

YARN is All You Knit

Encoding Multiple Semantic Phenomena with Layers

Siyana Pavlova, Maxime Amblard, Bruno Guillaume

Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

{firstname.lastname}@loria.fr

Abstract

In this paper, we present the first version of `YARN`, a new semantic representation formalism. We propose this new formalism to unify the advantages of logic-based formalisms while retaining direct interpretation, making it widely usable. `YARN` is rooted in the encoding of different semantic phenomena as separate layers. We begin by presenting a formal definition of the mathematical structure that constitutes `YARN`. We then illustrate with concrete examples how this structure can be used in the context of semantic representation for encoding multiple phenomena (such as modality, negation and quantification) as layers built on top of a central predicate-argument structure. The benefit of `YARN` is that it allows for the independent annotation and analysis of different phenomena as they are easy to “switch off”. Furthermore, we have explored `YARN`’s ability to encode simple interactions between phenomena. We wrap up the work presented by a discussion of some of the interesting observations made during the development of `YARN` so far and outline our extensive future plans for this formalism.

Keywords: semantics, semantic framework, formalisation, layered semantic representation

1. Introduction

Current semantic representation formalisms can be split into two broad categories - those inspired by Logic (Kamp and Reyle, 1993; Montague, 1970), and those stemming from a graph-based perspective (Banarescu et al., 2013; Abend and Rappoport, 2013; White et al., 2016; Van Gyssel et al., 2021). While powerful in terms of encoding, logic-based representations can be difficult to read without prior training in Logic. Graph-based ones, on the other hand, are easier to read, but often lack when it comes to expressing scope or being compositional.

In this work, we aim to find a way to “mediate” between the two and find a representation which is both powerful in terms of encoding as the first group is, but also easier to read and annotate, as the second group is. Thus, we focus on the differences stemming from the logic-based vs graph-based view. Reviews of further differences between various deep-syntax or semantic representation formalisms can be found in the literature (Žabokrtský et al., 2020; Abend and Rappoport, 2017; Pavlova et al., 2023b; Giordano et al., 2023).

We present here the first version of a new semantic representation formalism, `YARN` (from “laY-ered meANing RepresentatioN”), with a predicate-argument structure (PA-structure) based on Abstract Meaning Representation (AMR) (Banarescu et al., 2013), and a layered approach to encode semantic phenomena. We provide proof-of-concept examples which demonstrate how the layered structure can be used to encode phenomena such

as negation, modality, temporality and quantification, and how they can interact with each other. Considering the interactions between diverse phenomena presents a challenge that existing formalisms do not explicitly address. This question is undeniably complex, yet significant if we aspire to provide a realistic outlook on the practical application of representations. Our initial tests for the cases of modality and temporality show a promising start for `YARN`’s ability to model these.

The main motivation for our approach is to allow the user of the formalism to focus on phenomena that they are interested in exploring, by allowing them to “switch off” the ones they are not interested in as to not clutter the representation. This gives the opportunity to encode specific properties needed for a general interpretation, but still anchor in a global meaning representation. The main contributions of this article are (1) to position the importance of considering the modelling of several semantic phenomena at the same time (2) as well as their interactions in order to (3) propose a rich representation that remains accessible for annotation and use.

In section 2, we present some of the existing semantic representation formalisms which are currently the most developed and have a similar outlook. To fully present the representation power of layers, in section 3, we provide the formal definition for our formalism, followed by annotation examples in section 4. In section 5, we provide a number of discussion points regarding our proposal, as well as aspects concerning the broader topic of semantic representation. This is followed by an outline for our future work in section 6.

2. Semantic Formalisms

In this section we outline some existing semantic representation formalisms that we later compare to our proposal. We focus here on AMR as our proposal uses its PA-structure as a base. We then describe Uniform Meaning Representation (UMR) (Van Gysel et al., 2021) as it is an extension of AMR that addresses many of its shortcomings. Finally, we mention Discourse Representation Theory (DRT) (Kamp and Reyle, 1993) as an example of a logic-based formalism.

Abstract Meaning Representation (AMR) (Banarescu et al., 2013) is a formalism that is meant to be simple enough to allow for large-scale annotation. As such, it focuses on the PA-structure of a sentence, annotating core arguments of each predicate according to PropBank’s (Palmer et al., 2005) predicates and argument roles, as well as a closed set of non-core roles, to annotate additional arguments such as `time`¹, `location` or `manner`. However, to keep the simplicity, many semantic phenomena such as tense, plurality or scope are not accounted for. AMR has been developed with English in mind and does not claim to be universal. That being said, AMR-annotated datasets exist in multiple languages.

Uniform Meaning Representation (UMR) (Van Gysel et al., 2021) is currently the broadest extension of AMR and can be considered a formalism in its own right. It combines a number of AMR extensions proposed over the years (Donatelli et al., 2018; Pustejovsky et al., 2019) to annotate phenomena such as temporal information, aspect, quantifier scope and co-reference. One of UMR’s goals is to keep the simplicity and ease of annotation of AMR, while enriching the set of phenomena it accounts for. UMR is a relatively new formalism and no large corpora exist yet, but annotation work is underway, including annotation procedures for low resource languages.

Discourse Representation Theory (DRT) (Kamp and Reyle, 1993) was introduced with the idea of preserving the principles of compositionality introduced by Montague (Montague, 1970, 1973) while making the representation more accessible. One of the main contributions of DRT is to consider the semantic contribution of an utterance or one of its components as a function that updates the general representation. In this way, it takes into account the process of representation construction. Based on logical representation, it makes the concept of scope explicit by means of boxes containing information in the form of predicates. The logical relationships between

¹When temporal adverbials are present as separate surface tokens.

them are encoded in such a way that a semantic structure emerges, a structure that is useful, for example, for expressing the accessibility of the variables used. This structure is also extended for discourse representation with SDRT (Asher and Lascarides, 2003). The combination of semantic representation, logical properties and readability makes it a useful formalism. A large corpus of DRT-annotated data exists in the form of the Parallel Meaning Bank (PMB) (Abzianidze et al., 2017). There is also a recent proposal to simplify the notation to foster easier annotation (Bos, 2021).

3. Description

We propose a structure with a central graph, representing the PA-structure, on top of which various layers can be defined. We provide examples to demonstrate how layers can be used to encode semantic phenomena, be it by interacting with nodes in the graph, or between themselves.

We follow the neo-Davidsonian tradition of placing a variable at the centre of the representation, representing the event being described (Davidson, 1967; Parsons, 1990). Our goal is to represent the semantics of an event, encompassing its core PA-structure, and modifiers in a readable and as simple as possible framework.

3.1. Formal Definition

Here is the formal mathematical definition. A YARN is an 8-tuple $\langle S, V, F, E, \hat{E}, E_{F\hat{E}}, E_{\hat{E}V}, E_s \rangle$ where:

- S and V are sets of vertices
- F is a set of features
- E is a set of edges between pairs of vertices $v_1, v_2 \in V$
- E_{FV} , which we will also call \hat{E} , is a set of edges between a feature $f \in F$ and a vertex $v \in V$
- $E_{F\hat{E}}$ is a set of edges between a feature $f \in F$ and an edge $e \in \hat{E}$
- $E_{\hat{E}V}$ is a set of edges between an edge $e \in \hat{E}$ and a vertex $v \in V$
- E_s is a set of edges between a pair of vertices $s_1, s_2 \in S$

We can imagine a layer-based solution using hypergraphs instead, as they are sufficiently expressive, but in order to maintain direct readability we prefer this solution.

One way of approaching these definitions is to consider that the central element of the representation is a simple graph around the predicate defining the main event. This gives us a very readable base representation. To avoid making the representation more cumbersome, we don't modify it directly, but allow other information to be added in the form of layers. These new elements lead to the use of new objects that operate either on the graph nodes or on the layer edges.

3.2. From Definition to Semantic View

With the formal definition given, let us look at how YARN can be applied to semantic representation. The vertices from S can be thought of as *event* nodes, with one defined for each event in the text².

The vertices of V and edges of E can be thought of as the ones used in the graphs of AMR and AMR-derived representations. Vertices in V represent predicates and concepts. Edges in E represent core argument roles. In this part of the representation, our focus is on the core concept that constitutes the central event. The sub-categorisation in a meaning bank helps to identify the mandatory arguments, as for AMRs. Consequently, the representation is lucid and easy to comprehend.

However, concepts representing non-core arguments and their modifiers are not always linked to the main predicate (see Figure 5). This results in elements of V and E making up the PA-structure of the sentence, which is a connected component within the graph, but also a number of (smaller) connected sub-graphs for some of the non-core arguments. Thus, the resulting graph formed by V and E is not necessarily connected.

The vertices of F represent various semantic phenomena, such as temporality, quantification and modality. The vertices are connected by lines that run between the feature nodes and V nodes, resembling strands of yarn. Each phenomenon is assigned a colour to simplify the reading.

The edges in \hat{E} , $E_{F\hat{E}}$ and $E_{\hat{E}V}$ are used to represent the linking between the semantic phenomena being annotated and the predicates and concepts of the sentence, as well as between the semantic phenomena themselves. We will see in section 4 how these three different types of edges are used for the different phenomena.

Finally, edges in E_S represent relations between different events, which can be thought of as representing discourse relations.

²In this paper, we mainly use simple texts, each containing only one event. See section 5 for a discussion on annotating more complex examples.

3.3. Hello World Example

To bridge the formal definition and the pictorial examples that will follow, we will present the first such example also in the mathematical notation that follows naturally from the formal definition.

Let us consider the sentence “*I found a newspaper*”. Its formal representation, ensuing from the YARN definition is the following:

$$\begin{aligned} S &= \{S_1\} & V &= \{find-01, i, newspaper\} \\ F &= \{temp, quant\} \\ E &= \{(find-01, ARG0, i), \\ &\quad (find-01, ARG1, newspaper)\} \\ E_{FV} &= \{(quant, \exists, newspaper), \\ &\quad (temp, past, find-01)\} \\ E_{F\hat{E}} &= \emptyset & E_{\hat{E}V} &= \emptyset & E_S &= \emptyset \end{aligned}$$

We note that for these examples we have chosen PropBank (Palmer et al., 2005) as our predicate sense and argument role bank, utilising Propbank's Frame Files (Choi et al., 2010) to collect the relevant senses and argument roles.

Figure 1 is the YARN graphical representation of the same sentence. The PA-structure of the sentence is a graph which appears in the dotted box, where nodes are predicates and concepts, and edges are relations between a predicate and its arguments. For this sentence, as the vertices of type V , we have the predicate `find-01` and two concepts, `i` and `newspaper`, representing the two arguments of the predicate. The two arguments are linked to the predicate via two labeled edges of type E , annotating their argument roles as `ARG0` and `ARG1`, respectively. Up to this point (and in this example, but not in general), the PA-structure of YARN coincides with the entire AMR.




Figure 1: YARN for “*I found a newspaper*”, featuring temporality and quantification.

In addition, we have a vertex $S_1 \in S$, that represents the *event*, to which two features of type F are linked: `temp` for temporality and `quant` for quantification. An edge of type \hat{E} , labeled `past` links the `temp` to the main predicate `find-01`, indicating the event happened in the past. Another edge of type \hat{E} , labeled `exists` links `quant` to `newspaper`, in-

dicating existential quantification. For readability’s sake, we put a box around the PA-structure, but this is not a part of the formal representation.

Here, we limit the representation to two features to demonstrate their operation. YARN’s modularity is advantageous since only specific semantic aspects of modelling can be considered. If the analysis also encompasses others, for example modalities, a modality feature can be added with a new “thread of yarn”. We can selectively activate the features that interest us.

4. Towards Multi-Layered Examples

In this section, we demonstrate how the structure described in section 3 can be used to encode various semantic phenomena, with the help of a number of examples. In the following we will concentrate only on a handful of semantic phenomena, but enough to cover the ways to combine different types of elements of YARN. For the sake of clarity, we will use only the graphical representation of the formalism from here on, but the mapping from these representations to the formal one is direct.

Figure 2 is the YARN for “I couldn’t find the newspaper”. The PA-structure for this sentence is the same as in Figure 1, but differs from the AMR of the sentence, where there are additional nodes and edges to account for the possibility and the negation. Aside from the PA-structure, we have annotated three phenomena: modality, introduced by *could*, negation introduced by *n’t*, and the temporality of the main predicate. For each, a feature that connects to S_1 is added. *Could* indicates *possibility*, so we add an E_F edge linking *modal* with the corresponding label to the main predicate. The possibility is then negated, with an unlabeled $E_{F\hat{E}}$ edge from *neg* to the *possibility* edge. Finally, since the impossibility was in the past, we add an $E_{F\hat{E}}$ edge from *temp* to the *possibility* edge, with a label *past*.

Temporality classes include *past*, *present* and *future* for now and modality classes: *possibility* and *necessity*. We kept this simple as the choice of classes for each phenomenon is not the focus of this work. These will be extended and made to account for different granularities across languages via lattices (Van Gysel et al., 2019). This approach has already been adopted for meaning representations by UMR.

To discuss $E_{\hat{E}V}$, we will use the example in Figure 3. This very simple example helps us to showcase how YARN deals with a classical logical issue, quantification. Figure 3 is one of the possible representations for the sentence “Every cow ate an apple”. This sentence has two quantifiers - universal for *cow* and existential for *apple*, thus giving rise to scope ambiguity. Two readings ex-

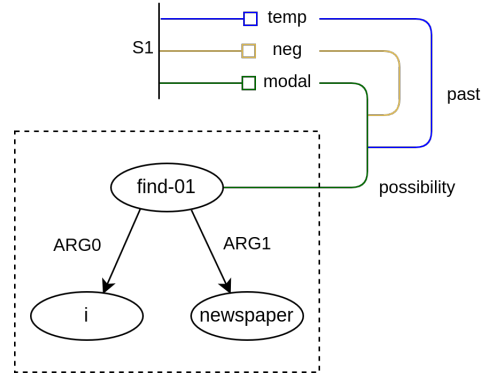


Figure 2: YARN for “I couldn’t find the newspaper”, featuring temporality, negation and modality.

ist: one where “every cow” takes wider scope, encoding the meaning where every cow ate a different apple, and one where “an apple” takes wider scope - where all cows ate the same apple. The representation in Figure 3 is for the latter. Here, aside from the *temp* feature, we have introduced a *quant* feature, linked as usual to S_1 . An edge of type E_F links *quant* to the wider-scope taking entity, namely *apple*, labeled with the appropriate quantifier, in this case \exists . Finally, an edge of type $E_{\hat{E}V}$ is introduced linking the \exists edge to the narrower scope entity *cow*. The appropriate label, \forall , is given to this edge. Thus, when annotating multiple quantifiers in a representation, we introduce a *quant* feature, then link it to the outermost scope-taking argument. Moving inwards, each argument is linked to the previous scope-defining edge.

A key issue in representing quantifiers semantically is the potential for combining scopes in various orders. The fundamental inquiry is whether every cow consumes an apple that is its own, or whether every cow consumes the same apple. Undoubtedly, pragmatics directs us towards a preferred interpretation for cows and many apples. Figure 4 provides an illustration of how quantifier scopes can be reversed. It appears that the type of link used for the quantifiers has changed, with the link for \exists deriving from an element of type $E_{\hat{E}V}$ and the link for \forall from type E_F . The entity now taking wider scope is the *cow*. The continuous link in the graphical representation represents the wider scope, while the link starting from the junction circle represents narrow scope. This approach fully utilises the expressiveness of YARN enabling the retention of readability whilst explicitly addressing logical constraints.

To represent some non-compulsory arguments, such as *manner* or *location*, we propose a solution as the one in Figure 5. Here, we have the representation of the sentence “Every cow ate an apple in the garden”. In addition to the two predicate-specific arguments of *eat-01*, an op-




Figure 3: YARN for “Every cow ate an apple”, featuring quantification and temporality. Reading where all cows ate the same apple.

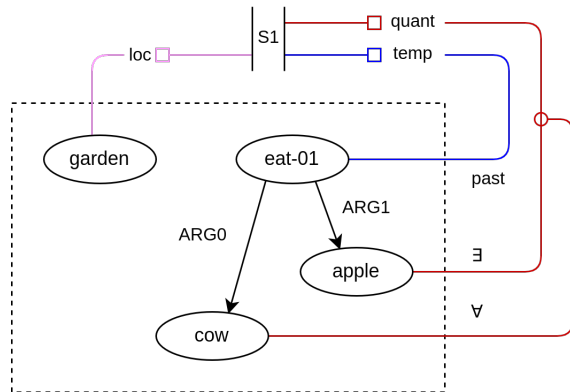


Figure 5: YARN for “Every cow ate an apple in the garden”, featuring an additional argument for location.




Figure 4: YARN for “Every cow ate an apple”, featuring quantification and temporality. Reading where every cow ate a different apple.

tional argument for `location`, namely “in the garden”, is specified. In AMR, optional arguments are attached to the predicate using the so called `non-core` roles in the same manner as predicate-specific ones. In our proposal, we annotate some of them as separate nodes (or subgraphs, in the case of more complex modifiers) that appear in the same box as the main predicate. To specify that argument’s role, we introduce a feature of type F , and an unlabeled edge of type E_F from the feature to the argument. In the example in Figure 5, a new feature `loc` is added that links S_1 to the node `garden` that has been added to the predicate box. The same can be done for other kinds of modifiers that are typically annotated with non-core roles in AMR.

In the preceding examples, we demonstrated how to formulate the control component that characterises the event, how to append non-compulsory parameters, and how the yarn principle straightforwardly encompasses distinct semantic phenomena, by integrating aspects of scope. The addition of a variable $s \in S$, which stands

for the event, is a beneficial realisation for modelling other occurrences, including those that are conventionally encountered in discourse representation. Figure 6 is a sample representation of the sentence “I entered the room, because the phone rang”, where we have a causal relation between the two events “enter” and “ring”. For simplicity, we have chosen not to show the quantification annotation here. However, this is entirely possible and would result in a more extensive sample, allowing the reader to select certain features for analysis. We introduce an edge of type E_S , labeled `CAUSE`, between the two event nodes S_1 and S_2 . This example shows that this representation also provides a solution for annotating discourse relations, taking the representation beyond semantics. Discussion of labelling the links between the S type elements that construct a higher level structure is beyond the scope of this work. Common discourse theories, such as SDRT (Asher and Lascarides, 2003) or RST (Mann and Thompson, 1986), can be utilised. YARN remains theory agnostic. In YARN, variables can represent elementary discourse units (EDUs) which creates a structure that covers the entire document, similar to SDRT. Alternatively, we can introduce relations between specific elements, as is done in RST.

Having seen the definition of the structure in subsection 3.1, and the examples above demonstrating how each can be used in the context of semantic representation, we sum up the characteristics for each element of our YARN 8-tuple in the context of semantic annotation. These elements demonstrate the technical nuances of formalisation, which can be linguistically interpreted.

- Edges in E and E_S are directed. For the rest of the edges, while there is an implicit direction - from a feature $f \in F$ towards either a vertex or another edge, or from an edge in E_{FV} , $E_{F\hat{E}}$ or $E_{\hat{E}V}$ towards a vertex, there is




Figure 6: YARN for “I entered the room, because the phone rang”, featuring a discourse relation between two events.

no need to draw it.

- Not all relations need to be labeled, only the information needed to disambiguate the interpretation is required.
- Each element f of F is linked to the *event* node of the representation, which means that it can be considered as the reification variable on which all the properties are applied.
- V is not a closed set. While users can choose the specific lexicon for the predicate senses and even the concepts, there is no restriction to do so - predicate senses and concepts can be used freely as the context in which the formalism is used requires it.
- F is a closed set of semantic phenomena. We have only briefly addressed some of these principles, deferring in-depth discussion to the future.
- Labels in E are a closed set – they can come from semantic role lexicons like VerbNet (Kipper et al., 2008), or be the set of core + non-core roles from AMR, just as two possible examples.
- Labels in E_F are a closed set, made up of the subsets defined for each feature. For these, we’ll use lattices for each feature, inspired by (Van Gysel et al., 2021).
- Quantifiers are expressed as labels on \hat{E} or $E_{\hat{E}V}$ edges. In the examples, we have solely employed common quantifiers, but it is possible to expand the list to encompass additional forms of more precise quantification.

These formal characteristics are crucial since they enable us to have a controlled representation that can be projected onto a logical representation by means of a simple algorithm, as in the case of DRT, while at the same time offering an adaptation to the needs of the linguistic representation without

mixing up the different elements, allowing us to focus on particular points.

5. Discussion

Here we discuss some of the observations we made while designing YARN. These concern features of the formalism itself and more general observations and questions about the kinds of phenomena we may want to include moving forward.

5.1. Yarn interactions

The introduction of these “strands of yarn” gives us a flexible structure for the representation and introduces a level where new interactions are possible, either through their non-explicitation with the sub-specification or through the swapping of relations.

Some semantic representation formalisms, such as Minimal Recursion Semantics (Copestake et al., 2005) allow for underspecification, for example in the case of scope ambiguity. While not present in the current version, as part of future work, we intend to offer the possibility for underspecification in the representation. In Figure 7 we provide a graphical representation of one possible solution for scope ambiguity/underspecification. As can be seen in the figure, the scope precedence between “an apple” and “every cow” has not been resolved. When the precedence becomes apparent from context, the representation can be updated in order to accommodate for that. As part of our future work, we intend to formalise this new structure, and investigate if the same one can be used for other types of underspecification.




Figure 7: “Every cow ate an apple” - YARN with underspecified scope.

In the example in Figure 2 we saw that the layers belonging to different phenomena can be stacked on top of one another. This property is not commutative. In Figure 2, modality is applied to the main predicate, followed by a negation applied to the modality in order to encode “couldn’t do

something]”. Compare this to the example in Figure 8, where the modality and negation have been swapped. If we try to build the meaning from the representation, we end up with “Maybe/[It is possible that] I did not find the newspaper”.

As our focus lies on the formal aspects of the framework, we will reserve the analysis and discussion of the stands of yarns’ interactions for future research. For instance, we have yet to exhaustively investigate how modality and temporality interact and significantly alter the resulting interpretation.




Figure 8: YARN of “Maybe I did not find the newspaper”.

Our goal with this proposal is to really separate the various semantic phenomena into different layers, and the focus is to make them easier to see at first glance. This is the reason why we prefer to disconnect optional arguments from the core PA-structure where that is linked to semantic phenomena. However, the mapping between this disconnected version and the AMR-style version is preserved, at least for what we have tested so far: simple combinations of the modifiers for time and location. Transformation functions between our “disconnected” version and the “connected” AMR-style one will be the subject of future work, allowing an easier transition between our formalism and AMR-based structures.

5.2. Comparison to other formalisms

AMR is the basis for our representation, but, as we have shown, YARN differs from it substantially in the way it encodes semantic phenomena. Thanks to this, we are able to achieve a richer representation, without making the reading too complex.

Bos (2020) points out that the proposal to extend AMR for quantification and scope presented in (Pustejovsky et al., 2019) suffers from the so called *bound variable problem*: the way sentences such as “Every snake bit itself” are represented, can be interpreted also as “Every snake bit every snake”. Since UMR (Van Gysel et al., 2021) im-

plements the above-mentioned extension, it faces the same issue. To tackle this, we need different representations for the two sentences. The YARN for “Every snake bit itself” is straightforward: in Figure 9a, the presence of a single quantifier indicates that the same entity participates both as an ARG0 and ARG1 in each occurrence of the biting event. Thus, for “Every snake bit every snake”, we perforce need two quantifiers. Figure 9b shows such a representation. However, we face another issue here: we cannot tell whether the wide-scope taking entity is the *biter* or the *bitee*. In the example with two universals, however, explicit resolution of the ambiguity is not required: the two options have different representations, but, when resolved, the interpretation is the same. However, in a case where we have this same representation but different quantifiers, this information is necessary. One possible solution for the same case but with different quantifiers, would be to include the name of the argument to which it is applied to the quantification edge, for instance $\forall : \text{ARG0}$. It does, however, still seem strange, to have two quantifiers pointing to the same entity (or set of entities). What is more, if we explore the other three quantifier combinations for two entities, we see that two universals is a special case: in all other cases, it is more natural to have separate entities for each participant. Thus, we propose the solution in Figure 9c. Here, we split the *biters* and *bitees* into two separate nodes and acknowledge, via a new type of edge (the dotted link in the figure) that in this special case it happens that the two nodes refer to the same set of entities, via the = sign. This new type of relation is yet to be formalised in the next version of YARN. This can be useful not only here, but also for linking co-referents (see Figure 10). Lastly, the representation in Figure 9b can still remain, purely for ease of readability, with the caveat that it is simply a visualisation, equivalent to the one in Figure 9c and in the background, the latter is the canonical form.

Aside from this, for formalisms where all events remain in the same graph, the representation becomes difficult to follow, especially for longer texts, as can be seen with UMR (Zhao et al., 2021). We believe our solution to represent each event as its own substructure makes our representation more easily readable, even for larger texts, making it easy to spot the phenomena applying to each event and the interactions between events thanks to the discourse-style relations between them.

This is illustrated by the annotation in Figure 10 for the sentence “I couldn’t find the newspaper until you told me where it suddenly appeared”. Here, we have three events: *finding*, *telling* and *appearing*, each represented by its own S-type node. Thanks to these, we can easily differentiate the features that apply to each event, and also track

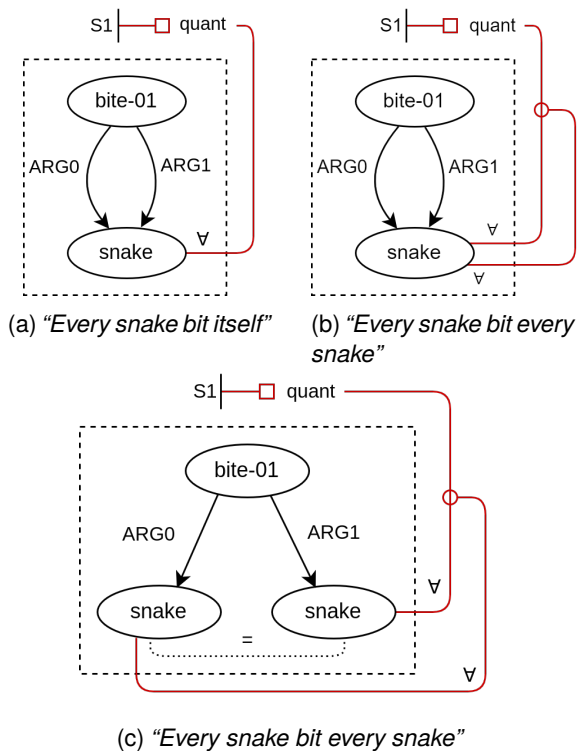


Figure 9: YARN for snake examples

the relationships between events. This example also illustrates that while subordinate clauses introduced by subordinating conjunctions such as “before”, “until” or “because” can be modelled by edges from E_S (as we see between $S1$ and $S2$), where a predicate permits it, they can be modelled by argument roles from E (as with $t_{\text{ell-01}}$ from $S2$ and the empty node marking the *location* from $S3$). Finally, by using the same type of edges as the ones proposed for the equality between sets in Figure 9c, we can also model co-reference, as shown by the links between the *i*’s in $S1$ and $S2$, and *newspaper* in $S1$ and *it* in $S3$.

It is worth noting that YARN also draws inspiration from DRT, although it is not a direct representation of it. As a result, the algorithmic principle is used to convert the representation into a standard logical formula, which induces the same structure as DRT. The crucial elements primarily lie in “stands” of quantification. The conversion from one to the other is a task for future work. Furthermore, thanks to the S nodes, we can imagine an expansion of the depiction from a standpoint comparable to SDRT’s expanding of DRT.

5.3. Some broader questions

We can have a broader discussion on what constitutes an event and how events are deduced from the surface form of a sentence. If we take the sentence from the WSJ (Paul and Baker, 1992), “Edmond Pope tasted freedom today for the first

time in more than eight months.”, we have the main event E , “tasted [freedom]”, but also a reference point to something that happened “more than eight months ago”. Thus, we may ask whether apart from E , we also have another (possibly static) event, E' of “having tasted freedom”, or “having been free” more than eight months ago. We may argue whether such implication (the one of E') is or should be part of the semantic representation of the sentence, or whether we should only annotate events that appear explicitly in the sentence.

Our annotation experiments so far demonstrate that it is easy to extend the formalism with new semantic phenomena. Adding a new one so far has consisted in adding a new feature to F and deciding on the appropriate type of relations to use for edges of that layer. To test our proposal, we are currently analysing sentences in English from the Parallel Universal Dependencies (PUD) corpus³. Although our proposal is currently robust, we may encounter issues when annotating more complex phenomena or sentences. It is yet to be determined, following the expansion of the range of observable occurrences, whether working on a larger dataset will help to test the ease of use of the framework. The annotation of a substantial corpus will enable us to assess the capabilities of YARN.

6. Conclusion and Future Work

In this paper, we presented the first version of YARN, a proposal for encoding multiple semantic phenomena with layers. The framework differs from others in that it maintains a logical structure, while remaining clear to the reader. The incorporation of diverse levels allows for the comprehensive modelling of various phenomena, whilst still maintaining their distinctiveness and potential interconnections.

We have shown, through examples, that our initial annotations show a promising structure that manages to encode difficult phenomena and keep the representation visually simple. Analysis is further aided by the fact that “switching off” layers is straightforward. We have highlighted interesting discussion points that were raised during the design of our formalism, and have outlined the future work directions for this project.

As we have shown, the thus proposed structure, YARN, is capable of representing a range of semantic phenomena, namely: temporality, modality, negation, quantifier scope. While not presented here due to space limits, we have also tested definiteness, number and questions. In the preceding section, we presented various view-

³https://github.com/UniversalDependencies/UD_English-PUD

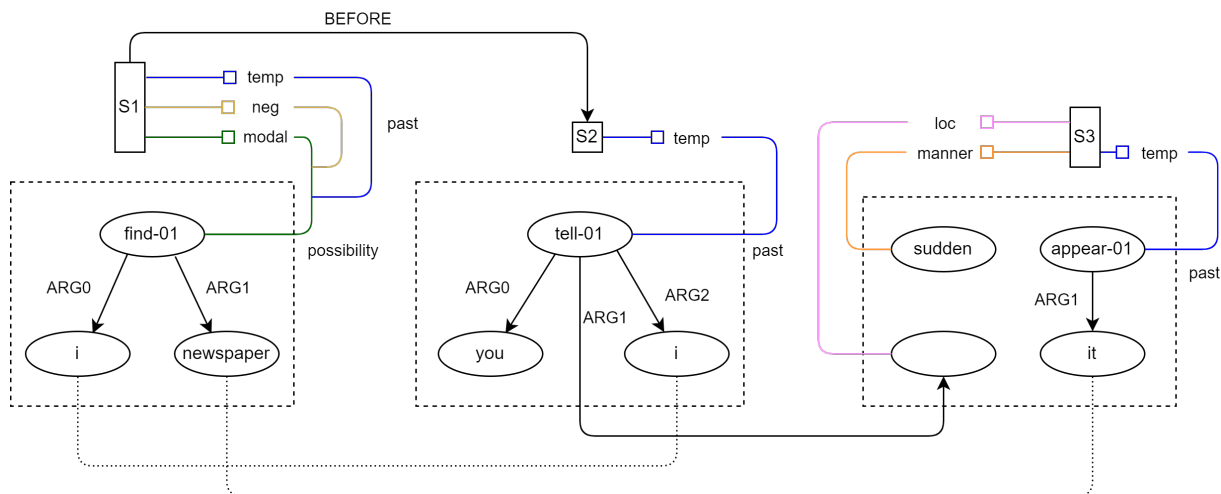


Figure 10: YARN of “I couldn’t find the newspaper until you told me where it suddenly appeared”.

points on the evolution of specific phenomena. We now go back to more general aspects.

As a first step, in our future work, we plan to add more phenomena to the formalism, such as comparison, gender, predicates whose arguments are events (such as “begin”, “stop”), etc. For each phenomenon, a set of possible classes will be defined. We do not intend to limit the classes, but rather allow lattices as presented in (Van Gysel et al., 2019) and used in UMR in order to enhance comparison between languages without limiting the possible classes to those available in a specific language.

In parallel to this, we will formalize annotation guidelines and develop annotation tools, with the help of which to carry out annotation experiments.

As mentioned earlier, one of our goals is to provide a formalism where “switching off” layers is simple, which is a major difference from others such as UMR. This is straightforward in cases where layers do not interact with each other (as in Figure 1). However, in more complex cases such as in Figure 2, the process is not straightforward. Removing the `temp` layer would not necessarily affect the `modal` layer as it is attached on top of it. However, what would it mean for the `temp` feature and the interpretation of the whole representation if only the `modal` feature were to be removed? Understanding this interaction and defining procedures on how to “switch off” a layer that interacts with other layers will be the subject of another future work.

As we want our representation to be able to “communicate” with both logic-based and established graph-based formalisms, we envision two further future work directions: (1) make explicit the formal procedure to convert a YARN into first-order logic and vice-versa, and (2) creating transformation systems between ours and other graph-based formalisms, in the spirit other transformation-based comparison works (Hersh-

covich et al., 2020; Pavlova et al., 2022, 2023a).

Finally, we want to propose a textual representation format for YARN, in the spirit of the PENMAN notation (Matthiessen and Bateman, 1991), widely used for AMR, and AMR-derived formalisms. We expect having such a representation will be useful for developing parsing algorithms for our formalism, both with symbolic and hybrid approaches.

7. Ethical Considerations

While Universality is one of the desired features for the presented meaning representation, we note that there is likely an inherent bias towards phenomena which are more prevalent in occidental linguistic culture, and English in particular, which is the main language we have used so far for YARN’s development. While we have not had the chance to do this for the current version of the formalism, we acknowledge that a more thorough study and discussion of non-occidental languages is necessary for a less biased representation. This is further affected by our use of PropBank, a sense lexicon, an equivalent of which is not available for the majority of world languages. Thus, we also need to employ strategies for either a resource agnostic resource development or follow UMR’s steps in proposing strategies on how to build and extend such resources for low-resource languages.

Acknowledgments

We would like to thank the anonymous reviewers for their feedback and comments. Part of this work has been funded by *Agence Nationale de la Recherche* (ANR, fr: National Research Agency), grant number ANR-20-THIA-0010-01.

8. Bibliographical 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. Association for Computational Linguistics.
- Omri Abend and Ari Rappoport. 2017. [The state of the art in semantic representation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 77–89, Vancouver, Canada. Association for Computational Linguistics.
- Lasha Abzianidze, Johannes Bjerva, Kilian Evang, Hessel Haagsma, Rik van Noord, Pierre Ludmann, Duc-Duy Nguyen, and Johan Bos. 2017. [The Parallel Meaning Bank: Towards a multilingual corpus of translations annotated with compositional meaning representations](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 242–247, Valencia, Spain. Association for Computational Linguistics.
- Nicholas Asher and Alex Lascarides. 2003. *Logics of Conversation*. Studies in Natural Language Processing. Cambridge University Press.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. [Abstract Meaning Representation for sembanking](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.
- Johan Bos. 2020. [Separating argument structure from logical structure in AMR](#). In *Proceedings of the Second International Workshop on Designing Meaning Representations*, pages 13–20, Barcelona Spain (online). Association for Computational Linguistics.
- Johan Bos. 2021. Variable-free discourse representation structures. *Semantics Archive*.
- Ann Copestake, Dan Flickinger, Carl Pollard, and Ivan A Sag. 2005. Minimal recursion semantics: An introduction. *Research on language and computation*, 3(2):281–332.
- Donald Davidson. 1967. The logical form of action sentences. In Nicholas Rescher, editor, *The Logic of Decision and Action*, pages 81–95. University of Pittsburgh Press.
- Lucia Donatelli, Michael Regan, William Croft, and Nathan Schneider. 2018. [Annotation of tense and aspect semantics for sentential AMR](#). In *Proceedings of the Joint Workshop on Linguistic Annotation, Multiword Expressions and Constructions (LAW-MWE-CxG-2018)*, pages 96–108, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Bastien Giordano, Cédric Lopez, and Immeuble Le. 2023. Mr4ap: Meaning representation for application purposes. In *The 15th International Conference on Computational Semantics (IWCS 2023)*.
- Daniel Hershcovich, Nathan Schneider, Dotan Dvir, Jakob Prange, Miryam de Lhoneux, and Omri Abend. 2020. [Comparison by conversion: Reverse-engineering UCCA from syntax and lexical semantics](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2947–2966, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Hans Kamp and Uwe Reyle. 1993. *From discourse to logic: Introduction to modeltheoretic semantics of natural language, formal logic and discourse representation theory*. Dordrecht. Kluwer.
- Karin Kipper, Anna Korhonen, Neville Ryant, and Martha Palmer. 2008. A large-scale classification of english verbs. *Language Resources and Evaluation*, 42:21–40.
- William C. Mann and Sandra A. Thompson. 1986. Relational propositions in discourse. *Discourse processes*, 9(1):57–90.
- Christian MIM Matthiessen and John A Bateman. 1991. Text generation and systemic-functional linguistics: experiences from english and japanese. *Communication in Artificial Intelligence Series*, 19(1).
- Richard Montague. 1970. English as a formal language. *Logic and philosophy for linguists*.
- Richard Montague. 1973. The proper treatment of quantification in ordinary english. In *Approaches to natural language: Proceedings of the 1970 Stanford workshop on grammar and semantics*, pages 221–242. Springer.

- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. [The Proposition Bank: An annotated corpus of semantic roles](#). *Computational Linguistics*, 31(1):71–106.
- Terence Parsons. 1990. *Events in the Semantics of English: A Study in Subatomic Semantics*. MIT Press.
- Douglas B. Paul and Janet M. Baker. 1992. [The design for the Wall Street Journal-based CSR corpus](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York, February 23-26, 1992*.
- Siyana Pavlova, Maxime Amblard, and Bruno Guillaume. 2022. [How much of UCCA can be predicted from AMR?](#) In *Proceedings of the 18th Joint ACL - ISO Workshop on Interoperable Semantic Annotation within LREC2022*, pages 110–117, Marseille, France. European Language Resources Association.
- Siyana Pavlova, Maxime Amblard, and Bruno Guillaume. 2023a. [Bridging Semantic Frameworks: mapping DRS onto AMR](#). In *The 15th International Conference on Computational Semantics (IWCS 2023)*, Nancy, France.
- Siyana Pavlova, Maxime Amblard, and Bruno Guillaume. 2023b. Structural and global features for comparing semantic representation formalisms. In *The 4th International Workshop on Designing Meaning Representation*.
- 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. Association for Computational Linguistics.
- Jens E. L. Van Gysel, Meagan Vigus, Pavlina Kalm, Sook-kyung Lee, Michael Regan, and William Croft. 2019. [Cross-linguistic semantic annotation: Reconciling the language-specific and the universal](#). In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 1–14, Florence, Italy. Association for Computational Linguistics.
- Jens EL Van Gysel, Meagan Vigus, Jayeol Chun, Kenneth Lai, Sarah Moeller, Jiarui Yao, Tim O’Gorman, Andrew Cowell, William Croft, Chu-Ren Huang, et al. 2021. Designing a uniform meaning representation for natural language processing. *KI-Künstliche Intelligenz*, 35(3-4):343–360.
- Aaron Steven White, Drew Reisinger, Keisuke Sakaguchi, Tim Vieira, Sheng Zhang, Rachel Rudinger, Kyle Rawlins, and Benjamin Van Durme. 2016. Universal decompositional semantics on universal dependencies. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1713–1723.
- Zdeněk Žabokrtský, Daniel Zeman, and Magda Ševčíková. 2020. [Sentence meaning representations across languages: What can we learn from existing frameworks?](#) *Computational Linguistics*, 46(3):605–665.
- Jin Zhao, Nianwen Xue, Jens Van Gysel, and Jinho D. Choi. 2021. [UMR-writer: A web application for annotating uniform meaning representations](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 160–167, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

9. Language Resource References

- Choi, Jinho D. and Bonial, Claire and Palmer, Martha. 2010. *PropBank Frame Files*. distributed by Linguistic Data Consortium. PID <https://github.com/proppbank/proppbank-frames/>.