

SOCIOFILLMORE: A Tool for Discovering Perspectives

Gosse Minnema¹, Sara Gemelli², Chiara Zanchi², Tommaso Caselli¹, Malvina Nissim¹

1. University of Groningen, The Netherlands

2. University of Pavia, Italy

g.f.minnema@rug.nl

Abstract

SOCIOFILLMORE is a multilingual tool which helps to bring to the fore the *focus* or the *perspective* that a text expresses in depicting an event. Our tool, whose rationale we also support through a large collection of human judgments, is theoretically grounded on frame semantics and cognitive linguistics, and implemented using the LOME frame semantic parser. We describe SOCIOFILLMORE’s development and functionalities, show how non-NLP researchers can easily interact with the tool, and present some example case studies which are already incorporated in the system, together with the kind of analysis that can be visualised.

1 Introduction

Descriptions of the very same event can vary widely. Sometimes completely different versions of a situation can be reported, while other times more subtle differences emerge by the way in which such situations or episodes are depicted by choosing specific natural language expressions. Figure 1 illustrates this: in the first sentence the car crash is lexicalized by the noun “collision” leaving the entire dynamic unknown and suggesting that the cyclist may have some responsibility. On the contrary, the second sentence uses the verb “hit” with a subject (the agent) and an object (the patient) making more transparent how the event happened and who is responsible for it.

This phenomenon is known as *framing* or *perspectivization*, and can happen in any discourse either in full awareness or, more often, unconsciously (Horst, 2020). Politics, for instance, is the prime arena where intentional and biased framing takes place (Iyengar, 1994; Semetko and Valkenburg, 2000; Entman, 1993; Matthes, 2012), but this occurs also in other domains, such as sports. Indeed, representation alternatives, namely the different choices that language allows to describe the same event (not only at the lexical but also at the

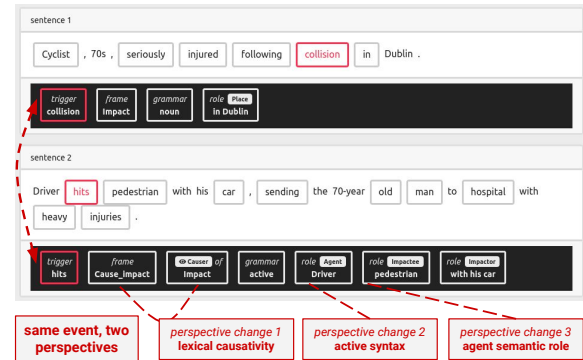


Figure 1: Analysis from SOCIOFILLMORE showing linguistic markers indicating the perspective changes in two descriptions of the same event. Words in boxes indicate triggers of semantic frames in the sentence.

syntactic and pragmatic level), are key to expressing and understanding the ideological power of discourse (Haynes, 1989).

Different theoretical frameworks can guide the study of framing and perspectivization in discourse, such as Critical Discourse Analysis (CDA) (Fairclough, 2010; Van Dijk, 1995) and Frame Semantics (Fillmore, 1985; Baker et al., 2003). While NLP tools based on some of such frameworks do exist to potentially support large-scale text analysis of perspectivization, and more specifically of Fillmore’s frame semantics (Xia et al., 2021), they are (i) recent, thus not yet established as analysis tools for specific perspectivization problems outside of the NLP community; (ii) technical, so that their adoption is basically impossible for the non-experts, who would though benefit from them. On the other hand, cognitive linguists have carried out cognitive linguistic analyses on the discourse regarding social issues and events (e.g. Pinelli and Zanchi (2021)), but these analyses usually imply manual scrutiny of data and, accordingly, deal with small datasets.

We fill this gap with the development of SOCIOFILLMORE, a user-friendly multilingual

tool based on frame semantics which allows users to conduct large-scale analyses of text by highlighting the perspectivization strategies they adopt.¹

2 Context & Evidence

In this section, we highlight the theoretical frameworks that have guided the design of SOCIOFILLMORE and the empirical evidence we have collected to support its application to written corpora.

2.1 Frames, Constructions, & Perspectives

Being able to depict or report about events is part of the broader human ability of storytelling (Boyd, 2009; Gottschall, 2012). One of the key properties of telling a story is the presence of a *focus*, or a perspective (Bal, 1997), which is embedded and intrinsic into every communication act, as shown in Figure 1. Furthermore, the lexical units in a discourse are powerful access points to complex conceptual structures of encyclopedic knowledge and perspectives. This vision is at the core of Fillmore’s frame semantics (Fillmore, 1971, 1985, 2006), and encoded into the FrameNet project (Baker et al., 2003). The availability of a computational resource and the strong linguistic orientation of Frame Semantics are the major driving reasons for preferring this framework over other existing ones.

Frames are powerful devices that express perspectives and they strongly interact with linguistic constructions (Langacker, 2006). In some cases, the construction used to present the events can even lead the evocation of different frames. This is clearly illustrated in Figure 1: the noun `collision` evokes the frame IMPACT, while the verb `hit` in an active voice evokes the frame CAUSE_IMPACT. The change of frames associated with different constructions triggers different perspectives and, in this case, also different responsibilities of the participants of the event.

In cognitive linguistic terms, these different constructional options are named *construals* (Langacker, 1991). Frames and construals are socially, culturally, and discursively constructed and constitutive (Dirven et al., 2007): they mirror our ideology, beliefs, and stereotypes and in turn contribute to building and enhancing them. Magnifying the connections between these two elements is a way to support users (e.g., social scientists, journalists,

¹An online demo and docker images of the app are available at <https://osf.io/8kh3d/>.

linguists, media studies scholars, among others) to identify and study the perpetration of biases and power structures in discourse.

2.2 Empirical Evidence

As a rationale for the validity of our tool two pieces of evidence are needed. The first concerns the feasibility and accuracy of (multilingual) frame semantic parsing, since any reasoning over the significance of finding one frame activated rather than an alternative one is meaningful only if frames can be accurately detected. For this, we ran an evaluation of LOME (Xia et al., 2021), a multilingual end-to-end frame parsing system which is the backbone of our tool, and found that it indeed produces reliable analyses. Details of the fine-tuning procedure and the evaluation are in Section 3.1.

The second piece of evidence concerns the relationship between frames and perspectives from a cognitive viewpoint, and more specifically on the human perception of semantic frames and/or of the interaction between syntactic and lexical semantic expressions of agentivity. Is it true that certain frames and/or constructions are associated with agentivity more than others? To test this, we ran a questionnaire where participants had to express judgements about their perception of the *focus* (i.e., *on which participant or entity is the main focus on the sentence?*) on a set of 400 sentences extracted from a large corpus on femicides in Italian (details in §4.1 and Appendix A). Judgements are expressed on a 5 point Likert-scale for four dimensions, namely focus on: ‘the murderer’, ‘the victim’, ‘an object’ (e.g., a weapon), or ‘an abstract concept or emotion’ (e.g., jealousy). The selected frames are reported in Table 1, together with the results, i.e., perception scores for each frame-construction pair (averaged over participants and sentences; every pair had approximately the same number of ratings). The highest level of focus on the murderer is found with the KILLING frame evoked by an active transitive construction; this makes sense, since this is the only situation in which the presence of a Killer role is required both syntactically and semantically.² On the other hand, constructions perceived as placing a high focus on the victim are found across all frames ex-

²Note that its syntactic realization appears to have a large influence on the perceived focus placed on the agent: on average, passive constructions evoking KILLING have a ‘murderer’ score of almost two points lower than their active counterparts.

frame/construction	murderer**	victim**	object	concept / emotion*
CATASTROPHE				
nonverbal	1.319	2.713	0.760	2.190
DEAD_OR_ALIVE				
nonverbal	1.195	3.387	1.386	1.993
vrb:unaccusative	1.983	3.529	1.566	1.539
DEATH				
nonverbal	0.967	3.247	1.507	1.914
vrb:unaccusative	1.867	3.921	1.690	1.286
EVENT				
nonverbal	1.431	1.503	1.186	2.339
vrb:impersonal	1.169	2.201	1.309	1.949
KILLING				
nonverbal	2.007	2.387	1.032	1.673
other	2.410	2.345	1.198	1.663
vrb:active	3.897	2.659	1.570	1.651
vrb:passive	1.947	3.425	1.491	1.315

Table 1: Average scores for survey question “the main focus is on X”. Legend: ‘vrb’ = verbal construction; ‘*’ = differences between frame-construction pairs are significant at $\alpha = 0.05$, ‘**’ = significant at $\alpha = 0.001$ (Kruskal-Wallis non-parametric H-test). Cells with a value > 2.5 are highlighted in green.

cept EVENT.³ While the analysis presented here is limited to a specific domain, there is a clear pattern in the perception scores of frames and constructions, granting sufficient ground for treating their automatically detected presence as a good proxy for how sentences perspectivize events in terms of foregrounding and backgrounding participants.

3 SocioFillmore

SOCIOFILLMORE consists of two parts: on the back-end side, there is a series of linguistic analysis components, and on the front-end side, there is a number of components for interacting with the user.

3.1 Linguistic Analysis Components

The linguistic analysis components are a combination of existing models and resources, linked by a set of rule-based bridging components.

LOME At the core of SOCIOFILLMORE is a frame semantic parser for annotating texts with FrameNet-based semantic frames and roles. While the rest of our architecture is agnostic as to what specific model is used, we decided to use LOME (Xia et al., 2021) as this is (i) one of very few available models capable of producing end-to-end frame analyses (i.e., taking raw text as input, without pre-specifying predicates to annotate), and (ii) based on

³This is consistent with the fact that, in FrameNet, the EVENT frame does not include any ‘core’ semantic roles apart from Place and Time, whereas all the other included frames include a Patient-like role that likely corresponds to the victim.

XLM-R, it is the only model we are aware of that supports zero-shot multilingual FrameNet analysis. Zero-shot multilingual predictions (given English-only annotated) data are very useful given the complexities of Multilingual FrameNet and the limited availability of training data in languages other than English. In an effort to evaluate and improve LOME’s multilingual capabilities, in Minnema et al. (2021), we tested LOME against an existing benchmark for Italian (the 2011 Frame Labeling over Italian Texts Task [FLAIT], Basili et al. 2013), and experimented with several methods for exploiting the limited available training data for Italian for improving on this. Interestingly, in a zero-shot setting, LOME underperformed versus the previous state-of-the-art SVM-based model, (Croce et al., 2013) by 24 percentage points (57% vs 81%) on the benchmark’s frame detection task (i.e., predict semantic frames given gold predicates), but outperformed it on the frame boundary and argument classification tasks (i.e. predicting role spans and labels given predicates and frames). In our cross-lingual training experiments, we achieved best results training LOME on the concatenation of the available corpora for English and Italian, with a frame detection score much closer to the previous state of the art (77%). However, in a small-scale manual annotation experiment on texts from the RAI femicides corpus (see §4.1) focusing on a limited set of frames, we found that the zero-shot LOME model performed substantially better than the cross-lingually trained version, which seems to be largely due to a drop in predicate detection performance (i.e., given raw text, find all predicates that evoke frames); this can be explained by the nature of the available Italian annotations, which cover only one frame per sentence, making it hard for the model to learn which lexical units evoke frames. Thus, while being far from perfect, zero-shot LOME seems to be the best currently available option in practice for automatic frame annotation.

Syntactic Analysis The second main step of the SOCIOFILLMORE analysis pipeline is syntactic analysis. For each frame structure (i.e., a semantic frame together with its semantic roles) identified by LOME, we extract three types of syntactic information through a combination of a UD parse obtained with spaCy (Honnibal et al., 2020), the FrameNet database, and a set of hand-written rules: (i) syntactic construction, (ii) role-dependency links, and (iii) predicate-root status. We distinguish between

several types of syntactic constructions; Table 2 lists these constructions in increasing order of participant foregrounding (nonverbal and impersonal constructions do not require any event participant to be syntactically expressed, unaccusatives and passives require only a Patient-like argument, while actives also require an Agent-like role).

Constructions are classified based on three criteria. First, we look at the part-of-speech tags produced by the UD parser, in order to separate non-verbal from verbal constructions. Second, if the construction is verbal, we use FrameNet for determining which core semantic roles are required for the construction: if the triggered frame is EVENT, the construction is classified as impersonal; all other frames in FrameNet have been manually annotated as being either ‘active’ or ‘non-active’ based on the presence or absence of an Agent-like participant in the definition of the frame. This information is used to classify verbal constructions associated with non-active frames as unaccusative. Finally, semantically active frames are classified as instantiating either passive or active constructions based on the syntactic features taken from the dependency parse (e.g., finite or infinitive verb form, presence of passive auxiliary, etc.).

Two additional types of extracted syntactic information are role-dependency links and predicate-root status. The former are labels that indicate how a semantic role label is expressed syntactically relative to the frame trigger, and are extracted from the dependency tree by a rule-based algorithm that traverses the dependency tree, starting from the frame trigger, until it encounters a token that is included in the role argument span, or until a pre-set number of maximum traversal steps has been reached. Some possible role-dependency links are “Event:nsubj↓” (nominal subject, e.g. in *the [event] happened*), “Killer:*” (self-referring link, e.g. in *the [assassin] of JFK*), and “Suspect:↑-nsubj↓” (subject of an intermediate node, e.g. in *the [prisoner] remains in detention*). For the purposes of perspective analysis, role-dependency links can provide a useful extra layer on top of construction information: for example, it can help us distinguish between agent-centered and event-centered nonverbal constructions (e.g. *assassin*, with a self-referential Killer role, vs. *homicide*, without a Killer role) or detect active constructions in which the main focus is on an inanimate cause rather than on an animate agent (*the [accident]*

Construction	Semantic Roles	Examples
non-verbal	none	<i>The *murder* of MLK</i> <i>A *deadly* accident</i>
vrb:impersonal	none [Event-like]	<i>It *rained*</i> <i>The event *occurred*</i>
vrb:unaccusative	Patient-like	<i>The victim *died*</i> <i>He *fell* off the stairs</i>
vrb:passive	Patient-like	<i>She *was found* in her house</i> <i>The cyclist *was hit* by a car</i>
vrb:active	Agent-like [Patient-like]	<i>The girl *walked* to school</i> <i>The police *arrested* the man</i>

Table 2: Syntactic construction types used by SOCIOFILLMORE. Semantic roles between square braces are mandatory in a subset of constructions within the type. ‘vrb’=verb-based constructions

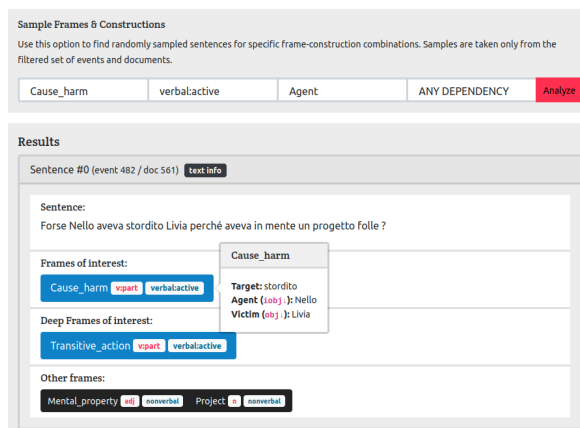


Figure 2: Explorer mode: visualize sentences matching specific linguistic features

killed him vs. *The [murderer] killed him*). On the other hand, predicate-root status refers to the position of the frame trigger in the syntactic tree, and serves as a proxy for how ‘central’ the construction is within the sentence. Verbal constructions are classified as ‘roots’ only if they are the root nodes of the dependency tree (i.e. are the main verb in the sentence), and nonverbal constructions are classified as ‘roots’ only if they are the subject of the root node. The intuition behind this is that the main verbal construction in a sentence, or nonverbal constructions closely related to it, are more under focus than other constructions. For example, in *the homicide happened* (root) versus *he was arrested for homicide ten years later* (non-root), *homicide* is foregrounded in the former but backgrounded in the latter.

3.2 User Interaction

SOCIOFILLMORE has two usage modes: one for exploring and analyzing existing, pre-processed corpora, and one for interactively exploring frame-based perspective analysis. The former mode is targeted towards domain experts interested in in-depth analysis of specific phenomena, while the latter is targeted to a broader community of (social) scientists who are unfamiliar with frame semantics but would like to learn about using linguistic frames and constructions for analyzing how events can be framed through language.

Corpus Explorer The corpus explorer is illustrated in Figure 2. It consists of three parts: first, the user can select relevant semantic frames from a pre-defined set (or add additional frames) that correspond to the perspective-taking phenomena that they want to investigate. Second, the user can specify which subset of the corpus they would like to analyze by adding document-based and event-based filters. Document-based filters (e.g. publication date, news outlet) can be added for any kind of corpus, while event-based filters (e.g. location of event, participants involved) can only be used for corpora that provide event-level metadata (i.e., each document is linked to some structured event representation). This second type of corpus is especially attractive for perspective analysis because it allows for investigating how similar real-world events are conceptualized in different documents that reference these events. Finally, having selected frames and filtered the corpus, the user can analyze the corpus in two ways: one can get global descriptive statistics over the selected part of the corpus, or visualize annotated sentences from the corpus. A wide range of descriptive statistics are available, for example, simple frequencies of semantic frames and constructions, frequencies of role dependencies per frame, and frame frequencies plotted as a function of time elapsed between the event occurrence and publication of the documents referencing it. On the other hand, when visualizing annotations from the corpus, there are two options: the user can either select specific documents from the corpus and analyze them sentence-by-sentence, or make a selection of linguistic features of interest (e.g. a combination of frames, constructions, and role dependencies) and request to see randomly sampled sentences matching these features.

As of now, we have implemented the corpus explorer for four different corpora on three domains

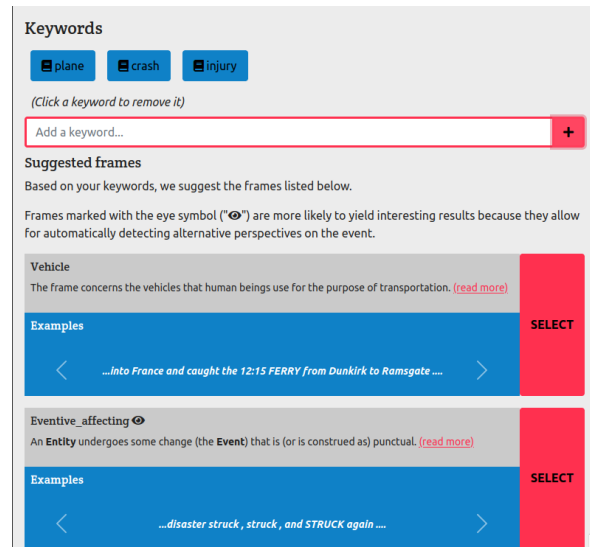


Figure 3: Interactive mode: keyword search interface

in two languages (femicides and migration in Italian, and traffic crashes in Dutch). We are planning to add additional corpora in the future, and also welcome contributed corpora from others. Adding an additional corpus requires some amount of expertise in NLP and FrameNet as well as in the domain of interest, and involves pre-processing the corpus and its metadata, running LOME and the SOCIOFILLMORE linguistic pipeline over the corpus, and adding corpus-specific logic (e.g. document/event filters) to the explorer interface. For future work, we are planning to develop a graphical UI to streamline this process and make it more accessible for users without a technical background.

Interactive Mode The interactive mode of SOCIOFILLMORE aims to make frame-based perspective analysis more accessible to people without a specific background in frame semantics. To this end, we provide three main features, as shown in Figure 3: a step-by-step interface with examples guiding the user through the stages of perspective analysis (event definition, frame selection, and document visualization), an interactive frame selection tool, and an on-demand version of LOME and the SOCIOFILLMORE linguistic pipeline with simplified annotation visualization. The most novel of these features is the interactive frame selection tool, which is meant to help users who are not familiar with the FrameNet database to find frames that are relevant for the event type that they would like to analyze. The frame selection tool consists of two components: an embedding-based keyword search function, and a rule-based algorithm for

automatically finding frames that provide alternative perspectives on selected frames. The former of these makes use of ‘bag-of-LU’ frame embeddings that are computed using a similar method to that proposed in [Alhoshan et al. \(2019\)](#). Keyword searches are performed by retrieving GloVe vectors for the specified keywords, and then finding the frame embeddings with the top-N closest cosine distance to the centroid of the set of keyword embeddings. The suggested frames are then displayed to the user along with their definition and example sentences retrieved using the NLTK FrameNet API ([Bird et al., 2009](#)). Alternative representations for the creation of the frame and keyword embeddings (e.g., BERT embeddings) could be easily integrated in the tool. Complementary to the keyword search system, we use frame-to-frame relations to automatically add additional frames that provide alternative perspectives on the frames that the user specified.

4 Case studies

SOCIOFILLMORE has been applied (so far) to three case studies: Italian news reporting on femicides, Italian news reporting on migration, and Dutch news reporting on traffic crashes. In each of these cases, we target events where there is a potential imbalance of power between the actors involved and the attribution of responsibility for the happening to (at least) one of the participants of the event. This makes the study of the perspectives associated with the reporting of these events very suitable to investigate how responsibility is framed in news reports of such events, and where potential representation biases may emerge.

4.1 Femicides

For femicides, the domain of SOCIOFILLMORE developed most extensively to date, two corpora are currently available in the exploration tool. The first has been compiled by the CRITS research team at RAI (Radiotelevisione Italiana), composed by 2,734 news articles from 31 different Italian news sources, reporting on 182 femicides perpetrated between 2015 and 2017 ([Belluati, 2021](#)). The corpus is enriched with metadata (time, news source, etc.) and for each femicide event multiple news articles are available. For our analysis we selected 15 frames based on the examples in [Pinelli and Zanchi \(2021\)](#) - see Appendix B.

In [Minnema et al. \(2021\)](#), we applied

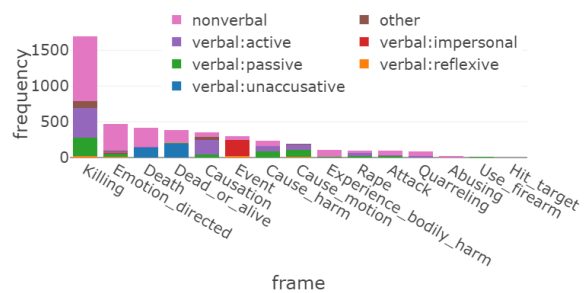


Figure 4: Femicides analysis: frequency of frames split by syntactic construction. Figure exported from SOCIOFILLMORE.

SOCIOFILLMORE on a randomly chosen 200K word subcorpus (10% of all events) of the RAI corpus. The main findings are shown in Figure 4. As expected, KILLING is by far the most frequent typical frame, followed by EMOTION_DIRECTED and DEATH. Concerning syntax, the *nonverbal* constructions (i.e., the predicate is either a noun or an adjective) are dominant across many frames, while *verbal:active* constructions are much rarer, as well as *verbal:passive* and *verbal:unaccusative*.

The combination of syntactic constructions and semantic frames is the key to magnifying perspectives. In particular, we observe that 60% of the instances of the frame KILLING are associated with constructions that foreground the victims and background the perpetrators. This reaches 79% of cases when the frame used to present the event is DEATH. In terms of perspective analysis, this indicates a bias in framing femicides as events where a killing takes place but no one is actually responsible for it.

4.2 Other domains

SOCIOFILLMORE has been productively used as well for studying the framing of traffic crashes and migrations in the Dutch and Italian media, respectively. [Zanchi et al. \(2021\)](#) used the tool for automatically identifying frames contributing to either dehumanizing or humanizing migrants in newspaper headlines (e.g. reporting on ‘waves’ of migrants, and thus collectively conceptualizing them as a non-human mass entity, vs. reporting on migrants as single and intentional individuals), and also compared the change in these frames over time relative to statistics about newly arrived migrants in Italy. We are also involved in ongoing work, in collaboration with the author of the original paper, aiming at reproducing the findings on traffic

framing reported in Te Brömmelstroet (2020).

5 Conclusions

SOCIOFILLMORE is a multilingual tool that we have developed for studying perspectives in written text, grounded in Frame Semantics and Cognitive Linguistics. Through an interactive mode, the tool is easily accessible to non-experts, too. We support the rationale for the validity of our tool through a rigorous evaluation of the frame semantic parser at the core of our tool and a collection of human judgement on the connection between frames and perspectivization. The tool is available as a web interface (and as a docker release), it supports multiple languages, and already integrates a few large-scale case studies which can be browsed for research, but also for further understanding of the tool's functionalities.

Ethical Statement

One of the key properties of SOCIOFILLMORE is its being agnostic on whether a perspective should be considered “good” or “bad”. In this respect SOCIOFILLMORE is **not a prescriptive tool** on how news should be reported but rather a support tool that helps to magnify misuse of frames and biases that may mirror and strengthen asymmetric power dynamics existing in our societies.

SOCIOFILLMORE is based on state-of-the-art NLP technologies. While these tools achieve very good performances (also in zero-shot multilingual settings), none of them can be considered to reach nor mimic humans. The tools are based on powerful machine learning algorithms but they are far away from being artificial intelligent agents. We recommend **caution** when using SocioFillmore since margins of errors are present. At the same time, additional tests are needed before deploying SOCIOFILLMORE as an integrated tool or service that citizens or professionals may use for purposes other than research.

One of the services of SOCIOFILLMORE is corpus-assisted language analysis. The outcome of this service is highly sensitive to the data that are feeded to the tool. This requires users to pay particular attention to the curation of the data that will compose their corpus. Results on the presence of frames and bias are a direct consequence of what is input to the tool. The case studies we have illustrated are based on carefully curated corpus collections conducted by experts. We thus recommend

that users of SOCIOFILLMORE should accompany the presentation of their results with a **documentation of their data collections** using tools such as data statements (Bender and Friedman, 2018) or data sheets (Gebru et al., 2021).

References

- Waad Alhoshan, Riza Batista-Navarro, and Liping Zhao. 2019. [Semantic frame embeddings for detecting relations between software requirements](#). In *Proceedings of the 13th International Conference on Computational Semantics - Student Papers*, pages 44–51, Gothenburg, Sweden. Association for Computational Linguistics.
- Collin F. Baker, Charles J. Fillmore, and Beau Cronin. 2003. The structure of the FrameNet database. *International Journal of Lexicography*, 16(3):281–296.
- Mieke Bal. 1997. *Narratology: Introduction to the theory of narrative*. University of Toronto Press.
- Roberto Basili, Diego De Cao, Alessandro Lenci, Alessandro Moschitti, and Giulia Venturi. 2013. Evalita 2011: The frame labeling over italian texts task. In *Evaluation of Natural Language and Speech Tools for Italian*, pages 195–204, Berlin, Heidelberg. Springer Berlin Heidelberg.
- M. Belluati. 2021. *Femminicidio. Una lettura tra realtà e interpretazione*. Biblioteca di testi e studi. Carocci.
- Emily M Bender and Batya Friedman. 2018. Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc."
- Brian Boyd. 2009. *On the origin of stories*. Harvard University Press.
- Danilo Croce, Emanuele Bastianelli, and Giuseppe Castellucci. 2013. Structured kernel-based learning for the frame labeling over Italian texts. In *Evaluation of Natural Language and Speech Tools for Italian*, pages 195–204, Berlin, Heidelberg. Springer Berlin Heidelberg.
- René Dirven, Frank Polzenhagen, and Hans-Georg Wolf. 2007. Cognitive linguistics, ideology, and critical discourse analysis. In *The Oxford handbook of cognitive linguistics*.
- Robert M Entman. 1993. Framing: Towards clarification of a fractured paradigm. *Journal of Communication*, 43.

- Norman Fairclough. 2010. *Critical Discourse Analysis: The Critical Study of Language (Second Edition)*. Routledge.
- Charles J. Fillmore. 1971. Subjects, speakers, and roles. *Synthese*, 21(3-4).
- Charles J. Fillmore. 1985. Frames and the semantics of understanding. *Quaderni di semantica*, 6(2):222–254.
- Charles J. Fillmore. 2006. Frame semantics. In D. Geeraerts, editor, *Cognitive Linguistics: Basic Readings*, pages 373–400. De Gruyter Mouton, Berlin, Boston. Originally published in 1982.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92.
- Jonathan Gottschall. 2012. *The storytelling animal: How stories make us human*. Houghton Mifflin Harcourt.
- John Haynes. 1989. *Introducing stylistics*. Allen & Unwin Australia.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. *spaCy: Industrial-strength Natural Language Processing in Python*.
- Dorothea Horst. 2020. Patterns ‘we’ think by? critical cognitive linguistics between language system and language use. *Yearbook of the German Cognitive Linguistics Association*, 8(1):67–82.
- Shanto Iyengar. 1994. *Is anyone responsible?: How television frames political issues*. University of Chicago Press.
- Robert W. Langacker. 1991. *Foundations of Cognitive Grammar, vol. II: Descriptive application*. Stanford University Press, Stanford.
- Ronald W. Langacker. 2006. Dimensions of defocusing. In Tasaku Tsunoda and Taro Kageyama, editors, *Voice and grammatical relations: In honor of Masayoshi Shibatani*, page 115–138. John Benjamins Publishing.
- Jörg Matthes. 2012. Framing politics: An integrative approach. *American behavioral scientist*, 56(3):247–259.
- Gosse Minnema, Sara Gemelli, Chiara Zanchi, Viviana Patti, Tommaso Caselli, and Malvina Nissim. 2021. Frame semantics for social NLP in Italian: Analyzing responsibility framing in femicide news reports. In *Proceedings of the Eighth Italian Conference on Computational Linguistics*. Milan, January 26–28, 2022 (postponed).
- Erica Pinelli and Chiara Zanchi. 2021. Gender-based violence in Italian local newspapers: How argument structure constructions can diminish a perpetrator’s responsibility. *Discourse Processes between Reason and Emotion: A Post-disciplinary Perspective*, page 117.
- Holli A Semetko and Patti M Valkenburg. 2000. Framing European politics: A content analysis of press and television news. *Journal of communication*, 50(2):93–109.
- Marco Te Brömmelstroet. 2020. *Framing systemic traffic violence: Media coverage of Dutch traffic crashes*. *Transportation research interdisciplinary perspectives*, 5.
- Teun A. Van Dijk. 1995. Discourse analysis as ideology analysis. In C. Schäffner and A.I. Wenden, editors, *Language and Peace*. Harwood Academic Publishers, Amsterdam.
- Patrick Xia, Guanghui Qin, Siddharth Vashishtha, Yunmo Chen, Tongfei Chen, Chandler May, Craig Harman, Kyle Rawlins, Aaron Steven White, and Benjamin Van Durme. 2021. *LOME: Large ontology multilingual extraction*. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 149–159, Online. Association for Computational Linguistics.
- Chiara Zanchi, Serena Coschignano, and Gosse Minnema. 2021. Numeri in arrivo anziché uomini e donne in partenza: Come i frame ci rendono inconsapevolmente ingiusti [‘Numbers in arrival instead of men and women in departure: How frames render us unconsciously unjust’]. In *Notizie ai margini: IX rapporto Carta di Roma 2021 [‘News at the margins: 9th report of the Rome Charter’]*. <https://www.cartadiroma.org/>.

A Questionnaire

The questionnaire has been conducted using the platform Qualtrics.

Participants have been recruited from several universities in Italy ($N = 239$; Male = 86; Female = 153). Each participant was presented with 50 sentences and was asked to express a judgement on Likert-type scale from 0 to 5 for the perceived “amount” of focus placed on four dimensions, namely: ‘the murderer’, ‘the victim’, ‘an object’ (e.g., a weapon), or ‘an abstract concept or emotion’ (e.g., jealousy).

The sentences were selected by first automatically annotating the corpus with semantic frames and construals, determining a set of semantic frames that correspond to different ways of conceptualizing the femicides, selecting the set of most frequent construals for each frame, and then, for

each frame-construal pair, randomly sampling sentences containing at least one instance of the pair from the corpus.

The frames that we selected were, in order of increasing level of detail and presence of event participants:

- EVENT, e.g., *the incident occurred*;
- CATASTROPHE, e.g., *the tragedy cost the life of ...*;
- DEAD_OR_ALIVE, e.g., *the victim was found dead*;
- DEATH, e.g., *the victim died at the hands of... and*
- KILLING, e.g., *the man murdered his wife*.

Each of these frames can occur in various syntactic configurations, but only KILLING can be evoked by a transitive verbal construction (actively or transitively used).

While it is possible that, in some cases, instances of these frames refer to another event referenced in the texts (e.g. KILLING could also refer to the perpetrator committing suicide, or to some other type of secondary murder; EVENT could also refer to other type of 'events' or 'incidents' mentioned in the text), but, from a manual inspection of the data, this seems to be fairly rare. We informed participants of the possibility of anomalous sentences occurring in the survey and instructed them to mark these as 'irrelevant' and to not assign any points to them.

The questionnaire has been approved by the Ethical Board of the University of Groningen. Participants were compensated with 5 euro. Each participant was asked to provide judgments on a set of 50 sentences. Participation is anonymized (there is no link between the questionnaire answers and the participants.) and we limit the collection of personal data to last name, initials, bank account and address only for payment purposes. After the payment, all personal data of the participant is deleted.

A screenshot of the original instructions provided to the participants is presented in Figure A.1. Translations in English of the instructions is given below:

INSTRUCTIONS:
Dear participant,
the following questionnaire
is part of a project related to

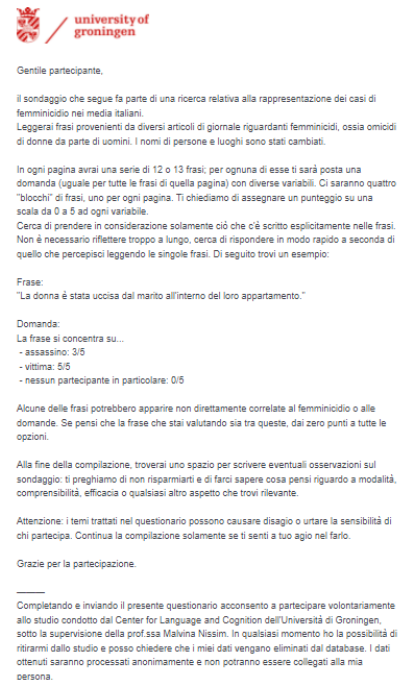


Figure A.1: Original instruction given to the participant in the questionnaire for validating SOCIOFILLMORE.

the representation of femicides in the Italian media. You will read sentences extracted from newspaper articles about femicides, i.e. murders in which a man kills a woman. The real names in the sentences have been changed.

In each page you will be presented with a series of 12 or 13 sentences; there are four different "stacks" of sentences, one for each page. For each sentence you will be asked a question (the same of every sentence in that page) followed by different variables. You will be asked to rate every variable on a scale from 0 to 5.

While doing the task, try to consider just what the sentence explicitly expresses. There is no need to think for too long, try to answer fast according to what you perceive reading the

sentences. Example: *The woman was killed by her husband at home.*
Question: The sentence focuses on...

- the murder: 3/5
- the victim: 5/5
- nobody in particular: 0/5

After the data analysis, the participants will receive an e-mail with a link to an online meeting in which we will briefly present the outcomes of the research.

Warning: the topics in the questionnaire could make you feel uncomfortable; please continue only if you feel at ease. Thank you for your time!

By filling in the following questionnaire I consent to participate voluntarily in the study conducted by the Center for Language and Cognition of the University of Groningen, supervised by prof. dr. Malvina Nissim. I can withdraw my participation at any time and have the data obtained through this study returned to me, removed from the database or deleted. The data obtained during this study will be processed anonymously and will therefore not be able to be traced back to me.

B Frames and semantic roles

In Tables [B.1](#), we provide the set of semantic frames and a mapping between their associated semantic roles and the main participants in the femicides.

frame	role:perpetrator_like	role:victim_like	role:cause_like
Abusing	Abuser	Victim	-
Attack	Assailant	Victim	-
Causation	Causer	Affected	Cause
Cause_harm	Agent	Victim	Cause
Cause_motion	-	-	-
Dead_or_alive	-	Protagonist	Explanation
Death	-	Protagonist	Cause
Emotion_directed	-	-	-
Event	-	-	-
Experience_bodily_harm	Experiencer Body_part	-	-
Hit_target	Agent	Target	-
Killing	Killer	Victim	Cause
Quarreling	-	-	-
Rape	Perpetrator	Victim	-
Use_firearm	Agent	Goal	-

Table B.1: Femicides: mapping frames, participants, and roles