

# Frame Semantic Patterns for Identifying Underreporting of Notifiable Events in Healthcare: The Case of Gender-Based Violence

Lívia Dutra<sup>1,2</sup>, Arthur Lorenzi<sup>1,3</sup>, Laís Berno<sup>1</sup>, Franciany Campos<sup>1</sup>,  
Karoline Biscardi<sup>1,4</sup>, Kenneth Brown<sup>1</sup>, Marcelo Viridiano<sup>1</sup>, Frederico Belcavello<sup>1</sup>,  
Ely Matos<sup>1</sup>, Olívia Guaranha<sup>3</sup>, Erik Santos<sup>3</sup>, Sofia Reinach<sup>3</sup>, Tiago Timponi Torrent<sup>1,5</sup>

<sup>1</sup> Federal University of Juiz de Fora | FrameNet Brasil, <sup>2</sup> University of Gothenburg,

<sup>3</sup> Vital Strategies Brasil, <sup>4</sup> Federal University of Minas Gerais,

<sup>5</sup> Brazilian National Council for Scientific and Technological Development – CNPq

livia.vicente.dutra@svenska.gu.se, {alorenzi, oguaranha, esantos, sreinach}@vitalstrategies.org,

{lais.berno, franciany.campos, kenneth.cyrill, marcelo.viridiano}@estudante.ufjf.br,

karolbiscardi@let.grad.ufmg.br, {fred.belcavello, ely.matos, tiago.torrent}@ufjf.br

## Abstract

We introduce a methodology for the identification of notifiable events in the domain of healthcare. The methodology harnesses semantic frames to define fine-grained patterns and search them in unstructured data, namely, open-text fields in e-medical records. We apply the methodology to the problem of underreporting of gender-based violence (GBV) in e-medical records produced during patients' visits to primary care units. A total of eight patterns are defined and searched on a corpus of 21 million sentences in Brazilian Portuguese extracted from e-SUS APS. The results are manually evaluated by linguists and the precision of each pattern measured. Our findings reveal that the methodology effectively identifies reports of violence with a precision of 0.726, confirming its robustness. Designed as a transparent, efficient, low-carbon, and language-agnostic pipeline, the approach can be easily adapted to other health surveillance contexts, contributing to the broader, ethical, and explainable use of NLP in public health systems.

**Keywords:** FrameNet, Digital Healthcare Surveillance, Gender-Based Violence

## 1. Introduction

Healthcare surveillance is a major challenge for health professionals, even in the context of the growing digitalization of medical records and the widespread use of advanced tools for data analysis. In a review paper on the topic, [Aliabadi et al. \(2020\)](#) point out, among the major challenges in implementing adequate surveillance, the difficulty in analyzing unstructured data, the low quality of information in e-medical records, as well as privacy and security issues. [Nazi and Peng \(2024\)](#), in turn, list nine challenges for the use of LLMs in healthcare, which include: in the dataset gathering phase, (a) presence of Personally Identifiable Information (PII), (b) possible privacy/security breaches and (c) bias introduced by uneven data distributions; in the model training phase, (d) lack of access to the workings of blackbox models and (e) carbon footprint; and, finally, in the application phase, lack of (f) transparency and (g) explainability, (h) ethical risks and (i) hallucinations.

Within the context of surveillance, Gender-Based Violence (GBV) remains a persistent challenge worldwide, affecting millions of women and girls in different cultural and social contexts. Global surveys suggest that more than one in three women who have ever been married or partnered have experienced physical or sexual violence from an in-

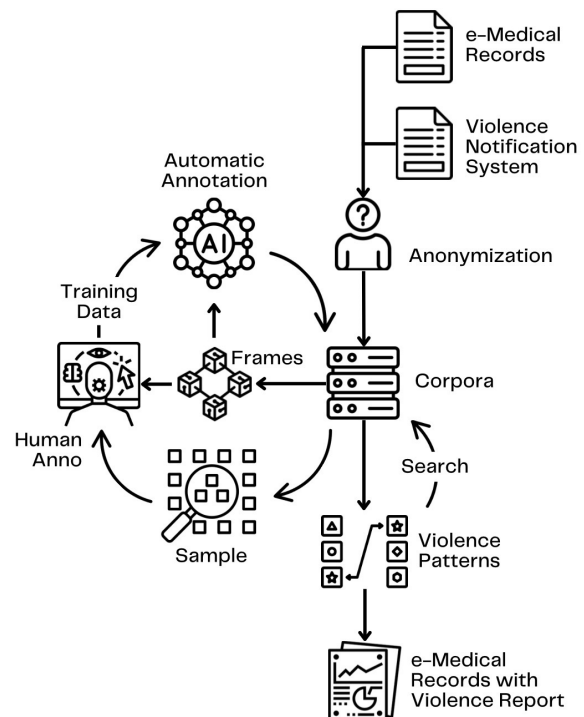


Figure 1: Schematization of the methodology

timate partner ([WHO, 2024](#)). Various studies treat Gender-Based Violence (GBV) as a public health is-

sue (Garcia-Moreno and Watts, 2011; Sweet, 2014; Öhman et al., 2020), highlighting the importance of the healthcare system in the fight against GBV. In Brazil, underreporting is still a major issue when it comes to GBV cases. According to Garbin et al. (2015) and Kind et al. (2013), underreporting is usually due to excessive workload of healthcare teams, lack of knowledge about the importance of notification, fear of retaliation, as well as difficulties in identifying which injuries and conditions are related to an episode of violence.

This paper presents an innovative methodology, outlined in Figure 1, to improve healthcare surveillance in general and to tackle GBV underreporting specifically, by leveraging the modeling of frames covering the domains of healthcare (Dutra et al., 2023) and violence (Larré and Torrent, 2024) in Brazilian Portuguese to identify GBV patterns in open text fields of e-medical records. We present the steps involved in the modeling and pattern definition process and evaluate the precision of the proposed methodology against a sample of 5,000 anonymized sentences extracted from e-medical records of primary care visits. We also conduct an error analysis and discuss its impact in feeding back the pattern-definition effort. Results show a precision of 0.726 of the proposed method in identifying GBV and 0.600 in the fine-grained classification of the type of violence.

Our contributions are threefold: (i) we introduce an innovative explainable and transparent pipeline for the identification of health events in e-medical records and report its energy consumption and carbon footprint; (ii) we evaluate the methodology for the challenging context of GBV in healthcare surveillance; and (iii) we design the methodology to be generalizable to healthcare surveillance needs beyond GBV.

In the remainder of this paper we briefly introduce FrameNet, the framework under which the domains of healthcare and violence were modeled, and describe the modeling process employed for the two domains in Section 2. Next, Section 3 shows how the manual and automatic annotation of the open-text fields in the e-medical records was carried out. Section 4 focuses on the construction of frame-semantic patterns and on the methods used for their initial validation. Section 5 presents the setup used to evaluate the methodology, while Section 6 presents and discusses the results.

## 2. Frame-Based Modeling of the Healthcare and Violence Domains

Developed by Charles Fillmore in the late 1970s, Frame Semantics is based on the idea that the meaning of a particular lexical item is understood against the evocation of an associated conceptual

scene. For its part, a frame is a structure of related concepts organized around a specific context in which certain words function as shortcuts to such a frame. That is, by using them, participants in a conversation not only instantly access the shared knowledge of that situation but also classify elements within it, assuming that the general context is already understood by all involved (Fillmore, 1982).

As an implementation of Frame Semantics, Berkeley FrameNet (BFN) is a resource developed in 1997 to analyze the English lexicon. Since the inception of this project, the FrameNet model has been adapted to create similar resources for multiple other languages, such as Brazilian Portuguese: the FrameNet Brasil (FN-Br) (Torrent et al., 2022). The core principle of any framenet analysis is that "meaning is relativized to scenes" (Fillmore, 1977). Simply put, the meaning of any word is determined by a background situation defined by the participants and props involved in it. Therefore, framennets, which are structured according to the principles of Frame Semantics, theoretically have the capacity to computationally represent diverse aspects of context (Torrent et al., 2022).

Every frame consists of a definition and a set of Frame Elements (FEs). The definition provides a general description of the scene the frame represents, typically referencing the FEs. These FEs, in turn, are classified as core FEs, mandatory to the frame's meaning, and non-core FEs, encoding circumstantial information. To illustrate the structure of a frame, consider the `Experience_bodily_harm` frame in (1), which models a scenario of physical injury.

- (1) **Experience\_bodily\_harm**  
**Definition:** An `Experiencer` is involved in a bodily injury to a `Body_part`. Often an `Injuring_entity` on which the `Experiencer` injures themselves is mentioned.  
**Core Frame Elements:**  
`Body_part`: The location on the body of the `Experiencer` where the bodily injury takes place.  
`Experiencer`: The being or entity that is injured.  
**Peripheral Frame Elements:**  
`Injuring_entity`: The `Experiencer` injures her/himself on an `Injuring_entity`.  
`Severity`: The extent to which the `Experiencer` is affected by the injury.  
**Frame-to-Frame Relations:**  
is inchoative of: `Being_hurt`  
inherits from: `Undergoing`  
is used by: `Cause_harm`

**Lexical Units:** *injury.n; break.v; bruise.v...*

Note that the definition mentions the main elements involved in the frame: the EXPERIENCER, the BODY\_PART and the INJURING\_ENTITY, the two first being core, and the last being peripheral, since this frame assumes the perspective of the person suffering the injury. Each frame element receives a definition that is specific to the frame. Moreover, because FrameNet is a net, not a list of frames, relations between the target frame and other frames in the network are also given for each frame. In the case of the `Experience_bodily_harm` frame, it inherits the `Undergoing` frame, meaning it is a more specific type of event in which an ENTITY is affected by some EVENT. It is also presupposed by the `Cause_harm` frame, meaning that, for someone to cause a harm on someone else, this last person has to experience a bodily harm. This is represented by the using relation indicated in (1). Finally, there is an alternation to our exemplar frame where a stative framing is adopted: the `Being_hurt` frame, which is in an inchoative\_of relation with `Experience_bodily_harm`.

FrameNet domain modeling aims to build a structured cognitive representation of the key concepts within a topic. Each domain comprises frames representing events, participants, and entities connected through frame-to-frame relations. Following Costa (2020), we adapted and expanded this methodology for our study. The next subsections describe the modeling process.

**Anonymization and Corpora Compilation :** The first step in modeling the domains was to compile and anonymize a representative corpus for each domain. The open text fields of both e-medical records from e-SUS APS and violence notifications from SINAN were made available by the city of Recife in Brazil.<sup>1</sup> All open text fields from 3,8 million records were split into 21 million sentences using Trankit, a transformer-based NLP toolkit, (Nguyen et al., 2021).

All sentences used in the research underwent a three-step anonymization process: automatic, semiautomatic, and manual validation. The automatic process used three Named Entity Recognition models (Souza et al., 2019; Guillou, 2021; Cunha and Ramalho, 2022), regular expressions, and a fuzzy search based on known venues to detect potential PII. A text span was anonymized only if at least two techniques agreed that it contained sensitive data. The semi-automatic step

<sup>1</sup>e-SUS APS is the system used by the Brazilian Universal Public Healthcare System in primary care units. For the e-medical records, the SOAP notes (Subjective, Objective, Assessment and Plan) were used. SINAN is the National System for Violence Notification.

used frequency lists of Brazilian Portuguese words and common names to identify potential personal names to accelerate manual validation. Lastly, manual validation was done only over specific samples of the corpus. The goal was to evaluate the process and to guarantee full anonymization of sentences that would be made available for the researchers doing frame-based annotation or evaluating semantic patterns.

The corpora compiled for domain modeling comprised 206,000 words distributed into 43,000 sentences for the e-medical records, and 172,000 words distributed into 16,000 sentences for the violence notifications.

**Domain analysis and Scenario Frame Creation:**

From a sample of the corpora, the study of the domains was conducted to identify the possible participants, actions, places, and objects that could and/or should be present in each of the domains. Based on this study, the first frame for each of the domains could be modeled: the scenario frame. This is a generic frame whose role is to organize the other frames in the domain. In the case of the Health and Violence domains, existing frames were adjusted to occupy this position. The previous frame `Medical_interaction_scenario` was restructured to be broader and became the `Healthcare_scenario` frame. Meanwhile, the `Violence` frame was adjusted to become the `Violence_scenario` frame.

**Candidate Term Extraction:** In the next step, the *Keywords — Terminology extraction* feature of Sketch Engine (Kilgarriff et al., 2014) was used to extract the terms—single or multiword expressions—most relevant to the domains. This feature calculates a score that highlights the terms that are especially prominent in a focus corpus; thus, the higher the score, the more characteristic or representative the term is of the specific domain. For example, terms such as *abuso.n* 'abuse.n' or *soco.n* 'punch.n' stood out as more characteristic of the Violence domain, compared to terms such as *árvore.n* 'tree.n' or *casa.n* 'house.n'. To accomplish that, the anonymized corpora were compared against a generic corpus, the Portuguese Web 2020. The extracted list of terms consisted of the 2,000 candidates most likely to belong to the domain.

**Lexicographic Behavior Analysis:** The lexicographic behavior of the candidate terms in the list was analyzed and potential lexical units (LUs) were grouped together with the aim of identifying those belonging to existing frames or suggesting new frames that needed to be modeled to better structure the domain. For example, a large number of candidate LUs denoted physical assaults such

as *machucar.v 'hurt.v'* that could be directly linked to the *Cause\_harm* frame in the violence domain. However, this frame covers the perspective of the person causing the harm, not that of the victim, which led to the modeling of a new *Being\_hurt* frame. A considerable number of candidate LUs could also be associated with frames that did not belong to the domains, but that could be of interest to the research objective, such as those that described relations between family members or acquaintances and were linked to the *Kinship* and *Personal\_relationship* frames.

**Frame Creation:** The next step was to model the newly suggested frames. In this process, two approaches were followed: the bottom-up approach, using the study of the lexical items as its starting point, and the top-down approach, where the researcher's intuition is the starting point, as described in [Ruppenhofer et al. \(2016\)](#). For the Healthcare domain, 17 new frames were modeled, six frames were adjusted, and 12 existing frames were linked to the domain, resulting in 35 frames. For the violence domain, five new frames were modeled, one was adjusted, and 42 existing frames were linked, totaling 48 frames.

**Frame Relation Definition:** Once the frames for each domain were defined, they were connected through frame-to-frame relations. The relations were defined between frames both in and outside the domains when needed. Following [Ruppenhofer et al.](#), there are eight types of relation between frames: inheritance, subframe, use, causative of, inchoative of, precedence, perspective of and see also. Figure 2 illustrates the process by showing the relations between some of the frames in the Violence domain.

**Lexical Unit Creation:** Once the network of frames is in place, the lemmas grouped as possible LUs were associated with the frames in the database. This step consists of first adding the lemmas and word forms—including not only inflectional, but also spelling variations, abbreviations and acronyms typically found in e-medical records—to the database management and annotation tool ([Torrent et al., 2024](#)), so that their instantiations in the corpora could be identified in the annotation phase. In total, the Violence domain comprises 1,764 LUs, while the Healthcare domain counts 2,776 LUs.

**Ternary Qualia Relation Definition:** Although frame-to-frame relations play an important role in the model, they are not able to specify which LUs in each frame are to be related to which ones in another frame and also what is the nature of such

a relation. To address this issue, Ternary Qualia Relations were defined. These relations are based on the Generative Lexicon Theory ([Pustejovsky, 1995](#)), which argues that the meaning of a word can be perspectivized from four types of qualia: agentive, formal, constitutive, and telic. [Torrent et al. \(2024\)](#) expanded the theory to have such relations mediated by frames, as illustrated in Figure 3 for the LU *facar.n 'knife.n'*. Over 4,500 relations have been created between the LUs associated with the frames in the Healthcare and Violence domains.

### 3. Data Annotation

Annotation was divided into two phases: human and automatic. Given the nature of the data, as an additional layer of protection for the privacy of patients, a version of the annotation platform with restricted access to the team members was created. Anonymized sentences of e-medical records were then uploaded to this instance of the platform to be annotated. The first goal of the manual annotation was to validate the modeling of the domains. The annotation allows for a general analysis that determines the adequacy of the model, especially with regard to the sufficiency and adequacy of frames in terms of coverage and structure. The second goal of this process was to create a domain-specific training dataset for an automatic semantic role labeler. Both processes are described in detail next.

#### 3.1. Human Annotation

The human annotation process was conducted by seven annotators trained in FrameNet methodology ([Ruppenhofer et al., 2016](#)). The annotation process focused on the semantic content—instead of abiding by the traditional three-layer semantic and syntactic annotation usually employed by FrameNet—which means that the annotation consisted of tagging the frame evoked by the target LU and its frame elements. Also, the expansion of the syntactic locality of the annotation was used in some cases, so as to keep track of core frame elements expressed at the beginning of very long sentences or fragments where automatic sentence segmentation fails, typical of e-medical records, where healthcare professionals may not follow punctuation conventions. The annotation was performed in 2,352 sentences, which yielded 6,313 annotation sets—each of which has one target LU—for the Healthcare domain and 8,309 for the Violence domain.

#### 3.2. Automatic Annotation

Automatic annotation used an instance of LOME ([Xia et al., 2021](#)). LOME is a system developed for multilingual information extraction that in-




Figure 2: Part of the frames in the Violence domain

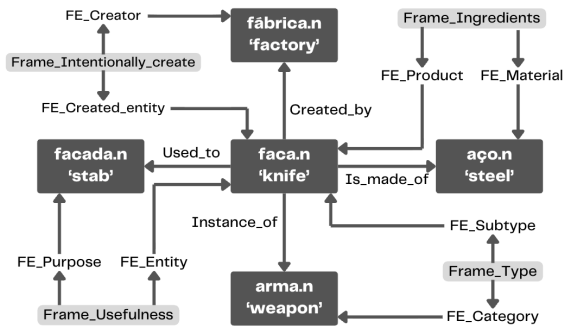


Figure 3: Example of Ternary Qualia Relations

cludes FrameNet semantic parser based on XLM-RoBERTa (Conneau et al., 2020). This parser was re-trained from scratch, having the human annotation described in Section 3.1 as part of its training dataset in addition to more than 130,000 other annotation sets from previous domain-general annotations available in the FN-Br database. The newly trained parser<sup>2</sup> has in its tag set over 200 new frames that expand on FrameNet 1.7 — the database version used by Xia et al. (2021).

The performance of the LOME instance trained for the project was evaluated and achieved a micro-F1 of 50.68 on the exact matching metric, which is slightly lower than the 56.34 of the original implementation. This reduction in micro-F1 for this metric is considered small given the more challenging nature of the training set: it contains domain specific, noisy clinical texts in a less resourced language—Brazilian Portuguese—and a larger set of labels of frames and frame elements.

Before feeding the whole corpus of 21 million sentences to LOME, Trankit was used for tokenization, lemmatization and part-of-speech tagging. The tokenized sentences were fed to LOME and the lemmas, in combination with part-of-speech and the predicted frames were used to infer lexical units, as the FrameNet parser does not predict those out of the box. The whole process resulted in 59 million automatic annotation sets, an average of 2,81 sets

<sup>2</sup><https://github.com/frameNetBrasil/span-finder>

per sentence. The computational cost of the full annotation pipeline – parser retraining (7h51m) and inference over 21 million sentences (62h3m) – on a single NVIDIA A30 GPU (165W TDP) amounted to approximately 23.06 kWh of energy and an estimated 1.98 kgCO<sub>2</sub>eq, assuming a PUE of 2.0 and a regional grid carbon intensity of 86 gCO<sub>2</sub>eq/kWh for the Centro-Sul region of Brazil (Electricity Maps, 2025), following the estimation methodology of Lacoste et al. (2019).

As an elicited example of the resulting automated semantic annotation process, consider the sentence in (2), which was annotated for the `People_by_health_condition`, `Statement`, `Symptoms`, `Body_parts`, `Cognitive_connection`, and `Emotion_directed` frames, as shown in (2a-f), respectively. Lexical units evoking each frame are shown in boldface, frame element names are in superscript.

- (2) Patient reports pain in the chest associated with household stress.
- [Patient]**<sup>Patient</sup> reports pain in the chest associated with household stress
  - [Patient]**<sup>Speaker</sup> **reports** [pain in the chest associated with household stress]<sup>Message</sup>
  - [Patient]**<sup>Patient</sup> reports **pain** [in the chest]<sup>Body\_part</sup> [associated with household stress]<sup>Descriptor</sup>
  - Patient reports pain in the **[chest]**<sup>Body\_part</sup> associated with household stress
  - Patient reports [pain in the chest]<sup>Concept\_1</sup> **associated** [with household stress]<sup>Concept\_2</sup>
  - Patient reports pain in the chest associated with [household]<sup>Reason</sup> **stress**

#### 4. Construction of Semantic Patterns of GBV

Referring back to the definition of frames, they describe structures that represent contextual meaning and enable the systematic identification of linguistic

and semantic patterns that may reveal experiences of violence. Therefore, in order to better understand the scenarios reported in medical records, as well as to look for patterns that could lead to possible unreported and overlooked victims of GBV, an analytical methodology was devised into two phases: **pattern proposition** and **validation through corpus querying**.

#### 4.1. Pattern proposition

The first phase consisted of manually analyzing texts to propose frame-based patterns. It was carried out through the study of the same anonymized sample used in the human annotation that resulted in the dataset used for training LOME. This process was done using two methods.

First, a group of eight possible scenarios covering types of violence was selected to be investigated. They were general violence, psychological violence, sexual abuse, negligence, infant labor, violence against LGBTQIA+ groups, violence against indigenous people and violence against people with disabilities. Subsequently, the LUs that could describe aspects in such cases of violence were selected and grouped to be used in a search for occurrences in the data. Once an LU was identified as present in the record, the report was individually studied to consider: (1) the frame evoked; (2) the frame elements present; (3) the lemmas marked as the frame elements; and (4) the frames those lemmas could be evoking. Based on that analysis, a possible pattern would then be defined.

To illustrate this process, consider the theme of general violence. The lemma *bater.v 'hit.v'* was chosen as a potential LU capable of indicating instances of such cases. Then this lemma was searched for in the corpus sample, which revealed examples similar to (3).

- (3) Paciente relata que o marido bate nela quando fica nervoso.  
*Patient reports that husband hits her when he gets angry.*

In this case, the LU evokes the `Cause_harm` frame, whose core frame elements are `AGENT`—the person causing the Victim's injury—and `VICTIM`—the being or entity that is injured. Accordingly, the FE `AGENT` corresponds to *marido.n 'husband'*, while the FE `VICTIM` corresponds to *paciente.n 'patient'*. The LU *marido.n 'husband'*, in turn, evokes the `Personal_relationship` frame. Therefore, we can link the aggressor as someone related to the victim, and a conceptual pattern can be identified: *Physical violence by a family member or person related to the victim*, where the FE `AGENT` in the `Cause_harm` frame is filled by an LU that evokes the `Personal_relationship` frame.

A second method to propose patterns consisted of the study of frames in the violence domain. The aim of this method was to further explore the non-core frame elements and how they could also carry relevant information based on the frames the lexical material they labeled were evoking. Each violence-domain frame was analyzed to determine which core and non-core frame elements were relevant when filled by specific lexical units. This allowed for the exclusion of some of the frames given its more generic character. Once the analysis was performed, the sample was again checked.

At the end of both processes, 19 patterns were defined within the eight possible scenarios of violence initially envisioned. This indicates that for some types of violence, more than one pattern was designed. For example, the pattern of *Physical violence by a family member or person related to the victim* was defined under the 'general violence' topic, which is based on the perspective of the agent who causes harm to a victim. However, a second pattern was defined under the same topic, focusing on the victim's perspective—that of experiencing an injury—activated by the `Experiencing_bodily_harm` frame. The decision to keep a pattern for the next phase was based on the number of occurrences in the sample and/or the researchers' intuition.

#### 4.2. Validation through corpus querying

The second phase of pattern construction consisted of querying the 19 patterns against the full corpus of 21 million sentences with automatic annotations. Two criteria were established to select useful patterns. For the first, any pattern with less than 30 matches in the corpus had to be discarded. This was the case for six patterns. The second criterion was based on manual inspection of the results (or a sample when they had more than a hundred matches): when patterns had too many false positives, they had to be discarded as they require further refinement and could not be useful as is. This was the case for five additional patterns.

The remaining eight patterns were kept for analysis and are listed in Section 6. They cover general violence, sexual violence and negligence. All patterns designed to identify cases of violence against LGBTQIA+, indigenous, and people with disabilities groups had less than ten matches, while queries related to infant labor and psychological violence had too many false positives. In the case of psychological violence, patterns related to fear covered fear of many things that are not necessarily related to violence, while threats were not common. Psychological violence is discussed further in Section 6.2.

## 5. Evaluation Design

To evaluate the precision of the proposed methodology in identifying reports of GBV in the e-medical records, the eight specific fine-grained patterns of GBV defined in Section 4 were queried against the automatically annotated corpus of sentences and the 5,000 matching sentences extracted. This sample was manually validated one more time to guarantee its full anonymization. Duplicated sentences were removed from the sample, resulting in a total of 4,186 texts.

Human evaluation of the the sample was conducted by six trained annotators, which were familiar with the phenomenon of GBV and with how patterns were defined. Evaluators had to indicate, in the annotation, whether the fragment was an exact instance of the pattern under consideration or not. An exact matching occurs when all parts of the pattern are present in the sentence and the patient is the victim. Partial correspondences with the pattern were judged as a non-matching instances. For example, for the pattern that specifies *injury by family member or person related to the victim*, cases where it was not possible to ensure that an aggressor was identified in the record or that they were not related to or acquainted with the victim were judged as non-matching. Similarly, if the pattern was referencing a person who was not the patient or if it was part of larger pattern such as 'attempt of', it was treated as non-matching. The goal of this round of evaluation is to capture the precision of the methodology in the fine-grained identification of cases of GBV.

Following this first round of evaluation, the same team of annotators performed an error analysis. For this second round, annotators had to indicate whether the mismatches identified in the first round were instances of GBV reports, nevertheless belonging to a different pattern, having partial match or describing, for example, an speculation of violence. In other words, they had to evaluate whether the system correctly identified a case of GBV that would be relevant for a health surveillance team, even if it incorrectly classified its fine-grained type. For this the second round, text fragment batches were shuffled among annotators so that one the same annotator would never judge a fragment they had already analyzed in the first round.

## 6. Results

### 6.1. Precision of the Pattern Identification System

Table 1 shows the results obtained for the first round of evaluation for the eight patterns of GBV described in Section 4.

In general, patterns that specify both the violence subtype and the aggressor were those with lower precision. There are two relevant explanations for this fact. The first is that violence subtypes are usually evoked by a wide range of verbs, increasing the chances of the semantic parser failing to fully learn the semantic properties of these frames. The second is that complex patterns such as relation to aggressor have more points where it can make a mistake, which reduces its matching chance.

### 6.2. Error Analysis

The error analysis demonstrated that a considerable part of the non-matching cases, were, in fact, cases of GBV reporting. In most cases, the errors in the first evaluation round were related to a partial correspondence between the pattern and the text fragment being evaluated. Consider, for example, the *physical violence by family member or person related to the victim* pattern, which presented the highest increase in precision post-error analysis, as shown in Table 2.

Among the 345 cases that were classified as non-matching in the first round but had some report GBV, 18 different types of violence were found. These were manually classified after the second round of evaluation to identify possible patterns of errors. Although distributed in many types, 90% of errors were classified in one of the six most common forms of violence. Table 3 shows the distribution of these most frequent types found during error analysis and groups the remaining into the 'Other' category.

These results may indicate that the base frame of this pattern (*Cause\_harm*) may be capturing more than just physical violence. In many instances, the semantic parser labeled this frame on verbs such as *curse*, *despise* or verbs with modification *attack* (*verbally*). This sort of nuance points towards the need of more specific examples on the parser's training set to restrict *Cause\_harm* annotation to physical examples.

Also note that psychological violence is the most common type in error analysis. As noted in Section 4, the psychological violence patterns initially created were not found in the original sample. The three patterns established to identify psychological violence were:

1. The LU *psicológico.a* 'psychological' evoking the *Domain* frame, when its *PREDICATE FE* was filled by an LU evoking the *Experience\_bodily\_harm*, *Cause\_harm*, or *Violence\_scenario* frames.
2. The LUs *ameaça.n* 'threat' and *ameaçar.v* 'to threaten', evoking the *Risk\_situation* frame, where the *DANGEROUS\_ENTITY FE* was filled by an LU evoking either the *Kinship* or *Personal\_relationship* frames.

Pattern	Correct	Error	Precision
Abandonment of child/elderly	92	27	0.773
Abuse by family member or person related to the victim	483	140	0.775
Abuse of child/elderly	92	17	0.844
Injury by family member or person related to the victim	493	386	0.560
Physical violence by family member or person related to the victim	574	1013	0.362
Sexual violence by family member or person related to the victim	57	43	0.570
Sexual abuse	362	90	0.800
Violence by family member or person related to the victim	525	66	0.888
<b>Overall</b>	<b>2678</b>	<b>1782</b>	<b>0.600</b>

Table 1: Precision in the fine-grained identification of GBV patterns.

Pattern	Correct	Error	Precision
Abandonment of child/elderly	95	24	0.798
Abuse by family member or person related to the victim	511	112	0.820
Abuse of child/elderly	94	15	0.862
Injury by family member or person related to the victim	593	286	0.675
Physical violence by family member or person related to the victim	919	669	0.579
Sexual violence by family member or person related to the victim	76	24	0.760
Sexual abuse	400	52	0.885
Violence by family member or person related to the victim	553	38	0.936
<b>Overall</b>	<b>3241</b>	<b>1220</b>	<b>0.726</b>

Table 2: Precision in the identification of GBV reports by pattern

Violence type	Count	Percentage
Psychological	110	31.9%
Witnessed	75	21.7%
Aggressive patient	42	12.2%
Verbal	33	9.6%
Abandonment	25	7.2%
Self-inflicted	25	7.2%
Other	35	10.2%
<b>TOTAL</b>	<b>345</b>	<b>100%</b>

Table 3: Distribution of types of violence found in the error analysis performed for the *physical violence by family member or person related to the victim* pattern

- General LUs evoking the `Fear` frame, where the `STIMULUS` FE was filled by an LU evoking either the `Kinship` or `Personal_relationship` frames.

An analysis of the 110 cases of psychological violence identified during the error analysis shows interesting reasons why the patterns were not found when querying the corpus or had too many false positives 4. First, among the 110 cases, only one features the LU *psicológico.a 'psicológico'* evoking the `Domain` frame. Second, although 33 cases refer to threats, only in approximately half the cases there is an LU evoking the `Risk_situation` frame. For the other half of the cases, the threat is introduced by LUs such as

*dizer.v 'say'*, which evoke the `Statement` frame. The notion of threat is inferred by the nature of the content of the reported speech, which would add sensible complexity to the pattern. A very similar situation occurs with the cases where the psychological violence is based on fear, which total 22 instances. In approximately one third of the cases, there is no explicit LU evoking the `Fear` frame.

The same analysis showed, however, that almost half the cases of psychological violence—49 instances—are reports of gaslighting, where the patient is the victim. Those reports are very diverse in terms of frames, and also represent a challenge for the pattern identification methodology.

## 7. Conclusion and Future Work

This study presented an innovative methodology for identifying underreporting of notifiable health events, focusing on Gender-Based Violence (GBV) through frame-based modeling and semantic parsing. The results show that this linguistically grounded framework effectively captures a wide range of GBV reports with good precision (average of 0.726).

While the patterns defined in this paper are specific to GBV, the core methodological contribution, *i.e.*, semantic parsing grounded in FrameNet combined with pattern definition over automatic annotations, is generalizable to other health surveillance tasks. This capability rests on two properties

of FrameNet. First, beyond the domain-specific healthcare and violence frames modeled by [Dutra et al. \(2023\)](#), FrameNet's broad network of general-domain and abstract frames already provides conceptual coverage that other tasks can make use of. Second, the methodology for creating new frames, lexical units, and frame-to-frame relations is well-established and has been applied with methodological rigor across decades of FrameNet development. Extending the system to tasks such as suicide prevention or substance abuse requires defining new patterns — and, only where necessary, new domain-specific frames — within this existing and principled infrastructure.

The same explicit, human-readable structure that makes FrameNet extensible also makes it interpretable. This quality makes the approach particularly suitable for healthcare applications: outputs can be easily understood and validated by professionals, even when the model errs. Such transparency strengthens the system's reliability and reinforces its practical value in real-world settings.

The methodology also aligns with the realities of Brazil's public healthcare infrastructure, where limited computational resources often constrain large-scale AI systems. The models employed have a lower carbon footprint and can be deployed on modest infrastructure while maintaining data governance and an ethical implementation. As reported in Section 3.2, the entire pipeline — from parser retraining to inference over 21 million sentences — consumed approximately 22.4 kWh and emitted an estimated 1.93 kgCO<sub>2</sub>eq, confirming the low environmental cost of deploying task-specific small language models in resource-constrained settings.

Future work will focus on developing frame-based risk models and automating the discovery of semantic patterns to reduce manual effort and enhance scalability.

## 8. Ethics and limitations

The research presented in this paper was approved by the Research Ethics Committee of the Federal University of Goiás (CAAE: 64733922.3.0000.5083; Approval number: 6.126.995).

**PII Removal:** The handling of the data used in the research reported in this paper required rigorous attention since, given the nature of violence notifications and e-medical records, the information logged in them is extremely sensitive and has the potential to lead to the identification of the patients. Moreover, some of the patients whose records were used are potential victims of violence and their identification could, apart from violating data privacy regulations, jeopardize their safety. Therefore, to ensure the protection of this information, all members in-

involved in the research signed confidentiality agreements. Also, as described in Section 2, all Personally Identifiable Information (PII) for every patient was removed before the team involved both in data annotation for frames and in evaluation of the performance of the pattern identification system had access to any piece of text obtained from violence notification and/or e-medical records. Moreover, access to the data was restricted even among team members, meaning that only anonymized samples were passed on to researchers.

**Annotator's compensation:** All annotators are members of the research team. For the annotation described in Section 3.1, annotation tasks were part of the work/research plan of team members, which was compensated through monthly stipends that were at least equivalent to the minimum wage defined by local regulations. For the human evaluation of the pattern recognition task, every annotator is a member of the research team and co-author of the paper. With the exception of one volunteer annotator, all annotators are funded by the project with a monthly stipend at least 38% higher than the local minimum wage.

**Implications of GBV identification to victims:** Although all data used in this research originate from a historical series ranging from 2016 to 2022, we recognize the potential of this methodology to be implemented as an alert system for healthcare professionals in real time. Such a system may expose victims to additional risks and serious analysis should be made in regards to which ranks of healthcare surveillance professionals should have clearance to access data originated from such an alert system. The research team has conducted extensive research on this issue and has been consulting with specialists in the fields of personal data regulation before any further step towards the use of the methodology as an alert system can be pursued.

**Technical limitations:** This methodology was developed as part of a broader project aimed at improving the use of health data to prevent tragic outcomes in Brazil. One of its key objectives is to better estimate underreporting of health-related events. To this end, frame-based modeling, semantic parser training, and the formulation and evaluation of linguistic patterns were conducted using texts in Brazilian Portuguese. Although the work focused on this language, the methodology itself is largely language-agnostic. Since only multilingual language models were employed, it can be applied to other languages with little or no adaptation. Nevertheless, the lack of thorough testing in languages other than Portuguese and the potential

bias of frames and lexical units toward Brazilian Portuguese remain a notable limitation.

## 9. Acknowledgments

This work was supported by the Patrick J. McGovern Foundation's acceleration program, the José Luiz Setúbal Foundation, and the Instituto Galo da Manhã. We also express our gratitude to our partners from the Recife Municipal Health Department — Luciana Caroline, Marcella Abath, Natalia Barros, and Yana Lopes — for their valuable collaboration and continuous support throughout this project. Laís Berno is funded by the Coordination for the Improvement of Higher Education Personnel (CAPES – grant 88887.964418/2024-00). Franciany Campos and Kenneth Brown are funded by the Research Support Foundation of the State of Minas Gerais (FAPEMIG - grant 5.45/2021, scholarship IDs 15366 and 11667, respectively). Tiago Torrent is a grantee of the Brazilian National Council for Scientific and Technological Development (CNPq – grant 311241/2025-5).

## 10. Bibliographical References

- Ali Aliabadi, Abbas Sheikhtaheri, and Hossein Ansari. 2020. [Electronic health record-based disease surveillance systems: A systematic literature review on challenges and solutions](#). *Journal of the American Medical Informatics Association*, 27(12):1977–1986.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexandre Diniz da Costa. 2020. [A tradução por máquina enriquecida semanticamente com frames e papéis qualia](#). Ph.D. thesis in Linguistics, Universidade Federal de Juiz de Fora, Juiz de Fora.
- Luís Filipe Cunha and José Carlos Ramalho. 2022. [Ner in archival finding aids: Extended](#). *Machine Learning and Knowledge Extraction*, 4(1):42–65.
- Lívia Dutra, Arthur Lorenzi, Lorena Larré, Frederico Belcavello, Ely Matos, Amanda Pestana, Kenneth Brown, Mariana Gonçalves, Victor Herbst, Sofia Reinach, Renato Teixeira, Pedro de Paula, Alessandra Pellini, Cibele Sequeira, Ester Sabino, Fábio Leal, Mônica Conde, Regina Grespan, and Tiago Torrent. 2023. [Building a frame-semantic model of the healthcare domain: Towards the identification of gender-based violence in public health data](#). In *Anais do XIV Simpósio Brasileiro de Tecnologia da Informação e da Linguagem Humana*, pages 338–346, Porto Alegre, RS, Brasil. SBC.
- Electricity Maps. 2025. Carbon intensity of electricity: Brazil Centro-Sul (BR-CS). <https://app.electricitymaps.com/map/zone/BR-CS/12mo/monthly>. 12-month average, accessed March 2025.
- Charles J. Fillmore. 1977. The case for case reopened. In Peter Cole and Jerrold M. Sadock, editors, *Syntax and Semantics, vol. 8: Grammatical Relations*, pages 59–81. Academic Press, New York.
- Charles J. Fillmore. 1982. Frame Semantics. In Linguistics Society of Korea, editor, *Linguistics in the morning calm*. Hanshin Publishing Co., Seoul, South Korea.
- Cléa Adas Saliba Garbin, Isabella de Andrade Dias, Tânia Adas Saliba Roviada, and Artênio José Ísper Garbin. 2015. [Desafios do profissional de saúde na notificação da violência: obrigatoriedade, efetivação e encaminhamento](#). *Revista Ciência & Saúde Coletiva*, 20(6):1879–1890.
- Claudia Garcia-Moreno and Charlotte Watts. 2011. [Violence against women: an urgent public health priority](#). *Bulletin of the World Health Organization*, 89:2–2.
- Pierre Guillou. 2021. [NER-BERT-Base-Cased-pt-lenerbr: BERT-based NER model for Portuguese \(legal domain\)](#). <https://huggingface.co/pierreguillou/ner-bert-base-cased-pt-lenerbr>. Accessed: 2023-08-30.
- Adam Kilgarriff, Vít Baisa, Jan Bušta, Miloš Jakubíček, Vojtěch Kovář, Jan Michelfeit, Pavel Rychlý, and Vít Suchomel. 2014. [The Sketch Engine](#). *Lexicography*, 1(1):7–36.
- Luciana Kind, Maria de Lourdes Pereira Orsini, Valdênia Nepomuceno, Letícia Gonçalves, Gislaíne Alves de Souza, and Monique Fernanda Félix Ferreira. 2013. [Subnotificação e \(in\) visibilidade da violência contra mulheres na atenção primária à saúde](#). *Cadernos de Saúde Pública*, 29:1805–1815.
- Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. 2019. [Quantifying the carbon emissions of machine learning](#). *arXiv preprint arXiv:1910.09700*.

- Lorena Larré and Tiago Torrent. 2024. [Modelagem baseada em frames para identificação do léxico da violência de gênero](#). In *Anais do XV Simpósio Brasileiro de Tecnologia da Informação e da Linguagem Humana*, pages 403–412, Porto Alegre, RS, Brasil. SBC.
- Zabir Al Nazi and Wei Peng. 2024. [Large language models in healthcare and medical domain: A review](#). *Informatics*, 11(3).
- Minh Van Nguyen, Viet Dac Lai, Amir Pouran Ben Veyseh, and Thien Huu Nguyen. 2021. [Trankit: A light-weight transformer-based toolkit for multilingual natural language processing](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 80–90, Online. Association for Computational Linguistics.
- James Pustejovsky. 1995. *The Generative Lexicon*. MIT Press, Cambridge, USA.
- Josef Ruppenhofer, Michael Ellsworth, Miriam Petruck, Christopher Johnson, and Jan Schefczyk. 2016. *FrameNet II: Extended Theory and Practice*. International Computer Science Institute (ICSI).
- Fábio Souza, Rodrigo Nogueira, and Roberto Lotufo. 2019. [Portuguese Named Entity Recognition using BERT-CRF](#). *arXiv preprint arXiv:1909.10649*.
- Patricia L Sweet. 2014. [Every bone of my body: Domestic violence and the diagnostic body](#). *Social Science & Medicine*, 122:44–52.
- Tiago Timponi Torrent, Ely Edison da Silva Matos, Frederico Belcavello, Marcelo Viridiano, Maucha Andrade Gamonal, Alexandre Diniz da Costa, and Mateus Coutinho Marim. 2022. [Representing Context in FrameNet: A Multidimensional, Multimodal Approach](#). *Frontiers in Psychology*, 13.
- Tiago Timponi Torrent, Ely Edison da Silva Matos, Alexandre Diniz da Costa, Maucha Andrade Gamonal, Simone Peron-Corrêa, and Vanessa Maria Ramos Lopes Paiva. 2024. [A flexible tool for a qualia-enriched FrameNet: the FrameNet Brasil WebTool](#). *Language Resources and Evaluation*, pages 1–29.
- WHO. 2024. [World Health Organization Health Topics: Violence Against Women](#). Accessed: October 6, 2025.
- 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.
- Ann Öhman, Marika Burman, Maria Carbin, and Kerstin Edin. 2020. [The public health turn on violence against women: analysing swedish health-care law, public health and gender-equality policies](#). *BMC Public Health*, 20:1–12.