
Scales and Hedges in a Logic with Analogous Semantics

H.R. Schmidtke

SCHMIDTKE@ACM.ORG

Independent / Postfach 11 01 29 Schwerin, 19001 Germany

S. Coelho

ARASCOELHO@GMAIL.COM

Faculty of Medicine, University of Lisbon / Av. Prof. Egas Moniz, 1649-028 Lisbon, Portugal

Abstract

Logics with analogous semantics, such as Fuzzy Logic, have a number of explanatory and application advantages, the most well-known being the ability to help experts develop control systems. From a cognitive systems perspective, such languages also have the advantage of being grounded in perception. For social decision making in humans, it is vital that logical conclusions about others (cognitive empathy) are grounded in empathic emotion (affective empathy). Classical Fuzzy Logic, however, has several disadvantages: it is not obvious how complex formulae, e.g., the description of events in a text, can be (a) formed, (b) grounded, and (c) used in logical reasoning. The two-layered Context Logic (CL) was designed to address these issues. Formally based on a lattice semantics, like classical Fuzzy Logic, CL also features an analogous semantics for complex formulae. With the Activation Bit Vector Machine (ABVM), it has a simple and classical logical reasoning mechanism with an inherent imagery process based on the Vector Symbolic Architecture (VSA) model of distributed neuronal processing. This paper adds to the existing theory how scales, as necessary for adjective and verb semantics can be handled by the system.

1. Introduction

Logics with analogous semantics have a number of advantages over conventional set-theoretical semantics. A primary advantage is that a semantics based on vector spaces can be related to a representation of physical reality in terms of vector spaces (Gärdenfors, 2000). In a literal sense, a spatial semantics over a coordinate space yields an image of the world. The *symbol grounding* problem (Harnad, 2003) – vital for understanding higher cognitive functions –, becomes easier with such a – multidimensional – image semantics, since an image can be compared to an image generated by means of sensory data (Schmidtke, 2020b).

While questions regarding higher cognitive functions, such as *symbol grounding* or our sense of *meaning*, used to be philosophical rather than practical questions of immediate concern, advances in autonomous vehicles have brought the lack of meaning in current computer systems, including AI systems, into the public debate. Context Logic (CL Schmidtke, 2021a,b,c, 2020b,a, 2018b,a, 2016, 2014, 2013, 2012; Schmidtke & Beigl, 2011, 2010; Schmidtke & Woo, 2009; Schmidtke et al., 2008; Hong et al., 2007) as a cognitively motivated logic with analogous semantics is a promising candidate for addressing this shortcoming. Together with its biologically inspired reasoner/imager – the *Activation Bit Vector Machine* (ABVM) it becomes feasible to deliver a concept of meaning

that more closely resembles a cognitively plausible inner world model semantics. This paper adds a discussion on scales to the theory and illustrates that CL and the ABVM can be applied to ethical problems, such as the popular Trolley Problem discussed in the literature on ethics for autonomous vehicles (e.g., Sütfield et al., 2017).

Three key aspects are important for applying the theory to the ethical domain: first, how temporal notions, as used in the literature on decision making systems and planning, can be defined in the formalism; second, how scales, including ethical value scales, can be derived; and third, how the research can be related to human social reasoning. The first point connects the theory to existing domain theories on ethical reasoning and illustrates how to add analogous semantics to knowledge-based systems based on these domain theories. The second point is relevant to illustrate that the analogous semantics are analogous and cognitively adequate in specific key aspects, such as context dependency and the ability to apply linguistic means that allow the derivation of further values on a scale. The third point allows us to suggest the theory as a fine-grained tool to, on the one hand, better understand human pathologies and, on the other hand, better analyze and predict social and jurisdictional aspects of different types of AI systems.

Conventional extensional set-theoretical semantics of first order logic (FOL) conceived as a formalization of scientific language and thought following Frege (1879) was a crucial advancement for scientific inquiry and progress (Whitehead & Russell, 1912): it provides mathematical rigidity to notions of scientific reasoning and a binary decision, true or false, sufficient for many purposes, such as when reviewing a scientific paper. From the perspective of cognitive science and commonsense reasoning, however, conventional set-theoretical semantics are unsatisfactory. Commonsense suggests that, e.g., a fictional novel creates a dynamical model within the mind of the reader, an inherent meaning. This commonsense notion of meaning is of crucial relevance as we move from lab experiments to autonomous robotic systems performing what we would call with a human actor unethical actions. We are at a crossroads: do we either expand our theory of semantics or do we re-educate our ethical sense to accept what is technically feasible under a previous paradigm?

An early example of a family of logics with atom-level analogous semantics are Fuzzy Logics (Zadeh, 1975, 1988). The original conception of Fuzzy Logic addressed the issue by assigning analogous values to statements, such as “this apple is red,” e.g., by calculating the average percentage of red in the RGB value of an area called *the apple* in a picture. The result is a value in the set $[0, 1]$ and by definition directly corresponds to reality. Accordingly, it can be used to control systems in a quasi-linguistic manner (Passino et al., 1998). There are three primary disadvantages with the original conception if we interpret it cognitively for characterizing the semantics of adjectives. First, adjectives like “red” are *context dependent*. A red wine and a red car would be imagined as having a different type of red. Likewise, we change our conception of what size is referred to as “tall” when we hear the speaker is talking about a child or a woman. Second, linguistic research on the meaning of adjectives suggests that the unary, predicative use with the positive in “ x is tall” is semantically derived from the binary, comparative use in “ x is taller than y ” (Bierwisch, 1989). Most importantly however, the analogous semantics of conventional Fuzzy Logic regards only the level of atoms. The semantics for compound sentences discussed in the literature focusses mostly on derivatives of classical $[0, 1]$ -valued semantics, that is: the meaning of a sentence is a value in $[0, 1]$, which is still far away from a dynamical model, an analogous sentence semantics. The class of semantics for Fuzzy

Logic is larger, however, as shown by Hájek (1998). All lattice structures, of which those based on the totally ordered $[0, 1]$ are only a subset, have the basic properties. However, it was less clear what practical use within the Fuzzy Logic framework there could be for arbitrary lattices, e.g., what the lattice spanned by the subsets of a finite set with the subset-relation could contribute.

But lattice structures have been the fundament of another, earlier research thrust to avoid the issues of a set-theoretical foundation of logic: *mereology*. Leśniewski's Protothetic (Srzednicki & Stachniak, 2012) aimed at building the fundamentals of logic on binary relations *part of* to fill the role of the set-theoretical subset-relation. However, the direct application of mereology to provide analogous semantics to complex statements, e.g., in the form of images resembling Venn-diagrams (Sowa, 2008) only generates visualizations of the set-theoretical semantics. The images may be of pedagogical use to learn set-theoretical semantics, but they do not resemble the represented world in the same strong sense in which Fuzzy Logic allows the reconstruction of the color of the red apple in our initial example. They also cannot help us to understand, for instance, the vivid mental model and strong feeling of disgust evoked when reading a description of how a driver was mutilated by his autonomous car.

The CL family of languages, like mereology, leverages a partial order (\sqsubseteq , which can be read as an abstract notion of “part of” as in mereology or as “sub-context”) and lattice structures at its core, but the image generation mechanism is based on an activation measurement process that has both neuronal and logical roots. The visualization capabilities arise directly from its semantics and allow, e.g., the reconstruction of complex layouts from spatial language in a uniquely defined manner. CL has been developed since 2006 as a logical language at the boundary between perception and reasoning. A wealth of results have been discovered in the past 15 years (Schmidtke, 2021a,b,c, 2020b,a, 2018b,a, 2016, 2014, 2013, 2012; Schmidtke & Beigl, 2011, 2010; Schmidtke & Woo, 2009; Schmidtke et al., 2008; Hong et al., 2007). While we do not assume any familiarity with CL in this paper and will, for purposes of completeness, outline the language, its semantics, reasoner, imager, and so on, it is obviously beyond the scope of one paper to discuss all details and prove all the theorems again. The interested reader is referred to the respective publications.

Most recently, it has been confirmed that CL can be understood as a member of the Fuzzy Logic family, and that it can even be further extended into a Fuzzy Context Logic (FCL) that unites the respective perspectives (Schmidtke, 2021c). However, many interesting questions remain open in this regard, as the connection between this new, more powerful logic and CL's analogous semantics remains to be studied. If FCL has an analogous semantics for complex formulae like CL, conventional Fuzzy Logic would have one as well.

Structure of the Article The paper starts with a short introduction of the background theories of Context Logic and the ABVM (Section 2). We then explicate how the theory can be applied in the ethical domain (Section 3), before we move to discussing the importance of ethical grounding in Section 4 by comparing two social disorders, one associated with unethical behavior (psychopathy) due to impaired grounding of ethical concepts and one not associated with unethical behavior (autism) where grounding is intact. Section 5 concludes the paper summarily emphasizing the importance of grounding.

2. Background

This section provides a brief introduction to Context Logic Section 2.1 and the Activation Bit Vector Machine (ABVM) reasoner and imager Section 2.2. Given the limited space, the reader in doubt about a particular aspect is referred to the respective previous publications (Schmidtke, 2021a,b,c, 2020b,a, 2018b,a, 2016, 2014, 2013, 2012; Schmidtke & Beigl, 2011, 2010; Schmidtke & Woo, 2009; Schmidtke et al., 2008; Hong et al., 2007).

2.1 Context Logic

Formally, CL is a two-layered logic.

1. *Context terms* \mathcal{T}_C are defined over a set of variables \mathcal{V}_C :
 - TY1 Any context variable $v \in \mathcal{V}_C$ and the special symbols \top and \perp are atomic context terms.
 - TY2 If c is a context term, then its complement ($\sim c$) is a context term.
 - TY3 If c and d are context terms then the intersection ($c \sqcap d$) and sum ($c \sqcup d$) are context terms.
2. *Context formulae* \mathcal{F}_C are defined as follows:
 - FY1 If c and d are context terms then $[c \sqsubseteq d]$ is an atomic context formula.
 - FY2 If ϕ is a context formula, then $(\neg\phi)$ is a context formula.
 - FY3 If ϕ and ψ are context formulae then $(\phi \wedge \psi)$, $(\phi \vee \psi)$, $(\phi \rightarrow \psi)$, and $(\phi \leftrightarrow \psi)$ are context formulae.
 - FY4 If $x \in \mathcal{V}_C$ is a variable and ϕ is a formula, then $\forall x : \phi$ and $\exists x : \phi$ are context formulae.

In the following, we leave out brackets as far as possible applying the following precedence: \sim, \sqcap, \sqcup for term operators and $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$ for formula operators. The scope of quantifiers is to be read as maximal, i.e., until the first bracket closes that was opened before the quantifier, or until the end of the formula. Brackets around atomic formulae are used for easier visual separation between term layer and formula layer.

The language's syntax gives rise to a hierarchy of sub-languages: $CLA \subset CL0 \subset CL1$, with the subset relation holding both syntactically as well as semantically. The CLA fragment (atomic CL) allows only \wedge , CL0 (propositional CL) allows any construction without quantifiers (FY2, FY3), CL1 (first order CL) adds quantifiers (FY4). We sketch the semantic properties axiomatically (Schmidtke, 2021c, 2012) and characterize \sqsubseteq as a partial order, i.e., as reflexive (1), antisymmetric (2), and transitive (3) (cf. Schmidtke, 2021c, for a more detailed treatment):

$$[\alpha \sqsubseteq \alpha] \tag{1}$$

$$[\alpha \sqsubseteq \beta] \wedge [\beta \sqsubseteq \alpha] \rightarrow \forall \xi : ([\alpha \sqsubseteq \xi] \rightarrow [\beta \sqsubseteq \xi]) \wedge ([\xi \sqsubseteq \alpha] \rightarrow [\xi \sqsubseteq \beta]) \tag{2}$$

$$[\alpha \sqsubseteq \beta] \wedge [\beta \sqsubseteq \gamma] \rightarrow [\alpha \sqsubseteq \gamma] \tag{3}$$

where $\alpha, \beta, \gamma, \xi$ are schema variables. The transitivity axiom (3) allows us to move any complex context term, represented with an arbitrary unused context variable x , to the right-hand side. If a context β is subcontext of a context γ , then a context x , subcontext of β must be subcontext of γ :

$$[\beta \sqsubseteq \gamma] \rightarrow ([x \sqsubseteq \beta] \rightarrow [x \sqsubseteq \gamma]) \tag{4}$$

This characterization allows us to move any complex context term to the right hand side. It is then sufficient to characterize the context term operators only with respect to their occurrence on the right hand side.¹

$$[\top \sqsubseteq \alpha \sqcap \beta] \leftrightarrow [\top \sqsubseteq \alpha] \wedge [\top \sqsubseteq \beta] \quad (5)$$

$$[\top \sqsubseteq \alpha \sqcup \beta] \leftrightarrow \forall \xi : \exists \chi : [\chi \sqsubseteq \xi] \wedge ([\chi \sqsubseteq \alpha] \vee [\chi \sqsubseteq \beta]) \quad (6)$$

$$[\top \sqsubseteq \sim \alpha] \leftrightarrow [\alpha \sqsubseteq \perp] \quad (7)$$

The formula (5) states that if the current context \top is a subcontext of the intersection of α and β , then it is both a subcontext of α and of β . In contrast, if the current context \top is a subcontext of the sum of α and β (6) it need not be completely in either α or β – e.g., Russia is in Eurasia, the sum of Europe and Asia, but not in Europe or in Asia. Thus (6) can only demand that any subcontext has a subcontext in α or β – e.g., all parts of Russia have parts that are in Europe or in Asia. If \top is in the complement of α , α is outside of the current context (7).

The converse (\sqsupseteq), equality ($=$), overlap (\circ), and a non-empty variant (\sqsubseteq) can be defined as:

$$[\alpha \sqsupseteq \beta] \stackrel{\text{def}}{\Leftrightarrow} [\beta \sqsubseteq \alpha] \quad [\alpha = \beta] \stackrel{\text{def}}{\Leftrightarrow} [\alpha \sqsubseteq \beta] \wedge [\beta \sqsubseteq \alpha] \quad (8)$$

$$[\alpha \circ \beta] \stackrel{\text{def}}{\Leftrightarrow} \neg[\alpha \sqcap \beta \sqsubseteq \perp] \quad [\alpha \sqsubseteq \beta] \stackrel{\text{def}}{\Leftrightarrow} [\alpha \circ \beta] \wedge [\alpha \sqsubseteq \beta] \quad (9)$$

The converse (\sqsupseteq) relation is useful to specify opposites, such as *north* and *south* (Schmidtke, 2020b). Two contexts α and β are equal ($=$) iff they are mutually subcontexts of each other. Overlap (\circ) between α and β means, that the intersection of α and β is not empty. The non-empty variant of \sqsubseteq (\sqsubseteq) is defined by α being part of β and overlapping it. The definitions of \sqsupseteq and $=$ (8) are still in CLA. But the use of \neg for \circ and thus \sqsubseteq (9) requires at least CL0, i.e., a CL0 reasoner (Schmidtke, 2021a). It is a basic property of lattice theory that using the conjunctive term operator \sqcap , further p.o. relations with corresponding variants can be constructed from the sublattices under elements. E.g., *spatial-part-of* or *north-of* can be modeled as p.o. relations based on contexts P and N (Schmidtke & Beigl, 2011). For any $\xi, \alpha, \beta, \gamma$, reflexivity (10), antisymmetry (11) and transitivity (12) properties of a derived relation represented by a context x follow immediately from the p.o. properties of \sqsubseteq together with the properties of \sqcap (Schmidtke, 2012).

$$[\alpha \sqcap x \sqsubseteq \alpha] \quad (10)$$

$$[\alpha \sqcap x \sqsubseteq \beta] \wedge [\beta \sqcap x \sqsubseteq \alpha] \rightarrow [\beta \sqcap x = \alpha \sqcap x] \quad (11)$$

$$[\alpha \sqcap x \sqsubseteq \beta] \wedge [\beta \sqcap x \sqsubseteq \gamma] \rightarrow [\alpha \sqcap x \sqsubseteq \gamma] \quad (12)$$

The proofs follow by (5) and (4) via (cf. Schmidtke, 2012, for details):

$$[\alpha \sqcap x \sqsubseteq \beta] \leftrightarrow [\alpha \sqcap x \sqsubseteq \beta \sqcap x] \quad (13)$$

1. The formulation is advantageous as it illustrates how a decidable conventional reasoner can operate over context logic formulae. This reasoning procedure can be considered as cognitively motivated in so far as it proceeds by zooming into a context (left hand side of \sqsubseteq), moving it to the right hand side (4), and then addressing its parts. We see that the process will terminate after finitely many steps as the number of term operators is decremented with each step (see Schmidtke et al., 2008, for a detailed proof of decidability).

The intersection of α and x is subcontext of β if and only if the intersection of α and x is subcontext of the intersection of β and x (13). This holds because any subcontext that is in the intersection of α and x is trivially also in x . We can define a *contextualization* syntax as a shorthand.

$$x : [\alpha \sqsubseteq \beta] \stackrel{\text{def}}{\Leftrightarrow} [\alpha \sqcap x \sqsubseteq \beta]. \quad (14)$$

With respect to x , α is subcontext of β , iff the subcontext of α in x is subcontext of β (14). For instance, a (physical) conference is *spatially* a subcontext of a certain convention center and *temporally* in a certain month. But containment relations are not the only partial order relations, e.g.: the resulting trajectory of a billiard ball is caused by the angle, speed, and spin at which it hit another ball, and this, in turn, was caused by the angle, speed, and position at which it was hit by the billiard cue. The reflexive variant of causation, whatever other properties causation may have, is antisymmetric and transitive, i.e., belongs to the class of a partial orders.

Together with the existential quantifier \exists , CL1 allows any other relation to be constructed including non-p.o. relations, such as the instance-of relation between an object o and a class c :

$$\exists e : [e \sqsubseteq isi] \wedge [o \sqcap e \sqsubseteq a_1] \wedge [c \sqcap e \sqsubseteq a_2]. \quad (15)$$

We can express that o is an *instance of* c in CL by saying that there is a subcontext e of *isi* (intuitively, the edge e of the graph *isi*) of so that o overlaps e in a_1 (first arguments or ends of edges) and c overlaps e in a_2 (second arguments or tips of edges). Formally, this construction is a *tuple generator*, with which arbitrary relations can be constructed.

We can write both types of relational constructions as well as other derived relations in the conventional manner of standard FOL by using Context Logic schemata, e.g.:

$$starts(\alpha, \beta) \stackrel{\text{def}}{\Leftrightarrow} [\alpha \sqcap c \sqsubseteq \beta] \wedge [\alpha \sqcap t \sqsubseteq \beta] \quad (16)$$

$$isi(\alpha, \beta) \stackrel{\text{def}}{\Leftrightarrow} \exists e : [e \sqsubseteq isi] \wedge [\alpha \sqcap e \sqsubseteq a_1] \wedge [\beta \sqcap e \sqsubseteq a_2]. \quad (17)$$

We can use square brackets to indicate a relation is a partial order, thus further shortening (14):

$$x[\alpha, \beta] \stackrel{\text{def}}{\Leftrightarrow} [\alpha \sqcap x \sqsubseteq \beta]. \quad (18)$$

We thus have derived the conventional syntax of *predicate expressions* demonstrating that these can be considered internally complex constructions. We gain the advantage of reducing the number of axioms required for basic and compound transitive relations, such as the above *starts* (see below) and have shown that CL does not replace conventional predicate logic but adds a way to further analyze and better understand its atomic formulae.

2.2 Analogous Semantics of Context Logics

If we compare classical propositional logic with FOL, a key distinction of FOL is its semantics' dependence on external set-theoretically specified structures. Truth in propositional logic, in contrast, does not depend on additional structures. A formula with n variables has exactly 2^n possible assignments.

In recent decades, a larger number of other decidable logics have been identified, especially in the family of Modal Logics and Description Logics. A particularly interesting discovery regards the status of partial order relations yielding decidable logics as well (Kieroński & Tendera, 2018). P.o.s are both historically relevant, at the heart of Intuitionistic Logic (Chagrov & Zakharyashev, 1997) and Mereology (Srzednicki & Stachniak, 2012) as well as the core of modern type and class based object-oriented reasoning systems (Brachman & Levesque, 2004). Propositional Context Logic (CL0) is the decidable fragment of Context Logic (Schmidtke, 2021c) featuring \sqsubseteq as the single partial order relation Schmidtke (2021c). Without CL0's negation, like intuitionistic negation a more powerful tool than propositional logic negation, the more basic CLA is even equivalent to propositional logic. This allows us to bring CLA reasoning down to a neuronal bit-level reasoning procedure, where we can also connect it to high-level theories of neuronal computation.

Vector Symbolic Algebras (VSA, Kanerva, 2009) are a cognitively plausible yet computationally parsimonious model of neuronal network operation (Gayler, 2006). Apart from computational advantages, such as low energy requirements, they are also a logically interesting format. The basis of a traditional VSA system are long binary random vectors – although other variants exist. Traditional VSA systems focus, on the one hand, on modelling cognitively plausible associative learning, retrieval, and reasoning mechanisms, and, on the other hand, on the replication of symbolic computing mechanisms, such as pointers and tuples. The Activation Bit Vector Machine (ABVM, Schmidtke, 2021a,b, 2020b, 2018b,a) leverages the same infrastructure, but only uses logical operations: as the operation retaining similarity, we use the bit-wise *or* (\vee), which behaves in the same way as the averaging sum on sparse vectors.

Focus and Filter If we interpret a VSA logically, we can note that it behaves similarly to a truth table, the standard semantics of Propositional Logic. If a binary vector a encodes one formula's possible models and a vector b encodes another formula's models and the vector operation $a \& b$ does not have any 1s, at all, then a and b probably contradict each other. We cannot be certain, because we did not check systematically, but if both vectors were generated in a random manner and are large enough, then there is a high probability that we encountered each possible combination, as in a truth table. Moreover, the distinct rows of the truth table will have been encountered a proportional number of times in the randomized variant. More formally, we have a Monte Carlo Simulation of the truth table method and thus a linear probabilistic SAT and #SAT reasoner. This system is very fast but like human working memory has tight limits, such as the human 7 ± 2 items (Baddeley, 1994; Miller, 1956). If we operate with binary vectors of length κ , the parameter κ corresponds to the number of Monte Carlo samples taken. With n variables in a formula, we have approximately $\kappa/2^n$ occurrences of a combination. If we chose $\kappa < 2^n$, we necessarily could not cover all combinations. We can, however, already cover the human $n = 7 \pm 2$ using the truth table method with $\kappa = 2^{7 \pm 2}$, i.e., $32 \leq \kappa \leq 512$. With the random method, the probability to miss a combination of n variables is

$$\frac{2^n}{\kappa}, \text{ i.e., for } n = 7 \pm 2 \text{ and } \kappa = 10000 : \frac{32}{10000} \leq p \leq \frac{512}{10000},$$

i.e., $0.0032 \leq p \leq 0.0512$ (Schmidtke, 2021b).

CLA reasoning maps well to logical bit vector reasoning. With a spatial intuition for \sqsubseteq as part-of between regions, we can consider any position i on vectors a as a sampling point, which can either

be inside ($a_i = 1$) or outside of a ($a_i = 0$). The operation $\&$ provides us with a focus mechanism: $a_i \& b_i = 1$ holds in those parts of b that are also in a , so $a \& b$ is the b part of a or the a part of b . Using negation (symbol: $!$), we can filter out parts we do not want to consider: $!a_i \& b_i = 1$ holds for points in b outside a . We can alternatively say $!a \& b$ subtracts a from b and obtain a filter mechanism. We can thus explain $\&$ and $!$ as arising from perceptual focus and foregrounding mechanisms. With $|$ leveraged for associative/similarity inference, all the binary operators ($!$, $\&$ and $|$) can be explained as being individually useful for fundamental perceptual and memory mechanisms.

Partial Order A corresponding region-based partial order can be derived from Propositional Logic entailment. A formula a entails a formula b , iff all assignments that make a true also make b true. In terms of sampling a and b : if all points i in a ($a_i = 1$) are also in b ($b_i = 1$) – while points outside a can be either in or outside b – then the p.o. relation holds ($[a \sqsubseteq b]$). This can be phrased in terms of negation and conjunction as $(a_i \& !b_i) = 0$ for all positions i , or, alternatively, $!(a_i \& !b_i) = 1$ for all positions i . Note the difference between vector operators ($!$, $\&$, $|$) which we derive from binary operators and which applied to two vectors yield a vector, in contrast to the entailment relation which applied to two vectors yields either *yes* or *no*, i.e., a truth value. In terms of a classical logical semantics interpretation in VSAs, the only results we thus would be interested in would be whether $(a \& !b)$ yields the 0-vector, consisting only of 0s or not. We have a cognitively plausible linear probabilistic SAT reasoner for formulae with a small number of variables, such as 7 ± 2 . This is also the bridge between the term layer and the formula layer in CL. We can describe different partial orders with the focus mechanism (Schmidtke & Beigl, 2011). With the above intuition about \sqsubseteq , we also obtain arbitrary many other partial relations, such as *spatial part of* or *north of*. We can consider $p \& x$ to mean a focus on the positions i where $x_i = p_i = 1$.

In order to encode an entire network of relations, we can use the operator $\&$ also for reading the logical operator \wedge , given that it is closely related to \sqcap (5). Note that, like for intuitionistic logic, we will need more for \neg and \vee (CLO), but for CLA the bit vector logical operations are sufficient. For a knowledge base (KB) $\phi = n : [a \sqsubseteq b] \wedge n : [b \sqsubseteq c] \wedge \dots$, we obtain $\phi = !(a \& n \& !b) \& !(b \& n \& !c) \& \dots$ as the vector encoding of the KB. We can query the KB, e.g., for $n : [a \sqsubseteq c]$ by asking for the encoding of the query $q = !(a \& n \& !c)$, whether $! \phi | q$ is the 1-vector.

Relation to Conventional Set-Theoretical Semantics Summarizing we can see that the VSA semantics is a probabilistic variant of a classical set-theoretical semantics for CLA conceived of as a fragment of FOL with a single relation \sqsubseteq . Note, that with infinitely long deterministically ordered vectors, we obtain a fragment of FOL with a unique, inherent semantics.

From Binary Vectors to Analogous Representations We can do even more with the focus and filter mechanism. If we look at $\phi \& a$ we focus on all information a formula ϕ has about the object a . If we look at $\phi \& n$ we focus on all information ϕ has about the relation n . Moreover, $\phi \& n \& a$ allows us to focus on a with respect to the n -relation. Resuming the previous example $\phi = !(a \& n \& !b) \& !(b \& n \& !c) \& \dots$ and looking at vectors $\phi \& n \& a$, $\phi \& n \& b$, $\phi \& n \& c$, we see that with ϕ we remove – i.e., set to 0 – all positions i where $\phi_i \& n_i \& a_i = 1$ and $b_i = 0$ and all positions where $\phi_i \& n_i \& b_i = 1$ and $c_i = 0$. In other words, all positions, where $\phi_i \& n_i \& c_i = 1$ have $\phi_i \& n_i \& b_i = 1$ and $\phi_i \& n_i \& a_i = 1$, and all positions, where $\phi_i \& n_i \& b_i = 1$ have $\phi_i \& n_i \& a_i = 1$. This means that there can only be more (or equally many) 1s in $\phi \& n \& a$ than in $\phi \& n \& b$, which in

turn has more 1s than ϕ and c . Generalizing, $|\phi \& n \& x|$ the number of 1s in $\phi \& n \& x$ yields a numerical representation of the ordering regarding the north-aspect of x , a rough north-coordinate (cf. Schmidtke, 2021a, 2020b,a, 2018b, for larger examples and a more detailed discussion). Obviously, we can do the same for any number of relations. That is, we can take any two relations, e.g., north and east or size, but also, e.g., the health-dimension, its related emotional and ethical dimensions as well as the two temporal dimensions, and relate them to yield analogous values. These values can be fed backwards along the perceptual pathway to components lower in the pathway. The result of reasoning can be felt or seen, that is, *imagined*.

3. Application in the Ethical Domain

In this section, we illustrate the process of domain modelling with CL, with particular focus on the ethical domain. As a concrete example, we use the popular Trolley problem (Foot, 1967). We highlight a key step of modeling with CL, the disassembly of compound relations into grounded partial order core relations for the example of the temporal domain.

3.1 Natural Language Fragment

Table 1 shows an excerpt of the small language fragment we use. We focus on simple predication and simple action sentences. To discuss a more realistic and concrete example with complex ethical and emotional dimensions we leveraged the following simple description (Schmidtke, 2021a) of the philosophical Trolley problem (Foot, 1967):

A trolley is moving down a track. If the agent pulls a lever, the trolley will move down a side track killing one person. If the agent does not pull the lever, the trolley will continue down the track killing five people.

While we do not go into the details of the ethical and emotional processing of participants in experiments on this particular ethical problem (JafariNaimi, 2018), we will return to the example and ethical and related emotional dimensions below (Sections 3.2-5). It is sufficient to note here, that pain is a core physical sensory experience related to loss of health (Table 1) and associated at a deep cognitive level, early in both development and evolution, with existential fear. The evaluation of something as causing pain, i.e., reducing health, or alleviating pain, i.e., felt as increasing health, is tied in social animals to the emotional/social dimensions of enmity and benevolence, respectively. The ethical dimensions (good and evil) are directly tied to the emotional/social dimensions. Ethical reasoning, however, requires the level of alternative scenarios (CL0), actions, and agency (CL1). At the core of ethical reasoning is reasoning about alternatives (Schmidtke, 2021a) as well as the awareness of an agent of their own position in the evaluation frame in dependence on their actions. An agent denying the existence of alternatives or their own agency in creating harm to another has cognitive issues at this higher level. However, an agent lacking the ability to tie harm to another to the emotional/social dimensions of enmity lacks a more fundamental capability found early in development and evolutionarily in many other mammals: *affective empathy*, the capability to feel the pain of the other on a sensory dimension of compassion (Batson, 2009), in the case of a participant reading the Trolley Dilemma, the grounding of abstract symbols in emotion.

Table 1. Excerpt from basic Vocabulary and Semantic Classes with Dimensional Meanings

Type	Dimension	Comp.	Type	Dimension	Comp.
Structural	-	is	Verb (i)	spatial (obj-ext.)	move
Structural	-	are	Verb (i)	spatial (obj-ext.)	run
PP (static)	north-south (+)	north	Verb (i)	spatial (obj-ext.)	go
PP (static)	north-south (-)	south	Verb (i)	spatial (obj-ext.)	continue
PP (static)	east-west (+)	east	Verb (i)	health (+, neg to avg)	recover
PP (static)	east-west (-)	west	Verb (t)	spatial (obj-ext.)	drive
PP (static)	size (+)	large	Verb (t)	spatial (subj-arm)	pull
PP (static)	size (-)	small	Verb (t)	health (+, neg to avg)	heal
PP (static)	left-right (-)	left	Verb (t)	health (-)	harm
PP (static)	left-right (+)	right	Verb (t)	health (min)	kill
PP (static)	left-right (*)	side (adv.)	Verb (t)	poss. space (subj, -)	give
PP (dynamic)	up-down (+)	up	Verb (t)	poss. space (subj, +)	receive
PP (dynamic)	up-down (-)	down	NP	-	a/the trolley
PP (dynamic)	to-from (+)	to	NP	-	a/the track
PP (dynamic)	to-from (-)	from	NP	-	an/the agent
Aspect	contains-during (+)	V-ing	NP	-	a/the lever
Tense	before-after (+)	will V	NP	-	one person
			NP	-	five people

Whether a certain direction is considered primary (+) or converse (-) direction depends on the dimension. Measures of size, for instance, have a distinct positive direction, as size is a ratio attribute (Suppes & Zinnes, 1963) with a distinct zero and positive direction determined by physical reality. Other dimensions, such as the positive direction for left-right may be determined by characteristics of the speaker/listener such as handedness or direction of writing.

Space Both spatial containment and spatial ordering dimensions can be covered by p.o.s. Moreover, the same absolute space can be conceptualized as having different dimensionality across contexts. The formalization of granular mereogeometry (Schmidtke, 2016) shows that both granularity and dimensionality can be conceived of as linearizations of the spatial containment p.o. illustrating the expressive power of the CL framework.

Linguistically, the CLA level is sufficient for simple predicates, such as “A is north of B.” For eight of the nine static PP components in Table 1 a simple p.o., $\langle dim \rangle^\sigma$, where $\sigma \in \{+, -\}$, is sufficient to capture the semantics:

$$\langle subj \rangle (is|are) \langle pred^\sigma \rangle \langle obj \rangle \mapsto \langle dim \rangle^\sigma [\langle subj \rangle, \langle obj \rangle] \quad (19)$$

E.g.: “A is north_of B.” is mapped to $north_of[A, B]$, i.e., by (18) to $north_of : [A \sqsubseteq B]$ which is equivalent to $[north_of \sqcap A \sqsubseteq B]$ by (14). The dynamic PPs occur in conjunction with verbs. Action sentences require the more powerful CL1 framework, more specifically the ability to generate contexts that describe different situations or states of the world (Schmidtke, 2021a).

Time We can describe temporal succession and containment as specified by tense (τ) and aspect (α) of a verb in terms of two dimensions. By combining ordering relations *before-after*, derived

relation name	causation (c)	containment (t)
$core(i, j)$	$i \sqcap c = j \sqcap c$	$i \sqcap t \sqsubset j$
$starts(i, j)$	$i \sqcap c \sqsubset j$	
$finishes(i, j)$	$j \sqcap c \sqsubset i$	
$icore(i, j)$	$i \sqcap c = j \sqcap c$	$j \sqcap t \sqsubset i$
$istarts(i, j)$	$i \sqcap c \sqsubset j$	
$ifinishes(i, j)$	$j \sqcap c \sqsubset i$	
$qcore(i, j)$	$i \sqcap c = j \sqcap c$	$j \sqcap t = i \sqcap t$
$qstarts(i, j)$	$i \sqcap c \sqsubset j$	
$qfinishes(i, j)$	$j \sqcap c \sqsubset i$	
$qovlc(i, j)$	$i \sqcap c = j \sqcap c$	$i \sqcap t \circ j$
$ovlc(i, j)$	$i \sqcap c \sqsubset j$	
$iovlc(i, j)$	$j \sqcap c \sqsubset i$	
$meets^*(i, j)$	$i \sqcap c = j \sqcap c$	$\neg[i \sqcap t \circ j]$
$before(i, j)$	$i \sqcap c \sqsubset j$	
$after(i, j)$	$j \sqcap c \sqsubset i$	




Figure 1. Temporal relations based on p.o. relations. Left: fifteen jointly exhaustive and pairwise disjoint simple relations from two ordering relations of potential causation (c) and temporal containment (t). Here, $i \sqcap t \circ j$ is an abbreviation for $[i \sqcap t \circ j] \wedge \neg[j \sqcap t \sqsubseteq i] \wedge \neg[i \sqcap t \sqsubseteq j]$. Right: visualized examples for orderings for potential causation (context c , circle point symbol) and interval containment (t , lines between boxes).

from causation (c), and *contains-during* derived from local temporal interval containment (t) one can obtain a jointly exhaustive and pairwise disjoint relation system similar to that of Allen (1983).² Figure 1 shows the definitions (left) and illustrates how we can think about the difference between c and t (right). While t provides interval containment (horizontal lines between boundary boxes), c contributes the ordering properties (symbolized by a circle). We can think of c as contributing a point-based representation of directedness, whereas t focusses on temporal parts and overlap. Without further axioms, the two notions are independent. Not depicted in Figure 1 on the right are the inverses (*icore*, *istarts*, and *ifinishes*) and causation variants for intervals of the same extension in t (*qcore*, *qstarts*, and *qfinishes*). Of these 15 relations 9 are transitive: *core*, *starts*, *finishes*, *icore*, *istarts*, *ifinishes*, *qcore*, *qstarts*, and *qfinishes*. For the three relations with overlap (right, middle) we could further distinguish whether any starts/finishes parts of an overlapping interval are contained or not, which would produce five instead of three relations.³ One can think about the point-based visualization for c as indicating a temporal “point of no return,” a time point after which a certain

2. Note, that we do not define causation based on *before-after*, but, just the opposite, define *before-after* to be derived from the more immediate causation. The abstract time line is – even one abstraction step further away – formally a linear extension, i.e., a linear or total ordering derived from the relation shown here. Following the CL program, we would conjecture humans learned to, and children learn to, recognize causation before developing *before-after* and before developing a concept of linear time.

3. The notion of *meets** produced by this system of relations differs considerably from, in particular, that of Allen (1983) both in meaning and role for the system. Note, moreover, that the analogous semantics handle two different relations

process, which may have been going on for some time, such as the trolley approaching the people on the main track, can no longer be stopped. This notion is highly relevant in ethical problems. For example, if the agent thinks too long about the problem while standing at the lever, the trolley’s passing is the point of no return. In an ethically unambiguous situation, where the side track is clear, we would consider the agent complicit in any injuries caused by inaction. Similarly, we can consider, e.g., global warming actions that target a date after the projected point of no return, as a choice of inaction.

We can then describe the action verb sentence semantics in terms of situation or event parameters (Schmidtke, 2021a) e_n, e_{n+1} , where e_n is the current event, e_{n+1} is the newly introduced event, and the aspect $\alpha\tau$ is one of the relations in Figure 1, e.g.: an intransitive verb with a prepositional adjunct (20) as in “a trolley is moving down a track” or a transitive verb (21) as in “the trolley kills a person.”

$$\begin{aligned} \langle subj \rangle \langle verb_i^{\sigma_v}(a, t) \rangle \langle pred^{\sigma_p} \rangle \langle obj^p \rangle \mapsto \\ \exists e_{n+1} : \alpha\tau(e_n, e_{n+1}) \wedge \dots \wedge subj(e_{n+1}, \langle subj \rangle) \wedge \langle dim_v \rangle^{\sigma_v}[e_n, e_{n+1}] \wedge \\ \langle dim_p \rangle^{\sigma_p}[e_n, e_{n+1}] \wedge obj^p(e_{n+1}, \langle obj^p \rangle) \end{aligned} \quad (20)$$

$$\begin{aligned} \langle subj \rangle \langle verb_t^{\sigma}(a, t) \rangle \langle obj \rangle \mapsto \\ \exists e_{n+1} : \alpha\tau(e_n, e_{n+1}) \wedge \dots \wedge subj(e_{n+1}, \langle subj \rangle) \wedge \\ obj(e_{n+1}, \langle obj \rangle) \wedge \langle dim_v \rangle^{\sigma_v}[e_n, e_{n+1}]. \end{aligned} \quad (21)$$

For the former example (20), the aspect $\alpha\tau(e_n, e_{n+1})$ of “is moving” resolves to the containment relation: the event of the trolley moving (e_{n+1}) contains the current situation (e_n): ($e_n \sqcap t \sqsubseteq e_{n+1}$). *Move* does not have any additional dimensional semantics ($\langle dim_v \rangle^{\sigma_v}[e_n, e_{n+1}]$), like, e.g., *enter* or *climb*, but there is a prepositional adjunct “down the track,” specifying that the next situation is *downwards* (on the track): $\langle dim_p \rangle^{\sigma_p}[e_n, e_{n+1}]$ resolves to *down*[e_n, e_{n+1}].

Other Dimensions As examples of other dimensions, the size dimension large-small illustrates that adjectives derivable from their comparative meaning (Bierwisch, 1989) fit well into the same pattern as the spatial and temporal ordering mechanisms. The health dimension illustrates the use of distinct points in the meaning of verbs, the intransitive “recover”, for instance, describes health increasing from a value below the average to the average, as does the transitive use of “heal.” A minimal state is referenced in “kill.” No distinct points are given in “harm”: an exceptionally healthy individual may be harmed and still remain above average. The possessive space referred to in the meanings of “give” and “receive”, has the subject as a given reference point from which or to which the object is moved by the action. With the dynamic PPs based on “to” and “from” respectively a corresponding goal/source of the object can be specified in the sentence.

3.2 Scales and Hedges

So far, we obtained orderings of objects with the ABVM, but one may ask whether we can generate distinct points on a scale, e.g., for the semantics of *small* as smaller than average in a context.

c and t as independent. That is, we arrive at 2D coordinates for time intervals as in the diagrammatic approach of Kulpa (2001) with the axis for t representing the length of the intervals and the axis for c representing the time line.

Moreover, can we distinguish *small* from *somewhat small* or *very small*? For the ethical domain, under what circumstances will we be horrified by a conclusion, so as to be motivated to act?

We can discuss these questions again with the above example: $\phi = !(a \& n \& !b) \& !(b \& n \& !c) \& \dots$. We get coordinate values for the minimum ($|\phi \& x \& 0| = 0$) and maximum ($|\phi \& x \& 1| = |\phi \& x|$) for a relation x and can locate a representation for the mean at $|\phi \& x|/2$. Every dimension automatically comes with such a rudimentary scale. We can thus, e.g., represent death as minimal health for the semantics of “kill,” as suggested above, and a small car as a car of less than size $|\phi \& x|/2$. The average size in this case does not require a representation of a distinctly conceptualized average sized car. We can compute a representation of a small car as $< |\phi \& x|/2$. It, inter alia, depends on the co-text and context: a large car among students may be a small car in a street scene. This is a desirable property in so far as the representation of the high amount of context dependency in human natural language use is still a challenge for NLP systems. Moreover, we do not even need to assume that the brain possesses a division mechanism. Any object that is not specifically mentioned in the text with respect to a relation, appears in the depiction at the center, because, without specification, a random number of its 1-positions are removed when a random object r , i.e., a random sequence of bits, is queried $|\phi \& r|$. If we know about an object a that it is larger than average (size context s), we can construct this situation in ϕ by generating a random vector r for one-time use and position a with $!(r \& s \& !a)$ in the larger-than-average half of the space, as the expression $!(r \& s \& !a)$ removes bit positions from ϕ that are on s but outside of a .

Generalizing, further positions on the scale can be constructed. For instance, $!(r_1 \& r_2 \& s \& !a)$ to move a closer to the average ($r_1 \& r_2$ has only 50% of the 1s of r_1 , i.e., less bits of s outside a are removed). In contrast, $!((r_1 | r_2) \& s \& !a)$ moves a further away ($r_1 | r_2$ has 50% more 1s than r_1). Turning around the relation as with opposites, we can locate an object also below the average, as for “ a is small” (= smaller than average): $!(a \& s \& !r)$. We thus also obtain hedges, like “very” and “somewhat” as well as semantically related relations that differ with respect to their position on the size scale, such as “tiny” and “huge.”

If we consider the ethical dimension of the Trolley Problem, we can now see how the system will evaluate the two options it is given: to kill more or less people. Both options cause pain to others. If actions causing pains to others evoke an unpleasant sensory experience in the system, e.g., disgust or anger at the agent of such an action, it will see itself as more or less disgusting, whatever answer it gives. Asking the system to decide between the two options has the only effect of making the system “feel” bad about itself: the position it sees itself in before the decision is higher on the moral scale than the position after any of the two actions, whatever it decides. The urge in human readers is to somehow derail the trolley and remain without having to carry the guilt of either of the bad alternatives. But we can also see, how the average response influences the decision. It is the degree to which the own decision is evaluated as below average that influences the strength of disgust. Where the average car is small, a small car is seen as an average car, while a normal-sized car may appear as large.

It is inevitable that climate change will ultimately enforce a different average level of comfort. Our legacy will be judged by this standard.

4. Discussion: Grounded Ethics

An intelligent machine that does no harm may be the dream of future drivers that have autonomous cars at hands. The inhibition of harmful deeds as a good moral practise may have its roots in the way we represent or imagine others' internal state or, in other words, the way we empathize. A sad face may be a source of sorrow and a catalyst of caring gestures whereas discovering a wicked motivation in a friend may awaken our aggressive instincts. Empathy may lead us to different actions and, thus, it is important to characterize how our imagination can put us on the right track.

Empathy comprises a wide range of different phenomena and, therefore, it is difficult to find a consensual definition (Batson, 2009). However, the current literature seems keen to state that empathy is composed by two distinct, although related, concepts: cognitive empathy and affective empathy (Davis, 1983). Cognitive empathy refers to our inferences about others' mental states or to our ability to perceive their intentions, motivations and expectations. Affective empathy concerns our ability to share an other's feelings and emotions (De Vignemont & Singer, 2006; Paulus et al., 2013). As simulation theory posits, it conveys the simulation or representation of the other's emotional experience in ourselves (Batson, 2009). Nevertheless, affective empathy does not mean emotional contagion, the vicarious feelings of an others' emotional state. In affective empathy, we are conscious of the other representational state as distinctive from our own (De Vignemont & Singer, 2006; Paulus et al., 2013).

Adopting the perspective of an other requires an act of imagination and shapes the kind of affective experience they may have (Engelen, 2011). For instance, we can imagine how the other is feeling (*imagine other*) or imagine how we would feel in the other's place (*imagine self*). These differences may be critical when it comes to adopting prosocial behaviours. In fact, neuroimaging studies have shown that brain systems supporting memory and imagination may shape empathy (Gaesser, 2013). Some authors include the emotions of compassion and distress in the definition of affective empathy (Hodges & Myers, 2007). Empathic concern or compassion, namely feeling pity when imagining an other's emotions, and empathic distress, meaning the feelings of anxiety and unease at an other's suffering, are reactions to emotional sharing. They promote pro-social behaviour, facilitating socially desirable actions, leading to caring behaviour and inhibiting harmful actions (Decety & Cowell, 2014, 2015). A question that may arise from this debate is what dimensions of empathy or acts of imagination bolster ethical behaviour. The answer to this question may be found in pathologies that present impairments on empathy, represented by an imbalance between cognitive empathy and affective empathy dimensions, such as psychopathy and autism (Smith, 2006, 2009).

Psychopathy Psychopathy is a condition characterized by anti-social behaviour, insensibility towards an other's signs of distress (e.g. sadness or fear) and incapacity to feel certain emotions such as fear, shame, guilt or remorse (American Psychiatric Association, 2013; Hare, 2003). Psychopaths show an inability to recognize distress cues and a lack of compassion or empathic concern. These deficits lead them to inflict serious harm on others, and even go through a life of crime (Blair, 1995). However, their manipulative and charming behaviour point to an understanding of others' emotions and motives, although they remain affectively callous. As a consequence, psychopaths are said to detain a high *cognitive* empathy and a low *affective* empathy (Smith, 2006). This means that they are capable of understanding their victims' mental state, despite their inability to experience their

victims' emotions. Such imbalance between cognitive empathy and affective empathy may inspire aggressive and competitive behavior but be at failure to induce cooperation and to avoid damage in others (Smith, 2006). A recent study showed that psychopaths are severely impaired in imagine-other perspective rather than in imagine-self perspective (Decety et al., 2013). In other words, they are better at representing their own emotions when imagining themselves in an other's place than at simulating an other's emotions in themselves. The failure to feel certain emotions, such as fear, prevents them from representing those emotions in their emotional system, narrowing the emotional experience (Bird & Viding, 2014).

Autism Autism is a neurodevelopmental disorder characterized by communication and relational problems, repetitive behaviour and disturbances in empathy (American Psychiatric Association, 2013). Autistic people are said to lack cognitive empathy, that is they have trouble in understanding others' mental states, motives or intentions (although it has been recognized that this condition improves with time, for instance, they become able to pass on false believe tests). Contrary to psychopaths, they are able to feel basic emotions such as happiness, anger, sadness and fear. Remarkably, in spite of a low cognitive empathy, their affective empathy is intact and, according to some authors, even surfeits (Smith, 2009). Therefore, they are able to feel another person's emotional state, although many times they can neither explain it nor understand it. In virtue of their ability to empathize with an other's emotional state, they exhibit signs of distress at others' distress and show signs of compassion and concern at others' disturbance. This imbalance between cognitive empathy and affective empathy may produce a caring and cooperative behaviour and avoid harmful tendencies (Smith, 2006). In the case of autism, there is not a sharp contrast between imagine-other perspective and imagine-self perspective (Bird & Viding, 2014), turning people more receptive to others' distress.

The model of autism empathy may be suitable to adopt if we want to build a machine that does no harm. If the machine is capable to represent in itself basic emotions such as others' fear, it will restrain itself from harmful actions. A possible objection to this architecture is that a machine conceived in this kind of framework, when facing utilitarian dilemmas, such as the trolley dilemma, probably will produce a peculiar behaviour. Some authors foresaw an enhancement in utilitarian decision-making in populations with empathy impairments, such as autism and psychopathy, for different reasons (Patil, 2015; Vyas et al., 2017).

Grounding in Disgust Disgust is an emotion characterized by avoidance of situations perceived as unclean or unhealthy. When feeling disgust, individuals tend to display lower levels of aggression and to restrain themselves from committing moral wrongdoings, as these actions are faced as unclean (Tindell, 2019). Therefore, it is expectable that psychopathic subjects, who are often involved in criminal acts and exhibit high levels of aggression, show diminished feelings of disgust (Aharoni et al., 2012). On the other hand, studies regarding individuals with autism, who present compassionate behavior, showed that they are responsive to disgusting situations (Zalla et al., 2011).

Grounding via Empathy People with autism exhibit distress at experiencing another person's suffering and show compassion, a sign that their behavior is moral (Goetz et al., 2010; Nussbaum, 2003). They even engage in practical actions designed to reduce the perceived suffering (Attwood, 2015; Smith, 2009). According to some researchers, the moral behavior performed by people with

autism is a result from a preserved affective empathy (Bollard, 2013; Smith, 2009). In contrast, people with deficits on affective empathy, such as individuals with psychopathy, show no compassion and do not engage in moral behavior (Blair, 2007).

5. Outlook and Conclusions

This article illustrated how the CL language with analogous semantics can be applied to ethical decision making. We showed how concepts of temporal reasoning from the literature on decision making and planning can be characterized in CL. We, moreover, demonstrated how the positive form of adjectives can be derived from the comparative form within the ABVM, and how we can in a similar way also represent hedges. For illustration purposes, we used a small natural language fragment and a description of an ethical dilemma to illustrate details of the lexicon mechanism for a neuronal logical reasoner realizing the computation of analogous semantics for atomic CL expressions. The article focussed on the specification of adjectives and verbs, showing that we can represent distinct scales for arbitrary sensory dimensions, with context-dependent average, minimum, maximum, and a mechanism to construct intermediate points. The results suggest that a grounded reasoner is a key component of human working memory (WM) required for higher cognition. We conjecture, moreover, that the many benefits of grounded reasoning created sufficient evolutionary pressure to further enlarge WM capacity, but that the complexity limits of the core #SAT reasoner presented a strict boundary, leading to the evolution of additional step-wise depth-first reasoning mechanisms, capable of representing alternatives (\vee , \neg , CL0, which has a decidable reasoning procedure) and object persistence (\exists , CL1 or FOL).⁴

We compared the notion of grounding in the proposed logical cognitive system with respect to the domain of ethical reasoning in human beings, where a lack of grounding of ethics in empathy has been proposed to be a key component for psycho-social disorders leading to unethical decision making and studied in depth for public health purposes. Interpreting these results into the proposed architecture, we can model the ethical notion of *evil* as grounded in the sensory dimensions of disgust and, via empathy, pain. Generalizing, we can conclude that current autonomous vehicles even with ethical logics added will, without analogous semantics, be likely to exhibit anti-social behavior similar to that of humans with impaired grounding of ethical concepts. Leveraging logics with analogous semantics could help in this regard. Given the scalability issues of a core #SAT mechanism for empathy and the high scalability of perceptual capacities, we would caution that large-scale AI systems would suffer from an imbalance of considerably higher perceptual and reactional capabilities than capacity for benevolence and compassion. The current global race to develop the most powerful reactive AI systems for centrally processing vast amounts of perceptual user data may come at a dire price. Small-scale AI, in contrast, could excel in this regard.

4. The CL0/CL1 processing is detailed in Schmidtke (2021a). A video of the system processing the full Trolley example can be viewed at: <https://logical-lateration.appspot.com/video.html>. The video shows how the CL0/CL1 reasoner circumvents the limitations of the CLA reasoner starting from a clean slate in every exploration of an alternative, thus producing a sequence of images, an inner movie, for the visuospatial domain rather than a static image.

Acknowledgments

The authors wish to thank the Hanse-Wissenschaftskolleg, Delmenhorst for funding and supporting this collaboration. We are grateful to the four reviewers of this paper for valuable comments.

References

- Aharoni, E., Sinnott-Armstrong, W., & Kiehl, K. A. (2012). Can psychopathic offenders discern moral wrongs? A new look at the moral/conventional distinction. *Journal of abnormal psychology, 121*, 484.
- Allen, J. F. (1983). Maintaining knowledge about temporal intervals. *Communications of the ACM, 26*, 832–843.
- American Psychiatric Association (2013). *Diagnostic and statistical manual of mental disorders (5th ed.)*. Technical report, American Psychiatric Association, Washington DC.
- Attwood, T. (2015). *The complete guide to asperger's syndrome, revised edn*. London: Jessica Kingsley.
- Baddeley, A. (1994). The magical number seven: Still magic after all these years? *Psychological Review, 101*.
- Batson, C. D. (2009). These things called empathy: eight related but distinct phenomena. In J. Decety (Ed.), *The social neuroscience of empathy*, 3–15. MIT press.
- Bierwisch, M. (1989). The semantics of gradation. *Dimensional adjectives, 71*, 261.
- Bird, G., & Viding, E. (2014). The self to other model of empathy: Providing a new framework for understanding empathy impairments in psychopathy, autism, and alexithymia. *Neuroscience & Behavioral Reviews, 47*, 520–532.
- Blair, R. J. R. (1995). A cognitive development approach to morality: investigating a psychopath. *Cognition, 57*, 1–29.
- Blair, R. J. R. (2007). Empathic dysfunction in psychopathic individuals. In F. D. Farrow & R. Woodruff (Eds.), *Empathy in mental illness*, 3–16. Cambridge: Cambridge University Press.
- Bollard, M. (2013). Psychopathy, autism and questions of moral agency. In A. Perry & A. Yankowski (Eds.), *Ethics and neurodiversity*, 238–259. Cambridge: Cambridge Scholars Press.
- Brachman, R. J., & Levesque, H. J. (2004). *Knowledge representation and reasoning*. Morgan Kaufmann.
- Chagrov, A., & Zakharyashev, M. (1997). *Modal logic*. Oxford University Press.
- Davis, M. (1983). Measuring individual differences in empathy: evidence for multidimensional approach. *Journal of Personality and Social Psychology, 44*.
- De Vignemont, F., & Singer, T. (2006). The empathic brain: how, when and why? *Trends in cognitive sciences, 10*, 435–441.

- Decety, J., Chen, C., Harenski, C., & Kiehl, K. A. (2013). A fmri study of affective perspective taking in individuals with psychopathy: imagining another in pain does not evoke empathy. *Frontiers in Human Neuroscience*, 7, 1–12.
- Decety, J., & Cowell, J. M. (2014). The complex relation between morality and empathy. *Trends of Cognitive Science*, 18, 337–339.
- Decety, J., & Cowell, J. M. (2015). Empathy, justice, and moral behavior. *AJOB neuroscience*, 6, 3–14.
- Engelen, E.-M. (2011). Empathy and imagination. *American Philosophical Association – Pacific Division Meeting* (pp. 1–13).
- Foot, P. (1967). The problem of abortion and the doctrine of the double effect. *Oxford Review*, 5.
- Frege, G. (1879). *Begriffsschrift, eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. L. Nebert.
- Gaesser, B. (2013). Constructing memory, imagination, and empathy: a cognitive neuroscience perspective. *Frontiers in Psychology*, 3, 1–6.
- Gärdenfors, P. (2000). *Conceptual spaces: The geometry of thought*. Cambridge, MA: MIT Press.
- Gayler, R. W. (2006). Vector Symbolic Architectures are a viable alternative for Jackendoff’s challenges. *Behavioral and Brain Sciences*, 29, 78–79.
- Goetz, J. L., Keltner, D., & Simon-Thomas, E. (2010). Compassion: an evolutionary analysis and empirical review. *Psychological bulletin*, 136, 351–374.
- Hájek, P. (1998). *Metamathematics of fuzzy logic*, volume 4 of *Trends in Logic*. Springer Science & Business Media.
- Hare, R. D. (2003). *Manual for the revised psychopathy checklist*. Technical report, Multi-health systems.
- Harnad, S. (2003). The symbol grounding problem. In *Encyclopedia of cognitive science*. Macmillan and Nature Publishing Group.
- Hodges, S. D., & Myers, M. W. (2007). Empathy. *Encyclopedia of social psychology*, 1, 297–298.
- Hong, D., Schmidtke, H. R., & Woo, W. (2007). Linking context modelling and contextual reasoning. *4th International Workshop on Modeling and Reasoning in Context (MRC)* (pp. 37–48). Roskilde University.
- JafariNaimi, N. (2018). Our bodies in the trolley’s path, or why self-driving cars must *not* be programmed to kill. *Science, Technology, & Human Values*, 43, 302–323.
- Kanerva, P. (2009). Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors. *Cognitive computation*, 1, 139–159.
- Kieroński, E., & Tendera, L. (2018). Finite satisfiability of the two-variable guarded fragment with transitive guards and related variants. *ACM Transactions on Computational Logic (TOCL)*, 19, 1–34.
- Kulpa, Z. (2001). Diagrammatic representation for interval arithmetic. *Linear Algebra and its Applications*, 324, 55–80.

- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review*, 63, 81.
- Nussbaum, M. C. (2003). *Upheavals of thought: The intelligence of emotions*. Cambridge: Cambridge University Press.
- Passino, K. M., Yurkovich, S., & Reinfrank, M. (1998). *Fuzzy control*. Addison-Wesley.
- Patil, I. (2015). Trait psychopathy and utilitarian moral judgement: The mediating role of action aversion. *Journal of Cognitive Psychology*, 27, 349–366.
- Paulus, F. M., Müller-Pinzler, L., Westermann, S., & Krach, S. (2013). On the distinction of empathic and vicarious emotions. *Frontiers in Human Neuroscience*, 7, 196.
- Schmidtke, H. R. (2012). Contextual reasoning in context-aware systems. *Workshop Proceedings of the 8th International Conference on Intelligent Environments* (pp. 82–93). IOS Press.
- Schmidtke, H. R. (2013). Path and place: the lexical specification of granular compatibility. In M. Dimitrova-Vulchanova & E. van der Zee (Eds.), *Motion encoding in language and space*, Explorations in Language and Space, 166–185. Oxford University Press.
- Schmidtke, H. R. (2014). Context and granularity. In P. Brézillon & A. Gonzalez (Eds.), *Context in computing: A cross-disciplinary approach for modeling the real world*, 455–470. Springer.
- Schmidtke, H. R. (2016). Granular mereogeometry. *Formal Ontology in Information Systems: Proceedings of the 9th International Conference (FOIS 2016)* (pp. 81–94). IOS Press.
- Schmidtke, H. R. (2018a). A canvas for thought. *Procedia Computer Science*, 145, 805–812.
- Schmidtke, H. R. (2018b). Logical lateration – a cognitive systems experiment towards a new approach to the grounding problem. *Cognitive Systems Research*, 52, 896 – 908.
- Schmidtke, H. R. (2020a). Logical rotation with the Activation Bit Vector Machine. *Procedia Computer Science*, 169, 568–577.
- Schmidtke, H. R. (2020b). Textmap: A general purpose visualization system. *Cognitive Systems Research*, 59, 27–36.
- Schmidtke, H. R. (2021a). Multi-modal actuation with the Activation Bit Vector Machine. *Cognitive Systems Research*, 66, 162 – 175.
- Schmidtke, H. R. (2021b). Reasoning and learning with Context Logic. *Journal of Reliable Intelligent Environments*, 7, 171–185.
- Schmidtke, H. R. (2021c). Towards a fuzzy context logic. In C. Volosencu (Ed.), *Fuzzy systems*. IntechOpen.
- Schmidtke, H. R., & Beigl, M. (2010). Positions, regions, and clusters: Strata of granularity in location modelling. *KI 2010* (pp. 272–279). Springer.
- Schmidtke, H. R., & Beigl, M. (2011). Distributed spatial reasoning for wireless sensor networks. *Modeling and Using Context* (pp. 264–277). Springer.
- Schmidtke, H. R., Hong, D., & Woo, W. (2008). Reasoning about models of context: A context-oriented logical language for knowledge-based context-aware applications. *Revue d’Intelligence Artificielle*, 22, 589–608.

- Schmidtke, H. R., & Woo, W. (2009). Towards ontology-based formal verification methods for context aware systems. *Pervasive 2009* (pp. 309–326). Springer.
- Smith, A. (2006). Cognitive empathy and emotional empathy in human behavior and evolution. *The Psychological Record*, *56*, 3–21.
- Smith, A. (2009). The empathy imbalance hypothesis of autism: a theoretical approach to cognitive and emotional empathy in autistic development. *The Psychological record*, *59*, 489–510.
- Sowa, J. F. (2008). Conceptual graphs. In F. van Harmelen, V. Lifschitz, & B. Porter (Eds.), *Handbook of knowledge representation*, 213–237. Elsevier.
- Szrednicki, J. J., & Stachniak, Z. (Eds.). (2012). *Leśniewski's systems protothetic*, volume 54 of *Nijhoff International Philosophy Series*. Springer Netherlands.
- Suppes, P., & Zinnes, J. (1963). Basic measurement theory. In R. Luce, R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology*, 1–76. New York: John Wiley & Sons.
- Sütfeld, L. R., Gast, R., König, P., & Pipa, G. (2017). Using virtual reality to assess ethical decisions in road traffic scenarios: Applicability of value-of-life-based models and influences of time pressure. *Frontiers in Behavioral Neuroscience*, *11*, 122. From <https://www.frontiersin.org/article/10.3389/fnbeh.2017.00122>.
- Tindell, C. N. (2019). *Examining psychopathy: Relationships with disgust, moral severity, and gender*. Master's thesis, Texas Woman's University.
- Vyas, K., Jameel, L., Bellesi, G., Crawford, S., & Channon, S. (2017). Derailing the trolley: everyday utilitarian judgments in groups high versus low in psychopathic traits or autistic traits. *Psychiatry research*, *250*, 84–91.
- Whitehead, A. N., & Russell, B. (1912). *Principia mathematica*. University Press.
- Zadeh, L. A. (1975). Fuzzy logic and approximate reasoning. *Synthese*, *30*, 407–428.
- Zadeh, L. A. (1988). Fuzzy logic. *Computer*, *21*, 83–93.
- Zalla, T., Barlassina, L., Buon, M., & Leboyer, M. (2011). Moral judgment in adults with autism spectrum disorders. *Cognition*, *121*, 115–126.