matrix and each of the typological distance matrices. Results are presented in Table 1. We find that saliency distance correlates weakly with both geographic and genetic distances based on $r \in [-1, 1]$ scores, but p -values are statistically significant. On the other hand, we find no significant correlation with featural distance.

4.3 Cross-lingual Variation of Entity Saliency

Next, we investigate which entity categories exhibit cross-lingual variation in saliency. In other words, which entities are universally salient or non-salient? The saliency of which ones varies the most across languages?

Let $\mathcal{I}_o \subseteq \mathcal{I}$ identify the subset of images where entity o is visible according to our manual image annotation mentioned in Section 3.4. Then we calculate the global saliency of an entity o and language l as the mean saliency across images $i \in \mathcal{I}_o$.

We plot the resulting global saliency scores in Figure 2: for each entity category, we report their distribution across languages. First, we find that certain entities are universally salient: for instance, ANIMAL.N.01 (0.693), FOOD.N.01 (0.669), and PERSON.N.01 (0.660). Remarkably, these include the only two entities with animacy (Dahl et al., 2000). Conversely, we found that some entities are universally non-salient, e.g., BAG.N.01 (0.109), TIMEPIECE.N.01 (0.093), and JEWELRY.N.01 (0.070). These broadly fall into fashion accessories and apparel.

Finally, Figure 2 lets us identify entities whose global saliency exhibits the highest standard deviation across languages, i.e., SKY.N.01 (0.109), PLANT.N.02 (0.069), MOUNTAIN.N.01 (0.063).




Figure 2: Violin plot of the distribution of saliency scores across languages for each entity category. Saliency scores are averaged across all images containing the entity. Silver lines indicate quartiles.

This suggests that the inclusion of entities commonly located in the background in the descriptions varies between languages, consistent with findings from previous small-scale, controlled studies (e.g., Miyamoto et al., 2006).

4.4 Case Study: English versus Japanese

While the results in Section 4.3 inform us of global cross-linguistic tendencies, a finer-grained analysis could try to determine the differences in saliency between specific pairs of languages. Nevertheless, this raises the concern that the data collected for XM3600 may be insufficient, as it employed only two annotators per language. Thus, any observed effect would reflect the preferences of those individuals rather than a broader trend.

To mitigate this potential drawback, we focus on a case study for a language pair, namely English–Japanese, for which a significantly larger parallel image captioning dataset is available. Specifically, the MSCOCO (Lin et al., 2014) dataset comprises a substantial collection of English image captions, encompassing approximately 120K images, each accompanied by an average of five captions. STAIR-captions (Yoshikawa et al., 2017) is a Japanese version of MSCOCO, wherein Japanese

annotators have provided original captions, without resorting to translation, for each image, yielding five captions per image. Overall, about 2100 annotators contributed to the STAIR-captions dataset.

Similarly to Section 4.3, we define the global saliency for a specific entity category in each language as the mean saliency across all images in MSCOCO in the captions for that language. For each root synset, we calculate the ratio between the English and Japanese per-language saliency. Table 2 plots the top five root synsets in terms of saliency ratio, in both directions. While we only report results for MSCOCO due to the small sample size of XM3600, we do note that similar rankings were found with XM3600, where the synsets with the highest saliency ratios between Japanese and English were CLOTHING.N.01 (3.78) and BODY_PART.N.01 (3.52). As expected, the concepts where we record the highest differences between Japanese and English are also among those with the least saliency cross-linguistically, as they lie at the bottom of Figure 2.

Statistical significance. To assess statistical significance, we conduct a comparison between the lists of global saliency values associated with each

Synset	Saliency ratio
	en/ja
WRITING_IMPLEMENT.N.01*	2.06
KITCHEN_UTENSIL.N.01*	1.43
BUILDING.N.01*	1.42
FURNITURE.N.01*	1.39
BODY_OF_WATER.N.01*	1.27
	ja/en
CLOTHING.N.01*	2.09
BODY_PART.N.01*	1.51
MOUNTAIN.N.01*	1.42
MUSICAL_INSTRUMENT.N.01	1.32
BAG.N.01*	1.18

Table 2: Top saliency ratios in favor of English (top) and Japanese (bottom). * denotes statistical significance.

synset in the two languages. We use the Wilcoxon signed-ranked test (Wilcoxon, 1992) with Bonferroni correction. All entities in Table 2 except MUSICAL_INSTRUMENT.N.01 show significance.

Qualitative Analysis: Clothing. We now zoom in further onto a specific category, showcasing how our method can serve as an exploration tool for identifying preliminary findings that can later be validated through controlled experiments. We focus on clothing-related entities, which showed the highest saliency ratio between Japanese and English. To gain deeper insights into the origin of this disparity, we conduct the statistical significance test for all the explicit synsets that are descendants (see Section 3.2) of the CLOTHING.N.01 synset. We find that only the SHIRT.N.01 synset was significantly more salient in Japanese in both MSCOCO and XM3600. This indicates that differences in high-level categories may be driven by specific (culturally motivated) entities.

4.5 Granularity

Given a hierarchy of categories and their corresponding nominal expressions, from the most generic to the most specific, the level chosen for describing entities is in principle arbitrary. For instance, when describing an image of an armchair, annotators can choose to label it as *furniture*, *chair* or *armchair*, sorted by their depth in the hierarchy in descending order. Rosch et al. (1976) introduced the concept of *basic level categories* which are the




Figure 3: Distribution of depths in the WordNet synset hierarchy across languages.

categories in a hierarchy preferred over more specific subordinates or broader superordinates. If indeed such basic levels exist for a specific language, we would expect the distribution of the levels of mentioned entities to form a two-tailed distribution, with the basic level being used more frequently than the very specific or very general concepts.

We use the depth in the WordNet tree as a proxy for the granularity of an entity category. Figure 3 plots a histogram of synset depths of entities mentioned in each language. Although the number of mentions varies across languages at different levels, all distributions are two-tailed with most mass concentrating in the centre (levels 5–10 in the WordNet hierarchy), in support of the basic-level category hypothesis.

4.6 Entity Count

Does the location of an image affect the number of entities mentioned in its descriptions, due to its familiarity or complexity of the environment? To answer this question, we use the location indicators provided for each image in XM3600, where 100 images per language were captured in regions where that language is spoken. For each language, we first compare the average entity mentions for the 100 images from its relevant regions versus the other 3000 images from abroad.

Figure 4a illustrates that most languages mention around the same number of entities regardless of familiarity, i.e., whether the image was taken at

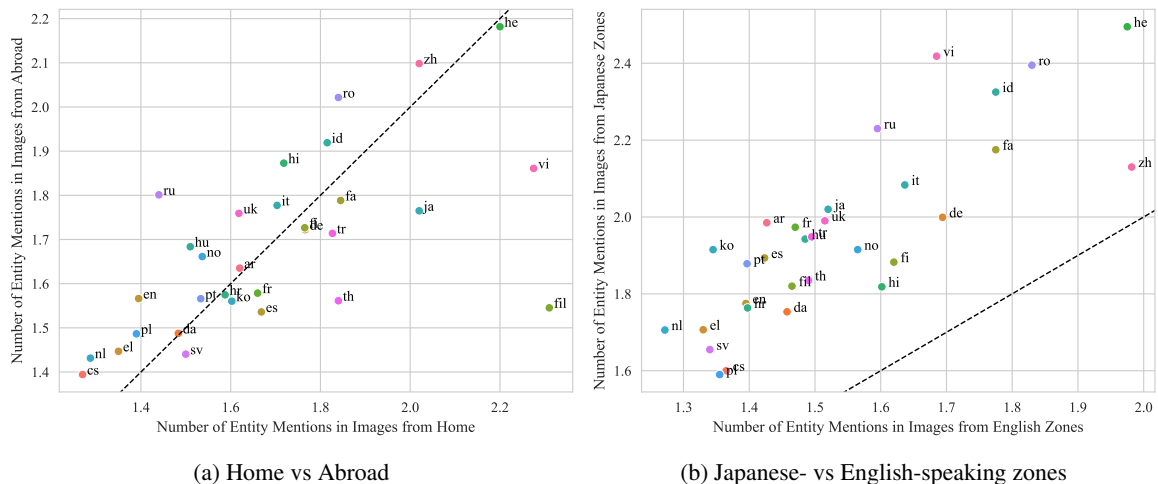


Figure 4: Average number of entities mentioned by speakers of 31 languages in captions of images captured in different locations. The dashed line indicates identity.

home or abroad. The correlation between these two numbers is strong, with Pearson’s $\rho = 0.64$, and significant, with $p < 10^{-4}$. Major exceptions, however, are found in a group of languages including Filipino, Vietnamese, Japanese, and Thai, which resort to significantly more entities when describing images taken at home. Notably, these languages are spoken in (South) East Asia, suggesting a potential areal effect.

In addition, we validate one of the studies of Miyamoto et al. (2006), namely that patterns of attention are (at least in part) afforded by different environments in images (e.g., Japan vs America). In particular, we compare the average number of entities mentioned in XM3600 pictures captured in Japanese- vs English-speaking zones. Figure 4b shows that Japanese images consistently elicited more entity mentions in all 31 languages (including both Japanese and English), supporting Miyamoto et al. (2006)’s finding that speakers of both languages count more entities in Japanese images.

5 Conclusion

We presented the first large-scale empirical study of how the saliency of entities in image descriptions varies across languages. Starting from an existing dataset with culturally diverse images and typologically diverse languages, CrossModal3600, we developed an automated method to extract entity mentions from multilingual captions and we annotated images with the corresponding entities.

We then conducted a series of experiments on the resulting dataset, finding that:

1. Languages in the same family or area have a mild tendency to mention the same entities, whereas no such effect occurs for featural similarity between languages.
2. The saliency of some entities is universally high (animate beings), universally low (clothing accessories), or varies significantly across languages (landscape).
3. We verify Rosch et al. (1976)’s theory of basic-level categories, finding that languages universally prefer entities corresponding to synsets in the middle of the WordNet hierarchy (depths 5-10).
4. The number of entities mentioned is affected by the environment of the image location (e.g., Japan vs Anglosphere), as argued by Miyamoto et al. (2006), rather than its familiarity.

Although we tested our method solely on the XM3600 dataset, which is notable for its cultural diversity, we anticipate that the findings will generalise to other datasets such as MSCOCO, which is more Anglocentric. This expectation is supported by our analysis in Section 4.6, where we observed consistent trends across languages in the number of entities mentioned, regardless of the annotator’s familiarity.

Our method provides an alternative to small-scale, highly controlled experiments, but is meant to complement them, rather than replace them. In

fact, while carefully designed experiments require significant effort in recruiting and training participants, our method is more cost-effective and faster. Therefore, it can serve as an exploratory tool to identify interesting phenomena for later thorough study, and as a large-scale verification for previous findings. We demonstrated both use cases in this work.

Limitations

Translation quality. Translating captions from all languages into English offers certain advantages (see Section 3.1). Nevertheless, there are also some drawbacks. One concern is that translation quality may differ across languages, potentially affecting our findings. Additionally, if a specific entity category does not exist in English, it might be translated into a more general category, which could influence the analysis in Section 4.5. Lastly, translation may alter cultural nuances, potentially affecting the distribution of mentioned objects.

Nonetheless, we opted to include translation in our pipeline. Previous studies that avoided translating captions into English (e.g., van Miltenburg et al., 2017) were limited to a small number of languages, as they relied on the availability of language-specific analysis tools like parsers. This highlights a trade-off between maintaining original captions and expanding language coverage. Given our focus on cultural diversity, we prioritised broader language coverage in this trade-off.

Recognising this limitation, we took several measures to address the issues outlined above. First, we excluded low-resource languages from our analysis due to their low translation quality, as discussed in Section 4. Second, we manually verified that the entities identified by our pipeline were indeed present in the images, as detailed in Section 3.4.

Number of annotators per language. As mentioned in Section 4.4, the XM3600 data collection involved only two annotators per language. Consequently, the observed phenomena for a specific language might be significantly influenced by the personal preferences of its annotators. To mitigate this effect, we focused primarily on cross-linguistic patterns observed across all 31 languages. In one Section (4.4), where we did study individual languages, we used an additional, much larger dataset.

Image diversity. While examining the XM3600 images, we observed that many images depict the

same event. For example, 17 images taken in locations where Ukrainian is spoken (17% of all Ukrainian images) are from a single event (a large protest), all showing people protesting. This raises the concern that the limited variety in the scenes depicted by XM3600 images might impact our results. Nonetheless, we use XM3600 as it is the only multilingual and culturally diverse image-caption dataset that is currently available.

Ethics Statement

We use publicly available resources in our experiments, in accordance with their license agreements. The datasets are fully anonymised and do not contain personal information about the caption annotators or any information that could reveal the identity of the photographed subjects.

Social impact. While our findings explore different cultures, it is important to emphasise that cultures are constantly evolving. These results should be viewed as a “snapshot” of these cultures at a specific moment in time, rather than as representations of permanent characteristics.

References

- M Benjamin. 2019. Teach you backwards: An in-depth study of google translate for 108 languages. *Kamusi Project International*.
- Serge Caparos, Lubna Ahmed, Andrew J Bremner, Jan W de Fockert, Karina J Linnell, and Jules Davidoff. 2012. Exposure to an urban environment alters the local bias of a remote culture. *Cognition*, 122(1):80–85.
- Osten Dahl et al. 2000. Animacy and the notion of semantic gender. *Trends in linguistics studies and monographs*, 124:99–116.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **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. Association for Computational Linguistics.
- Maxim Enis and Mark Hopkins. 2024. From llm to nmt: Advancing low-resource machine translation with claude. *arXiv preprint arXiv:2404.13813*.
- Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza

- Fierro, Katerina Margatina, Phillip Rust, and Anders Sogaard. 2022. [Challenges and strategies in cross-cultural NLP](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6997–7013, Dublin, Ireland. Association for Computational Linguistics.
- Rachael E Jack, Caroline Blais, Christoph Scheepers, Philippe G Schyns, and Roberto Caldara. 2009. Cultural confusions show that facial expressions are not universal. *Current biology*, 19(18):1543–1548.
- Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing Wang, Shuming Shi, and Zhaopeng Tu. 2023. Is chat-gpt a good translator? yes with gpt-4 as the engine. *arXiv preprint arXiv:2301.08745*.
- 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 *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer.
- Karina J Linnell, Andrew J Bremner, Serge Caparos, Jules Davidoff, and Jan W de Fockert. 2018. Urban experience alters lightness perception. *Journal of Experimental Psychology: Human Perception and Performance*, 44(1):2.
- Patrick Littell, David R. Mortensen, Ke Lin, Katherine Kairis, Carlisle Turner, and Lori Levin. 2017. [URIEL and lang2vec: Representing languages as typological, geographical, and phylogenetic vectors](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 8–14, Valencia, Spain. Association for Computational Linguistics.
- Fangyu Liu, Emanuele Bugliarello, Edoardo Maria Ponti, Siva Reddy, Nigel Collier, and Desmond Elliott. 2021. [Visually grounded reasoning across languages and cultures](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10467–10485, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Nathan Mantel. 1967. The detection of disease clustering and a generalized regression approach. *Cancer research*, 27(2_Part_1):209–220.
- Takahiko Masuda and Richard E Nisbett. 2001. Attending holistically versus analytically: comparing the context sensitivity of japanese and americans. *Journal of personality and social psychology*, 81(5):922.
- George A. Miller. 1992. [WordNet: A lexical database for English](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York, February 23–26, 1992*.
- Holger Mitterer, Jörn M Horschig, Jochen Müsseler, and Asifa Majid. 2009. The influence of memory on perception: Its not what things look like, its what you call them. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(6):1557.
- Yuri Miyamoto, Richard E Nisbett, and Takahiko Masuda. 2006. Culture and the physical environment: Holistic versus analytic perceptual affordances. *Psychological science*, 17(2):113–119.
- Edoardo Maria Ponti, Helen O’Horan, Yevgeni Berzak, Ivan Vulić, Roi Reichart, Thierry Poibeau, Ekaterina Shutova, and Anna Korhonen. 2019. [Modeling language variation and universals: A survey on typological linguistics for natural language processing](#). *Computational Linguistics*, 45(3):559–601.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online. Association for Computational Linguistics.
- Eleanor Rosch, Carolyn B Mervis, Wayne D Gray, David M Johnson, and Penny Boyes-Braem. 1976. Basic objects in natural categories. *Cognitive psychology*, 8(3):382–439.
- Marshall H Segall, Donald T Campbell, and Melville J Herskovits. 1963. Cultural differences in the perception of geometric illusions. *Science*, 139(3556):769–771.
- Sawa Senzaki, Takahiko Masuda, and Keiko Ishii. 2014. When is perception top-down and when is it not? culture, narrative, and attention. *Cognitive Science*, 38(7):1493–1506.
- Sawa Senzaki, Takahiko Masuda, Akira Takada, and Hiroyuki Okada. 2016. The communication of culturally dominant modes of attention from parents to children: a comparison of canadian and japanese parent-child conversations during a joint scene description task. *PloS one*, 11(1):e0147199.
- Ashish V. Thapliyal, Jordi Pont Tuset, Xi Chen, and Radu Soricut. 2022. [Crossmodal-3600: A massively multilingual multimodal evaluation dataset](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 715–729, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Emiel van Miltenburg, Desmond Elliott, and Piek Vossen. 2017. [Cross-linguistic differences and similarities in image descriptions](#). In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 21–30, Santiago de Compostela, Spain. Association for Computational Linguistics.
- Frank Wilcoxon. 1992. Individual comparisons by ranking methods. In *Breakthroughs in statistics: Methodology and distribution*, pages 196–202. Springer.

- Jonathan Winawer, Nathan Witthoft, Michael C Frank, Lisa Wu, Alex R Wade, and Lera Boroditsky. 2007. Russian blues reveal effects of language on color discrimination. *Proceedings of the national academy of sciences*, 104(19):7780–7785.
- Andre Ye, Sebastin Santy, Jena D Hwang, Amy X Zhang, and Ranjay Krishna. 2023. [Cultural and linguistic diversity improves visual representations](#). *ArXiv preprint*, abs/2310.14356.
- Yuya Yoshikawa, Yutaro Shigeto, and Akikazu Takeuchi. 2017. [STAIR captions: Constructing a large-scale Japanese image caption dataset](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 417–421, Vancouver, Canada. Association for Computational Linguistics.
- 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 for Computational Linguistics*, 2:67–78.

A Synset List

The list of WordNet synsets used in this study is as follows.

Root synsets.

ANIMAL.N.01,
BAG.N.01,
BODY_OF_WATER.N.01,
BODY_PART.N.01,
BUILDING.N.01,
CLOTHING.N.01,
ELECTRONIC_EQUIPMENT.N.01,
FOOD.N.01,
FURNITURE.N.01,
HAND_TOOL.N.01,
HOME_APPLIANCE.N.01,
JEWELRY.N.01,
KITCHEN_UTENSIL.N.01,
MOUNTAIN.N.01,
MUSICAL_INSTRUMENT.N.01,
PERSON.N.01,
PLANT.N.02,
PLAYTHING.N.01,
SKY.N.01,
SUN.N.01,
TABLEWARE.N.01,
TIMEPIECE.N.01,
VEHICLE.N.01,
WEAPON.N.01,
WRITING_IMPLEMENT.N.01

Implicit and explicit synsets. Please see supplementary materials for the full list of implicit and explicit synsets used in this study.

B Synset Extraction

Here, we elaborate on the synset extraction algorithm. Specifically, we delve into the steps for synset identification and resolving ambiguities.

B.1 Synset Identification

To identify synsets, we utilise the WordNet ontology tree but make some modifications to adjust it to image captions:

- We remove certain synsets representing meanings not used by annotators in captions. For example, the word *snake* is mapped to the meaning SNAKE.N.01 (the animal), but also to the meaning SNAKE.N.02 (a treacherous person). As the second meaning is never employed in image captions, we excluded it from the tree.

- We unify certain synsets, such as FOOD.N.01 (e.g., MILK.N.01) and FOOD.N.02 (e.g., BAKED_GOODS.N.01) as we would like to consider them as a single category.
- The WordNet ontology encodes hyponymy–hypernymy relations between synsets. When synsets are conceptually placed under one root synset but are more visually relevant to another, we relocate them to the latter. For example, COUPLE.N.01 (a pair of people) is located under the GROUP.N.01 subtree. For our purposes we would like to consider mentions of couples as persons and we therefore move it to the PERSON.N.01 subtree.

C Annotation Guidelines

In this section, we describe the annotation guidelines provided in each of the datasets used in this study.

XM3600. The annotation guidelines of XM3600, taken from the original paper, were as follows.

To guide your caption generation, imagine that you are describing the image to a visually impaired friend. The caption should explain the whole image, including all the main objects, activities, and their relationships. The objects should be named as specifically as practical: For example when describing a young boy in a picture, “young boy” is preferred over “young child”, which in turn is preferred over “person”.

Note: the goal is to generate captions that would be labeled as “Excellent” under the Rating guidelines above, but raters should not copy captions from the first phase. We want the raters to generate the captions on their own.

We outline here a procedure that you should try and follow when writing your image caption. Note that not all these steps may be applicable for all images, but they should give you a pretty good idea of how to organize your caption. We will make use of the first image in the table below (the one with the young girl smiling). Note: It is acceptable to make assumptions that are reasonable as long as they don't contradict the information in the image (e.g.: in the second image below, we use “families” in captions 1 and 3 because there seems to be a mix of children and adults though it

is not perfectly clear. So it is a reasonable assumption to make and nothing in the image contradicts it. However it is also ok to use “people”.)

1. Identify the most salient objects(s)/person(s) in the image; use the most informative level to refer to something (i.e., “girl” rather than “child” or “person”); in the example image: “girl”
2. Identify the most salient relation between the main objects; example “girl standing in front of the whiteboard”
3. Identify the main activity depicted; in the example image: “smiling” as an activity (note that this can also be an attribute of the girl), or “standing” as an activity
4. Identify the most salient attributes of the main object(s)/person(s)/activity(es); in the example image: “smiling” and “young” as attributes for the girl
5. Identify the background/context/environment in which the scene is placed; in the example image: “classroom”
6. Put everything together from steps 1-5 above; for the example image: “a smiling girl standing in a classroom”, or “a young girl smiling in a classroom”.

MSCOCO. The following instructions were provided to MSCOCO annotators:

- Describe all the important parts of the scene.
- Do not start the sentences with “There is”.
- Do not describe unimportant details.
- Do not describe things that might have happened in the future or past.
- Do not describe what a person might say.
- Do not give people proper names.
- The sentences should contain at least 8 words.

STAIR-Captions. The following instructions were provided to STAIR-Captions annotators:

1. A caption must contain more than 15 letters.
2. A caption must follow the da/dearu style (one of writing styles in Japanese).
3. A caption must describe only what is hap-

pening in an image and the things displayed therein.

4. A caption must be a single sentence.
5. A caption must not include emotions or opinions about the image.