CHAPTER 11

Discovering Spatial Referencing Strategies in Environmental Narratives

Simon Scheider
Department of Human Geography and Spatial Planning, Utrecht University, the Netherlands

Ludovic Moncla
Univ Lyon, INSA Lyon, CNRS, UCBL, LIRIS, UMR5205, F-69622, France

Gabriel Viehhauser
Department of Digital Humanities, University of Stuttgart, Germany

In an environmental narrative, authors describe their perception and experience of a landscape in a form that enables the reader to follow their track in mind. The narrative resembles a consecutive set of snapshots of space viewed from a particular angle, either from an imagined ego on the track or from other locations specified relative to identifiable landmarks. This allows a reader to embed the authors' journey into the landscape, even if this landscape or journey are merely imagined (Tuan, 1991).

A key problem in analysing and comparing such narratives is the ability to *geo-reference* the places referred to in the text (Scheider and Purves, 2013)

How to cite this book chapter:
Scheider, Simon, Ludovic Moncla, and Gabriel Viehhauser (2022). "Discovering Spatial Referencing Strategies in Environmental Narratives." In: *Unlocking Environmental Narratives: Towards Understanding Human Environment Interactions through Computational Text Analysis.* Ed. by Ross S. Purves, Olga Koblet, and Benjamin Adams. London: Ubiquity Press, pp. 211–232. DOI: https://doi.org/10.5334/bcs.k.

(cf. the chapter about spatio-temporal linking of narratives by Tim Baldwin in this volume). Geo-references not only allow texts to be put on a map but also to be segmented into situated scenes. Furthermore, places can be linked across documents, so that it becomes possible to track a given environment across different perspectives, diverse authors, or even different literary periods. For instance, suppose we would like to assess the environmental change of a landscape such as the *Moors in Scotland*, based on comparing narrative settings of historical novels with contemporary travel literature. To do this, we need to know when narrators talk about the same place.

One of the major methodical challenges is that environmental narratives seldom refer to the environment in terms of place names. Rather, spatial references are often *indirect*, that is, relative to perceived objects, and the difficulty is that automated *geocoding* tools and Natural Language Processing (NLP) strategies currently struggle with any references beyond place names (Purves et al., 2018; Chen, Vasardani, and Winter, 2018; Stock, 2014; Scheider and Purves, 2013). The diversity of strategies narrators have at their disposal to refer to a location have been subject of empirical study by cognitive linguists and anthropologists such as Levinson (Levinson, 2003) or Palmer (Palmer, 2002) across different language communities. A *frame of reference (FoR)* is a strategy for describing a given location relative to diverse sets of objects, including the perceiving ego or salient landmarks. Understanding this strategy is needed for a reader to comprehend the meaning of diverse locative expressions, ranging from egocentric ones such as 'the mountain in front of me' to allocentric ones such as 'the place where the river Ba flows into Loch Laidon', and from rather precise absolute references, such as 'ten miles north-east of Loch Laidon', to relatively vague descriptions without any directional hint, such as 'away from the waterway'. Frequently, authors of narratives also speak about a location only in a metaphorical way (Talmy, 1996), using *fictive motion* to move an imagined ego through a landscape. This is reflected, for example, in expressions such as 'the trail runs along the lake'. Also, temporal references can play an important role in spatial referencing (Tenbrink, 2011).

How can we discover FoRs in environmental narratives? The relevance of FoR for spatial referencing and geographic knowledge discovery has been known for a long time (Mark et al., 1999; Burenhult and Levinson, 2008). Qualitative models of spatial information were developed in the past precisely with an eye on such cognitive frames of reference (Clementini, Di Felice, and Hernández, 1997). From the viewpoint of geographic information retrieval (GIR) (Purves et al., 2018), the task of discovering FoRs in narratives has only been looked at sporadically in the past. It is apparent that this requires more than just building formal FoR models (Clementini, 2013), extracting parts-of-speech (PoS), spatial relation words without context (Stock and Yousaf, 2018), or the recognition of named entities (NER). What is needed includes, to the very least:

1. *Extracting those PoS* from a text that are needed for identifying the type of FoR
2. *Identifying the referencing strategy* (type of FoR + parameters) used by the speaker
3. *Georeferencing the parameters* used in the FoR
4. *Transforming the target location* into coordinate space, taking account of vagueness

While some research has recently been done to address the latter two challenges (Chen, Vasardani, and Winter, 2018; Scheider et al., 2018; Stock and Yousaf, 2018), the first two challenges about geoparsing are seldomly taken into focus (Moncla et al., 2014; Vasardani et al., 2012; Stock and Yousaf, 2018). In particular, it is still unclear which kinds of reference strategies need to be distinguished for environmental narratives, and to which degree they can be extracted from texts based on state-of-the-art geoparsing methods.

In this chapter, we illustrate how FoRs can be automatically discovered in Scottish narratives, and we test the quality of such discovery. We first explore a range of referencing strategies which occur in environmental narratives, without any pretence at completeness. We then assess how well both human annotators and geoparsers can be used to discover these strategies in three sample texts. Our goal is to support people interested in automated alignment and mapping of environmental narratives beyond place names. We will finally discuss to what extent the method is useful for this purpose.

11.1 Mountaineering in Scotland

As a literary basis for exploration, we selected two mountaineering texts which are narrative descriptions of a given landscape, namely *Rannoch moor* in Scotland. W. H. Murray's 1957 book *Undiscovered Scotland* talks about a hike through the moor in Chapter 17 'The Moor of Rannoch' (Murray, 2003). Fifty years later, in 2007, R. Macfarlane describes a similar trip through the moor in his book *The Wild Places* (Macfarlane, 2008, p. 73). While fictional texts do not necessarily aim at an explicit description of their settings, but rather evoke them implicitly (Viehhauser and Barth, 2017), *non-fictional travel literature* may be more likely to contain sophisticated references to actual landscape in a way that, we believe, reflects the diversity of spatial referencing in narrative texts. The texts serve us both as a source for discovering FoR diversity, as well as a source for evaluating the quality of annotation and geoparsing. For external validation of geoparsing, we used in addition a *fictional* text that describes a travel through Scotland, namely R. L. Stevenson's *Kidnapped*[1]. The intention is

[1] https://www.gutenberg.org/files/421/421-0.txt

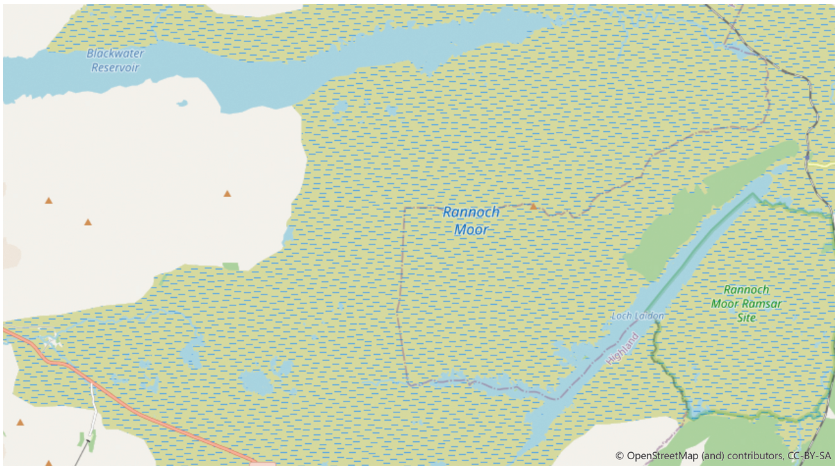

Figure 11.1: Rannoch Moor on Open Street Map (OSM). © Open Street Map contributors.

to identify the right referencing strategy used in these texts, in order to approximate the localisation of the many implicit places mentioned in the trip. While the latter are localisable only relative to toponyms such as 'Loch Laidon' and 'Loch Ba', the toponyms themselves can be easily georeferenced with standard geodata sources such as Open Street Map (Figure 11.1).

11.2 Referencing Strategies in Environmental Narratives

Within the sample mountaineering texts, we first explored the contained locative expressions (Herskovits, 1985), interpreting them in terms of known frames of reference or geometric strategies. In doing so, our intention was to capture the particular referencing strategy that might be used to technically reconstruct or approximate the referenced location in geographic coordinate space (c.f. Stock, 2014).

Types of FoR were proposed in Levinson (2003), Pustejovsky, Moszkowicz, and Verhagen (2011), Frank (1998), Clementini (2013) and Tenbrink (2011). Levinson's original set of frames (Levinson, 2003) mainly focuses on descriptions captured by *Euclidean coordinate axes*. This includes the construction of coordinate axes on some perceived ground object, and the localisation of figure objects along these axes. While exploring the texts, we quickly realised however that the richness of referencing strategies in mountaineering texts goes well beyond such strategies, exploiting also qualitative (Freksa, 1991), metric or topological relations, several ground objects, as well as metaphorical strategies including fictive motion (Talmy, 1996). A more comprehensive annotation framework in this respect is *ISO-Space* as proposed by Pustejovsky et al. (2011), which is based on Spatial-ML (Anderson et al., n.d) and includes motion events and corresponding paths, as well as mereo-topological relations, such as 'inside', 'outside', 'overlap' (Herring, Mark, and Egenhofer, 1994). While this approach acknowledges the relevance of diverse spatial and temporal concepts in spatial referencing, it treats qualitative spatial relations as a superclass of metric relations, and seems to be restricted to frames having coordinate axes on a single ground object.

For our purpose, we preferred a less strong spatial commitment. First, we suggest to regard the diversity of referencing strategies on a par, similar to Stock (2014), and not merely as parameters of the same kind of frame. This means we treat the different geometric bases of a referencing strategy as independent from each other. So, for example, using a referencing strategy based on distance does not necessarily imply any Euclidean axes or even a metric, because the assumption of Euclidean space are not needed to define a distance or a metric[2]. Furthermore, qualitative relations, such as 'inside', do not necessarily need to be interpreted as boundary cases of metric relations. And finally, we include the possibility of a multitude of ground objects. Second, we take seriously the observation that a given spatial referencing strategy, though in itself well defined, may be expressed in language in diverse and unforeseen ways, forcing us to take the context of an expression into account (Stock and Yousaf, 2018; Herskovits, 1985)[3]. For example, the preposition 'at' can have different meanings in different contexts (Vasardani et al., 2012). Our intention is therefore to test the quality of geoparsing rules which can take the context of an expression into account.

Based on the cognitive strategies we encountered in the two mountaineering text sources, we distinguish the following frame categories (cf. Figure 11.2): *Euclidean frames (EF), Zonal frames (ZF), Topological frames (TF), Linear construction frames (LCF) and Betweenness frames (BF)*. These categories directly reflect different ways how the corresponding locative expressions could be geometrically translated into the coordinate space of a map[4]:

1. [(EF)] *Euclidean frames* (*EF*) cover the well-known types as proposed in linguistic literature (Levinson, 2003), see also Scheider et al. (2018) and Frank (1998). These frames are used to denote target locations using axes in a coordinate system centered on a "ground" object, such that

---

[2] Formally, distances, metric spaces, topologies and Euclidean spaces can all be considered independent from each other (Worboys, 1996).

[3] For this very reason, Stock and Yousaf (2018) used a case-based learning approach.

[4] See Stock (2014) for a more comprehensive list of possible strategies. Note that many of our strategies can be mapped to this list.

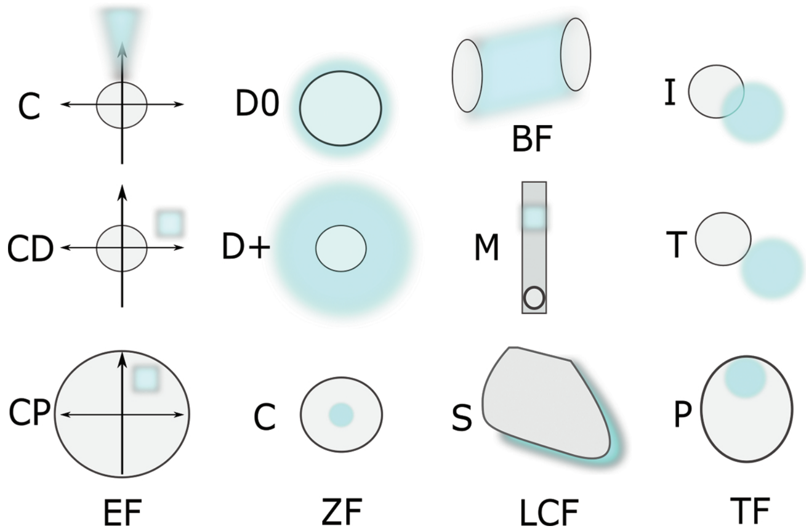

Figure 11.2: Illustration of types of referencing strategies used in environmental narratives. Target locations are indicated in turquoise, ground objects in grey. See text for the meaning of labels.

the main axis is oriented using some perceptual cue, such as a cardinal direction[5]. We only consider *allocentric* frames, where the ground object is some element of the landscape, and not the ego, since egocentric descriptions did not occur in our example texts. We distinguish the following subtypes: *cardinal direction* (EF C), *cardinal distance* (EF CD), *cardinal part* (EF CP), and *gravitation axis* (see illustrations in Figure 11.2).

"... ^(EFC)^ [the mountains]$_{(target)}$ [west of]$_{(orient)}$ [the moor]$_{(ground)}$"

2. ^(ZF)^ *Zonal frames* (*ZF*) are used to denote locations purely based on distance measured with respect to a ground object. We distinguish *Zonal frame with zero distance (ZF D0)*, *Zonal frame with distance modifier (ZF D+)*, and *Zonal frame center (ZF C)*.

[5] *Absolute* frames are oriented by cardinal directions, while *intrinsic* frames are oriented by the shape of the ground object (such as 'in front'). In environmental narratives, absolute frames seem the dominating orientation strategy, which is why we did not add other sub-types here.

"(*ZFD0*) At [Rannoch station]$_{(ground)}$ we stepped down from the train ..."

3. (*LCF*) *Linear construction frames* (*LCF*) are used to denote locations on a one-dimensional path. We distinguish two subtypes depending on how this path is constructed: *Linear construction using shape (LCF S)* and *Linear construction using movement (LCF M)*.

   "... (*LCFS*) along [the east side of]$_{(orient)}$ [Loch Ba]$_{(ground)}$"

4. (*BF*) *Betweenness frames* (*BF*) are used to denote places located between two given ground objects. For example, in the following sentence, the location of bog streams is described using two other landscape elements:

   "(*BF*) Between [the lochs]$_{(ground)}$ and [the peat hags]$_{(ground)}$ [bog streams wriggle]$_{(target)}$"

5. (*TF*) *Topological frames* (*TF*) are used to denote locations which stand in a mereo-topological relation to ground, such as 'inside', 'outside', 'touch', 'overlap' or 'part of' (Varzi, 1996). Some of these strategies in addition make use of figured features of the ground object in order to further specify the location, which we call *specific part*. We distinguish *Topological frame with touch (TF T) or intersect (TF I)*, *Topological frame with unspecific part (TF P)* and *Topological frame with specific part (TF SP)*.

   'a short and twisted (*TFT*) [river]$_{(target)}$ linked [Loch Ba]$_{(ground)}$ to [Loch Laidon]$_{(ground)}$'

What we spotted from looking through these examples is that identifying FoR is challenging in particular for the following reasons: (1) The different parameters of the FoR may be distributed across several sentences, making it hard to keep track of them across sentences with NLP. (2) The keywords that may indicate a given type of FoR or parameter can change considerably. For example, a topological relation may be indicated by the word 'link' instead of 'touch', and a vague distance may be given by the word 'well' instead of 'near'. (3) In the case of linear frames, 'fictive motion' (Talmy, 1996) is often required to construct a path in terms of an imagined trail. Thus simulation is required to localise these target locations inside a Geographic Information Systems (GIS). (4) Frames can be nested. Which means that identifying a parameter in parts of speech depends on first identifying another frame, adding considerable complexity to the search task.

11.3 Geoparsing Frames of Reference

In this section, we explain our approach for geoparsing FoRs. We chose a geoparsing framework which addresses word embeddings on different syntactic levels of a sentence, including motion and space keywords.

11.3.1 Perdido parsing rules

The *Perdido Geoparser* annotates different types of information such as named entities (with a special focus on spatial entities), extended named entities (ENE), spatial relations and motion expressions (Gaio and Moncla, 2019). An ENE consists of several overlapping phrase levels where each level is embedded in the previous one. This concept is based on the fact that a proper name can be categorised as pure or descriptive and that the descriptive expansion associated with a name can change the implicit type of the considered entity. The Perdido geoparsing rules have been developed for three romance languages (i.e., French, Spanish and Italian) and they have been used in several research projects to retrieve geographic information and to reconstruct itineraries from texts (Moncla et al., 2019; Moncla et al., 2016). Gaio and Moncla, 2019 argue that for a fine-grained task such as marking, classifying and disambiguating named entities, it is essential to consider the geo-semantic information expressed in the context in order to solve classification and disambiguation issues. Rules implemented in the Perdido Geoparser using cascades of finite-state transducers are a computational synthesis of previous works on how language expresses space and motion (Talmy, 1983; Vandeloise, 1986; Aurnague, 2011). This geoparsing task is based on a bottom-up strategy where each level of embedded entities of the ENE is marked from the pure proper name to the complete expression. Additionally, it can distinguish between two types of spatial entities: 'absolute' referring to standard spatial entities and 'relative' referring to spatial entities associated with spatial relations. More complex expressions involving motion verbs, spatial relations (e.g., spatial prepositions, topological relations, cardinal relations) and spatial ENE are also identified and classified. Examples 9. and 10. show the type of information annotated by Perdido.

9. there is a path that [runs]$_{(motion)}$ [along]$_{(rel)}$ [Loch Laidon]$_{(place)}$ [to]$_{(rel)}$ [Kingshouse]$_{(place)}$ [on]$_{(rel)}$ [the [Glencoe]$_{(place)}$ road]$_{(spatial\ ENE)}$.
10. [From]$_{(rel)}$ [Pitlochry]$_{(place)}$ the route would be either [by]$_{(rel)}$ [Struan]$_{(place)}$ and [Kinloch-Rannoch]$_{(place)}$ or [by]$_{(rel)}$ [Tummelside]$_{(place)}$ [to]$_{(rel)}$ [Tummelbridge]$_{(place)}$ and [straight through to]$_{(rel)}$ [Rannoch]$_{(place)}$.

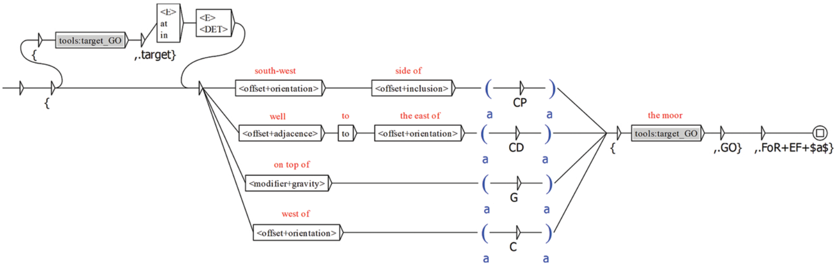

Figure 11.3: Transducer for parsing Euclidean frames.

11.3.2 FoR rules

In this work, we enriched a custom version of the Perdido Geoparser adapted for English texts. Our objective was to transform the existing rules and to add new ones in order to retrieve and classify FoR in environmental narratives. An example of an FoR parsing rule for Euclidean frames implemented using transducers is shown in Figure 11.3, and further ones are described below. A transducer is a local grammar defined as an automaton with an input and output alphabet. It is a type of finite-state machine that makes insertions, replacements and deletions in a text. The cascade of transducers of the Perdido Geoparser is implemented using the CasSys system (Friburger and Maurel, 2004) developed in the Unitex platform[6]. The grammar below mentions the main new transducers that have been added to the Perdido Geoparser for the recognition and classification of FoRs[7]. Transducers are implemented using graphs where each branch refers to a syntactic rule. For instance, for linear construction frames there are two patterns (as shown in the grammar below). Annotations are produced by the output alphabet of the transducers in brackets *{text.,semantic_tag}* with content after the comma referring to the semantic tags that will be transformed into XML elements. This grammar was developed based on a preliminary corpus analysis of the two mountaineering texts (see Section 11.1). The challenge is to build the most exhaustive and precise set of rules. Because of the ambiguity of natural language, however, the same syntactic rule may match different meanings, and may refer to FoRs but also to something else. This distinction cannot be expressed directly in the grammar because it needs external knowledge or context to understand the correct meaning. For this reason the developed grammar is incomplete and highly corpus and task dependent.

---

[6] https://unitexgramlab.org/

[7] For the full list, see our repository https://github.com/simonscheider/FoR.

TF SP = {Target} + modifier topological + GO

TF I = {Target} + modifier intersect + GO + modifier topological + GO

TF T = {Target} + modifier touch + GO + *and* + GO

TF P = {Target} + modifier inclusion + GO

EF CP = {Target + (*at*|*in*)} + modifier orientation + modifier inclusion + GO

EF CD = {Target + (*at*|*in*)} + modifier distance + modifier orientation + GO

EF G = {Target + (*at*|*in*)} + modifier gravity + GO

EF C = {Target + (*at*|*in*)} + modifier orientation + GO

ZF Dn = {Target} + modifier distance + GO

ZF Do = {Target} + modifier location + GO + *and* + GO

ZF C = {Target} + *in* + modifier central + GO

LCF S = {Target} + *along* + modifier orientation + modifier inclusion + GO

LCF M = {Target} + *from* + GO + *to* + GO

BF = {Target} + *between* + GO + *and* + GO

Target = spatial entity

GO = spatial entity + {separator + GO}

Several modifiers were used in our FoR grammar, which consist of lexicons or existing transducers executed at the beginning of the Perdido cascade and then already associated with semantic tags such as <offset+orientation>. These modifiers are expressed by keywords and spatial prepositions in the language. We consider nine types of modifiers: topological (*the mouth of*, *flows into*, …), intersect (*where*, …), touch (*linked*, …), inclusion (*in*, *inside*, *part of*, …), orientation (*north*, *south*, *north-east*, …), distance (*near by*, *away from*, *close to*, …), gravity (*on top of*, …), central (*middle of*, …), location (*at*, *where*, …).

11.4 Results and Quality of Geoparsing

In this section, we present and explain our validation results[8]. We start with inter-annotator agreement, and then discuss the quality of automatic parsing on the two mountaineering texts, as well as on an external text source that was not used for exploration and training.

[8] All raw resources can be accessed under https://github.com/simonscheider/FoR

11.4.1 Inter-annotator agreement

To assess the comprehensibility of our notion of FoRs, we performed an annotation task on Robert Macfarlane's 'The Wild Places' describing a hike through Rannoch moor (Macfarlane, 2008). For the purpose of this paper, we constrained ourselves to an exemplary approach that has to be expanded in the future, taking into account more texts.

In the vein of the annotation workflow outlined by Pustejovsky and Stubbs (2013), we understand the modeling process as an iterative cycle, in which concepts are tested empirically with the help of manual annotations that in turn serve as a base for a revision of the concepts. To assess the inter-subjectivity of the annotations we let different annotators annotate the same text and calculate their inter-annotator agreement.

For our first annotation round, we formulated guidelines that instructed the annotators to mark up the text with 'FrameOfReference', 'target' and 'groundObject' tags. 'FoR' tags were supposed to be classified according to our categories outlined above. Our first annotation round was performed by four annotators, amongst them one expert annotator from the project team. All annotations were carried out with the webAnno-tool as provided by the CLARIN-D web-service[9].

An error analysis of the first annotation round showed that the syntactic boundaries of the FoRs are very hard to define clearly enough, and thus hard to identify by all annotators, even if the FoR concepts themselves may have a clear definition. Annotators often had different ideas about where exactly a phrase containing a FoR would start or end, and therefore many correct identifications of a given frame type slightly overlapped within a sentence. Therefore, to report on the inter-annotator agreement, we abstain from calculating standard token-based kappa-metrics (Pustejovsky and Stubbs, 2013, pp. 126–134), but rather give the total number of FoRs annotated. Furthermore, we count every annotation that shows an overlapping match as an agreement. In total, the four annotators of the first annotation round classified 40 phrases as FoRs. In 10 cases, all of the annotators agreed, seven phrases were unanimously annotated by three annotators, 11 phrases by two annotators and in 12 cases only one out of the four annotators classified a phrase as a FoR.

To establish a more stable reference for comparison, we revised the expert annotation and defined it as a gold standard (which we later on also used for the comparison of the automatic detection of FoRs in Section 11.4.2). In total, the gold standard features 29 FoRs. Table 11.1 shows the precision and recall of three annotators of the first round compared to the gold standard.

A more in-depth analysis of the errors revealed that a high percentage of disagreement resulted from an insufficient distinction of our notion of FoR to the

---

[9] https://webanno.sfs.uni-tuebingen.de/

Annotator 1	Precision	64%
	Recall	79%
Annotator 2	Precision	70%
	Recall	55%
Annotator 3	Precision	59%
	Recall	34%

Table 11.1: Precision and recall of the first annotation round compared to the gold standard.

Annotator 4	Precision	64%
	Recall	62%
Annotator 5	Precision	65%
	Recall	72%
Annotator 6	Precision	56%
	Recall	79%

Table 11.2: Precision and recall of the second annotation round compared to the gold standard.

concept of motion: For our model, we want to exclude expressions that indicate a motion without referring to a place, whereas often annotators tended to annotate motion phrases as linear construction frames. Therefore, we revised our *guidelines*[10] on this behalf and explicitly told annotators not to tag motion expressions.

On the basis of these guidelines, we performed a second annotation round with three more annotators. However, even though the quality slightly improved in this round, the annotators still showed a significant amount of disagreement. In total, 52 phrases were classified as FoRs. In 15 cases all four annotators (including the gold standard annotation) agreed, in six cases three out of the four, in 12 cases only two and in 19 cases annotators stayed on their own in their decision to annotate a phrase as FoR. Table 11.2 shows precision and recall in comparison to the gold standard.

The rather low agreement shows that FoRs are a difficult concept that still needs clarification. Furthermore, it is apparent that more training is needed for interpreting FoRs in texts in a consistent way. A further refinement of the annotation guidelines remains a task for future work.

[10] Available under http://geographicknowledge.de/pdf/AnnGuiEnv.html.

11.4.2 Quality on mountaineering texts

For our parser experiments, we processed both McFarlane's 'The Wild Places' and Murray's 'Undiscovered Scotland' with the proposed FoR parsing rules implemented in the custom version of the Perdido Geoparser (as described in Section 11.3.2). In order to assess the quality of the automatic annotation, we compare the results with the gold standard annotation (see Section 11.4.1). For this purpose, we use three different metrics: precision, recall and Slot Error Rate (SER)(Makhoul et al., 1999). SER takes into account different types of errors: substitutions (S), insertions (I) and deletions (D). Substitution errors are of three kinds: wrong boundaries identification (B), wrong classification (T) and both (CT). Insertion errors refer to false positives (i.e., entities identified by the system that do not exist in the gold standard) and deletions errors refer to false negatives (i.e., entities existing in the gold standard that are not identified by the system). In addition to the precision and recall measures, the SER allows us to consider not only the identification of the expressions but also their classification.

Our gold standard for evaluation is composed of 69 FoR expressions, 76 ground objects and 10 target entities. Table 11.3 shows the number of each error type and the scores for different evaluation measures over the MacFarlane and Murray chapters. We notice that the distribution of all types of FoR expression is not homogeneous in our gold standard. For instance, we only have two betweenness FoR but 10 Euclidean or 22 zonal FoR. This implies that betweenness FoR results will not be representative and meaningful. However, one interesting observation is that all types of FoR have a rather high precision score.

Another interesting result is that the greatest number of errors refers to deletions (i.e., false negative) and implies a rather bad recall score (37,68%). This means that our system missed a lot of FoRs (44 over 69 for FoR expressions, 46 over 76 for 'Ground object' and eight over 10 for 'Target entities'). This can be explained by the fact that we built the rules based on a preliminary analysis of a very limited corpus. We thus will need many more examples in order to build more exhaustive rules. This also shows there might be a potential for machine learning-based approaches, however, this will also need a larger manually annotated dataset in order to train a model.

11.4.3 Precision on adventure novel

In order to measure the quality of automatic geoparsing on an external source, we ran Perdido over the first four chapters of the novel *Kidnapped* by R. L. Stevenson. These chapters describe David Belfours journey to his uncle's house in Scotland. Within the first 3311 sentences, Perdido found 80 occurrences of a reference frame, that is, one occurrence per 41 sentences. We went through all of these text snippets and manually evaluated their correctness, in order to measure the *precision*.

As a result, 67 of these 80 occurrences were correct, which is a precision of **84%**. The stacked bar chart in Figure 11.4 shows the distribution over the different types of frames, where grey bars indicate the wrongly annotated cases. From this diagram, we can see that zonal frames were detected most often, followed by topological and Euclidean frames. Betweenness frames and linear construction frames are most seldom. Regarding the precision per type, we can see that it differs largely between the different types of frames. Euclidean, betweenness and linear construction frames were detected without any false positives, and only three out of 32 annotations of topological frames were incorrect. This corresponds to a precision of **90%**. The precision of zonal frames is worse, but still **70%**. It seems thus our rules work rather well also on external text sources for the considered frame types. We expect however that recall, which we did not test on the novel, should be equally worse as in the training corpus. Also, note that zonal frames are not only the most challenging case but also the most frequently occurring type.

To spot the syntactic reasons for this pattern, we plotted the frequency of different modifiers and keywords (every word except the ground and target objects) within the annotated text snippets, for all true-positive as well as

	FoR						Ground objects	Target entities
	all	btw	eucl	linear	topo	zonal		
Gold standard	69	2	10	18	17	22	76	10
Correct	18	1	4	4	2	7	28	1
(B)	6	1	1	0	2	2	2	1
(D)	44	0	5	14	12	13	46	8
(I)	7	0	3	0	0	4	10	3
(T)	1	0	0	0	0	1	0	0
(CT)	1	0	0	0	1	0	0	0
SER	80,43%	25%	85%	77,78%	82,35%	84,09%	75%	115%
Recall	37,68%	100%	50%	22,22%	29,41%	45,45%	39.47%	20%
Precision	78,79%	100%	62,5%	100%	100%	71,43%	75%	40%
Classification precision	72,73%	100%	62,5%	100%	80%	64,29%	75%	40%
Boundaries precision	57,58%	50%	50%	100%	40%	57,14%	70%	20%

Table 11.3: Number of errors and evaluation scores for our gold standard corpus.

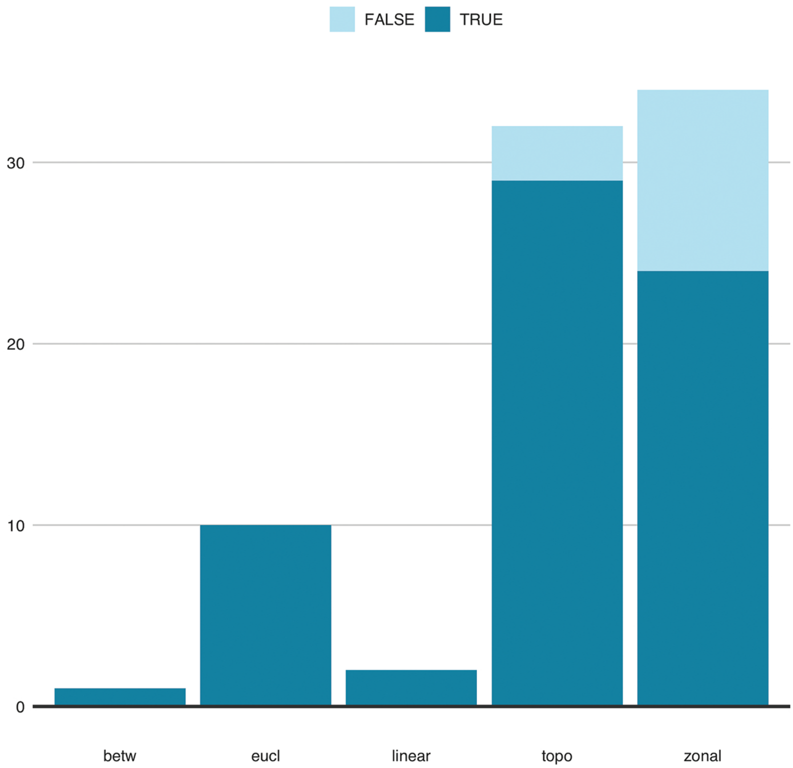

Figure 11.4: Precision of detecting different types of reference frames in the first 3311 sentences of the 1886 novel *Kidnapped*, including ZF ('zonal'), TF ('topo'), LCF ('linear'), EF ('eucl') and BF ('betw') frames.

false-positive annotations. Figure 11.5 shows a wordcloud for each type of frame over true- and false-positives. It can be seen there is no spottable difference in keywords between true and false zonal frames (Figure 11.5d), where locative prepositions 'where' and 'at' are used in both cases. An example for an erroneous annotation is:

'Looking $^{(ZF)}$ at [the shore]$_{(ground)}$ ... '

where the preposition 'at' is not locative, but used instead for indicating the direction of view. The topological frame errors (Figure 11.5b) have mostly to do with the motion indicator 'through', whereas a large diversity of expressions is correctly exploited in both the topological (Figure 11.5a) and the Euclidean case (Figure 11.5e).

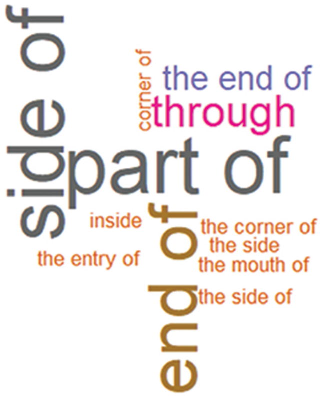

(a) Topological true.

(b) Topological false.

(c) Zonal true.

(d) Zonal false.

(e) Euclidean.

(f) BF and LCF.

Figure 11.5: Distribution of keywords and modifiers over true- and false-positive frame annotations.

11.5 Discussion and Future Work

Though our results are based on a limited corpus and thus are preliminary, we think they offer three main insights:

1. The suggested FoR typology seems to cover many relevant strategies in the chosen texts, however, it is almost certainly incomplete. We therefore expect that the diversity of cognitive spatial referencing strategies is far from exhausted by the suggested FoR model. Research that tests or extends this model is therefore needed. Furthermore, we believe that cognitive research (beyond text analysis) is needed in order to better understand which cognitive referencing strategy is actually used by readers when interpreting a text. The fact that a particular keyword is used in a particular context is only a very vague hint at the spatial cognitive strategy, and so text alone cannot be decisive. For this reason, empirical research combining spatial cognition and linguistics should focus on experiments that actually demonstrate and highlight the strategy hidden in the syntactic depths of environmental texts (Montello, 2009).
2. Even in case we might come up with clear, empirically validated models of FoR concepts, our study demonstrates that this does not yet mean clarity and ease of annotation for human annotators. As the rather low rate of inter-annotator agreement on types of FoR illustrates, human annotators frequently confuse actual movements with fictive motions and linear path references to space, as well as the strategies implied by mereology, topology, direction and distance in space. Another difficulty concerns the inherent vagueness of deciding which parts of speech should belong to a given frame (boundary errors). It seems that deciding on the precise sequence of words which denote a frame is hard, even if the frame used in a sentence might be easy to spot. This is true, by the way, for both human annotators and geoparsers. Annotators therefore need to be trained experts in order to be considered trustworthy producers of gold standards for information retrieval.
3. It is thus surprising that under these circumstances, the classification precision of a geoparser using transducer-based rules designed on a small training sample is considerably high (∼ 70–80%), both on training and test data (when disregarding boundary errors). The fact that recall, on the other hand, is so low (∼ 40%) shows that the main challenge of automatically discovering FoRs is not finding a robust model for a particular strategy, but rather handling the diversity of frames in their diverse natural language forms. In addition, we have a severe *cold-start problem*: In order to exploit machine learning to tackle diversity, we would need a large ground truth data set, which is lacking precisely for the reasons mentioned under point 2. It is thus an

important future task to establish a gold standard corpus of annotated FoRs of a sufficient size and variability.

What do these results tell us about the possibilities of automated spatial referencing and alignment of narrative texts? It seems that hand-crafted grammatical rules, as illustrated in this study, can in fact be used to reliably extract indirect spatial references, albeit only for a limited set of linguistic strategies. Once discovered, strategies allow approximating unnamed places of the described journey, for example, starting from a place in the middle of the moor, towards a place next to a Loch, and from there to a place near a certain mountain. Mapping of journeys in this way would allow us to find out to what extent the trails described by Murray, McFarlane and Stevenson really overlapped, and whether these authors saw the landscape from comparable vantage points. To this end, future work should investigate geometric approximations of any given FoR reference.

11.6 Conclusion

In this chapter, we have investigated the possibility of discovering indirect spatial references in environmental narratives, in order to align these narratives with the landscape features they describe. Based on two mountaineering texts over Scotland, we have explored the referencing strategies behind expressions used to localise the narrative within a landscape, and we have suggested a FoR typology on this basis, treating different geometric referencing strategies on a par. We have then designed rules and trained a transducer-based geoparser to automatically find these frames, before applying it to an external text for testing purposes. We compared the results with a manually annotated gold standard, which was also tested for inter-annotator agreement. Our results show that on the one hand, manual annotation of FoR types is surprisingly hard and annotators are in disagreement. On the other hand, classifier precision of the geoparser is considerably high on both the training and test data. The biggest challenge seems to be the low recall rate, which underlines our principle insight that discovering FoRs is primarily a problem of coping with the semantic and syntactic diversity of spatial referencing in texts. Yet, due to the considerably high precision, we believe automated parsing of indirect spatial references is possible and could be used to align documents and vantage points based on equivalent places described in these texts.

References

Anderson, Dave, Jade Goldstein-Stewart, Amal Fayad-Beidas, Dave Harris, Dulip Herath, Qian Hu, Janet Hitzeman, Seok Bae Jang, Inderjeet Mani, James Pustejovsky, et al. (n.d). *SpatialML: Annotation scheme for marking spatial expressions in natural language*. Tech. rep. Bedford.

Aurnague, Michel (2011). "How motion verbs are spatial: The spatial foundations of intransitive motion verbs in French". In: *Lingvisticae Investigationes* 34.1, pp. 1–34. DOI: 10.1075/li.34.1.01aur.

Burenhult, Niclas and Stephen C Levinson (2008). "Language and landscape: A cross-linguistic perspective". In: *Language Sciences* 30.2-3, pp. 135–150. DOI: 10.1016/j.langsci.2006.12.028.

Chen, Hao, Maria Vasardani, and Stephan Winter (2018). "Georeferencing places from collective human descriptions using place graphs". In: *J. Spatial Information Science* 17.1, pp. 31–62. DOI: 10.5311/JOSIS.2018.17.417.

Clementini, Eliseo (2013). "Directional relations and frames of reference". In: *GeoInformatica* 17.2, pp. 235–255. DOI: 10.1007/s10707-011-0147-2.

Clementini, Eliseo, Paolino Di Felice, and Daniel Hernández (1997). "Qualitative representation of positional information". In: *Artificial intelligence* 95.2, pp. 317–356. DOI: 10.1016/S0004-3702(97)00046-5.

Frank, Andrew U (1998). "Formal models for cognition–taxonomy of spatial location description and frames of reference". In: *Spatial cognition*. Berlin/Heidelberg: Springer, pp. 293–312. DOI: 10.1007/3-540-69342-4_14.

Freksa, Christian (1991). "Qualitative spatial reasoning". In: *Cognitive and linguistic aspects of geographic space*. Berlin/Heidelberg: Springer, pp. 361–372. DOI: 10.1007/978-94-011-2606-9_20.

Friburger, Nathalie and Denis Maurel (2004). "Finite-state transducer cascades to extract named entities in texts". In: *Theoretical Computer Science*. Implementation and Application of Automata 313.1, pp. 93–104. ISSN: 0304-3975. DOI: 10.1016/j.tcs.2003.10.007.

Gaio, Mauro and Ludovic Moncla (2019). "Geoparsing and geocoding places in a dynamic space context. The case of hiking descriptions". In: *The Semantics of Dynamic Space in French: Descriptive, experimental and formal studies on motion expression*. Ed. by Michel Aurnague and Dejan Stosic. Vol. 66. Human Cognitive Processing. Amsterdam: John Benjamins. Chap. 10, pp. 354–386. DOI: 10.1075/hcp.66.10gai.

Herring, John, David M Mark, and Max J Egenhofer (1994). *The 9-intersection formalism and its use for natural-language spatial predicates*. Tech. rep. Bedford.

Herskovits, Annette (1985). "Semantics and pragmatics of locative expressions". In: *Cognitive Science* 9.3, pp. 341–378. DOI: 10.1207/s15516709cog0903_3.

Levinson, Stephen C (2003). *Space in language and cognition: Explorations in cognitive diversity*. Vol. 5. Cambridge: Cambridge University Press. DOI: 10.1017/CBO9780511613609.

Macfarlane, Robert (2008). *The wild places*. London: Granta Books.

Makhoul, John, Francis Kubala, Richard Schwartz, Ralph Weischedel, et al. (1999). "Performance measures for information extraction". In: *Proceedings of DARPA broadcast news workshop*. Herndon, VA, pp. 249–252.

Mark, David M, Christian Freksa, Stephen C Hirtle, Robert Lloyd, and Barbara Tversky (1999). "Cognitive models of geographical space". In: *International Journal of Geographical Information Science* 13.8, pp. 747–774. DOI: 10.1080/136588199241003.

Moncla, Ludovic, Walter Renteria-Agualimpia, Javier Nogueras-Iso, and Mauro Gaio (2014). "Geocoding for texts with fine-grain toponyms: An experiment on a geoparsed hiking descriptions corpus". In: *Proceedings of the 22nd acm sigspatial international conference on advances in geographic information systems*. New York: ACM, pp. 183–192. DOI: 10.1145/2666310.2666386.

Moncla, Ludovic, Mauro Gaio, Javier Nogueras-Iso, and Sébastien Mustière (2016). "Reconstruction of itineraries from annotated text with an informed spanning tree algorithm". In: *International Journal of Geographical Information Science* 30.6, pp. 1137–1160. DOI: 10.1080/13658816.2015.1108422.

Moncla, Ludovic, Mauro Gaio, Thierry Joliveau, Yves-François Le Lay, Noémie Boeglin, and Pierre-Olivier Mazagol (2019). "Mapping urban fingerprints of odonyms automatically extracted from French novels". In: *International Journal of Geographical Information Science* 33.12, pp. 2477–2497. DOI: 10.1080/13658816.2019.1584804.

Montello, Daniel R (2009). "Cognitive research in GIScience: Recent achievements and future prospects". In: *Geography Compass* 3.5, pp. 1824–1840. DOI: 10.1111/j.1749-8198.2009.00273.x.

Murray, W.H. (2003). *Undiscovered Scotland*. Baton Wicks Publications, pp. 155–163.

Palmer, Bill (2002). "Absolute spatial reference and the grammaticalisation of perceptually salient phenomena". In: *Representing space in Oceania: Culture in language and mind*, pp. 107–157.

Purves, Ross S., Paul Clough, Christopher B. Jones, Mark H. Hall, and Vanessa Murdock (2018). "Geographic information retrieval: Progress and challenges in spatial search of text". In: *Foundations and Trends in Information Retrieval* 12.2-3, pp. 164–318. DOI: 10.1561/1500000034.

Pustejovsky, James, Jessica L Moszkowicz, and Marc Verhagen (2011). "Using ISO-Space for annotating spatial information". In: *Proceedings of the International Conference on Spatial Information Theory*.

Pustejovsky, James and Amber Stubbs (2013). *Natural language annotation for machine learning*. Sebastopol: O'Reilly, pp. 155–163.

Scheider, Simon and Ross Purves (2013). "Semantic place localization from narratives." In: *COMP@ SIGSPATIAL*, pp. 16–19.

Scheider, Simon, Jürgen Hahn, Paul Weiser, and Werner Kuhn (2018). "Computing with cognitive spatial frames of reference in GIS". In: *Transactions in GIS* 22.5, pp. 1083–1104. DOI: 10.1111/tgis.12318.

Stock, Kristin (2014). "A geometric configuration ontology to support spatial querying". In: *Connecting a Digital Europe Through Location and Place. The 17th AGILE International Conference on Geographic Information Science*.

Stock, Kristin and Javid Yousaf (2018). "Context-aware automated interpretation of elaborate natural language descriptions of location through learning from empirical data". In: *International Journal of Geographical Information Science* 32.6, pp. 1087–1116. DOI: 10.1080/13658816.2018.1432861.

Talmy, Leonard (1983). *How language structures space*. anglais. Berkeley Cognitive Science Report 4. Berkeley, CA, Etats-Unis: Cognitive Science Program, Institute of Cognitive Studies, University of California at Berkeley.

— (1996). "Language and space". In: vol. 21. Cambridge, Mass. London: MIT Press. Chap. Fictive motion in language and "ception", pp. 211–276.

Tenbrink, Thora (2011). "Reference frames of space and time in language". In: *Journal of Pragmatics* 43.3, pp. 704–722. DOI: 10.1016/j.pragma.2010.06.020.

Tuan, Yi-Fu (1991). "Language and the making of place: A narrative-descriptive approach". In: *Annals of the Association of American Geographers* 81.4, pp. 684–696. DOI: 10.1111/j.1467-8306.1991.tb01715.x.

Vandeloise, Claude (1986). *L'Espace en français. Sémantique des prépositions spatiales*. Paris: Editions du Seuil.

Varzi, Achille C (1996). "Parts, wholes, and part-whole relations: The prospects of mereotopology". In: *Data & Knowledge Engineering* 20.3, pp. 259–286. DOI: 10.1016/S0169-023X(96)00017-1.

Vasardani, Maria, Stephan Winter, Kai-Florian Richter, Lesley Stirling, and Daniela Richter (2012). "Spatial interpretations of preposition at". In: *Proceedings of the 1st ACM SIGSPATIAL International Workshop on Crowdsourced and Volunteered Geographic Information*. ACM, pp. 46–53. DOI: 10.1145/2442952.2442961.

Viehhauser, Gabriel and Florian Barth (2017). "Towards a digital narratology of space". In: *Digital Humanities 2017 Conference Abstracts* (McGill University & Université de Montréal Montreal Canada), pp. 643–646.

Worboys, Michael F (1996). "Metrics and topologies for geographic space". In: *Advances in Geographic Information Systems Research II: Proceedings of the International Symposium on Spatial Data Handling, Delft, pages A*. Vol. 7.