

It Is What It Tends to Do: Defining Qualitative Parameter Regions by Their Effects on Physical Behavior

Mihai Pomarlan¹, Guendalina Righetti² and John A. Bateman¹

¹*Department of Linguistics, University of Bremen, Bibliothekstraße 1, 28359 Bremen, Germany*

²*KRDB Research Centre for Knowledge and Data, Free University of Bozen-Bolzano, piazza Domenicani 3, 39100, Bolzano, Italy*

Abstract

Qualitative descriptions of object arrangements and behaviors (e.g. terms such as “near” or “fast”) can be interpreted as placing constraints on the values of the parameters they refer to, but they do so in a context-dependent way. As such, it is problematic to define such qualitative terms by value regions. Instead, in this paper we propose “functional” definitions: a qualitative term’s meaning is defined by the effect on behavior of having that quality. The resulting functional definitions are context-dependent. To obtain a more general semantics for a qualitative term, we use concept blending to combine its functional definitions coming from several contexts. We then describe how a general, functional definition can be specialized to a new context, how this can be useful in transferring existing knowledge, and illustrate this with an example, raising the possibility that appropriate specializations may be supported with image schemas.

Keywords

embodied cognition, cognitive robotics, concept blending, commonsense reasoning

1. Introduction

Interest in commonsense reasoning has grown in the past few years, especially with autonomous artificial agents moving into the physical world at scales which may pose dangers to human beings. We want our artificial agents to predict the outcome of their actions, to understand the world around them, and to do so in a way that we can check and understand ourselves.

Such computational models remain elusive, largely because it is difficult to elicit commonsense knowledge from humans – it is too obvious and too ingrained. However, we suspect there is also another problem coming from a fundamental tension between requirements placed on commonsense reasoning, whatever it turns out to be. On the one hand, it should “generalize”: commonsense principles should allow an agent mastering them to cope with situations it has not encountered before. This is, after all, the main point of investing in cognition. On the other hand, commonsense inference is extremely situated and dependent on context.

The Sixth Image Schema Day (ISD6), March 24–25, 2022, Jönköping University, Sweden

✉ pomarlan@uni-bremen.de (M. Pomarlan); guendalina.righetti@stud-inf.unibz.it (G. Righetti); bateman@uni-bremen.de (J. A. Bateman)

ORCID 0000-0002-1304-581X (M. Pomarlan); 0000-0002-4027-5434 (G. Righetti)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

Language provides many examples: a spoon may be near a cup, which contains coffee from a nearby store, which reflects the light coming from a near star called the Sun. Many orders of magnitude separate the distances mentioned, yet each was described as “near”. The same applies to every word we use to refer to some set of parameter values. “Hot” means one thing when applied to the weather, and another when applied to the cores of stars.

It has been argued that attempts to interpret spatial prepositions as describing, even qualitatively, an arrangement of objects in a context-free manner will quickly encounter problems [1]. But surrendering to context entirely and treating words such as “near” or “hot” as meaning entirely different things in different situations wastes commonalities between the various uses of a word and appears to miss cognitively relevant generalizations. We do not say “the spoon is near the cup” and “the Earth is near the Sun” merely by linguistic accident. More likely is that a familiarity with a kind of situation, such as arrangements of tableware, is leveraged to help understand another situation, such as the relative placements of celestial bodies. This transfer of knowledge from one situation to another suggests that qualitative labels for a parameter are paired with mechanisms to pick situation-appropriate regions of the parameter’s possible values. In this paper, we describe one method such a mechanism might use.

To begin, we will assume an agent has obtained, from (potentially simulated) experience, symbolic rules to predict behaviors of some situations. Such *predictive symbolic rules* map qualitative descriptions of initial arrangements to qualitative descriptions of observed behaviors, and provide an interpretable mechanism to predict the environment and adjust action. However, the rules offer a second benefit: a “functional” definition for a qualitative label, in the situations to which the rules apply. The qualitative label represents the parameter values that will make an arrangement, obeying these other qualitative constraints, *behave in this qualitative way*. Different situations, different arrangements and behaviors, will result in different, and potentially incompatible, functional definitions of the same qualitative label. To resolve these incompatibilities, we use concept-blending techniques and axiom weakening, to produce more general definitions that remain “functional” in the above sense: they define a qualitative label in terms of the more general behavior of a more general class of arrangements.

The more general functional definition of a qualitative label can then be rendered more specific once the particularities of a new situation are known. We conjecture that this will result in quicker learning for situations that are novel but in some sense similar to previously encountered ones, because the more general, functional definitions for qualitative labels and their effects on behavior will make knowledge transfer easier.

We illustrate our proposal with a running example, interpreting the qualitative label “heavy”. We assume a set of symbolic rules describing the behavior of heavy objects are already known to the agent, and describe how, via an algorithm similar to dialog-based concept blending [2], these rules can be combined into one more general definition “heavy”. We then produce a new, more specialized definition, intended to apply to a new kind of situation, i.e. one for which the agent does not yet possess a definition of heavy, nor a predictive rule.

Our example uses the Description Logic ALC, on which the current implementation concept blending operates, and so, due to expressivity limitations, does not then quite capture human intuitions. However, our approach is not limited to ALC, because the concept-blending algorithm is not limited to ALC [3]. To address the expressivity limitation, we intend to eventually employ Image Schema Logic (ISL) [4] for the writing of the predictive symbolic rules.

2. Related Work

2.1. Axiom Weakening and Concept Blending

Axiom weakening is a recently introduced technique for repairing inconsistent ontologies by *weakening*, instead of removing, their axioms. The main advantage of this technique is it preserves as much information as possible while maintaining the ontologies' consistency [5, 6, 7]. Often, axiom weakening relies on refinement operators, such as specialisation and generalisation. In [8] a concept refinement operator is introduced to generalise EL^{++} concepts in the context of concept invention. In [5] a similar line of work was extended to ALC axioms, proposed different algorithms to repair inconsistent ontologies, and analyzed their computational complexity. In [9], similar refinement operators are exploited for ontology aggregation. The previous definitions have also been extended to deal with $SROIQ$ constructs [3].

In what follows, we refer to the refinement operators formally introduced in [5]. Informally, a generalisation (resp. specialisation) operator is a function that takes a concept C and returns the set of the *super-concepts* (resp. *sub-concepts*) of C , given a reference ontology.

In our context, axiom weakening is taken into account to resolve incompatibilities that may arise when merging, or *blending*, functional symbolic rules learned in different, potentially conflicting, contexts. Since the functional rules are expressed in DL axioms, when an incompatibility arises, we exploit axiom weakening techniques to restore consistency.

As an example, our system may individuate different functional rules concerning the concept *heavy* in different situations – for instance, keeping paper on a table so that the air won't blow it away (stabilizing, preventing motion) vs. using something heavy to drive a nail through a piece of wood (causing/enabling motion against some resistance). When trying to merge functional definitions, inconsistencies may then arise in the definition of the concept *heavy*.

Related issues are analysed in the context of computational conceptual blending (CCB). As a theory of cognition, conceptual blending was proposed to model conceptual integration and creativity, which is seen as arising from the conceptual blend of different input spaces through analogical mapping [10]. Computational conceptual blending studies formal strategies to allow the integration of possibly conflicting input spaces. This process may be easy for human beings, thanks to the flexibility of human concepts, but is not trivial in AI (for an extreme case, see [11]). CCB often relies on the identification of shared structures between different input spaces, and on the identification of a generic space to steer the blending process (see e.g. [12]).

As observed above, refinement operators have been applied to conceptual blending and concept invention in [8]. [13] further carries out this line of research, by proposing a general workflow and a formal reconstruction of the conceptual blending process, including axiom weakening in the picture. Relatedly, but focusing on the literature of noun-noun combination in the context of cognitive linguistics, [2] proposes a computational treatment of *impossible combinations* in the context of formal ontologies through the procedure of axiom weakening. In line with this work, we sketch an extension of the algorithm proposed there, to manage, through generalisations allowed by the axiom weakening procedure, definitions of concepts emerging from different learning contexts (see section 5).

2.2. Natural Language Understanding and Commonsense Reasoning

There is an active search for computational implementations of commonsense reasoning, to enable artificial agents to cope with the physical and social world of humans in a human-like way. This goal however remains elusive. Even when considering simple uses of spatial language, one discovers an apparent need for extremely expressive logical techniques [14]. Further, if it is to be useful for the activities of a physical agent, symbolic inference must connect somehow to numeric descriptions of the agent's environment and actions. In the case of spatial language, such links can be provided via potential fields or probability distributions which are meant to capture the appropriateness of pairing a qualitative label (e.g. "left of") with a numerical values [15, 16, 17]. However, such probability distributions need to be tuned or learned [18] and they are highly context dependent. Despite several proposals in the literature, it remains unclear how they should best be adjusted when context changes.

There is also good evidence that human spatial descriptions are sensitive to functional aspects of arrangements of objects [19, 20], or at least, that functional aspects are a more reliable guide to predicting human language use than references to context-independent spatial arrangements, even when these are allowed to be probabilistic. This has led to efforts to formally model spatial relations at a fairly abstract, functional level [1, 21], and it has been argued that, by separating an abstract, semantic level of formal modelling from contextualization, one can avoid some of the apparent difficulties in formalizing commonsense inference [22].

Although the works cited here refer to spatial language, it appears to us that similar conclusions apply to qualitative labels in general: functional characterizations, rather than references to context-independent regions of possible values, are a better fit to how humans use language. Further, such functional characterizations open avenues for generalizing, and transferring, situation-dependent knowledge to new situations, which we investigate here.

3. From Simulated Experience to Qualitative Behavior Prediction Rules

Here we will briefly summarize our previous work [23] on how an agent can acquire symbolic prediction rules. One prerequisite is ability to act on or simulate an environment. The agent must have some vocabulary with which to describe object properties, arrangements, and behaviors, and concepts in this vocabulary are linked to "generative models" understood as joint probability distributions over qualitative labels and values for the parameter associated to the labels.

Generative models allow a bridge between symbolic, qualitative descriptions of an arrangement or behavior, and numeric descriptions of same. This is necessary because to run a simulation an agent needs to provide exact values for all parameters, and likewise the behavior observed in the (more or less continuous) world of the agent will be reported as numeric data. By using generative models, it is possible to go from qualitative to numeric descriptions via sampling, and vice-versa by asking what qualitative hypothesis best fits the observed numeric evidence. Our method in [23] samples parameter values and runs simulations for an "antecedent" – a qualitative description of an arrangement and action – and obtains "consequents" – qualitative descriptions of behavior. Antecedents are mapped to consequents resulting from the

simulations sampled from those antecedents, yielding a set of predictive symbolic rules. A heuristic selects which antecedents to sample next, based on expectations of new behavior.

The generative models themselves are not updated at any point of the method, and so they are the kind of context-independent interpretations of qualitative labels that we wish to go beyond. A problem with context-independent interpretations of qualitative labels is they make predictive symbolic rules fragile: a rule may not apply for the entire region of values that are deemed likely by the generative models attached to the qualitative terms in the rule's antecedent. Therefore, the rules and the generative models must be learned together, which is what inspired us to treat the learned predictive rules also as a way to functionally define the qualitative labels.

Another apparent limitation of the method from [23] is the use of simulation. Critics of simulation as a tool for cognition exist [24, 25], and while we do not entirely agree with them, we acknowledge that, among other problems, there are limits to what an agent can simulate. However, our method in this paper does not depend on simulation; the experience from which predictive rules have been obtained can be real experience in the physical world.

3.1. Running Example: Predictive Symbolic Rules involving 'Heavy'

Suppose we have a robotic agent attempting to understand how the physics of a household environment operates. It may be interested in cooking popcorn, or making sure papers don't get blown by the wind, or driving nails into wooden boards. Through its own experimentation, or perhaps from being told by someone else, the agent arrives at the following predictive rules, which we first express here informally for the various situations:

- (Lid) A heavy lid placed on the opening of a pot will stop popcorn from getting out;
- (Paperweight) A heavy item placed on a sheet of paper will prevent it from being blown by the wind;
- (Hammer) A hit from a heavy tool, like a hammer, will drive a nail into a wooden board.

To capture the above by expressions in a formal language is not trivial, and may require some fairly expressive logics such as ISL [4]. For our examples here we will present approximations as prediction rules formalised in ALCI (ALC with inverse roles). The inverse properties in the prediction rules are inverted again, i.e. be used as themselves, in the functional definitions in the next section, which we then feed into the concept blending algorithm.

Let us then consider the following predictive rules for the situations described above (and we assume the agent was either capable of discovering them from learning from experience and combining those results with some prior knowledge, or by being instructed):

“Lid: if something is a place where a heavy object is put, then it permits only motions that keep still relative to some container”:

$$\exists placedAt - .(\exists isQualityOf - .Heavy) \sqsubseteq \forall permitsUnimpeded. \\ (\exists relativeStillnessTo.(\exists isClassifiedBy.Container))$$

“(Paperweight): an object upon which a heavy object is placed will not be immersed and moved by a physical medium”:

$\exists placedAt-.(\exists isQualityOf-.Heavy) \sqsubseteq (\neg(\exists immersedInAndMovedBy.(\exists isClassifiedBy.Medium)))$

“Hammer: something hit by a heavy object is not kept still by some physical medium”:

$\exists placedAt-.(\exists isQualityOf-.Heavy) \sqsubseteq \neg(\exists immersedInAndKeptStillBy.(\exists isClassifiedBy.Medium))$

We assume the agent has prior knowledge, e.g. that objects exist and may play different roles in a situation, qualities belong to objects, heavy is a quality and in particular a quality relating to mass and so on. The OWL files for our running example are available online.¹

4. Numeric vs. Functional Definitions of Qualitative Labels

As briefly discussed in section 3, one can define the meaning of a qualitative label for a parameter (such as “heavy” for mass) in terms of a probability distribution over a space of possible values. Clearly however, such an understanding of a qualitative label would not match human intuition. We take spatial language to be best understood in functional terms, and intuitively the same applies also to qualitative attributes describing other physical parameters. For a human being, and one supposes, for other agents acting in the world, the important information communicated by a qualitative label is not a range of values but *a disposition for a particular behavior*. The qualitative label answers a “what does it do?”, rather than “how much is it?” question.

Therefore, it appears useful to define qualitative labels for parameters in functional terms, i.e. of resulting behaviors of arrangements in which some object is described by that qualitative label. A generic form for the predictive rules of section would be 3:

$$\forall O, Q : (Arrangement(O, \dots) \wedge isQualityOf(Q, O) \wedge x(Q)) \rightarrow Behavior(O, \dots)$$

where *Arrangement* and *Behavior* are logical formulas describing constraints on the initial state of some objects and their ensuing behavior; these formulas can have several variables but here we focus on one of them, *O*, which stands for one of the objects which has a quality *Q* of type *x*.

To obtain a functional definition for a qualitative label, we first rearrange the predictive rule:

$$\forall O, Q : x(Q) \rightarrow ((isQualityOf(Q, O) \wedge Arrangement(O, \dots)) \rightarrow Behavior(O, \dots))$$

and then stipulate that this is a definition for the label *x*:

$$\forall O, Q : x(Q) \leftrightarrow ((isQualityOf(Q, O) \wedge Arrangement(O, \dots)) \rightarrow Behavior(O, \dots))$$

Such a functional definition should be interpreted as situation specific, and then for that situation it serves as a guide to select, or learn, what parameter values correspond to the qualitative label.

4.1. Running Example: Situation-Dependent Functional Definitions

The predictive symbolic rules of section 3.1 constrain what behavior should be observed given some arrangement and object properties. By rewriting the OWL axioms such that they are now about the quality of an object, and stipulating that if an object placed in an arrangement results

¹https://github.com/mpomarlan/ISD6_HeavyBlends

in the associated behavior, then the object is heavy, we obtain the following definitions:

“Lid: heavy is the quality of objects such that wherever they are placed, that place allows only motions that keep still relative to some container”:

$$\text{Heavy} \equiv \text{MassQuality} \sqcap \forall \text{isQualityOf} . (\forall \text{placedAt} . (\forall \text{permitsUnimpeded} . (\exists \text{relativeStillnessTo} . (\exists \text{isClassifiedBy} . \text{Container}))))))$$

“Paperweight: heavy is the quality of objects such that whatever they are placed on, that object is not immersed and moved by a physical medium”:

$$\text{Heavy} \equiv \text{MassQuality} \sqcap (\forall \text{isQualityOf} . (\forall \text{placedAt} . (\neg (\exists \text{immersedInAndMovedBy} . (\exists \text{isClassifiedBy} . \text{Medium}))))))$$

“Hammer: heavy is the quality of objects such that whatever they collide with, that collided object is not kept still by a physical medium”:

$$\text{Heavy} \equiv \text{MassQuality} \sqcap \forall \text{isQualityOf} . (\forall \text{placedAt} . (\neg (\exists \text{immersedInAndKeptStillBy} . (\exists \text{isClassifiedBy} . \text{Medium}))))$$

The above functional definitions for ‘heavy’ are intended as situation specific, not as axioms that can coexist in a single, situation-independent, ontology.

5. General Functional Definitions from Concept Blending

Supposing an agent has experienced various situations, and obtained from each a definition of some qualitative label, it might ask itself whether some more general, situation-independent definition exists. It is interesting to look at combinations of the concept definitions, but not all combinations are interesting. A mere enumeration of known definitions is not interesting, because it says nothing about a situation not yet on the list. A simple conjunction of definitions may either result in an unsatisfiable concept, or at least in an unduly restricted one.

We have looked at concept blending for this, in the dialog-based approach of Righetti et al. [2]. In brief, by starting from two “incompatible” ontologies, the algorithm finds a blend (or a combined ontology) through a turn-based procedure which allows combining the axioms of the different ontologies according to a given preference order, and to weaken them until the result is satisfiable. In our case, the axioms to combine are the various situation-dependent definitions of a qualitative label known to an agent; they will all be of the form $X \equiv C$, where X is a concept name for the qualitative label we wish to define, and C is a concept expression. The set of these axioms is denoted by Q . By following [2], the agent’s background knowledge is denoted by O_{init} .

Because we are looking to combine only axioms of a given structure, as opposed to several ontologies that may contain any number of axioms, some simplifications to the concept blending algorithm are needed. In particular, we do not need preference orders on the axioms in Q . Also, while several ways of weakening an axiom are possible, we are only interested in weakenings obtained by a generalization operator γ_O applied to the concept expressions that define X .

However, inconsistency is too strong a test in our application. Consistent situation-dependent definitions are unlikely to be equivalent, and a conjunction of them will remove possibilities

from the agent’s consideration that were nonetheless feasible in particular situations. We therefore weaken to obtain more permissive axioms. That is:

- if we must combine axioms $X \sqsubseteq C$ and $X \sqsubseteq D$,
- and $C \sqcap \neg D$ or $D \sqcap \neg C$ are satisfiable under the conditions imposed by O_{init} ,
- then we replace the two axioms by $X \sqsubseteq E$, where E is a generalization of $C \sqcup D$ with respect to the background knowledge encoded in ontology O_{init}

With the above modifications, the algorithm becomes the one shown in Algorithm 1.

Algorithm 1 Combination(O_{init}, Q, X)

<pre> while $1 < Q$ do $\{X \sqsubseteq C; X \sqsubseteq D\} \leftarrow \text{RandomPickAxiomPair}(Q)$ if Satisfiable$_{O_{init}}(C \sqcap \neg D)$ or Satisfiable$_{O_{init}}(D \sqcap \neg C)$ then $E \leftarrow \gamma_{O_{init}}(C \sqcup D)$ else $E \leftarrow C$ end if $Q \leftarrow (Q - \{X \sqsubseteq C; X \sqsubseteq D\}) \cup \{X \sqsubseteq E\}$ end while return $O_{init} \cup Q$ </pre>	<p>▷ Assumption: Q contains only axioms of the form $X \equiv C$</p> <p>▷ $E = \gamma_O(C)$ is s.t. $O \models C \sqsubseteq E$</p> <p>▷ If we get here, $C \equiv D$</p>
---	--

5.1. Running Example: a General Definition of ‘Heavy’

The different definitions of “heavy” given in section 4.1 give us the set of axioms to combine using the modified concept blending algorithm, and the ontology relative to which axiom weakening is performed will be the agent’s other knowledge. We note that there is still a random component in the blending algorithm in terms of what weakening to select through the generalization operator γ , because there are many possible generalizations.

For our running example – which includes the axioms from previous sections as well as a few axioms about how certain behaviors arise from the interaction of forces, axioms which are listed in in the example’s github repository – a generic definition of “heavy” comes out as

$$Heavy \equiv MassQuality \sqcap (\forall isQualityOf. (\forall placedAt. (\exists isForcedBy. SignificantForce)))$$

or, informally, “heavy is the quality of objects that exert significant force wherever they are placed”. This definition subsumes the situation dependent ones, and can be obtained from the background knowledge in our example OWL files about when a force is significant.

6. Obtaining New Situation-Specific Functional Definitions from General Ones

A benefit of having a general definition for a qualitative label is the possibility to specialize it once a new situation is encountered. The specialized functional definition can then be employed to inform selection of parameter values for one’s own actions or understand what other agents describing this situation mean. It also can be used to generate new prediction rules for the

new situation, by reversing the rewrite procedure described in section 4 and thus, one hopes, accelerate an agent's learning of how to deal with a situation.

In our case we start from some axiom $X \equiv C$ and replace it by $X \equiv D$ where $D \sqsubseteq C$ as given by some ontology O describing the agent's background knowledge and whatever knowledge of the new situation it has obtained so far.

6.1. Running Example: Specializing Heavy to New Situations

So far, our example agent has encountered situations where heavy objects were used as lids to keep contents in, as weights to keep objects put, or as colliders to drive objects through media, and arrived at the conclusion that heavy is a quality of objects such that wherever they are placed, some significant force is exerted.

Suppose it then is told about two new situations, one in which the significant force to be considered is one that impedes the motion of a trajector, and one in which the significant force is one that prevents an object from containing parts in a stable way. Loosely speaking, these are the situations of a traveller who may be encumbered by their luggage, or of an object that may be destroyed by impact with or pressure from another. Via specialization, the following situation-specific definitions are obtained:

$$\text{Heavy} \equiv \text{MassQuality} \sqcap \forall \text{VisQualityOf} . (\forall \text{placedAt} . (\exists \text{hasMotion} . (\exists \text{impededMovingRelativeTo} . (\exists \text{isClassifiedBy} . \text{Background}))))))$$

and respectively

$$\text{Heavy} \equiv \text{MassQuality} \sqcap (\forall \text{VisQualityOf} . (\forall \text{placedAt} . (\neg (\exists \text{immersesAndKeepsStill} . (\exists \text{isClassifiedBy} . \text{Substance}))))))$$

The details are available in our example files.

6.2. The potential role of Image Schemas

Image schemas are recurring structures establishing fundamental patterns of cognition, formed since childhood from our bodily experience. In Cognitive Science, image schemas are identified as conceptual building blocks which allow reasoning about the world and move therein [26].

The role of image schema in the context of Computational Conceptual Blending has been analyzed in [4, 27]. According to [10], a blend is constructed by mapping the shared features of different (mental) input spaces into a common, *generic space* in a selective way. The blend will then develop its emergent structure, which derives from the combination of the projected features. While the selection of the relevant features is done by humans in a seemingly effortless way, it is a non-trivial problem when applied in the context of automated Computational Conceptual Blending. The selection of different features can indeed lead to quite different outcomes, and different projections can thus generate better or worst blends. Image Schemas have been thus applied in this context to steer the search for generic spaces, by helping in identifying the relevant features in the input spaces [4, 27].

The approach proposed in this paper differs from the standard approach to computational conceptual blending, because it does not rely on the identification of a shared structure between

the input spaces - i.e. it does not require the identification of a generic space to steer the combination process. However, as already mentioned, the axiom weakening procedure exploited here relies on the random selection of different possible generalizations choices, and can thus lead to different outcomes. Of course, among different possibilities, certain weakenings are more interesting than others, but at the moment the algorithm is subject to a random selection. Similar to what is done in the context of standard conceptual blending for the selection of meaningful generic spaces, image schemas could play a role here to guide the weakening procedure in the selection of interesting generalisations. This is however matter for future work.

7. Conclusions and Future Work

We have argued in this paper for the benefit of what we call functional definitions of qualitative labels for physical parameters. Rather than committing to a context- and situation-independent numerical region, functional definitions describe, in a situation-dependent way, what parameter values are appropriate. This comes from the relationship between functional definitions for a qualitative label, and behavior prediction rules where that qualitative label appears.

Further, we have illustrated how situation-dependent functional definitions for a qualitative label can be combined into a more general one via concept blending, and in turn this general definition be specialized for new situations. This allows us to formally model how an agent might, through experience, acquire something resembling human commonsense knowledge – that is, knowledge about how its environment behaves, knowledge that is situated and context-dependent, but which nevertheless allows itself to be adapted to new situations.

This process of generalization is not entirely deterministic – the concept combination algorithm often has several, equally feasible choices available during its operation. Perhaps this is to be expected, and it would be an interesting line of research to look into whether human beings can agree on general-purpose meanings for qualitative labels. We suspect that they cannot, and that some variation will exist precisely because generalization is underconstrained.

Nevertheless, generalized, functional understandings of labels such as “heavy”, “near”, “fast” etc. can be useful even if several agents do not agree on these definitions’ exact content. This is because the important feature of functional definitions as pursued in this paper is that they ask questions – e.g., if “heavy” is that which exerts a significant force, what is a significant force? These questions are the means by which the general definition can attach to aspects of a new situation and give an agent a tentative understanding of qualities that would be appropriate for it. In future work, we will look into ways to quantify this knowledge transfer, in particular with regards to how well can an agent predict behaviors in a new situation, and select parameter values for its own actions to cope with it. This can be done by comparing an agent’s performance when starting to learn to cope with a situation from scratch versus when making a tentative guess based on specializing a functional definition.

A limiting factor in the examples we have shown here is the ALC formalism, and we plan to address this by investigating how Image Schema Logic [28, 4] could be used to write the prediction rules and functional definitions for qualitative labels. Further, we expect that connecting our approach to existing ontologies of commonsense knowledge, spatial language [21], and image schemas will provide a wealth of good material for our method, as opposed to the

admittedly somewhat artificial running example shown here.

Acknowledgments

The research reported in this paper has been supported by the FET-Open Project 951846 “MUHAI – Meaning and Understanding for Human-centric AI” (<http://www.muhai.org/>) funded by the EU Program Horizon 2020 as well as the German Research Foundation DFG, as part of Collaborative Research Center (Sonderforschungsbereich) 1320 “EASE - Everyday Activity Science and Engineering”, University of Bremen (<http://www.ease-crc.org/>). The research was conducted in sub-projects *P01 Embodied semantics for the language of action and change: Combining analysis, reasoning and simulation*

References

- [1] J. A. Bateman, J. Hois, R. Ross, T. Tenbrink, A linguistic ontology of space for natural language processing, *Artificial Intelligence* 174 (2010) 1027–1071.
- [2] G. Righetti, D. Porello, N. Troquard, O. Kutz, M. Hedblom, P. Galliani, Asymmetric hybrids: Dialogues for computational concept combination, in: *Formal Ontology in Information Systems: Proc. of the 12th International Conference (FOIS 2021)*, IOS Press, 2021.
- [3] R. Confalonieri, P. Galliani, O. Kutz, D. Porello, G. Righetti, N. Troquard, Towards even more irresistible axiom weakening, in: S. Borgwardt, T. Meyer (Eds.), *Proc. of the 33rd International Workshop on Description Logics (DL 2020)*, volume 2663 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2020.
- [4] M. M. Hedblom, *Image Schemas and Concept Invention: Cognitive, Logical, and Linguistic Investigations*, Cognitive Technologies, Springer Computer Science, 2020.
- [5] N. Troquard, R. Confalonieri, P. Galliani, R. Peñaloza, D. Porello, O. Kutz, Repairing ontologies via axiom weakening, in: S. A. McIlraith, K. Q. Weinberger (Eds.), *Proc. of the Thirty-Second AAAI Conference on Artificial Intelligence*, AAAI Press, 2018, pp. 1981–1988.
- [6] F. Baader, F. Kriegel, A. Nuradiansyah, R. Peñaloza, Repairing description logic ontologies by weakening axioms, *CoRR abs/1808.00248* (2018).
- [7] J. Du, G. Qi, X. Fu, A practical fine-grained approach to resolving incoherent OWL 2 DL terminologies, in: J. L. et al. (Ed.), *Proc. of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, ACM, 2014, pp. 919–928.
- [8] R. Confalonieri, M. Eppe, M. Schorlemmer, O. Kutz, R. Peñaloza, E. Plaza, Upward refinement operators for conceptual blending in the description logic \mathcal{EL}^{++} , *Ann. Math. Artif. Intell.* 82 (2018) 69–99.
- [9] D. Porello, N. Troquard, R. Peñaloza, R. Confalonieri, P. Galliani, O. Kutz, Two approaches to ontology aggregation based on axiom weakening, in: J. Lang (Ed.), *Proc. of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018*, ijcai.org, 2018, pp. 1942–1948.
- [10] G. Fauconnier, M. Turner, Conceptual integration networks, *Cognitive science* 22 (1998) 133–187.

- [11] M. Hedblom, G. Righetti, O. Kutz, Deciphering the cookie monster: A case study in impossible combinations, in: ICCG, 2021.
- [12] M. Eppe, E. Maclean, R. Confalonieri, O. Kutz, M. Schorlemmer, E. Plaza, K.-U. Kühnberger, A computational framework for conceptual blending, *Artificial Intelligence* 256 (2018) 105–129.
- [13] R. Confalonieri, O. Kutz, Blending under deconstruction, *Ann. Math. Artif. Intell.* 88 (2020) 479–516.
- [14] E. Davis, Qualitative spatial reasoning in interpreting text and narrative. *spatial cognition and computation*. forthcoming, 2013. doi:10.1.1.295.4482.
- [15] K.-P. Gapp, Basic meanings of spatial relations: Computation and evaluation in 3d space, in: *Proceedings of the Twelfth AAAI National Conference on Artificial Intelligence, AAAI'94*, AAAI Press, 1994, p. 1393–1398.
- [16] M. Beetz, L. Mösenlechner, M. Tenorth, CRAM – A Cognitive Robot Abstract Machine for Everyday Manipulation in Human Environments, in: *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 1012–1017.
- [17] M. Pomarlan, J. A. Bateman, Embodied functional relations: a formal account combining abstract logical theory with grounding in simulation, in: *Formal Ontology in Information Systems*, IOS Press, Amsterdam, 2020. To appear.
- [18] F. Stulp, A. Fedrizzi, L. Mösenlechner, M. Beetz, Learning and Reasoning with Action-Related Places for Robust Mobile Manipulation, *Journal of Artificial Intelligence Research (JAIR)* 43 (2012) 1–42.
- [19] K. R. Coventry, S. C. Garrod, *Saying, seeing, and acting: The psychological semantics of spatial prepositions*, Psychology Press, New York, 2004.
- [20] L. Carlson, E. van der Zee, *Functional Features in Language and Space: Insights from Perception, Categorization, and Development*, 2005.
- [21] J. A. Bateman, Gum: The generalized upper model, *Applied Ontology* (2021). doi:10.3233/AO-210258.
- [22] J. A. Bateman, Space, Language and Ontology: A Response to Davis, *Spatial Cognition & Computation* 13 (2013) 295–314. doi:10.1080/13875868.2013.808491.
- [23] M. Pomarlan, M. M. Hedblom, R. Porzel, Panta rhei: Curiosity-driven exploration to learn the image-schematic affordances of pouring liquids, in: *29th Irish Conference on Artificial Intelligence and Cognitive Science (AICS)*, 2021.
- [24] E. Davis, G. Marcus, The scope and limits of simulation in cognitive models, *CoRR* abs/1506.04956 (2015). URL: <http://arxiv.org/abs/1506.04956>. arXiv:1506.04956.
- [25] E. Davis, G. Marcus, The scope and limits of simulation in automated reasoning, *Artificial Intelligence* 233 (2016) 60–72.
- [26] M. Johnson, *The body in the mind: The bodily basis of meaning, imagination, and reason*, University of Chicago press, 2013.
- [27] M. M. Hedblom, O. Kutz, F. Neuhaus, Image schemas in computational conceptual blending, *Cognitive Systems Research* 39 (2016) 42–57.
- [28] M. M. Hedblom, O. Kutz, T. Mossakowski, F. Neuhaus, *Between Contact and Support: Introducing a Logic for Image Schemas and Directed Movement*, Springer International Publishing, 2017, pp. 256–268.