

Towards a Science of Definition

Anne M. Cregan^{1,2}

¹Artificial Intelligence Group
School of Computer Science and Engineering
University of New South Wales

²Knowledge Representation and Reasoning Program
Kensington Laboratory, Sydney
National ICT Australia

223 Anzac Parade Kensington, NSW 2052, Australia

annec@cse.unsw.edu.au

Abstract

The vision of the Semantic Web is to provide machine-processable meaning for intelligent applications. Whilst knowledge representation structures like ontologies now have well-developed formalisms, the issue of determining or specifying exactly what it is that they represent is still not well-understood. However, it is crucial for validation, merging and alignment, as we cannot possibly hope to judge the accuracy or applicability of a representation structure without a clear specification of what it is intended to represent. This being the case, we must either accept that our representations will have a limited applicability and lifespan, or develop methods by which we can define our terms in a robust and standardized way. Building on philosopher Richard Robinson's analysis, it is argued that 'definition' is in fact the isolation of territories within conceptual landscapes, using the four mechanisms: example, semantic relation, analysis, and rule. These mechanisms are related to cognitive processes like abstraction and categorization. We speculate that there is a common semantic ground which forms the initial basis for symbol grounding, and is then extended through cognitive mechanisms. Some starting points for identifying common semantic ground and points of divergence from it are suggested.

Keywords: Semantic Web, Ontology, Definition.

1 Why do we need a Science of Definition?

A key consideration in constructing ontologies, or indeed any form of knowledge representation, is to be clear about what it is, exactly, that the elements of the representation are to represent, and to be able to communicate this precisely. Without this, any claim that the structure is a representation at all is tenuous at best. Yet, whilst formal knowledge structures and ontology languages are now well-developed in terms of their ability to express the relationships between and attributes of their elements, and to apply reasoning algorithms on this basis, there is currently no accepted methodology for

“grounding” the elements themselves: that is, of objectively establishing exactly what the elements of the ontology are representing. However, when attempting to validate, align or merge formalisms, this information is crucial. We cannot hope to determine whether any one ontology, let alone an alignment or merging of more than one ontology, is correct, without judging whether the referring relations are correct, in the sense of matching the intentions of the authors of the ontology(ies).

But how are such intentions to be clearly stated? As there is currently no clear process for this kind of specification, such a judgement usually requires human intervention, often in the form of a dialogue or negotiation between the various parties responsible for each ontology. As is to be expected, the exercise is often labour-intensive, and it may be quite difficult to establish exactly what is meant by a term if the underlying mind-sets are divergent.

Yet the vision of the semantic web is to make meaning machine-accessible and processable. Therefore, we need to establish a method by which we can formally state the meanings of terms, so that no further human intervention is required. Such a methodology would constitute a “science of definition”

However, the principles by which one can relate a term, whether a symbol, or a word from natural language, unambiguously and unequivocally to a thing, have not yet been established. Or stated another way, there is not yet a sufficient understanding of the representation relationship. In fact, many are unsure whether it is even possible to establish such principles. It is often claimed (eg Thomas 2005) that definition is an art, not a science, and surveying the field up to this point, one has to agree. But if this is always to be the case, we are looking at a very challenging road ahead for the semantic web, if multiple parties must engage in personal negotiations every time interoperability of metadata needs to be achieved.

The purpose of this paper, then, is to investigate whether there can in fact be a science of definition, and if so, on what principles would it be based? The paper is structured as follows: Section 2 explores what it means to define something, and the different methods by which this may be done, drawing heavily on philosopher Richard Robinson's analysis (1950). Section 3 argues that Robinson's methods isolate territories within conceptual landscapes, using the mechanisms of example, semantic relation, analysis, and rule, and relates these to underlying cognitive processes like abstraction and

categorization. It claims that common semantic ground forms an initial basis for symbol grounding, which is later extended through cognitive mechanisms. Starting points for investigating this common semantic ground and the points of divergence from it are suggested. Conclusions as to the plausibility of a science of definition are drawn in Section 4.

2 Dimensions of Definition

It is commonly accepted that in setting up any formal system, such as an ontology, one must carefully define one's terms. Therefore, let us firstly examine the concept of definition itself, to see what it has to offer us, and alternately, how it might mislead us in our endeavour.

The concept of definition has had a long and tortured history dating back to the ancient Greeks, and in its long history, many strands have become tangled together, so that 'definition' has come to mean many things to many people. In untangling the various strands, the work of Oxford philosopher Richard Robinson (1950) is particularly enlightening, and this section draws heavily on Robinson's analysis to isolate the dimensions of definition and draw attention to some salient points.

Robinson recognized that the word 'definition' had been used by many people for many different things, often in ways which were inconsistent, even by the same person. From an analysis of usages of the term, he distilled several key distinctions, which we will refer to here as dimensions.

2.1 Dimension A: What is Being Related?

The first distinction drawn by Robinson is that 'definition' has been used in attempts to relate completely different kinds of entities. He found that 'definition' has been used variously to describe relations between:

A1. Words and other Words

Eg " *Let i be the square root of minus 1.* "

A2. Words and Things

Eg " *'Red' is the colour of my shirt.* "

A3. Things and Things

Eg " *$E = mc^2$* "

Robinson refers to A1 and A2 as 'Nominal Definition' and to A3 as 'Real Definition'. Although he refers to 'words' in the categories above, he intends this to include all written and spoken symbols. By analysing the three kinds of relation, Robinson concludes that what is attempted under the guise of 'definition' is a fundamentally different activity in each of the three cases. Whilst all three activities are valuable and useful, Robinson ultimately concludes that only A2 should lay claim to being called "definition".

Robinson's reasoning may be summarised as follows: A relation between words and other words (A1) is a substitution relation which may give brevity and convenience, but makes no logical contribution : it could be omitted and make no difference to the representation

or logical structure, merely rendering it more verbose and therefore more difficult to understand.

On the other hand, a relation between things and things (A3) is the main means by which scientific understanding is gained: its very purpose is to make logical contributions. Robinson refers to A3 as 'real definition', but his ultimate conclusion is that it should not be referred to as "definition" at all, because it is in fact the analysis, synthesis and improvement of concepts. That is, it is the activities of finding and describing structure within concepts, finding and describing concepts as part of larger structures, and improving the fidelity of these finding and describing activities. It encompasses the purpose of scientific endeavour very well, but does not tell us what we *mean* by something. For example (mine not Robinson's), whilst scientific theory may constantly revise and improve our understanding of the nature of light, it does not tell us what is meant by 'light': this is something that is established before the scientific investigations begin. If we establish that light corresponds to a certain band of wavelengths in the electromagnetic spectrum, this is certainly useful to know, but does not define what light is: further down the track we might hypothesise that light is energetic subatomic particles called 'photons' and not just a wave, and yet we are still talking about exactly the same entity 'light'.

Thus in Robinson's 'real definition', we are not defining a thing, but attempting to describe its relationships with other things. Whilst we may create a formal system with the aim of capturing logical or empirical relations between things, we must necessarily in such a system use symbols to represent these things. Indicating what the symbols refer to extends outside the system itself to the domain of things through Word-Thing relations (A2).

Robinson's conclusion is that only relations between words and things (A2) give information about what a symbol is referring to, or denoting. He recommends that only this activity should be considered to be "Definition". He then goes on to discuss further ways in which Word-Thing relations may be distinguished.

2.2 Dimension B: (Word:Thing) Are we Reporting an Existing Use of a Word, or Establishing a New Usage?

Robinson's analysis of "Word-Thing" (A2) definition revealed two variants:

B1. Lexical : the reporting of an existing or past actual usage of a term eg " *'compound fracture' is a medical term meaning a broken bone which punctures the skin* "

B2. Stipulative : the specification of how a term is to be used in future in a particular context eg " *where for the purpose of our analysis 'grue' means 'green before the year 2020 and blue after the year 2020'* "

Lexical definitions seek to pinpoint what was meant by some word to someone at some time. These definitions

reflect the way language is used: the definitions may be intuitive or systematic, but often are not, as they have an organic, evolutionary aspect to them as the word meanings shift and adapt over time. Dictionaries may be considered to be compilations of lexical definitions, as they are intended to assist listeners/readers and speakers/writers in using the language being defined.

Obviously when contemporary dictionaries are published, there is an assumption that the entries will continue to be accurate at least for the foreseeable future, but lexical definition should always be recognised as being in the context of who used the term and when: it is an historical report. Meanings shift over time, words come in and out of usage, whole languages come and go. Dictionary definitions of words are not god-given or absolute: they are a social agreement at a particular time and place amongst particular people, and thus belong to history rather than science.

However, in particular fields such as mathematics, or where a high level of precision is needed, *stipulative definition* may be used to specify that a certain term will be used only to mean a certain precise thing within a given context. This may be a broadening or narrowing of an existing sense of a term; a completely new usage of an existing term; or may involve the creation of a totally new term. It is a specification of how the author intends to use the term in future within a particular work or context, and binds the author to making good on that intention.

Whilst Robinson makes it clear that he considers 'definition' to be a relation between Words and Things, there is an interesting point to be made that Robinson does not identify explicitly: do we define words, or things? It seems that lexical definition takes an existing word and attempts to circumscribe the things to which it applies, whilst stipulative definition circumscribes some area within the domain of things and assigns it to a particular word. Robinson gives his exposition of methods of definition (see Section 2.3) independently of this distinction: he does not differentiate between methods for defining words and for defining things. Each of Robinson's methods has been shown diagrammatically in an effort to illustrate this point, and we will later argue for a theory of definition in which mapping in both directions is essential.

2.3 Dimension C: (Word:Thing) Methods of Definition

Robinson argues that all words are definable, in the sense of it being possible, at least theoretically, to establish what they indicate(d) to some one at some time. In fact, for something to qualify as a word at all, it must have had some meaning to someone at some time, otherwise it would be only a noise or a meaningless sequence of letters. Thus, in Robinson's view, there are no indefinable words.

Let us briefly cover Robinson's arguments against indefinability. Firstly, some have argued that since dictionaries merely relate words to other words, certain words must be indefinable: there must be some indefinable core vocabulary. Robinson points out

however, that there are non-verbal methods of definition, such as the ostensive method, and that by including such methods all words are definable. Using these methods it is possible to go beyond the lexicon, and we will argue that these non-verbal methods are essential for the grounding of language and critical for the initial stages of (first) language learning.

Secondly, it has been argued that some particular concepts of a philosophical nature, like "beauty" or "good" are intrinsically indefinable. Robinson's response is that this is not claimed because we cannot establish what they refer to, but because we are encountering difficulties when attempting to understand their structure and essential nature, which may (perhaps) be unanalysable. Thus, such a claim is confusing 'real definition' with 'nominal definition', but unanalysable does not equal indefinable. Although one may be unable to establish valid Thing:Thing (A3) relations for such concepts, we may certainly establish what the word refers to by using a Word:Thing (A2) method of definition. (It is my personal belief that such unanalysable concepts are of fundamental importance in understanding the ultimate motivators of human behaviour, in a manner akin to George Kelly's personal construct psychology (1955).

In all, Robinson identified seven methods by which Word:Thing "definition" could be established. He did not claim the division was "exhaustive, exclusive, or the only useful one": each method is simply a means of communicating to a learner what the word refers to, ie indicating what the word means. Not all methods are applicable in all cases, and in some cases a combination of methods might be needed. The seven methods of Word:Thing definition Robinson described follow.

C1. The Method of Synonyms

This method gives the meaning of a word by giving the learner a synonym for the word, with which she is already familiar eg " 'Chien' means 'dog' ".




Figure 1: The Method of Synonyms

Whilst the method requires a Word-Word correspondence (A1), Robinson includes it as a Word-Thing method (A2), as the synonym is being used to identify a thing.

C2. The Method of Analysis

This method consists of referring the learner to a thing by giving an analysis of it: eg "An 'octagon' is a polygon with eight sides". If learner knows the meanings of the other words, and is able to construct the concept in the way intended by the speaker, the learner will then know both what constitutes an octagon, and what the word 'octagon' refers to.




Figure 2: The Method of Analysis

This method thus has the advantage of imparting knowledge of the thing via the analysis, at the same time as it giving the referring relation for the word. However, it is useless in cases where analysis is not applicable, for instance for words referring to particulars. Particulars, such as a person, cannot be defined by a list of specific characteristics, no matter how long, as it could always logically belong to some other particular.

C3. The Method of Synthesis

This method identifies what a word refers to by indicating the relation of the thing indicated to other things eg “ ‘orange’ is the colour in between yellow and red”.




Figure 3: The Method of Synthesis

The thing is indicated by saying where it can be found, or what causes it: it is assigned its place in a system of relations and synthesised as a part of a whole with other things. Robinson notes that this method is indicative only; it is like saying “John is the tallest man in the room” to identify which is John. He therefore seems to imply that in order to be applicable, this method requires a background context, which seems to involve some intended frame of reference with a particular structure. Robinson also notes that the method enables great precision in indicating things where the two previous methods would fail, and is particularly useful for primary modalities of sense like ‘green’, ‘soft’, ‘sweet’ and ‘middle C’.

C4. The Implicative Method

This method defines a word by putting it in a context that defines its sense: one may determine its meaning via the meanings of the other words eg “ A square has two diagonals, each of which divides the square into two right-angled isosceles triangles”.




Figure 4: The Implicative Method

If one knows the meanings of all the other words, the meaning of ‘diagonal’ is implied. In this method the word being defined is used rather than mentioned. Consequently, unlike the other methods, this method is not based on an equivalence relation. Usually however, the statement could be reworded into another method of definition which is based on an equivalence relation.

C5. The Denotative Method

This method defines a word by mentioning known examples of things the word applies to. eg “ A ‘bird’ is things like swans, robins, geese and hens.”




Figure 5: The Denotative Method

Sometimes a thing is exactly a finite collection of examples, so a denotative definition is complete: for instance, this is the case connectives in symbolic logic, which are wholly defined by their truth tables. Usually though the list is incomplete and the learner is left to deduce an underlying connotation: that is, to take the necessary steps to abstract the concept, so that a previously unseen example could be classified. It would seem that such definitions would often require an iterative process of trial and error involving revision, broadening, narrowing etc, in order to ascertain the exact denotation of the word. In this example, the learner might initially hypothesise, for instance, that a bird is anything that has wings, and subsequently revise this notion when informed that a bat is not a bird.

C6. The Ostensive Method

Unlike the five methods above, this method does not rely on words alone but indicates a thing by drawing attention to it in some way, for instance by pointing to an actual thing, drawing the thing, or using demonstrative words either in context or in the absence of the physical object eg “ ‘Geese’ are the kind of bird we saw yesterday at the lake”.




Figure 6: The Ostensive Method

In other respects, the ostensive method resembles the denotative method as it works through giving examples.

C7. The Regular Method

Robinson points out that the six methods above are suitable for names, where a name is “a word appointed to mean always some one and the same thing, whether a particular or a general thing”. However, some words are not names, for example, eg demonstrative words like ‘this’, ‘I’, ‘yesterday’; conjunctions and logical operators like ‘and’, ‘not’, ‘but’ ‘who’, and words that do not have any common quality, but all have a particular relationship in common eg “Swiss citizenship”. In this cases the regular method is used: the meaning of these words is given by rule eg “ *I is to be used by each utterer to indicate himself*”. In practice, the rules may be either given explicitly, or deduced by the learner through exposure to examples.

3 Charting the Conceptual Landscape

We will argue that underlying Robinson’s methods of definitions are four fundamental mechanisms which are mental in nature. These are: example, analysis, semantic relation to some already grounded thing, and rule. Firstly however, let us clarify whether Robinson’s methods of definition actually relate words with things, or with our concepts of things.

3.1 Things or Concepts?

Robinson’s methods of definition were stated to be Word-Thing (A2) relations. Although Robinson did not note it, by working through some examples as presented in the previous section, it seems clear that the mapping can go either way: some methods map from Words to Things; others from Things to Words, and some seem to do both at the same time.

However, are the Word-Thing relations really mapping words to and from the domain of things, or are they actually mapping words to and from the domain of concepts of things?

If we can establish our claim that Robinson’s methods are underpinned by the four mechanisms of example, analysis, semantic relation to some already grounded thing, and rule, then clearly the latter three are mental operations and must therefore be mapping words and concepts. In the case of example, if the thing is not physically present, then the learner must be relating a word to some mental placeholder that is standing for the thing. On the other hand, if the thing is physically present, there must still be some mental placeholder if the mapping with the word is to last when the thing is no longer physically present.

Therefore, we claim that Robinson’s methods relate words not to things, but to concepts of things. We will now proceed to illustrate that his methods of definition are in fact instruments for isolating or marking out a piece of a conceptual landscape using the four mechanisms. Some of the mechanisms also involve mental operations to construct and transform conceptual landscapes prior to such an isolation.

3.2 Definition as the isolation of areas in the conceptual landscape

Our argument is as follows: Firstly, let us recognize that there is no relationship between symbols and things, other than that mediated by us humans. A tree is only a ‘tree’ if humans make it so. The symbol-making process involves having some mental representation of a thing, and giving it a name of our choosing. The name then has ‘meaning’ by virtue of its connection to the conceptual landscape. By setting up the same mappings between names and mental representations in others, we set up a basis for communicating the contents of individual mental landscapes: once the mapping between word and conceptual landscape is set up in the mind of the listener, the word may be used by a speaker to evoke the corresponding mental territory in the mind of the listener. Once evoked by a word, the mental territory provides a frame of reference for processing other information provided by the speaker. In this way, the process can be bootstrapped: once some initial mappings are achieved, they may be used in constructing others.

The method gets off the ground initially because most human mental landscapes are similar enough prior to the commencement of language learning that we can identify common landmarks, and use them as reference points. Consequently, language learning is the process of associating symbolic markers (words) with areas in the conceptual landscape. The process commences in early childhood, where language learning draws heavily on the ostensive method: conceptual representations of things directly experienced are mapped to words, under the guidance of a language speaker, usually a parent. By tracing the stages of language learning, we are able to identify four mechanisms and explain why it is that Robinson’s methods of definition are effective.

3.2.1 Examples

- **Acquiring rudimentary abstraction**
- **Acquiring symbolisation**

The process of language learning commences by relating example stimuli to symbols. Initially, the conceptual landscape arises naturally through interactions with the world, and is directly related to physical sensations. The infant perceives a collection of colours, lines, sounds and smells which arise in clusters, and learns through sensorimotor experience that these clusters are associated with particular physical objects.

The infant then learns to recognise components of the clusters which provide a similar sensory experience eg redness. By identifying and isolating common elements of the naturally arising conceptual landscape, the infant begins to develop a rudimentary ability to abstract from experience. Through the association of example stimuli with spoken words, via a process of trial, error and correction, the infant can learn that ‘swan’ is a particular cluster of sensory stimuli, and that ‘red’ is a specific component present in several sensory stimuli.

In this way, words become associated with a mental constructs, and the infant learns that symbols may be used to denote an element of their conceptual landscape: the ability to symbolize is acquired. Having experienced success in identifying and using particular components of sensory stimuli, like colours, these dimensions become reinforced within the conceptual landscape, and the infant gives more attention and weight in future.

3.2.2 Semantic Relations

- **Extending rudimentary abstraction to acquire the ability to recognise and use semantic relations**
- **Acquiring Categorisation**

The basic mechanism of abstraction learnt earlier is then used to extend the conceptual landscape, and language itself is a major enabler in this process. The naturally arising conceptual features now marked out and associated with words are used to build higher order mental constructs, further extending the conceptual landscape.

For instance, the ability to abstract common features from various sensory situations facilitates the learning of basic semantic relations like 'between', 'above' etc. The child develops the ability to recognise semantic relations and use them as links between different mental territories. Through experience, trial and error, more sophisticated mental territories are able to be marked out based on such semantic relations, and along the way, words become the markers and signposts of these territories.

The ability to understand and use semantic relations leads to the ability to deduce rudimentary rules which support the categorisation of previously unseen examples. For instance, the denotative bird example requires the marking out of a new mental territory - 'bird' - to contain existing territories already associated with words which were mapped to sensory clusters : 'swan', 'hen' etc. Children learn by experience that in such cases there is usually some reason to group the things together, perhaps the possession of some common feature(s), like having wings. Note that at this stage, categorization is based not on just one dimension of sensory experience like 'redness', but on an emergent quality based on several dimensions eg 'has wings', so semantic relations are required.

3.2.3 Analysis

- **Acquiring completely abstract concepts**
- **Using semantic relations within abstract domains**

As the conceptual landscape is built up, it is possible to move away from physical stimuli altogether and use the previously abstracted concepts as raw material to build new totally abstract concepts, with language assisting the process. This is the essence of Robinson's analysis method, which gives a method for building a new construct based on existing ones.

In the examples for 'octagon' and 'diagonal', we are given instructions for building these mental constructs by using the existing mental constructs 'polygon' and

'square'. In this way, we are able to move step by step away from the naturally arising domain of sensory stimuli into purely conceptual domains like mathematics.

Robinson's synthesis method works because the mental landscape for some underlying concept has already been marked out, and we can use our knowledge of its features to isolate or build a new construct within it or relative to it, using semantic relations. This may be either pre- or post-analysis. In the 'orange' example, the word 'between' indicates not just the presence of a semantic relation, but the existence of an underlying idea with structure: in this case, a continuum or scale by which colours are organised: we must draw on our conceptual representation of colour and how it works to isolate a mental construct within it.

Thus analysis requires the use of semantic relations within conceptual domains that have already been mapped out. Note, conversely, that Robinson's synonym method uses an established syntactic relation to a conceptual territory to evoke the mental territory and map a new symbol to it.

3.2.4 Rules

- **Acquiring the ability to circumscribe conceptual territories, and use them to classify new examples**

The rule method is a method for marking out a piece of an existing conceptual territory, which may be at any level of abstraction: the rules give borders that define the territory, which may then be applied explicitly to determine whether a mental construct falls inside the conceptual territory or outside it.

3.3 Common Semantic Ground

We have shown that words are intermediaries for communicating mental landscapes, and are effective because the underlying landscapes can be marked out by reference to common features, or what we might call 'common semantic ground'. The argument assumes that initially there are common reference points which are naturally arising, enabling the process to commence so that we can all acquire basic concepts like 'redness' and 'swan'. It seems reasonable to postulate this, as without some commonality of experience, there would be no basis on which to develop communication at all, and perhaps, one might argue, no reason to communicate. In the example discussed, we were able to pair a set of stimuli with a word ('swan'), and to isolate a particular element of several sets of stimuli with a word ('red'). This does not necessarily mean that everyone has exactly the same subjective experience of these stimuli; however, everyone has to have the same *kinds* of mappings in order for this to work.

For instance, if I experience that 'orange' is, on some dimension of my sensory experience, between the concepts I have learnt as 'red' and 'yellow', but you experience it as being between the concepts 'hot' and 'cold', then we will be unable to understand each other. But if you experience it as between 'red' and 'yellow', it would not matter whether your subjective experience of

these concepts was like my subjective experience of 'hot' and 'cold', as long as there was some homomorphic mapping between our experiences, such that the internal semantic relations hold in both cases. Thus the theory requires some predictable, common kind of relation between the stimuli and the nature of the conceptual landscape it evokes across individuals, but does not require that each individual has exactly the same subjective experience.

The issue of homomorphic semantic relations becomes clearer when we consider more complex phenomena, like events and higher order constructs. With events, for instance, if you perceive time as a linear sequence moving in a particular direction and I perceive it as a jumbled sequence with no particular direction, then our conceptual landscapes cannot be homomorphic, and we will not be able to understand each other: words like 'before' and 'after' will make sense to you but not to me, or at least not in the same way, as I will not be able to abstract the common feature which marks these out within my conceptual landscape. Consequently, I will also be unable to use past, present or future tense in language in any consistent way.

In extremely abstract scenarios, like law or politics, it is clear that without agreed common ground, accurate communication is impossible. In this case, the common ground is not naturally arising but is almost entirely constructed: clearly a legal or political system does not arise as a natural part of the conceptual landscape, but is based on many abstractions and analytic constructs. Even to define 'guilt' and 'innocence' requires a great deal of pre-existing mental architecture, and given exactly the same information, does not lead to predictable categorisation of the guilty and the innocent by different individuals, (eg different jurors) because things like value judgements and the relative importance of various factors, which may be weighted differently by different individuals, come into play.

So conceptual landscapes clearly do have significant individual and cultural differences, according to the particular ways the culture or the individual has built up higher order constructs. But at the same time, there is clearly a basic level where there is common ground, otherwise we would be unable to communicate, and totally unable to create highly abstract shared domains like political and legal systems.

Therefore, in order to establish a science of definition, we must firstly establish what and where the common ground is, and alternately, where and how the points of departure from it arise. In this way, we can establish a basis for marking out conceptual landscapes in a systematic way, which will support our endeavour of establishing robust methods for definition.

As a first pass to identifying a shared conceptual landscape at the most primitive level, let us consider what could be common to all peoples of all cultures? In the endeavour of uncovering the nature of the underlying common conceptual landscape, we suggest the following avenues may be fruitful for investigation:

a. Human perceptual and cognitive apparatus :

There is commonality in the way we perceive sensory stimuli, as evidenced by the ability to identify specific deficits such as colour blindness or hearing loss in individuals, and the ability to identify a common underlying topography with which we can explain different mappings of perceptual words: for instance colour terms in different languages. Gardenfors' theory of conceptual spaces (2000) identified the structure of shared spaces underlying perception of colour and taste, for instance.

b. "Basic level" categories:

Rosch's work (1976) showed that there is a level of basic categories of objects determined by overall gestalt perception without distinct feature analysis. These basic level categories are predicated on common methods of sensori-motor handling.

c. Challenges and Experiences of surviving in the physical environment:

Johnson (1987) showed that there are common patterns of organisation arising from the physical environment which he called 'image schema'. These schema, which include frame-like templates for experiential gestalts such as 'container' and 'path', arise directly from physical experience and the need to operate effectively within it, and might be postulated to be common to all peoples.

d. Early stages of Cognitive Development :

Piaget (1972), for instance, was able to identify several common stages of normal cognitive development. Each stage is characterised by the distinctions and realisations the child is able to make: eg differentiating self from objects; realizing things continue to exist when no longer present ('object permanence'), etc. The cognitive stages pass from sensori-motor through to formal operational. The earlier stages, through to concrete operational, which is usually achieved by age 12, may be postulated to be common, whilst the highest stage (formal operational) is characterized by purely abstract thinking and is likely to be divergent.

e. Words common to all languages:

If the same words arise ubiquitously across languages, they would seem to mark out some common conceptual territory. Anna Wierzbicka's work on Semantic Primes (1996) identifies 61 such words including 'I', 'you', 'this' and basic notions of time and space, which she has found to be present in all the languages she has investigated.

f. The "semantic core" of English:

In the 1930's, CK Ogden was able to identify a core vocabulary of 800 words of English by which, he claimed anything of a practical everyday nature could be communicated. Ogden made the reduction by eliminating verbs to basic operations like 'have' 'do' and 'be', and reducing nouns by means of a dimensional analysis he called the 'panoptic method': a word like 'puppy' could be removed as it could be expressed with 'young' on the age dimension in combination with 'dog'. In all, Ogden identified twenty semantic dimensions.

3.4 Higher Order Constructs and Points of Departure

In 'Metaphors We Live By', Lakoff & Johnson provide a convincing argument that higher level understandings and abstractions develop as metaphors from lower ones. For instance, an image schema like a path, learnt in relation to the physical domain, may then applied to an abstract, non-physical domain such as planning to achieve a goal. Investigation of the abstraction process, and the foundations on which higher order mental constructs are built, may assist in identifying the points of divergence in conceptual landscapes.

Some divergences may be due to lack of differentiation, or to underlying value judgements. On the first count, for instance, there is evidence that autistics do not differentiate between animate and inanimate objects, and thus do not develop the folk concepts of belief and desire which are used to understand behaviour, so they never develop this part of the conceptual landscape. With more abstraction there are more possible differentiations, and if people lack the ability or interest to make these, particular abstract domains, like higher mathematics, will be inaccessible. On the second count, some ideologies, value judgements or social beliefs like 'all men are created equal' may be pivotal in constructing the mental landscape, and conceptual landscapes may diverge from these kinds of ontological commitments.

It is worth noting that each individual's understanding of the world is necessarily built on certain assumptions: without them we cannot possibly hope to build many of our most common concepts. We should not apologize for or hide this, but openly acknowledge and identify them. However, we should also attempt to use them sparingly. Particularly when attempting stipulative definition, which involves building and marking out new conceptual territories, we should attempt to use as few ontological commitments as possible: we certainly want to be able to represent mental activities such as beliefs, hypotheses and value judgements at the appropriate level, but should be careful not to build these in at the foundations.

4 Conclusion

In conclusion, let us review whether a science of definition is achievable in light of our discussion. Development of a 'science of definition' would require an understanding of common semantic ground and the ways in which divergences from it are made. We have made some progress in delineating the ways in which the conceptual landscape is naturally structured, and the means by which abstract constructs may overlaid or built within it. Taking a cognitive viewpoint, we have suggested some paths for investigation to determine where common semantic ground may be located and how and where divergences from it arise. We have mentioned some avenues, including empirical studies, which may be of use in inferring the elements and structure of the conceptual landscape. By putting together the various strands, it may well be possible to describe an underlying shared conceptual space which would provide a viable basis for robust methods of definition.

5 Acknowledgements

This work is part of the NICTA-wide priority challenge "From Data to Knowledge". I would like to thank my supervisors for allowing me the freedom to pursue the ideas presented in this paper. I would also like to thank James Nicholson, without whom none of this would be possible, and my brother John Cregan and soon-to-be sister-in-law Nadja for their love and support.

6 References

- Gärdenfors, P. (2000): *Conceptual Spaces: The Geometry of Thought*. Cambridge Massachusetts, MIT Press.
- Johnson, M. (1987): *The Body in the Mind: The Bodily Basis of Meaning, Imagination and Reason*. Chicago, University of Chicago Press.
- Kelly, G.A. (1955). *The Psychology of Personal Constructs*. New York: Norton.
- Lakoff, G., & Johnson, M. (1980): *Metaphors We Live By*. Chicago, University of Chicago Press.
- Ogden, C.K. (1930): *Basic English: A General Introduction with Rules and Grammar*. London: Paul Treber & Co., Ltd.
- Piaget, J (1972): *The Principles of Genetic Epistemology* New York: Basic Books.
- Robinson, R. (1950): *Definition*. London: Oxford University Press.
- Rosch, E., Mervis, C.B., Gray, W.D., Johnson, D.M., & Boyes-Braem, P. (1976): Basic objects in natural categories. *Cognitive Psychology*, **8**: 382–439.
- Thomas, S. (2005) : *The Importance of Being Earnest about Definitions*. BTW 2005, 560-577. Available at <http://www.informatik.uni-trier.de/~ley/db/conf/btw/btw2005.html#Thomas05>
- Wierzbicka, A. (1996): *Semantics: Primes and Universals*. Oxford: Oxford University Press.