

Lutma: a Frame-Making Tool for Collaborative FrameNet Development

**Tiago Timponi Torrent^{1,2}, Arthur Lorenzi¹, Ely Edison da Silva Matos¹,
Frederico Belcavello¹, Marcelo Viridiano¹, Maucha Andrade Gamonal³**

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 LETra, Graduate Program in Linguistic Studies, Federal University of Minas Gerais

{tiago.torrent, ely.matos, fred.belcavello}@ufjf.br, {arthur.lorenzi, barros.marcelo}@estudante.ufjf.br,
mgamonal@ufmg.br

Abstract

This paper presents Lutma, a collaborative, semi-constrained, tutorial-based tool for contributing frames and lexical units to the Global FrameNet initiative. The tool parameterizes the process of frame creation, avoiding consistency violations and promoting the integration of frames contributed by the community with existing frames. Lutma is structured in a wizard-like fashion so as to provide users with text and video tutorials relevant for each step in the frame creation process. We argue that this tool will allow for a sensible expansion of FrameNet coverage in terms of both languages and cultural perspectives encoded by them, positioning frames as a viable alternative for representing perspective in language models.

Keywords: FrameNet, Collaborative Frame Creation, Perspective, Multilinguality

1. Introduction

As we have claimed elsewhere, “language can be vague, messy and variable because humans cooperate while using it” (Torrent, 2021), and communicative cooperation critically involves sharing culture, and perspectives on it and inviting our interlocutors to reconstruct them as needed. We claim that computational models of language should embrace the characteristics of human language and not try to overcome them by the mathematical manipulation of linguistic form patterns extracted from large datasets alone (Bender and Koller, 2020). Moreover, we claim that the FrameNet model (Fillmore and Baker, 2010; Torrent et al., 2018b) is a good candidate for representing those characteristics, provided that it is extended to cover more languages and dialects.

Doing so is not trivial if one considers how FrameNets have been built so far – see, among others, Fillmore et al. (2003), Ohara et al. (2004), Torrent and Ellsworth (2013) and Dannélls et al. (2021). All of them relied on a relatively small team of intensively trained linguists, who invested a considerable amount of time building machine-readable frames and associating linguistic material with them based on corpus evidence representing a small number of languages and an even smaller number of variants within one same language. Due to this *modus operandi*, FrameNets usually present limited coverage, which is frequently pointed out by NLP practitioners as a reason for not using them in their applications. Nonetheless, the knowledge accumulated in the past three decades of FrameNet development allows for a methodological turn, presented in this paper in the form of a software tool: Lutma¹.

Lutma is semi-constrained, tutorial-based tool for fostering a community of distributed frame builders who

will be able to enrich FrameNet with more diverse perspectives grounded on their own languages and language variants, enhancing its coverage and representativeness. In the remainder of the paper we start, in section 2, by making the case in favor of a larger, multilingual and multidialectal FrameNet, by contrasting FrameNet with Large Language Models (LLMs). Next, we present Lutma in section 3. Section 4 presents an example of culturally grounded frame creation. Finally, section 5 closes the paper by presenting the current limitations and future developments planned for collaborative FrameNet building.

2. The Case for a Collaborative Frame Building Tool

In recent years, discussions about perspectives have flourished in NLP, motivated by a variety of reasons, among which the problems of reducing multiple labels into a single ground-truth label, the ambiguous nature of many NLP tasks and bias encoded on large language models (Aroyo and Welty, 2015; Artstein and Poesio, 2008; Bender et al., 2021). Integrating multiple perspectives into models, however, is far from trivial, as specific model characteristics must be taken into consideration.

In white-box systems, one can take many different approaches to prevent bias and guarantee representation of multiple perspectives, namely through pre- and post-processing, but also by altering the models’ internals (Ntoutsis et al., 2020). For that reason, one can argue that they are particularly useful in cases where curators and developers want to make sure that the model – or the dataset – are not encoding bias (Criado and Such, 2019).

In the case of black-box models, the strategies used in supervised and unsupervised models to introduce multiple perspectives are very distinct. For the former –

¹<https://lutma.frame.net.br>

represented by Machine Translation (MT), Question Answering (QA), Named Entity Recognition (NER) and many other models – Basile et al. (2021) defend the adoption of *data perspectivism*, moving away from gold standard datasets and instead adopting different points of view for each object in the data. In practical terms, this variation is obtained by assigning the annotation task of a single data record to multiple human subjects. The way those different perspectives are integrated into a model determines whether the implementation follows a *weak* or a *strong perspectivist* approach. Any system that aggregates – using majority, average, etc. – those different perspectives into a single label is considered to follow the *weak perspectivist* method. When, instead, the model is adapted to handle the outputs from multiple annotators, it is classified as a *strong perspectivist* one. This added complexity, and consequently the amount of work, is outweighed by the ethical principles being followed, and in many cases, also leads to better performance (Basile et al., 2021).

Unsupervised methods, on the other hand, pose different types of challenges. We turn to them next.

2.1. Perspective in LLMs

Unsupervised methods are primarily represented by LLMs, such as BERT (Devlin et al., 2019), GPT-3 (Raffel et al., 2020) and T5 (Brown et al., 2020), which cannot integrate multiple perspectives into their training by expanding the number of gold standard labels per object. Rather, pre-existing bias in the data needs to be addressed during corpus curation and preprocessing.

Bender et al. (2021) state that models trained on large and uncensored text datasets encode biases that lead to unethical technology, also calling for investment to curate those datasets as a means to avoid such biases. Rogers (2021) claims that curation already takes place in the datasets feeding language models, and poses the question of what type of curation would be most effective to avoid harmful biases and improve models' abilities to understand language.

One key aspect of LLMs is that, regardless of whether data curation takes place or not, perspective understanding is treated as a byproduct of the mathematical manipulation of linguistic form. In other words, the machine's role is to "figure out" different perspectives solely from variations in form, while the researcher is responsible for making sure that the training data is representative of those perspectives. Within this framework, whether or not multiple points of view are being considered by the model depends on the dataset size and the quality of the data sample, much like the overall model performance.

We, in turn, claim that, instead of manipulating only input data, NLP systems should integrate a cognitively and culturally-oriented model that represents alternative perspectives on the meaning of linguistic forms.

We also claim FrameNet to be the most suitable model for doing that.

2.2. Perspective in FrameNet

Perspective is a core aspect of Frame Semantics (Fillmore, 1982; Fillmore, 1985). Directly related to the classic Fillmorean proposition that *meanings are relativized to scenes* is the idea that different perspectives may be taken on that scene (Fillmore, 1977). The classic example revolves around the perspectives on the commercial transaction event. Fillmore (1977) demonstrates that different English verbs – namely *buy*, *sell*, *cost* and *charge* – adopt the perspectives of the different participants in the scene – Buyer, Seller, Goods and Money. Each perspective structures the scene in a particular fashion, so that some of the participants that are core to one perspective – e.g. Buyer and Goods in the perspective lexicalized in *buy.v* – may not be central to others – e.g. in the perspective lexicalized in *charge.v*, to which the Seller and the Money are central.

As the implementation of Frame Semantics, FrameNet models frames in terms of the elements in them, the coreness status of those element and the relations established between frames. The model also includes the lexical items evoking the frames. In this context, there are three aspects of FrameNet that may contribute to enrich and diversify language models:

Cognitively-based: FrameNet was initially proposed as lexical database inspired in Frame Semantics (Ruppenhofer et al., 2016). Lexical units are linked to frames, which in turn are linked to other frames. According to Frame Semantics, in order to understand a single frame, one needs to understand the structure in which it fits (Petrucci, 1996). Frames are schematic representations of concepts based on recurring experiences against which the meanings of lexical units are relativized (Fillmore, 1977). This means that, instead of representing words with a single vector, or a list of senses, they can be associated with multiple activation patterns in the network. It also means, as demonstrated by Torrent et al. (2022), that FrameNet structure captures contextual information, namely commonsense knowledge.

Socially contextualized: Since frames are schematic representations of concepts, they are also used to represent socially construed entities and events, for example. The existence of specific types of frame relations also facilitates the creation of frames that may represent particular views on the same event. For example, research on Japanese noun-modification constructions by Matsumoto (2010) demonstrates that the societal grounding of the interaction and the purpose of the discourse influence on the grammar of noun-modification, claiming that interactional frames play a central role in the comprehension of those constructions.

Multilingual: Research on Frame Semantics adopting a contrastive multilingual approach has demonstrated that different languages may lexicalize different per-

spectives on a given scene, one of the parade examples being the study of verbs of emotion in Spanish vs. English (Subirats-Rüggeberg and Petruck, 2003). Because there are frameNets under development for a number of languages, those differences are also captured via either the Global FrameNet Shared Annotation task (Torrent et al., 2018a) or the Multilingual FrameNet database alignment (Gilardi and Baker, 2018; Baker and Lorenzi, 2020). Such efforts allow for the construction of a single database supporting lexical units from any language, but at the same time not restricted to universally applicable frames. Even if some culturally specific frames are created, they can still be linked to the global network of frames. One of the challenges in this undertaking is to merge resources that were built independently for years. Nonetheless, for those working with low-resource languages (LRLs), Global FrameNet and Multilingual FrameNet are of great help, since they allow users to focus on frame-evoking units in their language, instead of modeling frames and their relations. This possibility of focusing on the language itself rather than on the underlying frame structure can be an important tool to reduce the gap of LRLs in NLP (Magueresse et al., 2020; Cruz and Cheng, 2020; Lakew et al., 2020).

Despite the three features described above, FrameNet still lacks coverage in terms of both number of languages and cognitive domains included in the model. Therefore, a fourth feature must be pursued if FrameNet is to be presented as an alternative to representing perspective in language models:

Community-based: Expanding contributions to FrameNet language resources is the main solution to increase the speed at which those resources evolve. And although some challenges come with a community-based approach, it brings an overall positive balance to research pursuits. Having a global community increases the possibility that contributors from varied backgrounds work together. This directly impacts the quality of the resource, as more languages will be supported and also new frames from specific cultural backgrounds will be created.

Implementing this feature is the purpose of Lutma, which is presented next.

3. Lutma: a Frame-Maker Tool

Lutma is one of the steps towards building an extended FrameNet resource in which language is not isolated from human cognition and social backgrounds. The project is part of the Global FrameNet effort, a collaboration between labs and affiliate researchers of twelve different countries, with the goal of facilitating the sharing of findings and research data, as well as building partnerships for the development of novel research. The way Lutma differentiates itself from the various tools used by the FrameNets around the world is that it has two design goals not shared by the others: first, that frames and LUs must have a clear indication of

the languages/cultures in which they belong; second, that the user experience must be aimed towards people interested in FrameNet, but with little or no training on frame creation. These goals are both aligned with the idea of building a cognitively-based and collaborative language resource.

The main challenge of extending FrameNet lies on the fact that, until recent, it has been mostly reliant on specialists. To address this challenge, frame creation in Lutma follows a linear approach, much like a wizard, where to advance, users are asked a certain amount of information about the records. The idea of having separate steps for different pieces of data during the process also allows the system to run consistency and redundancy checks, making sure that users who are not experienced with the concepts can create frames with adequate quality.

Frame creation is separated into two execution flows, one for lexical and the other for non-lexical frames. Both are presented in the flow diagram in Figure 1. The first one starts with a lemma search: the user inputs the system with the part-of-speech and language of a lemma that will evoke the frame they want to create. The system checks if this lemma already exists in the database and, if not, searches for synonyms using Open Multilingual WordNet (OMWN) (Bond and Foster, 2013). If any of the synonyms is an LU in the database, the evoked frames are displayed and the user can decide whether a new frame needs to be created or if a new LU will be created for an existing frame. A second check is executed, but instead of looking for synonyms, Lutma searches for words in other languages with similar spelling in the same OMWN synset. This is a last measure to prevent redundancy, taking advantage of the existing data in other languages, although the multilingual synsets are limited to a small subset of languages.

When a new frame needs to be created, the user is guided to the next screen where they select its root type. They can choose from event, entity, relation, attribute, state or undefined (when it doesn't fit any of the previous). The selected type is used by the system to make suggestions of possible frame elements during the process, e. g., "Direction" or "Material" for relation and entity frames, respectively. In the next step, the user needs to fill out frame names and definitions and once again the system checks for duplicate names and if the name follows certain standards, such as scenario frame names ending with "_scenario" or state frames having names following the "Being_x" or "x_state" patterns.

In the following screen, users can create relations between their new frame and existing frames in Lutma's database. When creating those relations, the system allows users to choose whether they want to map FEs from the other frame into the frame (in the case of an inheritance relation, this is required). After storing those relations, the next step consists of creating FEs. At this point, the user's frame may already have some FEs be-

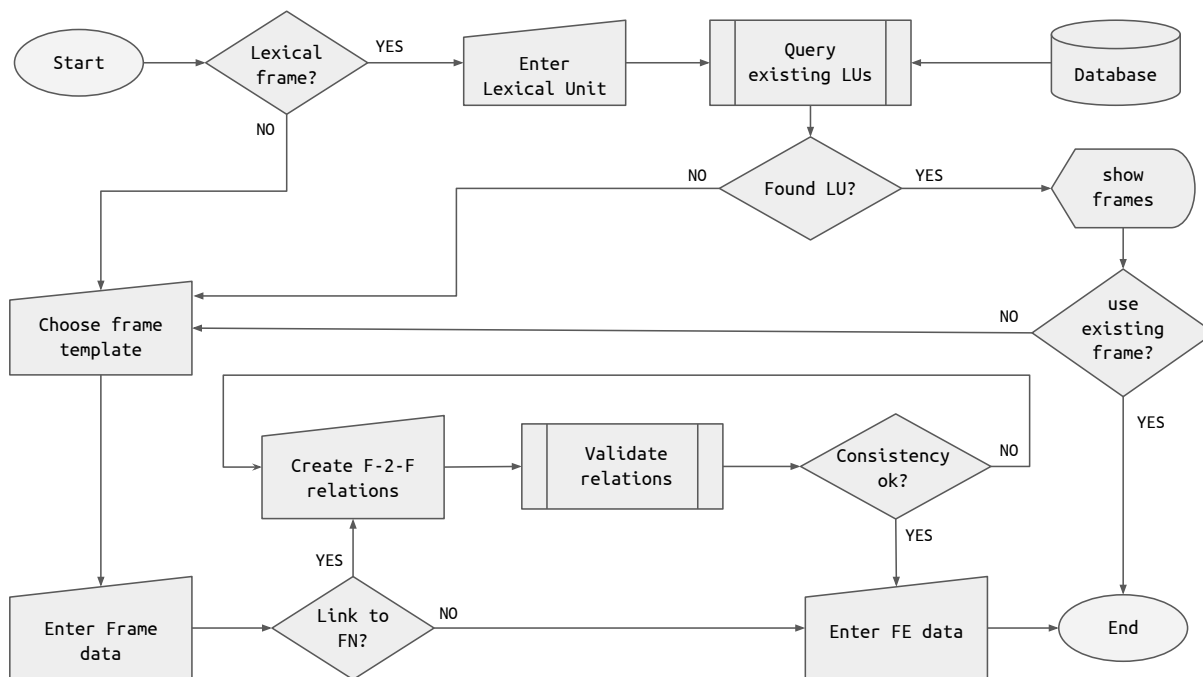


Figure 1: Lutma's frame creation flow. The diagram encompasses the creation of lexical and non-lexical frames.

cause of the frame relations, but Lutma still suggests more FEs based on the frame type. The rest of the FEs can be manually created and at least one is required to proceed. In the final step, before a summary is shown, the system asks for FE relation information. The creation summary displays all of the information related to the frame and to finish creation users must provide an example sentence for the lemma that was informed at the beginning of the process, as well as inform if it incorporates one of the FEs. After that, the frame is registered in the database.

The flow for non-lexical frames shares most of the steps with the lexical one. The first difference is the absence of a lemma search (because the frame will not be evoked by lexical material). The second is the need to inform the frame language in frame type selection step. We opted for this solution to guarantee that those frames would not be treated as universal, despite the fact that they can be associated to multiple languages. These two execution flows separated into steps, along with the automatic quality checks run by the system, improve user experience by reducing the chances that a mistake will be made. However, when trying to build a community around this type of resource, we expect users with varying knowledge about FrameNets. For that reason, Lutma also integrates tutorials into its interface. Those tutorials can be categorized into two types, those about the system's interface, that explain how users can achieve certain goals, and those about Frame Semantics and FrameNet. The former are automatically displayed at the first time an user logs in to Lutma, while the latter can be accessed using the interface elements at any step.

These theoretical and practical tutorials found in all

system screens can be further divided into two types. The simpler ones are displayed in the form of dialog boxes that are rendered every time a user clicks on one of the UI elements. For example, when searching for a lexical unit in the database, a user can click on the "Enter lemma" label of the search field to open a dialog card explaining what is the definition of lemma used in Lutma. Those small texts are useful for users that want to remember how certain concepts are defined in FrameNet. They can also be useful to users with different backgrounds in Linguistics, facilitating the comprehension of how those same concepts are represented in Lutma.

When users need more information than provided in the dialog boxes, the system presents a link to a video tutorial on the specific subject. These videos are always paired with one of the previously mentioned dialog boxes as a way of presenting a longer discussion for the concept. They are also tailored to a broader audience and explain essential concepts in a simple manner, using various examples. In total, six videos were produced, ranging from three to nine minutes in duration. Most videos address more than one topic relevant to the process of frame creation and, because of that, are linked to different parts of the systems, but effectively covering all of the dialog boxes. For future work, these tutorials could be expanded even further, including more topics and references to relevant publications. Last but not least, to make sure that the data created by collaborators will benefit other contributors or projects, we opted for a copyleft license, namely GPLv3². With this licensing scheme, not only the data is made acces-

²www.gnu.org/licenses/gpl-3.0.html

sible for any interested party, but improvements made outside of its ecosystem can be potentially reintegrated into the database.

Since the project is quite new, there are still limitations that need to be addressed before a full release for the community. The final section discusses those and summarizes our contributions. Before turning to them, though, we present an example demonstrating how culturally specific frames can be created in Lutma and integrated to the existing FrameNet frames.

4. The Brazilian Way Frame

To provide an example of how Lutma may aid in the expansion of FrameNet to include culturally grounded frames, while linking them to the existing database, we present the creation of the `Brazilian_way` frame, evoked by LUs such as *jeitinho.n*, *malandragem.n* and *gingado.n* in Brazilian Portuguese (br-pt). Those LUs can be literally translated into English (en) as ‘little way’, ‘trickery’ and ‘waddle’, respectively. Their culturally grounded meaning is quite different though.

According to DaMatta (1986), *jeitinho* is characterized as the space Brazilians find between what one can do and what one cannot do within a normative system. Such a system can be institutionalized in the Judiciary, or may correspond to implicit social norms that should be followed by everyone. When the concept of *jeitinho* is brought into play, one seeks to solve some private problem by adopting some behavior that makes the solution easier and/or faster. Such a behavior is inadequate under the strict observation of the norm regulating the problem-solving task. Nonetheless, by bringing *jeitinho* into play, such inadequacy is relativized.

The fact that the main lexical unit in this frame is in the diminutive form is not coincidental. The br-pt expression *dar um jeito* corresponds roughly to the en verb *fix* in sentences like *Alguém deu um jeito no problema do visto*, meaning that someone fixed the visa problem. On the other hand, the expression *dar um jeitinho*, by using the diminutive form of *jeito.n*, introduces a sense of empathy and proximity. Hence, in sentences like *Alguém deu um jeitinho no problema do visto*, what is being said is that someone found a non-standard, possibly illegal way of solving the visa problem. This way of solving the problem may involve a favor being granted by some authority on that matter or even the corruption of such an authority. Moreover, Schroder and Silva (2020) demonstrate that the conceptualization of *jeitinho.n* involves the idea of making rules flexible via an exchange of favors.

Given this scenario, if we want to create a frame for *jeitinho.n*, we would start by searching the Global FrameNet Database for this LU. Since there is no frame for this LU in br-pt or for any translation of it in another language, we proceed to the frame creation process. First, we select the *event* root type, since, as the example sentence in the previous paragraph shows, this LU tends to occur with sup-

port verbs and indicates an action taken by someone towards solving a problem. Next, we name the frame as `Brazilian_way` and connect it to the `Attempting_and_resolving_scenario` frame in FrameNet. We will then map FE in the latter to the ones we are in the process of creating. Hence, the AGENT FE is mapped to the INTERESTED_PARTY and the GOAL, MANNER and other non-core FEs in the mother frame are repeated in the newly created frame. Because *jeitinho.n* requires the conceptualization of an AUTHORITY being convinced – or corrupted – and of some NORM being violated, we add those two core FEs by clicking the *Create New FE* button. Next, additional non-core FEs typically occurring in eventive frame may be added. Finally, we edit the frame definition, and the description and coreness status of some FEs. Appendix A shows how this process is performed in Lutma step by step.

5. Limitations and Outlook

Despite already being deployed, Lutma is not yet a finished project. As with most software, there is room for improvement in regards to user experience and interface. Those issues will be dealt with after we receive more user feedback. There are also important functionalities that still need to be designed and developed, and because of that, in its current state, the system has some limitations.

One of them has to do with the fact that there are no tools that could aid users in assessing the quality of newly created frames. This process is not objective either, since previous research has shown how frame annotation can be ambiguous (Burchardt et al., 2006). Interestingly, this also means that the subsystems implemented to reduce redundancy during the frame creation process could benefit from multiple perspectives. When those subsystems fail and thus, a user creates a redundant frame, or one with overall less quality, a reasonable solution is the adoption of a wiki-like approach, where users can see edits and open discussions to determine the best course of action.

One final point worth considering is that Lutma’s metalanguage is English, meaning that only users proficient in this language will be able to contribute. Naturally, this can be circumvented by allowing users to translate attributes of frames, LUs, FEs and any other entities. However, this would also mean that users would spend less time actually creating those entities. For now, we have decided to leave English as the metalanguage, taking into consideration that other factors can also restrict potential users, even though we have no control over them (e.g. Lutma’s contributors are most likely people interested in fields such as Frame Semantics or NLP). Even given its limitations, Lutma is a step in an effort of scaling up FrameNet, without sacrificing the model’s advantages, especially in regards to perspective in NLP, to the extent that it facilitates contributions from any language and from non-specialist users.

6. Acknowledgements

The development of Lutma was funded by an Anneliese Maier Research Award provided to Mark Turner by the Alexander von Humboldt Foundation. Torrent’s research is funded by CNPq grant 315749/2021-0.

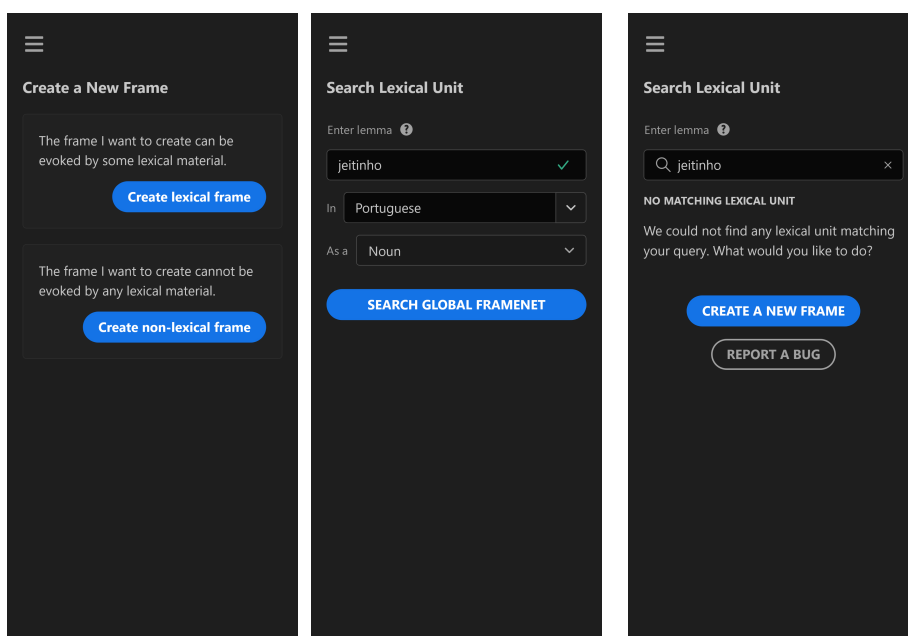
7. Bibliographical References

- Aroyo, L. and Welty, C. (2015). Truth is a lie: Crowd truth and the seven myths of human annotation. *AI Magazine*, 36(1):15–24.
- Artstein, R. and Poesio, M. (2008). Inter-coder agreement for computational linguistics. *Computational linguistics*, 34(4):555–596.
- Baker, C. F. and Lorenzi, A. (2020). Exploring Crosslinguistic Frame Alignment. In *Proceedings of the International FrameNet Workshop 2020: Towards a Global, Multilingual FrameNet*, pages 77–84, Marseille, France, May. European Language Resources Association.
- Basile, V., Cabitza, F., Campagner, A., and Fell, M. (2021). Toward a Perspectivist Turn in Ground Truthing for Predictive Computing. *arXiv preprint arXiv:2109.04270*.
- Bender, E. M. and Koller, A. (2020). Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online, July. Association for Computational Linguistics.
- Bender, E. M., Gebru, T., McMillan-Major, A., and Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623.
- Bond, F. and Foster, R. (2013). Linking and Extending an Open Multilingual Wordnet. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1352–1362.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Burchardt, A., Erk, K., Frank, A., Kowalski, A., Padó, S., and Pinkal, M. (2006). The SALSA Corpus: a German Corpus Resource for Lexical Semantics. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC’06)*, Genoa, Italy, May. European Language Resources Association (ELRA).
- Criado, N. and Such, J. M. (2019). Digital discrimination. *Algorithmic Regulation*, pages 82–97.
- Cruz, J. C. B. and Cheng, C. (2020). Establishing Baselines for Text Classification in Low-resource Languages. *arXiv preprint arXiv:2005.02068*.
- DaMatta, R. (1986). *O que faz o brasil, Brasil? [What makes brazil, Brazil?]*. Rocco, Rio de Janeiro, Brazil.
- Dannélls, D., Borin, L., and Heppin, K. F. (2021). *The Swedish FrameNet++: Harmonization, integration, method development and practical language technology applications*, volume 14. John Benjamins Publishing Company, Amsterdam.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (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, June. Association for Computational Linguistics.
- Fillmore, C. J. and Baker, C. (2010). A frames approach to semantic analysis. In *The Oxford handbook of linguistic analysis*.
- Fillmore, C. J., Petruck, M. R., Ruppenhofer, J., and Wright, A. (2003). Framenet in action: The case of attaching. *International journal of lexicography*, 16(3):297–332.
- 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.
- Gilardi, L. and Baker, C. (2018). Learning to Align across Languages: Toward Multilingual FrameNet. In Tiago Timponi Torrent, et al., editors, *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Paris, France, may. European Language Resources Association (ELRA).
- Lakew, S. M., Negri, M., and Turchi, M. (2020). Low resource neural machine translation: A benchmark for five african languages. *arXiv preprint arXiv:2003.14402*.
- Magueresse, A., Carles, V., and Heetderks, E. (2020). Low-resource languages: A review of past work and future challenges. *arXiv preprint arXiv:2006.07264*.
- Matsumoto, Y. (2010). Interactional frames and grammatical descriptions: The case of japanese noun-modifying constructions. *Constructions and Frames*, 2(2):135–157.
- Ntoutsis, E., Fafalios, P., Gadiraju, U., Iosifidis, V., Nejdil, W., Vidal, M.-E., Ruggieri, S., Turini, F., Papadopoulos, S., Krasanakis, E., et al. (2020). Bias in data-driven artificial intelligence systems—An introductory survey. *Wiley Interdisciplinary Re-*

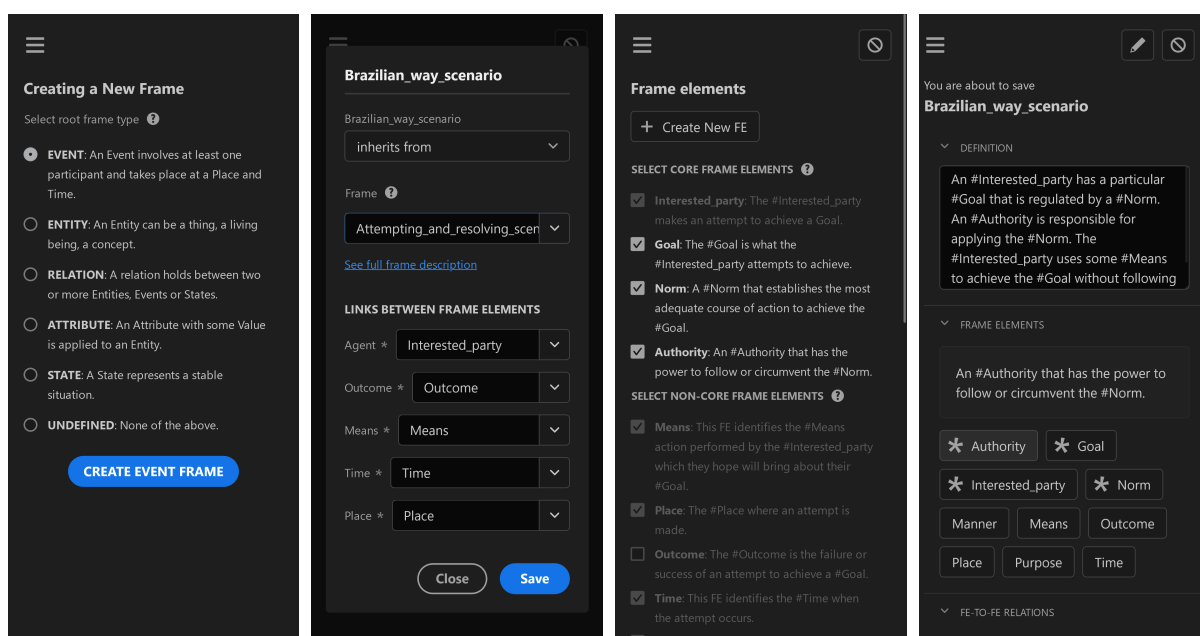
- views: *Data Mining and Knowledge Discovery*, 10(3):e1356.
- Ohara, K. H., Fujii, S., Ohori, T., Suzuki, R., Saito, H., and Ishizaki, S. (2004). The Japanese FrameNet project: An introduction. In *Proceedings of LREC-04 Satellite Workshop "Building Lexical Resources from Semantically Annotated Corpora" (LREC 2004)*, pages 9–11.
- Petruck, M. R. L. (1996). Frame Semantics. In Jef Verschueren, et al., editors, *Handbook of pragmatics*, pages 1–8. John Benjamins, Philadelphia.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. (2020). Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research*, 21(140):1–67.
- Rogers, A. (2021). Changing the World by Changing the Data. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2182–2194, Online, August. Association for Computational Linguistics.
- Ruppenhofer, J., Ellsworth, M., Schwarzer-Petruck, M., Johnson, C. R., and Scheffczyk, J. (2016). Framenet ii: Extended theory and practice. Technical report, International Computer Science Institute.
- Schroder, U. A. and Silva, R. d. C. (2020). O 'jeitinho brasileiro' a partir de uma perspectiva cognitivo-interacional [The 'Brazilian jeitinho' from a cognitive-interactional perspective]. *Estudos da Língua(gem)*, 18(2):117–134.
- Subirats-Rüggeberg, C. and Petruck, M. R. (2003). Surprise: Spanish FrameNet! In Eva Hajičová and Anna Kotěšovcová and Jiří Mirovský, editor, *Proceedings of the Workshop on Frame Semantics, XVII International Congress of Linguists (CIL)*, Prague. Matfyzpress, Matfyzpress.
- Torrent, T. T. and Ellsworth, M. (2013). Behind the labels: Criteria for defining analytical categories in Framenet Brasil. *Veredas-Revista de Estudos Linguísticos*, 17(1):44–66.
- Torrent, T. T., Ellsworth, M., Baker, C., and Matos, E. E. d. S. (2018a). 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., Lage, L., Laviola, A., Tavares, T., Almeida, V., and Sigiliano, N. (2018b). Towards continuity between the lexicon and the construction in Framenet Brasil. *Constructicography: Construction development across languages*, 22:107.
- 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.
- Torrent, T. T. (2021). What a Cognitive Linguist Means by Meaning and Why It Could Impact Research in Natural Language Processing. CogSci Blog. Cognitive Science Society.

Appendix A - User Interface Screenshots

This section presents screenshots of Lutma’s UI as presented to the user during the creation of the *Brazilian_way* frame as described in section 4. Since it is a mobile-first UI, we present how it is rendered in smartphones. The focus on mobile experience was a decision made by the team to make sure that users that do not have access to a desktop computer could also contribute. This decision also influenced the design of most screens. Some processes were split into multiple screens, which in the end, also makes it easier for unfamiliar users to work with some complex FrameNet concepts.



(a) Frame creation screen. (b) Lexical unit search screen. (c) Search result screen after no existing frame is found.



(a) Frame type selection screen. (b) Frame relation creation dialog. (c) Frame element list screen, including non-core FEs suggested by Lutma. (d) Summary screen displayed at the end of the execution flow.