

**Why Radical Enactivism is not Radical
Enough:
A Case for Really Radical Enactivism**

J.C. van den Herik

Master Thesis

Theoretical Philosophy

Faculty of Philosophy

Erasmus University of Rotterdam

May 6, 2014

Number of words: 24946

Student number: 303354

Supervisor: prof.dr. F.A. Muller

Advisor: prof.dr. J. de Mul

Contents

Preface	iv
Thesis Outline	vi
List of Figures	vii
List of Tables	vii
Abbreviations	vii
1 Radical Enactivism	1
1.1 Cognitivism	1
1.2 Enactivism	2
1.3 Radical Enactivism and Other Enactivisms	4
1.4 A Very Short History of Radical Enactivism	5
1.5 The Notion of Content	6
1.6 The Problem of Naturalising Content	6
1.7 The Hard Problem of Content	7
2 Basic Minds and Superminds	9
2.1 Basic Minds: From Teleosemantics to Biosemiotics	9
2.1.1 Why Teleosemantics cannot Account for Content	9
2.1.2 Biosemiotics: Teleosemantics Without the Semantics	10
2.1.3 Going Off-line: Imaginative Recreations	11
2.2 Linguistic Superminds	13
2.2.1 Davidsonian Interpretationalism	13
2.2.2 Thinking in Natural Language	15
2.2.3 The Linguistic Supermind	16
3 Linguistic mediation, propositional attitudes, and propositions	18
3.1 The Necessity of Linguistic Mediation for Propositional Attitudes	18
3.2 Language as a Medium and as an Intentional Object	21
3.3 The Nature of Propositions	23

4	Why Language is not an Abstract Symbol System	25
4.1	Language as an Abstract Symbol System	25
4.2	Davidsonian Interpretationalism and ‘How it Works’	28
4.3	Objections to an Abstract Symbol System-View	29
4.3.1	Repetition and Reification	29
4.3.2	From Denotation to Connotation	31
4.3.3	Language and Dynamics	34
4.3.4	Language as an Extension of Social Behaviour	36
4.4	The Complex Interplay of Symbols and Dynamics	39
4.4.1	First-Order Linguaging and Second-Order Language	39
4.4.2	Linguistic Symbols as Replicable Constraints	41
5	Conclusion	44
	References	47

Preface

As philosophers, we are in the business of promoting possibilities.

Hutto & Myin (2013, x–xi)

Human beings are capable of a truly extraordinary range of interactions with their environments and with each other. We can track and catch a ball that somebody throws towards us, enjoy the soothing experience of classical music, navigate through a crowd of people, discuss the meaning of life, etc. How on Earth do we do it? The phenomena that are picked out here by examples are collectively known as *cognition*, which comprises such processes as perception, memory, reasoning, calculating, problem solving, decision making, and using and understanding language. When we ask ourselves how cognition is possible, we ask what processes underlie these abilities. In other words, we want to know what makes intelligent behaviour possible.

Traditionally, philosophers have assumed that any form of cognition presupposes the ability to *represent* the world. Only when a creature knows *that* something is the case, can it act accordingly. And only when the creature's representation of the world is true (or accurate), will its actions be effective. Cognition was thus thought to be crucially dependent on the acquisition of knowledge. The main intuition behind this conviction is that thinking is always thinking *about* something, and the only way thinking can be about something is through representing it. *Cognitivism* is the main proponent of this representational approach to cognition.

Over the past twenty-five years, however, a revolutionary spirit has been going through the cognitive sciences and the philosophy of mind. According to Ramsey (2007, xv), 'cognitive science has taken a dramatic anti-representational turn'. Perhaps the most radical anti-representational account of cognition is Radical Enactivism (Hutto & Myin, 2013; Hutto, 2006e), according to which the emergence of cognitive capabilities has to be explained in terms of the dynamic unfolding of whole-bodied active engagement with the environment. Non-linguistic, or *basic cognition*, does not rely on knowing that certain things are the case, but has to be explained in terms of a history of previous interactions which have shaped the organism's structure. *Thinking* is still conceived to be about the world, but this aboutness, or intentionality, is not a property of some inner mental state (the representation), but is rather expressed in a creature's behaviour.

At the same time Radical Enactivism draws an upper limit to explanations of cognition in non-representational terms. Linguistic cognition is thought to rely on using public symbols that are used to 'stand for' something else, and are therefore representations. To think of language in representational terms it to think of speaking a language as consisting primarily in the activity of stating true things about states of affairs, and take these as premises in acts of reasoning. Language is conceived of as an individual achievement that is dependent on knowing the meaning of symbols. Linguistic symbols are thought to be *detached* representations, that is, they are disconnected from more basic modes of interaction with the world and each other.

The meaning that a person ascribes to a symbol can be ascertained through a process of *radical interpretation*. By prompting assent or dissent from a person, both the logical relations between words and what they are about is uncovered. In order to get to know the meaning of a symbol, we therefore have to abstract away from the concrete act of uttering

it, as the logical relations between the symbols can only be determined on the basis of patterns of use. This representational mode of thinking is thought to be so distinct from basic modes of responding that they are envisaged to take place in a separate mind: the linguistic supermind, a ‘softwired virtual machine’. To deny the representational nature of linguistic cognition is to advance *Really Radical Enactivism*, which is a bridge too far for Radical Enactivism.

In this thesis, I challenge this distinction between basic and linguistic cognition in order to pave the way for Really Radical Enactivism. I take Hutto’s Radical Enactivism as a starting point, as it promotes a radical non-representational view of basic cognition, but extend this non-representational view to language. The main argument is that the distinction is based on a particular view of language that turns out to be untenable on closer inspection. I argue that this semio-centric view on human language robs us of the possibility to explain how our linguistic capabilities work. The public symbols of natural language are not representations, but evoke a history of dynamic interaction with the world and with each other, and therefore cannot be disconnected from basic cognitive capacities. Language emerges from more basic social coordination and remains constitutively dependent on these basic forms of interaction. Rather than representing the world, the primary function of language is *doing* things with words. If we want to know *how* this is possible, we have to look at the concrete situations of linguistic communication, which involve a lot more than the mere exchange of symbols (think for instance of gestures, facial expressions, voice dynamics). To try to fix the meaning of a symbol through a process of radical interpretation is to conflate the results of what we do with symbols with properties of the symbols themselves.

Once we realise that speaking is an extension of social behaviour, there is no principled distinction between basic and linguistic cognition in terms of representation (see Table 1). This realisation opens the door for Really Radical Enactivism.

	Cognitivism	RadEn	RRadEn
Basic cognition is representational	✓	X	X
Linguistic cognition is representational	✓	✓	X

Table 1: Representation in Cognitivism, Radical Enactivism (RadEn), and Really Radical Enactivism (RRadEn).

The conceptual shift from Radical Enactivism to Really Radical Enactivism in thinking about linguistic cognition thus mirrors the conceptual shift from Cognitivism to Radical Enactivism in thinking about basic cognition. The representational presuppositions turn out to be untenable on closer inspection, as we turn our attention to what is actually going on in the concrete unfolding of basic or linguistic interaction with the world and each other that we want to explain.

Thesis Outline

The main aim of Section 1 is to get to grips with the radical element of Radical Enactivism (RadEn): the conviction that basic cognition is to be explained in a thoroughly non-representational fashion. I first contrast the Enactivist approach with the traditional representational approach to cognition. Thereafter, I introduce RadEn's notion of representation in terms of content. I then describe RadEn's worries about naturalising content by relying on the notion of *information*, in terms of which the contemporary debate is set. I end this section with Hutto and Myin's formulation of the Hard Problem of Content.

In Section 2, I introduce Radical Enactivism's distinction between basic, contentless, cognition and linguistically mediated thought, which is mirrored in the distinction between basic minds and linguistic superminds. Basic cognition is described in terms of biosemiotics: an Enactivist makeover of teleosemantics—an approach that accounts for content in terms of functions of organisms. I further describe how RadEn describes cognition that deals with objects that are absent in terms of recreative imaginings. I then introduce Hutto's account of linguistic content, which is based on Davidsonian interpretationalism: content can only be ascribed through a process of radical interpretation, which functions as an *intensionality test* that can only be passed by the fine-grained responding of linguistically competent creatures .

In Section 3, I raise some initial worries for RadEn's distinction between basic and linguistic cognition. I argue that linguistic mediation is not necessary for adopting a propositional attitude *if* we accept the intensionality test as a criterion. Thereafter, I argue that even if we adopt a linguistically mediated propositional attitude, the symbols of natural language are not the intentional objects of our attitudes. It is only when we talk *about* language that we shift our attention to the linguistic symbols and they become intentional objects. After that, I argue that, using the the resources available to RadEn, it will be impossible to make sense of the notion of a proposition.

In Section 4, I trace the problems stated in Section 3 to problematic assumptions about language. I argue that the root problem is that Hutto conceives of our linguistic competence as being made possible by an abstract system of symbols, and that conceiving of language in this way robs us of the possibility to explain how language works. I show that Davidson, who is the main source of inspiration for Hutto's account of language, never meant his theory of meaning as a theory that could also explain our linguistic abilities. Although Hutto extends the Davidsonian approach to include the acquisition of words through *social triangulation*, this extension does not explain how language works. After that I give some general objections to thinking of language in terms of an abstract symbol system. I argue that the main function of language is not to represent the world, but to coordinate behaviour.

These considerations will provide the basis for showing the implications of relinquishing an abstract symbol system view of language and the alleged distinction between the basic and supermind, which I consider in the final Section of this thesis.

List of Figures

1	Comparison of the Cognitivist linear and modular view of cognitive processes, according to which sensory (<i>S</i>), cognitive (<i>C</i>) and action (<i>A</i>) processes are sequential and distinct, and the sensorimotor loop. The sensorimotor loop is based on Stewart (2010, Figure 1.1, 3).	4
2	The referential triangle, in which both the child and the teacher respond similarly to a worldly object (solid lines), the teacher recognises that the child's responses are similar to his own, and the child in its turn recognises that its responses are similar to the teacher's (dashed lines). Based on Hutto (1999, Figure 17, 118).	28
3	Semantic triangle. Adapted from Kravchenko (2007, Fig. 1, 656).	32
4	Kanzi using the lexigram board to communicate with Sue Savage-Rumbaugh. Photo: greatapetrust.org.	38

List of Tables

1	Representation in Cognitivism, Radical Enactivism (RadEn), and Really Radical Enactivism (RRadEn).	v
---	--	---

Abbreviations

ActCR	Action Coordination Routines
CR_{Re}Im	Complex Recreative Imagining
LoT	Language of Thought
MoP	Mode of Presentation
RadEn	Radical Enactivism
ReIm	Recreative Imagining
TNL	Thinking in Natural Language Hypothesis

1 Radical Enactivism

The aim of this Section is to get to grips with the radical element of Radical Enactivism (RadEn): the conviction that basic cognition is to be explained in a thoroughly non-representational fashion. As Enactivism was conceived as an alternative to Cognitivism, I begin this Section by sketching it (§1.1). Thereafter, I introduce Enactivism (§1.2). I then discuss RadEn’s place in the Enactivist landscape by showing its differences to other strands of Enactivism (§1.3). After this, I show how RadEn challenges representational accounts of cognition. This challenge is put forward in terms of content, which I introduce first (§1.5). Thereafter I describe the problems RadEn diagnoses in naturalising content by relying on the concept of information (§1.6). Finally, I end this section with the challenge that Hutto and Myin formulate for anybody who wants to explain cognition in terms of content. They coin this challenge the Hard Problem of Content (§1.7).

1.1 Cognitivism

The founders of Enactivism, Varela et al. (1991), trace the Cognitivist approach to the formative years of cognitive science (1943–1953), when it was still called *cybernetics*. At the time digital computers and information theory were formed, the cyberneticians wanted to create a science of the mind in which mental phenomena were explained in terms of mathematical formalisms and mechanisms. They put forward the idea that mental activity is best understood in computer-terms. Following in their footsteps, the early Cognitivists in the second half of the 1950s took computational processes operating on symbolic representations as the defining characteristic of cognitive processes. The key notion in explaining cognitive processes is therefore *representation*: mental symbols represent *because* they ‘stand for’ certain aspects of the world. *Re*-presenting the world in certain ways consists in the application of conceptual templates or representational schema’s to sensory inputs, which are usually conceived of in terms of information (Hutto & Myin, 2013, 13). To perceive a tomato as round and red, for instance, can be explicated by the application of the mentalistic concepts ‘round’ and ‘red’ to the sensory inputs. This approach implies that intelligent behaviour presupposes the ability to truthfully represent the world as being certain ways: only when the tomato is actually round and red will the actions that are based on these representations be effective.

At the same time, the digital computer was taken as a mechanical model for the human brain, where cognition was located. The link between the physical operation—to which digital computers are limited—and the semantic level of the symbolic representations was explained by the encoding of semantic distinctions between symbols in the syntax of a symbolic language of thought. In this way, the physical operation of a brain can be constrained by the semantic value of the physical symbols, without a direct manipulation of the semantic content of these symbols. This point of view is captured in the slogan ‘no computation without representation’, for it is only when these syntax-driven computations are effectively constrained by *semantic*, that is, representational encodings, that we would describe these processes *as* computations. If they were not semantically constrained, these ‘computations’ would be merely arbitrarily juggling around of symbols.

The Cognitivist view thus presupposes three irreducible levels (Varela et al., 1991):

1. the physical level, which is described by physics, biochemistry and neurobiology;
2. the symbolic level, which is not reducible to the physical level because the relation between symbols and their physical realisation is arbitrary;
3. the semantic or representational level, which is also irreducible, because the same semantic value can be realised in different symbolic forms.

Accordingly, Varela et al. (1991, 42) give the following characterisation of Cognitivism:

Question 1: What is cognition?

Answer: Information processing as symbolic computation—rule-based manipulation of symbols.

Question 2: How does it work?

Answer: Through any device that can support and manipulate discrete functional elements—the symbols. The system interacts only with the form of the symbols (their physical attributes), not their meaning.

Question 3: How do I know when a cognitive system is functioning adequately?

Answer: When the symbols appropriately represent some aspect of the real world, and the information.

As Cognitivism turned to the digital computer as a model for human cognition, thereby detaching the symbolic level from its material instantiation, the sensorimotor apparatus was distinguished from the truly cognitive processes which take place in a central cognitive core. There was no longer any need to include eyes, ears, or noses into accounts of cognition, as they only form the intermediaries between the world and cognition—understood as computations on symbols. The motor system faced a similar fate, for our bodies are for manipulating physical objects, not the symbols of which cognition consists (Hutchins, 1995).

On this Cognitivist view, the mind is pictured as passively awaiting sensory input from its environment. Cognitive processes structure these inputs, and, based on the thus-obtained representations of the world, combined with goals or intentions, control the motor processes. The relevant causal flows are therefore conceived of as linear, starting with perceptual processes, then moving on to central cognitive processes where the representational symbols reside, and finally leading to action. Hurley (1998) coins this view of perceptual input modules, central cognitive processes, and behavioural output modules the *Classical Sandwich* model of cognition.

1.2 Enactivism

Enactivism was first put forward by Varela et al. (1991) as an alternative to Cognitivism, which they thought had ‘virtually nothing to say about what it means to be human in everyday, lived situations’ (Varela et al., 1991, xv): the abstract formalisms of Cognitivism do not do justice to our everyday experience. Enactivism’s main aim is to show how cognitive structures emerge from the embodied processes that constitute a history of recurrent interactions between organisms and their environment. Thinking is something we, as embodied organisms, *do*. It depends constitutively on our bodies and our environment. We are not brains in vats, which can be described in isolation, but living bodies in a world (Cosmelli & Thopson, 2010). Besides being *embodied*, cognition is also *extended*:

an explanation of our cognitive abilities has to feature both our bodies and their environments. Rather than describing cognition in abstract symbolic terms, Enactivism thus aims for a *re-enchantment of the concrete* (Varela, 1995).

Enactivism is part of a larger embodied and extended movement in the cognitive sciences (Kiverstein & Clark, 2009), which wants to overturn Cognitivism’s hegemony. Approaches in this movement differ in how they conceive of the claim that cognition is embodied. For some, embodied and extended cognition provides new conceptual tools to extend the framework of Cognitivism—for instance by also allowing bodily encoding of representations alongside neural ones—, whereas others feel that a more radical break with tradition is required (Menary, 2010). What sets Enactivism apart is that it *explicitly* wants to incorporate from the start consciousness, emotion, and interaction with the world, without abstracting from biological details. The Enactivist interpretation of embodiment is thus of the radical sort. The classical distinction between functions and implementations (e.g. Putnam’s (1967) idea of the *multiple realizability*) is denied, which entails that cognition cannot be understood in separation from its biological constitution. The arbitrary relation between the three levels (physical, symbolic, semantic) which Cognitivism relies on to explain our thinking is denied.

To oppose the representational commitments of Cognitivism, Enactivism denies the modularity of perception, cognition, and action. There are a number of empirical observations that motivate the relinquishment of perception and action as distinct and peripheral modules. I briefly consider two (see Hurley, 1998, Ch. 14 for a more extensive review of these findings). *Firstly*, there is neurophysiological evidence—coming for instance from the study of mirror neurons—that the same neuron, or group of neurons, can be activated both in performing *and* viewing the same action, and thus play roles in both perception *and* action (e.g. Di Pellegrino et al., 1992). *Secondly*, there is a growing conviction that feedback or circular causal flows are essential to cognition. Varela et al. (1991), for instance, state that empirically, if an area *A* of the nervous system has connections to area *B*, then area *B* also has reciprocal connections to area *A*. This implies that the nervous system, rather than passively waiting for stimuli from the environment, can self-engage activity that is not constrained (entirely) by sensorimotor processes. This implies that rather than being two disjoint modules that sandwich central cognitive processes, perception and action are actually ‘so intimately bound together that they may be considered products of one cerebral function’ (Trevarthen, 1968, 391).

If the modularity of perception and action is denied, so is the passing on of representations between these modules. The alternative to the Classical Sandwich model is the sensorimotor loop (e.g. Stewart, 2010). In conceiving of living systems as having these sensorimotor loops, we acknowledge that sensory inputs are used to guide actions whilst at the same time actions modify the environment and/or the relation of the organism to its environment, thereby in turn modifying the sensory input. The loop is closed and thereby sets up a dynamic system (see figure 1).

According to Enactivism, we have to look at the history of an organism to understand how these sensorimotor loops were forged. This is explained by the concept of *structural coupling*, which is more than a mere exchange of matter and energy between two systems. If two systems share a history of interaction, their structures become congruent (Maturana & Varela, 1987): the organism is shaped by *and* shapes its environment (cf. Maturana,

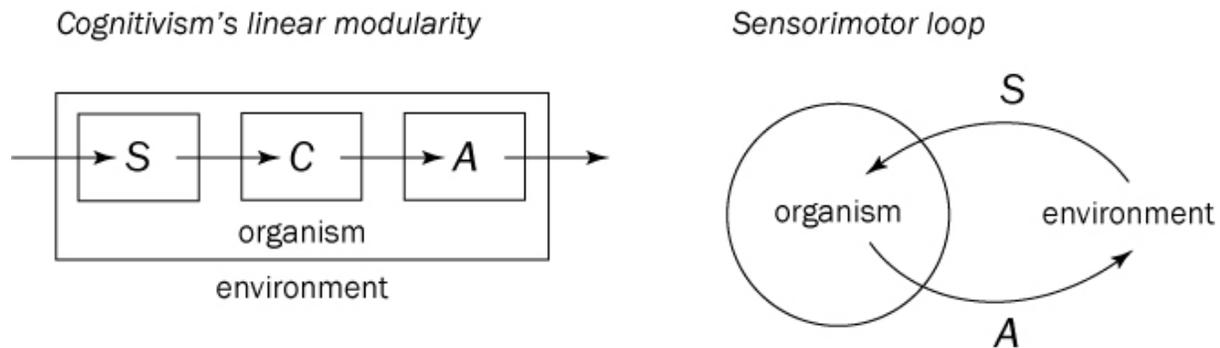


Figure 1: Comparison of the Cognitivist linear and modular view of cognitive processes, according to which sensory (*S*), cognitive (*C*) and action (*A*) processes are sequential and distinct, and the sensorimotor loop. The sensorimotor loop is based on Stewart (2010, Figure 1.1, 3).

1975). In other words, the cognitive structures of an organism depend constitutively on a history of dynamic interactions with its environment in which they were modified. Cognition is thus thought *to emerge* from the dynamic interactions between an organism and its environment. Accordingly, the Enactive approach consists of two points: ‘(1) perception consists in perceptually guided action and (2) cognitive structures emerge from the recurrent sensorimotor patterns that enable action to be perceptually guided’ (Varela et al., 1991, 173). Cognition then becomes enaction: ‘A history of structural coupling that brings forth a world’ (Varela et al., 1991, 206)

1.3 Radical Enactivism and Other Enactivisms

From this basic definition of what cognition is, different kinds of Enactivism have been developed. Two main strands can be identified (Torrance, 2006; Thompson, 2005). Firstly, *autopoietic Enactivism*, which is a broad approach to the nature of the mind, as expounded most prominently in the work of Thompson (2007); and secondly, *sensorimotor Enactivism*, which is focussed on the nature of perception (e.g. Hurley, 1998; Noë, 2004, 2009). Although Radical Enactivism (RadEn) is closely allied to these approaches,¹ there are two main points which distinguish it from the two aforementioned strands.

Firstly, RadEn aims to bring Enactivism ‘into direct conversation with accounts of basic cognition of the sort favoured by analytical philosophers of mind and more traditional cognitive scientists’ (Hutto, 2011, 45). This is achieved through a biosemiotic account of basic cognition, which is basically an Enactivist makeover of teleosemantics. I explicate this account in Section 2.

Secondly, RadEn holds that the other strands of Enactivism are not radical enough in their denial of representationalism (see Hutto (2005) and Hutto and Myin (2013, Ch. 2) for this diagnosis). Only when these, often implicit, commitments to traditional ways of thinking about cognition are abandoned, can Enactivism unlock its full potential.

¹See Hutto (2005) and Hutto and Myin (2013, 23–31) for a comparison of RadEn and sensorimotor Enactivism; and Hutto (2011) and Hutto and Myin (2013, 32–25, 76–78) for a comparison of RadEn with autopoietic Enactivism.

Accordingly, RadEn ‘not only retains the enactivist’s hostility to cognitivism, it defines itself in such terms’ (Menary, 2006, 11). What makes RadEn radical is its insistence on the embodiment of basic cognition (Hutto, 2013). This is apparent from the *Embodiment Thesis*, which states that ‘basic cognition [is equated with] concrete spatio-temporally extended patterns of dynamic interaction between organisms and their environments’ (Hutto & Myin, 2013, 5). This equation entails that the only explanation that can be given for the emergence² of cognition is the history of an organism, broadly construed, that is, including the organism’s evolutionary and cultural history. Hutto and Myin coin this the *Developmental-Explanatory Thesis*. In Section 2.1, I show how this leads RadEn to an explanation of basic cognition in terms of direct, that is, not representationally mediated, responsiveness to worldly offerings.

Moreover, as this interaction is to be understood as sensorimotor loops which constitutively include both the environment and the organism, there is no principled distinction between inner organismic responses and outer, environmental contributions. This leads to a radical interpretation of the extension of cognition: minds are not merely extended, but extensive. Rather than there first being a mind which can then extend itself into the environment, minds always already depend constitutively on the environments³.

In the rest of this Section, I look into the reasons for Radical Enactivism’s root and branch denial of the representational nature of of basic cognition.

1.4 A Very Short History of Radical Enactivism

Radical Enactivism has been expounded by Daniel Hutto in several books and a few dozen articles. The roots of Radical Enactivism are found in his diptych *Presence of Mind* (1999) and *Beyond Physicalism* (2000), where he first develops the distinction between contentless basic cognition and linguistic cognition that is contentful. In these two books he does not describe his approach as Enactivist. The term ‘Radical Enactivism’ is first coined in his 2005 paper ‘*Knowing What? Radical versus conservative Enactivism*’. In this paper, Hutto claims that other strands of Enactivism remain under the spell of old habits of thinking, and are thus too conservative. In the edited volume *Radical Enactivism* (2006), Hutto proposes an Enactivist approach to emotion and perception. In the replies to his critics in the same volume, the distinction between basic minds and linguistic minds is further developed. In his *Folk Psychological Narratives* (2008), the ideas developed in *Radical Enactivism* are further systematised and used to propose a narrative account of our folk psychological abilities. The book *Radicalizing Enactivism* (2013) is co-authored by Hutto and Myin, and focusses exclusively on basic cognition.

²Hutto and Myin do not discuss the notion of emergence, although they do state that ‘Sentience and sapience *emerge* through repeated processes of organismic engagement with environmental offerings’ (Hutto & Myin, 2013, 8, emphasis mine). Hutto (2008, 2006e) also does not explain what he means by emergence. One might assume that RadEn relies on autopoietic Enactivism’s conception of emergence (e.g. Di Paolo et al., 2010). I do not go into the metaphysics of emergence in this thesis.

³Hutto and Myin (2013) make a possible exception for phenomenality, which they claim could have a narrow basis. As I will not be addressing questions regarding the ontology of phenomenal experience, I disregard this subtlety.

1.5 The Notion of Content

The main thesis of Hutto and Myin's (2013) book *Radicalizing Enactivism* is that basic cognition has to be understood without invoking representations, which is cashed out in the notion of content. A state is contentful if it has specified conditions of satisfaction, or truth-bearing properties (the prime example is that of a proposition). For most cognitive scientists and philosophers, this very idea is radically misguided. The existence of contentful representations is often seen as a necessary condition for intelligent behaviour. Only when these representations are true with regards to the actual states of affairs can cognitive processes, understood as computational processes operating on these representational symbols, lead to effective acting in the world. Cognitivism is thus based on classical philosophy of language (Love, 2004): the mind is seen as a filing cabinet containing abstract language-like symbols that correspond to states of affairs in the world.

This, however, leaves the adherents of this view with an important task: to explain where these language-like symbols and their content comes from. Moreover, this explanation has to impart a real explanatory contribution of the content to a larger cognitive economy. These contentful symbols must therefore 'have the function of *saying or indicating* that things stand *thus and so*' (Hutto & Myin, 2013, 62). In other words, if the representational gloss can easily be eliminated without invalidating the explanation of the phenomenon, the explanation in terms of representations does no real work and is thus superfluous.

If a cognitive scientist or a naturalistically inclined philosopher wants to use the notion of content, they have to explain where it comes from in a way that is naturalistically credible. The prime candidate for explaining the genesis of content in the contemporary debate is the concept of information. Organisms are thought to 'pick up', or 'extract', information from their environments, which they then encode in a form which allows the mind to engage in computational processes on these informational contents. But in order for this account to work, informational contents have to be literally picked up. According to Hutto and Myin, to think that by only appealing to information one can explain what content is, is to be fundamentally misguided.

1.6 The Problem of Naturalising Content

Hutto and Myin (2013, 66–67) start from the assumption that the only scientifically credible form of information depends on the notion of *covariance*. This is the notion of information that philosophers rely on in order to explain informational content. When a perceived state of affairs, or signal, correlates to a distal state of affairs, it is thought that the signal carries information about the state of affairs. We can think for instance of tree rings: normally, they correlate with the age of the tree in a lawful way, and therefore, if we know the number of tree rings, we can come to know the age of the tree. The following general definition of information-as-covariance can then be given: 's's being F carries information about t's being H iff the occurrence of these states of affairs covary lawfully, or reliably enough' (Hutto & Myin, 2013, 66). The final qualification entails that Hutto and Myin want to target all theories that explain content in terms of covariance

relations, regardless of whether they conceive of these covariance in strong terms, e.g. lawful dependence, or in weaker terms, as indicated by ‘reliable enough’.

The question is now whether this covariance relation is the right kind of relation to constitute content, that is, instantiate truth-bearing properties. Or, put another way, does the obtaining of a covariance relation entail that one state of affairs ‘indicates’ or ‘says’ anything about the state of affairs with which it correlates. An important observation here is that covariance relations are symmetric, whereas indicating relations are not (Hutto, 2008, 48). Think for instance of the tree rings. Although they covary with the age of the tree, they do not ‘say’ anything about the age of the tree. If, for some reason or another, we find a tree in which the tree rings did not covary with its age, we would not say that the tree rings were false. They would simply fail to correlate with the tree’s age in this case.

The same line of reasoning holds if one of the correlating states of affairs is within an organism (Hutto & Myin, 2013, 67). If we were to rely solely on covariance to constitute content, the organism’s mental states would not only represent the world, the world would also represent the organism’s mental states. Even if no organism is involved, content would still be constituted by covariance relations, which entails that ‘the world would be everywhere silently referring or marking truths about itself (a strange metaphysics indeed).’ (Hutto, 2008, 49).

If this line of reasoning is sound, this entails that relying solely on the concept of information to explain the contentfulness of mental states is not enough. This is what Hutto & Myin (2013) coin the *Covariance doesn’t Constitute Content principle*. If informational content is not to be found in the world, it cannot literally be picked up. Hutto & Myin (2013, 73) thus adhere to the *No Acquired Content Principle*. To bring content into an explanation of cognition is an additional theoretical move that requires an justification (Hutto, 2008, 49).

1.7 The Hard Problem of Content

The Hard Problem of Content is to find a naturalistic explanation of what content is. The standard ploy of relying on information is unsuccessful: it can only account for information-as-covariance, but not for information-as-content. The Hard Problem of Content poses those who want to explain content in terms of information for a dilemma. Either one gives up on the notion of content, or one accepts that ‘contentful properties exist, even if they don’t reduce to, or cannot be wholly explained in terms of, covariance relations’ (Hutto & Myin, 2013, 68). Taking the latter option entails that contentful properties are like qualia: they may exist alongside physical properties but cannot be reduced to them. However, this radically alters the naturalistic project: it becomes the discovering of bridging laws between physical objects and content. Alternatively, of course, one might show the falsity of the Covariance Doesn’t Constitute Content Principle, or offer an alternative naturalistic explanation of content.

In this thesis, I will not further defend the Hard Problem of Content against the multitude of proposals about what content might be, and how these proposals might solve or evade The Hard Problem of Content, as Hutto and Myin (2013) spend most of their book on. I will rather accept the contentlessness of basic cognition at face value,

that is, I accept that basic minds do not represent conditions that the world may be in. In the next Section, I look at how RadEn circumvents the Hard Problem of Content by giving a description of basic minds without content, whilst at the same time denying that linguistic thought is contentless.

2 Basic Minds and Superminds

Radical Enactivism makes a strict distinction between basic cognition⁴, or ‘basic visceral responding’ (Menary, 2010, 7), which is contentless, and linguistically mediated thought that is contentful and therefore representational. In this section, I expound how RadEn describes both basic cognition (§2.1) and linguistic cognition (§2.2).

2.1 Basic Minds: From Teleosemantics to Biosemiotics

In the previous section, we have seen that content cannot be acquired. This entails that something more is needed before we can legitimately speak of contentful cognition. Teleosemantics aims to provide this extra ingredient by introducing information-exploiting organisms with proper biological functions. The basic idea is that content is fixed by the needs of the organism: ‘a device will have the teleofunction of representing Xs if it is used, interpreted, or consumed by the system because it has the proper function of representing Xs’ (Hutto & Myin, 2013, 76). The notion of proper function is based on the evolutionary history of an organism. Rather than asking what the organism is *disposed* to do, that is, give a purely mechanical explanation, we ask what it is *supposed* to do in *historically normal conditions*, that is, in those conditions in which the function evolved.

One of the favourite examples in the literature is a frog that snaps at black moving objects (e.g. Hutto & Myin, 2013, 117). Although it is disposed to snap at all black moving objects (e.g. shadows, experimenters waving black pieces of paper, etc.), it is supposed to snap at flies, for it is the subsequent ingestion of flies that explains its evolutionary success. The evolutionary benefits of eating flies thus explains the normativity inherent in the use of *supposed*: the frog is *meant* to be intentionally directed at flies (Hutto, 2006a, 141), as that is the reason that explains its having this particular function. In this way, the frog can be intentionally directed towards the black objects it snaps at, based on its evolutionary history. But can these teleofunctions also explain intensional, that is, contentful directedness?

2.1.1 Why Teleosemantics cannot Account for Content

What would constitute content on such a teleosemantic account? Hutto and Myin point out that two-tiered teleofunctional accounts (e.g. Dretske, 1988) have trouble in specifying what the contributions are of, on the one hand, the organism’s interpretative responses, and, on the other hand, the worldly offerings it is directed at, in determining the content. They therefore combine Millikan’s (2005, 100) acknowledgement that ‘the content is determined, in a very important part, by the systems that interpret it’, with the No Acquired Content principle. The resulting *Strengthened Millikan Maneuver* (Hutto & Myin, 2013, 75) shifts the explanation in terms of content-*consuming* mechanisms to one in terms of content-*creating* mechanisms. In other words, if there is no content in the absence of responding organisms, then the organism has to be solely responsible for the constitution of content.

⁴Hutto and Myin (2013, x) take *basic cognition* to be narrower than *basic mentality*. Basic cognition does not imply phenomenality, whereas basic mentality may exhibit it.

But according to Hutto and Myin, these functions of creatures are also unable to constitute content. The root problem is that biology, by itself, is unable to account for intensional directedness. In order for an organism to be intensionally directed at a worldly offering, it must apprehend the worldly offering under a *Mode of Presentation* (MoP) (Hutto, 2008, 43), that is, it must represent it under a certain guise or description (Hutto & Myin, 2013, 79). The selectionist explanations given by teleosemanticists, however, are extensional. The reason for this is that the processes of natural selection that shaped organisms’ teleofunctions do not care about truth. When appealing to evolution, the proliferation of biological functions can be explained only in terms of their contribution to reproductive success. Creatures have evolved to have particular functions that contribute to their survival in historically normal conditions, not to represent the objects of these functions in truth-evaluable ways. The success or failure of biological functions is therefore unable to constitute content (cf. Hutto, 2008, 66–71).

This does not negate the fact that an organism can be *extensionally misaligned* in conditions that are not historically normal—when a frog snaps at a black piece of paper waved around by an experimenter, for instance. Extensional misalignment, however, is not the same thing as being *wrong*, and it therefore, by itself, gives us no reason to describe a creature as having a mental state that is false.

2.1.2 Biosemiotics: Teleosemantics Without the Semantics

The fact that the selectionist explanations are unable to provide an intensional description of organisms’ cognitive antics, does not entail that they fail to describe an organism’s intentional directedness. RadEn therefore talks of *biosemiotics*⁵ instead of *teleosemantics*. This intentional directedness cannot be ascribed to subparts of the organism in question. Take the example of the frog once more: performing the function of snapping at flies is a whole-bodied action, that relies constitutively on, *inter alia*, the frog’s tongue, brain, and eyes, but also on the fly being present. Hutto (2008, 57) therefore states that the ‘intentionality they exhibit is an attitude of the whole organism expressed in their behaviour; it is neither a property of the signs themselves nor of organismic inner states’. Organisms are therefore informationally sensitive to certain natural signs which *directly* guide their behaviour, without the need to posit mental representations.

Although the natural signs to which organisms are responsive in this way have to correlate to the distal state of affairs for which this behaviour has been selected in historically normal conditions, the organism itself need not register, or know, *that* these correlations hold. A direct responsiveness to these signs themselves is enough to explain the evolutionary gain, which makes a description in terms of content superfluous.

According to Hutto (2008), this biosemiotic account of basic cognition can be understood as scripted patterns of action and reaction. He coins these *Action Coordination Routines* (ActCRs), which can be hard-wired⁶ as well as learned. They involve the following parts (Hutto, 2008, 51):

⁵Hutto and Myin (2013) use the word teleosemiotics instead of biosemiotics (Hutto, 2006e, 2008). Both terms denote the same Enactivist makeover of teleosemantics.

⁶Di Paolo, Rohde, & De Jaegher (2010) similarly express worries about the reduction of normativity to purely evolutionary mechanisms. They state that the idea of ‘phylogenetic invariant yardsticks’ (Di Paolo et al., 2010, 45) is at odds with the Enactivist idea of describing the emergence of cognitive structures in

- Step 1:** Paradigmatic recognitional element involved in X
- Step 2:** Paradigmatic outward expression of X
- Step 3:** Paradigmatic bodily changes and feeling of those changes
- Step 4:** Paradigmatic motivational response involved in X
- Step 5:** Paradigmatic action out of X

We cannot individuate these non-verbal intentional attitudes based on their contents, but solely on what they are directed at, that is, what class of objects, extensionally speaking, triggers the appropriate ActCR. They are ‘a kind of embedded, embodied, and enactive *know-how*, not a form of thinking *that*’ (Hutto, 2008, 65).

Although intentional directedness is thus conceived as being directed at an object without a Mode of Presentation, this does not entail that organisms are directed at objects ‘as such’. Although objects might appear a certain way to organisms when they are directed at them, ‘talk of things “appearing a certain way” needs to be unpacked in terms of actions and possible actions’ (Hutto, 2006a, 140). Hutto thus denies that ‘appearing a certain way’ has anything to do with a description or MoP, as the former is based in responding, whereas the latter is a conceptual ability.

To quickly recap RadEn’s view on basic cognition: an organism’s indexical responsiveness to natural signs allow it to be directly guided by them. The selectionist explanations of teleofunctions can only account for *intentional* directedness, that is, not for *intensional* directedness. Teleosemantics thus offers no solution to the Hard Problem of Content. Fortunately, no such solution is required to explain basic cognition. As non-verbal organisms do not rely on mental states that ‘say’ or ‘indicate’ anything about the world, their interpretative responsiveness to the world cannot be true or false. They can merely be extensionally misaligned through a failure of biological proper functions. Before we direct our attention to the role of linguistically mediated thought in RadEn, I first briefly introduce RadEn’s account of non-linguistic cognition that is directed at objects that are not here-and-now.

2.1.3 Going Off-line: Imaginative Recreations

The account of basic cognition as described so far is severely limited. The object at which a creature is directed has to be present in order for a creature to be directed at it. This kind of cognition is called *on-line cognition*. To be directed at objects that are absent, as for instance in the case of remembering an object, is called *off-line cognition*. The indexical nature of the ActCRs entails that ActCRs can only be triggered by worldly offerings in the immediate here-and-now environment. But surely there are non-verbal creatures

terms of an organism’s history. This critique seems to hold even if the selectionist explanation is merely used to ascribe intentional attitudes to organisms. Hutto (2006c, 161) talks for instance of ActCRs as being ‘*pre-scripted*’ and ‘well worn hand-me-downs’. These ActCRs are therefore described as invariants being passed on from one generation to the next. But this sits uncomfortably with Hutto and Myin’s Developmental-Explanatory Thesis, according to which cognition can be explained by ‘nothing more, or other, than the history of *an organism’s* previous interactions’ (Hutto & Myin, 2013, 8, emphasis mine). The selectionist pressures that both shape and explain the proliferation of ActCRs surely fall outside the scope of a *particular* organism’s history of interaction. It seems that RadEn thus still has to explain the interaction between phylogenetic and ontogenetic time-scales, that is, how are the passed-on ActCRs re-enacted in ontogenetic development.

that behave in ways that are more complex, and thus cannot be explained by a direct responsiveness to their current environments. Accordingly, Hutto (2008, 80) proposes an account of Recreative Imaginings (ReImS) to account for basic off-line cognition.

These ReImS are perceptual *re-enactments*, that is, the perceptual mechanisms operating in off-line mode. Hutto (2006c, 167) states that ‘non-standard triggers might result in non-standard responses; images rather than actions’. We can think of these non-standard triggers as associative memories for instance. These ‘images’, however, are not representations. If anything, they are ‘presentations’ (Hutto, 2006c, 168), for they can be objects of attention.

It has to be noted that Hutto’s use of the word ‘image’ can be misleading here. As we have seen, things appearing a certain way in on-line perception has to be unpacked in terms of (possible) actions, which is congruent with thinking of perception and action in terms of sensorimotor interaction with an environment. It is unclear, however, whether Hutto takes these ‘images’ as understandable in terms of actions and possible actions, or that he takes an act of perceptual re-enactment to constitute a novel type of cognition. The metaphor of the ‘image’ seems to indicate a presentation that is detached from possible action. This is contradictory with the Enactivist insight that perception and action cannot be separated. Hutto (2006c, 163) explicitly endorses this view: ‘In all, like Hurley, I take it to be a mistake to think of basic *perceptual acts* in traditional input and output terms.’ If we thus relinquish the classical sandwich view of the mind—the modularity of perception, cognition, and action—we can no longer make sense of inputs or outputs in online-cognition. And if there are no perceptual inputs in cases of online perception, this makes the discussion in terms of non-standard inputs and outputs objectionable in the case of perceptual re-enactments. In other words, *if* we think of perceptual enactments in sensorimotor terms, *then* perceptual *re-enactments* should have the same sensorimotor nature. As Hutto explicitly acknowledges that things appearing a certain way is based, not only on occurrent actions, but also on possible actions, there seems to be no problem in taking the appearance of perceptual re-enactments in terms of possible, or suppressed, actions. For the rest of this thesis I therefore interpret ReImS to be sensorimotor in nature in this sense, and accordingly refer to them as *sensorimotor re-enactments*.

According to Hutto, although these images can ‘resemble’ worldly offerings (or, as I would say based on the considerations in the previous paragraph: recreatively imagining an object resembles actually perceiving an object), they fail to represent them. This is because the resemblance relation is not a referential one. Hutto (2008, 82) gives the example of identical twins: although they look alike, they do not stand in a referential relation to one another. ReImS further lack the appropriate compositional structure to account for logical reasoning. These off-line modes of thought can thus only sponsor ‘associative and analogical modes of thought’ (Hutto, 2006c, 168), or ‘nonlogical instrumental thinking’ (Hutto, 2008, 83), which Hutto calls protological reasoning. A defining feature is that this kind of reasoning is always tied to a particular domain, as opposed to the kinds of abstract reasoning afforded by the manipulation of symbols. Moreover, a capacity for this kind of off-line sensorimotor re-enactments does not entail that these creatures can be said to harbour beliefs about the world. They might picture the world in a certain way, that is, recreatively imagining it, but ‘it makes no sense to talk of believing an image,

only believing propositions will do.’ (Hutto, 2008, 82). Although ReImS can thus explain more complex behaviour than ActCRs, they do not allow creatures to be intensionally directed at their environment.

Moreover, ReImS can explain a creature’s ability to reidentify particulars (cf. Hutto, 2008, 72–3), without resorting to conceptual identification: ‘Recognising berries in the sense of being able to re-identify them is one thing, recognising them as berries—i.e. conceptually identifying them and categorising them “as such” is quite another; it is something that must wait on the development of linguistic abilities.’ (Hutto, 2006c, 170). I return to these re-identifying abilities in the discussion in Section 3.1.

2.2 Linguistic Superminds

RadEn does not want to throw out content as an explanatory device across the board. Linguistically mediated thought is obviously contentful according to Hutto and Myin (2013). Our initiation into the realm of language constitutes a novel way of responding that is wholly unlike the basic modes of interaction, even of the off-line variety:

off-line iconic thinking [based on recreative imaginings] *is still embodied* and visceral because of its re-enactive character, it is therefore quite unlike the kind of reflective or neutral responding afforded by the manipulation of *abstract amodal* symbols. . . . [this is] a consequence of the capacity to represent objects and situations using complex linguistic signs (Hutto, 2008, 98, emphases mine).

Using language, we can be directed at propositions (Hutto, 2008, 87), that is, certain states of affairs under a Mode of Presentation (MoP), instead of just at worldly offerings.

The main motivation for securing the special status of linguistically mediated propositional attitudes seems to be Hutto’s project of giving a narrative account of our folk psychological abilities. In this thesis, I will not elaborate on the link between RadEn and Hutto’s (2008) so-called *Narrative Practice Hypothesis*. But it is important to realise that for this project, Hutto relies on a belief-desire psychology, which allows us to act according to *reasons*, rather than *causes*. For this belief-desire psychology, it is crucial that both beliefs and desires interrelate logically. As ReImS are not composed of recurring meaningful subparts, they do not allow for this logical interrelation between desires and beliefs. This can only be the case if a compositional semantics is in place. Non-verbal creatures are thus not *true believers* according to Hutto⁷. Although we can describe their behaviour in terms of beliefs and desires, by adopting Dennett’s (1987) *intentional stance*, this does not entail that they actually have the mental states we attribute to them. For deciding whether a creature can be counted as a true believer, a more stringent test is needed: the intensionality test.

2.2.1 Davidsonian Interpretationalism

To ascribe propositional attitudes, Hutto (1999, 111) relies on Davidsonian interpretationalism. If we want to know whether a creature is capable of harbouring propositional

⁷In a recent article, Hutto (2013) *does* talk about contentless beliefs. Whether or not we call intentional attitudes beliefs seems to be a mere stipulation though.

attitudes, we have to administer an *intensionality test*, which is essentially *radical interpretation*. This test is based on the observation that—contrary to intentional attitudes—propositional attitudes are opaque. They can be individuated not only based on their broad contents, that is, their truth conditions, but also on their narrow contents, that is, their MoPs. This entails that, in the case of linguistically mediated propositional attitudes, ‘coreferential substitution . . . is not automatically permitted’ (Hutto, 2008, 46). For instance, knowing that ‘next Tuesday is the first day of April’ does not entail knowing that ‘next Tuesday is April Fools’ Day’, though the truth-conditions of these two propositions are the same. Only when these MoPs are in play, can we talk about thought in a content-involving way (Hutto, 2008, 43), and it is only through ‘structured vehicles of thought’ (Hutto, 2008, 87), as provided by natural language, that creatures can be directed at states of affairs under an MoP, that is, have propositional attitudes.

The intensionality test thus ‘essentially reveals in what way, and to what degree, a subject’s words or thoughts is [sic] resistant to co-referential substitution’ (Hutto, 1999, 114), by prompting assent or dissent from this subject. The fact that a person might assent to ‘next Tuesday is the first day of April’, but not to ‘next Tuesday is April Fool’s Day’ tells us *how*, that is, under what MoP, that person conceives of ‘next Tuesday’. Moreover, it allows us to chart the pattern in a person’s thought. The propositional content of a particular utterance is partially determined by the place it occupies in the ‘logical geography’ (Hutto, 1999, 114) or ‘intensional cartography’ (Hutto, 2008, 67) of that person, which is uncovered in the process of radical interpretation.

In this way, we can *prima facie* uncover what a person means by a particular utterance. However, radical interpretation by itself is not enough to accomplish this charting task, for it only allows us to specify which sentences a subject takes to be true, which is an extensional relation. The theory of truth arrived at through this process can thus only allow for uncovering ‘the structure, not the content, of propositional thought’ (Hutto, 1999, 129). To arrive at a theory of meaning, Hutto relies on two principles of charity as formulated by Davidson.

The first is *the principle of coherence*, according to which a radical interpreter has to assume that a speaker is normally correct in making assertions, and that these assertions are about features of the world. The motivation for this principle is that a subject’s use of words has to be consistent in order to attribute any meaning to them. If a person was to be thoroughly inconsistent in his linguistic responses, we would be unable to interpret his words using the intensionality test. Such a person would therefore fail to say anything at all. However, no speaker will turn out to be perfectly consistent in her verbal responses.

To deal with these inevitable inconsistencies, Hutto further relies on the *principle of correspondence*, that is, an interpreter has to take ‘the speaker to be responding to the same features of the world that he (the interpreter) would be responding to under similar circumstances’ (Davidson, 1991, 158). We take inconsistencies to be *mistakes* on the part of the speaker. Moreover, this principle allows to bootstrap our radical interpretative practice. Since the propositional content of such an attitude is determined in part by its place in the logical geography, the ascription of propositional content is holistic in nature on this interpretationalist understanding. In other words, it is impossible to ascribe a single propositional attitude to a subject (Hutto, 1999, 133). It makes no sense for instance to only ascribe the belief (which is a propositional attitude) that ‘the Earth is round’ to a

person, without ascribing a host of other beliefs—such as ‘the Earth is a planet’, ‘we live on the Earth’, etc.—that together define *what* is meant by ‘the Earth is round’.

However, using this principle of charity, we can start from the defeasible assumption that the subject we want to interpret responds to the world in similar ways as we do. We can therefore attribute content to an utterance based on what we ourselves would say in that particular situation. For example, the field-linguist would take *gavagai* to mean ‘rabbit’, and not ‘detached-rabbit-part’ or ‘time-slice-of-rabbit’ *because* that is how *she* would respond to the event of seeing a rabbit.

Apart from its position in the logical geography, content is therefore also determined by referential properties of meaningful subparts of sentences. The referential properties are also an important difference for Hutto between basic non-verbal responding and verbal responding. It is only in the latter case that signs are used to stand for something else, and thus serve as a representation. It is therefore the combination of a compositional semantics with recurring meaningful subparts that stand in logical relations to one another, and the denotational content (Hutto, 2008, 87) of these subparts, that makes a radical interpretation of a subject’s speech, and thus the ascription of propositional attitudes possible.

2.2.2 Thinking in Natural Language

As we have seen, propositional thought requires a creature to apprehend a state of affairs using structured vehicles of thought. Traditionally, two candidates can be identified that can account for this, either a Language of Thought (LoT) or natural languages. Hutto denies the possibility of a LoT because it entails basic cognition that is contentful (see Hutto, 2008, 88–95 for a criticism of LoT), and thus, following Davies, he promotes the *Thinking in Natural Language Hypothesis* (TNL). According to this hypothesis, the ‘sentences of natural language *are* the vehicles that enable propositional thought’ (Hutto, 2008, 95), that is, they are ‘the required medium for thinking and speaking truths (or falsehoods for that matter)’ (Hutto, 2008, 88)

The compositional semantics of natural language make propositionally based thought possible. Because sentences consist of recurring meaningful subparts, in contradistinction to the ‘images’ in sensorimotor re-enactments, radical interpretation is possible. Only because of this internal structure can we characterise language by a finite set of atomic sentences, which, together with compositional rules, allows for the production of an potentially infinite number of sentences by us, finite beings (Hutto, 1999, 127).

The TNL has several implications. Firstly, the linguistic symbols that enable propositional thought are not ‘stored in the mind’ (Hutto, 2008, 124), as follows from accepting a LoT. They are rather public items such as words. This has the added advantage that we can explain how we can be causally related to language. Unlike abstract entities, the sentences of natural language are, in Dennett’s (1998) terms, *manipulanda*. That is, sentences made up of symbols of natural language are concrete spatio-temporally extended particulars with which we can causally interact (Hutto, 2008, 88).

Secondly, the TNL entails that we first have to learn to think outside our heads before this process can be internalised. ‘Thinking in our heads’ is thus an achievement that

depends both logically and developmentally on our ability to use the public symbols of natural language to communicate with others.

If we make a quick head-count, the TNL leaves us with three entities which are involved when a subject has a propositional attitude: (1) the subject itself, (2) the concrete symbols of natural language, and (3) the contents expressed by these symbols. Accordingly, Hutto (2008, 87) sees a propositional attitude as a three-place relation: ‘thinkers stand in relation to sentences of some kind and this, in turn, is what allows them to adopt various psychological attitudes towards the complex contents that those sentences express.’ Hutto calls these complex contents by their traditional name, that is, propositions. Schematically, the relation is thus as follows: subject \rightarrow sentence \rightarrow proposition. The concrete symbols of natural language act as a medium between the subject and propositions. However, Hutto (2008, 87) does not want to pass judgement on the ontological status of these propositions. I will return to status of propositions in Section 3.3.

2.2.3 The Linguistic Supermind

According to Hutto, linguistically mediated thinking brings a whole new way of thinking into being. He borrows Frankish’ (2004) distinction between the *basic mind* and the linguistic *supermind*. This supermind is built atop, or supervenes on, the basic mind. As it is only made possible by an initiation into public language, we have to conceive of the supermind as a ‘kind of softwired virtual machine’ (Hutto, 2008, 97), not unlike Dennett’s (1991, 214ff.) virtual Joycean machine. However, a more dynamic role is envisaged for the supermind. Adopting a linguistically mediated attitude towards a proposition entails an active conscious commitment to take this proposition to be true, in the case of belief. Hutto, again following Frankish, takes the supermind to be a premising machine. The supermind not only allows us ‘to speak as if the proposition were true but to reason as if it were—to take it as a premise’ (Frankish, 2004, 91).

A combination of this idea of ‘premising’ with the compositional semantics of natural language is what makes possible ‘deliberative acts of explicit practical or theoretical reasoning—those of the classical deductive, inductive, and abductive variety’ (Hutto, 2008, 96). Hutto thus views linguistically capable creatures to be ‘truth-seeking’ and ‘means-end reasoning’. Although this reasoning is based on basic protological reasoning using ReImS, it is the complex ‘abstract amodal symbols’ that ‘represent objects and situations’ (Hutto, 2008, 98) that make domain-general reasoning possible. For Hutto, the two main functions of the supermind, and therefore of language itself, are thus *first* to represent the world in linguistic terms, and *second* to use these representations as premises in acts of reasoning.

Moreover, the supermind is conceived of as sponsoring ‘central’ cognitive process: ActCRs ‘operate quite independently of *central cognition*: in predating language-mediated forms of thinking and believing, they would have to.’ (Hutto, 2008, 63, emphasis mine). As the ActCRs do not rely on content, they are ‘*naturally* cognitively impenetrable’ (Hutto, 2008, 63). Supermental thinking, on the other hand, is ‘consciously accessible, indeed transparent to us’ (Hutto, 2008, 97). ‘Hard-wired’ ActCRs therefore operate ‘outside’ the direct control of the soft-wired premising machine. The advent of linguistically mediated thought does not entail a complete cognitive refit, as our basic modes of responding remain largely untouched by our superminds. Hutto views the supermind as some sort of abstract

problem-solver that only kicks into gear when a problem is encountered that the basic mind cannot handle:

In the human case, [basic minds] often allow us to navigate by autopilot as it were, but sometimes they are thwarted. Tackling problems that the basic mind cannot handle is a job for the supermind. However, contrary to what might be expected, going over to “manual control” is to switch into low-gear thinking. It is a shift to a slow, careful deliberative mode that is serial, sequential and fragile. (Hutto, 2008, 99)

However, superminds are not wholly powerless with regards to basic modes of thinking. In some cases they can control, override, or overcome the cognitively impenetrable ActCRs. We can, for instance, use propositional knowledge of our basic modes of responding to consciously try to overcome them (Hutto, 2008, 63, 123).

It is important to notice that this definition of linguistically mediated thought in terms of the manipulation of *abstract amodal symbols* entails that the supermind is no longer embodied according to Hutto. If we conceive of the linguistic mind as a virtual machine, we can abstract away from the embodied processes from which this kind of thinking emerged. *Prima facie*, the idea of the supermind thus conflicts with the Embodiment Thesis, one of the basic assumptions of RadEn (Hutto & Myin, 2013, 5). In the rest of this thesis I therefore take a closer look at our linguistic abilities to see whether this *disembodied* view of them can be upheld.

In the following two Sections, I critically examine the distinction RadEn draws between basic contentless mind and contentful linguistically mediated superminds, and argue that it is implausible. In particular, I focus on the following two assumptions that underlie the distinction between the two types of mind:

The Necessity of Linguistic Mediation. *Only through linguistic mediation, that is, using the structured vehicles of natural language as a medium of thought, can a creature adopt propositional attitudes.*

Language Consists of Abstract Amodal Symbols. *Linguistic cognition takes place in a softwired virtual machine that operates on abstract amodal symbols, which, because of their place in a logical geography combined with referential properties, allow for the fixation of their meaning through the process of radical interpretation.*

In Section 3, I look at the first of these assumptions, that is, the role of the metaphor of linguistic mediation in relation to propositional attitudes. In Section 4, I look into the alleged properties of these linguistic symbols.

3 Linguistic mediation, propositional attitudes, and propositions

The three-place relation Hutto envisages a propositional attitude to be (subject \rightarrow natural language sentence \rightarrow proposition), entails that linguistic mediation is a necessary condition for the adoption of propositional attitudes. It is unclear, however, exactly how Hutto envisages this mediating role of language. One would suppose that a creature that has a propositional attitude is directed at a proposition that is expressed by the symbols of natural language. Hutto, however, sometimes seems to take the symbols of natural language themselves as the intentional object: ‘Mastery of language reveals the existence of, what we might for convenience call, *linguistic objects of attention* and co-attention.’ (Hutto, 2006d, 200, emphasis mine).

This conception of propositional attitudes raises three questions:

1. Is linguistic mediation necessary for adopting a propositional attitude *if* we take the intensionality test as a criterion?
2. In what way, if any, are linguistic symbols objects of attention in adopting propositional attitudes?
3. What is a proposition if we take the structured vehicles of intensional thought to be the symbols of natural language?

In what follows I focus on these questions in turn. I argue that linguistic mediation is not necessary for adopting a propositional attitude *if* we accept the intensionality test as a criterion (§3.1). Thereafter, I argue that even if we take the metaphor of language as a medium seriously, we have to accept that, in most linguistic communication, it is not the medium at which we are directed (§3.2). Finally, I argue that, using the resources available to Radical Enactivism it will be impossible to make sense of the notion of a proposition (§3.3).

3.1 The Necessity of Linguistic Mediation for Propositional Attitudes

In looking into the necessity of linguistic mediation for adopting propositional attitudes, it is crucial to see that there are two ways of interpreting what it is to have a propositional attitude: (1) ‘taking a sentence to be true’ (call this a *sentential attitude*); and (2) ‘taking things to be the case’ (Ratcliffe, 2009). *Prima facie*, it seems perfectly reasonable to state that we can ‘take things to be the case’ *without* ‘taking a sentence to be true’. Ratcliffe (2009, 389–390, first emphasis mine) gives the following example:

When I see my young son crawling towards me, *my belief that* ‘Thomas is in front of me’ does not first of all take the form of recognising that the sentence ‘Thomas is in front of me’ is true. Rather, I take it *to be the case* that Thomas is in front of me. Regardless of whether and to what extent the formation of this conviction requires linguistic ability, the object of the conviction is not first and foremost a sentence.

Here Ratcliffe describes himself as having a non-linguistically mediated conviction, or belief, with the content ‘Thomas is in front of me’. It seems that if we are to rely on

the intensionality test this pre-verbal conviction has content. A radical interpreter could for instance ask Ratcliffe a number of questions, like, ‘is your son in front of you?’, ‘is a member of the species *homo sapiens sapiens* in front of you?’ etc. In this way, we could find out under what Mode of Presentation he conceives of the current states of affairs. It might be objected that this process of radical interpretation brings language into play, thereby vindicating the idea that content *does* rely on linguistic mediation. It is certainly true that we could not administer the intensionality test without relying on linguistic means. But what we are trying to do in this case, is to ‘fix the content’ of Ratcliffe’s pre-verbal attitude, not of the words he uses to describe this. In other words, we are not positioning a particular linguistic unit in his logical geography, but rather, we are inquiring under what guise, or description, Ratcliffe conceived of the state of affairs *before* he made any explicit linguistic judgements

The following dilemma can now be formulated for Hutto’s position. This belief can either be a function of the basic mind—in which case the description of this conviction in linguistic terms is unwarranted because basic cognition is contentless—or it can be a function of the supermind. The latter option can be dismissed instantly. As supermental thinking is cognitively transparent according to Hutto, we would know if we were to make a linguistic judgement. The former option, however, is also unappealing. It entails that if Ratcliffe were to utter the sentence ‘Thomas is in front of me’, this linguistic act would change his cognitive take on the current state of affairs. The supermind would take over from the basic mind, thereby transforming his contentless intentional attitude into a contentful intensional attitude. But this conclusion is phenomenologically implausible: uttering the sentence ‘Thomas is in front of me’ is not an additional judgement.

This does not entail that making explicit judgements never changes a subject’s cognitive take on a situation. However, the fact that we sometimes reason about a state of affairs to arrive at a novel way of takings things to be the case does not imply that we always have to change our cognitive take of a states of affairs when we say something about it. Moreover, I need not claim that adopting such a non-linguistically mediated attitude towards a state of affairs is thoroughly detached from linguistic abilities. For one can still hold that the ability to adopt these attitudes that can pass the intentionality test is developmentally dependent on linguistic abilities, without making the stronger claim that each attitude is dependent on linguistic mediation.

The alternative I shall be proposing in the next Section is that the distinction between basic responding and linguistic responding cannot be drawn as Hutto wants to. The ascription of content by a radical interpreter does not illuminate how linguistic cognition works. Before I sketch this alternative, I first further investigate Hutto’s idea of linguistic mediation as a necessary condition for adopting propositional attitudes.

Let us take a look at how Hutto deals with examples like these. He follows Millikan in claiming that re-identification abilities do not rely on identify judgements, that is, on classifying or categorising objects (Hutto, 2006b, 84). I ignore the technicalities of how Hutto takes such an account of re-identification to function, for I need not, and do not wish to challenge it: we can accept that *basic* re-identification abilities—that is, those of non-verbal creatures—do not rely on conceptual judgements, whilst still allowing that these re-identification abilities *can* be conceptual in the case of linguistically competent

human beings. To understand the distinction, let us look at the example Hutto (2006b, 84–5) gives here, and what conclusions he draws from this:

Or consider the way in which I re-identify the milkman, again fallibly, by hearing the distinctive sounds he makes when opening the gate and clinking the milk bottles at a particular *time* each morning. . . . My point is that my ability to re-identify substances in this way is the *basis for*, not the product *of*, judgements of the sort “There’s the milkman”, or “It’s the milkman again” (or ‘More milkman!’). In my own case and that of other linguistically competent humans, *but not that of non-verbals*, I can use concepts to make and express judgements about the very same situation. . . . We can form judgements, but our basic perceptual responses, like those of other animals, do not involve them—I do not form content-involving ‘thoughts’ about the milkman in order to re-identify him by sensory means.

The following distinction between the re-identification abilities of non-verbals and our re-identification abilities can thus be made: a non-verbal creature, on the account of Hutto, can only re-identify a particular, whilst we can use these re-identification abilities to form conceptual judgements⁸. The question now is whether we can, in our case, make this distinction. It seems at least very plausible that when I re-identify the milkman *as such*, I do not *first* re-identify him in an indexical contentless way, *before* judging that what I re-identified is the milkman. However, Hutto’s strict distinction between the two kinds of cognition seems to entail just this implausible two-step process.

We could imagine for instance that, upon hearing the characteristic sounds of the milkman, I go outside to pay my milk bill, only to discover that the postman was tripping over the milk bottles. My consequent utterance ‘I thought you were the milkman’ does not entail that I previously said to myself ‘There’s the milkman’, that is, that I explicitly formed a linguistic judgement. One way out would be to say that my belief was implicit. But this is incompatible with how Hutto conceives of the supermind. For we would then have to assume that the supermind was forming linguistic judgements outside of our awareness, which conflicts with the transparency Hutto imparts on the supermind.

My taking it to be the case that the milkman was at my door thus need not be linguistically mediated. Moreover, this example shows that my non-linguistically mediated conviction that the milkman was at my doorstep is truth-evaluable: I was wrong in taking it to be the case that the milkman was at my door. Even if we accept that re-identification abilities are non-conceptual in their basic mode, it thus does not follow that they are always non-conceptual in the case of linguistically competent human beings. The basic idea is this: *once* the relevant conceptual abilities are in place, some re-identification abilities are altered, changing from non-conceptual to conceptual.

In §4.3.2 I propose a general approach to conceptual abilities that allows for these non-linguistically mediated attitudes. If we accept this possibility, the relation between the basic mind and the supermind is significantly altered. The linguistically grounded conceptual abilities ‘spill over’ into basic modes of interacting with the world. This idea is expressed by Rudd (2006, 70) in his reply to Hutto (2006e):

⁸Hutto (1999, Chap. 5) describes how our non-conceptual re-identification abilities, when combined with joint attention mechanisms, can explain how non-conceptual basic responding can lead to conceptual thinking. He relies there on the Davidsonian idea of social triangulation. I return to this account in Section 4.2.

So the contribution of conceptuality (language, culture etc.) to the animal level of sensibility is not like placing a layer of icing on a cake; it is more like adding a coloured dye to water, which then becomes permeated through and through by the added element.

Hutto resists this conclusion, because he fears that if we deny that, even in the case of linguistically capable humans, basic responding is non-conceptual, ‘we will be unable to make best sense of such phenomena as perceptual illusions, aspect-dawning, mental imagery and the fine-grainedness of experience.’ (Hutto, 2006b, 86). He gives the following argument based on our colour perception: we can visually discriminate between millions of colours, but we are unable to re-identify all these particular colours. Our discriminatory abilities thus far outweigh our re-identification abilities. However, so the argument goes, if perception *was* conceptual, we would think that this re-identification would be relatively straight-forward, as the discrimination abilities would be sponsored by concepts. Therefore, colour perception is not conceptual. So it seems that for Hutto there are only two options: either perception is thoroughly conceptual from the start, or we accept Hutto’s story on how conceptual abilities arise from non-conceptual abilities and adhere to the strict distinction between them (Hutto, 2006b, 86).

However, there seems to be a middle way, the one I’ve been arguing for. It is crucial to see that this line of reasoning does neither lead to the conclusion that all of our basic cognition is always contentful, that is, that concepts are necessary for perceiving, nor does it imply that we can put all of what it is like to experience the world into words. The acceptance of the fact that once we have conceptual abilities, we can wield these without explicitly using the word associated with this concept, that is, without linguistic mediation, does not imply that all of our perception is conceptual. For example, I can see *that* a particular thing is green. However, my seeing that something is green does not entail that I can re-identify that *particular shade of green* reliably. It only entails that I perceive of it as being green, and that, were a situation to arise in which this would be appropriate, I could linguistically express *that* I take it to be green.

To sum up: I take it that for linguistically competent creatures, it is possible to have an intentional attitude that *can* pass Hutto’s intensionality test *without* being linguistically mediated. This makes the strict distinction between basic contentless responding and contentful supercognition implausible. ‘Taking things to be the case’ does not reduce to ‘taking a sentence to be true’, and can be achieved without linguistic mediation. In what follows I argue that even linguistically mediated propositional attitudes need not be sentential attitudes.

3.2 Language as a Medium and as an Intentional Object

As we have seen, Hutto takes a creature that has a propositional attitude to be directed at a proposition expressed by natural language, which in turn represents a particular states of affairs. However, to further complicate things, Hutto sometimes seems to take the symbols of natural language as intentional objects:

Mastery of language reveals the existence of, what we might for convenience call, *linguistic objects of attention* and co-attention. Children powerfully augment and

extend their possibilities for engaging with others when they learn how linguistic signs can serve both to designate other objects and states of affairs *as well as* being the *focus of attention*' (Hutto, 2006d, 200, first and third emphasis mine).

The question now is whether the symbols of natural language are *objects of attention* in acts of using and understanding language. Remember that we distinguished between two interpretations of a propositional attitude, firstly, 'taking a sentence to be true', and secondly, 'taking things to be the case'. If we conceive of natural language as a medium, as Hutto does, the first interpretation is unwarranted. When we make a linguistic statement about a particular state of affairs, the intentional object of our attitude is mostly not the sentence we utter. When I for instance tell someone that 'The Erasmus University is in Rotterdam', the focus of my attention is not that sentence. Rather, I intend to say something *about* the Erasmus University *using* words.

If natural language is indeed a medium, this entails that in normal conditions it does not show up in our awareness as something we are directed at, that is, it is *phenomenologically transparent*. It is important to note that this is a different kind of transparency than Hutto attributes to supermental thinking. For him, transparency means that our linguistic thought is consciously accessible to us. The phenomenological transparency I am referring to here entails that, although our linguistic thought is consciously accessible, the focus of our attention is most of the time not directed at our utterances whilst we utter them.

This point is comparable to Heidegger's distinction between *readiness-to-hand* (*Zuhanden*) and *presence-at-hand* (*Vorhanden*). His example is that of a hammer:

The less we just stare at the hammer-thing, and the more we seize hold of it and use it, the more primordial does our relationship to it become, and the more unveiledly is it encountered as that which it is—as equipment. The hammering itself uncovers the specific 'manipulability' of the hammer. The kind of Being which equipment possesses—in which it manifests itself in its own right—we call 'readiness-to-hand'. (Heidegger, 1927, 15: 98)

In hammering, the hammer is not encountered as an object, but as something with which one can do something. It is a medium, something that one uses to engage in the act of hammering. The same point can be made with regards to language. A linguistically competent human being does not usually encounter her own or her conversation partner's words as objects of attention, that is, as being present at hand. When talking, we encounter language as ready-to-hand, that is, we are not directed at it as a concrete object with specifiable and objective properties, but as something we can use to communicate.

It helps here to think of the act of speaking as vocal gesturing. Just as in regular gesturing, it is 'the process of *doing* the gesture, *the actual motion*, that is thinking' (Steffensen, 2011, 192). Take the example of ostensive gesturing: when I point at an object, I am directed at that object, and wish to direct someone else's attention towards that object. My hand with which I do the pointing is not the intentional object of the pointing gesture, nor does someone else who just looks at my hand 'gets' what the pointing is supposed to do. Similarly, when I talk about something, the focus of my attention is at that which I talk about, rather than at the words I use to talk about it.

And just as when our relation to the hammer is significantly altered when it breaks, as it suddenly appears present-at-hand as an object, the same shift in attention can occur in our relation to language. This happens for instance when communication breaks down, that is, if we do not know what a speaker means by a particular word. We can then ask the speaker what he means, by quoting his original words, treating them in an object-like fashion (‘what do you mean by “X”’). This is in line with the distinction Tarski proposed between object-language and metalanguage: it is only when we direct our attention at pieces of language, that language itself becomes the object (object-language) we are directed at, and about which we can then talk (in a meta-language). Of course there are plenty of other examples of language showing up as the intentional *object*, including *inter alia* writing poetry, exegesis of texts, and, notably, doing *philosophy of language*.

When we communicate linguistically in normal circumstances, language is used as a medium, that is, we are not intentionally directed at utterances. Once we can speak, we can also shift our attention toward the medium, and speak *about* language. It thus seems wrong to think of ‘linguistic objects of attention’ in describing what subjects are directed at in a normal conversation. Uttering a sentence is not the same thing as taking a sentence to be true, just as knowing *how to use* language is not to know things *about language*. In other words, linguistic competence is a form of know-how, not a form of knowing that (Van Elk, Slors, & Bekkering, 2010). The former ability must come first, for it is only when we can talk at all that we can talk about language.

Before I discuss the implication of this distinction between the two relations we can have towards language in the next Section, I first look at the possibility of making sense of the notion of a proposition on Hutto’s account in the rest of this Section.

3.3 The Nature of Propositions

As we have seen, Hutto (2008, 87) does not want to pass judgement of the ontological status of propositions. However, whilst introducing the notion of content as relying on a Mode of Presentation in the context of basic cognition, Hutto (2008, 44) states that ‘Fregean thoughts, “senses” (*Sinn*) or their equivalent—that is, contents—’ can be either ‘intellectually “graspable” objective entities existing in public space’, or ‘private, inner mental representations’ (c.f. Hutto, 2006d, 210n3). We have already seen that the latter option, which is essentially a Language of Thought, is explicitly denied by Hutto (2008, 88–95) as it entails contentful basic cognition.

The former option, however, also seems to be incompatible with Hutto’s proposal. Remember that the process of radical interpretation—together with the coherence and correspondence principles of charity (§2.2.1)—aims to fix propositional contents with regard to (1) their place in a subject’s logical geography, and (2) the referential properties of the meaningful subparts of a sentence (their denotational content). This does not sit well with thinking of communication as involving the mutual grasping of *the same* thought content by several individuals, which the graspable abstract object-view of propositions entails. Moreover, the idea that having a contentful thought essentially involves the ‘grasping’ of an such an abstract entity—that is, one that exists before it is grasped—seems to be incongruent with the Developmental-Explanatory Thesis (p. 5), which entails that cognitive structures emerge from a history of recurrent interaction.

Both traditional approaches to understanding propositions are thus unavailable to the Radical Enactivist. Note that on these traditional accounts, it makes sense to say that one is directed at a proposition, as these are objects (either public and abstract, or private and mental) and can thus be *intentional* objects. But what then would be an alternative conception of a proposition, such that it is both congruent with Radical Enactivism *and* can be the intentional object of an attitude? Can we even make sense of the notion of a proposition in this context?

Just after his refusal to pass judgement on the ontological status of propositions, Hutto (2008, 87) says that ‘my interest is only in the idea that in order to have a content-involving propositional attitude one must have facility with complex linguistic forms—sentences of some kind or other.’ Being directed at a proposition thus is nothing more than being able to use language (Hutto, 2008, 96). However, it now seems unclear what the proposition, as an additional entity besides the symbols of natural language, is supposed to be. If we exclude both public abstract objects and private mental objects, the only alternative seems to be the public concrete objects. But obviously, in linguistic communication there are no other kinds of concrete object besides the linguistic symbols themselves and the objects to which they allegedly refer.

We have to realise that these traditional interpretations of propositions rely on a code-metaphor of linguistic symbols (Love, 2004). Traditionally, these symbols are thought of as meaningless by themselves. It is only because they are translated to private mental objects (or public abstract ones) that *are* intrinsically meaningful and representational, that they acquire meaning. On these accounts, meaning thus properly exists on a level that is distinct from the symbols we use to refer to these meanings. If we, however, take the language of thought to be natural language, as Hutto does, the code-metaphor becomes ill-fitting. As we thus lose the ability to specify ‘meaning’ in terms of mental objects that are intrinsically meaningful, it becomes unclear what the proposition expressed by these concrete linguistic symbols is supposed to be.

The notion of a proposition thus becomes unintelligible. The meaning of linguistic symbols are not additional objects that exists over and above them at which we can be directed. However, as we have seen, in linguistic communication we are also not intentionally directed at the symbols themselves. Something has got to give if we want to get out of this predicament. In the next Section, I therefore look at how we have ended up here, and how we should reconceptualise language to be congruent with Enactivism.

4 Why Language is not an Abstract Symbol System

In this Section, I look into how we arrived at the problems we faced in the previous Section. I argue that the root problem is that Hutto conceives of our linguistic competence as being made possible by an abstract system of symbols, and that conceiving of language in this way robs us of the possibility to explain how language works. If we disregard the cognitive and interpersonal dynamics from which our symbolic capabilities arose, we make ‘language into a medium that renders understanding, speech, and thought independent of living human beings’ (Cowley, 2011b, 2): it becomes disembodied. I therefore first argue that Hutto does conceive of language in this fashion (§4.1). Thereafter, I show that Davidson, who is the main source of inspiration for Hutto’s account of language, never meant his theory of meaning as a theory that could also explain our linguistic abilities (§4.2). After that I give some general objections to thinking of language in terms of an abstract symbol system (§4.3). I then introduce a novel conception of a symbol, according to which first-order dynamic unfolding of communication situations is constrained by second-order constructs, which consist of socio-culturally selected verbal patterns (§4.4). These considerations will provide the basis for showing the implications of relinquishing an abstract symbol system view of language for the notion of content, and the alleged distinction between the basic and supermind which I consider in the final Section of this thesis.

4.1 Language as an Abstract Symbol System

First of all, it has to be noted that accepting *both* the Thinking in Natural Language hypothesis (TNL) *and* the idea that the supermind manipulates abstract amodal symbols is *prima facie* inconsistent. In order to see why, we have to remember that the TNL entails that the structured vehicles of linguistic thought are those of natural languages. If one speaks, one produces sound waves. These symbols are thus first and foremost concrete and spatio-temporally extended phenomena that are used as symbols. This was also one of the main reasons why Hutto opted for the TNL over the LoT (Hutto, 2008, 88).

Note that shifting attention from overt speech acts to intraverted speech acts does nothing to change the concrete nature of the linguistic symbols. In line with the discussion of recreative imaginings (§2.1.3), these re-enactments of linguistic symbols are sensorimotor in nature. The difference between speaking and thinking in one’s head, or *intraverted speech*, then boils down to the suppression of outward responses⁹ (Bottineau, 2010).

The crucial part in Hutto’s idea that linguistic symbols become abstract seems to be that they stand for something else, that is, that they refer to, or denote a certain state of affairs or object. This is apparent from how he describes the origin of content by example

⁹One can convince oneself of this by a little experiment: first say a sentence out loud. Then ‘say’ the same sentence, still moving one’s lips, but suppress the vocalisation. Finally, suppress even the moving of the lips. One can still ‘hear’ one’s voice in one’s head. When intravertedly speaking to oneself, one can for instance change one’s tone of voice (try saying the sentence, while suppressing outward vocalisation and movement of the lips, first in an angry tone of voice, and then in a friendly tone of voice.), ‘speak’ faster or slower, ‘speak’ in a particular dialect etc.

of the symbolic antics of a chimpanzee named Sheba. The task that experimenters put to Sheba was to select between two piles of candies, one bigger than the other. The twist is that when she chose the bigger pile, she would get the smaller one, and vice versa. When confronted with this choice, Sheba would *always* select the larger pile. Hutto (2008, 98) describes that ‘much to her manifest frustration, Sheba could not overcome this tendency even after she apparently grasped the experimenter’s rule.’ However, through years of intensive training, Sheba managed to form a link between numerals and the quantities of candy. When confronted with the numerals, she *was* able to select the smaller numeral, which led to her receiving the larger pile of candy. Now Hutto (2008, 265n14, emphasis mine) concludes that:

Unlike cases of purely indexically or iconically based responding, this really is a situation in which a simple public sign is being used to stand for something, denotationally and referentially. *Here we might begin to talk of content legitimately.*

Although Hutto makes it clear that learning a handful of symbols is a far cry from actual linguistic abilities, this quote makes the reliance of Hutto’s account on denotational properties clear. As Sheba only has very primitive symbolic abilities, we can conclude that for Hutto denotation is the most primitive symbolic operation, for it is this *standing for something else* which marks the boundary between basic iconic modes of responding and symbolic abilities. The problematic move here is that Hutto seems to imbue linguistic symbols *themselves* with these referential properties. This is apparent from the fact that the difference between basic and linguistic cognition is described as being able ‘to think using detached representations’ (Hutto, 2008, 60) that, crucially, are *context-invariant* (Hutto, 1999, 124).

The referential properties of linguistic symbols are described as being conventional in nature. Whilst discussing the origins of language, Hutto (2008, 237, emphasis mine) describes this processes as follows:

it was the combining of a generative capacity with lexical abilities that would have been the eventual basis for the emergence of a symbolic protospeech—a protospeech of a kind that would have required forging of *conventional links* between vocal gestures and their referents, converting the former into *conventionalized symbols*; these would have been the ultimate basis on which the language makers developed *stable symbol systems*.

It can now be made clear what Hutto means by abstract symbols: the concrete symbols of natural language are thought to be tokens of an abstract type. That Hutto essentially thinks of language as such a stable system of abstract symbols is further apparent from his use of the intensionality test. As we have seen, the intensionality test positions particular words in a logical geography. The relations between words we thus discover are not relations between concrete linguistic symbols (tokens), but between words abstracted from their context of utterance: the words, *as types*, are positioned, not particular utterances of words. The intensionality test is thus conceived of as some sort of measurement process, leading to an explicit *theory* of meaning for a particular person’s language. The metaphor of intensional cartography therefore expresses the following picture: we survey a person’s linguistic-cognitive landscape, and draw a logical chart of the relations we encounter in doing so.

As we have seen, this logical geography can ‘only lay the ground for understanding the *structure*, not the content, of propositional thought’ (Hutto, 2008, 129). By using principles of charity (§2.2.1) we assume that the speaker whose words we are trying to interpret responds to the same worldly offerings as we would, and can thus come to know how the logical structure of one’s thought relates to the world, that is, what the referential relations are between certain words and worldly objects. This turns the logical chart into an intensional chart.

By treating language in this way we thus abstract away from the context of utterance and treat the concrete symbols of natural language *as if* they were something *in themselves*. The following quote of Hutto (2008, 87, emphases mine) further confirms this interpretation:

Sentences have the right form and content to be truth-evaluable; *they* are true or false depending on how things stand with the world. If the world obliges, if the relevant state of affairs obtains, then *the proposition expressed by a given sentence will be true*. If not, *the believed proposition is false*. It is no accident that the relevant state of affairs must be picked out disquotationally.

Here we see that the sentence itself, rather than the utterance of a sentence, is taking to be the bearer of truth. Sentences themselves are the things that are true or false. We have to realise that this theory of meaning is given, following Davidson, in Tarski T-sentences of the form “‘*p*’ is true in \mathcal{L} iff *p*’ (Hutto, 1999, 127–8). That is, the relevant state of affairs is picked out disquotationally. In the process of radical interpretation, we take our interpretee to be speaking this language \mathcal{L} , and, together with the principles of charity, this T-sentence therefore implies that ‘whenever a person utters “*p*”, he means *p*’. Prior to a person uttering a particular sentence, given that we have a T-theory for his language, we thus already know what the sentence means, and to what worldly objects its meaningful subparts refer.

The general picture is thus as follows: when we learn a language, we learn *that* a particular word *stands for* a particular worldly object. The most primitive linguistic operation is therefore denotation. A word is conceived of as an abstract type, for we not just learn that this particular word, understood as a concrete token, refers to a worldly offering, but that any instantiation of this word will have this property in virtue of *being the same word*. Because other people similarly know that this word stands for this object, we can use it to communicate. *How* we conceive of the denoted object can be ascertained by radical interpretation.

However, if we want to explain how creatures can understand these abstract symbols *as* denoting, we need an account of how a creature understands this denotational relation and instantiates it in speaking and understanding a language. That is, we need an account of how this denotational relation between an abstract linguistic symbol and its referent class is cognitively realised in a creature. In other words, we want to know how the supermind works.

4.2 Davidsonian Interpretationalism and ‘How it Works’

Hutto (1999, 123) is surely aware of the tension that exists between an interpretationalist account of meaning and an explanation of our linguistic competence, as is apparent from his quoting Davidson (1990, 325):

The approach to the problems of meaning, belief, and desire which I have outlined is not, I am sure it is clear, meant to throw any direct light on how in real life we come to understand each other, nor how we master our first concepts and our first language.

Hutto (1999, 115–119) claims to advance on Davidson by being able to show how the acquisition of words works, and therefore, what the origins of our conceptual abilities are. Following Davidson, he explains this in terms of social triangulation: a teacher and a child respond similarly to worldly objects, the teacher recognises that the child’s responses are similar to his own, and the child in its turn recognises that its responses are similar to the teacher’s. Based on this *referential triangle* (see Figure 2), the child can learn a word for this worldly object. Of course, this requires a host of abilities, like the ability for non-linguistic joint attention for instance (Hutto, 2008, 137). I shall not go into how exactly this functions, as a detailed description of the foundations of this referential triangle falls outside the scope of the current discussion.

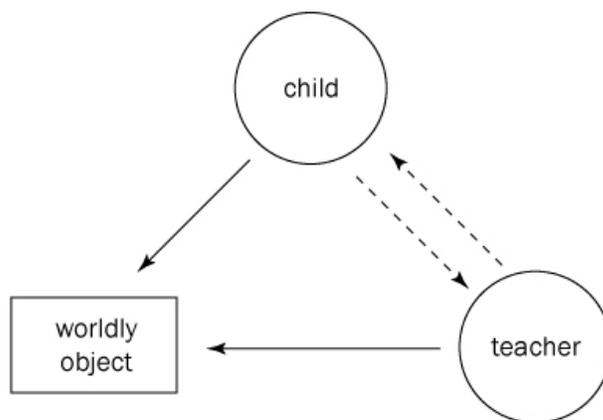


Figure 2: The referential triangle, in which both the child and the teacher respond similarly to a worldly object (solid lines), the teacher recognises that the child’s responses are similar to his own, and the child in its turn recognises that its responses are similar to the teacher’s (dashed lines). Based on Hutto (1999, Figure 17, 118).

The only point I want to make here is that Hutto (1999) supplements the idea of triangulation with an account of contentless basic cognition. In this way, he claims to be able to explain why our responses would be similar to one another’s: *we have the same inherited biosemiotic¹⁰ directedness to the world*. Our non-conceptual abilities thus allow us to bootstrap this process of triangulation by which a child can learn *that* a word stands

¹⁰Hutto (1999) talks of a ‘modest biosemantics’. His later renouncement of the semantic aspects these basic modes of responding, leading him to his biosemiotic account, does not have an impact on this line of argument. I therefore gloss over this distinction.

for a worldly object. I return to this idea of social triangulation, and why it is unable to account for referential properties of symbols in §4.3.2.

For it is not just the acquisition of words that Davidson thought was not explained by his theory of meaning. The trouble runs deeper: the enterprise of giving an explicit theory of truth—or meaning, when augmented with principles of charity—was not meant to throw light on *how* a person’s linguistic competence *works*. It is rather a tool that *we*, as interpreters interested in *the* meaning of utterances of a person, have to employ. It is a description, rather than an explanation of how linguistic competence works (Davidson, 1986, 438). In Hutto’s account of language, however, these two projects—giving an explicit theory of meaning through radical interpretation on the one hand, and explaining our linguistic competence on the other hand—are conflated. As we have seen, this is apparent from the metaphor of intensional cartography: the logical relations between words are *uncovered*, that is, not constructed, in the process of radical interpretation. However, the question how this intensional chart is realised in a particular person is left unanswered.

This leads us to the question whether the Davidsonian interpretationalist framework Hutto adheres to can explain how our linguistic competence works, as Hutto surely wants to do with his account of the supermind. In the next subsection, I argue that it is not only left unanswered, but that we will be unable to answer it within an interpretationalist framework because this framework itself rests on problematic assumptions about language.

4.3 Objections to an Abstract Symbol System-View

Why can a theory as proposed by Hutto not explain how our linguistic competence works in concrete linguistic communication? As we have seen in this Section so far, the crucial point is that Hutto treats the symbols of natural language *as if* they themselves have specific properties which allow for communication: words denote worldly objects and sentences express propositions which are true or false. This is made possible by conceiving of concrete linguistic tokens as instantiations of abstract types. Hutto’s picture of language is thus *representational*: sentences re-present certain states of affairs in truth-evaluable ways. Here I focus on two arguments against this view on language. *Firstly*, I show that we are unable to decide what the criteria are to decide whether two tokens of a word fall under *the same type* (§4.3.1). *Secondly*, I argue that language is not denotational, but connotational in nature, by closely looking into what a child can learn in the social triangulation in terms of which Hutto conceives the ontogenesis of symbolic activities (§4.3.2). Thereafter, I show how language relies on our basic modes of interacting with each other, that is, on the dynamic unfolding of a communication situation. I give an example of how voice dynamics contribute non-trivially to the sense invested in a linguistic symbol (§4.3.3). Finally, I argue that the primary function of language is not to represent the world, but to coordinate our behaviour (§4.3.4).

4.3.1 Repetition and Reification

First, let us look into the alleged abstract nature of linguistic signs. Love (2004, 539) states that this idea ‘ultimately derives from the ease with which the possibility of repetition conduces to reifying a repetitium’. When we hear someone say a word, we can repeat this

word. The first and second utterances of this word now seem to be two tokens of the *same* word, leading to the conclusion that there must be some third thing, the abstract type.¹¹

However, deciding sameness of concrete linguistic symbols is not a straightforward matter. Take the word ‘bank’ for instance, which could mean either a financial institution or a side of a river.¹² This places the radical interpreter for a choice: do we treat ‘bank’ as one word in the logical geography we are constructing, or as two distinct words, say *bank*₁ for the financial institution and *bank*₂ for the riverside? In this instance, opting for the latter seems the best option as the two words are clearly distinguishable in their meaning. In other words, the same form ‘bank’ stands for two distinct words.¹³

The process of radical interpretation continues, and our interpretee utters the sentence ‘she banked a cheque’. Similar problems arise: do we introduce a novel word into our logical geography (e.g. *bank_{verb}*), or do we take this to be an instantiation of *bank*₁? Once again, if we opt for the former, where do we position the verb bank as used in ‘I banked the eight ball off two cushions’? Perhaps we should introduce a new word once again. As the process of radical interpretation continues, there will be more and more of these decisions to be made. The Oxford Dictionary¹⁴ lists four main definitions of ‘bank’ (the two nouns and two verbs we already encountered), which are further subdivided into twelve definitions, some of which are again subdivided leading to a total of twenty-six definitions of the word ‘bank’. Does this entail the existence of the words *bank*₁ all the way to *bank*₂₆? Or do we want to count the definitions ‘A transverse slope given to a road, railway, or sports track’ and ‘The land alongside or sloping down to a river or lake’ as the same? It seems that no non-arbitrary criteria can be formulated to decide these matters. As this example shows, simply relying on the formal properties of the allegedly repeated sign is unsatisfactory.

It is important to realise that this example generalises to all instances we might classify as a recurrence of *the same* bit of language. One can ask the question for instance, whether the sentence ‘It’s nice weather today’ is the same sentence when uttered sincerely on a sunny day, or with a sarcastic tone of voice during a hurricane; or whether ‘*It’s* nice weather today’ is the same sentence as ‘*It is* nice weather today’. The mere fact that the form is the same gives us no reason to assume that the alleged contents expressed will be the same.

Although we might thus have the intuitive idea that our language consists of readily re-identifiable and separable units called words—each of which is coupled to a meaning that is, in principle, specifiable—, deciding exactly when a word is repeated will be impossible

¹¹This reification of linguistic symbols is greatly facilitated by the fact that we can write our utterances down. Looking at written language can, however, be misleading. To see why this is the case, it is important to realise that written language postdates spoken language, both ontogenetically and socioculturally. Although the invention of the text might have had transformative effects on language, we can only understand written language relative to spoken language. This is still apparent from *how* we read. We do not look at the written forms and instantly determine their ‘meaning’, but rather recreatively imagine the phonological experience. For a discussion of distorting effects of the *written language bias* see Linell (2005).

¹²This example is based on an example by Love (2004).

¹³Alternatively, we could say that the form ‘bank’ expresses two distinct contents. The word ‘bank’ would then express the contents *bank*₁ in the case of a financial institution, and *bank*₂ in the case of the riverside. This alternative formulation will however procure the same difficulties for the radical interpreter, for in either way, the same form ‘bank’ must take up two positions in the logical geography.

¹⁴Consulted through <http://www.oxforddictionaries.com/>.

in real-life linguistic encounters. In Love’s (2004, 540) words: ‘We may be inclined vaguely to conceive of our language as made up of determinately reified linguistic entities, but this conception happily co-exists in our minds with any amount of banal evidence to the contrary.’ The reification of words thus consists on a *naïve realism* with regards to linguistic forms (Cowley, 2011a).

Relinquishing this naïve realism has profound effects on the enterprise of radical interpretation. If we are unable to specify what the relata are to be in the alleged logical geography, that is, if we are unable to identify the linguistic units we are supposed to be locating, the process of radical interpretation grinds to a halt. In other words, if the decision about how many words we take ‘bank’ to stand for is arbitrary, then the logical geography based on these decisions will be as well. And if the logical geography is arbitrary, it will not be able to perform its function of ‘fixing’ the contents of an utterance. At the same time, we cannot deny that language consists of these recurring forms. In our example, ‘bank’ is surely a recurrent form. What I deny here is that we can extract a determinate amount of distinct uses for a particular form, that is, we cannot ascertain ‘the meaning’ of distinct uses of the form ‘bank’ by abstracting away from concrete utterances. In other words, there is no content to be ‘fixed’ through a process of radical interpretation.

In terms of the previous discussion, it is the act of directing our attention to language as opposed to using a word, that leads to the illusion that there must be an abstract entity of which a concrete utterance is an instantiation. It is only when we conceive of linguistic symbols *as if* they are something in themselves, that is, independent of them being used in concrete acts of linguistic communication, that the view arises that our language consists of readily re-identifiable forms that each have a specifiable meaning, or a set of specifiable meanings.

4.3.2 From Denotation to Connotation

Let us now turn our attention to the alleged determinate referential or denotational properties of linguistic symbols. As we have seen, Hutto wishes to distinguish between indexical basic cognition and intensional linguistic cognition, where reference is supposed to be fixed and context-invariant (Hutto, 1999, 124). The following question however is left unanswered: in what way does a person know *that* a certain word stands for—denotes, refers to—a (set of) worldly offerings?

The traditional way of spelling this out would be to rely on a LoT which is thought to be intrinsically representational. Communication is made possible because a speaker encodes her mental representation into a linguistic form, and the hearer subsequently decodes this form into the same mental representation. Symbols thus denote worldly objects *because* they are encodings for mental representations that are thought to have intrinsic referential properties (see figure 3). On these traditional accounts, words only get their referential properties in virtue of the intrinsic referential properties of mental representations. If we now cut out the middle man in this picture, that is, the mental representations, the relation between the symbols of natural language and their referents becomes questionable (Kravchenko, 2007).

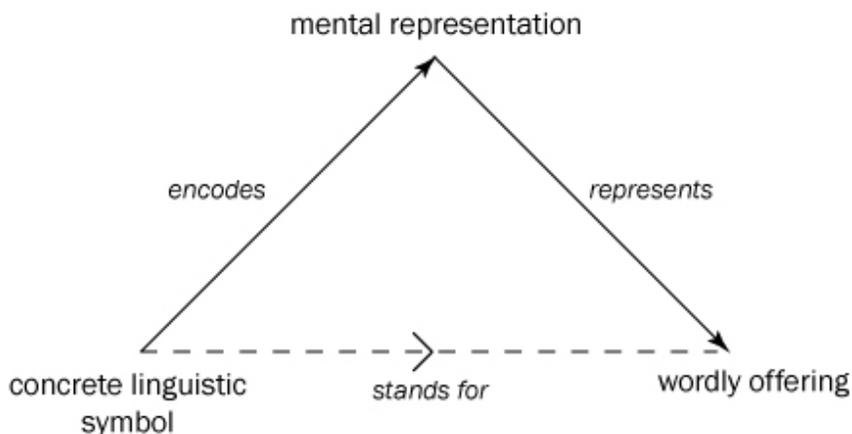


Figure 3: Semantic triangle. Adapted from Kravchenko (2007, Fig. 1, 656).

Supermental thinking by itself can offer no explanation, as its medium *consists only of* the concrete symbols of natural language, either in the world, that is, as sound waves, or as sensorimotor re-enactments of these concrete symbols in the case of ‘thinking in one’s head’. As supermental thinking is supposed to be transparent, we would know if additional entities were involved. Knowledge of the conventionalised referent of linguistic symbols is therefore not to be found ‘in the supermind’. By relinquishing the idea of an LoT we are thus unable to understand *how* linguistically capable creatures can know *that* a certain words stand for (a class of) worldly offerings.

To get us out of this predicament, let us look at the process of social triangulation as Hutto envisages it in detail. I use the same example as earlier, that is, a child learning the word for table. Once the referential triangle is set up (the teacher responds to the table, the child responds similarly to the table, and they both recognise that they respond similarly to the table), the teacher can start teaching the child the word. To do so, she directs its attention at the table, and utters the word ‘table’. By doing this for some time, whilst rewarding the child when it tries to mimic her utterance, the child starts to associate the table, as it conceives of it, with the word. Once the association is formed, the child can begin to use its newly-acquired word.

However, the way the child makes use of the word might be surprising. For we have to remember that one of the assumptions of Hutto was that the teacher conceives of the table similarly as the child does. But this by itself is not obvious. We might for instance imagine that the child is always fed whilst sitting at the table, and it only sits at the table to be fed. Like Pavlov’s dogs, the child’s mouth might water when it is seated at the table. Based on this previously formed association, which surely is a part of how the child conceives of the table, it might use the word ‘table’ to indicate that it is hungry.

From this example, three things might be learned, which I discuss in what follows. Firstly, based on this social triangulation, a child can only associate a linguistic symbol with how it already conceives of the worldly offering, that is, not with the worldly offering *as such*, or as it is conceived by the teacher. Secondly, the linguistic communication takes place against a backdrop of cognitive and interpersonal dynamics which contribute in non-

trivial ways to the meaning of utterances (§4.3.3). Thirdly, children start out by doing things with words, rather than by representing the world linguistically (§4.3.4).

Firstly, the process of social triangulation only allows the child to associate words with what it can already re-identify or recreatively imagine, that is, not with objects as such. As we have seen, these re-identification abilities form the basis for conceptual, linguistic abilities. However, this implies that denotation cannot be a primitive linguistic operation, for this would require a consensus both on denotant and the denoted (Maturana, 1978, 50). It is the consensus on the denoted which is problematic. In our example, both child and teacher would have to conceive of the table similarly. But this is an impossible condition, since how they conceive of the table is determined by their respective histories of dynamic interaction, which are never equal, as is made clear by the Developmental-Explanatory hypothesis (p. 5).

It is important to realise that how pre-verbal infants conceive of worldly offerings, that is, how they actively engage with and recreatively imagine them, is not unstructured before they are initiated into the realm of language. Non-verbals surely have complex associations between Recreative Imaginings (ReIm)s¹⁵. Through the recurrent simultaneous occurrence of different perceivable events, the ReIm)s of these events form a multi-modal associative network¹⁶. As these ReIm)s are sensorimotor re-enactments, they are not merely images. Pavlov’s dogs, for instance, did not just form an ‘image’ of food in their minds when they heard the bell ring, if at all. Their reaction through association was fully embodied. The triggering of one ActCR can thus trigger a cascade of associated recreative imaginings. I shall call such conceptual nodes in a network of associations a *Complex Recreative Imagining* (CReIm). In such a CReIm, no ReIm ‘originally took precedence or acquired unique significance as compared with others’ (Kravchenko, 2007, 653). That is, any ReIm which forms a part of the CReIm can trigger it. There is thus no single ReIm that acts as a ‘label’ for a CReIm.

We should be careful not to reify either ReIm)s or CReIm)s, by reminding ourselves of their nature: as sensorimotor re-enactments, they are, just as on-line sensorimotor loops, spatio-temporally extended dynamic processes. Recreatively imagining a table is thus not calling an image of a table before one’s mind, but rather, to re-enact the embodied

¹⁵This account of association between ReIm)s is based on Kravchenko (2007), which I translated into the conceptual framework we have been using so far. Kravchenko (2007) uses the term *elementary representation* (ER), which he defines as ‘relative neuronal activities’, which are ‘determined by the structure of the nervous system at the moment at which the behaviour is enacted and, therefore, is *circumstantial by nature*. It follows that adequate behaviour is necessarily only the result of structural matching between an organism as a dynamic system and its medium (which compromises other organisms structurally matched to *their* medium).’ (Kravchenko, 2007, 653). We can think of these ERs essentially in terms of ActCRs and their ReIm)s—understood as sensorimotor re-enactment. This is because the Developmental-Explanatory hypothesis, when combined with the extensive nature of basic cognition—as we have seen, two assumptions of RadEn—, entails that according to RadEn a creature’s cognitive structures are shaped by a history of structural coupling with the environment. I therefore take ERs to be similar to ReIm)s. Accordingly, I shall call Kravchenko’s ‘Complex Representations’ Complex Recreative Imaginings (CReIm)s.

¹⁶Recreative Imaginings are of course not limited to the visual modality. We can also recreatively imagine the sensorimotor experience of smells, sounds, etc. These networks can include ReIm)s from all these modalities.

experience of seeing a table, that is, to actively engage with an imagined table (Myin & Zahidi, 2012). In other words, as sensorimotor *re-enactments*, ReImS rely on the same cognitive structures that sensorimotor loops rely on.

Words, understood as concrete entities, are in essence not different from other stimuli. When we say that the word ‘table’ is associated to tables for the child, we can therefore not understand this as a process of labelling. Just as the sight of a table might trigger the ReIm of the word ‘table’, perhaps leading to an utterance of the word, hearing the word ‘table’ triggers an CReIm of tables. The word therefore does not denote a concept, it rather becomes part of a network of structured CReImS that is already present. Hearing a word activates a network of memories (Thibault, 2011), which are sensorimotor in nature. We thus conceive of linguistic symbols as ‘operators of reminiscence’ (Bottineau, 2010, 283). Our conceptual abilities are thus connotational rather than denotational in nature (Kravchenko, 2007).

Words, as readily re-identifiable signs, allow us to organise our experience of the world. Because they are easily perceivable as being the same in varying contexts, they function as excellent reminiscence operators. They therefore allow for a much stronger associative bond than other worldly offerings. The fact however, that the word is easily perceived as similar also brings forth ‘the nominalistic illusion that any word does revolve around a firm conceptual core’ (Bottineau, 2010, 283).

We are now in a position to understand the idea of a non-linguistically mediated propositional attitudes we encountered earlier. Let us return to the example of re-identifying the milkman. Any worldly offering, whether online perceived or off-line recreatively imagined, that is associated with the concept (the CReIm) of the milkman can activate this CReIm. This CReIm includes, *inter alia*, the word ‘milkman’. We thus always re-identify the milkman ‘as such’ *once* our conceptual abilities are in place, even if we do not *say* so, either out loud or to ourselves. And although it is only because we can re-identify particulars in *similar ways* to one another—that allows for the bootstrapping of the process of social triangulation—this does not entail that we conceive of these particulars exactly the same. The word thus does not come to *stand for* the object; it is rather *associated* with the ReIm’s that pertain to the object, that is, it is associated with the object as we already conceive of it on a pre-verbal level, which is based solely on our history of interaction—understood as structural coupling—with the object.

4.3.3 Language and Dynamics

The second lesson to be learned from the example of social triangulation is that linguistic communication always takes place against a background of dynamics that are already present. Besides the cognitive dynamics just discussed, there are also interpersonal dynamics, that is, dynamics that only emerge in the unfolding of an—at least dyadic—communication situation. One can think for instance of synchronising of bodily movements (Thibault, 2011), voice dynamics, gaze and task-oriented modes of action (Steffensen, 2011, 191) and emotional coordination (Rączaszek-Leonardi, 2009, 162). Just as linguistic communication relies on the presence of cognitive dynamics, it also relies on the presence of these interpersonal dynamics. Moreover, these interpersonal dynamics contribute non-trivially to the sense of verbal patterns.

One example of the reliance on interpersonal dynamics in the case of spoken linguistic communication is given by Cowley (2009, 70), who describes the social interaction between an Italian father, daughter and mother:

Briefly, while at dinner, the conversation turns to the peas they are eating. The daughter asks if they were grown in the garden which, for her mother, is unfortunate. She points out, emphatically, a ‘certain person’ did not cut pea-poles (by implication, she could not grow peas). Evidently, she is ‘getting at’ her husband. Having held his counsel for a long conversational time, he then reminds her that he did cut pea-poles. Instead of falling silent (or apologising), she mounts a new attack. Yes, he did cut pea-poles but, she claims, they were longer than the room in which they are sitting (later said to be more than two metres). At this point, father and daughter *ridicule her, good naturedly and in harmony*. Why ascribe *this* sense?

Father and daughter do so by using the Italian word *oeu*, a response cry. What is interesting is how father and daughter use their voice dynamics to enact this sense. To understand how these voice dynamics contribute to the sense of what is being said, we have to ‘turn up the microscope’ (Thibault, 2011, 18), that is, look at how events unfold at the pico-scale, which comprises events ranging from milliseconds to fractions of seconds. Cowley (2009, 71) describes the dynamic unfolding of the voice dynamics as follows:

First, the father, whose mean utterance-initial pitch is 127 Hz interjects in *his wife’s pitch range*. His utterance starts at 180 Hz half an octave up and, in 400 milliseconds, has reaches [sic] the same pitch (to within 4.5 Hz) at which his wife ceased speaking (220 Hz). By ‘using’ a rise-fall tone (conventional in such settings) he directs ridicule—not at what she says—but (as we say) *at her person*. Then, 300 milliseconds after her father starts speaking, their daughter joins in (with an audible smile). She too matches the final pitch of her mother’s utterance and, so, harmonises with her father (on the same pitch contour). Their vocal chords vibrate at the same rate for about 500 milliseconds. When the father stops, his daughter’s pitch drops slightly and she gives a little laugh (at lower volume). Manifestly she has enjoyed the events.

In this communication situation, a lot is going on that is not captured when we direct our attention merely to the verbal patterns. To a person who is unable to attune to the dynamic unfolding of voice dynamics, the whole sense of the encounter would be lost. *Simply knowing the meaning of the words that the participants use in this situation does not allow one to know what is happening*. What is important in this encounter is not so much *what* is said, that is, the alleged content of the words used, but the subtle interbodily dynamics that unfold in real-time (Steffensen, 2011, 193). It is the attunement to the dynamic unfolding of voice dynamics in a triadic communication situation that determines the sense invested in the word ‘*oeu*’.

To further illuminate this point, remember the trivial example we mentioned earlier, when we asked the question whether the sentence ‘It’s nice weather today’ is the same sentence when uttered sincerely on a sunny day, or with a sarcastic *tone of voice* during a hurricane. If a person were to answer to the sarcastic utterance by stating that the other person is surely wrong, that it is not nice weather at all, we would have to conclude that he did not *understand* what the other person was saying. His failure to understand, once again, lies not in his ignorance of the meaning of the words; rather, it lies in him

missing an attunement to both the context and the dynamic unfolding of voice dynamics, that is, to how people *use* words to communicate by relying on dynamic aspects of the communication situation.

In the referential triangle from which symbolic capabilities emerge, a host of dynamic phenomena are thus unfolding. Here I looked into the role of voice dynamics, but similar arguments can be mounted for other dynamic phenomena such as gaze dynamics and (facial) gestures. The symbols of natural language rely on these interpersonal as well as cognitive dynamics. Therefore, verbal patterns radically underdetermine the sense of a particular utterance (Cowley, 2011a). Before I turn our attention to how symbols are related to the real-time dynamic aspects of linguistic communication, there is one final lesson to be learned from our example.

4.3.4 Language as an Extension of Social Behaviour

The third lesson that can be learned from the example of social triangulation is that ‘children start off by *doing* things with words’ (Rączaszek-Leonardi, 2009, 170). Once the child has associated the word ‘table’ with tables as it conceives of them, the child can use the word to indicate to others that it is hungry. When we are talking, we are thus ‘doing something’ rather than merely representing some state of affairs (Rączaszek-Leonardi & Kelso, 2008). Language allows us to coordinate our behaviour in ways that are not available to non-verbals (e.g. Cowley, 2011b; Rączaszek-Leonardi, 2009). Thibault (2011, 2) stresses this ‘centrality of coacting agents who extend their worlds and their own agency through embodied, embedded processes of languaging behaviour.’ Learning to talk is thus not an individual achievement, that is, a creature learning to use a language system by coming to know *that* symbols stand for worldly offerings. Language ‘must be based on social coordination’ (Cowley, 2011b, 11), that is, on dynamic coordination processes—interpersonal dynamics—that are already present.

To see how this distinction between language as an extension of social behaviour and language as representation changes the ways we perceive of what language is, let us return to the example of Sheba. For it is telling how Hutto conceives of the symbolic abilities of this chimpanzee. Remember that Sheba was able to select the smaller pile of candies if she had to choose between numerals, while she always selected the larger pile when confronted with actual candies. According to Hutto (2008, 98) ‘representing the object of her desire in a more abstract way gave her a new means of achieving an old end. She was able to use the symbol as a go-between—as a detached representation of the food quantity—and this freed her from her “more natural,” nonsymbolic embodied mode of responding.’ Hutto (2008, 99) describes her acquisition of this ability as follows: ‘Her training required an extremely long regime in order to establish the relatively simple associative link between different quantities of food and symbols for these quantities.’

Hutto thus conceives of Sheba’s accomplishment as an individual achievement. It was because she managed to ‘get’ that symbols stand for food quantities that she was able to curb her natural responses. Steffensen (2011), however, notes that the two experimental situations have very different interactional dynamics. In the case when Sheba had to select between two piles of candies, a chimp-candy relation unfolds. However, the introduction of numerals, as the alleged neutral go-between, brings a host of interactional dynamics

into play. Even in the absence of the experimenters, the ‘numeral marks a history of human-chimp coupling’ (Steffensen, 2011, 196–197). Interaction with symbols thus evokes a history of complex social coupling from which these symbolic capabilities emerged. Accordingly, ‘the candy problem is solved, not by Sheba alone, but cooperatively by Sheba and the human experimenter, drawing on a history of training’ (Steffensen, 2011, 197).

Shanker and King (2002) describe how incorporating this idea of language as an extension of social behaviour into the ‘training regimes’ of apes can lead to remarkable improvements in linguistic abilities. Traditionally, ape language research was based on techniques that are used in behavioural modification therapy. As was the case with Sheba, apes were trained to associate a certain sign with a worldly offering by rewarding them with food. Ape language researcher Sue Savage-Rumbaugh however, decided to move away from this behaviour modification paradigm, in which apes are taught arbitrary symbol object-pairings, to the use of linguistic symbols in ordinary social circumstances. Rather than focussing on what went on in the apes heads, Savage-Rumbaugh shifted attention to what apes could do with words (Savage-Rumbaugh & Lewin, 1994, 49).

By making use of a *lexigram board*¹⁷ that feature symbols that are useful in day to day interaction, Savage-Rumbaugh was able to teach apes linguistic abilities that were previously deemed impossible for apes to achieve (see figure 4). The most remarkable fruit of Savage-Rumbaugh’s work with apes are the linguistic abilities of the bonobo Kanzi. He mastered the use of over 250 symbols on a lexigram board, which he used ‘purposefully, without cuing or imitation, to do such things as refer to objects and locations in the immediate, present surroundings, as well as to other that are “absent,” and even, to comment on events that occurred in the past, ask questions, play games, or simply provide information (both requested and unsolicited)’ (Shanker & King, 2002, 614). Moreover, he could comprehend spoken English sentences, recognising at least 650 words, showing comprehension levels similar to a two year old human child.

What is interesting is how Kanzi acquired these linguistic abilities. Originally, his adopted mother Matata was the object of study, and Kanzi merely tagged along. He saw how Savage-Rumbaugh was trying to get his mother to use the lexigram board. Around the age of six months, Kanzi ‘became mesmerized by the keyboard’ (Savage-Rumbaugh & Lewin, 1994, 129). Around the age of two, Kanzi started to use lexigrams to request and name worldly offerings. Crucially, he learned all of this ‘without any direct instruction’ (Shanker & King, 2002, 616). Accordingly, Savage-Rumbaugh and her team decided to ‘abandon any and all plans of [formally] teaching Kanzi and simply to offer him an environment that maximized the opportunity for him to learn as much as possible’ (Savage-Rumbaugh & Lewin, 1994, 137). The lexigram board was redesigned to incorporate only symbols that would be useful in day-to-day activities, such as ‘the names of foods, caregivers, other apes, locations in the forest, toys and games’ (Shanker & King, 2002, 616). Rather than submitting Kanzi to a painstaking ‘training regime’, he was thus able to

¹⁷A lexigram board is a portable interface that features buttons marked by symbolic inscriptions (lexigrams), that, when pressed produce audible wordings. As chimpanzees are unable to control their vocalisations to produce vocal gestures as humans can, they have to rely on such a device to produce utterances. For an interactive virtual reproduction of such a lexigram board, as used by Savage-Rumbaugh, see <http://www.iowaprimatlearning.org/interactive-lexigram/interactive-lexigram/index.html>



Figure 4: Kanzi using the lexigram board to communicate with Sue Savage-Rumbaugh. Photo: greatapetrust.org.

acquire language in pretty much the same way as human children normally do. That is, in a dyadic relation with a caregiver, in which the focus is laid on what one can do with language¹⁸, rather than on how symbols represent the world. One of the key factors in this process of language acquisition was that Kanzi ‘was aware that we employed the keyboard as a means of communication and apparently felt keenly motivated to do so as well’ (Savage-Rumbaugh & Lewin, 1994, 139).

The role of the experimenters in this developmental process is crucial. Kanzi’s linguistic abilities emerged from a history of social dyadic interaction with his caregivers. These encounters were structured by the caregivers, for instance by the use of the lexigram board. Kanzi’s abilities therefore rely on the interpersonal (or inter-hominid) dynamics that were at play in this history of social interaction. In Shanker and King’s (2002, 612) words: ‘what Savage-Rumbaugh was doing was a crucial factor in what [the ape] was doing, just as what he was doing was a crucial factor in what she was doing. The manner in which each of them was acting was part of their shared history together and the evolving dynamics of their co-regulated interactions.’ It is thus not merely tasks and tools that ‘made this cognitive and communicative development possible; the presence of an unusually responsive and emotionally attuned caregiver was absolutely crucial’ (Shanker & King, 2002, 613). Moreover, what licenses us to attribute linguistic skills to Kanzi is how he interacts with his caregivers by responding to their linguistic acts and uttering sentences through manipulating the lexigram board, that is, *not* by an alleged mental transformation (such as the construction of a linguistic supermind).

¹⁸For instance, one of the first things Kanzi did with the lexigram board initiating a game, by selecting the ‘apple’ lexigram followed by the ‘chase’ lexigram before running off with the apple with a play-grin on his face

It is therefore by replacing the behaviour modification paradigm with a dyadic communication situation, that Savage-Rumbaugh was able to achieve these remarkable results. Rather than being explicitly taught *that* certain symbols stand for certain worldly object, Kanzi learned *how* to use language by Savage-Rumbaugh encouraging him, and, crucially, herself, to co-regulate their activities through language. Linguistic abilities are thus not adequately captured when we describe them as knowing that a symbol stands for a worldly offering, which is how Hutto envisages the role of the supermind. It is an extension of social behaviour, and thus remains reliant on our already present social coordination of activity and basic modes of interacting with the world.

4.4 The Complex Interplay of Symbols and Dynamics

In this section I have shown that if we conceive of language merely as a collection of abstract symbols, the logical relations between them, and their conventionally determined denotational properties, we cannot understand how linguistic activity is possible. We should therefore shift emphasis from how philosophers and linguists have described language to how linguistic communication is possible, whilst realising that ‘the focus on forms blinds us to how language is anchored in human life’ (Cowley, 2011b, 5). Hutto’s insistence that there is a clear-cut distinction between linguistic cognition and basic cognition entails that there is a clear-cut distinction between aspects of one’s behaviour that are linguistic and aspects of it that are not. However, in practice, this distinction between linguistic and non-linguistic behaviour is impossible to make. The example we have seen of the father and daughter ridiculing their wife/mother, shows that by abstracting away from the dynamics of a communication situation, in this case the voice dynamics, we are unable to determine the sense invested in the verbal patterns in the concrete act of linguistic communication. So the dynamics of a communication situation, which are deemed non-linguistic on the abstract symbol-view, matter to the sense of what is communicated.

Moreover, as we have seen, learning how to speak is not an individual achievement. As the example of Kanzi makes strikingly clear, linguistic behaviour emerges in social coordination and should be seen as building on already present interpersonal and cognitive dynamics, that is, as a way of co-regulating action. Language is thus *dialogical* in nature (Cowley, 2011a), and therefore cannot be explained solely in how individuals engage in languaging events. However, a description of language purely in terms dynamics terms seems a step too far. Language is unmistakably symbolic in nature. In the rest of this section I look at how we should reconceptualise language in such a way that we can do reconcile both dynamic as well as symbolic aspects of linguistic communication.

4.4.1 First-Order Languaging and Second-Order Language

In order to get to grips with the complex interplay of symbolic and dynamic aspects of language, Love (2004, 530) proposes the following distinction:

a language is a second-order cultural construct, perpetually open-ended and incomplete, arising out of the first-order activity of making and interpreting linguistic *signs*, which in turn is a real-time, contextually determined process of investing behaviour or the products of behaviour (vocal, gestural or other) with semiotic significance.

First-order *linguaging*—a neologism coined by Maturana (1970)—is, in the words of Thibault (2011, 7) ‘whole-body sense-making activity that enables persons to engage with each other in forms of coaction and to integrate themselves with and to take part in social activities that may be performed either solo or together.’ This first-order linguaging relies, as we have seen, not only on symbolic vocalisations, but also on interpersonal and cognitive dynamics. Language is therefore a ‘heterogeneous set of physical, cognitive and social activities that unfold in real-time’ (Steffensen, 2011, 193).

At the same time, our linguistic abilities are only possible because we can draw on a rich sociocultural history. As we have seen, Kanzi only acquired his linguistic abilities because his environment was structured—in this case wilfully by his caregivers/experimenters—in ways as to promote his linguistic skills. For human language acquisition, the same principle holds. In human communities, the linguistic behaviours of others structure the environments in which infants develop. Language is thus a crucial part of the ecological niche of developing humans, which makes a description of linguistic capabilities in individualist terms highly suspect. In the words of Steffensen (2011, 195) ‘the “individual cognizer” is a *contradictio in adiecto* when it comes to *linguistic cognition*.’ Once we accept that minds are extensive and rely, at least in the case of linguistic cognition, on the active engagement with symbols refined through selection pressures operating on the time-scale of sociocultural evolution, we cannot help but conclude that our thinking is shaped by language, just as language is shaped by our thinking. Our thinking is therefore never purely or own, but determined, in part, by the structure of the linguistic communities in which we find ourselves. This conclusion is in line with the concept of structural coupling that we encountered in explaining basic cognition.

Accordingly, Steffensen (2011) makes the analogy between language and beaver dams: just as the beaver dam outlives the beavers that contribute in its construction, the evolution of human language occurs at time-scales that transcend individual life-spans. Those phenomenologically stabilised patterns that we find no use for will get selected out, whereas new patterns emerge to solve novel coordination problems. Crucially, this evolutionary process of selection occurs ‘not blindly, but as a result of behaviour that emerges from doing things with other people’ (Steffensen, 2011, 194). Second-order symbolic—that is, *lexicogrammatic*—patterns thus emerge from first-order linguaging events and are shaped by sociocultural evolutionary processes. Based on the interactional needs and motives of a population, first-order linguaging dynamics such as vocal tract gestures are selected on and, in turn, come to serve as symbolic constraints on first-order linguaging (Thibault, 2011).

First-order linguaging events thus take place within a wider context of language evolution. Cowley (2011a) states that this entails that linguistic events cannot be traced to a localised cause, that is, to an individual’s knowledge of a language system (knowing-*that* certain words stand for certain worldly offerings). Our linguistic capabilities cannot be explained solely in terms of the here-and-now-environment, but depend on the history of previous interaction, both socioculturally as well as ontogenetically. For just as first-order linguaging depends on the slower time-scale of language evolution, language evolution itself depends on, or consists of nothing more than, recurrent events of first-order linguaging. Language therefore cannot be explained, in Whitehead’s (1926) terms, on the *assumption of simple location* (Cowley, 2011a). Language is *non-local*, and, accordingly,

cannot be localised in the linguistic supermind of an individual ‘language user’. A synchronic description of an alleged language system is thus always insufficient. In the words of Rączaszek-Leonardi & Kelso (2008, 6):

Instead, therefore, of searching for “essences” of symbols at [one] time scale one might rather ask: Which variables were important enough in the relation between the environment of an individual or in the relations between individuals such that they were encompassed by symbols?

Moreover, we should be careful not to conceive of second-order language, that is, the selection on verbal patterns on the time-scale of sociocultural evolution, as a monolithic structure. We should rather conceive of it as a meshwork of bio-socio-ecological relations (Steffensen, 2011, 203). To give an obvious example, architects do not talk in the same way as structural engineers do. This can be explained by the different coordination problems that these ways of talking evolved to deal with, roughly, to design aesthetic and functional buildings on the one hand and making sure that these do not collapse on the other hand. But to flesh out the concept of a meshwork, one can also think of the rapid evolution of slang, of inside jokes that only make sense to a group of friends, or of the often highly idiosyncratic ways of talking that evolve in a romantic relation.

4.4.2 Linguistic Symbols as Replicable Constraints

How should we understand symbols once we realise that linguistic communication is aimed at the coregulation of action, emerges from the complex cognitive and interpersonal dynamics of first-order languaging events, and relies on physical symbols—sound waves in talk? Building on the biosemiotic work of Howard Pattee, Rączaszek-Leonardi (2011) conceives of linguistic symbols as replicable constraints. To understand what this constraining relation entails, Rączaszek-Leonardi & Kelso (2008, 194) give the example of grasping a mug:

To picture arising of coordinated action in a somewhat simplified fashion, look at your hand, wave it, wave your fingers: it is a system with many degrees of freedom, constrained by the physical structure of your bones, joints, and muscles (this constraining is probably also ‘carefully’ shaped by evolution, which happens on a different time scale). When unconstrained, it can move freely in many directions limited only by its physical structure. However, when an action is performed, such as grasping a mug, some movements of some parts of this multielement system have to act together in a spatially and temporally organised fashion. It is important to realise that the limitations imposed on the hand’s degrees of freedom are not specified once and forever. One does not just switch on some automatic function called ‘grasp a mug’. The action, and therefore the binding together of the degrees of freedom, has to be adapted, on-line, to the actual demands of the environment, such as the actual dimensions and shape of the mug, temperature of the liquid inside, etc. Still another important fact is that the specification of action to fit different situations has to be incomplete: one can grab a mug, even a specific mug at a specific distance, in many different ways: e.g., with a palm turned outside, with a hand holding a pen, etc. Therefore, flexibility of action depends on leaving some degrees of freedom unbound.

The main idea is thus as follows: just as our intention to grasp a mug constrains the degrees of freedom of our hands, albeit incompletely, the public system symbols of natural language constrain cognitive and interpersonal dynamics, based both on the ontogenetic history of an individual as well as the ecological niche shaped by sociocultural evolution in which this ontogenetic development took place. A simple example might clarify this constraining role of language in linguistic communication. When I ask a friend what she had for dinner last night, I evoke her memories. If I had not asked her, she might have been thinking about anything; my question thus constrains her thinking. We therefore conceive of a linguistic symbol as ‘a mind-guiding piece of voicing’ (Bottineau, 2010, 274).

These symbols can arise because some aspects of a dynamic interaction are perceived to be stable. Exactly how this process would have taken place in our socio-cultural history is outside the scope of this thesis. I do think however, that Hutto’s mimetic abilities hypothesis might prove illuminating for this account (see Hutto, 2008, Ch. 11 & 12). On this account, recreative imagining abilities underlie mimetic skills that could account for the stabilising of interpersonal dynamics through recurrent mirroring of each other and aspects of the environment (e.g. wiggling one’s arm to indicate the presence of a snake).

It is important to realise that ‘signhood’, that is, being a symbol, is not an intrinsic property of an object, but is dependent on whether it is used as such. In the words of Love (2004, 531): ‘Signhood is conferred on a sign—on what thereby becomes a sign—if and when human beings (or other semiotically competent creatures) attach a signification to it that goes beyond its intrinsic physical properties’. All aspects of our behaviour which are in principle perceivable as ‘the same’, that is, which are re-identifiable, can thus be enlisted to function semiotically: gestures, vocal gestures, voice dynamics, environmental props, etc. Note that one does not need to have explicit knowledge of *how* these dynamics are used. As we have seen in the example of the father and daughter ridiculing their wife/mother, the rise-fall tone the father used to ridicule is conventional in nature. The fact however, that the father is able to use these voice dynamics to impart a certain sense on his words, does not entail that he can describe how he does this. Just as one can know how to ride a bike without knowing *what* one does to keep the bike upright, the father can use the conventional rise-fall tone to convey a certain sense without having explicit knowledge of *how* his voice dynamics contribute to this sense. Once again, knowing how to communicate linguistically is a form a know-how.

When we conceive of symbols as replicable constraints, they thus ‘do *not* “map like a code” to the consequences of their actions’ (Rączaszek-Leonardi, 2011, 308). They are naturally context-dependent, as their constraining influence on dynamics depends on which dynamics are present at the time of their acting on a dynamic system. As physical symbols that are replicable through mutual mirroring, linguistic symbols can be selected based on the outcomes of their constraining relation to dynamic systems, that is, the effectiveness in inter-individual coordination (Rączaszek-Leonardi, 2011).

It has to be noted that constraining relations come in different degrees: it can be very strong, for instance in the case of scientific constructs, or very small, for instance in the case of indexical words such as ‘this’. We can even construct languages in which we artificially exclude all dynamics by explicitly defining the meanings of symbols and the rules according to which they have to be manipulated. Think for instance of formal languages or the operation of digital computers. The rules for manipulation are given by

logic, and symbols *are* used denotationally: we can reach consensus on both the denotant as well as the denoted, for instance ‘let p be a cat’. However, the symbols in such formal systems are not grounded (c.f. Harnad, 1990) and thus are not meaningful.

This novel conception of the relation between symbols and dynamics gives an answer to this *symbol grounding problem*, by fundamentally changing the way we conceive of ‘the meaning’ of a symbol. Rather than describing this in abstract denotational terms, ‘meaning is what an expression does in a situation’ (Rączaszek-Leonardi, 2011). Because of the context-dependency, the meaning of a symbol will thus be determined both by the historical selection pressures that acted on that particular symbol—spanning the time-scales of sociocultural linguistic evolution and ontogenetic development—*and* the situation in which it is used. The dynamics thus not only contribute to the meaning, they constitute the meaning. The meaning of a symbol is nothing more than how it constrains the unfolding of interpersonal and cognitive dynamics.

5 Conclusion

In this Thesis, I have argued against the idea that our linguistic abilities can be characterised by conceiving of language as a collection of symbols which have specifiable meanings and determinate denotational content. Language is rather an extension of social behaviour, based on associative abilities rather than denotation. Language is something we do. Second-order symbols emerge from first-order languaging, which is the dynamic unfolding of social interaction processes. Their meaning is not specifiable without reference to the cognitive and interpersonal dynamics from which they arose, as they can only constrain those dynamics that are present. To think of language as consisting of the concrete symbols of natural language, that is, to give up the idea that symbols encode mental representations, therefore entails that no principled distinction can be made between aspects of our behaviour that are linguistic, and aspects that are not: language is a radically heterogeneous phenomenon. It allows us to coordinate our behaviour, and derivatively, our thinking, in ways that are unattainable for non-verbal creatures. As our linguistic abilities are only made possible by these second-order symbols that were forged on time-scales that transcend our own lifetimes, language is a non-local phenomenon. Learning to talk is therefore not merely an individual achievement, but is better described as a person adapting to an ecological niche which is formed by the linguistic behaviour of others.

This conceptual shift in thinking about language mirrors the conceptual shift that Radical Enactivism proposes for basic cognition. We have to substitute a static description in terms of representational symbols for an explanation which takes the unfolding of dynamics of a situation into account, in order to do justice to both basic as well as linguistic cognition. The representationalist model is thus not only unable to explain our basic cognition, it is also insufficient for explaining our linguistic abilities. It is only when we relinquish this model—which is based on notions of truth and logic (computation)—*completely* that we can unlock Enactivism’s full potential. This entails that, just as when we consider basic cognition we ought to opt Radical Enactivism over conservative Enactivism, when we consider language we ought to opt for Really Radical Enactivism over Radical Enactivism. Only by becoming truly radical, that is, by a root and branch denial of the model that wishes to explain our thinking in terms of abstract context-invariant symbols that represent, can we do justice to the complex phenomena that the dynamic unfolding of our thinking is.

In the rest of this closing Section, I look into what relinquishing this model entails for our understanding of language, and how this Really Radical account differs from Hutto’s Radical Enactivist approach.

The core of language is neither truth nor representation. As Hacker (2014) notices, in his words, the *meaning-calculus view of language* has lost track of what it means to give the meaning of a symbol. Remember that Hutto states that the relevant state of affairs has to be picked out disquotationally. The reason for doing so is ascertaining whether a particular sentence is true or false. Thus we say that “*p*” is true in \mathcal{L} iff *p*’, or to give a concrete example: “‘Roses are red’ is true in English iff roses are red’. But this tells us exactly nothing about what ‘Roses are red’ *means* in any ordinary sense of the word. In Hacker’s (2014, §3) words, reiterating a person’s words will not tell us anything about the

meaning of them, ‘any more than Englishman can make themselves better understood abroad by loudly repeating what they said’.

But it is this very process of reiteration that lays at the basis of the questionable reification of linguistic symbols. It is only when we direct our attention at the symbols that we use in acts of languaging, and we can find no other way of talking about them then by quoting them, that the illusion of a single conceptual core—‘the meaning’—of a word arises. However, by trying to specify this meaning in terms of truth and logic, that is, in terms of denotation and logical geographies, we tend to lose sight of what we normally mean by the meaning of a word. When I ask a person what he means by his words, I expect him to answer me in words that differ from the ones he originally used, and to enlighten me by given an analogy, an example, drawing something on a piece of paper, miming, etc.

Once we define the meaning of a symbol as *the socioculturally selected constraining of first-order interpersonal and cognitive dynamics*, we can see that denoting is not a primitive linguistic operation, but one of the things we *can do* with language. By thus letting go of truth and representation as the primary driving forces behind our linguistic abilities, we also evade the strict distinction between linguistic and non-linguistic aspects of our behaviour. Languaging together is a heterogeneous set of activities that encompass much more than the mere exchange of symbols.

Language is not an individual achievement. Language is not something that we do by ourselves, like breathing. It is something we do together. In languaging together, we enlist everything that we can in order to communicate. Even when these dynamics are artificially removed—by relying for instance on techniques such as the written word—and it seems as though we can only rely on the linguistic symbols themselves, our efforts to make sense of what we encounter still evoke a history of previous encounters. As a non-local phenomenon, a concrete act of languaging, such as reading, talking, and even ‘talking to oneself’—that is, thinking—marks an ontogenetic and sociocultural history. Even in our most intimate thinking, there is thus ‘always an anonymous third party present’ (Døør, 1998, 40, as cited in Steffensen (2011)). Once we realise that linguistic symbols are not encodings for mental representations, but socioculturally selected constraints on first order languaging, we have to conclude that our thoughts are never solely our own.

For although language greatly enhances our possibilities, it also constrains them. The words we use determine what we can say. The linguistic ecological niche in which we are born is already structured in remarkably complex ways. By acknowledging these structuring selection forces on language that extend beyond ourselves, we let go of the last vestiges of the default internalist assumption, according to which *thinking* is something which happens exclusively in our minds. Language and its meaning is therefore not properly localisable in an individual using a language—understood as a stable system of abstract and context-invariant symbols—as Hutto’s idea of radical interpretation as an intensionality test could lead us to believe. It is a spatio-temporally extended phenomenon which we keep alive (the seizing to be of a particular language is known for instance as language death, linguistic extinction or even linguicide) through re-enacting it in concrete acts of languaging, and which constantly evolves in order to provide answers to new problems we encounter. Our linguistic abilities simply cannot be seen in separation of these larger

time-scales. A child that grows up with no human contact does not learn to speak its own language; it remains mute.

By focussing only on the symbolic aspects of language, and explaining them in terms of representation and truth, Hutto paints a picture of language that leaves us with the same problem that Enactivism originally diagnosed in Cognitivism. For just as the founders of Enactivism realised that classical Cognitivism had ‘virtually nothing to say about what it means to be human in everyday, lived situations’ (Varela et al., 1991, xv), we have seen that Hutto’s account of our linguistic abilities has virtually nothing to say about what it means to actually talk to another person. To make contact with the phenomenon of language anew, we have to look at what actually goes on in the dynamic unfolding of concrete acts of languaging. This entails to give up notions such as ‘*the* meaning of a word’ and ‘*the* proposition expressed by a sentence’. Once we realise that words are no more than guiding pieces of voicing, we have no choice but to be really radical in our denial of both representationalism and intellectualism, that is, we have to opt for Really Radical Enactivism.

References

- Bottineau, D. (2010). Language and enaction. In J. Stewart, O. Gapenne, & E. Di Paolo (Eds.) *Enaction: Towards a New Paradigm for Cognitive Science*, (pp. 267—306). Cambridge, MA: The MIT Press.
- Cosmelli, D., & Thopson, E. (2010). Embodiment or envatment?: Reflections on the bosily basis of consciousness. In J. Stewart, O. Gapenne, & E. Di Paolo (Eds.) *Enaction: Towards a New Paradigm for Cognitive Science*, (pp. 361—386). Cambridge, MA: The MIT Press.
- Cowley, S. J. (2009). Langauge flow: Opening the subject. *Cognitive Semiotics*, 4, 63—91.
- Cowley, S. J. (2011a). Distributed langauge. In S. Cowley (Ed.) *Distributed Language*, (pp. 185—210). Amsterdam/Philadelphia: John Benjamins.
- Cowley, S. J. (2011b). Taking a language stance. *Ecological Psychology*, 23, 1—25.
- Davidson, D. (1986). A nice derangement of epitaphs. In E. Lepore (Ed.) *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson.*, (pp. 433—446). Oxford: Basil Blackwell.
- Davidson, D. (1990). The structure and content of truth. *Journal of Philosophy*, LXXXVIII(6), 279—328.
- Davidson, D. (1991). Three varieties of knowledge. In P. Griffiths (Ed.) *A.J. Ayer: Memorial Essays.*, (pp. 153—166). Cambridge, Ma.: Cambridge University Press.
- Dennett, D. C. (1987). *The Intentional Stance*. Cambridge, MA.: The MIT Press.
- Dennett, D. C. (1991). *Consciousness Explained*. New York: Little Brown and Co.
- Dennett, D. C. (1998). Reflections on language and mind. In P. Carruthers, & J. Boucher (Eds.) *Language and Thought: Interdisciplinary Themes*, (pp. 284—294). Cambridge, Ma.: Cambridge University Press.
- Di Paolo, E. A., Rohde, M., & De Jaegher, H. (2010). Horizons for the enactive mind: Values, social interaction, and play. In J. Stewart, O. Gapenne, & E. Di Paolo (Eds.) *Enaction: Towards a New Paradigm for Cognitive Science*, (pp. 33—88). Cambridge, MA: The MIT Press.
- Di Pellegrino, G., Fadiga, E., Fogassi, I., Gallese, V., & Rizzolatti, G. (1992). Understanding motor events: a neurophysiological study. *Experimental Brain Research*, 91, 176—180.
- Døør, J. (1998). *Moralske Meditationer. To Essays*. Odense: Odense Universitet.
- Dretske, F. (1988). *Explaining Behavior: Reasons in a World of Causes*. The MIT Press.
- Frankish, K. (2004). *Mind and Supermind*. Cambridge, Ma.: Cambridge University Press.

- Hacker, P. M. S. (2014). Two conceptions of language. *Erkenntnis*, *online*, n.p.
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, *42*, 335–346.
- Heidegger, M. (1927). *Being and Time*. Oxford: Basil Blackwell. Translated by J. Macquarrie and E. Robinson, 1967.
- Hurley, S. (1998). *Consciousness in Action*. Cambridge, MA.:Harvard University Press.
- Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, MA: The MIT Press.
- Hutto, D. D. (1999). *The Presence of Mind*. Philadelphia/Amsterdam: John Benjamins.
- Hutto, D. D. (2000). *Beyond Physicalism*. Philadelphia/Amsterdam: John Benjamins.
- Hutto, D. D. (2005). Knowing *what?* radical versus conservative enactivism. *Phenomenology and the Cognitive Sciences*, *4*, 389–405.
- Hutto, D. D. (2006a). Against passive intellectualism: Reply to crane. In R. Menary (Ed.) *Radical Enactivism: Intentionality, Phenomenology and Narrative*, (pp. 121–150). Amsterdam/Philadelphia: John Benjamins.
- Hutto, D. D. (2006b). Both bradley and biology: Reply to rudd. In R. Menary (Ed.) *Radical Enactivism: Intentionality, Phenomenology and Narrative*, (pp. 81–105). Amsterdam/Philadelphia: John Benjamins.
- Hutto, D. D. (2006c). Embodied expectations and extended possibilities: Reply to goldie. In R. Menary (Ed.) *Radical Enactivism: Intentionality, Phenomenology and Narrative*, (pp. 157–178). Amsterdam/Philadelphia: John Benjamins.
- Hutto, D. D. (2006d). Four herculean labours: Reply to hobson. In R. Menary (Ed.) *Radical Enactivism: Intentionality, Phenomenology and Narrative*, (pp. 185–222). Amsterdam/Philadelphia: John Benjamins.
- Hutto, D. D. (2006e). Unprincipled engagement: Emotional experience, expressions and response. In R. Menary (Ed.) *Radical Enactivism: Intentionality, Phenomenology and Narrative*, (pp. 13–38). Amsterdam/Philadelphia: John Benjamins.
- Hutto, D. D. (2008). *Folk Psychological Narratives: The Sociocultural Basis of Understanding Reasons*. Cambridge, Ma.: The MIT Press.
- Hutto, D. D. (2011). Philosophy of minds new lease on life: autopoietic enactivism meets teleosemiotics. *Journal of Consciousness Studies*, *18*(5-6), 44–64.
- Hutto, D. D. (2013). Radically enactive cognition in our grasp. In *The Hand, an Organ of the Mind: What the Manual Tells the Mental*, (pp. 227–252). Cambridge, MA.: The MIT Press.
- Hutto, D. D., & Myin, E. (2013). *Radicalizing Enactivism. Basic Minds without Content*. Cambridge, Ma.: The MIT Press.

- Kiverstein, J., & Clark, A. (2009). Mind, embodied, embedded, enacted: One church or many? *Topoi*, 28, 1—7.
- Kravchenko, A. V. (2007). Essential properties of language, or, why language is not a code. *Language Sciences*, 29, 650—671.
- Linell, P. (2005). *The Written Language Bias in Linguistics: Its Nature, Origins and Transformations*. London: Routledge.
- Love, N. (2004). Cognition and the language myth. *Language Sciences*, 26, 525—544.
- Maturana, H. R. (1970). *Biology of cognition*. Urbana: University of Illinois.
- Maturana, H. R. (1975). The organization of the living: A theory of the living organization. *International Journal of Man-Machine Studies*, 7(3), 313—332.
- Maturana, H. R. (1978). Biology of language: The epistemology of reality. In G. A. Miller, & E. Lenneberg (Eds.) *Psychology and Biology of Language and Thought: Essays in Honor of Eric Lenneberg*, (pp. 27—63). New York: Academic Press.
- Maturana, H. R., & Varela, F. J. (1987). *The tree of knowledge: The biological roots of human understanding*. Boston: Shambhala Publications.
- Menary, R. (2006). What is radical enactivism? In R. Menary (Ed.) *Radical Enactivism: Intentionality, Phenomenology and Narrative*, (pp. 1—12). Amsterdam/Philadelphia: John Benjamins.
- Menary, R. (2010). Introduction to the special issue on 4e cognition. *Phenomenology and the Cognitive Sciences*, 9, 459—463.
- Millikan, R. G. (2005). *Language: A Biological Model*. Oxford University Press.
- Myin, E., & Zahidi, K. (2012). Het bereik van het mentale. van uitgebreid naar omvattend geheugen (the extent of mind. from extended to extensive memory). *Tijdschrift voor Filosofie*, 74, 103—127.
- Noë, A. (2004). *Action in Perception*. Cam.
- Noë, A. (2009). *Out of our Heads*. Hill and Wang.
- Putnam, H. (1967). Psychological predicates. In W. H. Capitan, & D. D. Merrill (Eds.) *Art, Mind, and Religion*, (pp. 37—48). University of Pittsburgh Press.
- Ramsey, W. (2007). *Representation Reconsidered*. Cambridge, MA.: Cambridge University Press.
- Ratcliffe, M. (2009). There are no folk psychological narratives. *Journal of Consciousness Studies*, 16(6—7), 379—406.
- Rączaszek-Leonardi, J. (2009). Symbols as constraints: The structuring role of dynamics and self-organization in natural language. *Pragmatics & Cognition*, 17(3), 653—676.

- Rączaszek-Leonardi, J. (2011). Language as a system of replicable constraints. In *Laws, Language and Life: Howard Pattee's classic papers on the physics of symbols*, (pp. 295—333). Dordrecht: Springer.
- Rączaszek-Leonardi, J., & Kelso, J. A. S. (2008). Reconciling symbolic and dynamic aspects of language: Toward a dynamic psycholinguistics. *New Ideas in Psychology*, *26*(2), 193—207.
- Rudd, A. (2006). Unnatural feelings: A non-naturalistic perspective on the emotions. In R. Menary (Ed.) *Radical Enactivism: Intentionality, Phenomenology and Narrative*, (pp. 65—80). Amsterdam/Philadelphia: John Benjamins.
- Savage-Rumbaugh, S., & Lewin, R. (1994). *Kanzi: The Ape at the Brink of the Human Mind*. John Wiley & Sons, Inc.
- Shanker, S. G., & King, B. J. (2002). The emergence of new a paradigm in ape language research. *Behavioral and Brain Sciences*, *25*, 605—656.
- Steffensen, S. V. (2011). Beyond mind: An extended ecology of languaging. In S. Cowley (Ed.) *Distributed Language*, (pp. 185—210). Amsterdam/Philadelphia: John Benjamins.
- Stewart, J. (2010). Foundational issues in enaction as a paradigm for cognitive science: From the origin of life to consciousness and writing. In J. Stewart, O. Gapenne, & E. Di Paolo (Eds.) *Enaction: Towards a New Paradigm for Cognitive Science*, (pp. 1—32). Cambridge, MA.: The MIT Press.
- Thibault, P. (2011). First-order languaging dynamics and second-order language: The distributed language view. *Ecological Psychology*, *23*, 1—36.
- Thompson (2005). Sensorimotor subjectivity and the enactive approach to experience. *Phenomenology and the Cognitive Sciences*, *4*(4), 407—427.
- Thompson, E. (2007). *Mind in life*. Cambridge, MA: Harvard University Press.
- Torrance, S. (2006). In search of the enactive: Introduction to special issue on enactive experience. *Phenomenology and the Cognitive Sciences*, *4*(4), 357—68.
- Trevarthen, C. B. (1968). Two mechanisms of vision in primates. *Psychologische Forschung*, *31*, 299—337.
- van Elk, M., Slors, M., & Bekkering, H. (2010). Embodied language comprehension requires an enactivist paradigm of cognition. *Frontiers in Psychology*, *1*(234), 1—9.
- Varela, F. J. (1995). The re-enchantment of the concrete: Some biological ingredients for a nouvelle cognitive science. In L. Steels, & R. Brooks (Eds.) *The Artificial Life Route to Artificial Intelligence*, (pp. 11—22). Hove, UK: Lawrence Erlbaum Associates.
- Varela, F. J., Thompson, E., & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: The MIT Press.

Whitehead, A. N. (1926). *Science and the Modern World*. Cambridge: Cambridge University Press.