

Towards an Italian Corpus for Implicit Object Completion

Agnese Daffara¹ and Elisabetta Jezek¹

¹ University of Pavia, Corso Strada Nuova, 65, Pavia, Italy

Abstract

This study centers on the creation of an Italian corpus designed for the task of Implicit Object Completion. In this corpus, every sentence contains a token [MASK] denoting the position of the Object's head, along with the annotation of a Gold Standard filler word. The completion of the Object is conceived as a masked word task, theoretically executable by a BERT-based transformer model. In the next phase of the project, this task will be applied to a range of Italian language models, and their performance will be assessed. Overall, this project seeks to offer insights into the capabilities and constraints of such models in successfully completing Implicit Objects within various contexts.

Keywords

BERT, Implicit Object, masked word

1. Introduction

When coming across the verb-argument structure of a sentence, individuals have the cognitive ability to comprehend its meaning by forming a semantic representation of the situation in their minds. Even in cases where one argument is implicit, they are still capable of understanding the overall sense, thanks to the verb's inherent lexical meaning and the neighbouring words. The Distributional Hypothesis, as proposed by Harris (1954) and Firth (1951), suggests that it is possible to infer the meaning of a word purely on the basis of the context.

In the field of Natural Language Understanding, Artificial Intelligence must replicate this ability in order to reconstruct the scenario of the event, specifically identifying its semantic participants. Given the requirement for a computational model to fill in the missing information, we propose that this task can be conceived and construed as a masked word completion task, for which transformer-based technologies such as BERT (Devlin et al., 2019) have proven to be the most suitable.

This paper focuses on building an Italian corpus for this specific purpose while hinting at the same time at the forthcoming work of evaluation.

The corpus centers on verbs that exhibit an Optional Object, i.e. an Object that can be Implicit or Explicit. The ontological set of verbs on which the corpus is constructed is presented in section 4. Following these verbs' ambivalent possibility of expressing or implying the Argument, the corpus is divided into two datasets: on one side, an IMPLICIT dataset of sentences with Implicit Objects; on the other side, with a contrastive role, an EXPLICIT dataset of sentences containing Explicit Objects.

Our decision to create two different datasets is motivated by the idea of observing the differences in

the performance of the models: do they perform better when the original Object is Implicit? This issue is grounded in the findings of a prior guide experiment conducted by Ye et al. (2020); according to their results, the model's performance would be notably improved when fine-tuned on an IMPLICIT dataset, because of the greater richness of contextual information available. We want to investigate if such observations can be generalized to our experiment.

Regarding the annotation of the masked word, the two datasets are treated differently. In the IMPLICIT dataset, we inserted two [MASK] tokens right after the verb or the adverb, and allowed the model to generate either a single Noun, as in Table 1, sentence 1., or, if not found, a Noun Phrase (NP) consisting of a Determiner plus a Noun. Further explanation of this possibility can be found in section 5. Furthermore, a Gold Standard (GS) Noun representing the optimal completion of the Object's head position was annotated aside each sentence, together with the type of omission (see section 2 for theoretical references).

In the EXPLICIT dataset, on the other hand, we removed the Explicit Object's nominal head, consisting of one word, and we annotated it as a GS. Two examples of annotation for sentences 1. and 2., belonging respectively to the IMPLICIT and the EXPLICIT dataset, are provided in Table 1.

1. Da quel 26 dicembre non vuole più bere [MASK][MASK] né lavarsi. 'Since that December 26, they no longer want to drink [MASK][MASK], nor wash themselves.'
2. Infilo la fetta velocemente nel sacchetto delle fragole e tiro un sospiro, bevo un [MASK] di caffè. 'I slide the slice quickly into the strawberry bag and let out a sigh, I drink a [MASK] of coffee.'

Table 1
Example of annotation on sentences 1. and 2.

CLiC-it 2023: 9th Italian Conference on Computational Linguistics, Nov 30 – Dec 02, 2023, Venice, Italy.

✉ agnese.daffara01@universitadipavia.it (A. Daffara);

jezek@unipv.it (E. Jezek)



Copyright © 2023 for this paper by its authors. The use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).
CEUR Workshop Proceedings (CEUR-WS.org)

dataset	sent	GS_obj_head	omission_type
IMP.	1.	acqua 'water'	LD
EXP.	2.	sorso 'sip'	-

As referenced above, the corpus will allow us to undertake a selection of BERT Italian models and systematically evaluate their performance on the task of Implicit Object Completion, which we define in section 6.

We firmly believe that an annotated Italian dataset containing masked Optional Objects, their categorisation and their corresponding Gold Standard completions, as well as the subsequent experiment and evaluation, will greatly contribute to research endeavors in the field of NLP.

2. Related Work

Previous computational works have primarily focused on the task of Implicit Argument Detection rather than on the mere completion of a masked Object. SemEval 2010 task 10 (Ruppenhofer et al., 2009) introduced a variety of approaches aimed at detecting the semantic participants to the event, specifically identifying the Null Instantiations of the Arguments. The term Null Instantiation was introduced by Fillmore (1982) within the theory of Frame Semantics. In fact, most of the proposals relied on this theoretical background, adopting as a starting point the Framenet dataset and annotation.

While the rise of transformer-based models has brought significant improvement for this task, as shown, for example, by Zhang et al. (2020), it still remains an interesting and challenging issue.

For what concerns the Italian language, we identified a potential gap in the literature on the computational detection and processing of Implicit Arguments, probably due to the lack of annotated corpora designed for this task. It is, therefore, of utmost importance to investigate this topic and create new computational suitable resources.

Significant progress has been made in the training of BERT-based Italian models, including ALBERTo (Polignano et al., 2019), UmBERTo (Parisi et al., 2020), GilBERTo (Ravasio and Di Perna, 2020), and the distilled Italian version of DistilBERT called BERTino (Muffo and Bertino, 2020). Thanks to the availability of this generation of open-source BERT models, the masked word task has been applied to a variety of different linguistic and cognitive topics, such as the study of Agentivity and Telicity (Lombardi and Lenci, 2021) or connectives (Albertin et al., 2021). In particular, our study consistently builds upon the prior application of the masked word Task to the semantic topic of Logical Metonymy by Ye et al. (2020).

The existing linguistic literature has extensively explored the concept of Implicit Argument and the phenomenon of Argument omission in Italian. Notably, Cennamo (2017) proposed a meticulous comparative analysis of the parameters involved in this process. In our study, we adopt the notion of "Defaulting", first introduced by Pustejovsky (1995) and further refined by Jezek (2018). Following

Fillmore's distinction between Definite and Indefinite Null Instantiation, we delineate Pragmatic Defaulting (PD) as the omission of the Object based on contextual cues and Lexical Defaulting (LD) as the omission of the Object licenced by the core meaning of the verb. Overall, it is undeniable that both the contextual cues and the semantics encoded into the verb contribute to the possibility of implying and reconstructing an Argument, and we believe it is necessary to consider this difference when studying Implicit Objects.

3. Data Preparation

As a first step towards the corpus preparation, we established a set of 30 verbs that allow for their Object to remain implied. We refer to this set with the term 'ontology' since it contains the basic verbal structures of reference for the building of the corpus. Our selection of such verbal structures draws upon the resource T-PAS (Jezek et al., 2014), a repository of Typed Predicate-Argument Structures (T-PAS) which was developed at the University of Pavia in collaboration with the Bruno Kessler Foundation in Trento (I) and the Masaryk University in Brno (CZ) by adopting a corpus-driven methodology.

Each pattern in T-PAS corresponds to a distinct contextual meaning of the verb (Predicate), plus the list of all the possible semantic participants to the event associated with that specific meaning (Arguments). Notably, T-PAS not only captures information concerning the syntactic structure but also provides insights into the semantic types of the Arguments. This resource is a valuable foundation for our data collection, as it also annotates (in round brackets) the potential for exhibiting an Implicit Argument for each structure. An example of three patterns displayed on the online T-PAS website for the verb 'bere' ('drink') (including a metonymic use) is given in Figure 1.

1	[Animate] bere [(Beverage (birra caffè tè bibita bevanda aperitivo cocktail liquore vino acqua latte grappino birretta spritz mojito birrozza tisana cappuccino cioccolata whisky vodka rum rum rhum cognac pozione elisir sangue liquido acqua))]
1.m	[Human] bere [(Container (bicchiere bottiglia)) [Business Enterprise = Producer] [Quantity (sorso goccio)]] [Human] ingerisce, assume [(Container (bicchiere bottiglia)) (che contiene [Beverage]), [Business Enterprise = Producer] (che produce [Beverage]), [Quantity (sorso goccio)) (che è una porzione di [Beverage])]
2	[Human] bere [Human] ingerisce, assume una certa quantità di bevande alcoliche

Figure 1. Example of three patterns for the verb 'bere' ('drink') in T-PAS.

From the comprehensive dataset of patterns available, we first identified the ones containing one or more Optional Arguments, using a simple RegEx match search to detect round brackets. Afterwards, we isolated 'fundamental' verbs (verbi fondamentali) according to the Nuovo Vocabolario di Base della Lingua Italiana (NVdB) (De Mauro, 2016). These particular verbs were chosen due to their presence in 90% of Italian texts, making them a suitable representative set for constructing an ontology.

We then conducted a cleansing process that excluded causatives, passives, and idiomatic

expressions, as well as other multiword expressions or subpatterns with relatively infrequent occurrences and less common meanings. We constantly consulted the online version of the NVdB to decide whether a structure was fundamental or not. This cleansing process finally yielded a comprehensive list of 324 patterns with an Optional Argument, spanning across 213 distinct verb types.

We finally proceeded to further narrow down our focus, isolating the structures with an Optional Object. The final ontology comprehends 30 different verbs, corresponding to 60 T-PAS patterns and over 50 possible Object’s Semantic Types, which represents a consistent variety. The detailed list is provided in Appendix A, whereas a summary of the quantities of verbs and patterns contained in T-PAS can be found in Table 2.

Table 2
Verb types and patterns in T-PAS

List	Verb types	Patterns
Tot.	1000	5392
Fundamental	492	
Fundamental with	213	324
Optional Arg. (clean)		
Fundamental with	30	60
Optional Obj. (clean)		

4. Data Collection

After the assessment of the ontology, our attention turned to collecting the sentences. The resource of reference is the T-PAS dataset, comprising 252,943 Manually Annotated Corpus Instances. All the sentences were selected from the It-Wac reduced corpus (Baroni and Kilgarriff, 2006) and annotated with the corresponding T-PAS number, denoting the specific semantic pattern being used in that sentence. For example, as shown in Figure 1, when the verb ‘bere’ (‘drink’) was used with a metonymic Object, like ‘sorso’ (‘sip’), it was tagged with T-PAS number 1m, while, when it was found without an Object and implying an alcoholic drink, it was tagged with T-PAS number 2.

After isolating the T-PAS structures contained in our ontology, we proceeded to manually select the sentences for the corpus. We removed those with a Noun as an Object and preferred those with a linear order (the Object following the Noun). In our pursuit of a more extensive and diverse dataset, especially concerning the variety of Objects, we also conducted searches in the whole It-Wac reduced corpus through the Sketch Engine online platform (Kilgarriff et al., 2014). Eventually, 40 sentences were selected for each verb. The resulting 1200 sentences are divided into the two datasets, each containing 600 sentences, as illustrated in Table 3.

Table 3
Structure of the corpus

	IMP. dataset	EXP. dataset	Whole corpus
Verbs	30	30	30

Sentences for verb	20	20	40
Tot. sentences	600	600	1200

5. Annotation

The annotation process was handled differently for the IMPLICIT and the EXPLICIT dataset.

In the IMPLICIT dataset, the token [MASK] was manually inserted after the verb or the adverbial modifier, in order to signal the position to be filled by the model. However, we’ve observed that when the model encounters only one masked word, it tends to generate either mass Nouns or plural Nouns due to the lack of a Determiner. This presents a significant limitation in the evaluation process. The chosen approach involves the following steps: 1. Initially, annotate two [MASK] tokens to indicate the positions of the Determiner and the Noun. 2. Instruct the model to first look for a Noun for the second position. (3) If a Noun is not generated, proceed to search for a Determiner. (4) Generate a new sentence containing that Determiner, which will later be filled with a Noun.

As a following step in the annotation of the IMPLICIT dataset, the GS word constituting the optimal Object’s head was manually inserted on the basis of the pragmatic context and the strength of the possible collocations. This value, achieved using the LogDice metric, can be obtained by querying Sketch Engine on the ItTenTen20 Italian corpus¹, as shown in Figure 2.

The last step of the IMPLICIT dataset’s annotation regards the type of omission. As already mentioned, we adopted the classification proposed by Jezek (2018) following Pustejovsky (1995) between Pragmatic and Lexical Defaulting. This categorization serves as a valuable tool during the final evaluation of the model, enabling an assessment of its performance across different kinds of omission.

objects of "bere"		
acqua bere acqua	42,723	10.8 ...
caffè bere un caffè	18,188	10.7 ...
vino bere vino	13,413	9.9 ...

Figure 2: Collocation scores of verb ‘bere’ (‘drink’) plus Object on ItTenTen20.

For what concerns the EXPLICIT dataset, the Object’s head was manually detected, removed and replaced by the token [MASK]. Subsequently, it was annotated aside the sentence. Note that by removing just a single word, these sentences retain their rich syntactic context, displaying the modifiers of the removed word. Such cues may improve the models’ ability to detect the original filler. As the EXPLICIT dataset primarily has a contrastive function, we anticipate that comparing results from both datasets will help determine whether the model’s output is closer to the original when it receives a significant amount of

¹ <https://www.sketchengine.eu/ittenten-italian-corpus/>

syntactic information or, conversely, when the context is semantically richer, as seen in Implicit Object sentences.

6. Task Definition

We define Implicit Object Completion as the task of substituting the masked Object in a sentence, previously marked with the token [MASK], with the most appropriate word or filler. When tested on each sentence, the transformer model is expected to produce the word that best fits the context of the sentence.

However, alternative outputs are possible, potentially encompassing other Parts of Speech. As an example, we employed the online demo of bert-base-italian-cased, made accessible by the MDZ Digital Library team (dbmdz) at the Bavarian State Library on Hugging Face². The model generated the most probable candidates for sentence 1.. Predictably, the first output was the punctuation sign "," and the expected Nouns were found in lower positions, as depicted in Figure 3. In order to mitigate this issue and ensure more accurate results, during the model's interrogation, we implemented a two-step filter that isolates Nouns. In particular, we exclusively considered the Noun with the highest probability score.

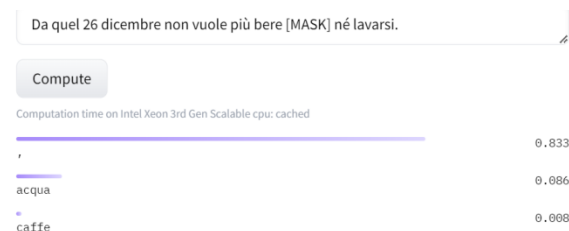


Figure 3: Outputs of the model bert-base-italian-cased for sentence 1.

7. Evaluation

An issue in the design of the task is the possibility of getting a synonym or a word that is only partially correct or doesn't perfectly align with the Gold Standard.

The Theory of Prototypes, as proposed firstly by Rosch (1973), posits that within a semantic category, certain members are more representative of the category's core meaning. In contrast, less central members demonstrate greater variability and may deviate further from the core concept. By taking into consideration both the Theory of Prototypes and the Distributional Hypothesis (cited in section 1), during the evaluation phase, we will systematically calculate the similarity score (sim) between the output word and the Gold Standard completion, corresponding to the cosine between the two word vectors. This value will be obtained by running the Python library SpaCy

² <https://huggingface.co/dbmdz/bert-base-italian-cased>

(Honnibal and Montani, 2017) on the Italian model it_core_news_lg, a large language model with a size of 541 MB³.

An example of output of the model bert-base-italian-xxl-cased (bert-base-xxl), the bigger version of bert-base-italian-cased from dbmdz, and its relative annotation for sentence 1. and 2. is shown in Table 4.

Table 4
Example of sentence 1. and 2. bert-base-xxl outputs

data set	sent	GS_obj_head	bert-base-xxl	sim
IMP.	1.	acqua	'water' acqua 'water'	1.0
EXP.	2.	sorso	'sip' bicchiere 'cup'	0.62

With these annotation parameters, we aim to extend our linguistic analysis beyond the model's ability to complete the cloze test by providing the right word. Instead, we will also investigate the model's capability to effectively cluster words within the same domain.

8. Discussion and Results

This paper discusses the ongoing construction of a corpus specifically tailored for the task of Implicit Object Completion. This resource contains sentences exhibiting both Implicit and Explicit Objects, thus enabling the assessment of two distinct datasets that will be treated separately.

In the IMPLICIT dataset, the position for the Noun/NP is signaled by manually inserting two tokens [MASK] right after the verb or the adverb. The GS Object's head is manually added, considering both the context of the sentence and the general strength of the verb-Object collocation, which can be quantified through the typicality score on the ItTenTen20 corpus (Jakubíček et al., 2013) using Sketch Engine. Additionally, in the case of Implicit Objects, we provide information about the type of omission, which may either depend on the contextual cues (Pragmatic Defaulting) or the lexical verbal root (Lexical Defaulting). On the other hand, the annotations within the EXPLICIT dataset include the manual identification of the Object's nominal head, which is substituted with the token [MASK] and annotated aside the sentence.

The forthcoming second phase of our project involves an in-depth analysis of the outputs generated by a selection of the primary BERT Italian models. As a metric for the evaluation, we will adopt cosine similarity. This value measures the similarity between the output word provided by the model and the GS word, thus measuring the ability of the model to generate a filler which is semantically close to the original. As an example of a comparison between two models, consider the results of bert-base-xxl and umberto-commoncrawl-cased-v1 (UmBERTo), on sentence 2., which are reported in Table 5.

³ https://github.com/explosion/spacy-models/releases/tag/it_core_news_lg-3.7.0

Table 5
Comparison of two models' outputs on sentences 2.
from the EXPLICIT dataset

sent	GS_obj_head	bert-base-xxl	sim	UmBERTo	sim
2.	sorso 'sip'	bicchie re 'cup'	0.62	cucchiaino 'teaspoon'	0.4

Bert-base-xxl returns a slightly higher score, as the vectors of 'bicchiere' ('cup') and 'sorso' ('sip') have a higher cosine similarity than those of 'cucchiaino' ('teaspoon') and 'sorso' ('sip'). Although both the models fail to understand the exact word and categorize the filler as a [CONTAINER] rather than a [QUANTITY], both the results are satisfactory and plausible. More results for the EXPLICIT dataset can be found in Appendix B.

In conclusion, we expect the results to raise a number of theoretical questions and possible investigations. By conducting this analyses, we will compare the models' performance on a novel topic and investigate their ability to identify the semantic category of the Objects, while effectively clustering words within the same domain. In addition, the annotation of the type of omission will allow further insights on the importance of the context in reconstructing Implicit Objects.

Acknowledgements

The authors gratefully acknowledge the contributions of the anonymous reviewer for CLiC-it 2023. Their insights and feedback have significantly enhanced the quality of this paper.

Upon completion of the corpus, the complete dataset will be hosted on a public GitHub repository in accordance with the FAIR principles.

References

- G. Albertin, A. Miaschi, D. Brunato, On the Role of Textual Connectives in Sentence Comprehension: A New Dataset for Italian, in: Proceedings of the Eighth Italian Conference on Computational Linguistics CLiC-it 2021, Milano, 2021.
- M. Baroni and A. Kilgarriff, Large Linguistically-Processed Web Corpora for Multiple Languages, Demonstrations (2006) 87-90.
- M. Cennamo, Object omission and the semantics of predicates in Italian in a comparative perspective, In: L. Hellan, A.L. Malchukov, M. Cennamo (eds.), Contrastive Studies in Verbal Valency, Benjamins, Amsterdam, 2017, pp. 251-273.
- T. De Mauro (ed.), Il Nuovo vocabolario di base della lingua italiana, Internazionale, 2016. URL: <https://dizionario.internazionale.it/>
- J. Devlin, M. Chang, K. Lee, K. Toutanova, 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, Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 4171-4186.
- C. J. Fillmore, Frame semantics. In *Linguistics in the Morning Calm*, Hanshin Publishing Co, Seoul, 1982, pp. 111-137.
- J.R. Firth, Modes of Meaning, *Papers in Linguistics* 1934-1951 (1951) 190-215.
- Z.S. Harris, Distributional Structure, *word* 10, 2-3 (1954) 146-162.
- M. Honnibal, I. Montani, SpaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing, 2017. URL: <https://spacy.io/>
- M. Jakubíček, A. Kilgarriff, V. Kovář, P. Rychlý, V. Suchomel, The TenTen Corpus Family, 2013. URL: <https://www.sketchengine.eu/ittenten-italian-corpus/>
- E. Jezeq, Partecipanti impliciti nella struttura argomentale dei verbi, In: S. Dallabrida, P. Cordin (eds.), *La Grammatica delle Valenze*, Franco Cesati, Firenze, 2018, pp. 55-71.
- E. Jezeq, B. Magnini, A. Feltracco, A. Bianchini, O. Popescu, T-PAS; A resource of Typed Predicate Argument Structures for linguistic analysis and semantic processing, in: Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14), European Language Resources Association (ELRA), Reykjavik, Iceland. 2014, pp. 890-895. URL: <https://tpas.sketchengine.eu/>
- A. Kilgarriff, V. Baisa, J. Bušta, M. Jakubíček, V. Kovář, J. Michelfeit, P. Rychlý, V. Suchomel, The Sketch Engine: ten years on, *Lexicography*, 1 (2014) 7-36 URL: <https://www.sketchengine.eu/>
- A. Lombardi, A. Lenci, Agentività e telicità in GiLBERTO: implicazioni cognitive, in: Proceedings of the Eighth Italian Conference on Computational Linguistics CLiC-it 2021, Milano, 2021.
- M. Muffo, E. Bertino, BERTino: an Italian DistilBERT model, in: Proceedings of the Seventh Italian Conference on Computational Linguistics CLiC-it 2020, Bologna, 2020.
- L. Parisi, S. Francia, P. Magnani, Umberto: An Italian language model trained with whole word masking, 2020. URL: <https://github.com/musixmatchresearch/umberto>
- M. Polignano, P. Basile, M. De Gemmis, G. Semeraro, V. Basile, ALBERTo: Italian BERT Language Understanding Model for NLP Challenging Tasks Based on Tweets, in: Proceedings of the Sixth Italian

Conference on Computational Linguistics CLiC-it
2019, Bari, 2019.

J. Pustejovsky, *The Generative Lexicon*. Cambridge,
MA. MIT Press, 1995.

G. Ravasio, L. Di Perna, *GilBERTo: An Italian
pretrained language model based on RoBERTa*. URL:
<https://github.com/idb-ita/GilBERTo>

E. Rosch, *Cognitive representations of semantic
categories*, *Journal of Experimental Psychology:
General*, 104,3 (1975) 192-233.

J. Ruppenhofer, C. Sporleder, R. Morante, C. Baker, M.
Palmer, *SemEval-2010 Task 10: Linking Events and
Their Participants in Discourse*, in: *Proceedings of the
Workshop on Semantic Evaluations: Recent
Achievements and Future Directions (SEW-2009)*,
Association for Computational Linguistics, Boulder,
Colorado, 2009, pp. 106-111.

B. Ye, J. Tu, E. Jezek, J. Pustejovsky, *Interpreting
Logical Metonymy through Dense Paraphrasing*, in:
*Proceedings of the Annual Meeting of the Cognitive
Science Society*, 44, 2022.

Z. Zhang, X. Kong, Z. Liu, X. Ma, E. Hovy, *A Two-Step
Approach for Implicit Event Argument Detection*, in:
*Proceedings of the 58th Annual Meeting of the
Association for Computational Linguistics*,
Association for Computational Linguistics, 2020, pp.
7479-74.

Appendix A: Ontology of selected verbs and patterns from T-PAS

verb	n.	pattern
ascoltare	1 1m	[Human] ascoltare ([Sound] [Musical Performance {concerto}]) [Human1] ascoltare ([Sound Maker] [Medium] [Musical Composition] [Tv Program] [Event] [Part of Language] [Speech Act] [Narrative] [Human2 = Singer Musician] [Human Group = Band])
attendere	1 1m	[Human] [Institution] attendere ([Event]) (che [Event]) [Human] [Institution] attendere ([Human2] [Vehicle] [Time Point {data}] [Document {visto passaporto}])
bere	1 1m 2	[Animate] bere ([Beverage {birra caffè tè bibita bevanda aperitivo cocktail liquore vino acqua latte grappino birretta spritz mojito birrozza tisana cappuccino cioccolata whisky vodka rum rhum cognac pozione elisir sangue liquido acqua}]) [Human] bere ([Container {bicchiere bottiglia}] [Business Enterprise = Producer] [Quantity {sorso goccio}]) [Human] bere ([Alcool])
cantare	1 2 3	[Human] cantare [Human] cantare [Human] cantare ([Musical Composition {canzone canto inno brano testo salmo}])
chiamare	6 6m	[Human1] chiamare ([Human2] [Institution {polizia}]) [Human] chiamare ([Number] [Device {telefono}] [Location {call center}] [Vehicle {ambulanza}])
combattere	2 3	[Human] combattere ([War {guerra battaglia}]) [Human1] [Human Group1] combatte ([War]) (con contro [Human2] con contro [Human Group2])
condurre	5	[Human] condurre ([TV Program])
consumare	2	[Human] [Human Group] [Machine] [Device] consumare ([Energy] [Gas] [Inanimate])
correre	3	[Human = Runner Pilot] correre ([Competition {maratona palio rally}])
cucinare	1 2	[Human] cucinare ([Food] [Meal {pranzo cena}]) [Human] cucinare
dirigere	2 3	[Human] dirigere ([Musical Performance {concerto}]) [Human = Director] dirigere ([Movie])
disegnare	1 2 3 4	[Human] disegnare ([Image] [Physical Entity]) [Human] disegnare ([Inanimate]) [Human] disegnare ([Document {fumetto comics copertina}]) [Human] disegnare
fumare	1 2	[Human] fumare ([Drug {sigaretta pipa sigaro marijuana}]) [Human] fumare
giocare	3 4	[Human] giocare [Human] [Human Group = Team] giocare ([Competition {partita}] {mano set stagione tempo})
guadagnare	1	[Human] guadagnare ([Money])
guidare	1	[Human] guidare ([Road Vehicle])
leggere	2 3 8	[Human] leggere [Human] leggere ([Document]) [Human1] leggere ([Document]) (a [Human2])

mangiare	1 2 3	[Human] mangiare ([Food {cibo carne pane uovo pizza panino gelato biscotto torta bistecca hamburger salsiccia salame polpetta frutta mela verdura banana riso patata carota formaggio minestra insalata polenta zuppa antipasto spaghetti pasta patatina panettone brioche piadina cornetto focaccia pasticcino pappa pasto biada}}) [Human] mangiare ([Food] {cibo}) [Human] mangiare ([Meal])
ordinare	1 2	[Human] ordinare ([Artifact]) [Human] ordinare ([Food] [Beverage] [Meal])
pagare	4	[Human] pagare ([Abstract Entity {conseguenza debito errore}])
perdere	7	[Human] [Human Group] perdere ([Competition])
pregare	1 2 4	[Human] pregare ([Deity]) [Human1] [Institution] pregare ([Deity]) (per [Human2]) [Human1] pregare ([Human2]) di [Activity]
preoccupare	2	[Anything] preoccupare ([Human])
provare	5	[Human = Artist] [Human Group = Artist] provare ([Artwork])
respirare	1	[Animate] respirare ([Vapor])
scrivere	1 2 4 6 7	[Human] scrivere [Human] scrivere ([Part of Language]) [Human] scrivere ([Document]) [Human] scrivere ([Document]) (a per [Human2]) [Human = Writer] scrivere
servire	6 7	[Human] servire ([Food] [Meal]) (Manner) [Human1 = Waiter] servire ([Food] [Meal]) a [Human2 = Customer]
suonare	3 2 1 5	[Human] suonare ([Musical Instrument]) [Human] suonare [Human = Artist] suonare ([Musical Composition {canzone brano pezzo concerto}] {musica}) [Human] suonare ({il campanello} {il citofono}) (alla {porta})
tirare	3	[Human = Football Player] tirare ([Ball])
vincere	1	[Human] [Human Group] vincere ([Activity {gara competizione festival elezioni}] [War])

Appendix B: Example of results for the EXPLICIT dataset

The following table reports an example of the outputs of two models, italian-BERT-xxl-cased and UmBERTo. The models were tested on the 20 sentences from the EXPLICIT dataset with the verb 'bere' ('drink'). The column 'sent' displays the sentence with the masked word, corresponding to the Object's NP's head. The removed word is shown in the column 'GS_obj_head'. The columns 'bert-base-xxl' and 'UmBERTo' report the outputs of the models. The similarity score between the output word and the GS word is shown in the columns 'sim'.

sent	GS_obj_head	bert-base-xxl	sim	UmBERTo	sim
Tutto meno che restare a guardare la televisione a bere [MASK] e divorare patatine.	[birra]	birra	1	birra	1
Giganti americani del valore di Theodore Dreiser, Ernest Hemingway e Thornton Wilder, quando erano stanchi della routine andavano nei locali a bere [MASK] di whisky e a sentire la grande musica per ricercare la giusta ispirazione.	[litri]	litri	1	fiumi	0,14
Grazie anche per avermi fatto bere l' [MASK] di barbabietola e per avermi fatto svegliare tutte le mattine alle 6 per prendere la pappa reale...	[estratto]	acqua	0,36	acqua	0,36
Bere due [MASK] di latte di soia o mangiare una tazza di tofu è responsabile di livelli ematici di Isoflavoni che possono essere 500 o 1000 volte più elevati dei normali livelli di Estrogeni nelle donne.	[tazze]	bicchieri	0,64	bicchieri	0,64
Da circa 3 settimane Federico ha cominciato a bere [MASK] parzialmente scremato alta qualità e sin dai primi giorni ha mostrato di gradire il nuovo alimento.	[latte]	latte	1	latte	1
La disposizione potrebbe essere utile anche nelle altre stagioni dell'anno, specie se si vieta ai minori di bere [MASK].	[alcool]	alcolici	0,69	alcolici	0,69
Bevi ogni giorno [MASK] in abbondanza è infatti la quinta regola della sana alimentazione, che invita a bere , ma rigorosamente acqua e non altre bevande, frequentemente e in piccole quantità.	[acqua]	acqua	1	acqua	1
Per il pubblico in generale e per i giovani studenti c'è come al solito il padiglione 9 aperto gratuitamente in cui si può fare shopping, bere del [MASK] e rilassarsi oppure informarsi su ciò che accade nella fiera.	[vino]	vino	1	caffè	0,54
Pruneddu, che forse aveva bevuto qualche [MASK] di troppo prima di affacciarsi sulla porta del bar, non si è accorto che il proprietario e le altre persone presenti avevano organizzato una castagnata per stare insieme a bere un po' di vino.	[bicchiere]	bicchiere	1	bicchiere	1
Ma io non bevo [MASK] e gioco a freccette mentre dico parolacce.	[alcolici]	alcolici	1	alcolici	1
Davanti al Castello c'è il Ritz, dove Mordecai e Florence spesso andavano a bere un [MASK].	[drink]	caffè	0,61	caffè	0,61
E poi pensa un po' che GiPo ormai non può più dir niente perchè ha bevuto la [MASK] ed è morto.	[cicuta]	birra	0,33	birra	0,33
Dopo aver recuperato un maglione e bevuto un buon [MASK] al cardamomo, rientriamo in chiesa per l'ora del silenzio.	[caffè]	caffè	1	tè	0,77
Continuai a bere in silenzio il [MASK], mentre il sole che tramontava tingeva di rosso il cielo.	[the]	vino	- 0,16	vino	- 0,16
Gli uruguayani, come gli argentini, bevono moltissimo [MASK], un the fatto con le foglie secche della pianta omonima, sorseggiato da una piccola zucca attraverso una cannuccia di metallo, la bombilla.	[mate]	tè	0,34	caffè	0,32
La vita umana andrebbe rispettata, ma non sentirti mai in colpa di	[sangue]	cibo	0,39	sangue	1

bere [MASK] umano: è una cosa naturale.					
Infilo la fetta velocemente nel sacchetto delle fragole e tiro un sospiro, bevo un [MASK] di caffè.	[orso]	bicchiere	0,62	cucchiaino	0,4
Questi giovani, ci scommetterei, han bevuto [MASK] ed ascoltato musica rock come gli altri coetanei, cosa é scattato ad un certo punto nel loro animo per un totale stravolgimento e per abbracciare un'ideologia perversa?	[coca cola]	birra	0,6	birra	0,6
Un turista che chiede un caffè in tazza, molto lungo e con latte - dice il barman Umberto - dissimula la voglia di bere un [MASK] e pagarlo come caffè ristretto.	[cappuccino]	caffè	0,69	caffè	0,69
Il cubetto grande è molto richiesto soprattutto sul mercato spagnolo; nei locali il consumatore vuole bere un [MASK] in un bicchiere grande (generalmente un tumbler alto) e gradisce che lo stesso gli venga presentato colmo di distillato.	[whisky]	cocktail	0,54	cocktail	0,54