

The Case for Perspective in Multimodal Datasets

Marcelo Viridiano¹, Tiago Timponi Torrent^{1,2}, Oliver Czulo³,
Arthur Lorenzi Almeida¹, Ely Edison da Silva Matos¹, Frederico Belcavello¹

¹ FrameNet Brasil Lab, Graduate Program in Linguistics, Federal University of Juiz de Fora

² Brazilian National Council for Scientific and Technological Development – CNPq

³ Institute for Applied Linguistics and Translatology, Universität Leipzig

{barros.marcelo, arthur.lorenzi}@estudante.ufjf.br, {tiago.torrent, ely.matos, fred.belcavello}@ufjf.br,
oliver.czulo@uni-leipzig.de

Abstract

This paper argues in favor of the adoption of annotation practices for multimodal datasets that recognize and represent the inherently perspectivized nature of multimodal communication. To support our claim, we present a set of annotation experiments in which FrameNet annotation is applied to the Multi30k and the Flickr 30k Entities datasets. We assess the cosine similarity between the semantic representations derived from the annotation of both pictures and captions for frames. Our findings indicate that: (i) frame semantic similarity between captions of the same picture produced in different languages is sensitive to whether the caption is a translation of another caption or not, and (ii) picture annotation for semantic frames is sensitive to whether the image is annotated in presence of a caption or not.

Keywords: multimodal datasets, annotation setup, multilingual annotation, multimodal annotation, perspective, frame semantic analysis

1. Introduction

Multimodal datasets that combine textual and visual information are gaining popularity in Natural Language Processing tasks such as multimodal machine translation (Specia et al., 2016; Elliott et al., 2017; Elliott, 2018), multimodal lexical translation (Lala and Specia, 2018), visual sense disambiguation (Gella et al., 2016), grounded representation of lexical meaning (Siberer and Lapata, 2014), and lexical entailment detection (Kiela et al., 2015). For all of those tasks – and especially for tasks of multimodal translation – the general claim is that using textual data in combination with the “ground truth” information provided by a visual mode would improve the performance of multimodal models, allowing them to surpass baselines and equivalent monomodal models.

In this paper, we describe the first steps into exploring what frame semantic analyses can tell us about multiperspectivity in multimodal datasets annotated in different languages and in different annotation settings. Based on the Flickr 30k dataset (Young et al., 2014) and its variants – Multi30k (Elliott et al., 2016) and Flickr 30k Entities (Plummer et al., 2015) – we conduct a set of experiments in which automatic frame semantic annotation for image captions and manual frame annotation for images are assessed for their semantic similarity. Image captioning and translation of captions were done by humans. Comparisons adopt both (1) a multilingual perspective with English and Portuguese originals and English-Portuguese translations of image descriptions and (2) multisetup perspective with English image annotations with or without image captions visible to annotators.

The contributions of this paper are two-fold:

- the multilingual setting is a first probing into how similar or different perspectives may be in different languages without digging deeper into how systematic differences might be.
- the multisetup perspective tests the assumption that visual information holds some sort of unbiased “ground truth”.

In the remainder of this paper, we discuss, in section 2, how Frame Semantics and its computational implementation – FrameNet – incorporate perspective to the core of semantic representations. Next, in section 3, we explain compilation, translation and annotation of the corpus used for the experiments devised in section 4. Results and discussion are presented in sections 5 and 6, respectively, while section 7 finalizes the paper.

2. Frame Semantics and Perspective

The main idea behind Fillmore’s frame semantics (Fillmore, 1982) is that human beings understand the meaning of a linguistic expression against the cognitive backdrop of a schematized scene, i.e. a *frame*. A frame is defined as “any system of concepts related in such a way that to understand any one of them you have to understand the whole structure in which it fits” (Petrucci, 1996, 1).

A central notion in frame semantics is perspective. One of Fillmore’s (1985) classic examples refers to the semantic distinction between *coast.n* and *shore.n* in English: Both refer to the `Relational_natural_features` frame, describing the stretch where land mass and sea water meet. While the former lexeme describes the view from the land, the latter is perspectivized from the point of view of the sea.

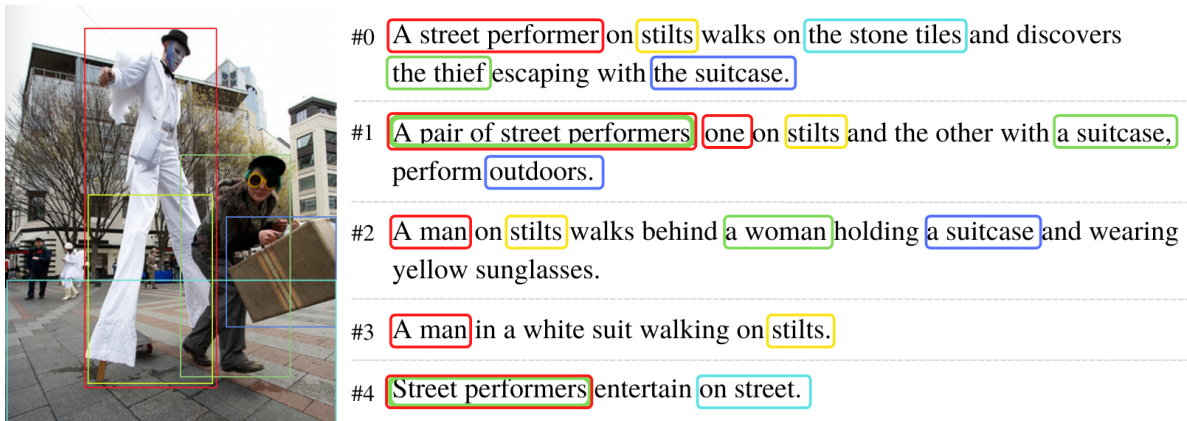


Figure 1: Captions and coreference chains from the Flickr 30k Entities dataset.

Berkeley FrameNet (Fillmore et al., 2003; Johnson et al., 2016) was the first lexicographic incarnation of Frame Semantics, building a frame-based lexicon to cover the general vocabulary of English. For each recorded frame that is lexically realizable, the database lists:

1. **Frame Elements (FE):** represent the corresponding participants and objects of a frame. Based on the type of supporting information they contribute, frame elements can be categorized as core and non-core. For the frame `Commerce_buy`, for instance, `BUYER` and `GOODS` are core FEs while `PLACE` is non-core.
2. A list of **Lexical Units:** The list of known lexical units (single or multiple words) that can evoke the frame. For the frame `Commerce_buy`, this includes *buyer.n*, *client.n* and *purchase.n*.
3. **Frame Relations:** Each frame is connected to related frames with edges denoting the kind of relation (eg. *precedence*, *inheritance*, etc.) that exists between the interconnected frames. `Commerce_buy` *inherits* from `Getting` and is a *perspective on* `Commerce_goods-transfer`

Currently, there are FrameNet projects for several languages including German, Japanese and Brazilian Portuguese. A part of these groups forms an initiative for multilingual research in frame semantics, Global FrameNet¹.

Lately, full-text annotation, translation analysis (Czulo, 2017; Torrent et al., 2018, i.a.) and multimodal annotation (Belcavello et al., 2020) have seen increasing interest in frame semantics, shifting the annotation focus. In lexicographic annotation, annotation practice usually is aimed at making clear decisions on categories. Annotation accuracy is often measured in relation to a gold standard or by inter-annotator agreement. A lesson from frame-semantic annotation, however, is that not all annotation cases can be decided and multiple

interpretations are possible. This issue has seen more attention in the above-mentioned more recent trends. In a phrase such as *Kinder, die dieses Jahr in die Schule gehen* ‘children who start going to school this year’, the annotation of *Schule* ‘school’ in this context could either be argued to be a `LOCALE_BY_USE` (a place with a certain purpose) or to refer to `EDUCATION_TEACHING` (a domain). In the end, one can wonder whether a decision for either of the two has an added benefit, or whether this actually represents the range of possible interpretations. Indeed, already in the early design of his theory, Fillmore points out that frames (in the terminology at that time ‘scenes’) are associated with and activate each other (Fillmore, 1975, 124). Considering this as well as the fact that frames are, in general, prototypical categories, they mostly cannot be understood as sharply distinct, necessarily discrete categories.

3. The Framed Multi30k Dataset

In recent years, several projects have been expanding the popular dataset for sentence-based image description Flickr30k (Young et al., 2014): a multimodal dataset containing 31,783 images of everyday activities, events and scenes, each paired with five different English captions providing clear descriptions of the salient entities and events. The Multi30k dataset (Elliott et al., 2016) is a multilingual expansion of the Flickr30K with five German original conceptual descriptions (Hodosh et al., 2013) crowdsourced independently of the original English captions. The German translations of the English captions were created by professional translators. The Flickr30k Entities dataset (Plummer et al., 2015) adds image-to-text relations by manually annotating bounding boxes that assign region-to-phrase correspondences, linking mentions to the same entities in images with lexical items in the five captions describing that image (Figure 1).

For the experiments reported in this paper, we rely on yet another extension of Flickr 30k: the Framed Multi30k dataset (Torrent et al., 2022), which augments both Flickr30k Entities and Multi30k datasets with (a) semantic annotation based on the network of

¹<https://www.globalframenet.org/>

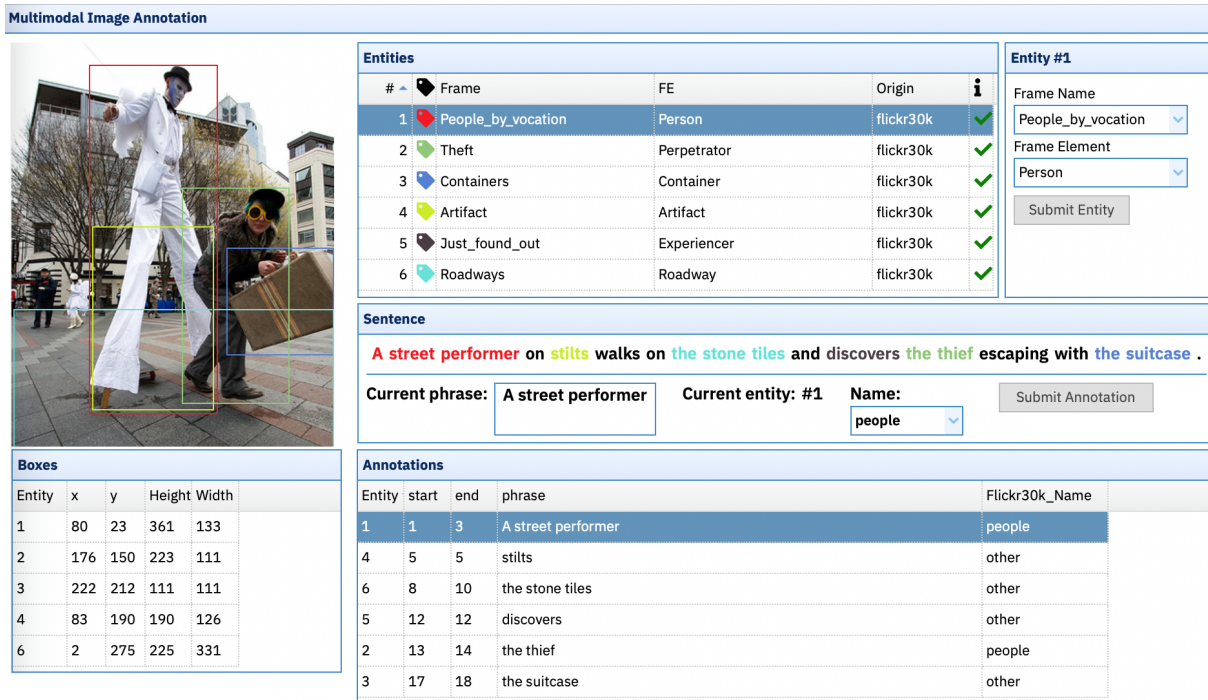


Figure 2: User interface for the multimodal annotation tool used for building the Framed Multi30k dataset.

frames and relations from FrameNet Brasil and (b) Brazilian Portuguese captions to the images. Framed Multi30k includes both English-Portuguese Translation (PTT) of the English Original captions (ENO) and Brazilian Portuguese Original descriptions (PTO) for each image in Flickr 30k. For the translation task, grad students majoring in translation studies were presented with one of the original English captions and instructed to translate the descriptions sticking as closely as possible to the English source sentence. For the task of creating original descriptions, native speakers of Brazilian Portuguese majoring in Language and Linguistics were presented with an image and instructed to write an original conceptual description (Hodosh et al., 2013). The instruction said to describe only what is depicted in the scene – its entities, their attributes, and relations – as opposed to providing additional background information that cannot be obtained from the image alone, such as about the situation, time, or location in which the image was taken.

As for the annotation of images, Framed Multi30k associates, via manual annotation, the bounding boxes in Flickr 30k Entities with frames and frame elements, using the annotation interface shown in Figure 2. Note that, based on the original Flickr 30k Entities dataset organization, the bounding boxes presented for annotation are only the ones grounded in the referential noun phrases found in the caption. This is to say that different captions of one same picture usually have different sets of bounding boxes.

Because Framed Multi30k is still under construction, for this paper, we selected a random sample of 2,000 images for which there already are both the PTT and

the PTO captions, in addition to the ENO captions. Those images and their corresponding captions comprise the dataset used in the experiments. Portuguese original captions are usually shorter than the original English captions, both in terms of the number of tokens (27,421 PTO x 38,881 ENO), types (3,747 PTO x 3,809 ENO), and characters (128,659 PTO x 152,837 ENO). When comparing the original descriptions with their translations to Brazilian Portuguese, we observe that, despite having fewer tokens (35,074 PTT x 38,881 ENO), translated sentences are more varied in terms of types (4,712 PTT x 3,809 ENO) and have a higher character count (165,724 PTT x 152,837 ENO). The properties of PTO and PTT captions are similar to the ones found for German original and translated captions when the Multi30k dataset was built (Elliott et al., 2016). Inflectional properties of German, in comparison to English, may in part be responsible for the differences; they may also reflect the oft-made observation that translations have a tendency to be longer than originals.

4. Experiments

Experiments were set up so as to assess frame semantic similarity across languages and across communicative modes. In both cases, we use the cosine similarity (CS) between the frame semantic representations generated either automatically by a semantic parser or manually, via annotation.

The CS algorithm used to measure frame semantic similarity between two annotations comprises three stages:

1. building an associate table using frames from FrameNet;

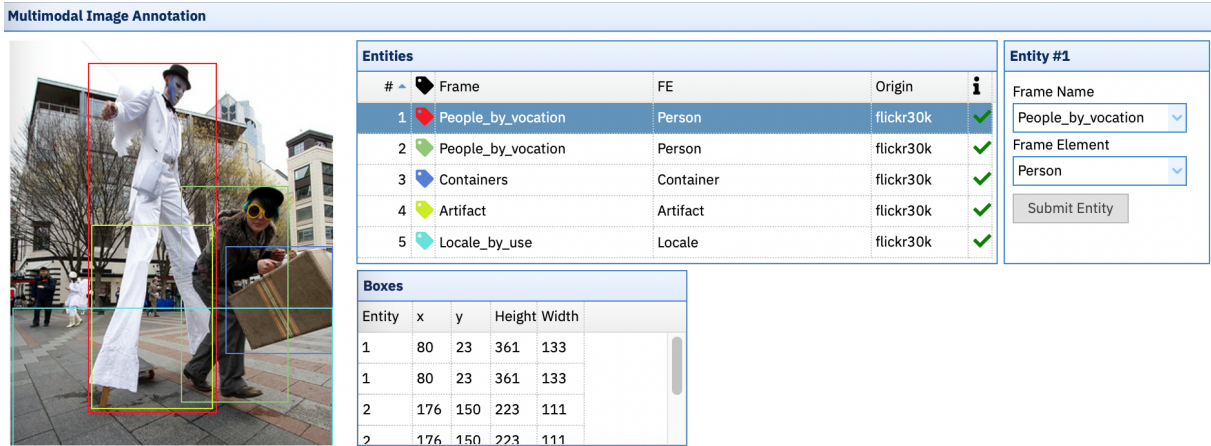


Figure 3: User interface of the multimodal annotation tool adapted for the second task.

2. building associative arrays associating each annotation – automatic or manual – with the directly evoked and related frames in the FrameNet database;
3. measuring cosine similarity between associative arrays.

In this application, the FrameNet network is represented as an Acyclic Directed Graph – the FN graph – where each node (frame) is associated with another node (frame) higher in the same hierarchy. The FN graph is used to build the associative table for each frame. This table indicates a relatedness metric and is calculated using Spread Activation (SA) (Gouws et al., 2010). The SA algorithm models an iterative energy propagation process from one or more nodes to other nodes in a graph in three stages: (i) pre-adjustment, (ii) spreading, and (iii) post-adjustment (Crestani, 1997). Before the spreading stage, the energy value for each node was calculated during the pre-adjustment stage. Energy decay was calculated for the value of the node so that this value is within the $[0,1]$ interval. The calculated value was then output to the neighboring nodes. Post-adjustment was not used, since the FN graph is acyclic and the FN hierarchies do not comprise many levels.

From the associative table of each frame present in an annotation, the associative arrays were built. Each index in the array corresponds to an associated frame and the value of the index indicates the activation level of that frame. When comparing two annotations, for example a1 and a2, the array for a1 is completed with the frames evoked by a2 but not evoked by a1, and vice versa. A zero value is computed for each of those “completion frames”. Finally, the relatedness between the two annotations is measured using the standard CS between two associative arrays.

CS of associative arrays were used as a metric for assessing semantic similarity of both captions in a crosslinguistic perspective, and communicative modes

in a crossmodal perspective. We describe the experimental designs used for each perspective next.

4.1. Semantic similarity across languages

To evaluate semantic similarity of image captions across languages, we compare the semantic frame representation of the original English captions (ENO) provided by the Flickr 30K dataset with the Brazilian Portuguese translations of those captions (PTT) and the original Portuguese captions produced for the same pictures (PTO). Comparisons are pairwise and are assessed using the cosine similarity between associative arrays generated for each caption. Therefore, three measures are extracted from this experiment, namely, cosine similarities for the ENO x PTT, ENO x PTO and PPT x PTO pairs.

Frame-evoking lemmas from all captions in each language were automatically retrieved using DAISY (Torrent et al., 2022), a disambiguation algorithm that uses the network of semantic frames, Frame Elements (FEs) and Lexical Units (LUs) from the FrameNet database to assign the correct frames to each lexical item based on the context provided by each sentence. By treating word forms, lexemes, LUs and frames as nodes in a graph, and attributing values to each node and also to each candidate lemma, the disambiguation algorithm uses a spread activation search method (Diederich, 1990; Tsatsaronis et al., 2007) to calculate the energy decay of each lemma as it propagates through the nodes in the network. This method takes into account not only how far a specific node is from the beginning nodes – meaning, how much energy is lost as connections need to be traversed in order to reach that specific node – but also how it is activated by liked neighboring nodes, receiving more energy, which helps determine its relative importance in the network.

Table 1 presents the differences among the multilingual corpora. Brazilian Portuguese translations have a lower average number of lemmas than the original English captions (19.44 ENO x 17.54 PTT) and also fewer frame-evoking lemmas (17.78 ENO x 12.38

PTT). Original descriptions in Portuguese also have a lower average number of lemmas than the English captions (19.44 ENO vs. 13.71 PTO lemmas) and also evoked fewer frames (17.78 ENO vs. 9.98 PTO frames). Considering the normalized number of frames per lemma, the original English caption ratios ($M = 0.92$, $SD = 0.33$) are significantly higher than the Portuguese translations ($M = 0.71$, $SD = 0.28$) and original descriptions ($M = 0.75$, $SD = 0.32$), with $t(3996) = 21.28$, $p < 0.001$ and $t(3996) = 16.35$, $p < 0.001$, respectively, using Welch’s t-test (Welch, 1947). This difference is a consequence of the broader lexical coverage of English in relation to Portuguese. It is worth noting, however, that this difference does not impact the validity of the comparisons between cosine similarities because the vectors representing frames evoked by a sample in English are always paired with samples of the other corpora. Any variation in a comparison between a translation and a original sentence in Portuguese is caused by their own differences.

Portuguese originals also have a higher frame:lemma ratio, although this difference is lower than the others ($t(3996) = -3.83$, $p < 0.001$). In this case, since both corpora are on the same language, the difference can be explained by the lexical choices of the annotators: the translation corpus contains 2,755 singletons, while the original Portuguese annotations have 2,118.

		ENO	PTT	PTO
Frames	avg. #	17.78	12.38	9.98
	stdev	8.04	5.96	4.86
Lemmas	avg. #	19.44	17.54	13.71
	stdev	6.47	6.21	5.48
Frame:Lemma	avg.	0.92	0.71	0.75
	stdev	0.33	0.28	0.32

Table 1: Counts and ratios for annotated frames and for lemmas

4.2. Semantic similarity across modes in different annotation setups

For assessing similarity of semantic annotation of different communicative modes and how it may be influenced by the annotation setup, the results of two annotation tasks for tagging bounding boxes in the Flickr 30k Entities dataset for frames and frame elements were used.

The first is the one originally designed for building the Framed Multi30k dataset, where native speakers of Brazilian Portuguese, who are also fluent in English, are assigned the task of enriching the multimodal dataset with FrameNet frame and frame element tags while analyzing image-caption pairings – see Figure 2. Before being assigned this task, annotators were trained on the guidelines and the quality of their an-

notation work was manually checked for a first batch comprised of six hundred annotations, one hundred images from each of the six annotators involved. After validating the control quality batch, a second batch with 275 images per annotator was assigned to them. For this first task, annotators were instructed to follow the FrameNet annotation guidelines and to assign a semantic frame and a frame element to the objects in the bounding boxes.

In the example annotation (Fig. 2), the lexical items “A street performer,” “stilts,” “the stone tiles,” “discovers,” “the thief,” and “the suitcase,” from the original English caption created to describe this image, are correlated to bounding boxes containing the entities referred to by those lexical items. For the first bounding box – colored red and correlated with the also red sentence segment “A street performer,” – the annotator assigned the frame `People_by_vocation`, which contains words for individuals as viewed in terms of their vocation, and the core frame element `PERSON`. For the second bounding box – colored green and correlated with the green sentence segment “the thief” – the annotator assigned the frame `Theft`, which is evoked by lexical items describing situations in which a perpetrator takes goods from a victim, and the core frame element `PERPETRATOR` – the person (or other agent) that takes the goods away.

The second annotation task was specifically devised for this paper. We asked a different group of five annotators to annotate a subset of 1,000 images within the original 2,000 images sample for frames and frame elements in the bounding boxes associated with each image. This time, however, annotators were not presented with the captions. In other words, they should apply FrameNet labels while having only the visual mode as reference. Given the nature of the visual corpus, annotators were instructed to only assign Entity frames and the related core frame element. The same manually created bounding boxes from Flickr 30k Entities were used to determine which objects from each image should be annotated. This time, however, all bounding boxes associated with each image were presented (Fig. 3). As in the first task, the quality of the annotations was assessed by manually checking a subset before allowing annotators to proceed to the full task.

5. Results

In the following subsections we present the results for each of the experiments described in section 4.

5.1. Similarity of semantic representations across languages

Average cosine similarity measures for each caption type pair – ENO x PTT, ENO x PTO and PTT x PTO – are presented in Table 2. The distributions of those similarities are shown in Figure 4.

Taking into consideration how the cosine similarity distributions shown in Figure 4 approximate a normal

	ENO		PTO	
	avg. cos	stdev	avg. cos	stdev
ENO	-	-	0.33	0.14
PTT	0.51	0.14	0.43	0.2
PTO	0.33	0.14	-	-

Table 2: Average cosine similarity between Associative_arrays built for the frame semantic representations of ENO, PTT and PTO captions

distribution and the almost equivalent variances, Student’s and Welch’s t-test were used to verify the significance of the differences, according to the variables variance. The cosine similarities between PTT and ENO ($M = 0.51$, $SD = 0.14$) were significantly higher than those between PTO and ENO ($M = 0.33$, $SD = 0.14$), $t(1998) = 41.78$, $p < 0.001$. Additionally, the average similarities between PTO and PTT ($M = 0.43$, $SD = 0.2$) are higher than those between PTO and ENO, with test statistic $t(1998) = 19.98$, $p < 0.001$. At the same time, those similarities are significantly smaller than those between PTT and ENO, with $t(1998) = -13.71$, $p < 0.001$.

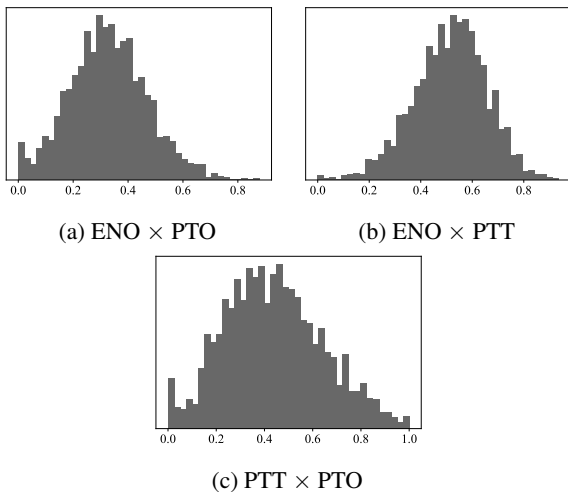


Figure 4: Distribution of cosine similarity values between ENO, PTT and PTO.

5.2. Similarity of semantic representations across modes in different annotation setups

Average cosine similarity measures and the standard deviation for image annotations compared to the original English textual annotations are presented in Table 3. The distributions of those similarities are shown in Figure 5. The annotations made with a reference caption in English (VWC), when compared against ENO, had higher cosine similarities ($M = 0.43$, $SD = 0.13$) than the ones annotated without any reference (VWoC)

($M = 0.38$, $SD = 0.12$), with Student’s t-test statistic $t(998) = 8.64$, $p < 0.001$.

	ENO	
	avg. cos	stdev
VWC	0.43	0.13
VWoC	0.38	0.12

Table 3: Similarity for image frame annotation setups with and without captions present.




Figure 5: Distribution of cosine similarity values between ENO, VWC and VWoC.

6. Discussion

The similarity coefficients from the multilingual comparison indicate that in terms of perspective of annotation, PTT seems to be somewhere between ENO and PTO, with a comparable distance to either. In some way, this was to be expected, given that the translation brief required translators to make a close rendition of the English original. With the cumulative analysis performed here, we cannot make the claim that this a clear case of *shining through* (Teich, 2003), i.e. a case of source text features being over-represented in the target text. For the domain of motion events, systematic framing differences between languages are well documented (Talmy, 2000; Slobin, 2004). In the image annotation, however, we will find a vastly broader annotation including people and artifacts, and we do not yet know whether any of the involved languages has a preference for a certain framing or interpretation of these categories in image descriptions. Also, the variation we see between ENO and PTO might be due to randomness, with PTT being an intended close rendition of ENO.

Besides a better understanding of potential framing preferences in a language, we also need to cross-check against variation that can happen within one language, or in other words: If we had two sets of annotations for English by different annotators, might they be as similar/distant as ENO/PTO or rather as ENO/PTT? While the numbers presented here may be first indicators, clearly more elaborate setups and evaluations are needed to dig deeper into the questions of framing preferences between languages and translation effects.

As to the latter, it is not a given that influence from the source language is the only factor to be taken into account: As shown, e.g., in (Vandevoorde, 2020), *normalization* (Baker, 1995), i.e. the (over-)adherence to target language conventions is another semantic effect that can be witnessed in translation, or re-framings due to an open list of factors such as those described, i.a., in (Slobin, 2005; Rojo and Valenzuela, 2013; Czulo, 2017; Ohara, 2020).

As for the image annotation setups, the experiments indicate that, under a model that takes perspective into account, labels assigned to images are not to be taken as some sort of “ground truth” representation. Annotators taking part in the experiment tended to frame the image according to the caption, which is shown by a significant higher cosine similarity between ENO and VWC. The very examples in Figures 2 and 3 give an indication of that: while the person holding the suitcase was annotated as a thief when in presence of the caption, they were tagged as a person in the absence of the caption clue.

This is an indication that the role of images as proxies for “ground truth” in multimodal datasets is, at the very least, limited, if one considers that meanings are relativized to perspectivized scenes, as pointed out by Fillmore (1977). The main issue at stake here is that, in general, image annotation in multimodal datasets involve the assignment of labels from a categorization system that does not encode perspective.

In cases where the annotation is carried out by humans, which is the case for the Flickr 30k Entities dataset, the categories available for annotation of the bounding boxes are very coarse-grained and include only: people, clothing, body parts, animals, vehicles, instruments, scene, not visual and other. In this scenario, the distinction between thief and person, for instance, would be subsumed under the “people” tag. The scenario is even more concerning when we consider that, among the 559,767 bounding boxes in the Flickr 30k Entities dataset, 182,136 (32.53%) received the “people” tag and 138,658 (24.77%) received the “other” tag. In the cases where image annotation is conducted automatically, using computer vision algorithms such as YOLOv3 (Redmon and Farhadi, 2018), for example, the core problem remains unchanged. Such systems are trained on datasets such as MS-COCO (Lin et al., 2014) or Open Images (Kuznetsova et al., 2020). In both cases, the categories used for tagging bounding boxes are organized in ontologies that do not encode perspective either. Last but not least, even in the case of Open Images, where, on top of the categories assigned to each image, attributes and relations can also be assigned as triplets, it is not made clear whether the different perspectives on them are encoded.

Because the way humans interpret the entities in an image may be influenced by the text accompanying it – as the results in Table 3 suggest –, the current configuration of image annotation systems may limit the

role of images in downstream tasks such as multimodal machine translation. Even in Flickr 30k, where captions are conceptual descriptions of images, the relation between image and text is not equivalent to that of an absolute fact and one possible description of it. Images may too accommodate different perspectives and accounting for those differences is key for assuring that the relation between the image and the text is preserved, for instance, in a translated sentence.

7. Conclusion(s) and further work

The experiments and analyses described in this contribution have produced two results:

1. Frame semantic similarity for image captions in different languages are sensitive to whether a description is a translation or not.
2. Semantic similarity was also influenced by the annotation setup in that presenting captions with images to be annotated produced higher semantic similarity across modes, indicating that image description data cannot be assumed to be an “independent source” or “ground truth” of semantic information.

The former result needs further investigation in order to test for its generalizability. This would include extending the experiment to different language pairs as well as comparing such results with variation within a language, e.g., with multiple captions per image in one language. Also, the question of whether there are language-specific framing preferences at play can only be answered by means of deeper analyses of the frame semantic annotation of the original captions.

The latter result in a sense is analogous to the observation that the translation of image captions seems to carry, in parts, influence from an “anterior” - for the case of translation, this the description to be translated, with regard annotation setup, shown to annotators. If priming is to be assumed as a relevant factor in both cases, then the observation does not come as a surprise, but at the same time commands caution for future annotation setups. Priming as a factor could be tested for in order to corroborate this assumption.

8. Acknowledgements

Authors acknowledge the support of the Graduate Program in Linguistics at the Federal University of Juiz de Fora. Research presented in this paper is funded by CAPES PROBRAL grant 88887.144043/2017-00 and CNPq grants 408269/2021-9 and 315749/2021-0. Viridiano’s research was funded by CAPES PROBRAL PhD exchange grant 88887.628830/2021-00.

9. Bibliographical References

Baker, M. (1995). Corpora in translation studies: an overview and some suggestions for future research. *Target*, 7(2):223–243. Hier werden die verschiedenen Typen von Korpora diskutiert.

- Belcavello, F., Viridiano, M., Diniz da Costa, A., Matos, E. E. d. S., and Torrent, T. T. (2020). Frame-based annotation of multimodal corpora: Tracking (a)synchronies in meaning construction. In *Proceedings of the International FrameNet Workshop 2020: Towards a Global, Multilingual FrameNet*, pages 23–30, Marseille, France, May. European Language Resources Association.
- Crestani, F. (1997). Application of spreading activation techniques in information retrieval. *Artificial Intelligence Review*, 11(6):453–482.
- Czulo, O. (2017). Aspects of a primacy of frame model of translation. In S. Hansen-Schirra, et al., editors, *Empirical modelling of translation and interpreting*, number 6 in Translation and Multilingual Natural Language Processing, pages 465–490. Language Science Press, Berlin.
- Diederich, J. (1990). Spreading activation and connectionist models for natural language processing. 16(1):25–64.
- Elliott, D., Frank, S., Sima'an, K., and Specia, L. (2016). Multi30K: Multilingual English-German image descriptions. In *Proceedings of the 5th Workshop on Vision and Language*, pages 70–74, Berlin, Germany, August. Association for Computational Linguistics.
- Elliott, D., Frank, S., Barrault, L., Bougares, F., and Specia, L. (2017). Findings of the second shared task on multimodal machine translation and multilingual image description. *arXiv preprint arXiv:1710.07177*.
- Elliott, D. (2018). Adversarial evaluation of multimodal machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2974–2978, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Fillmore, C. J., Johnson, C. R., and Petruck, M. R. (2003). Background to framenet. *International journal of lexicography*, 16(3):235–250.
- Fillmore, C. J. (1975). An alternative to checklist theories of meaning. In *Annual Meeting of the Berkeley Linguistics Society*, volume 1, pages 123–131.
- Fillmore, C. J. (1977). The case for case reopened. In P. Cole et al., editors, *Syntax and Semantics Volume 8: Grammatical Relations*, pages 59–81. Academic Press.
- Fillmore, C. J. (1982). Frame Semantics. In Linguistics Society of Korea, editor, *Linguistics in the morning calm*. Hanshin Publishing Co., Seoul, South Korea. pages: 111 – 138.
- Fillmore, C. J. (1985). Frames and the semantics of understanding. *Quaderni di Semantica*, 6(2):222–254.
- Gella, S., Lapata, M., and Keller, F. (2016). Unsupervised visual sense disambiguation for verbs using multimodal embeddings. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 182–192, San Diego, California, June. Association for Computational Linguistics.
- Gouws, S., van Rooyen, G.-J., and Engelbrecht, H. A. (2010). Measuring conceptual similarity by spreading activation over Wikipedia’s hyperlink structure. In *Proceedings of the 2nd Workshop on The People’s Web Meets NLP: Collaboratively Constructed Semantic Resources*, pages 46–54, Beijing, China, August. Coling 2010 Organizing Committee.
- Hodosh, M., Young, P., and Hockenmaier, J. (2013). Framing image description as a ranking task: Data, models and evaluation metrics. *Journal of Artificial Intelligence Research*, 47:853–899.
- Johnson, C. R., Schwarzer-Petruck, M., Baker, C. F., Ellsworth, M., Ruppenhofer, J., and Fillmore, C. J. (2016). Framenet: Theory and practice. Technical report, International Computer Science Institute.
- Kiela, D., Rimell, L., Vulić, I., and Clark, S. (2015). Exploiting image generality for lexical entailment detection. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 119–124, Beijing, China, July. Association for Computational Linguistics.
- Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Kolesnikov, A., Duerig, T., and Ferrari, V. (2020). The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*.
- Lala, C. and Specia, L. (2018). Multimodal lexical translation. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May. European Language Resources Association (ELRA).
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. (2014). Microsoft coco: Common objects in context. In David Fleet, et al., editors, *Computer Vision – ECCV 2014*, pages 740–755, Cham. Springer International Publishing.
- Ohara, K. (2020). Finding corresponding constructions in English and Japanese in a TED talk parallel corpus using frames-and-constructions analysis. In *Proceedings of the International FrameNet Workshop 2020: Towards a Global, Multilingual FrameNet*, pages 8–12, Marseille, France, May. European Language Resources Association.
- Petruck, M. R. L. (1996). Frame Semantics. In Jef Verschueren, et al., editors, *Handbook of pragmatics*, pages 1–8. John Benjamins, Philadelphia.
- Plummer, B. A., Wang, L., Cervantes, C. M., Caicedo, J. C., Hockenmaier, J., and Lazebnik, S. (2015). Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models.

- In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 2641–2649.
- Redmon, J. and Farhadi, A. (2018). Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.
- Rojo, A. and Valenzuela, J. (2013). Constructing meanings in translation: The role of constructions in translation. In Ana Rojo et al., editors, *Cognitive linguistics and translation: advances in some theoretical models and applications*, number 23 in Applications of Cognitive Linguistics, pages 283–310. De Gruyter Mouton, Berlin.
- Silberer, C. and Lapata, M. (2014). Learning grounded meaning representations with autoencoders. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 721–732, Baltimore, Maryland, June. Association for Computational Linguistics.
- Slobin, D. I. (2004). The many ways to search for a frog: linguistic typology and the expression of motion events. In S. Strömquist et al., editors, *Relating Events in Narrative: Typological Perspectives*, pages 219–257. Lawrence Erlbaum Associates, Mahwah, N.J.
- Slobin, D. (2005). Relating Narrative Events in Translation. In Dorit Diskin Ravit et al., editors, *Perspectives on Language and Language Development: Essays in Honor of Ruth A. Berman*, pages 115–129.
- Specia, L., Frank, S., Sima'an, K., and Elliott, D. (2016). A shared task on multimodal machine translation and crosslingual image description. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 543–553, Berlin, Germany, August. Association for Computational Linguistics.
- Talmy, L. (2000). *Toward a cognitive semantics: Vol. II: Typology and process in concept structuring*. MIT Press, Cambridge, MA.
- Teich, E. (2003). *Cross-linguistic variation in system and text. A methodology for the investigation of translations and comparable texts*, volume 5 of *Text, Translation, Computational Processing*. Mouton de Gruyter, Berlin/New York.
- Torrent, T. T., Ellsworth, M., Baker, C., and Matos, E. E. d. S. (2018). The Multilingual FrameNet Shared Annotation Task: a Preliminary Report. In Tiago Timponi Torrent, et al., editors, *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, pages 62–68, Paris, France, may. European Language Resources Association (ELRA).
- Torrent, T. T., Matos, E. E. d. S., Belcavello, F., Viridiano, M., Gamonal, M. A., Costa, A. D. d., and Marim, M. C. (2022). Representing context in framenet: A multidimensional, multimodal approach. *Frontiers in Psychology*, 13.
- Tsatsaronis, G., Vazirgiannis, M., and Androutopoulos, I. (2007). Word sense disambiguation with spreading activation networks generated from the-sauri. In *IJCAI*, volume 27, pages 223–252.
- Vandevoorde, L. (2020). *Semantic differences in translation: Exploring the field of inchoativity*. Language Science Press. OCLC: 1229399854.
- Welch, B. L. (1947). The generalization of ‘student’s’ problem when several different population variances are involved. *Biometrika*, 34(1-2):28–35.
- Young, P., Lai, A., Hodosh, M., and Hockenmaier, J. (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.