Natural Language meets Spatial Calculi

Joana Hois and Oliver Kutz

SFB/TR 8 Spatial Cognition
University of Bremen, Germany
{joana,okutz}@informatik.uni-bremen.de

Abstract. We address the problem of relating natural language descriptions of spatial situations with spatial logical calculi, focusing on projective terms (orientations). We provide a formalism based on the theory of $\mathcal{E}$-connections that connects natural language and spatial calculi. Semantics of linguistic expressions are specified in a linguistically motivated ontology, the Generalized Upper Model. Spatial information is specified as qualitative spatial relationships, namely orientations from the double-cross calculus.
This linguistic-spatial connection cannot be adequately formulated without certain contextual, domain-specific aspects. We therefore extend the framework of $\mathcal{E}$-connections twofold: (1) external descriptions narrow down the class of intended models, and (2) context-dependencies inherent in natural language descriptions feed back into the representation finite descriptions of necessary context information.

1 Introduction

We are aiming at a formal specification of connections between linguistic representations and logical theories of space. Language covers various kinds of spatial relationships between entities. It can express, for instance, orientations between them ("the cat sat behind the sofa"), regions they occupy ("the plant is in the corner"), shapes they commit to ("the terrace is surrounded by a wall"), or distances between them ("ships sailed close to the coast"). Formal theories of space also cover various types of relations, such as orientations [1], regions [2,3], shapes [4], or even more complex structures, such as map hierarchies [5]. Compared to natural language, spatial theories focus on one particular spatial aspect and specify its underlying spatial logic in detail. Natural language, on the other hand, comprises all of these aspects, and has thus to be linked to a number of different spatial theories. This linking has to be specified for each aspect and each spatial logic, identifying relevant information necessary for a linking or mapping function. This process involves contextual as well as domain-specific knowledge.

Our overall aim is to provide a general framework for identifying links between language and space as a generic approach to spatial communication and independent of concrete kinds of applications in which it is used. It should be applicable to any spatial context in connection with human-computer interaction, be it a geographic applications for way-finding and locating, city guides using maps, home/office automation applications, paths and spatial guidance, or architectural design planners. In particular, rather than attempting to integrate the most general spatial theories, we propose to use, in a modular way, various specialised (qualitative) spatial logics supporting dedicated and optimised reasoning algorithms.

In this paper, we analyse the linking between natural language and one specific aspect of space, namely orientation information for static spatial situations. We concentrate on static descriptions throughout this article, because dynamic descriptions (as they are defined in the linguistic ontology) do not differ from static descriptions with respect to their orientation-based locations: in "I am going to the left" and "The stove is to the left" the "to the left" refers to the same leftness in terms of the orientation. Moreover, most information about locatives are given by static descriptions of locations rather than dynamic movement [6].

We define links between language and a spatial theory concerning orientations, showing examples of linguistic projective terms, such as "A is to the right of B", "A is sitting to B's left", or "A is straight ahead". These types of terms are specified in a linguistic ontology, the Generalized Upper Model [7], and linked with necessary non-linguistic information of the orientation calculus [8]. In order to apply this representation to spatial maps, we introduce spatial orientations according to four basic projective, two-dimensional directions (left, right, front, back), which are distinguished and formalised. In particular, spatial entities are reducible to points and refer to material objects with finite dimensions.

We will introduce the linguistic ontology and its representation of spatial relationships in the next section. In Section 3, the connection between linguistic semantics, the double-cross calculus and relevant link-related aspects will be analysed using natural language examples. Finally, in Section 4, we will introduce an extension of the framework of $\mathcal{E}$-connections to formalise all these aspects in a modular way, which can be represented as a structured logical theory in the system HETS for heterogeneous specification.

2 Linguistic Spatial Semantics

Natural language groups spatial relations into different categories according to certain aspects, which can be related to specific spatial theories that deal with these aspects. A linguistic categorisation of spatial relationships on the basis of linguistic evidence, empirical research, and grammatical indications has been developed in detail in the Generalized Upper Model GUM [7,9], a linguistically motivated ontology. Linguistic ontologies structure language into groups of categories and relations by their semantics, i.e. categories are not based on lexemes but meanings. As a formal theory, GUM is axiomatised in first-order logic, parts of which can also be expressed in description logics (DLs) such as $\mathcal{SROIQ}$ [10], underlying the Web Ontology Language OWL 2.0. GUM's signature, i.e. its set of non-logical symbols, contains *categories* (unary predicates) and *relations* (binary predicates).

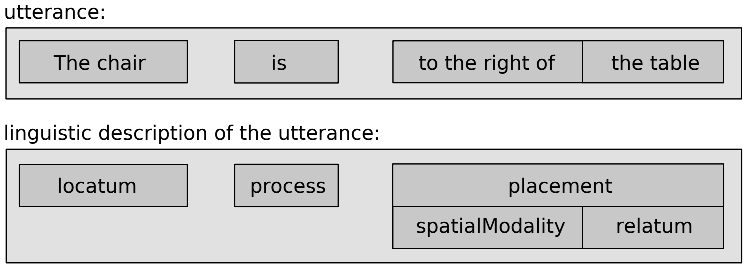

Fig. 1. Relations in GUM of an utterance example of a static spatial situation

GUM captures linguistic semantics of spatial expressions while nevertheless rendering this organisation independently of specific spatial logics. Also, its categorisation is not simply based on groups of spatial prepositions, but based on linguistic characteristics of spatial relations, grammatically or inherently, linguistic evidence and empirical data. Therefore, the development of GUM has been carried out with respect to empirical results in human computer interaction [11,7,12] and general linguistic research [13,14,15,16,17]. Utterances of spatial situations are specified as instances in GUM. We refer the reader to an overview of GUM in [9] and specific spatial components in [7].

2.1 Linguistic Specifications in the Generalized Upper Model

An utterance expressing a static spatial description is instantiated in GUM as a SpatialLocating. This category is a subclass of Configuration, a category that represents activities or states of affairs usually expressed at the level of the clause. They are defined according to their possible relations within the ontology, i.e. defined by entities that participate in the activity or state of affair. In principle, a single static description is specified by an instance of Configuration. Specific parts of the description (what, where, who, how, etc.) are specified by instances of Element, and their roles within the Configuration are specified by instances of relations (actor, manner, process, attribute, etc.).

Subcategories of Configuration that represent spatial activities or conditions are divided into static spatial situations and dynamic spatial situations. In the following, we will concentrate on the former, the SpatialLocating. This GUM category defines at least the following three relations:

1. The relation locatum in GUM relates the SpatialLocating to its located object within the spatial linguistic description. In the example “the chair is to the right of the table” (see Fig. 1), “the chair” is at a specific spatial position and represents the *locatum* [7] (also called the “referent” in [13]), i.e. the entity that is located somewhere.
2. The relation processInConfiguration relates the SpatialLocating to its process, the action or condition entity, which is usually expressed by a verbal group, indicating tense, polar and modal aspects [17]. In the example in Fig. 1, the process corresponds to “is”.

3. The relation placement relates the SpatialLocating to the location of the locatum. This location is represented by the GUM category GeneralizedLocation. It refers to "to the right of the table" in the example. A GeneralizedLocation specifies the spatial position of a locatum and consists of a spatial term, e.g. a spatial preposition, and an entity that corresponds to the reference object. Hence, the GeneralizedLocation defines two relations: spatialModality (spatial relation) and relatum (reference object). In the example, the spatialModality is expressed by "to the right of" and the relatum is expressed by "the table". The relatum, however, may remain implicit in natural language discourse [12], such as in the example "the chair is to the right", i.e. to the right of an undefined relatum, be it the speaker, listener or another entity. In case multiple relata are described together with the same spatial modality, they fill the relation relatum as a collection.

Binding the relatum and the spatialModality in the placement relation is rather a design issue than a logical constraint. This encapsulation allows convenient combinations of multiple locations expressed within one configuration: in the example "The plant is in the corner, by the window, next to the chair.", one SpatialLocating defines three placements. This is even more important as soon as placements are modified by expressing spatial perspectives, spatial accessibility, extensions or enhancements of the spatial relation. The utterance "The plant is to the *front left* of the chair, *right here* in the corner." combines two relations (front and left) with respect to one relatum (the chair), while a second relatum (in the corner) is combined with possible access information (right here). Moreover, modifications that are encapsulated together with the placement are easier to compare in case of re-use of spatial placements, e.g. throughout a dialogue discourse. Moreover, the GeneralizedLocation retains its structure independently of the configuration. It is equally specified in "he goes *to the right of the chair*" (dynamic spatial configuration) and "he stands *to the right of the chair*" (static spatial configuration), related by different relations (destination and placement).

Types of spatial relationships between locatum and reference objects are described by the category SpatialModality. Linguistically, this category corresponds to a preposition, an adverb, an adjective, or parts of the verb. It is subdivided into several categories that are primarily grouped into (1) relations expressing distance between entities, (2) functional dependencies between entities, and (3) positions between entities relative to each other depending on particular properties of the entities (such as intrinsic front side, size, shape). There are, however, intersections between these three general groups. Subcategories that refer particularly to spatial relationships based on orientations are subsumed under ProjectionRelation, describing positions between entities relative to each other depending on particular orientation-based properties of the entities.

2.2 Orientation-Related Linguistic Spatial Relationships

Projective Relations are distinguished along their three dimensions and can be divided into horizontal and vertical directions [18]. In order to reason (and talk)

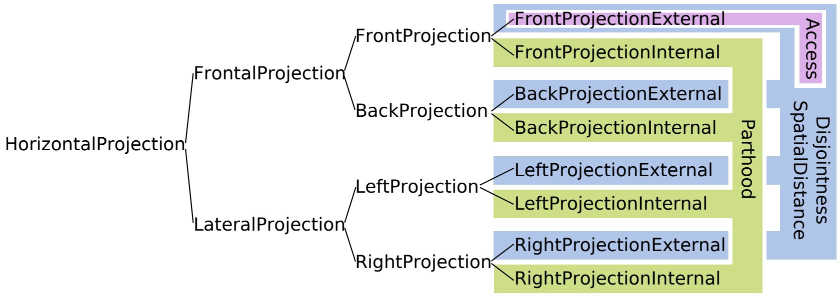

Fig. 2. Projective horizontal relations in GUM

about map-like representations, it suffices to concentrate on horizontal relations, which can be distinguished along lateral and frontal directions. Lateral projections comprise the directions left and right, frontal projections comprise front and back.

All four ontological categories of horizontal ground (atomic) projective relations, namely LeftProjection, RightProjection, FrontProjection, and BackProjection, can be expressed as an internal or external relationship [19]. Internal projective relations inherit from the category Parthood (topological) and refer to internal projections between locatum and relatum, such as "*A* is *in the left (part) of C*" or "*B* is *in the front of C*". External projective relations inherit from the categories Disjointness (topological) and SpatialDistance and refer to external projections between locatum and relatum, such as "*A* is *to the left of C*" or "*B* is *in front of C*" (compare Fig. 3). Furthermore, the category FrontProjectionExternal also inherits from the category Access, as external front projections imply functional access between locatum and relatum. An overview of the projective categories and their hierarchical dependencies in GUM are shown in Fig. 2. These categories are pairwise disjoint, for instance, FrontalProjection is disjoint with LateralProjection. They can, however, be *extended* (in GUM terminology), i.e. an instance of FrontProjectionInternal ("front") in "*A* is in the front left" is extended by an instance of LeftProjectionInternal ("left"). Spatial modalities can also be *enhanced* (in GUM terminology) by additional entities, e.g. distance information in "*A* is 10 meters to the left".

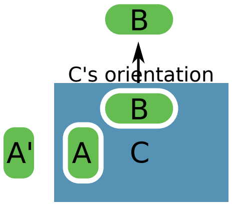

Fig. 3. Internal and external projective relations

Hence, GUM represents linguistic characterisations of orientations, which have to be associated with concrete spatial situations in order to yield a fully contextualised interpretation. In the next section, we will introduce an orientation-based spatial calculus and link this representation to GUM's projective categories. We will also identify missing aspects needed to minimise ambiguity in such a connection, namely context-dependent and domain-specific information.

3 Orientation Calculi and External Aspects

Spatial calculi address specific aspects of space, such as regions, orientations, shapes, etc., in order to provide formal representations as well as automatic reasoning techniques. General overviews for such representations are given in [20] and [21]. Calculi most relevant for mapping GUM's projective categories are those involving orientations since the linguistic projective relations described above refer to orientations within a spatial situation.[1] Many well known spatial calculi for orientations have been studied in the literature, among them are the double-cross calculus [8], the star calculus [22], the line segment-based calculus [23], or a model for positional relations[2] [24].

Such calculi are intended to be used for either static or dynamic relationships. They refer either to point-based or region-based spatial entities. They are based either on geometric or cognitive factors. The approach described in this paper maps orientations expressed in natural language to orientations represented in the double-cross calculus.

3.1 The Double-Cross Calculus

[8] introduces a ternary calculus of spatial orientations, the so-called double-cross calculus (DCC) [21]. In DCC, 15 relations are distinguished between an observer at position A, who is oriented (or moves) towards an entity at position B (compare Fig. 4). The 15 orientation relations are defined along three axes motivated by human cognitive characteristics: A and B are called the *perspective point* and the *reference point* respectively in [21]. They determine the front-back axis. Orthogonal to this axis are two further axes specified by A and B. Another entity located at some position C can then be described according to one of the 15 orientations.

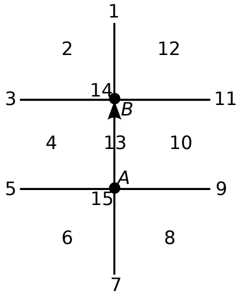

Fig. 4. DCC's 15 qualitative orientation relations according to [8]

Some of the correspondences between GUM and DCC are readily inferred: given an utterance, the perspective from where the relationship holds refers to an entity at A, the relatum refers to an entity at B, and the locatum refers to an entity located with respect to one of the 15 orientation relations determined by the spatial modality. The perspective, however, is often underspecified in utterances and might refer to the speaker, the listener, some other entity, or B. Which frame of reference [13] underlies the utterance is often not

---

[1] Although cardinal direction, i.e. north, east, south, west, are also related to orientations in some calculi, they are different from linguistic projective terms as introduced above and should thus be investigated separately.

[2] [24] use *projective* relations in their model, which do not correspond to linguistic projective relations as they are used in GUM (e.g. "surround", "inside", "outside" are not linguistic projective terms in GUM).

explicitly given. Also, in case the relatum is missing, B has to be inferred by other implicit or contextual information. The perspective (A) and the relatum (B) can even be identical: in this case, the locatum and the relatum are identical (i.e. $A = B$). The reference frame will automatically be intrinsic, and the orientation has to be determined by the intrinsic front.

Even if GUM's spatial relationships, then, are linked almost directly with DCC's orientations, especially by means of the inherent distinction between front/back and right/left projections, a missing perspective and relatum of an utterance have to be inferred and mapped to a DCC representation. What exactly these missing links are, and how an adequate mapping can be constructed by taking other information into account, is described in the following.

3.2 External Spatial Aspects in Linguistic Semantics

As GUM's linguistic specification is strongly based on concepts indicated by natural language, it does not entail enough information in order to map linguistic elements directly to entities of the spatial calculus. Hence, a mapping function from language to (models of) space needs additional information: [6] identifies eight parameters necessary to interpret a linguistic utterance. Among them are speaker and addressee, their locations and a view- or vantage point. Although [6] argues that orientations of speakers and addressees can be derived from their locations, this derivation is not specified in more detail, and as orientations are highly important in interpreting projective terms, our mapping has to specify them directly. Still missing are also intrinsic fronts of the reference object: projective linguistic terms can be interpreted along intrinsic orientations of objects independent of location and orientation of speaker or listener.

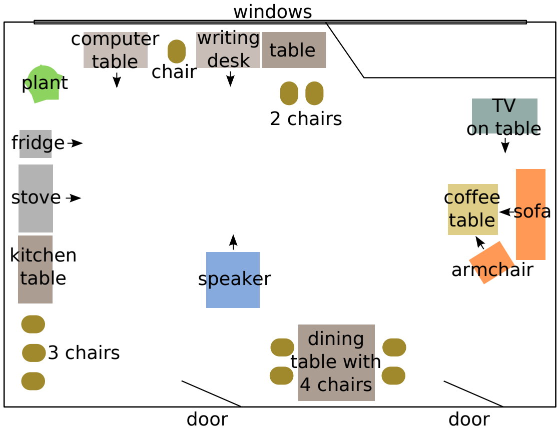

Fig. 5. Room layout of a scene description task, introduced in [7]. Arrows indicate intrinsic orientations of objects

Table 1. Example of utterances of native English speakers from spoken experiment and their representation in GUM. Utterances are cited without padding and pauses

utterance	locatum	spatialModality	relatum
1. the armchair is almost directly to my right	armchair	RightProjectionExternal	me
2. with the table just in front of it	table	FrontProjectionExternal	it
3. and diagonally behind me to the right is the table	table	BackProjectionExternal / RightProjection	me / –
4. the stove is directly to our left	stove	LeftProjectionExternal	us
5. and to the right of that is the fridge	fridge	RightProjectionExternal	that
6. there is a table to the right	table	RightProjectionExternal	–
7. further to the right a little bit in front is a living room	living room	RightProjection + FrontProjection	–
8. directly in front, there are two tables	tables	FrontProjectionExternal	–
9. from here the television is diagonally to the right	television	RightProjectionExternal (perspective: here)	–

Before we introduce links between corresponding linguistic and spatial entities, we start with examples of natural language utterances from a scene description. They motivate missing aspects not given in the utterance. The examples are taken from a series of experiments involving human-robot interaction in which participants were asked to explain a spatial situation to a robot. A detailed description of the experimental design is given in [7].

Fig. 5 shows the room layout of the experiment. Here, the position of speaker and listener coincides, i.e. they share the same perspective. In Table 1, an excerpt from the corpus data is given, in which participants refer to positions of objects in the room along their projective relationships. Although utterances from the corpus lack information about relatum and perspectives in general, such information is commonly omitted in natural language and has to be determined by other contextual or domain-specific factors. Even though positions of locatum, relatum, and perspective point have to be determined with respect to these external factors, links between projective spatial modalities and DCC relations can be defined in general: a concrete mapping, for instance, from a LeftProjection to the DCC orientations 2–6 is not affected by the position of A.

3.3 Non-linguistic Spatial Aspects of Projective Relations in GUM

The utterance "the armchair is (almost directly)[3] to my right" shows an example, where the locatum (armchair) is located to the right (RightProjectionExternal) of the speaker (relatum: me) (see Table 1). This utterance refers to an intrinsic frame of reference, where the perspective coincides with the relatum, i.e. the speaker, related to the position A in DCC. The locatum is then located at a

---

[3] Although GUM specifies modifications such as "almost directly" as enhancements of the spatial relation, we disregard them for a general mapping function, as they have minor impact on orientations (i.e. left does not become right).

point with one of the orientations 8–12 in the DCC, A and B are identical. In this example, information about the speaker's identity with "my (right)" and the frame of reference has to be added to the mapping function.

The next sentence "with the table just in front of it (the armchair)" also refers to an intrinsic frame of reference, but with the armchair as origin, i.e. the armchair refers to A in DCC (see also Fig. 6), which also coincides with B. In this case, the locatum (table) is located at a position with one of the orientations 1–4 and 10–14. Hence, information about the armchair's intrinsic front and the frame of reference have to be taken into account.

In case of a relative frame of reference as in "to the right of that (the stove) is the fridge", the perspective point A is indicated by the speaker, the reference point B is indicated by the relatum (stove), and the locatum (fridge) is indicated by a point that refers to one of the orientations 10–12 in DCC. Here, the frame of reference, the possibility of the stove having an intrinsic front and the perspective, i.e. the position of the speaker, are relevant for the mapping. If the relatum has no intrinsic front, it follows that a relative frame of reference applies. Otherwise, the choice of the underlying frame of reference is based on user preferences (extracted from the dialogue history) and the likeliness of intrinsic vs. relative frame of reference (according to the contextual descriptions).

In cases where the relatum is missing—e.g. the relatum of "further to the right" is omitted in Example 7—it is usually possible to determine its position by considering the preceding utterances. Hence, the sequence of utterances may give implicit information about missing entities in GUM's representation, and thus has to be considered throughout the construction of the mapping between GUM and DCC. Similarly, in Example 9, the given perspective "here" can either be interpreted as reference to the speaker or to the position that has just been described in a previous sentence, though a relative frame of reference can be assumed for explicit perspectives.

Given the corpus data, we conclude that the following parameters are involved in mapping the linguistic semantics of an utterance to a spatial situation:

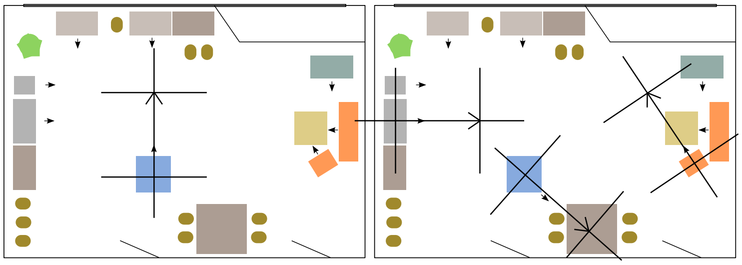

Fig. 6. DCC orientations of different entities: different perspectives cause different projective relationships. The DCC orientations in the left figure are based on the perspective of the speaker (participant), while the orientations in the right figure are based on intrinsic orientations of objects with intrinsic fronts and a changed orientation of the speaker. Objects are implicitly reduced to points defined by their centre

1. position and orientation of speaker and listener
2. reference system (relative or intrinsic) and origin (perspective)
3. domain-specific knowledge of entities (e.g. possibility of intrinsic fronts, their orientations and granularity)
4. dialogue history (sequence of utterances)

A linguistic representation in GUM together with the parameters can then be mapped to the location of the perspective point A, the reference point B and possible orientations towards the position of the located entity in DCC. The formalisation of this mapping is described in the following.

4 Multi-dimensional Formalisms and Perspectivism

The formation of multi-dimensional formalisms, i.e. formalisms that combine the syntax and semantics of different logical systems in order to create a new hybrid formalism, is a difficult and complex task in general, compare [25] for an overview. 'Classical' formalisms that have been used for formalising natural language statements involving modalities are counterpart theory [26] and modal predicate logics. However, both these formalisms, apart from being computationally rather difficult to deal with, are not particulary suited to deal with (qualitative) spatial reasoning as they do not, in their standard formulations, provide a dedicated spatial component, neither syntactically nor semantically. Similarly, the semantically tight integration that product logics of space and modality provide does not support the sometimes loose or unsystematic relationships that natural language modelling requires.

From the discussion so far, it follows that there are three desiderata for the envisaged formalisation:

1. To be able to represent various aspects of space and spatial reasoning, it needs to be multi-dimensional. However, in order to keep the semantics of the spatial calculi intact, the interaction between the formalisms needs to be loose initially, but also fine-tunable and controllable.
2. It needs to account for common sense knowledge that would typically be formalised in a domain ontology, and allow to restrict further the interaction between components.
3. It needs to account for context information not present in the representation using linguistic semantics.

The general idea of counterpart relations being based on a notion of *similarity*, however, gives rise to a framework of knowledge representation languages that seems quite well-suited to meet these requirements, namely the theory of $\mathcal{E}$-connections [27,28], which we sketch in the next section.

4.1 From Counterparts to $\mathcal{E}$-Connections

In $\mathcal{E}$-connections, a finite number of formalisms talking about distinct domains are 'connected' by relations relating entities in different domains, intended to capture different aspects or representations of the 'same object'. For instance, an 'abstract' object o of a description logic L_1 (e.g. an instance in GUM defining a linguistic item) can be related via a relation R to its life-span in a temporal logic L_2 (a set of time points) as well as to its spatial extension in a spatial logic L_3 (a set of points in a topological space, for instance). Essentially, the language of an $\mathcal{E}$-connection is the (disjoint) union of the original languages enriched with operators capable of talking about the link relations.

The possibility of having multiple relations between domains is essential for the versatility of this framework, the expressiveness of which can be varied by allowing different language constructs to be applied to the connecting relations. $\mathcal{E}$-connections approximate the expressivity of products of logics 'from below' and could be considered a more 'cognitively adequate' counterpart theory.

$\mathcal{E}$-connections have also been adopted as a framework for the integration of ontologies in the Semantic Web [29], and, just as DLs themselves, offer an appealing compromise between expressive power and computational complexity: although powerful enough to express many interesting concepts, the coupling between the combined logics is sufficiently loose for proving general results about the transfer of decidability: if the connected logics are decidable, then their (basic) connection will also be decidable. More importantly in our present context, they allow the heterogeneous combination of logical formalisms without the need to adapt the semantics of the respective components.

Note that the requirement of disjoint signatures of the *formal languages* of the component logics is essential for the expressivity of $\mathcal{E}$-connections. What this boils down to is the following simple fact: while more expressive $\mathcal{E}$-connection languages allow to express various degrees of *qualitative* identity, for instance by using number restrictions on links to establish partial bijections, they lack means to express 'proper' *numerical* trans-module identity.

For lack of space we can only sketch the formal definitions, and present only the two-dimensional case, but compare [28]: we assume that the **languages** $\mathcal{L}_1$ and $\mathcal{L}_2$ of two logics $\mathcal{S}_1$ and $\mathcal{S}_2$ are disjoint. To form a connection $\mathcal{C}^{\mathcal{E}}(\mathcal{S}_1, \mathcal{S}_2)$, fix a non-empty set of links $\mathcal{E} = \{E_j \mid j \in J\}$, which are binary relation symbols interpreted as relations connecting the domains of models of $\mathcal{S}_1$ and $\mathcal{S}_2$. The **basic $\mathcal{E}$-connection language** is then defined by enriching the respective languages with operators for talking about the link relations. A structure

$$\mathfrak{M} = \left\langle \mathfrak{W}_1, \mathfrak{W}_2, \mathcal{E}^{\mathfrak{M}} = (E_j^{\mathfrak{M}})_{j \in J} \right\rangle,$$

where $\mathfrak{W}_i = (W_i, .^{\mathfrak{W}_i})$ is an interpretation of $\mathcal{S}_i$ for $i \in \{1, 2\}$ and $E_j^{\mathfrak{M}} \subseteq W_1 \times W_2$ for each $j \in J$, is called an **interpretation** for $\mathcal{C}^{\mathcal{E}}(\mathcal{S}_1, \mathcal{S}_2)$. Given a concept C of logic $\mathcal{S}_2$, denoting a subset of W_2, the semantics of the basic $\mathcal{E}$-connection operator is

$$(\langle E_j \rangle^1 C)^{\mathfrak{M}} = \{x \in W_1 \mid \exists y \in C^{\mathfrak{M}} \ (x, y) \in E_j^{\mathfrak{M}}\}$$

Fig. 7 displays the connection of an ontology with a spatial logic for regions such as $\mathbf{S4_u}$, by means of a single link relation E which we might read as 'is the spatial extension of'.

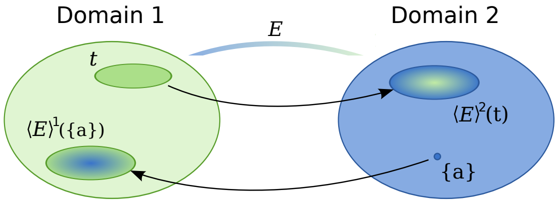

Fig. 7. A two-dimensional connection

As follows from the complexity results of [28], $\mathcal{E}$-connections add substantial expressivity and interaction to the component formalism. However, it is also clear that many properties related to (2) and (3) above can not directly be formalised in this framework. The next section sketches an extension to $\mathcal{E}$-connections that adds these expressive means, called perspectival $\mathcal{E}$-connections.

4.2 Perspectival $\mathcal{E}$-Connections

We distinguish three levels of interaction between the two representation languages $\mathcal{S}_1$ and $\mathcal{S}_2$:

1. *internal descriptions*: axioms formulated in the link language
2. *external descriptions*: axioms formulated in an external description language: reasoning over the same signature, but in a richer logic. They add interaction constraints not expressible in (1), motivated by general domain knowledge.
3. *context descriptions*: a class of admissible models needs to be finitely specified: here, not a *unique* model needs to be singled out in general, but a description of a class of models compatible with a situation (a context).

There are several motivations for such a modular representation: it (i) respects differences in epistemic status of the modules; (ii) reflects different representational layers; (iii) displays different computational properties of the modules; (iv) facilitates independent modification and development of modules; (v) allows to apply structuring techniques developed in algebraic specification theory; etc.

The general architecture of perspectival $\mathcal{E}$-connections is shown in Fig. 8. For an $\mathcal{E}$-connection of GUM with DCC, the internal descriptions cover the axioms of GUM and the constraint systems of DCC. Moreover, basic interactions can be axiomatised, e.g. mappings from GUM elements to DCC points need to be *functional*.

4.3 Layered Expressivity: External Descriptions and Context

The main distinction between external and contextual descriptions is not technical but epistemic. External descriptions are meant to enforce necessary interactions between ontological and spatial dimensions, while contextual descriptions add missing context information. The formal languages used to enforce these constraints will typically be different. Similar to conceptual spaces [30], they are intended to reflect different representational dimensions or layers of a situation.

$$\begin{tikzcd}
& \mathcal{O} \arrow[dr, dashed] & & \mathcal{D} \arrow[dl, dotted] & \\
\text{Contextual Descriptions} & & \mathcal{C}^{\mathcal{E}}(\mathcal{S}_1, \mathcal{S}_2)(\mathcal{D}, \mathcal{O}) & & \text{External Descriptions} \\
& & \mathcal{C}^{\mathcal{E}}(\mathcal{S}_1, \mathcal{S}_2) \arrow[u] & & \\
& & \text{Internal Descriptions} & &
\end{tikzcd}$$

Fig. 8. Architecture of Perspectival $\mathcal{E}$-connections

External Descriptions. An example, taken from [27], is the following constraint: “The spatial extension of the capital of every country is included in the spatial extension of that country”. This is a rather natural condition in an $\mathcal{E}$-connection combining a DL describing geography conceptually and a qualitative calculus for regions. Unfortunately, a basic $\mathcal{E}$-connection $\mathcal{C}^{\mathcal{E}}(\mathcal{ALCO}, \mathbf{S4_u})$ is not expressive enough to enforce such a condition. However, it can be added as an external description if we assume the external language allows quantification

$$\forall x \forall y \big(x \ \mathsf{capital\_of} \ y \rightarrow E(x) \subseteq E(y)\big)$$

In this case, the external description does not affect the decidability of the formalism, as shown in [28]. Of course, this is not always the case: the computational benefits of using $\mathcal{E}$-connections as the basic building block in a layered representation can get lost in case the external descriptions are too expressive. While a general characterisation of decidability preserving constraints is difficult to give, this can be dealt with on a case-by-case basis. In particular, the benefits of a modular design remain regardless of this issue.

Similarly to the above example, when combining GUM with DCC, assuming Φ axiomatises a LeftProjection (“left of”) within a SpatialLocating configuration, we need to enforce that elements participating in that configuration are mapped to elements of DCC models restricted to the five ‘leftness’ relations of DCC (see Section 3.2).

$$\forall x, y, z \ \Phi(x, y, z) \rightarrow \bigvee_{i=2}^{6} L_i(E(x), E(y), E(z))$$

This would be a typical external description for $\mathcal{C}^{\mathcal{E}}(GUM, DCC)$. Note that any internal description can be turned into an external one in case the external language is properly more expressive. However, the converse may be the case as well. For a (set of) formula(s) χ, denote by $\mathbf{Mod}(\chi)$ the class of its models. An external description Ψ may now be called **internally describable** just in case there is a finite set $\mathcal{X}$ of internal descriptions such that $\mathbf{Mod}(\Psi) = \mathbf{Mod}(\mathcal{X})$.

Contextual Descriptions. Assume an $\mathcal{E}$-connection $\mathcal{C} = \mathcal{C}^{\mathcal{E}}_{\mathcal{L}}(\mathcal{S}_1, \mathcal{S}_2)$ with link language $\mathcal{L}$ is given, and where $\mathbf{Sig}(\mathcal{C})$ denotes its signature, i.e. its set of non-logical symbols, including link relations. Moreover, assume $\mathfrak{S}$ is a finite set of **situations** for $\mathcal{C}$. Now, from an abstract point of view, a **context oracle** (or simply an oracle) is any function $\mathfrak{f}$ mapping situations for an $\mathcal{E}$-connection to a subclass of its models to pick out the class of models *compatible with a situation*:

$$\mathfrak{f}\colon \mathfrak{S} \longrightarrow \mathfrak{P}(\mathbf{Mod}(\mathbf{Sig}(\mathcal{C}))),$$

where $\mathfrak{P}$ denotes the powerset operation. This restricts the class of models for $\mathcal{C}^{\mathcal{E}}_{\mathcal{L}}(\mathcal{S}_1, \mathcal{S}_2)$ *independently* of the link language $\mathcal{L}$ and the external description language. For practical applications, however, we need to assume that these functions are computable, and that the classes $\{\mathfrak{f}(\mathfrak{s}) \mid \mathfrak{s} \in \mathfrak{S}\}$ of models they single out can be finitely described by a **context description language for** $\mathfrak{S}$. For combining GUM and DCC, the context description language simply needs to add the missing items discussed at the end of Section 3.3, i.e. fix the position of the speaker, the reference system, etc., relative to a situation $\mathfrak{s}$. Clearly, there are many options how to internalise the contextual information into an $\mathcal{E}$-connection. We have mentioned a language for specifying descriptions of finite models, but there are many other possibilities. For instance, [31] discuss several formal logics that have been designed specifically for dealing with contextual information, and compare their expressive power. Moreover, it might turn out that different contextual aspects require different logics or languages of context to be adequately formalised. Such problems, however, are left for future work.

4.4 Perspectival $\mathcal{E}$-Connections in Hets

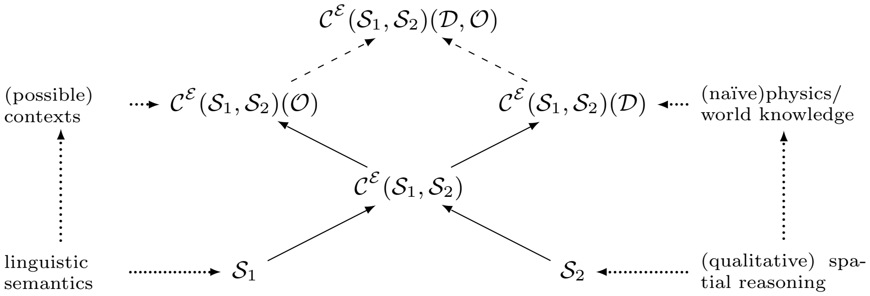

Fig. 9. Perspectival $\mathcal{E}$-connections as a structured theory in Hets

The Heterogeneous Tool Set Hets [32] provides analysis and reasoning tools for the specification language HetCasl, a heterogeneous extension of Casl supporting a wide variety of logics [33]. In particular, OWL-DL, relational schemes, sorted first-order logic $\mathbf{FOL}^{\mathsf{ms}}$, and quantified modal logic **QS5**, are covered. The DCC composition tables and GUM have already been formalised in CASL, and it has also been used successfully to formally verify the composition tables of qualitative spatial calculi [34].

As should be clear from the discussion so far, $\mathcal{E}$-connections can essentially be considered as many-sorted heterogeneous theories: component theories can be formulated in different logical languages (which should be kept disjoint or sorted), and link relations are interpreted as relations connecting the sorts of the component logics.[4]

Fig. 9 shows perspectival $\mathcal{E}$-connections as structured logical theories in the system HETS. Here, *dotted arrows* denote the extra-logical or external sources of input for the formal representation, i.e. for the description of relevant context and world-knowledge; *black arrows* denote theory extensions, and *dashed arrows* a pushout operation into a (typically heterogeneous) colimit theory of the diagram (see [35,36,37] for technical details).

5 Conclusions and Future Work

We have investigated the problem of linking spatial language as analysed in a linguistically motivated ontology with spatial (qualitative) calculi, by mapping GUM's projective spatial relationships to DCC's orientations. We concluded that various aspects important for this connection but omitted or not given explicitly in the linguistic semantics need to be added to the formal representation.

Moreover, we argued that these additional aspects can be divided into domain-specific (world-dependent) and contextual (situation-dependent) aspects. An approach for connecting all these heterogeneous modules into a structured heterogeneous theory is defined, called perspectival $\mathcal{E}$-connections.

Perspectival $\mathcal{E}$-connections now provide us with a formal framework for defining relationships between spatial language and calculi. This is not limited to the aspect of orientation discussed in detail in this paper. Rather, it can be carried out in the same way to deal with aspects covered by alternative orientation calculi, as well as calculi for distances, topology, shapes, etc. Here, the interplay between various such spatial calculi and GUM's respective treatment of the relevant non-projective spatial language has to be analysed.

Acknowledgements

Our work was carried out in the DFG Transregional Collaborative Research Center SFB/TR 8 Spatial Cognition, project I1-[OntoSpace]. Financial support by the Deutsche Forschungsgemeinschaft is gratefully acknowledged. The authors would like to thank John Bateman and Till Mossakowski for fruitful discussions.

---

[4] The main difference between various $\mathcal{E}$-connections now lies in the expressivity of the 'link language' $\mathcal{L}$ connecting the different logics. This can range from a sub-Boolean logic, to various DLs, or indeed to full first-order logic.

References

1. Moratz, R., Dylla, F., Frommberger, L.: A Relative Orientation Algebra with Adjustable Granularity. In: Proceedings of the Workshop on Agents in Real-Time and Dynamic Environments (IJCAI 05). (2005)
2. Casati, R., Varzi, A.C.: Parts and Places - The Structures of Spatial Representation. MIT Press (Bradford Books) (1999)
3. Cohn, A.G., Bennett, B., Gooday, J., Gotts, N.M.: Representing and reasoning with qualitative spatial relations. In Stock, O., ed.: Spatial and Temporal Reasoning. Kluwer Academic Publishers, Dordrecht (1997) 97–132
4. Schlieder, C.: Qualitative Shape Representation. In: Geographic Objects with Indeterminate Boundaries. Taylor & Francis, London (1996) 123–140
5. Kuipers, B.: The spatial semantic hierarchy. Artificial Intelligence **19** (2000) 191–233
6. Kracht, M.: Language and Space. (2008) Book manuscript.
7. Bateman, J., Tenbrink, T., Farrar, S.: The Role of Conceptual and Linguistic Ontologies in Discourse. Discourse Processes **44**(3) (2007) 175–213
8. Freksa, C.: Using orientation information for qualitative spatial reasoning. In Frank, A.U., Campari, I., Formentini, U., eds.: Theories and methods of spatio-temporal reasoning in geographic space. Springer, Berlin (1992) 162–178
9. Bateman, J.A., Henschel, R., Rinaldi, F.: Generalized Upper Model 2.0: documentation. Technical report, GMD/Institut für Integrierte Publikations- und Informationssysteme, Darmstadt, Germany (1995)
10. Horrocks, I., Kutz, O., Sattler, U.: The Even More Irresistible $\mathcal{SROIQ}$. In: Knowledge Representation and Reasoning (KR-06). (2006) 57–67
11. Shi, H., Tenbrink, T.: Telling Rolland where to go: HRI dialogues on route navigation. In: WoSLaD Workshop on Spatial Language and Dialogue, October 23-25, 2005, Delmenhorst, Germany (2005)
12. Tenbrink, T.: Space, time, and the use of language: An investigation of relationships. Mouton de Gruyter, Berlin (2007)
13. Levinson, S.C.: Space in Language and Cognition: Explorations in Cognitive Diversity. Cambridge University Press, Cambridge (2003)
14. Herskovits, A.: Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English. Studies in Natural Language Processing. Cambridge University Press, London (1986)
15. Coventry, K.R., Garrod, S.C.: Saying, seeing and acting. The psychological semantics of spatial prepositions. Essays in Cognitive Psychology. Psychology Press, Hove, UK (2004)
16. Talmy, L.: How language structures space. In Pick, H., Acredolo, L., eds.: Spatial Orientation: Theory, Research, and Application. Plenum Press, New York (1983) 225–282
17. Halliday, M.A.K., Matthiessen, C.M.I.M.: Construing Experience Through Meaning: A Language-Based Approach to Cognition. Cassell, London (1999)
18. Vorwerg, C.: Raumrelationen in Wahrnehmung und Sprache: Kategorisierungsprozesse bei der Benennung visueller Richtungsrelationen. Deutscher Universitätsverlag, Wiesbaden (2001)
19. Winterboer, A., Tenbrink, T., Moratz, R.: Spatial Directionals for Robot Navigation. In van der Zee, E., Vulchanova, M., eds.: Motion Encoding in Language and Space. Oxford University Press, Oxford (2008)

20. Cohn, A., Hazarika, S.: Qualitative Spatial Representation and Reasoning: An Overview. Fundamenta Informaticae **43** (2001) 2–32
21. Renz, J., Nebel, B.: Qualitative spatial reasoning using constraint calculi. In Aiello, M., Pratt-Hartmann, I., van Benthem, J., eds.: Handbook of Spatial Logics. Springer, Dordrecht (2007) 161–215
22. Renz, J., Mitra, D.: Qualitative direction calculi with arbitrary granularity. In Zhang, C., Guesgen, H.W., Yeap, W.K., eds.: PRICAI 2004: Trends in Artificial Intelligence. Volume 3157 of Lecture Notes in Computer Science., Springer (2004) 65–74
23. Dylla, F., Moratz, R.: Exploiting qualitative spatial neighborhoods in the situation calculus. In Freksa, C., Knauff, M., Krieg-Brückner, B., Nebel, B., Barkowsky, T., eds.: Spatial Cognition IV: Reasoning, Action, Interaction. International Conference Spatial Cognition 2004, Berlin, Heidelberg, Springer (2005)
24. Billen, R., Clementini, E.: Projective relations in a 3D environment. In Sester, M., Galton, A., Duckham, M., Kulik, L., eds.: Geographic Information Science, Berlin, Heidelberg, Springer (2006) 18–32
25. Gabbay, D., Kurucz, A., Wolter, F., Zakharyaschev, M.: Many-Dimensional Modal Logics: Theory and Applications. Number 148 in Studies in Logic and the Foundations of Mathematics. Elsevier, Amsterdam (2003)
26. Lewis, D.K.: Counterpart Theory and Quantified Modal Logic. In: Journal of Philosophy, Vol. 65. repr. in *The Possible and the Actual*, ed. Michael J. Loux, Ithaca 1979; also in: David K. Lewis, *Philosophical papers 1*, Oxford 1983. (1968) 113–126
27. Kutz, O.: $\mathcal{E}$-Connections and Logics of Distance. PhD thesis, The University of Liverpool (2004)
28. Kutz, O., Lutz, C., Wolter, F., Zakharyaschev, M.: $\mathcal{E}$-Connections of Abstract Description Systems. Artificial Intelligence **156**(1) (2004) 1–73
29. Cuenca Grau, B., Parsia, B., Sirin, E.: Combining OWL Ontologies Using $\mathcal{E}$-Connections. Journal of Web Semantics **4**(1) (2006) 40–59
30. Gärdenfors, P.: Conceptual Spaces - The Geometry of Thought. Bradford Books. MIT Press, Cambridge, MA (2000)
31. Serafini, L., Bouquet, P.: Comparing Formal Theories of Context in AI. Artificial Intelligence **155** (2004) 41–67
32. Mossakowski, T., Maeder, C., Lüttich, K.: The Heterogeneous Tool Set. In: TACAS 2007, Springer (2007) 519–522
33. CoFI: The Common Framework Initiative: CASL Reference Manual. Springer (2004) Freely available at `http://www.cofi.info`.
34. Wölfl, S., Mossakowski, T., Schröder, L.: Qualitative constraint calculi: Heterogeneous verification of composition tables. In: Proceedings of the Twentieth International Florida Artificial Intelligence Research Society Conference (FLAIRS 2007), AAAI Press (2007) 665–670
35. Kutz, O., Mossakowski, T., Codescu, M.: Shapes of Alignments: Construction, Composition, and Computation. In: International Workshop on Ontologies: Reasoning and Modularity (at ESWC). (2008)
36. Kutz, O., Mossakowski, T.: Conservativity in Structured Ontologies. In: 18th European Conf. on Artificial Intelligence (ECAI-08), IOS Press (2008)
37. Codescu, M., Mossakowski, T.: Heterogeneous colimits. In Boulanger, F., Gaston, C., Schobbens, P.Y., eds.: MoVaH'08. (2008)