

Refining Implicit Argument Annotation for UCCA

Ruixiang Cui and Daniel Hershcovich

Department of Computer Science
University of Copenhagen
{rc, dh}@di.ku.dk

Abstract

Predicate-argument structure analysis is a central component in meaning representations of text. The fact that some arguments are not explicitly mentioned in a sentence gives rise to ambiguity in language understanding, and renders it difficult for machines to interpret text correctly. However, only few resources represent implicit roles for NLU, and existing studies in NLP only make coarse distinctions between categories of arguments omitted from linguistic form. This paper proposes a typology for fine-grained implicit argument annotation on top of Universal Conceptual Cognitive Annotation’s foundational layer. The proposed implicit argument categorisation is driven by theories of implicit role interpretation and consists of six types: Deictic, Generic, Genre-based, Type-identifiable, Non-specific, and Iterated-set. We exemplify our design by revisiting part of the UCCA EWT corpus, providing a new dataset annotated with the refinement layer, and making a comparative analysis with other schemes.

1 Introduction

Semantic representation frameworks have been a major medium to understanding the nature of languages for NLP. Through these frameworks, researchers have been exploring linguistic phenomena such as quantification (Pustejovsky et al., 2019), coreference (Prange et al., 2019), and word sense (Schneider et al., 2018). However, most efforts were put into studying linguistic complexity superficially, rather than the more latent, implicit omission of arguments in an event. For instance, in the sentence “Just take the money!”, the addressee who should “take the money” is left out. Such omission cannot be recovered directly from the text in the way of gapping or ellipsis, but require a higher level of understanding and inference from the context. Traditional studies approach argument omission from different aspects, namely syntactically, semantically, or pragmatically. The interpretation of implicit roles varies to a great extent, from phonological deleted role during production (Perlmutter, 1968; Mittwoch, 1971; Pérez-Leroux et al., 2018) to timecourse reference omission from a psycholinguistic aspect (Garrod and Terras, 2000). However, few studies have explored the implicit role phenomenon in NLP.

In this paper, we propose a fine-grained cross-linguistically applicable implicit argument annotation typology as a refinement for Universal Conceptual Cognitive Annotation (Abend and Rappoport, 2013, UCCA) categories. The typology follows UCCA’s design concept, focusing on the semantic notion of Scene rather than linguistic form phenomena. The proposed implicit argument set contains six categories: Deictic, Generic, Genre-based, Type-identifiable, Non-specific, and Iterated-set. We refine the existing UCCA relation labels and add information to them, while keeping all categories from the underlying annotation. Our studies move UCCA, a semantic representation framework, to have less syntactic definitions of implicit arguments.

Based on the proposed typology, we conduct a pilot annotation study, including revisit and refinement of the UCCA EWT dataset,¹ and subsequently make a comparative analysis with the only other existing fine-grained implicit role annotation scheme, Fine-grained Annotations of Referential Interpretation Types (O’Gorman, 2019, FiGref).

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

¹https://github.com/ruixiangcui/UCCA-Refined-Implicit-EWT_English

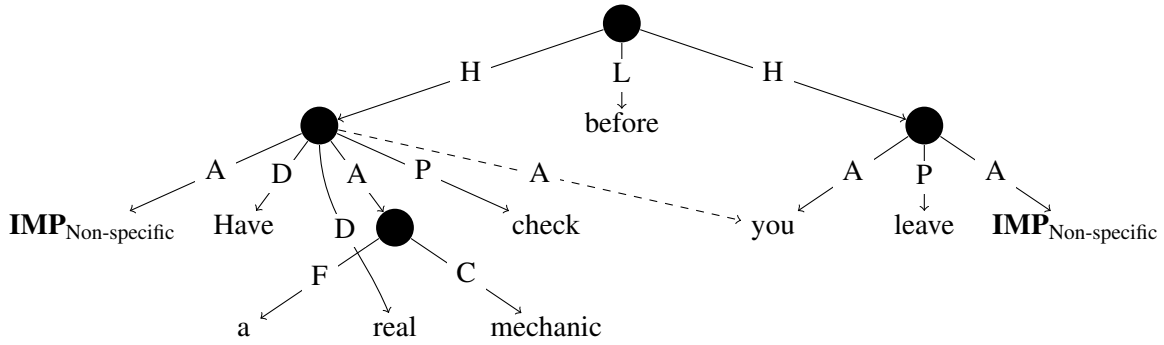


Figure 1: Example of UCCA graph: “Have a real mechanic check before you leave.”. Abbreviation of UCCA edge labels is explained in Table 1. The dashed line stands for *Remote* edge. In this case it is a coreference "you". An **IMP** represents an *Implicit* argument denoting a null-instantiated core element in its corresponding Scene.

A few studies have explored the possibilities to parse implicit arguments. Both Gerber and Chai (2012) and Cheng and Erk (2019) have developed parsers to recover implicit arguments for nominal predicates; Bender et al. (2011) parses against some linguistic constructions, two of which could license implicit arguments, that is, tough adjectives and verbal gerund; Elazar and Goldberg (2019) focus on resolving missing numeric fused-heads, which are implicit Centers in UCCA.

Providing the most up-to-date and fine-grained annotation of implicit arguments, our studies can potentially enhance natural language understanding if supplied with an adequate parser. For example, when companies conduct satisfaction analyses through web reviews, customers often express themselves colloquially in these reviews. Examples include “Serves bad ice cream, Joe’s is better” and “Near a nice district, bad service and expensive.” If these reviews are annotated with Genre-based implicit arguments, referring to the conventional omission of reviewees, algorithms can study which part of the reviews really refers to the companies rather than other entities, and make better predictions, despite the omission of subjects and non-standard language.

2 Universal Conceptual Cognitive Annotation

Universal Conceptual Cognitive Annotation is a semantic representation scheme whose design concept comes from the Basic Linguistic Theory typological framework (Dixon, 2010a; Dixon, 2010b; Dixon, 2012) and Cognitive Linguistics literature (Croft and Cruse, 2004). Abstracting away from syntactic forms, it aims at representing the main semantic phenomena in text while maintaining a low learning cost and rapid annotation by non-experts (Abend et al., 2017). Already providing datasets in English, French, and German, UCCA has demonstrated its cross-linguistic applicability in several languages and has become a popular target framework in multiple parsing tasks (Hershcovich et al., 2019b; Oepen et al., 2019). Abend et al. (2017) have also developed an open-source web-based annotation system, UCCAApp, which supports fast annotation for linguistic representations. The effectiveness and efficiency in annotating and refining UCCA have been proven by several studies (Shalev et al., 2019; Prange et al., 2019).

Participant	A	Linker	L
Center	C	Connector	N
Adverbial	D	Process	P
Elaborator	E	Quantifier	Q
Function	F	Relator	R
Ground	G	State	S
Parallel Scene	H	Time	T

Table 1: Legend of UCCA edge categories (Hershcovich et al., 2019a)

UCCA presents the meaning of a sentence with a directed acyclic graph (DAG) whose terminal nodes correspond either to surface lexical tokens or extra units representing implicit arguments. Non-terminal nodes correspond to semantic units that participate in some super-ordinate relation. Edges are labeled with the role of a child node related to its parent node. The basic notion in UCCA is Scene, describing

a state, action, movement, or some other relation that evolves in time. Each Scene involves one main relation (Process or State), and one or more Participants, including locations, abstract entities and sub-Scenes serving as arguments.

Furthermore, UCCA distinguishes *primary* edges, appearing explicitly in one relation, from *remote* edges, allowing a Scene to indicate its arguments by linking from another Scene. *Primary* edges form a tree while *Remote* edges allow reentrancy, forming a DAG. In some cases, an entity of importance in the interpretation of a Scene does not explicitly exist in the text. Hence, UCCA introduces the notion of *Implicit* Units to represent such kind of entity.

For instance, the sentence “Have a real mechanic check before you leave” in Figure 1 contains two Scenes, evoked by “check” and “leave”. The individual Scenes are annotated as follows:

1. “(You)_A have_D[a_F real_D mechanic_C]_A check_P IMP_{1A}”
2. “You_A leave_P IMP_{2A}”

“A mechanic” is a Participant in the first Scene, while “You” is Participant in both Scenes, as *Remote* constituting reentrancy in the first one (shown in dashed line), but explicit in the latter super-ordinate relations. In the first Scene, “Check” is the main Process used in the causative construction, requiring three Participants. While “You” refers to the customer making a request and “a real mechanic” is the service provider who should check something, the object that needs to be “checked” is missing. Therefore, we introduce an IMP A node to symbolize it. In the second Scene, “leave” is the Process meaning someone moves away from a source location. Although we state our little concern for non-core elements like location in this study in §3.2, the case is different for the “leaving” Scene since the source location is vital to the understanding of the departing action. For this reason, we add an IMP A to represent the place that is being left.

In UCCA’s foundational layer, only limited cases of implicit arguments have been annotated. The main focus is on omission licensed by certain grammatical structures,² whose notion is similar to *Constructional Null Instantiation* in FrameNet (Ruppenhofer et al., 2006). Two typical examples of such constructions are imperatives (forced omission of subjects) and passives (agent omission) in English:

1. Imperative: IMP_A Do_F n’t_D bother_P.
2. Passive: [The_F doctor_C]_A has_F already_T been_F paid_P IMP_A.

Several other kinds of constructions are mentioned in UCCA foundation layer guidelines, such as infinitive clause, gerund, and *thank* construction.

3. Infinitive clause: Is_F there_F [no_E other_E Verizon_C]_A IMP_A to_F go_P to_R [around_R downtown_C]_A?
4. Gerund: How_D addicting_S IMP_A going_P [to_R Fitness Unlimited_C]_A can_F be_F!”
5. “Thank” construction: IMP_A Thanks_P, John_A!

The annotation of implicit arguments in UCCA’s foundation layer is restricted to specific linguistic constructions, which is coarse, language-specific, incomplete, and unlike other distinctions in UCCA, is based on criteria of form rather than meaning. We are faced with the challenge of maintaining UCCA’s idiosyncrasy of differentiating *Remote* and *Implicit* while extending its boundary to include a rather refined categorisation for implicit roles.

²<https://github.com/UniversalConceptualCognitiveAnnotation/docs/blob/master/guidelines.pdf>

3 Implicit Roles In Meaning Representations

3.1 Background and Motivation

FrameNet (Ruppenhofer et al., 2006) is a source of inspiration for UCCA and Scenes can be seen as frame evocations. FrameNet developed Fillmore (1986)’s notion of null-instantiation into three types—Constructional Null Instantiation (CNI), licensed by grammatical constructions, Definite Null Instantiation (DNI), equivalent to core Frame Elements mentioned previously in text or inferrable from the discourse, and Indefinite Null Instantiation (INI), an element that is unknown and nowhere to be retrieved.

Nevertheless, this trichotomy treats unfavourably many cases where implicit roles occur, such as Free Null Instantiation (Fillmore and Kay, 1993, FNI;) and Identity of Sense Null Anaphora (Kay, 2004, ISNA;). FNI is neither restricted to definite nor indefinite null arguments, and ISNA is null instantiation within noun phrases. Lyngfelt (2012) even argues that the unclear definition of FNI leads to much false categorisation—some FNIs are unspecified, some are generic, and some should be considered DNI.

UCCA’s foundational layer mainly focuses on CNI, that is, the current annotated datasets only include grammatically licensed implicit arguments. So far, only a few corpora for implicit role labelling have been proposed, such as SemEval-2010 Task 10 (Ruppenhofer et al., 2009), Beyond NomBank (Gerber and Chai, 2010; Gerber and Chai, 2012), and Multi-sentence AMR (O’Gorman et al., 2018). But none of them is based on a more comprehensive fine grained implicit role characterisation theory aside from FiGref, refined on three corpora mentioned above for recoverability studies.

Although FiGref is not available to the public, O’Gorman (2019) has counterbalanced previous studies on implicit role description and synthesized an inventory of eleven interpretation types for implicit roles distinguished by their referential behaviours. They are Script-inferrable pragmatic, Salient/recent, Deictic, Remembered Roles of Event Reference, Implicit “Sloppy anaphora” and Bridging, Genre-based Default, Type-identifiable, Generic (“People in General”), Cataphoric, Low-information, and Iterated Events Implicit Roles.

Recoverable implicit roles fall into the category of *Remote* Participants in UCCA, which typically calls for coreference resolution (Prange et al., 2019). We discern the eleven types of implicit roles mentioned above, recoverable or not, and extract those types where only true *Implicit* arguments occur, that is, the argument cannot be explicitly recovered from text, but inference and non-specificity are allowed as they are aligned with the definition of implicit arguments of UCCA. In the following section, we will analyze these implicit role types and argue the appropriateness of the set of categories we choose.

For the sake of operability and consistency, we only focus on core arguments in Scenes where these arguments are essential to the meaning of corresponding predicates (Jackendoff, 1992; Jackendoff, 1997; Goldberg, 1992; Grimshaw, 1993). Elements such as location, time, and manner are of little interest in this study whilst they are able to appear as foundational units or *Remote* Participant in UCCA.

As UCCA distinguishes *Remote* edges from *Implicit* units, it is natural to take advantage of this property to account for argument recoverability. The definition of implicit arguments in UCCA, particularly for its strong emphasis on the inability of explicit recovering from text, is not strictly corresponding with the eleven implicit role types.

3.2 Forming UCCA Implicit Argument Typology

Table 2 shows the comparison of the primary eleven implicit role types and UCCA’s implicit argument set. Among these types of implicit roles in his inventory, four are definite implicit role constructions, viz. Salient/recent, Remember Roles, Script-inferrable, and Deictic. Only the last one of four, Deictic, we would consider a candidate category for UCCA’s *Implicit* arguments. Salient/recent roles, which can be directly found in the recent prior discourse, is the quintessential type of DNI, and they can be easily replaced by pronouns. Remember Roles and Script-inferrable roles, however, require a cognitive and reasoning process, as the referents can be understood through a common ground in the text shared between the speaker and the addressee, or are reflecting a different facet of the same or a subordinate event. Deictic roles, albeit explicit reference to the speaker or addressee, is an extra-linguistic and cannot

O’Gorman 2019	Definite	Indefinite	Edge Cases	UCCA’s Implicit Refinement
Salient/recent	✓			✗
Remember Roles	✓			✗
Script-inferable	✓			✗
Deictic	✓			Deictic
Cataphoric		✓		✗
Low-information		✓		Non-specific
Iterated Events		✓		Iterated-set
Bridging			✓	✗
Genre-based			✓	Genre-based
Generic			✓	Generic
Type-Identifiable			✓	Type-Identifiable

Table 2: The 11 implicit role types in O’Gorman 2019 and the set for UCCA’s implicit refinement.

be annotated as *Remote* Participant, since we are unable to retrieve them explicitly in the text. Therefore, we incorporate it in our set of *Implicit* arguments categories.

Three out of eleven implicit roles are marked as clearly indefinite arguments, namely Cataphoric, Low-information Arbitrary role, and Iterated Events Implicit Roles. Cataphoric, which Bhatia et al. (2014) define as “pragmatically specific indefinite”, is the only type we do not include in UCCA’s typology since it relies heavily on the interpretation of the discourse whether it will be referred to again. We would annotate it as *Remote* Participant if the role is mentioned in a later text, or Non-specific type if not so as not to complicate the reasoning process.

The other four, Bridging implicit roles, Genre-based Default, Generic and Type Identifiable, are regarded as edge cases. Once again, we will only admit the latter three in our typology. As far as we are concerned, bridging in ellipsis situations might not refer to the same referent conceptually. Nonetheless, it can be clearly resolved in the text. Therefore, it will be annotated as a *Remote* edge in UCCA.

We will focus on referents that do not appear anywhere in the text. Therefore, we follow the philosophy of UCCA and propose six categories of *implicit argument*, that is, Deictic, Generic, Genre-based, Type-identifiable, Non-specific and Iterated-set. In the next section, we will present and exemplify each one of them.

3.3 Categorisation Set for UCCA Implicit Arguments

3.3.1 Deictic

Deictic implicit arguments specifically refer to the speaker or the addressee in a sentence. In example 1, the second-person subject is exhorted to take a certain action, and such imperative construction allows the subject not to appear in the text explicitly. Shown in example 2, Deictic can also occur with certain interjections, where the subject as the speaker is habitually implicit.

- (1) Just_D ask_P them_A exactly_E what_C [you_A want_S (what)_A]_E *IMP*_{Deictic}.
- (2) [Thank you]_P guys_{S/G/A} *IMP*_{Deictic}.

It should be mentioned that only in certain languages is imperative likely to induce implicit arguments. In Romance languages such as Spanish, deictic information tends to be encoded morphologically due to person agreement (Ingram, 1971).

3.3.2 Generic

Generic implicit arguments denote “people in general” (Lambrecht and Lemoine, 2005). In example 3, the agent who “understands how this place has survived the earthquake” is not explicitly mentioned in the text, but it can be understood as it is the set of people in general. Example 4 can be construed as a gerund construction. “I” recommend taking a certain action. While the patient would not be specific, it conveys the message that “people in general” should follow such advice.

(3) It_F 's_F impossible_D to_F understand_P [how_C [[this_E place_C]_A has_F survived_P [the_F earthquake_C]_A]_E]_A *IMP_{Generic}*.

(4) I_A would_F recommend_P [not_D using_P [this_E company_C]_A *IMP_{Generic}*]_A.

3.3.3 Genre-based

Ruppenhofer and Michaelis (2010) found certain text genres, namely instructional imperatives, *labelese*, diary style, match report, and judgment-expressing quotative verbs, are closely linked with conventional omission. UCCA EWT corpus is based on online reviews of businesses and services by individuals. The review genre is so prominent across all dataset that it forms a pattern where the reviewers do not bother to mention the reviewees explicitly. In example 5 and 6, the review genre licenses the omission of the deliverer of the action “deliver” and server of the action “serve”, as they refer to the reviewees by default, because in both contexts it is the restaurants that are being reviewed.

(5) Delivery_P is_F [lightning_E fast_C]_D *IMP_{Genre-based} IMP_{Non-specific}!*

(6) [Great_D service_P *IMP_{Genre-based} IMP_{Generic}*]_H and_L [awesome_S prices_A]_H.

3.3.4 Type-identifiable

There exist some predicates allowing listeners to naturally think they “know” the omitted referents because of their high predictability. In example 7, the vague referent of “eat” can be understood from an inherent understanding of the listeners as “some kind of food”. In example 8, the thing that “I drive” is not mentioned. Instead, it comes from common sense that the referent should be a kind of vehicle. Whatever kind it is, the lack of explicit mention barely affects the understanding of the text.

(7) It_A is_F my_A favourite_S [place_C [(my)_A to_F eat_P *IMP_{Type-identifiable}*]_E]_A.

(8) I_A 'll_F drive_P [an_Q hour_C]_T [just_E for_R [their_S (volcano)_A]_E volcano_C]_A *IMP_{Type-identifiable}*.

3.3.5 Non-specific

Non-specific implicit arguments refer to the kind of referents that cannot be inferred or understood at all. Such required information absent from the context attributes to the low interpretability of the implicit arguments, leaving them non-specific. As in example 9 and 10, it is impossible to infer what is “delivered” or who “charged me” neither from common knowledge nor given context. Such kind of implicit arguments are commonly found in nominalization and passive because there are high possibilities that not all agents/patients are always mentioned despite the fact that they might be core frame elements.

(9) There_F is_F no_D delivery_P *IMP_{Genre-based} IMP_{Non-specific}*.

(10) I_A don_F 't_D think_P [I_A have_F ever_T been_F charged_P before_T *IMP_{Non-specific}*]_A.

3.3.6 Iterated-set

Iterated-set implicit arguments refer to a heterogeneous set of entities when the predicates are often an action that happens repeatedly, either iteratively or generically (Goldberg, 2001). For example, in sentence 11, the predicate “wait” implies high repetition, and the set of patients of “what/who I am waiting for” is so general that it does not hold any meaning beyond the context. As in example 12, the action “steal” designates a Scene where anything could be stolen, but “I” do not and will never steal. Unlike §3.3.4 Type-identifiable referring to a specific type of referents, the set of “things” in *Iterated-set* points to a vague set of entities to fill a role in a more functional way.

(11) I_A never_T wait_P [in_R the_F waiting_E room_C]_A [[more_C [than_R two_C]_C]_Q minutes_C]_T *IMP_{Iterated-set}*.

(12) I_A don_F 't_D steal_P *IMP_{Iterated-set}*.

3.4 Inherent Ambiguity and “Continuum” of Coreness

3.4.1 Category Priority

There are a few cases when it is difficult to choose between two categories. Since UCCA does not aspire to annotate all possible interpretations (Abend and Rappoport, 2013), the annotator should make a best guess and choose one option. The first one is between Deictic and Generic, shown in example 13. The second one is between Genre-based and Non-specific, as in example 14. To keep the annotation consistent and maintain as much information as possible, we always choose Deictic over Generic, and Genre-based over Non-specific if available.

(13) [The_F experience_C]_A [with_R every_Q department_C]_A has_F been_F great_D *IMP*_{Deictic}.

(14) I_A will_F definitely_D refer_P [[my_A friends_{A/S}]_C and_N [(my)_A family_{A/S}]_C]_A *IMP*_{Genre-Based}.

3.4.2 Nominalization As Occupation

It is a judgment call whether the patient of Process instantiated by a profession should be annotated at all. Even so it remains debatable which category such kind of implicit argument belongs to. In the current version of corpus, we will always annotate it as Type-identifiable. As in example 15, the patients of whom has been taught is unclear but neither require clarification. The Scene of teacher/teaching is annotated with a Type-identifiable implicit argument denoting a type of people receiving education.

(15) They_A are_F [very_E good_C]_D teachers_{A/P} *IMP*_{Type-identifiable}.

4 Refined Implicit Corpus

In furtherance of investigating the characteristics of UCCA’s implicit arguments, we piloted a study to revisit and refine part of English Web Treebank³ annotated with the UCCA foundational layer.⁴ 200 passages were randomly selected for experiment from the total 723 passages comprising the UCCA EWT dataset.

We use UCCAApp to carry out annotation. The process is divided into two stages. Firstly, we create passage-level review tasks to check the existing annotation whether they contain implicit arguments, and add missing arguments if necessary. Secondly, we split passages into sentences, create tasks with refinement layer and then annotate with corresponding fine-grained implicit categories. Since all the annotation works were undertaken by one single annotator, the dataset preferably serves as a demonstration of concept, and thus further measurement of inter-annotator agreement would be desired to establish a sounder dataset.

4.1 Revisiting Original EWT UCCA Dataset

The original implicit argument annotation in EWT UCCA corpus is restricted only to put concern on constructional null instantiations, and when a unit lacks a Center or a Process/State, which is out of the scope of this study. We only regard implicit argument whose category is Participant in UCCA as valid implicit in this research. Therefore, it is necessary to modify or add missing implicit arguments in the dataset. Considering the original corpus was annotated on passage-level, whereas our new dataset will be done on sentence-level, *Remote* edges across sentences will be treated by adding a new *Implicit* node under its origin Scene.

Table 3 shows the statistics before and after reviewing and refining according to the new UCCA Implicit Argument Typology. It can be seen that in the refined dataset, 116 out of 200 passages contain implicit arguments, 13 more passages than the original dataset, in which only 103 passages contain implicit arguments. Yielding an increase of 255%, the review process added 250 more valid implicit arguments in the corpus.

³<https://catalog.ldc.upenn.edu/LDC2012T13>

⁴https://github.com/UniversalConceptualCognitiveAnnotation/UCCA_English-EWT

	# Passages	# Passages w/ Implicit	# Sentences	# Sentences w/ Implicit	# Implicit (Valid)
Original	200 (out of 723)	103	306	111	153 (98)
Refined	200	116	393	221	415 (385)

Table 3: Statistics of the UCCA EWT dataset sampled passages before and after reviewing. The additional implicit arguments result from both reviewing original UCCA EWT and conducting new annotation according to our fine-grained typology. Implicit (Valid) denotes implicit argument whose role is Participant in UCCA.

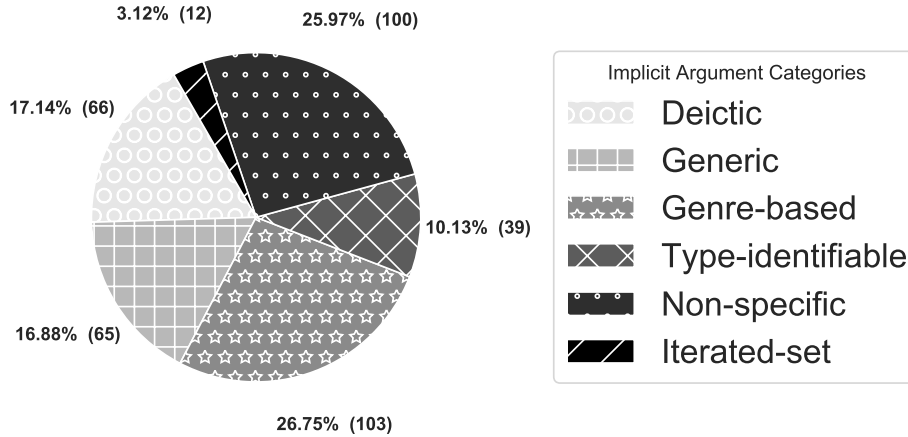


Figure 2: Statistics of pilot UCCA implicit dataset.

4.2 Statistics

Tokenized and split according to the Universal Dependencies English Web Treebank (Silveira et al., 2014; Hershovich et al., 2019a), this pilot corpus consists of 3702 tokens, 1411 nodes, and 4759 edges over 393 sentences. In total, 385 valid implicit arguments are found and annotated on 221 sentences. Figure 2 demonstrates each implicit argument category with its corresponding number in the pilot refined implicit corpus, and illustrates the percentage of each implicit category in the dataset. One can see that Genre-based and Non-specific are the two most frequent categories, both of which have more than 100 instances in the dataset, making up approximately 52% combined. They are followed by Generic and Deictic, and each occupies about 17%. Type-identifiable comes penultimate with 39 instances, while Iterated-set is the least frequent type, which merely has 12 instances, accounting for 3.12% of the whole corpus.

4.3 Comparisons to FiGref Annotation

	Type-Identifiable	Deictic	Generic	Non-specific	Genre-based	Iterated-set	Script-inferrable	Other	Invalid
Ours	9.79%	17.29%	17.29%	25.93%	25.93%	2.59%	\	\	\
FiGref	7%	5%	10%	4%	\	\	9%	12%	53%

Table 4: Relative frequency of annotated implicit types in UCCA’s refinement layer and FiGref’s annotation for non-recoverable roles in Multi-sentence AMR. The first four types are shared by both annotation corpora. The following two are exclusive to UCCA’s refinement layer. The last three are additionally introduced into FiGref’s set of interpretation types.

FiGref is annotated over Multi-sentence AMR training data, SemEval-2010 Task 10 training data, and Beyond NomBank. It contains 856 implicit roles classified into 14 types, which includes all 11 proposed interpretation types except Genre-based, and four more kinds of invalid roles to account for those implicit roles of low importance or deal with tricky occasions (Local-mention, Contextually-invalid, Non-

predictive and Low-importance). However, FiGref has a relatively low Cohen’s κ score (Cohen, 1960) of 55.2. With 14 types to distinguish, FiGref has a relatively high annotation complexity and ambiguity.

Comparatively, we do not annotate genuinely invalid implicit role such as these four low importance kinds, since we consider they are simply not core elements of event. As UCCA differs *Implicit* units from *Remote* edges, it naturally reduces the annotation complexity. Therefore, we do not have to distinguish and annotate these five types in O’Gorman 2019, viz.: Saisent/recent, Remember Roles, Script-inferable, Cataphoric and Bridging. We consider them belong to *Remote* in UCCA. We only need to attend to essential and unrecoverable implicit arguments within the set of six types we propose in §3.3.

Owing to the distinct annotation design and lack of statistics provided by FiGref, it is difficult to perform a comparative quantitative study between UCCA’s implicit refinement layer and FiGref. However, we are able to look into the relative frequency of annotated implicit types in UCCA and FiGref’s annotation for non-recoverable roles in Multi-sentence AMR shown in Table 4.

The distribution distinction can be possibly explained by the different domains of the corpora and their annotation methodology. We keep Genre-based type to account for the particular “review” genre in EWT dataset. Among the three types FiGref has introduced, We can see that invalid roles dominate the FiGref annotation with 53%. This is because a large amount of non-important interpretations of null-instantiation are taken into account in FiGref, whereas the implicit refinement annotation designed for UCCA is limited to essential important implicit arguments so as to lower annotation complexity and ambiguity.

5 Conclusion and Future Work

We proposed a novel typology for different implicit arguments in UCCA, which allows annotation with a relatively low cognitive load. Then we reviewed and refined part of the existing UCCA English Web Treebank dataset and piloted annotation of our guidelines with a refinement layer of fine-grained implicit arguments. It is currently the only published dataset with this kind of information.

Our work addresses a deep linguistic problem in implicit role interpretation, impossible to deal with current machine learning approaches, and provides an example for conducting research about it. It has the potential to benefit natural language understanding, especially in genre-specific context, e.g., social media and customer reviews. In this work we only use the existing definitions from O’Gorman (2019), testing their compatibility with the UCCA guidelines and their applicability to an UCCA annotated corpus. In future work we would also like to reason about and possibly redesign the criteria based on semantic and pragmatic considerations (Lyngfelt, 2012). We will also consider the usefulness of the refined scheme for practical NLP tasks, and possibly discard some of the categories, such as Iterated-set, which is the least frequent category in our corpus.

It is anticipated that our study will inspire tailored design of implicit role annotation in other meaning representation frameworks. Downstream tasks such as coreference resolution and human-robot interaction are also likely to benefit from reducing ambiguity and increase contextual understanding by explicitly modelling implicit arguments.

While the pilot design is promising, existing parsers that tackle implicit arguments only attend to certain aspect of linguistic phenomena. Therefore, for a rather comprehensive implicit argument annotation the like of UCCA, it is crucial to develop a parser that can emit implicit nodes dynamically and label them with fair accuracy.

Finally, work on reviewing the corpus by a second annotator is underway to validate the annotation quality by providing an inter-annotator agreement measurement. Future work will also expand the corpus and extend it to multiple languages.

Acknowledgements

We would like to thank the anonymous reviewers, Omri Abend and Dotan Dvir for their helpful feedback.

References

- Omri Abend and Ari Rappoport. 2013. Universal Conceptual Cognitive Annotation (UCCA). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 228–238, Sofia, Bulgaria, August. Association for Computational Linguistics.
- Omri Abend, Shai Yerushalmi, and Ari Rappoport. 2017. UCCApp: Web-application for syntactic and semantic phrase-based annotation. In *Proceedings of ACL 2017, System Demonstrations*, pages 109–114, Vancouver, Canada, July. Association for Computational Linguistics.
- Emily M. Bender, Dan Flickinger, Stephan Oepen, and Yi Zhang. 2011. Parser evaluation over local and non-local deep dependencies in a large corpus. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 397–408, Edinburgh, Scotland, UK., July. Association for Computational Linguistics.
- Archana Bhatia, Mandy Simons, Lori Levin, Yulia Tsvetkov, Chris Dyer, and Jordan Bender. 2014. A unified annotation scheme for the semantic/pragmatic components of definiteness. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 910–916, Reykjavik, Iceland, May. European Language Resources Association (ELRA).
- Pengxiang Cheng and Katrin Erk. 2019. Implicit argument prediction as reading comprehension. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6284–6291.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46.
- William Croft and D Alan Cruse. 2004. *Cognitive linguistics*. Cambridge University Press.
- Robert MW Dixon. 2010a. *Basic linguistic theory Volume 1: Methodology*, volume 1. Oxford University Press.
- Robert MW Dixon. 2010b. *Basic linguistic theory Volume 2: Grammatical topics*. Oxford-New York, Oxford University Press.
- Robert MW Dixon. 2012. *Basic linguistic theory: Further grammatical topics, Vol. 3*. Oxford: Oxford University Press.
- Yanai Elazar and Yoav Goldberg. 2019. Wheres my head? definition, data set, and models for numeric fused-head identification and resolution. *Transactions of the Association for Computational Linguistics*, 7:519–535.
- Charles John Fillmore and Paul Kay. 1993. *Construction grammar coursebook: Chapters 1 thru 11*. University of California.
- Charles J Fillmore. 1986. Pragmatically controlled zero anaphora. In *Annual Meeting of the Berkeley Linguistics Society*, volume 12, pages 95–107.
- Simon Garrod and Melody Terras. 2000. The contribution of lexical and situational knowledge to resolving discourse roles: Bonding and resolution. *Journal of memory and language*, 42(4):526–544.
- Matthew Gerber and Joyce Chai. 2010. Beyond NomBank: A study of implicit arguments for nominal predicates. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1583–1592, Uppsala, Sweden, July. Association for Computational Linguistics.
- Matthew Gerber and Joyce Y. Chai. 2012. Semantic role labeling of implicit arguments for nominal predicates. *Computational Linguistics*, 38(4):755–798.
- Adele E. Goldberg. 1992. The inherent semantics of argument structure: The case of the english ditransitive construction. *Cognitive Linguistics*, 3(1):37–74, January.
- Adele E Goldberg. 2001. Patient arguments of causative verbs can be omitted: The role of information structure in argument distribution. *Language sciences*, 23(4-5):503–524.
- Jane Barbara Grimshaw. 1993. *Semantic structure and semantic content in lexical representation*. Rutgers University.
- Daniel Hershcovich, Omri Abend, and Ari Rappoport. 2019a. Content differences in syntactic and semantic representation. 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 478–488, Minneapolis, Minnesota, June. Association for Computational Linguistics.

- Daniel Hershcovich, Leshem Choshen, Elior Sulem, Zohar Aizenbud, Ari Rappoport, and Omri Abend. 2019b. SemEval 2019 task 1: Cross-lingual semantic parsing with UCCA. In *Proc. of SemEval*.
- David Ingram. 1971. Toward a theory of person deixis. *Research on Language & Social Interaction*, 4(1):37–53.
- Ray Jackendoff. 1992. *Semantic structures*, volume 18. MIT press.
- Ray Jackendoff. 1997. Twistin’the night away. *Language*, pages 534–559.
- Paul Kay. 2004. Null instantiation of nominal complements. *Unpublished ms., University of California, Department of Linguistics, Berkeley*.
- Knud Lambrecht and Kevin Lemoine. 2005. Definite null objects in (spoken) french. *Grammatical constructions: Back to the roots*, 4:13.
- Benjamin Lyngfelt. 2012. Re-thinking fni: On null instantiation and control in construction grammar. *Constructions and frames*, 4(1):1–23.
- Anita Mittwoch. 1971. Idioms and unspecified np deletion. *Linguistic Inquiry*, 2(2):255–259.
- Stephan Oepen, Omri Abend, Jan Hajic, Daniel Hershcovich, Marco Kuhlmann, Tim O’Gorman, Nianwen Xue, Jayeol Chun, Milan Straka, and Zdenka Uresova. 2019. MRP 2019: Cross-framework meaning representation parsing. In *Proceedings of the Shared Task on Cross-Framework Meaning Representation Parsing at the 2019 Conference on Natural Language Learning*, pages 1–27, Hong Kong, November. Association for Computational Linguistics.
- Tim O’Gorman, Michael Regan, Kira Griffitt, Ulf Hermjakob, Kevin Knight, and Martha Palmer. 2018. AMR beyond the sentence: the multi-sentence AMR corpus. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3693–3702, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.
- Timothy J O’Gorman. 2019. *Bringing Together Computational and Linguistic Models of Implicit Role Interpretation*. Ph.D. thesis, University of Colorado at Boulder.
- David M Perlmutter. 1968. *Deep and surface structure constraints in syntax*. Ph.D. thesis, Massachusetts Institute of Technology.
- Jakob Prange, Nathan Schneider, and Omri Abend. 2019. Semantically constrained multilayer annotation: The case of coreference. In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 164–176, Florence, Italy, August. Association for Computational Linguistics.
- James Pustejovsky, Ken Lai, and Nianwen Xue. 2019. Modeling quantification and scope in Abstract Meaning Representations. In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 28–33, Florence, Italy, August. Association for Computational Linguistics.
- Ana Teresa Pérez-Leroux, Mihaela Pirvulescu, and Yves Roberge. 2018. *Direct Objects and Language Acquisition*. Cambridge Studies in Linguistics. Cambridge University Press.
- Josef Ruppenhofer and Laura A Michaelis. 2010. A constructional account of genre-based argument omissions. *Constructions and Frames*, 2(2):158–184.
- Josef Ruppenhofer, Michael Ellsworth, Myriam Schwarzer-Petruck, Christopher R Johnson, and Jan Scheffczyk. 2006. Framenet ii: Extended theory and practice.
- Josef Ruppenhofer, Caroline Sporleder, Roser Morante, Collin Baker, and Martha Palmer. 2009. 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)*, pages 106–111, Boulder, Colorado, June. Association for Computational Linguistics.
- Nathan Schneider, Jena D. Hwang, Vivek Srikumar, Jakob Prange, Austin Blodgett, Sarah R. Moeller, Aviram Stern, Adi Bitan, and Omri Abend. 2018. Comprehensive supersense disambiguation of English prepositions and possessives. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 185–196, Melbourne, Australia, July. Association for Computational Linguistics.
- Adi Shalev, Jena D. Hwang, Nathan Schneider, Vivek Srikumar, Omri Abend, and Ari Rappoport. 2019. Preparing SNACS for subjects and objects. In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 141–147, Florence, Italy, August. Association for Computational Linguistics.

Natalia Silveira, Timothy Dozat, Marie-Catherine de Marneffe, Samuel Bowman, Miriam Connor, John Bauer, and Chris Manning. 2014. A gold standard dependency corpus for English. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 2897–2904, Reykjavik, Iceland, May. European Language Resources Association (ELRA).