The place of ISO-Space in Text2Story multilayer annotation scheme

António Leal, Purificação Silvano, Evelin Amorim, Inês Cantante, Fátima Silva
Alípio Jorge, Ricardo Campos
CLUP-Universidade do Porto, INESC TEC
{jleal, msilvano, mhenri}@letras.up.pt, cantante.ines@gmail.com
{evelin.f.amorim, alipio.jorge}@inesctec.pt
ricardo.campos@ipt.pt

Abstract

Reasoning about spatial information is fundamental in natural language to fully understand relationships between entities and/or between events. However, the complexity underlying such reasoning makes it hard to represent formally spatial information. Despite the growing interest on this topic, and the development of some frameworks, many problems persist regarding, for instance, the coverage of a wide variety of linguistic constructions and of languages. In this paper, we present a proposal of integrating ISO-Space into a ISO-based multilayer annotation scheme, designed to annotate news in European Portuguese. This scheme already enables annotation at three levels, temporal, referential and thematic, by combining postulates from ISO 24617-1, 4 and 9. Since the corpus comprises news articles, and spatial information is relevant within this kind of texts, a more detailed account of space was required. The main objective of this paper is to discuss the process of integrating ISO-Space with the existing layers of our annotation scheme, assessing the compatibility of the aforementioned parts of ISO 24617, and the problems posed by the harmonization of the four layers and by some specifications of ISO-Space.

Keywords: ISO-space, multilayer annotation scheme, European Portuguese, news articles corpus.

1. Introduction

The consideration of spatial information, in general, and spatial relations, in particular, is one of the essential functions of natural language. The formal mapping of such information is challenging, notably due to the vastness and diversity of linguistic constructions that materialize them, with the consequent impact on the setting of relations' patterns (Kordjamshidi et al., 2010). However, this knowledge and its computational modelling are fundamental for many applications in computational linguistics and artificial intelligence, for which annotated datasets are essential (Pustejovsky et al., 2019). The existence of such datasets is still limited, with English being the dominant language.

In this context, some models of spatial information extraction have been proposed, working independently or in integration with other information domains (a.o., SpatialML (Mani et al., 2008); Language resource management — Semantic annotation framework: (ISO-24617-7, 2020)). They aim to map the spatial information of a text or other type of information modality in a formal representation that seeks to account for how humans conceptualize and process space in a wide variety of genres and domains (Gritta et al., 2018).

In the context of Text2Story project [1], which aims to extract narratives from news, represent them in intermediate data structures, and make these available to subsequent media production processes, we follow ISO 24617-7 (ISO-24617-7, 2020), which allows for the annotation of static and dynamic locations, and of a wide range of spatial relations. So far, only a few datasets have applied these ISO guidelines for annotating spatial information, none of which are in European Portuguese. We have selected this model because it proposes a framework that allow us to represent properly the spatial relations occurring in our corpus, contributing to a more effective subsequent visualization. Additionally, ISO-Space conforms with the multilayer semantic annotation scheme that combines and harmonizes three parts of ISO 24617-1/4/9, and that we have already designed (Silvano et al., 2021) to represent time and events, semantic roles, and referential information. This new model also promotes the balance between the amount of needed information to extract narratives and the load of the annotation process.

All in all, this paper aims to:

- put forward a multilayer annotation model;
- describe the integration of the spatial information and the procedures necessary for its harmonization with the existing scheme;
- identify some problems related to the harmonization process, and pointers to overcome them.

The remainder of the paper is organized as follows. Section 2 is dedicated to presenting some previous work about spatial annotation. Section 3 presents Text2Story ISO-based annotation framework, starting with the description of the harmonization of the temporal, referential and thematic layers 3.1 and proceeding to the description of the integration of the spatial layer 3.2. Section 3.3 explains some of the problems of combining ISO-Space to parts 1, 4, 9 of ISO 24617 and of applying it to the annotation of Portuguese data.

[1] https://text2story.inesctec.pt

Finally, some concluding remarks and future work are provided in Section 4.

2. Related work

During the past few years, there have been some new proposals for modeling spatial information, based on different theoretical, and domain-specific, approaches, which have advanced the state-of-art. Mani et al. (2008), for instance, propose SpatialML, an annotation scheme for marking up named and nominal references to places in natural language with geo-coordinates, whenever possible, and relationships among places in terms of a region calculus. Kordjamshidi et al. (2010) describe the task of spatial role labeling, which is language-independent, to identify and classify spatial arguments of spatial expressions in a sentence, and the mapping of the spatial relations established in natural language to formal representations by means of machine learning.

One of the best known and most comprehensive proposals is *Language resource management — Semantic annotation framework (SemAF) — Part 7: Spatial information (ISO-Space) Standard ISO/IEC TR 24617-7* (ISO-24617-7, 2020). ISO-Space, which has been under development for several years – one of the first papers to introduce was Pustejovsky et al. (2011) - aims at postulating "normative specifications not only for spatial information, but also for information content in motion and various other types of event in language" (ISO-24617-7, 2020), thus, enabling static and dynamic spatial annotations. The first edition, published in 2014, was revised in 2020 to conform with ISO 24617-6 (ISO-24617-6, 2016), incorporating some changes proposed along the years (for instance, Lee (2016)), and restoring the original proposal by Pustejovsky et al. (2012) and Pustejovsky and Yocum (2013), regarding event-paths (Pustejovsky et al., 2019).

There are, however, only a few studies and tasks that actually apply ISO-Space to corpus annotation. Some of those have been carried out during the process of building the standard, and they have been useful to identify some issues and to come up with solutions for them. For instance, Pustejovsky and Yocum (2013) describe the motion sub-corpus of ISO-SpaceBank with 50 entries from a travel blog, and its pilot annotation revealed some problems, which have led to the inclusion of further specifications in the standard. SpaceEval task (Pustejovsky et al., 2015) goes one step further aiming at automatic extraction of spatial information by means of supervised and semi-supervised machine learning systems. The evaluation results show that recognition of spatial entities and of *MOVELINK* are more easily achieved than of recognition of spatial relations between spatial entities. Since ISO-Space is an international standard directed to a wide range of languages, Lee et al. (2011) apply a first draft of ISO-Space to datasets from three typologically different languages, English (inflectional analytic), Korean (agglutinative) and Chinese (isolating), to ascertain its descriptive suitability. The version that was applied was considered largely suitable to describe spatial information, even though some modifications were required to deal with language specificities.

Other studies have demonstrated ISO-Space's scope and expressivity. On the one hand, the model can be used to represent spatial information in textual data, be they narratives (Pustejovsky and Yocum, 2013) or captions of images (Pustejovsky and Yocum, 2014) or even scene descriptions (Gaizauskas and Alrashid, 2019), and there are plans to transfer such annotations to three dimensions environments facilitating the reconstruction of scenes from text (Henlein et al., 2020). On the other hand, in terms of language coverage, although ISO-Space (ISO-24617-7, 2020) refers that the specification language for spatial annotation concerns English, and that its applicability to other languages is missing, Estarrona and Aldezabal (2018) discuss the viability of applying it to a Basque corpus. Despite the fact that the two languages differ, namely concerning prepositions, which in Basque, an agglutinative language like Korean, are postpositions, and that some adaptations are required, the authors conclude that this standard is adequate to codify spatial information in Basque.

Regardless of what has been accomplished, the reality is that ISO-Space, in particular the last edition from 2020, is not yet widely used in corpus annotations and/or in applications (like, for instance, ISO-TimeXL). According to Henlein et al. (2020), the reasons that explain why such situation is happening with this and other linguistic models for spatial information could be the model's complexity, and the lack of annotated datasets and of automated taggers. With our work, we intend to: (i) demonstrate that the ISO-Space is not so complex that it cannot be applied to a dataset from other language than English and that it cannot be integrated into a multilayer annotation scheme with models from other parts of ISO 24617; (ii) alleviate the shortage of annotated corpus in European Portuguese.

3. Text2Story annotation scheme

3.1. Temporal, referential and thematic layers

In order to extract the relevant information from our dataset, composed of news articles, we developed an interoperable multilayer semantic annotation that harmonizes three parts of the standard *Language resource management-Semantic annotation framework*: Part 1-Time and events (ISO-24617-1, 2012), Part 4- Semantic roles (ISO-24617-4, 2014) and Part 9- Referential annotation framework (ISO-24617-9, 2019) (Silvano et al., 2021). This task presented some obstacles, but we were able to overcome difficulties, with a constant work of annotation, revision and correction as proposed by MATTER's (Pustejovsky and Stubbs, 2012) subcycle, MAMA. This methodology allowed us to solve

problems and incongruities that were identified along the way. Furthermore, the use of BRAT (Stenetorp et al., 2012) as an annotation tool made it possible to update the scheme almost incrementally, maintaining, therefore, some of its parts, while erasing superfluous ones.

The first step was to model the types of structures as entity structures and link structures. The entity structures comprise events, times and participants, while the link structures incorporate the following links: *temporal, aspectual, subordination, objectal* and *semantic role*. Our model enables the annotation of an intricate network of relationships, because the different entity structures can be related among them by different types of link structures. For each structure type, subtypes were selected from the aforementioned parts of ISO-24617, rendering an annotation scheme with three layers: temporal, referential and thematic (semantic role labeling).

The temporal layer, based on ISO-24617-1, contains temporal information, and reconstructs the temporal sequencing of events, through the use of temporal links (*TLinks*). Despite following the standard, our scheme leaves out some tags and links that we considered unnecessary for the project's goal, that is, the construction of narratives' visualizations extracted from news [2]. Thus, we adopted the tags for events (*EVENT*) and times (*TIMEX3*). The former is used to mark eventualities, in the form of verbs, nominalizations, adjectives, pronouns, predicative constructions or prepositional complements. The remaining information about the eventualities is given by the following attributes: *class, part of speech, tense, aspect, verb form, mood, modality* and *polarity*. These attributes have the same values that were established for the Italian, in ISO-24617-1, with the exception of *mood*, to which we added the value *future*, and *modality*, complemented with the European Portuguese modal verbs *dever* (*must*), *poder* (*can*), *ter de* (*have to*) e *ser capaz de* (*be able to*). Regarding *TIMEX3* tag, we decided to keep the annotation scheme very simple, maintaining only two tags, relative to *type* (*date, time, duration, set*) and *value*, which refers to a specific value, manually inserted by the annotator, and two attributes, *temporal function*, to indicate that the *TIMEX3* expression corresponds to the *publication time*, and *anchor time*. Lastly, the extraction of the events timeline is accomplished by means of temporal links (*TLinks*), which represent relations between the events, events and times and between times. The aspectual and subordination links (*ALink* and *Slink*, respectively) encompass the pertinent aspectual and subordination information as proposed by ISO-24617-1.

The second layer refers to the identification and annotation of referential expressions, which, in Text2Story project, are represented by named entities referring to participants in the narrative. Although overall ISO-24617-9 is followed, considering that the scheme is multilayer, in order to avoid repetitive information and an overly complex annotation work, we have selected only the relevant tags for our purposes. For the discourse entity structures, the attribute *lexical head* has the values *noun* or *pronoun*. For the referential expression structures, the attributes for *domain* and *involvement* were integrated. The former introduces information about *individuation*, with the values *set, individual or mass*, retrieved from ISO-24617-9, and *types*, with values taken from several named entity classification typologies (PER (person), ORG (organization), LOC (location), OBJ (object), NAT (nature), OTHER). The *involvement* specifies the involvement of the participants in the event, and has the values *0, 1,* > 1*, all* or *undefined*, which were stipulated by us. Finally, with the purpose of representing nominal anaphora's mechanisms, we use the *objectal links* provided by ISO-24617-9, to identify the following relations: *objectal identity, part of, member of* and *referential disjunction*.

The third annotation layer is relative to semantic role labeling, and is grounded on ISO-24617-4. In this case, we integrated in our framework thematic relations between event structures and participant structures.

3.2. Spatial layer

This section describes the process of adding a spatial layer using ISO 24617-7 to the annotation scheme described previously.

In a nutshell, ISO-Space, in accordance with the general principles of ISO 24617-6 (ISO-24617-6, 2016), proposes the use of a set of entity structures, and a set of link structures, with attributes and values, which are employed to annotate spatial information. In ISO-Space, entity structures comprise participants, eventualities, and measures. Participant structures include two kinds of entities: (i) locational entities, such as places, paths (roads, rivers, etc.) and event-paths (motional trajectories), and (ii) non-locational entities that are involved in spatial relations. Eventuality structures include motion (i.e., events describing the displacement of an entity, or the change in some property or conformation of an object), and non-motion events. Measure structures provide quantitative information regarding some dimension of spatial entities or regarding some relation between spatial entities. As for link structures, they consist of four types: (i) qualitative spatial links, which express static relations between regions; (ii) orientational links, which express the spatial disposition or direction of an object; (iii) movement links, which relate an entity that moves to the trajectory that is being followed by that entity in the course of a motion event; and (iv) measure links, which connect an object to its measure. Measure structures and eventuality structures are always anchored in markables, i.e., word/expressions occurring in the

[2]For a more detailed account about the choices and their justification related to the annotation scheme described in this subsection, see Silvano et al. (2021).

text. However, participant structures are not always anchored in markables. In fact, ISO-Space (2020 version) proposes the use of "non-consuming" tags, i.e., tags without a markable, namely in the case of the entity structure *event-path*. As for link structures, they relate participant structures with each other.

From this ISO, we chose the following tags, attributes, and links: (i) for entities: *place, path, non-locational spatial entity, spatial relation, motion and non-motional eventualities* and *measure*, leaving out *event-path*; (ii) for links: *qualitative spatial link, movement link*, and *measure link*, leaving out *orientational link*.

Starting with entities, the *non-locational spatial entity* is subsumed, in our annotation scheme, into the entity structure *participant*, which has several attributes for the tag *type* (cf. section 3.1; e.g. PER, ORG). There was no need to formally introduce the notion of *non-locational spatial entity* in the existing annotation scheme. Although *place* tag was already in the list of attributes of the tag *type* of the entity structure *participant*, according to ISO 24617-7 one needed to distinguish the different types of places. The solution that we implemented included, in the list of attributes of the tag *type* of *participant structure*, the attributes of *place type* stipulated by ISO 24617-7 that were relevant to annotate the news articles. The restricted set of attributes for *place type* that we decided to add to the list of participants *type* already present in our annotation scheme is the following: *Pl_water*, *Pl_celestial*, *Pl_mountain*, *Pl_civil*, *Pl_country*, *Pl_mount_range*, *Pl_capital*, *Pl_region*, and *Pl_state*. In fact, we merged participant structures types proposed by ISO 24617-7 with tags that were already part of our annotation scheme, and that came from ISO 24617-9. Example (1) illustrates a locative participant structure.

(1) o bairro da Pasteleira Nova
the neighbourhood of-the Pasteleira Nova
the Pasteleira Nova neighbourhood
participant: lexicalHead=noun; individuationDomain=individual; type=plCivil; involvement=1

As for *paths*, the option was not to give any attribute of the ISO 24617-7 and simply to insert *path* also as an attribute of the tag *type* of participants. Overall, the already existing participant entity structures were enriched with the specification of different types of locative participants. There was no need to create new entity structures for locative participants.

Motions and *non-motional eventualities* were also easily included in our annotation scheme, as there were already entity structures for events. Therefore, besides the existing tags and attributes of the event structure defined by ISO24617-1 (cf. figure 1), we only added to the entity structure *event* the attribute *motion* with the tags *motion_literal, motion_fictive, motion_intrChange* (inherited from the attribute *motionSense* of the tag *motion* in ISO 24617-7) and *non-motion*. These attributes allow us to distinguish changes in the location of the *Figure* (*motion_literal*) and changes in the *Figure*'s configuration (*motion_intrChange*), both of them corresponding to event-type situations, from most cases of fictive motion (*motion_fictive*), which correspond to state-like eventualities. So, the different kinds of motion that are envisaged by ISO 24617-7 can be properly identified (as literal, fictive, or intrinsic change) and be opposed to all non-motional eventualities. We do not use any of the remaining attributes of the *motion* tag in ISO 24617-7 (for instance, *motionType* or *motionClass*), because such level of detailed information is not pertinent to our project's purposes, at least for the time being. Hence, for instance, all eventualities that receive the tag *motion* are interpreted as if there was a *motionClass=move*, that is, that specific eventuality is associated to a trajectory (or *event-path*) (being the subeventive structure underspecified; cf.(Pustejovsky, 2017)). Example (2) illustrates an event structure.

(2) os camponeses saíam das lavras
the peasants were leaving from-the fields
The peasants were leaving the fields
saíam: event: class=state; eventType=state; pos=verb; tense=past; aspect=progressive; polarity=pos; movement=motionLiteral

Finally, in order to encompass all sorts of spatial information, it was required to add two entity structures to our annotation scheme: *spatial relations*, and *measures*. *Measures* was introduced to deal with measurement relations, most of them corresponding to distances. Following ISO 24617-7, the markables that are identified as measures are further manually annotated with information regarding *measure value*, *measure unit* and *measure modification* (whenever required). Example (3) represents a measure structure.

(3) uma faixa de proteção inferior a 250 metros
a band of protection inferior to 250 meters
a protection zone under 250 meters
inferior a 250 metros: value=250; unit=metro; mod=inferior a

Spatial relations is associated with preposition and adverb-like markables. As defined by ISO 24617-7, this relation connects a *Figure* to a location or trajectory. *Spatial relation* in our annotation scheme only uses the tag *sRelation type*, with only two values: *topological* and *path defining*. Other values that are proposed by ISO 24617-7, namely *directional* and *topoDirectional*, appear to be unnecessary for different reasons. Regarding *Directional* value, it seems to be residual in the news texts that constitute our corpus. In fact, in the analysed news texts, there were no examples of directional prepositions, such as *em frente de* (*in front*

of), which justifies our option of not using *directional spatial relations*, nor *orientational links*, for now. In our corpus, *spatial relations* are mainly of the *topological* type, performing a vague location, which typically corresponds to *disjunction of tangential proper part* and *non-tangential proper part*. With respect to *topo-Directional*, this *sRelation type* does not correspond to any specific preposition or adverb in European Portuguese. The English preposition *on*, which codifies this relation, in European Portuguese is translated as *em*, which subsumes English prepositions *in* and *on*.

The *topological* value uses the nine attributes proposed by ISO 24617-7: *disconnected; externally connected; partial overlap; equal; tangential proper part; tangential proper part inverse; non-tangential proper part; non-tangential proper part inverse*; and *disjunction of tangential proper part* and *non-tangential proper part*. Nevertheless, the utility of using *tangential proper part inverse* or *non-tangential proper part inverse* is debatable, as the annotator can establish the relation from a participant structure A (the *Figure*) to a participant structure B (the *Ground*), or the other way round.

As for *path defining* value, we use the three attributes put forward by ISO 24617-7: *start*, *end* and *mids*. However, and this is a relevant adaptation of ISO 24617-7 to our annotation scheme, we have also merged the *goal defining* value with the attributes of the *path defining* value, so the tag *path defining* effectively has four attributes. The rationale for this option is related, for instance, to the fact that there is a frequently used preposition in European Portuguese, *para*, which can easily exhibit two readings: a mere directional reading (that can be translated by *towards*), and a goal reading (equivalent to English *to*). Accordingly, the preposition *para* can occur with both readings and, in most cases, the choice between the two is context-dependent. In other words, the goal preposition *para* easily allows also non-culminating readings, i.e., of canceling the event's culmination, when the Figure reaches the intended destination (Leal et al., 2018), contrary to other prepositions, such as *a* (*to*) and *até* (*up to*). Hence, in the presence of a sentence with *para*, it is easier for the annotator to mark this preposition with a *path-defining* value, with either the attribute *end*, or *goal-defining*. Example 4 represents a spatial relation structure.

(4) os camponeses saíam das lavras
the peasants were leaving from-the fields
The peasants were leaving the fields
d(as): spatialRelation: pathDefining=start

An important difference between our annotation scheme and ISO 24617-7 is the absence of *event-path* in the entity structures. In fact, there is a basic incompatibility problem between ISO 24617-7 and our annotation scheme: the former proposes entity structures that do not correspond to markables, whereas the latter presupposes that all entity structures are always linked to markables in the text. In order to overcome this major incompatibility, and to be able to accommodate the notion of *event-path* within our annotation scheme, since *event-paths* have triggers that correspond to motion verbs, we use those verbs as markables for building event structures and we associate to those event structures the relevant links. We do not include, in these event structures, any of the attributes of the *event-path* tag, namely *eventPath start*, *eventPath end*, or *eventPath mids*, and, instead, we resort to link structures to fill in this information. Adopting this solution means not discarding completely the notion of *event-path*. In fact, we make use of some of the information codified in its attributes to represent the spatial information present in our corpus.

The main reason not to include non-consuming tags, and just annotate lexical material concerns limitations of BRAT, the annotation tool that we use. Although this may not be the most valid reason, the following is more compelling. As mentioned before, the current proposal aims at integrating a spatial layer into a multilayer annotation scheme, which harmonizes parts 1, 4 and 9 of ISO 24617. In all of these parts, tags have always markables. So, if we wanted ISO-Space to fit the existing annotation scheme, we had to exclude non-consuming tags, at least at this stage of the project.

As far as the link structures are concerned, we deemed it best to use only three types: *qualitative spatial link, movement link*, and *measure link*. Since the information captured by the *orientational link* was not necessary to annotate our corpus, we discard it. The *qualitative spatial link* uses only two required attributes - *Figure* and *Ground*. The *relType* attribute, which is also required, is already present in the *spatial relation* structure. Accordingly, after the annotation of a markable with a *spatial relation* structure with a *topological* value, the annotator must connect that structure with a participant structure or an event structure by means of a *Figure* link, and with a (locative) participant structure by means of a *Ground* link. Example (5) illustrates the qualitative spatial links.

(5) Um homem de 20 anos foi detido pela PSP no bairro da Pasteleira Nova
A man of 20 years was arrested by-the PSP in-the neighbourhood of-the Pasteleira Nova
A 20-year-old man has been arrested by the PSP in the Pasteleira Nova neighbourhood.
QSLINK-Ground=n(o), o bairro da Pasteleira Nova
QSLINK-Figure=n(o), detido

The *movement link* was substantially adapted because, as mentioned earlier, *event-paths* are not part of our annotation scheme. Therefore, we took the *moveLink* trigger as the markable from where different relations that correspond to different attributes of the tag *moveLink* are established. So, in the presence of

an event structure tagged as *motion*, the annotator must connect that event structure: (i) to a participant structure by means of the *Figure* link (corresponding to the *moveLink figure*), (ii) to a spatial relation structure by means of the *spatialRelation* link, and (iii) to a (locative) participant structure with a *targetSpatialRelation* link (corresponding to the *moveLink ground*). These two links (*spatialRelation* link and *targetSpatialRelation* link) are not in ISO 24617-7. We added them to deal with the problem of not having the non-consuming tag *event-path*. The solution that we envisioned includes the following: the combination of (i) a markable that is a motion event with (ii) a participant that is a *Figure* of that event, (iii) a *spatial relation* that defines the initial, medial or final part of a trajectory, or just the direction of the motion, and (iv) the *Ground* of the spatial relation. Therefore, it is the combination of a motion event structure with a directional spatial relation and its *Ground* that functions as a substitute of the *event-path* notion of ISO 24617-7. Once again, this approach allow us to keep part of the information represented by this entity structure. In (6), we exemplify this solution.

(6) os camponeses saíam das lavras
the peasants were leaving from-the fields
The peasants were leaving the fields.
moveLink: figure=saíam, os camponeses moveLink: spatialRelation=saíam, d(as) moveLink: targetSpatialRelation= saíam, as lavras

Finally, the measure link uses the *relType* attributes (*distance, length, width, height* and *generalDimension*) and it connects an event structure or a participant structure to a measure structure by a *Ground* link (following ISO 24617-7, as "entities to which the measure value applies"). Example (7) shows how this annotation is performed.

(7) uma faixa de proteção inferior a 250 metros
a band of protection inferior to 250 meters.
A protection zone under 250 meters.
MLINK-Width=inferior a 250 metros, uma faixa de proteção

Figure 1 illustrates our annotation scheme.
In what follows, we show how Text2Story annotation scheme can be applied to represent the interpretation of example (8).

(8) Um homem de 20 anos foi detido pela PSP no [3] bairro da Pasteleira Nova
A man of 20 years was arrested by-the PSP in-the neighbourhood of-the Pasteleira Nova
A 20-year-old man has been arrested by the PSP in the Pasteleira Nova neighbourhood.

[3]In this example, the preposition *em* contracts with the definite article, as it happens frequently in Portuguese

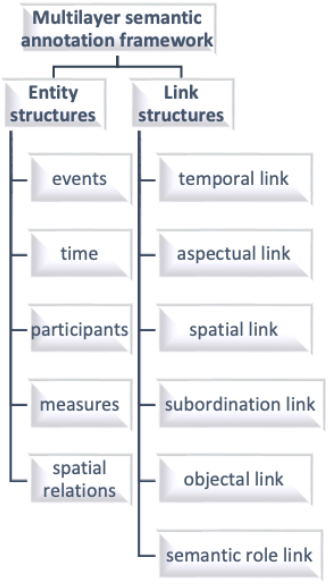

Figure 1: Text2Story multilayer annotation scheme.

STRUCTURES
um homem de 20 anos: participant: lexicalHead=noun; individuationDomain=individual; type=per; involvement=1
a PSP: participant: lexicalHead=noun; individuationDomain=individual; type=org; involvement=1
o bairro da Pasteleira Nova: participant: lexicalHead=noun; individuationDomain=individual; type=plCivil; involvement=1
detido: event: class=occurrence; eventType=transition; pos=verb; tense=past; aspect=perfective; polarity=pos; movement=nonMotion
n: spatialRelation: topological=disjunction-TTP-NTTP

LINKS
QSLINK-ground=n, o bairro da Pasteleira Nova
QSLINK-figure=n, detido
SR-patient=detido, um homem
SR-agent=detido, a PSP

From Example 8, it is possible to build a Discourse Representation Structure (DRS) as depicted in Figure 2. The event is described as a First Order Logic formula. The participants and spatial relations are represented as constants. Using this kind of notation makes it possible to reason over the events formulas, which is useful, for instance, to infer new relations between such elements. The Brat2Viz tool (Amorim et al., 2021) employed DRS as an intermediate language to aid in building a visual representation of the narrative components. The input of this tool is a human annotation text file, then a corresponding DRS is generated, and finally, a visual representation of the annotation is produced.

```
» EVENTS
# T2 (detido) -> a
# FOL:  exists a.(relationRole(qsfigure,T2) &
relationRole(patient,T2) & relationRole(agent,T2) &
event(a))
# DRS:  ([a],[relationRole(qsfigure,T2),
relationRole(patient,T2), relationRole(agent,T2), event(a)])

» ACTORS
# T1 -> Um homem de 20 anos
# T3 -> a PSP
# T5 -> o bairro da Pasteleira Nova

» SPATIAL-RELATION
# T4 -> n

» RELATIONS
# T4 - qsground - T5
# T4 - qsfigure - T2
# T2 - patient - T1
# T2 - agent - T3
```

Figure 2: DRS text built from the annotation of example 8

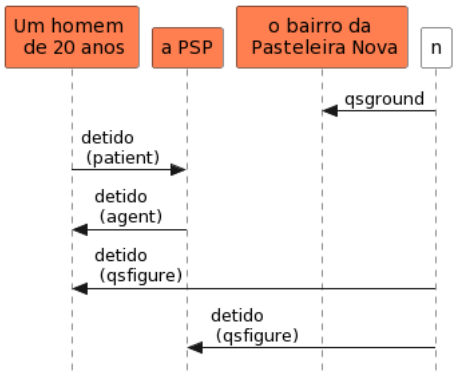

Figure 3: Message Sequence Chart for example (8)

Brat2Viz uses the Message Sequence Chart (MSC) diagram to represent the annotations, as in Figure 3.

Currently, in Text2Story project, the only options for the visualization of the annotation are either (MSC), or knowledge graphs (KG). A foreseen follow-up of this project is the improvement of visualization methods, namely with Venn diagrams and flowcharts. Both will put the events as the narrative element to follow in a timeline. In an MSC, participants are the elements in a timeline of events, and in a KG there is no timeline. Arranging the events in sequence to follow in a diagram also helps the reader to get a better perception of their relative position in the timeline.

At the moment, we are mainly concerned with defining a sound and comprehensive annotation scheme that can deal with the semantic content of news texts in a way that can be used in the development of visualization algorithms. As we are finishing up the annotation scheme, no inter-annotator agreement tasks were performed yet.

To sum up, the introduction of a spatial annotation level in our annotation scheme involved, in addition to the inclusion of some spatial-specific entity structures and link structures, the modification of pre-existing entity structures: (i) the participant structure attribute *type* was enriched with a value *path*, which corresponds to the notion of spatial entity *path* in ISO-Space, and with most of the values of the *placeType* attribute of the tag *place* of ISO-space; (ii) event structures were enriched with the attribute *motionSense* from the *motion* tag of ISO-Space. Simultaneously, the integration of the spatial layer enabled the removal of some (redundant) semantic role links, those that bear spatial nature: *location*, *initialLocation*, *finalLocation*, *path* and *distance*. In fact, as expected, spatial annotation was substantially improved with the replacement of locative semantic roles with spatial-specific annotation. For instance, there is now the possibility of expressing different relations that were previously included in the semantic role of *finalLocation*. They can now be marked as *end* or *goalDefining*, which gives rise to different entailments concerning the Figure's location at the end of the event. Additionally, measure structures and links can be utilized for a better characterization of participants: the semantic role link *amount* can was replaced by measure structures and links, with more information concerning *type, value* and *modification*, and with the possibility of being used not only to link a participant to an event, but also to describe a participant (cf. example (3) and (7)).

3.3. Some problems

In this section, we pinpoint some shortcomings of the implementation of ISO-Space to our annotation scheme.

As already stated, some problems arose when trying to combine ISO 24617-7 with our annotation scheme, which resulted from the harmonization of parts 1, 4 and 9 of ISO 24617. One of the major problems that we encountered was the incompatibility between the existence of "non-consuming" tags in ISO 24617-7 and our project's principle of annotating only lexical material occurring in text (also a principle of parts 1, 4 and 9 of ISO 24617), evidenced by the *event-path* issue that we discussed before. A similar issue is related to a very frequent structure in our corpus exemplified in (9).

(9) O caso deu-se na freguesia de Refogos de Basto, concelho de Cabeceiras de Basto, distrito de Braga.

The case happened-itself in-the parish of Refogos de Basto, municipality of Cabeceiras de Basto, district of Braga.

The case took place in the parish of Refogos de Basto, municipality of Cabeceiras de Basto, district of Braga.

In order to conduct spatial annotation, the annotator tags the spatial location of *the case* in a place called *Regofos de Basto*. That spatial location is expressed by the preposition *em* (*in*), heading the PP *na freguesia de Refogos de Basto*. The journalist adds further information about that place, as it is often done when the place in question is not part of the speakers "common knowledge". Typically, the journalist expands the sentence with appositive modifiers locating it in

a bigger (and more well-known) location. In some cases, those appositive phrases correspond to PPs headed by *em* (*in*). However, in most cases, as in (9), the preposition is absent (*concelho de Cabeceiras de Basto* and *distrito de Braga*). This poses a challenge to our annotation scheme, because a preposition (or an adverb) is obligatory to mark the spatial relation. The solution that we implemented was the following. Firstly, as our annotation scheme has different levels of annotation, and since one of those levels is the referential level, we can connect *na freguesia de Refogos de Basto*, annotated as a participant of the type *placeCivil*, to the participant *concelho de Cabeceiras de Basto*, also a participant of the type *placeCivil*, by means of a *partOf* objectal relation. The same sort of connection can be established between the participant *concelho de Cabeceiras de Basto* and the participant *distrito de Braga*, of the type *placeCivil*, as well. Secondly, since all cases of locations without preposition correspond to places that are a part of other locations, we decided to insert an inference rule in the DRS, which acts as an intermediate language to generate visualizations (cf. section 3.3). This inference rule (cf. (10)) states that whenever two locative participants are connected by a *partOf* relation, there is a topological spatial relation with the value *IN* (as defined by ISO 24617-7) that takes the first location in the text as its *Figure* and the second location as its *Ground*.

(10)

$$\forall x \forall y (LOCATION\, x \wedge LOCATION\, y \wedge PART\_OF\, x, y) \rightarrow \exists z (SPATIAL-RELATION\_IN\, z \wedge FIGURE\, x, y \wedge GROUND\, y, z)$$

In the process of harmonizing the different parts of ISO 24617, we have also observed some incongruities when we compare the parts of ISO 24617 that we used in our framework. We have already referred to one of them: the existence of non-consuming tags in ISO-Space, namely the *event-path* tag, contrary to happens in the other relevant parts. There are also some differences between the aforementioned parts of ISO 24617 regarding the level of granularity required to annotate the same type of entities, namely eventualities. As a matter of fact, ISO 24617-1 has only three values for the attribute *type* of event structure: *state*, *process*, and *transition*. However, in ISO-Space, the attribute *motionClass* of the *motion* tag has several values (*move*, *moveExternal*, *moveInternal*, *leave*, *reach*, *cross*, *detach*, *hit*, *follow*, *deviate*, and *stay*), which are related to the different subevent structures, and their consequences regarding the location of the *Figure*. Motion events are not the only ones exhibiting subevent structure, so it would maybe be congruent if ISO 24617-1 also contemplated a similar proposal of subeventive analysis for non-motion eventualities.

A final remark concerning ISO-Space more practical matters. This document could benefit (in clarity) if more definitions and examples were given regarding attributes and values. In fact, the scarcity of examples, together with the fact that they are all in English, poses some challenges when one tries to use it in the annotation of texts in languages that differ from English in distinct (and, sometimes, understudied) aspects.

4. Conclusion and Future Work

Representing and extracting information from texts is a challenging task. In our project, we aim to create visualizations from information extracted from news texts. To do so, we initially designed a multilayer annotation scheme comprising three levels of annotation (times and events, participants, and semantic roles) combining three parts of ISO SemAF (ISO-24617-1; ISO-24617-4; and ISO-24617-9). In this paper, we describe the efforts of improving our annotation scheme with spatial information by adding a fourth layer of annotation that uses ISO-Space (ISO-24617-7, 2020). To achieve a coherent annotation architecture, we had to select tags, attributes and values from ISO-Space that were relevant to the project's purposes, and to insert them into the pre-existing annotation scheme. This required not only adapting some of the proposals of ISO-Space, but also deleting and/or adjusting some of the tags/attributes of the first version of our annotation scheme.

Despite the fact that we discarded at this stage some of the ISO-Space tags and attributes (eg. type of motion), because either our dataset or annotation scheme did not require them, or because we needed to alleviate the annotation load within a four layer framework, overall the annotation as is performs in a satisfactory manner, being able to represent the pertinent information from our dataset.

In the future, it is our intention to validate Text2Story annotation scheme by using inter-annotator agreement and by generating different types of visualizations from narratives. Furthermore, more spatial information will be considered to enable the mapping of the locations referred in the texts to maps. This task will require adding more attributes to the *placeType*, such as *place latLong*. We will also conduct a detailed study about the specificities of European Portuguese regarding the expression of spatial information and assess to what extent ISO-Space can account for their representation. Since we recognize that the expressive power of ISO-Space goes beyond the annotation of news texts to extract where actions happen, we plan as well to expand the initial corpus of Text2Story project to encompass pre-processed texts that include a large array of spatial relations (for instance, texts depicting places, or travel descriptions), which will allow us to test all potentialities of ISO-Space.

5. Acknowledgements

This research is funded by the ERDF – European Regional Development Fund through the North

Portugal Regional Operational Programme (NORTE 2020), under the PORTUGAL 2020 and by National Funds through the Portuguese funding agency, FCT - Fundação para a Ciência e a Tecnologia within project PTDC/CCI-COM/31857/2017 (NORTE-01-0145-FEDER-03185).

6. Bibliographical References

Amorim, E., Ribeiro, A., Cantante, I., Jorge, A., Santana, B., Nunes, S., Silvano, M. d. P., Leal, A., and Campos, R. (2021). Brat2viz: a tool and pipeline for visualizing narratives from annotated texts. In *Proceedings of Text2Story-Fourth Workshop on Narrative Extraction From Texts held in conjunction with the 43rd European Conference on Information Retrieval (ECIR 2021)*.

Estarrona, A. and Aldezabal, I. (2018). Towards a Spatial Annotation Scheme for Basque based on ISO-Space. In Francesco Mambrini, et al., editors, *Proceedings of the Second Workshop on Corpus-Based Research in the Humanities (CRH-2)*, pages 13–21. Gerastree Proceedings, GTP 1, Dept. of Geoinformation, TU Wien.

Gaizauskas, R. and Alrashid, T. (2019). SceneML: A Proposal for Annotating Scenes in Narrative Text. In Harry Bunt, editor, *Proceedings of the 15th Joint ACL - ISO Workshop on Interoperable Semantic Annotation (ISA-15)*, pages 13–21.

Gritta, M., Pilehvar, M. T., and Collier, N. (2018). Which Melbourne? Augmenting Geocoding with Maps. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1285–1296, Melbourne, Australia, July. Association for Computational Linguistics.

Henlein, A., Abrami, G., Kett, A., and Mehler, A. (2020). Transfer of ISOSpace into a 3D Environment for Annotations and Applications. In *16th Joint ACL - ISO Workshop on Interoperable Semantic Annotation PROCEEDINGS*, pages 32–35, Marseille, May. European Language Resources Association.

ISO-24617-1. (2012). Language resource management- semantic annotation framework (semaf) - part 1: Time and events (semaf-time, iso-timeml). Standard, Geneva, CH.

ISO-24617-4. (2014). Language resource management- semantic annotation framework (semaf) - part 4: Semantic roles (semaf-sr). Standard, Geneva, CH.

ISO-24617-6. (2016). Language resource management- semantic annotation framework (semaf) - part 6:principles of semantic annotation (semaf principles). Standard, Geneva, CH.

ISO-24617-7. (2020). Language resource management-semantic annotation framework (semaf) - part 7: Spatial information. Standard, Geneva, CH.

ISO-24617-9. (2019). Language resource management- semantic annotation framework (semaf) - - part 9: Reference annotation framework (raf). Standard, Geneva, CH.

Kordjamshidi, P., Van Otterlo, M., and Moens, M.-F. (2010). Spatial Role Labeling: Task Definition and Annotation Scheme. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta, May. European Language Resources Association (ELRA).

Leal, A., Oliveira, F., and Silvano, P. (2018). Movimento ascendente e descendente em português europeu: os casos dos verbos subir e descer. In *Para Óscar Lopes: Estudos de Linguística*, pages 99–114. Afrontamento.

Lee, K., Fang, A. C., and Pustejovsky, J. (2011). Multilingual verification of the annotation scheme iso-space. In *2011 IEEE Fifth International Conference on Semantic Computing*, pages 449–458.

Lee, K. (2016). An abstract syntax for ISO-Space with its ¡moveLink¿ reformulated. In Harry Bunt, editor, *Proceedings of the LREC 2016 Workshop, 12th Joint ACL–ISO Workshop on Interoperable Semantic Annotation (isa-12)*, pages 28–37. Association for Computational Linguistics.

Mani, I., Hitzeman, J., Richer, J., Harris, D., Quimby, R., and Wellner, B. (2008). SpatialML: Annotation Scheme, Corpora, and Tools. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco, May. European Language Resources Association (ELRA).

Pustejovsky, J. and Stubbs, A. (2012). *Natural Language Annotation for Machine Learning - a Guide to Corpus-Building for Applications*. O'Reilly.

Pustejovsky, J. and Yocum, Z. (2013). Capturing motion in iso-spacebank. In *ACL 2013*.

Pustejovsky, J. and Yocum, Z. (2014). Image Annotation with ISO-Space: Distinguishing Content from Structure. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 426–431, Reykjavik, Iceland, May. European Language Resources Association (ELRA).

Pustejovsky, J., Moszkowics, J., and Verhagen, M. (2011). ISO-Space: The Annotation of Spatial Information in Language. In *Proceedings of the Sixth Joint ISO - ACL SIGSEM Workshop on Interoperable Semantic Annotation*, pages 1–9. Association for Computational Linguistics.

Pustejovsky, J., Moszkowics, J., and Verhagen, M. (2012). The current status of ISO-Space. In Harry Bunt, editor, *Proceedings of the Seventh Joint ACL–ISO Workshop on Interoperable Semantic Annotation (ISA-7), SRSL3 Workshop on Semantic Representation for Spoken Language, and I2MRT Workshop on Multimodal Resources and Tools*, pages 23–30.

Pustejovsky, J., Kordjamshidi, P., Moens, M.-F., Levine, A., Dworman, S., and Yocum, Z. (2015). SemEval-2015 Task 8: SpaceEval. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 884–894, Denver, Colorado, June. Association for Computational Linguistics.

Pustejovsky, J., Lee, K., and Bunt, H. (2019). The Semantics of ISO-Space. In Harry Bunt, editor, *Proceedings of the Fifteenth Joint ACL - ISO Workshop on Interoperable Semantic Annotation (ISA-15)*, pages 46–53.

Pustejovsky, J., (2017). *ISO-Space: Annotating Static and Dynamic Spatial Information*, pages 989–1024. Springer Netherlands, Dordrecht.

Silvano, P., Leal, A., Silva, F., Cantante, I., Oliveira, F., and Jorge, A. (2021). Developing a multilayer semantic annotation scheme based on ISO standards for the visualization of a newswire corpus. In *Proceedings of the 17th Joint ACL - ISO Workshop on Interoperable Semantic Annotation*, pages 1–13, Groningen, The Netherlands (online), June. Association for Computational Linguistics.

Stenetorp, P., Pyysalo, S., Topić, G., Ohta, T., Ananiadou, S., and Tsujii, J. (2012). Brat: a Web-based Tool for NLP-Assisted Text Annotation. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 102–107, Avignon, France, April. Association for Computational Linguistics.