

# Grounding Dynamic Spatial Relations for Embodied (Robot) Interaction

Michael Spranger<sup>1</sup>, Jakob Suchan<sup>2</sup>, Mehul Bhatt<sup>2</sup>, and Manfred Eppe<sup>3</sup>

<sup>1</sup>Sony CSL,  
3-14-13 Higashi Gotanda,  
141-0022 Tokyo

<sup>2</sup>Cognitive Systems, and  
Spatial Cognition Research Center (SFB/TR 8),  
University of Bremen, Germany

<sup>3</sup>IIIA-CSIC, Spain

**Abstract.** This paper presents a computational model of the processing of dynamic spatial relations occurring in an embodied robotic interaction setup. A complete system is introduced that allows autonomous robots to produce and interpret dynamic spatial phrases (in English) given an environment of moving objects. The model unites two separate research strands: computational cognitive semantics and on commonsense spatial representation and reasoning. The model for the first time demonstrates an integration of these different strands.

**Keywords:** *computational cognitive semantics, commonsense spatial reasoning, spatio-temporal dynamics*

## 1 Introduction

Commonsense spatio-linguistic abstractions offer a human-centred and cognitively adequate mechanism for the computational handling of spatio-temporal information in a wide-range of cognitive interaction systems, and cognitive assistive technologies [6]. Recently, there has been extensive work on static spatial relations such as “front”, “back”, “left” etc. We now have working models for the processing of static spatial relations [15, 14], the learning of spatial phrases [17], and their evolution [22]. At the same time, formalizations (e.g., logical, relational-algebraic) of space and development of tools for efficiently reasoning with spatial information is a vibrant research area within knowledge representation and reasoning (KR) [3, 18]. Commonsense spatial and temporal representations abstract from exact numerical representations by describing relations between objects using a finite set of relations. Spatial change is often modelled using *conceptual neighborhoods* [11]. Relations between two entities are conceptual neighbors if they can directly be transformed from one relation to another by continuous change of the environment. Researchers have investigated movement on the basis of an integrated theory of space, time, objects, and position [12] or defined continuous change using 4-dimensional regions in space-time [16]. Davis

[8] discusses the use of transition graphs for reasoning about continuous spatial change and applies them in physical reasoning problems. Qualitative spatial calculi have been integrated with the *situation calculus* family of action logics [4, 5]. A particular emphasis has been on the formal and computational characterisation of ‘space’ and dynamic ‘visuo-spatial’ problem-solving processes within a range of spatial assistance systems involving space and language [6].


On the other hand, there has been extensive work on modeling spatial relations for language. Models of proximal relations [14] based on proximity fields, for projective and absolute spatial relations based on prototypes [15] and group-based reference [29] have been proposed. Static spatial relations are interesting but they only cover a particular subdomain of spatial relations namely relations that do not encode temporal qualities. Recent models of dynamic spatial relations using semantic fields [10] or probabilistic graphical models [28] try to deal with temporal aspects of spatial relations. However, none of these models have attempted to integrate qualitative spatial reasoning systems which are able to deal with missing or unobserved data. This is somewhat of a surprise since the spatial reasoning community has provided sophisticated tools, frameworks and theories for modeling both static and dynamic spatial relations [3, 4].

This paper reports progress in combining the work on reasoning with cognitive semantics models of dynamic spatial relations. The focus of this paper is a model of dynamic spatial relations such as “across”, “over”, “into”, “out off” by integrating methods from cognitive computational semantics with commonsense spatial representation and reasoning techniques. The system presented here is capable of processing phrases such as “the yellow block *moves across* the red region”. The following sections will introduce the interaction scenario, followed by a more detailed description of the modules comprising the model.

## 2 Embodied Interaction

One of the key methodologies for research and validation of theories of grounding are situated interactions called *language games* [24]. Language games are interactions of two or more agents in the real (or in a simulated) world. Figure 1 shows the environment in which two robots interact. For the experiments presented in this paper, we used *Sony humanoid* robots. Robots are equipped with a vision system that fuses information from the robot’s camera with proprioceptive sensors distributed across the body. The vision system singles out and tracks objects [20]. Here, the environment contains four types of objects: *blocks*, *boxes*, *robots*. The vision system extracts the objects (as blobs) from the environment and computes a number of raw, continuous-valued features such as  $x$ ,  $y$  position, *width*, and *height* and colour values (YCbCr). Objects are tracked over time and assigned unique identifiers. So for instance, the green block has been given the arbitrary id `obj-12` by the left robot. The same robot identifies the white region as `reg-36`. The identifier stays the same for the period of time, where the robot is able to establish spatial-temporal continuity. We recorded a number of scenes with varying spatial combinations of objects, regions and landmarks.

A spatial language game follows a script between two randomly drawn agents from the population  $P$  of agents. One acts as the speaker, the other as the hearer. The agents see two scenes in succession that differ in terms of the movement of



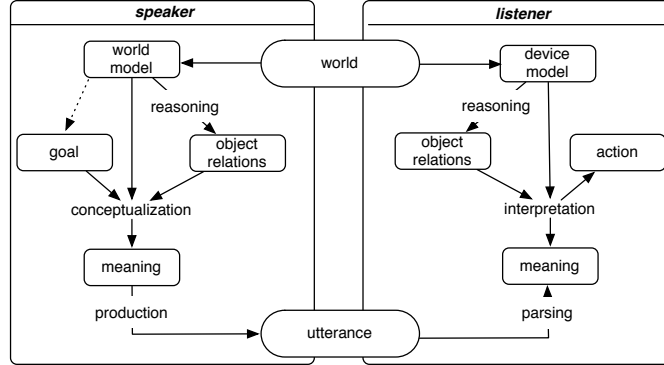
**Fig. 1.** Spatial setup. Objects move in an environment that features landmarks such as the box, and regions (colour patches on the ground). To the left the scene model extracted by the left robot is shown, to the right the scene model computed by the right robot is shown. The estimated movement of the block is signified through opacity. The starting point has a lower opacity (alpha) than the finish point. Regions are signified by the coloured rectangles. The blue square is the estimate of position and orientation of the box. Arrows signify robots. The lower centre arrow in each model stands for the own position (origin of the coordinate system). The other arrow visualises the estimated position and orientation of the other robot.

objects (but typically not in the number, type or colour of objects). For instance, in scene one an object might move from a yellow region to a white region over the red region. In the second scene, a similar object might move from the yellow region to the white region but without moving across the red region (see Fig. 1 as example).

1. The robots perceive two scenes and construe qualitative relations for each scene.
2. The speaker selects a scene from the two observed scenes, called the topic scene  $T$ .
3. The speaker *conceptualizes* a meaning comprised of dynamic spatial relations, and construal operations for discriminating  $T$ . E.g. the speaker computes that  $T$  is different from the second scene in that the yellow block crosses a red region.
4. The speaker tries to express the conceptualization using an English grammar. For instance, the robot might *produce* “the green block moves across the red region”.
5. The hearer *parses* the phrase using his English grammar, thereby computing the meaning underlying the phrase.
6. When the hearer was able to parse the phrase or parts of the phrase, he examines the two scenes to find the scene which satisfies the conceptualization. strategy (*interpretation*).
7. The hearer signals to the speaker which scene he thinks the sentence is about<sup>1</sup>.
8. The speaker checks whether the hearer selected the same scene as the one he had originally chosen. If they are the same, the game is a *success* and the speaker signals this outcome to the hearer. Otherwise, the game is a *failure*. In that case the speaker signals the topic  $T$  scene.

Agents are equipped with a range of software systems that allow them to participate successfully in interactions. The following sections detail the main components of the system 1) a qualitative spatial reasoner that provides qualitative

<sup>1</sup> Signalling is done by using extra-linguistic feedback. Two gestures are assigned the meaning scene-1 or scene-2. The knowledge which gesture refers to which scene is known to all participant agents.



**Fig. 2.** Schematic of the systems involved in autonomous processing of spatial scenes such as the one shown in Figure 1. Left are the processes the speaker runs, in order, to produce an utterance. Right shows the processes the hearer is running for understanding a phrase.

descriptions of dynamic scenes, 2) a cognitive semantics system that picks distinctive qualitative aspects of a dynamic scene in order to discriminate between scenes and 3) a construction grammar implementation of dynamic spatial phrase processing that allows to produce and parse correct English sentences. See Figure 2 for an overview of processing steps involved.

### 3 Reasoning for Dynamic Spatial Relations

Robots compute qualitative representations of space and motion. From these representations, spatio-temporal relations holding between spatial entities (objects) in the environment follow, i.e. *topology*, *orientation*, *movement*. The theory is implemented on top of CLP(QS) [5], which is a *declarative spatial reasoning framework* that can be used for representing and reasoning about high-level, qualitative spatial knowledge about the world.<sup>2</sup> CLP(QS) implements the semantics of qualitative spatial relations within a constraint logic programming framework (amongst other things, this makes it possible to use spatial entities and relations between them as native entities). In the following we describe the qualitative abstractions of space and motion and the thereon defined object relations, which serve as a basis to generate and interpret descriptions of the scene used by the robots in the description game.

#### 3.1 Qualitative Abstractions of Space and Motion

Based on perceived objects and regions represented by numerical features, the robots compute qualitative abstractions of space and motion (see Table 1) [26]. The qualitative abstractions are obtained from the sensory data by applying thresholds on the continuous feature values (position, size, angle) estimated from

<sup>2</sup> CLP(QS): A Declarative Spatial Reasoning System. [www.spatial-reasoning.com](http://www.spatial-reasoning.com)

$\Sigma$ Space	
<b>Topology</b>	$DC(p, q, t), PO(p, q, t), PP(p, q, t), PPI(p, q, t), EQ(p, q, t)$
<b>Extrinsic Orientation</b> (horizontal and in depth)	$left(p, q, t), overlaps\_left(p, q, t), along\_left(p, q, t), horizontally\_equal(p, q, t),$ $overlaps\_right(p, q, t), along\_right(p, q, t), right(p, q, t)$ $closer(p, q, t), overlaps\_closer(p, q, t), along\_closer(p, q, t), distance\_equal(p, q, t),$ $overlaps\_further(p, q, t), along\_further(p, q, t), further(p, q, t)$
$\Sigma$ Motion	
<b>Movement</b>	$approaching(p, q, t)$ and $receding(p, q, t)$

**Table 1.** Spatial relations used to describe the spatial configuration of a scene and the corresponding motion relations.

objects in the scene. The detected entities are represented by the basic domain primitives: *regions*, *points*, and *oriented-points*.

### 3.2 $\Sigma$ Space – Qualitative Spatial Relations

Spatial primitives are the basis for computing the spatial relations *topology*, and *extrinsic-orientation*.

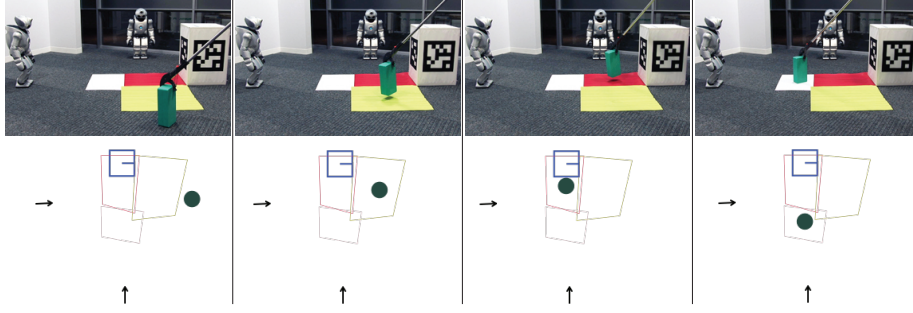
**Topology** We represent the connectedness of pairs of spatial primitives by the relations of the region connection calculus [7], using the RCC5 [7] subset in a ternary version, where the third argument represents the time point at which the relation holds.

**Extrinsic Orientation (Position)** We represent the extrinsic orientation (relative position) of two spatial primitives, with respect to the observer’s viewpoint, making distinctions on the *position* and the *size* of the spatial entities. For the described robot scenario, we use a 2-Dimensional representation that resemble Allen’s interval algebra [2] for each dimension, i.e. *vertical*, and *depth* (distance from the observing robot). However, in terms of visual perception, the interval relations that happen “instantaneously” (namely, *meets*, *starts*, and *finishes*) are irrelevant.

### 3.3 $\Sigma$ Motion – Qualitative Spatial Change

The dynamics of scenes are represented in terms of the *relative movement* of the objects [26].

**Relative Movement** The relative movement of pairs of spatial primitives is represented in terms of changes in the distance between the *centroids* of the entities. We represent these changes in terms of *approaching* and *receding* as defined below.



**Fig. 3.** Dynamic scene as observed by the robot to the left.

$$\text{approaching}(p, q, t) \leftrightarrow \exists t_1 t_2 (t_1 < t) \wedge (t < t_2) \wedge (\text{dist}(p, q, t_2) < \text{dist}(p, q, t_1)); \quad (1a)$$

$$\text{receding}(p, q, t) \leftrightarrow \exists t_1 t_2 (t_1 < t) \wedge (t < t_2) \wedge (\text{dist}(p, q, t_2) > \text{dist}(p, q, t_1)). \quad (1b)$$

Notice that the timepoint  $t$  falls within the open time interval  $(t_1, t_2)$ , which is assumed to be very small; therefore, these predicates are locally valid with respect to the time point [26].


### 3.4 Spatio-Temporal Object Relations

To describe perceived scenes in terms of spatio-temporal phenomena we combine the different aspects of the theory about *space* and *motion* providing a rich vocabulary about qualitative changes in the visual domain. This allows us to describe the ongoing interactions and operations between the physical entities represented by the spatial entities.

**Robots and objects** Spatial scenes (for the purpose of this paper) consist of blocks, regions and robots. For the objects in the scene we assume that they are all ridged and non-opaque. For the regions and robots, we assume, that they are static. This means that only blocks are considered movable, so any observed change in the scene is due to movement of the blocks. Additionally, robots are assumed to have a certain orientation (facing direction) and we define abstract objects to represent the robots field of view.

Object relations are defined on the relations of space and motion and temporal relations (similar to Allen's Interval Algebra [2]). These relations describe changes in the spatial configuration and motion of the objects and robots in the environment and are used by the robots to generate and interpret descriptions used in the discrimination game. E.g., a block  $B$  enters a region  $R$  at an time interval  $I$ .

$$\begin{aligned} \text{moves\_into}(\text{object}(B), \text{region}(R), I) \leftarrow & \text{approaching}(\text{object}(B), \text{region}(R), I_4) \wedge \\ & DC(\text{object}(B), \text{region}(R), I_1) \wedge PO(\text{object}(B), \text{region}(R), I_2) \wedge PP(\text{object}(B), \text{region}(R), I_3) \wedge \\ & \text{during}(I_2, I_4) \wedge \text{equal}(I, I_2) \wedge \text{meets}(I_1, I_2) \wedge \text{meets}(I_2, I_3). \end{aligned}$$



**Fig. 4.** Object Relations based on qualitative abstractions of space and motion used to generate and interpret scene descriptions

A complex interaction as e.g. a block B moving across a region R is then defined based on the basic interactions, i.e., *moves\_into*, *moves*, *moves\_out\_of* etc.

$$\begin{aligned}
 moves\_across(object(B), region(R), I) \leftarrow & moves(object(B), direction(Dir), I_1) \wedge \\
 & moves\_into(object(B), region(R), I_2) \wedge moves\_out\_of(object(B), region(R), I_3) \wedge \\
 & during(I, I_1) \wedge starts(I_2, I) \wedge finishes(I_3, I)
 \end{aligned}$$

Based on these object relations a description of the scene is generated by means of declarative logic programming. An exemplary description of the scene shown in Fig. 3 using the perspective of one of the robots (robot-a):

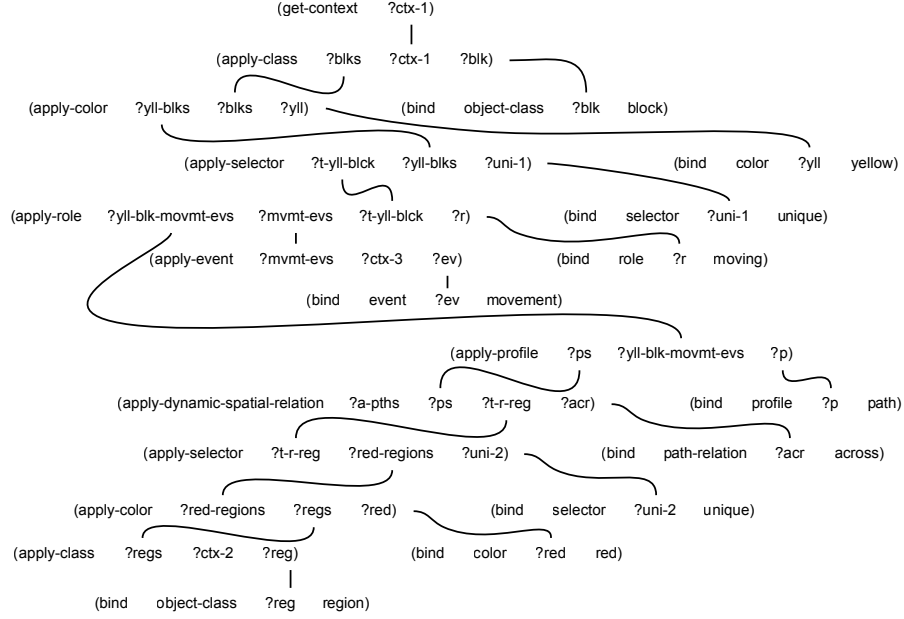
$$\begin{aligned}
 & moves(object(obj-12), direction(left), I_1) \wedge moves\_into(object(obj-12), region(reg-38), I_2) \wedge \\
 & moves\_across(object(obj-12), region(reg-38), I_3) \wedge moves\_into(object(obj-12), region(reg-37), I_4) \wedge \\
 & moves\_out\_of(object(obj-12), region(reg-38), I_5) \wedge moves(object(obj-12), direction(closer), I_6) \wedge \\
 & moves\_out\_of(object(obj-12), region(reg-37), I_7) \wedge moves\_into(object(obj-12), region(reg-36), I_8)
 \end{aligned}$$

These object relations (as depicted in Fig. 4) can be uttered with the following natural language description:

“The green block moves left, across the yellow region and enters the red region. It moves closer and leaves the red region at the bottom. The block enters the white region and stops within this region.”

## 4 Conceptualization and Interpretation

Robots use temporal object relations to pick a particular semantics discriminating or describing a scene. To represent the semantics of spatial phrases, we use a Cognitive semantics inspired formalism called Incremental Recruitment Language (IRL) [22]. The key idea behind this formalism is that the semantics of natural language can be modeled as a *semantic program* [13]. In IRL, the



**Fig. 5.** IRL-program representing the semantic structure of the phrase “the yellow block is moving across the red region”. The connections between the operations and bind-statements signify variables appearing in multiple places of the graph.

meaning of an utterance consists of an algorithm and data pointers that when executed by the hearer will lead him to identify the topic (i.e. topic scene).

Let us exemplarily consider the phrase “the yellow block moves across the red region”. The phrase consists of a dynamic spatial relation (across), concepts such as box and region and determiners. The example phrase encodes a particular strategy for conceptualization where the spatial relation is used in conjunction with a region. An interpreter of the phrase has to construe the path [27] of the object using the region and the particular dynamic spatial relation.

Figure 5 shows a graphical representation of the IRL-program underlying the phrase. Such programs consist of two things and links between them.

**Cognitive operations** represent algorithms used in conceptualization. They encode a particular cognitive function such as categorization using a spatial category, applying a selector or applying an object class such as region or box. Cognitive operations are identified by their name, e.g. **apply-class** and they have a set of arguments, which can be linked to other operations or semantic entities via variables (starting with ?).

**Semantic entities** are the data that cognitive operations work with. They can be prototypes, concepts and categories or more generally representations of the current context, as well as data exchanged between cognitive operations. They are introduced explicitly in the network via *bind-statements*.



For instance, the statement (`bind dynamic-spatial-relation ?acr across`) encodes the access to the agent-internal, dynamic spatial relation `across` which will be bound to the variable `?acr`.

#### 4.1 Evaluation

Evaluation is a process by which IRL-programs get executed. The process determines bindings for the variables in an IRL-program given a particular dynamic spatial scene or a number of spatial scenes. This process can fail, for instance, when a particular spatial scene does not fit the program. More precisely, evaluation can succeed, or fail, but all successful evaluations are also scored [21], as to how much the program fits the scene.

Let us assume the hearer wants to interpret the example phrase and has decoded the IRL-program in Figure 5. Evaluation proceeds as follows. First all bind-statements are evaluated, after which, for example, the variable `?blk` is bound to the concept `block` and `?acr` to the spatial category `across`. After that, the evaluation engine will try and find cognitive operations that can be evaluated. Probably, the first cognitive operation to be evaluated is `get-context` which binds the variable `?ctx-1` to each of the 2 scenes presented to the agent. Next `apply-class` identifies blocks and regions from the context. After that `apply-color` filters regions and blocks using colors. Then `apply-determiner` applies uniqueness constraints.

The cognitive operations are implemented by integrating ideas from prototype theory and spatial reasoning. The following gives a brief and incomplete overview of the inner workings of the most important operations<sup>3</sup>

**apply-event** (for our purposes here) applies the movement event descriptor and computes a set of events/trajectories that can be categorized as movements. For the scene discussed in Section 2, one movement event will be identified. The representation of these event includes information of the type (movement), event participants involved in the event (green block), as well as the time frame for it. Notice that at this stage no qualitative spatial information is used. There is a threshold for what is considered a movement vs. a non-movement and a classifier identifies those timeframes and objects that are considered part movements. For this paper only movement events are considered but the same mechanisms extend to complex events such as grasping, pushing etc.

**apply-role** filters events for their participants. Here, the green block has to be in the moving role of the event. All movement events which are not involving a moving green block are filtered out.

**apply-profile** is a cognitive semantics operation that focusses on aspects of a movement event, such as source, path or goal. Here the focus is on path, which means that the event representation is annotated with a particular focus on the movement part of the trajectory and not it's goal or starting point.

---

<sup>3</sup> This is a simplified description of the actual algorithm, cognitive operations implement multiple input/output patterns

`apply-dynamic-spatial-relation` checks whether a particular spatial relation such as `across` applies to the profiled aspect of the input movement events. Its input is a set of movement events. The output is a set of movement events that can be categorized by the spatial relation and additional information. For this, the operation queries the outcome of the spatial reasoner for the trajectories, i.e. objects, regions and their relations and tests whether it can find any trace of the spatial relation `across` as applying to the movement event. Here it has to check whether a particular movement events trajectory includes a particular region and whether this has region has been crossed.

## 4.2 Conceptualization strategies

The IRL-program in Figure 5 is part of a particular conceptualization strategy that can involve other dynamic spatial relations such as `into`, `outof` etc. Spatial conceptualization strategies involve more than just the choice of a spatial relation. Landmarks, perspective, frames of reference [29] are all important aspects of the construal of spatial reality. All of these are implemented in our system and can be used to produce sentences of significant complexity.

IRL includes mechanisms for the automatic and autonomous construction of IRL-programs. Agents use these facilities in two ways. First, when the speaker wants to talk about a particular scene, he conceptualizes an IRL-program for reaching that goal (see conceptualization in Figure 2). Secondly, a listener trying to interpret an utterance will construct and evaluate programs, in order, to find the best possible interpretation of the utterance (see interpretation in Figure 2). Interpretation and conceptualization are implemented as heuristics-guided search processes that traverse the space of possible IRL-programs. The basic building blocks for the search are IRL-programs packaged into *chunks* which are larger structures reflecting the standard semantics of some particular natural language, e.g. determined noun phrases in English. The IRL search process progressively combines chunks of IRL-programs into more and more complex IRL-programs. Each program is tested for compatibility with the goal of the agent, as well as the context.

## 5 Production and Interpretation of Spatial Phrases

Robots are communicating the conceptualisation of the topic scene using an English grammar. The grammar is implemented using a bidirectional Construction Grammar system called *Fluid Construction Grammar* (FCG) [25]. FCG uses one engine and a single grammar representation to support both production and interpretation. We implemented a spatial grammar comprising lexical items for basic concepts (e.g. block, box), events (e.g. move), as well as the spatial relations (e.g. along, across, into, out off). Additionally, we implemented a number of spatial grammar constructions. The system is similar to the one proposed in [19]. In total there are over 70 constructions.

FCG starts from a two sided feature structure. One side is for semantic features, another side for syntactic information. Constructions check if they find sufficient information to apply and subsequently if that is the case change the

structure by adding information or introducing hierarchy etc. In this paper the focus is on semantics and reasoning. Nevertheless we give a short overview of the constructions involved in translating IRL-programs into spatial utterances and back. The constructions are explained here as they would apply in production. That is, given an IRL-program (e.g. Figure 5), FCG will produce a phrase such as “the green block moves across the red region”.

**Lexical constructions** are bidirectional mappings between entities (bind statements in IRL-programs) and stems. For instance, there is a lexical construction for that maps (bind `dynamic-spatial-relation ?acr across`) to the stem “across”. Another example is the construction that maps (bind `object-class ?block block`) to “block” or the construction that maps (bind `color-category ?y11 yellow`) to “yellow”.

**Functional constructions** map each lexical item to a word class (also called lexical class). The idea is that the same concept can be used in different forms. For instance, a colour category can be used as an adjective such as in “the yellow block” or as a noun. Which one is used is determined by semantics. If the colour category is used as a modifier such as in the operation `apply-color` than its word class is adjective. The same is true for the spatial relation. Across is used here with the operation `apply-dynamic-spatial-relation` so that the dynamic-spatial-relation preposition construction would translate it to a preposition.

**Phrasal constructions** take into account the larger syntactic and semantic context. An example is the adjective-noun-phrase construction, which looks for an adjective and a noun as well as a particular structure in the IRL-program and adds phrasal information such as word order. Another example is the prepositional phrasal construction that combines a preposition and a noun phrase. For the discussed example, this combines the region noun phrase and the preposition across and adds word order. Other phrasal construction handle the verb phrase and the prepositional phrase, or combine the structure into one coherent phrase.

Applied to an example such as the one in Figure 5 the system produces the intended sentence “the yellow block moves across the red region”.

## 6 Discussion and Future Work

We tested the result of the current system on the initial scenes described in this paper. The system was able to correctly produce and interpret phrases allowing robots to communicate about scenes and discriminate between scenes with different temporal and spatial characteristics. Future work will have to test the system systematically on more scenes and perform a detailed analysis.

There are a number of possible extensions and future work on the system discussed in this paper. For example, the spatial reasoning system used in this paper has much more capabilities then employed for the purpose of this paper. The system can be used for abductive reasoning and explain missing or faulty observations [9]. In case of the robot scenario presented in this paper, this can be used to generate or interpret descriptions where only partial information are

available. E.g. consider the following scene: The object moves from the white to the yellow region, thereby passing a red region. Robot A perceives the whole narrative. It plays a description game with robot B, who is not able to see the red region because visibility is blocked. However, since it perceives the block leaving the yellow region and entering the white region, it could abduce that the red region was crossed.

Furthermore, the qualitative abstractions can be used to translate a description which is based on the point of view of one robot to the point of view of another robot. E.g. the description provided by the speaker says: “The green block moves from left to right.” In order to correctly understand this description, the hearer needs to be able to “imagine” the scene from the speaker’s point of view. To this end, reasoning about the robots perspective based on the position of the objects and the intrinsic-orientation of the other robot, can be used to produce or interpret viewpoint dependent descriptions of the scene. This has been studied as part of research into static locations [23], but so far has not been integrated with the system presented in this paper.

In conclusion, the system presented in this paper presents a fully working system able to interpret and produce natural language phrases with dynamic spatial relations. Importantly, this is a first step towards understanding the acquisition and evolution of dynamic spatial relations. The computational reconstruction of processing will allow us to study the learning of parts of the spatial grammar, lexicon, conceptual repertoire, and ultimately setup agent-based simulations, where we can study the evolution of dynamic spatial relations.

## References

1. Aiello, M., Pratt-Hartmann, I.E., Benthem, J.F.v.: Handbook of Spatial Logics. Springer-Verlag New York, Inc., Secaucus, NJ, USA (2007)
2. Allen, J.F.: Maintaining knowledge about temporal intervals. *Commun. ACM* 26(11), 832–843 (1983)
3. Bhatt, M., Guesgen, H., Wölfl, S., Hazarika, S.: Qualitative spatial and temporal reasoning: Emerging applications, trends, and directions. *Spatial Cognition & Computation* 11(1), 1–14 (2011)
4. Bhatt, M.: Reasoning about space, actions and change: A paradigm for applications of spatial reasoning. In: *Qualitative Spatial Representation and Reasoning: Trends and Future Directions*. IGI Global, USA (2012)
5. Bhatt, M., Lee, J.H., Schultz, C.: CLP(QS): A Declarative Spatial Reasoning Framework. In: *COSIT*. pp. 210–230 (2011)
6. Bhatt, M., Schultz, C., Freksa, C.: The ‘Space’ in Spatial Assistance Systems: Conception, Formalisation and Computation. In: Tenbrink, T., Wiener, J., Claramunt, C. (eds.) *Representing space in cognition: Interrelations of behavior, language, and formal models*. Series: *Explorations in Language and Space*. 978-0-19-967991-1, Oxford University Press (2013)
7. Cohn, A., Bennett, B., Gooday, J., Gotts, N.: Representing and reasoning with qualitative spatial relations about regions. In: Stock, O. (ed.) *Spatial and Temporal Reasoning*, pp. 97–134. Kluwer Academic Publishers, Dordrecht (1997)
8. Davis, E.: Qualitative reasoning and spatio-temporal continuity. In: Hazarika, S.M. (ed.) *Qualitative Spatio-Temporal Representation and Reasoning: Trends and Future Directions*, pp. 97–146. IGI Global, Hershey, PA (2012)

9. Dubba, K., Bhatt, M., Dylla, F., Hogg, D., Cohn, A.: Interleaved inductive-abductive reasoning for learning complex event models. In: *Inductive Logic Programming, LNCS*, vol. 7207, pp. 113–129. Springer Berlin / Heidelberg (2012)
10. Fasola, J., Mataric, M.J.: Using semantic fields to model dynamic spatial relations in a robot architecture for natural language instruction of service robots. In: *Intelligent Robots and Systems (IROS)*, 2013 IEEE/RSJ International Conference on. pp. 143–150. IEEE (2013)
11. Freksa, C.: Conceptual neighborhood and its role in temporal and spatial reasoning. In: Singh, M., Travé-Massuyès, L. (eds.) *Decision Support Systems and Qualitative Reasoning*. pp. 181–187. North-Holland, Amsterdam (1991)
12. Galton, A.: *Qualitative Spatial Change*. Oxford University Press (2000)
13. Johnson-Laird, P.N.: Procedural semantics. *Cognition* 5(3), 189–214 (1977)
14. Kelleher, J., Kruijff, G.J., Costello, F.: In: *ACL-44: Proceedings of the 21st International Conference on Computational Linguistics*. Morristown, NJ, USA
15. Moratz, R., Tenbrink, T.: Spatial reference in linguistic human-robot interaction: Iterative, empirically supported development of a model of projective relations. *Spatial Cognition & Computation* 6(1), 63–107 (2006)
16. Muller, P.: A qualitative theory of motion based on spatio-temporal primitives. In: Cohn, A.G., Schubert, L.K., Shapiro, S.C. (eds.) *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, Trento, Italy, June 2-5, 1998. pp. 131–143. Morgan Kaufmann (1998)
17. Regier, T.: The emergence of words: Attentional learning in form and meaning. *Cognitive Science* 29(6), 819–865 (2005)
18. Renz, J., Nebel, B.: Qualitative spatial reasoning using constraint calculi. In: *Handbook of Spatial Logics* [1], pp. 161–215
19. Spranger, M., Loetzsch, M.: In: Steels, L. (ed.) *Design Patterns in Fluid Construction Grammar*, pp. 265–298. John Benjamins
20. Spranger, M., Loetzsch, M., Steels, L.: A Perceptual System for Language Game Experiments. In: Steels, L., Hild, M. (eds.) *Language Grounding in Robots*, pp. 89–110. Springer (2012)
21. Spranger, M., Pauw, S.: Dealing with Perceptual Deviation - Vague Semantics for Spatial Language and Quantification. In: Steels, L., Hild, M. (eds.) *Language Grounding in Robots*, pp. 173–192. Springer (2012)
22. Spranger, M., Pauw, S., Loetzsch, M., Steels, L.: Open-ended Procedural Semantics. In: Steels, L., Hild, M. (eds.) *Language Grounding in Robots*, pp. 153–172. Springer (2012)
23. Spranger, M.: Evolving grounded spatial language strategies. *KI - Künstliche Intelligenz* 27(2), 97–106 (2013), <http://dx.doi.org/10.1007/s13218-013-0245-4>
24. Steels, L.: Evolving grounded communication for robots. *Trends in Cognitive Sciences* 7(7), 308–312 (2003)
25. Steels, L. (ed.): *Design Patterns in Fluid Construction Grammar*. John Benjamins (2011)
26. Suchan, J., Bhatt, M., Santos, P.E.: Perceptual narratives of space and motion for semantic interpretation of visual data. In: *Proceedings of International Workshop on Computer Vision + Ontology Applied Cross-Disciplinary Technologies (CON-TACT)*. ECCV 2014 – European Conference on Computer Vision (2014)
27. Talmy, L.: *Toward a cognitive semantics. Vol. 1: Concept Structuring Systems*. The MIT Press (2000)
28. Tellex, S., Kollar, T., Dickerson, S., Walter, M.R., Banerjee, A.G., Teller, S., Roy, N.: Approaching the symbol grounding problem with probabilistic graphical models. *AI magazine* 32(4), 64–76 (2011)
29. Tenbrink, T.: *Space, time, and the use of language: An investigation of relationships*, *Cognitive Linguistics Research*, vol. 36. Walter de Gruyter, Berlin, DE (2007)