

EventNet-ITA: Italian Frame Parsing for Events

Marco Rovera

Fondazione Bruno Kessler, Trento, Italy

m.rovera@fbk.eu

Abstract

This paper introduces EventNet-ITA, a large, multi-domain corpus annotated *full-text* with event frames for Italian. Moreover, we present and thoroughly evaluate an efficient multi-label sequence labeling approach for Frame Parsing. Covering a wide range of individual, social and historical phenomena, with more than 53,000 annotated sentences and over 200 modeled frames, EventNet-ITA constitutes the first systematic attempt to provide the Italian language with a publicly available resource for Frame Parsing of events, useful for a broad spectrum of research and application tasks. Our approach achieves a promising 0.9 strict F1-score for frame classification and 0.72 for frame element classification, on top of minimizing computational requirements. The annotated corpus and the frame parsing model are released under open license.

1 Introduction

Frame Parsing is a powerful tool for real-world applications in that it enables deep grasp of the meaning of a textual statement and automatic extraction of complex semantic descriptions of situations, including events, and their relations with the entities involved. To this effect, Frame Parsing can effectively be used for Event Extraction, as the two tasks share the common goal of recognizing and classifying argument structures of a target predicate. However, training supervised models for Frame Parsing is a data-intensive task, which is why comprehensive linguistic resources are available in few languages, thereby limiting further research and application in downstream tasks. Furthermore, most existing corpora are created by targeting the annotation of one single frame (or event) class per sentence. While lexicographically motivated, this procedure makes the training of automatic models and their application to real-world scenarios more complicated and burdensome.

The contribution described in this paper is twofold: first, we present EventNet-ITA (EvN-ITA) a large-scale, multi-domain corpus annotated *full-text* (see Section 4) with over 200 semantic frames of events (Fillmore and Baker, 2001) and 3,600 specific frame elements in Italian, also discussing the motivation behind its creation, the annotation guidelines and the covered domains; secondly, we introduce an efficient multi-label sequential approach for eventive Frame Parsing and evaluate it on the dataset. This work aims at providing the community with a solid, manually-curated corpus and a ready-to-use model for frame-based event extraction for Italian, thus filling an existing data gap. In fact, recent works in application fields like computational social science (Minnema et al., 2021, 2022) or historical NLP (Sprugnoli and Tonelli, 2017; Menini et al., 2023) showed how semantic frames can be used as a powerful textual analysis tool to investigate a wide range of societal and historical phenomena. The annotated dataset, along with its full documentation, is released to the community under open license (see Section 7). The envisioned application purpose of EvN-ITA is that of enabling accurate mining of events from large collections of documents, with focus on individual, social and, in a broad sense, historical phenomena. The paper is structured as follows: Section 2 discusses existing work in Frame Parsing and Event Extraction, with a subsection focused on Italian; Section 3 introduces our annotated corpus and describes the motivations and design decisions that guided its creation, while Section 4 focuses on the annotation procedure. In Section 5 we discuss the methodology for Frame Parsing, a transformer-based multi-label sequence labeling approach. In Section 6 we evaluate our methodology and discuss the results. Section 7 provides the reader with pointers for the dataset and model release, while Section 8 concludes the paper and highlights future directions of our work.

2 Related Work

The development of systems able to recognize and classify event mentions and their argument structure in text has been a long-term effort in computational linguistics and a variety of methods has been employed for the task of Event Extraction (Ahn, 2006; Liao and Grishman, 2010; Chen et al., 2015; Nguyen et al., 2016; Orr et al., 2018; Nguyen and Nguyen, 2019; Lu et al., 2021; Paolini et al., 2021). Event Extraction is the task of recognizing and classifying event mentions and entities involved in the event from a textual statement and it has seen applications in a wide range of fields, like social media analysis (de Bruijn et al., 2019), biomedical NLP (Li et al., 2019; Huang et al., 2020; Ramponi et al., 2020), history and humanities (Segers et al., 2011; Cybulska and Vossen, 2011; Sprugnoli and Tonelli, 2019; Lai et al., 2021; Rovera et al., 2021), as well as literary text mining (Sims et al., 2019). Although benchmark datasets exist, like Automatic Content Extraction (ACE) (Walker et al., 2006) for Event Extraction or TAC-KBP (Ellis et al., 2015) for multiple event-related tasks, they exhibit limitations in terms of size and domain coverage. Also, while they are well suited for evaluation campaigns, they have not been designed for use in real-world application tasks. Moreover, most of these corpora only exist for English, with few extensions for other languages (Ji et al., 2016).

2.1 Frame Parsing

Frame Parsing (Das et al., 2014; Swayamdipta et al., 2017, 2018) consists in recognizing, in a textual expression, a word or set of words (the *lexical unit*) as the predicate evoking a given frame and isolating the text spans that evoke the semantic arguments (*frame elements*) related to that frame. Frames are conceptual structures describing prototypical situations and their realizations in text. The reference linguistic resource for Frame Parsing in English is FrameNet (FN) (Baker et al., 1998; Fillmore and Baker, 2001; Ruppenhofer et al., 2006). In this work, we use Frame Parsing for extracting event frames. While event extraction initiatives have been based on a variety of models, approaches and schemes, which are not always interoperable or comparable, the advantage of using Frame Parsing for Event Extraction is the availability of an established resource, based on a unified, grounded theoretical framework (Fillmore et al., 1976). EvN-ITA differs from FN in that the latter is based on

lexicographic annotation (one target lexical unit per sentence), providing only a small subset of full-text annotated data (Ruppenhofer et al., 2016), whereas EvN-ITA has been annotated by design in a *full-text* fashion (see Section 4). Also, it is important to point out that EvN-ITA is not meant to be a comprehensive Italian version of the popular English FN. Instead, in this work we adopt part of the FN schema but focus exclusively on event-denoting frames, aiming at providing a large, self-contained and robust tool for frame-based Event Extraction in Italian.

2.2 Italian Event Extraction and Frame Semantics

As for Italian, the Frame Labeling over Italian Texts Task (FLAIT) was organized at EVALITA in 2011 (Basili et al., 2012). Moreover, Event Extraction in Italian was the object of the EVENTI evaluation campaign at EVALITA 2014 (Caselli et al., 2014), which focused on temporal processing and was based on the Ita-TimeBank schema (Caselli et al., 2011). Later on, Caselli (2018) experimented with the same dataset using a neural architecture and evaluated the impact of different word embeddings for event extraction. While Italian Event Extraction approaches have traditionally been based on the TimeML (Saurí et al., 2006) classification scheme, which provides 7 broad, temporal-oriented classes, more recently the necessity has emerged of a more fine-grained annotation schema for event classification, as discussed by Sprugnoli and Tonelli (2017). Supported by a survey involving historians, the authors investigated the application of event extraction on historical texts. Sprugnoli and Tonelli (2019) describe a specific schema, adapting semantic categories provided by the Historical Thesaurus of the Oxford English Dictionary (HTOED) (Kay et al., 2009), resulting in 22 topic-driven event classes, thereby moving towards developing a richer and at the same time finer-grained inventory of classes for representing events in text. As for frame semantics, on the other hand, Basili et al. (2017) and Brambilla et al. (2020) described a work in progress for the creation of IFrameNet, a large scale Italian version of FN, by using semi-automatic methods and manual validation for frame induction, with 5,208 sentences annotated with at least one lexical unit. However, the dataset has not been released so far. In fact, despite the considerable amount of work in lexical (Lenci et al., 2012a;

Jezek et al., 2014) and frame semantics (Tonelli and Pianta, 2008; Tonelli et al., 2009; Lenci et al., 2010, 2012b), Italian still lacks an extensive, publicly available linguistic resource for Frame Parsing.

3 Dataset

3.1 EventNet-ITA

In order to ensure multilingual compatibility, we employ a selection of event frames from FN (Baker et al., 1998; Ruppenhofer et al., 2006), where available. This way, 85% of EvN-ITA classes are mapped to FN schema, either by direct match (59%) or by subclassing (26%). In a minor number of cases (15% of the schema), where target phenomena are not covered in FN, an *ad hoc* frame class has been created. Frame-to-frame mappings between EvN-ITA and FN are provided in the documentation of the resource. Table 1 offers a quantitative description of the corpus. EvN-ITA counts

Annotated sentences	53,854
Tokens	1,583,612
Vocabulary (words)	97,512
Avg. sentence length (tokens)	29
Modeled event frames	205
Modeled frame elements	3,571
Lexical units	837
Frame instances	102,294
Frame element instances	180,279
Frame instances per sentence (avg.)	1.9
Examples per class (avg.)	491

Table 1: Statistics of the EvN-ITA dataset.

53,854 annotated sentences, including negative examples (see Section 4.2), 102,294 event instances and over 1.5 million tokens, annotated by an experienced annotator (native speaker) with background in Frame Semantics. Each frame class has on average 491 annotated examples.

The corpus - as well as the annotation schema - has been created with the purpose of covering historical narratives in a broad sense, but without committing to a specific textual genre. For this reason, as well as for creating a releasable corpus, sentences for the annotation set of EvN-ITA have been sampled from a subset of the Italian Wikipedia edition. In order to filter out irrelevant documents, i.e. documents not likely to contain events, we collected Wikipedia pages falling under the categories *Events*

*by country*¹ and *History by country*.² This choice ensures a wide variety of featured events, both temporally (from ancient history to the present days) and geographically.

Through standard pre-processing (tokenization, lemmatization and dependency parsing have been performed using TINT³ (Palmero Aprosio and Moretti, 2018)), a pool of sentences, arranged by lemma, was generated, from which to pick for the annotation set. Annotated sentences are drawn from 16,309 different Wikipedia articles.

3.2 Domain coverage

In the design phase of the resource, a manual analysis was made of existing corpora in multiple languages, in order to circumscribe the domains and classes to be modelled. Resources as Automatic Content Extraction (Doddington et al., 2004; Consortium et al., 2005), Event Nugget (Mitamura et al., 2015), the Historical Thesaurus of the Oxford English Dictionary (HTOED) (Kay et al., 2009) and FN (Baker et al., 1998) were reviewed and compared. FN is currently the most complete, rich and established existing resource and has been taken as reference for the development of EvN-ITA. This choice is motivated by the opportunities it offers in terms of reuse, coverage and possible multilingual extensions. EvN-ITA’s annotation schema covers 205 different event frames, each provided with a set of specific frame elements (unique modeled frame elements amount to 3,571), and has been extensively documented by providing, for each frame, its definition, the corresponding set of lexical units and frame elements associated to it. The distribution of classes, arranged by topic, is depicted in Figure 1. Beside conflict-related events, that hold a prominent place in historical accounts and journalistic narratives, we have taken care to extend the collection of event types to other aspects of the life of societies and individuals, such as legislative and legal processes, work, establishment of and membership in social organizations, life events, as well as events related to the arts, economic processes and cognitive processes such as decisions, skills, judgements, amongst others. In the design of the resource, attention has been paid also to maintain the balance between internal coherence and usability.

¹https://it.wikipedia.org/wiki/Categoria:Eventi_per_stato

²https://it.wikipedia.org/wiki/Categoria:Storia_per_stato

³<https://dh.fbk.eu/research/tint/>




Figure 1: Macro-topics covered in EvN-ITA (in brackets, the number of frames belonging to each domain).

ity of the class schema. This has been achieved in multiple ways:

- by including for each class, where existing, also its opposite (HIRING / FIRING, CREATE SOCIAL ENTITY / END SOCIAL ENTITY, VEHICLE TAKE OFF / VEHICLE LANDING);
- by making possible narrative chains (e.g. COMMITTING CRIME, ARREST, TRIAL, SENTENCING, IMPRISONMENT, CAPTIVITY, RELEASING);
- by providing couples of classes representing an event and the subsequent logical state (e.g. BECOMING A MEMBER / MEMBERSHIP, BECOMING AWARE / AWARENESS, MAKE ACQUAINTANCE / ACQUAINTANCE).
- by ensuring a certain degree of redundancy and perspective (BEING IN PLACE / TEMPORARY STAY, BEING BORN / GIVING BIRTH).

4 Annotation

The textual corpus, generated as discussed in Section 3.1, has been manually annotated by labeling event triggers with their frame class and predicate arguments with the corresponding frame element. Annotation was performed at the sentence level and was conducted frame-driven, by first selecting significant event frames for the domain and subsequently identifying the most relevant lexical units for each frame. Given a sentence, *any* lexical units in our schema and all related frame elements are annotated, producing as many layers as there are event mentions (*full-text* annotation). Figure

2 shows an example of full-text annotation from EvN-ITA.

4.1 Format

The IOB2 annotation format is being used, in which the B-tag identifies the first token of a span, the I-tag identifies all tokens inside the span and the O-tag all out-of-mention tokens. Discontinuous mentions are allowed, both for frames and for frame elements. The only constraint for event mentions is that they cannot overlap: each token in a sentence can denote at most one event type. This does not hold for frame elements: in fact, given a sentence with multiple frame occurrences, frame elements from different annotation sequences (i.e. belonging to different frames) can always overlap, hence a token may be labeled with more than one frame element tag⁴ (See Figure 2).

4.2 Guidelines

EvN-ITA is thoroughly documented, both in the form of general annotation guidelines (what to annotate) and at the annotation schema level (frame description, lexical units, frame elements).

As for lexical units, we exclusively focus on nouns, verbs and multi-word expressions. Although also other parts-of-speech (adverbs, for example) can be loosely event-evoking, this focus is motivated by practical reasons: nouns and verbs, along with multi-word expressions, are the most frequent triggers of event mentions in text and are characterized by a richer syntactic structure, which in turn is crucial for harvesting information related to frame elements. Nouns are annotated as event triggers only if they reference *directly* the occurrence of an event, but not when the reference is indirect, for example:

[...] *culminarono a Blois nel 1171 con la [morte **DEATH**] sul rogo di 31 ebrei.* (culminated in Blois in 1711 with the death at the stake of 31 Jews.)

Nell'aprile del 1700, Giovanni si ammalò terribilmente e si trovò quasi sul letto di [morte \emptyset]. (In April 1700, John fell terribly ill and was nearly on his deathbed.)

⁴In addition, a frame element may overlap with a frame mention if they belong to different annotation sequences, which is quite often the case. Only an frame-frame overlap is excluded.



Figure 2: An example of full-text annotation in EvN-ITA (English translation: *The construction of the Alvitian fortification dates back to the time of the Norman invasion.*).

As for verbs, we annotate the main verb but not the auxiliary.

*Appena ricostruita dalla devastazione del sisma la città fu [distrutta **DESTROYING**] nuovamente [...]* (As soon as it was rebuilt from the devastation of the earthquake, the city was destroyed again [...])

Multi-word expressions are annotated when they break compositionality, for example in *radere al suolo* (raze to the ground) or *aprire il fuoco* (opening fire), or in verbal periphrastic use *essere al corrente* (being aware) or *fare visita* (paying a visit). Frame mentions are annotated regardless of their factuality value, which means that also negated or hypothetical frame mentions must be annotated, as well as those introduced by modals. Conversely, in EvN-ITA we do not annotate as frame mention lexical units that are used with metaphorical meaning or in the form of rhetoric expression.⁵

EvN-ITA’s annotation schema consists of 205 frames and 837 lexical units, out of which 358 have at least 100 annotations each, 191 have a number of annotations comprised between 50 and 99, and 288 have a number of annotations comprised between 20 and 40. The annotation process has been oriented to keep the balance between frame completion and polysemy preservation. For this reason, we also annotated less frequent lexical units encountered in the corpus, resulting in a long queue of lexical units with less than 20 occurrences each. This strategy was adopted in order both to increase the flexibility of the resource and to set the stage for its future extension. Also, with the aim of improving robustness, for each lexical unit we annotated a number of *negative examples*, i.e. sentences in which the given lexical unit occurs without triggering any of the corresponding frames.⁶

Within the scope of this work, we consider as events any accomplishment, achievement or process, without distinction. The schema additionally

⁵More annotation examples for such cases are provided in Appendix A.

⁶See Appendix A for negative examples.

models a number of states (e.g. BEING IN PLACE, CAPTIVITY) and relations (LEADERSHIP, DURATION RELATION, POSSESSION). As for semantic roles, we referred to FN’s frame elements, with minor adaptations or additions, which in most cases tend towards increased specificity.

4.3 Inter-Annotator Agreement

EvN-ITA was annotated by one single native speaker annotator with a solid background in Frame Semantics. For this reason, particular attention has been devoted to assessing the robustness, consistency and intelligibility of the resource by means of inter-annotator agreement analysis. We therefore validated our schema and guidelines by re-annotating 2,251 sentences, spanning over 61 classes, with a second (native speaker) annotator. When selecting classes to be included in this validation set, we paid attention to include pairs or triplets of frames with high semantic similarity,⁷ in order to stress the test. We used two different metrics for assessing agreement: Jaccard Index (computed as the ratio between the number of items annotated with the same label and the sum of all annotated items) and Cohen’s Kappa (Cohen, 1960). The scores have been computed both at token level (relaxed) and at span level (strict). Results are reported in Table 2. Considering the high number of

	Jaccard		Cohen’s K	
	Token	Span	Token	Span
Lexical units	0.952	0.945	0.951	0.944
Frame elements	0.878	0.832	0.877	0.830

Table 2: Inter-annotator agreement scores.

different frames and frame elements in EvN-ITA, we observe that agreement values are high and indicate that the guidelines are sufficiently detailed in their description of the linguistic phenomena to be annotated. As expected, the annotation of frame elements has proven more challenging. A further manual analysis conducted on a sample of

⁷For example: BLAMING / ACCUSE, REPORTING / DENOUNCE, MARRIAGE / BEING_MARRIED.

245 sentence pairs with low agreement showed that disagreement had three main sources:

Ontological (67,5%) a textual span is recognized as frame or frame element by one annotator but not by the other;

Span Length (20,4%) annotators agree on the label but not on the exact span to annotate;

Classification (12%) annotators agree on the span to annotate but assign two different labels.

5 Methodology

A traditional approach for Frame Parsing, as well as for Event Extraction, is to break down the problem into sub-tasks (Das et al., 2014; Ahn, 2006), usually separating the steps of trigger identification, frame classification and argument extraction. However, a major downside of this approach, besides being more complex, is that it implies error propagation from higher-level sub-tasks downwards. Instead, we propose to learn all the tasks in one single step, allowing the model to simultaneously exploit tag relations on the time (sequence) axis and on the token axis. Thus, in this work Frame Parsing is approached end-to-end and is treated as a multi-label sequence tagging problem. The strength of this design option lies in its simplicity as it requires minimal pre-processing and does not imply the use of additional knowledge, as well as in its efficiency, as it minimizes computational requirements (see Section 5.2).

5.1 Preprocessing

The adoption of full-text annotation implies, at preprocessing time, the definition of each frame element as frame-specific, in order to avoid overlaps between frame elements with the same name but referring to different frames. In fact, many frames belonging to the same semantic area share a set of frame elements with the same name. For example, both motion frames FLEEING and MOTION_DOWNWARDS have a frame element called MOVER. In EvN-ITA, frame elements referring to the same *semantic role* (thus carrying the same name) but belonging to different frames are assigned different, frame-aware labels. Therefore, frame elements MOVER-FLEEING and MOVER-MOTION_DOWNWARDS will be assigned two different labels. This data encoding strategy, in return, allows us to minimize the need for post-processing (as each predicted frame element is implicitly linked to its frame) and enables the model

to learn relationships between multiple frame elements occurring on the same token/span.

5.2 Experimental Setup

Our frame parser aims at jointly extracting all frame mentions and all related frame elements in a target sentence. In other words, given an input sentence, each token must be labeled according to the event frame and/or frame element(s) it denotes. The underlying idea is to leverage mutual co-occurrence between frame (and frame element) classes, as certain frames typically tend to appear more often with, or have a semantic preference for, other frames.⁸ This way, the model is led to not only learn correspondences between a word and a given frame or frame element, but also local patterns of co-occurrence between different frame elements.

In order to provide a reliable performance assessment, we opted for an 80/10/10 *stratified* train/dev/test split, thus ensuring the same proportion of (frame) labels in each split. Moreover, we generate 4 folds from the dataset, the first used for hyperparameter search and the remaining three for evaluation. To this purpose we fine-tune a BERT model⁹ (Devlin et al., 2019) for Italian and show that the approach allows us to scale with thousands of (unique) labels without a remarkable computational and memory overhead. In this experimental setup we use MaChAmp, v 0.4 beta 2 (van der Goot et al., 2021), a toolkit supporting a variety of NLP tasks, including multi-label sequence labeling. We performed hyperparameter search by exploring the space with batch sizes between 8 and 256 and learning rates between 7.5e-4 and 7.5e-3. All other hyperparameters are left unchanged with respect to MaChAmp’s default configuration for the multi-sequential task.¹⁰ Overall, 64 configurations have been explored. The best hyperparameter values we found, according to the performance on the development set, are batch size of 64 and learning rate of 1.5e-3 and the resulting model has been used for the evaluation (Section 6). The training requires approximately 3.5 hours on an NVIDIA RTX A5000 GPU with 24 GB memory and 8192 CUDA cores.

⁸This assumption is mentioned in previous work (Liao and Grishman, 2010) and we verified it in our dataset by analysing frame relationships with several co-occurrence measures, such as Pointwise Mutual Information (see Appendix B).

⁹<https://huggingface.co/dbmdz/bert-base-italian-xxl-cased>

¹⁰<https://github.com/machamp-nlp/machamp/blob/master/docs/multiseq.md>

		Frames $n = 40$			Frame Elements $n = 200$		
		P	R	F1	P	R	F1
TEST	All classes	0.904	0.914	0.907	0.841	0.724	0.761
	All classes weighted	0.909	0.919	0.913	0.85	0.779	0.804
	Best n classes	0.974	0.982	0.978	0.938	0.912	0.923
	Worst n classes	0.811	0.808	0.806	0.72	0.441	0.516
	n most frequent classes	0.912	0.933	0.922	0.861	0.831	0.843
	n least frequent classes	0.865	0.871	0.865	0.781	0.493	0.575

Table 3: *Token-based (relaxed)* performance for multi-label sequential Frame Parsing (macro average, aggregate). Figures in bold represent the reference performance values for EventNet-ITA.

		Frames $n = 40$			Frame Elements $n = 200$		
		P	R	F1	P	R	F1
TEST	All classes	0.906	0.899	0.901	0.829	0.666	0.724
	All classes (weighted)	0.909	0.903	0.905	0.853	0.711	0.768
	Best n classes	0.975	0.976	0.975	0.937	0.867	0.898
	Worst n classes	0.81	0.789	0.796	0.673	0.398	0.476
	n most frequent classes	0.915	0.917	0.915	0.878	0.762	0.813
	n least frequent classes	0.879	0.866	0.87	0.743	0.441	0.529

Table 4: *Span-based (strict)* performance for multi-label sequential Frame Parsing (macro average, aggregate). Figures in bold represent the reference performance values for EventNet-ITA.

In terms of memory, the maximum requirement is 5 GB RAM.

6 Evaluation

In this section, we discuss the quantitative (Section 6.1) and qualitative (Section 6.2) performance of the multi-label sequence labeling approach on the EvN-ITA dataset.

6.1 Quantitative results

Evaluation results are reported in Table 3 and Table 4 in an aggregated fashion in order to provide the reader with different views on performance. Reported values have been obtained by separately computing the metrics class-wise on each fold, and then averaging the obtained scores. For each of the two groups of labels (frames and frame elements), beside the overall average performance, we provide the average of the n -best and n -worst performing classes and the average of the n most and least frequent classes in the dataset, on the three test sets, with $n = 40$ for frames and $n = 200$ for frame ele-

ments.¹¹ We also compute the macro average and the weighted macro average of all classes, the latter providing a more realistic view in a context of highly unbalanced label distribution. With a strict F1-score of 0.9 for frames and 0.724 for frame elements, our system shows very promising results for the task. Overall, the results show that, despite being fundamentally token-based, our multi-label sequence tagging approach proves effective also in the identification of (multiple) textual spans in a sentence, scaling well on a dataset involving a very high number of classes. This is further confirmed by the small delta between relaxed and strict performance values.

6.2 Error Analysis

To assess the potential of the proposed approach and the possible inconsistencies, we perform an error analysis on the test sets of the three folds,

¹¹In the case of frame elements, given their extremely skewed distribution, resulting in a long tail of rare labels, we also apply a threshold, taking into account only labels occurring at least 5 times in the span-based setting and at least 20 times in the token-based setting.

Gold	Predicted
BLAMING	ACCUSE
HOSTILE_ENCOUNTER	WAR
CONQUERING	OCCUPANCY
ACCUSE	BLAMING
REPLACING	TAKE_PLACE_OF
OCCUPANCY	CONQUERING
BUILDING	MANUFACTURING
CREATE_ARTWORK	TEXT_CREATION
KILLING	DEATH
REQUEST	QUESTIONING

Table 5: Top 10 prediction errors between event frames.

at token level. Since in a multi-label setting it is not always possible to establish a univocal correspondence between labels in the gold and predicted sets (given the possibility of multiple assignments on both sides), we proceed as follows: for each token, we filter out from both sets of labels (gold and predicted) the correctly matched labels. Based on this output, we focus on a subset of tokens, those labeled, in both sets, with exactly one label and we use it as an approximation for identifying most common errors. This allows us to focus on specific one-to-one label mismatches, both for event frames (Table 5) and for frame elements (Table 6). Considering only event frames, analysis reveals that only 4.8% of the identified mismatches involves two event labels, while 95.2% involves a mismatch between an event label and the O-tag (out-of-mention). This ratio becomes more balanced with regard to frame elements (39% and 61%, respectively). Also, the impact of errors referred to the IOB schema remains very low, amounting to 1.16% for event frames and 4.8% for frame elements.

Qualitatively, the analysis shows a clear pattern, namely that errors occur in most cases between frames with a high semantic similarity, like BLAMING/ACCUSE or HOSTILE_ENCOUNTER/WAR, which in some cases may be difficult to classify even for the human annotator. As for frame elements, errors occur mostly *a*) between the same frame element of two different event frames (for example MESSAGE-REQUEST vs. MESSAGE-QUESTIONING) or *b*) between frame elements that have a latent semantic correspondence in different frames (INTERLOCUTOR2-CONVERSATION vs. PARTY2-NEGOTIATION or REASON-BLAMING vs. OFFENSE-ACCUSE) or, still, *c*) between seman-

Frame element - Event frame
(G) REASON-BLAMING
(P) OFFENSE-ACCUSE
(G) INTERLOCUTOR2-CONVERSATION
(P) PARTY2-NEGOTIATION
(G) LOCATION-BEINGLOCATED
(P) RELATIVELOCATION-BEINGLOCATED
(G) MESSAGE-REQUEST
(P) MESSAGE-QUESTIONING
(G) RELATIVELOCATION-BEINGLOCATED
(P) LOCATION-BEINGLOCATED
(G) MESSAGE-ANSWER
(P) MESSAGE-REPLY
(G) ISSUE-TAKINGSIDES
(P) SIDE-TAKINGSIDES
(G) ARTWORK-CREATEARTWORK
(P) TEXT-TEXTCREATION
(G) EVALUEE-BLAMING
(P) ACCUSED-ACCUSE
(G) EXPLANATION-DEATH
(P) CAUSE-DEATH

Table 6: Top 10 prediction errors between frame elements (G = Gold, P = Predicted).

tically close frame elements within the same frame (EXPLANATION-DEATH vs. CAUSE-DEATH). These quite subtle error types further reveal how the multi-label sequence labeling approach is capable of learning cross-frame correspondences of frame elements, an aspect that we plan to further investigate in future work.

7 Dataset and Model Release

The EvN-ITA annotated dataset, along with its documentation, is being released upon request¹², under CC-BY-SA 4.0 license¹³. The model of the frame parser, described in Section 5.2, is available on Huggingface¹⁴.

8 Conclusion and Future Works

In this paper we presented EvN-ITA, a large corpus annotated with event frames in Italian, accompanied by an efficient multi-label sequential model for Frame Parsing, trained and evaluated on the

¹²The dataset can be requested by filling out the form at <https://forms.gle/qAgZsf4La9qdzETn6> or by emailing the author at eventnetita@gmail.com.

¹³<https://creativecommons.org/licenses/by-sa/4.0/deed.en>

¹⁴<https://huggingface.co/mrovera/eventnet-ita>

corpus. Future work includes extrinsic tests of the resource on new data from different textual genres and the reinforcement of the schema, in view of providing a wider domain coverage and increased adaptability of the model. Moreover, is our plan to employ EvN-ITA as a benchmark to investigate the performance of different methodologies and learning models for Frame Parsing, as well as to explore strategies for multilingual applications.

Limitations

The first limitation of this work lies in the unique source of the data, Wikipedia, that, if on the one end guarantees an ample variety of topics and types of events, on the other hand, from the linguistic point of view it sets a constraint on a homogeneous linguistic style. As mentioned above, this will be the focus of our future effort. Secondly, in case of multiple mentions of the *same* event frame in a given sentence (this case concerns 6% of the sentences in EvN-ITA), the currently adopted methodology does not support automatic linking of frame elements to the exact frame mention they refer to in the sentence. Future approaches will be geared to take this issue into account.

Acknowledgements

I owe my gratitude to Serena Cristoforetti, for her dedicated contribution to the inter-annotator assessment process, carried out as part of her curricular internship at the Digital Humanities (DH) group. I am equally thankful to Alan Ramponi and Sara Tonelli (DH) and to Enrica Troiano (Vrije Universiteit Amsterdam) for proofreading the manuscript and for their precious and constructive feedback. This work and the effort it represents are dedicated to the memory of Anna Goy.

References

- David Ahn. 2006. The stages of event extraction. In *Proceedings of the Workshop on Annotating and Reasoning about Time and Events*, pages 1–8.
- Collin F Baker, Charles J Fillmore, and John B Lowe. 1998. The berkeley framenet project. In *Proceedings of the 17th international conference on Computational linguistics-Volume 1*, pages 86–90. Association for Computational Linguistics.
- Roberto Basili, Silvia Brambilla, Danilo Croce, and Fabio Tamburini. 2017. Developing a large scale framenet for italian: the iframenet experience. *CLiC-it 2017 11-12 December 2017, Rome*, page 59.
- Roberto Basili, Diego De Cao, Alessandro Lenci, Alessandro Moschitti, and Giulia Venturi. 2012. Evalita 2011: The frame labeling over italian texts task. In *International Workshop on Evaluation of Natural Language and Speech Tool for Italian*, pages 195–204. Springer.
- Silvia Brambilla, Danilo Croce, Fabio Tamburini, Roberto Basili, et al. 2020. Automatic induction of framenet lexical units in italian. In *CEUR WORKSHOP PROCEEDINGS*, volume 2769. CEUR-WS.
- Tommaso Caselli. 2018. Italian event detection goes deep learning. *arXiv preprint arXiv:1810.02229*.
- Tommaso Caselli, Valentina Bartalesi Lenzi, Rachele Sprugnoli, Emanuele Pianta, and Irina Prodanof. 2011. Annotating events, temporal expressions and relations in italian: the it-time experience for the ita-timebank. In *Proceedings of the 5th Linguistic Annotation Workshop*, pages 143–151. Association for Computational Linguistics.
- Tommaso Caselli, Rachele Sprugnoli, Manuela Speranza, and Monica Monachini. 2014. Eventi: Evaluation of events and temporal information at evalita 2014. *EVENTI: Evaluation of Events and Temporal Information at Evalita 2014*, pages 27–34.
- Yubo Chen, Liheng Xu, Kang Liu, Daojian Zeng, and Jun Zhao. 2015. Event extraction via dynamic multi-pooling convolutional neural networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 167–176.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Linguistic Data Consortium et al. 2005. Ace (automatic content extraction) english annotation guidelines for events. version 5.4. 3. *ACE*.
- Agata Cybulska and Piek Vossen. 2011. Historical event extraction from text. In *Proceedings of the 5th ACL-HLT Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, pages 39–43.
- Dipanjan Das, Desai Chen, André FT Martins, Nathan Schneider, and Noah A Smith. 2014. Frame-semantic parsing. *Computational linguistics*, 40(1):9–56.
- Jens A de Bruijn, Hans de Moel, Brenden Jongman, Marleen C de Rooter, Jurjen Wagemaker, and Jeroen CJH Aerts. 2019. A global database of historic and real-time flood events based on social media. *Scientific data*, 6(1):311.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for*

- Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- George R Doddington, Alexis Mitchell, Mark A Przybocki, Lance A Ramshaw, Stephanie M Strassel, and Ralph M Weischedel. 2004. The automatic content extraction (ace) program-tasks, data, and evaluation. In *Lrec*, volume 2, page 1. Lisbon.
- Joe Ellis, Jeremy Getman, Dana Fore, Neil Kuster, Zhiyi Song, Ann Bies, and Stephanie M Strassel. 2015. Overview of linguistic resources for the tac kbp 2015 evaluations: Methodologies and results. In *Tac*.
- Charles J Fillmore and Collin F Baker. 2001. Frame semantics for text understanding. In *Proceedings of WordNet and Other Lexical Resources Workshop, NAACL*, volume 6.
- Charles J Fillmore et al. 1976. Frame semantics and the nature of language. In *Annals of the New York Academy of Sciences: Conference on the origin and development of language and speech*, volume 280, pages 20–32. New York.
- Kung-Hsiang Huang, Mu Yang, and Nanyun Peng. 2020. Biomedical event extraction with hierarchical knowledge graphs. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1277–1285.
- Elisabetta Jezeq, Bernardo Magnini, Anna Feltracco, Alessia Bianchini, and Octavian Popescu. 2014. T-pas: A resource of corpus-derived types predicate-argument structures for linguistic analysis and semantic processing. In *Proceedings of LREC*, pages 890–895.
- Heng Ji, Joel Nothman, H Trang Dang, and Sydney Informatics Hub. 2016. Overview of tac-kbp2016 trilingual edl and its impact on end-to-end cold-start kbp. *Proceedings of TAC*.
- Christian Kay, Jane Roberts, Michael Samuels, and Irené Wotherspoon. 2009. *Historical thesaurus of the Oxford English dictionary*. Oxford University Press.
- Viet Dac Lai, Minh Van Nguyen, Heidi Kaufman, and Thien Huu Nguyen. 2021. [Event extraction from historical texts: A new dataset for black rebellions](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2390–2400, Online. Association for Computational Linguistics.
- Alessandro Lenci, Martina Johnson, and Gabriella Lapesa. 2010. Building an italian framenet through semi-automatic corpus analysis. In *LREC*.
- Alessandro Lenci, Gabriella Lapesa, and Giulia Bonansinga. 2012a. Lexit: A computational resource on italian argument structure. In *LREC*, pages 3712–3718.
- Alessandro Lenci, Simonetta Montemagni, Giulia Venturi, and Maria Grazia Cutrulla. 2012b. Enriching the isst-tanl corpus with semantic frames. In *LREC*, pages 3719–3726.
- Diya Li, Lifu Huang, Heng Ji, and Jiawei Han. 2019. Biomedical event extraction based on knowledge-driven tree-lstm. 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 1421–1430.
- Shasha Liao and Ralph Grishman. 2010. Using document level cross-event inference to improve event extraction. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 789–797. Association for Computational Linguistics.
- Yaojie Lu, Hongyu Lin, Jin Xu, Xianpei Han, Jialong Tang, Annan Li, Le Sun, Meng Liao, and Shaoyi Chen. 2021. Text2event: Controllable sequence-to-structure generation for end-to-end event extraction. 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 2795–2806.
- Stefano Menini, Teresa Paccosi, Serra Sinem Tekiroğlu, and Sara Tonelli. 2023. Scent mining: Extracting olfactory events, smell sources and qualities. In *Proceedings of the 7th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 135–140.
- Gosse Minnema, Sara Gemelli, Chiara Zanchi, Tommaso Caselli, Malvina Nissim, et al. 2022. Sociofillmore: A tool for discovering perspectives. In *The 60th Annual Meeting of the Association for Computational Linguistics Proceedings of System Demonstrations*, pages 240–250. Association for Computational Linguistics.
- 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 *Italian Conference on Computational Linguistics 2021: CLiC-it 2021*. CEUR Workshop Proceedings (CEUR-WS.org).
- Teruko Mitamura, Yukari Yamakawa, Susan Holm, Zhiyi Song, Ann Bies, Seth Kulick, and Stephanie Strassel. 2015. Event nugget annotation: Processes and issues. In *Proceedings of the The 3rd Workshop on EVENTS: Definition, Detection, Coreference, and Representation*, pages 66–76.
- Thien Huu Nguyen, Kyunghyun Cho, and Ralph Grishman. 2016. Joint event extraction via recurrent neural networks. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 300–309.

- Trung Minh Nguyen and Thien Huu Nguyen. 2019. One for all: Neural joint modeling of entities and events. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6851–6858.
- Walker Orr, Prasad Tadepalli, and Xiaoli Fern. 2018. [Event detection with neural networks: A rigorous empirical evaluation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 999–1004, Brussels, Belgium. Association for Computational Linguistics.
- Alessio Palmero Aprosio and Giovanni Moretti. 2018. Tint 2.0: an all-inclusive suite for nlp in italian. In *Proceedings of the Fifth Italian Conference on Computational Linguistics (CLiC-it 2018)*.
- Giovanni Paolini, Ben Athiwaratkun, Jason Krone, Jie Ma, Alessandro Achille, Rishita Anubhai, Cicero Nogueira dos Santos, Bing Xiang, and Stefano Soatto. 2021. Structured prediction as translation between augmented natural languages. *arXiv preprint arXiv:2101.05779*.
- Alan Ramponi, Rob van der Goot, Rosario Lombardo, and Barbara Plank. 2020. Biomedical event extraction as sequence labeling. In *Proceedings of the 2020 conference on empirical methods in natural language processing (emnlp)*, pages 5357–5367.
- Marco Rovera, Federico Nanni, and Simone Paolo Ponzetto. 2021. Event-based access to historical italian war memoirs. *Journal on Computing and Cultural Heritage (JOCCH)*, 14(1):1–23.
- Josef Ruppenhofer, Michael Ellsworth, Myriam Schwarzer-Petruck, Christopher R Johnson, and Jan Scheffczyk. 2006. Framenet ii: Extended theory and practice.
- Josef Ruppenhofer, Michael Ellsworth, Myriam Schwarzer-Petruck, Christopher R Johnson, and Jan Scheffczyk. 2016. Framenet ii: Extended theory and practice. Technical report, International Computer Science Institute.
- Roser Saurí, Jessica Littman, Bob Knippen, Robert Gaizauskas, Andrea Setzer, and James Pustejovsky. 2006. Timeml annotation guidelines. *Version*, 1(1):31.
- Roxane Segers, Marieke Van Erp, Lourens Van Der Meij, Lora Aroyo, Jacco van Ossenbruggen, Guus Schreiber, Bob Wielinga, Johan Oomen, and Geertje Jacobs. 2011. Hacking history via event extraction. In *Proceedings of the sixth international conference on Knowledge capture*, pages 161–162.
- Matthew Sims, Jong Ho Park, and David Bamman. 2019. [Literary event detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3623–3634, Florence, Italy. Association for Computational Linguistics.
- Rachele Sprugnoli and Sara Tonelli. 2017. One, no one and one hundred thousand events: Defining and processing events in an inter-disciplinary perspective. *Natural Language Engineering*, 23(4):485–506.
- Rachele Sprugnoli and Sara Tonelli. 2019. Novel event detection and classification for historical texts. *Computational Linguistics*, 45(2):229–265.
- Swabha Swayamdipta, Sam Thomson, Chris Dyer, and Noah A Smith. 2017. Frame-semantic parsing with softmax-margin segmental rnns and a syntactic scaffold. *arXiv preprint arXiv:1706.09528*.
- Swabha Swayamdipta, Sam Thomson, Kenton Lee, Luke Zettlemoyer, Chris Dyer, and Noah A Smith. 2018. Syntactic scaffolds for semantic structures. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3772–3782.
- Sara Tonelli and Emanuele Pianta. 2008. Frame information transfer from english to italian. In *6th International Conference on Language Resources and Evaluation (LREC 2008)*.
- Sara Tonelli, Daniele Pighin, Claudio Giuliano, and Emanuele Pianta. 2009. Semi-automatic development of framenet for italian. In *Proceedings of the FrameNet Workshop and Masterclass, Milano, Italy*.
- Rob van der Goot, Ahmet Üstün, Alan Ramponi, Ibrahim Sharaf, and Barbara Plank. 2021. [Massive choice, ample tasks \(MaChAmp\): A toolkit for multi-task learning in NLP](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 176–197, Online. Association for Computational Linguistics.
- Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. Ace 2005 multilingual training corpus. *Linguistic Data Consortium, Philadelphia*, 57:45.

A Examples of annotation

While the full documentation of EvN-ITA, including annotation guidelines, frame-based descriptions and examples is being released along with the resource, in this section we provide more details about the annotation process.

As mentioned in Section 4, in EvN-ITA target parts-of-speech are nouns, verbs and multiword expressions. In real-world data, however, beside events expressed in positive, factual form, there are often cases that raise exceptions. In EvN-ITA, event mentions are annotated regardless of their factuality value, which means that also *negated*, *abstract*, *hypothetical* event mentions must be annotated, as well as those introduced by modals verbs. Conversely, we do not annotate as event mention those lexical units that are used with *metaphorical meaning* or in the form of *rethoric expression*. In the following, we provide some examples¹⁵:

1. *negated events*

[La Federazione russa, l'unico legale stato successore dell'Unione Sovietica PAYER], non ha mai riconosciuto le deportazioni degli estoni come un crimine e non ha [pagato PAY] [nessuna riparazione MONEY] [agli stati coinvolti BENEFICIARY].

La persistente segregazione razziale negli Stati Uniti d'America in tutto il profondo sud significò che [la maggior parte degli afroamericani MEMBER] non poteva [fare parte MEMBERSHIP] [dei Grand Jury GROUP] i quali – totalmente composti da bianchi – continuarono ad emanare verdetti discriminatori e palesemente ingiusti.

[Clary SPEAKER] però non può [ribattere REPLY] perché il suo cellulare squilla:

2. *hypothetical/possible events*

Sull'orlo di una [guerra WAR], la Russia comunicò riluttante a

Berlino e Vienna il suo consenso e, abbandonata a sé stessa, il 31 marzo, anche la Serbia si arrese.

Nel gennaio 2008 hanno iniziato a rincorrersi notizie, via via sempre più insistenti e accreditate, che [ad Albano Laziale PLACE], [in prossimità della discarica sita in località "Roncigliano" della frazione di Cecchina RELATIVE_LOCATION], sarebbe stato [realizzato BUILDING] [un inceneritore CREATED_ENTITY] per smaltire i rifiuti, in vista dell'imminente chiusura della discarica di "Roncigliano" e di quella di Malagrotta a Roma.

Il 22 novembre 1961 la polizia perquisì, senza risultato, il suo appartamento in cerca di [una pistola con cui MEANS] [Pasolini] avrebbe [rapinato ROBBERY], [il 18 sera TIME], [un distributore di benzina SOURCE] [di San Felice Circeo PLACE].

La bomba non esplose, altrimenti [la detonazione CAUSE] avrebbe potuto effettivamente [distruggere DESTROYING] [la nave PATIENT].

Il riluttante generale Raffaele Cadorna, per evitare che [Mussolini CAPTIVE] [cadesse nelle mani TAKING_CAPTIVE] [degli Alleati CAPTOR], rilasciò il salvacondotto necessario;

3. *events introduced by modals*

Verosimilmente, si può presumere che [egli MOVER] dovette [allontanarsi QUITTING_A_PLACE] [da Augsburg SOURCE] [in quanto cattolico EXPLANATION], dal momento che, con quanto sancito dalla pace di Augusta, era in vigore il principio del "cuius regio, eius religio".

Se l'indagine dimostrerà che sono stati commessi crimini di

¹⁵For the sake of readability, examples presented in this appendix are annotated with a single layer. Please see the documentation for more data samples.

guerra, [i responsabili DEFENDANT] dovranno essere trovati e [processati TRIAL] [conformemente alle norme in vigore BINDING_PRINCIPLE].

Conversely, we do *not* annotate as event mention occurrences used with *metaphorical meaning* or in the form of *rethoric expression*:

Questa è [la domanda Ø] che Malone, Russo e Montague hanno cominciato a [porsi Ø] una volta esaurita la spinta idealistica dei primi anni di lavoro e la loro risposta è stata:

Signor Presidente, il risultato delle elezioni in Israele ha [fornito la risposta Ø] della popolazione israeliana.

Il visitatore / studioso poteva [intraprendere Ø] così [un viaggio Ø] dal microcosmo (la chimica), attraverso gli elementi primi della natura, al macrocosmo (l'astronomia) nel torrino che concludeva il percorso.

Sostiene inoltre che, nonostante i problemi della filosofia della scienza e della ragione in generale, le "questioni morali" avranno [risposte Ø] oggettivamente giuste e sbagliate suffragate da fatti empirici su ciò che induce la gente a star bene e prosperare.

Finally, as mentioned in Section 4, in order to increase robustness, EvN-ITA contains many negative examples. Given a lexical unit, a negative example is an occurrence of the lexical unit which denotes a meaning not covered by the current schema. Negative examples are meant to improve the classifier's ability to work in an open-world setting and to generalize to extrinsic/unseen data.

Lexical unit: *istituire*

Positive example:

La richiesta venne accolta e il papa diede l'autorizzazione a [istituire CREATE_SOCIAL_ENTITY] [in Inghilterra PLACE] [un tribunale ecclesiastico CREATED_ENTITY] [per esaminare attentamente il caso PURPOSE],

Negative example:

venne così [istituito Ø], nel 46 a.C., il calendario giuliano.

B Association between event frames

In this section we present numerical evidence of association between events, mentioned in Section 5.2. As stated above, patterns of association between frames can be identified by computing their co-occurrence. We choose 5 event frames and list the first 5 most related event frames and the 5 most unrelated frames, using Pointwise Mutual Information (PMI).

Target: INVEST	
Related frames	PMI
BUY	2.87
GROWTH ON A SCALE	2.37
BUILDING	2.22
SELL	2.20
MANUFACTURING	1.79
...	
COMMUNICATION	-0.87
DEATH	-0.89
CONQUERING	-0.93
STATEMENT	-1.04
ATTACK	-1.38

Table 7: Correlation with the INVEST frame.

Target: ARRIVING	
Related frames	PMI
DEPARTING	1.94
ENCOUNTER	1.86
MOVE AWAY	1.82
REMAIN IN PLACE	1.78
MOTION DOWNWARDS	1.75
...	
BEING MARRIED	-1.55
DECREASE ON A SCALE	-1.58
ACQUITTAL	-1.59
EARTHQUAKE	-1.62
TAKING SIDES	-1.79

Table 8: Correlation with the ARRIVING frame.

Target: CREATE ARTWORK	
Related frames	PMI
PERFORMING ARTS	2.18
ASSIGN TASK	1.89
TEMPORAL ORIGIN	1.73
PUBLISHING	1.72
BEING LOCATED	1.67
...	
PURPOSE	-1.49
ROBBERY	-1.50
APPOINTING	-1.51
PROCESS END	-1.59
WAR	-1.70

Table 9: Correlation with the CREATE ARTWORK frame.

Target: TRIAL	
Related frames	PMI
ACQUITTAL	3.81
SENTENCING	3.59
VERDICT	3.12
ACCUSE	2.99
EXECUTION	2.80
...	
FLEEING	-1.29
BEING LOCATED	-1.36
CREATE ARTWORK	-1.37
BEAT OPPONENT	-1.39
AGREEMENT	-1.41

Table 10: Correlation with the TRIAL frame.

Target: COMMUNICATION	
Related frames	PMI
CONTACTING	2.64
QUESTIONING	2.09
ENCOUNTER	2.04
AWARENESS	2.02
GIVING	1.88
...	
EVENT ORDERING	-1.22
COUNTERATTACK	-1.24
APPOINTING ELECTION	-1.25
SUPPRESSING	-1.25
TAKE PLACE OF	-1.53

Table 11: Correlation with the COMMUNICATION frame.