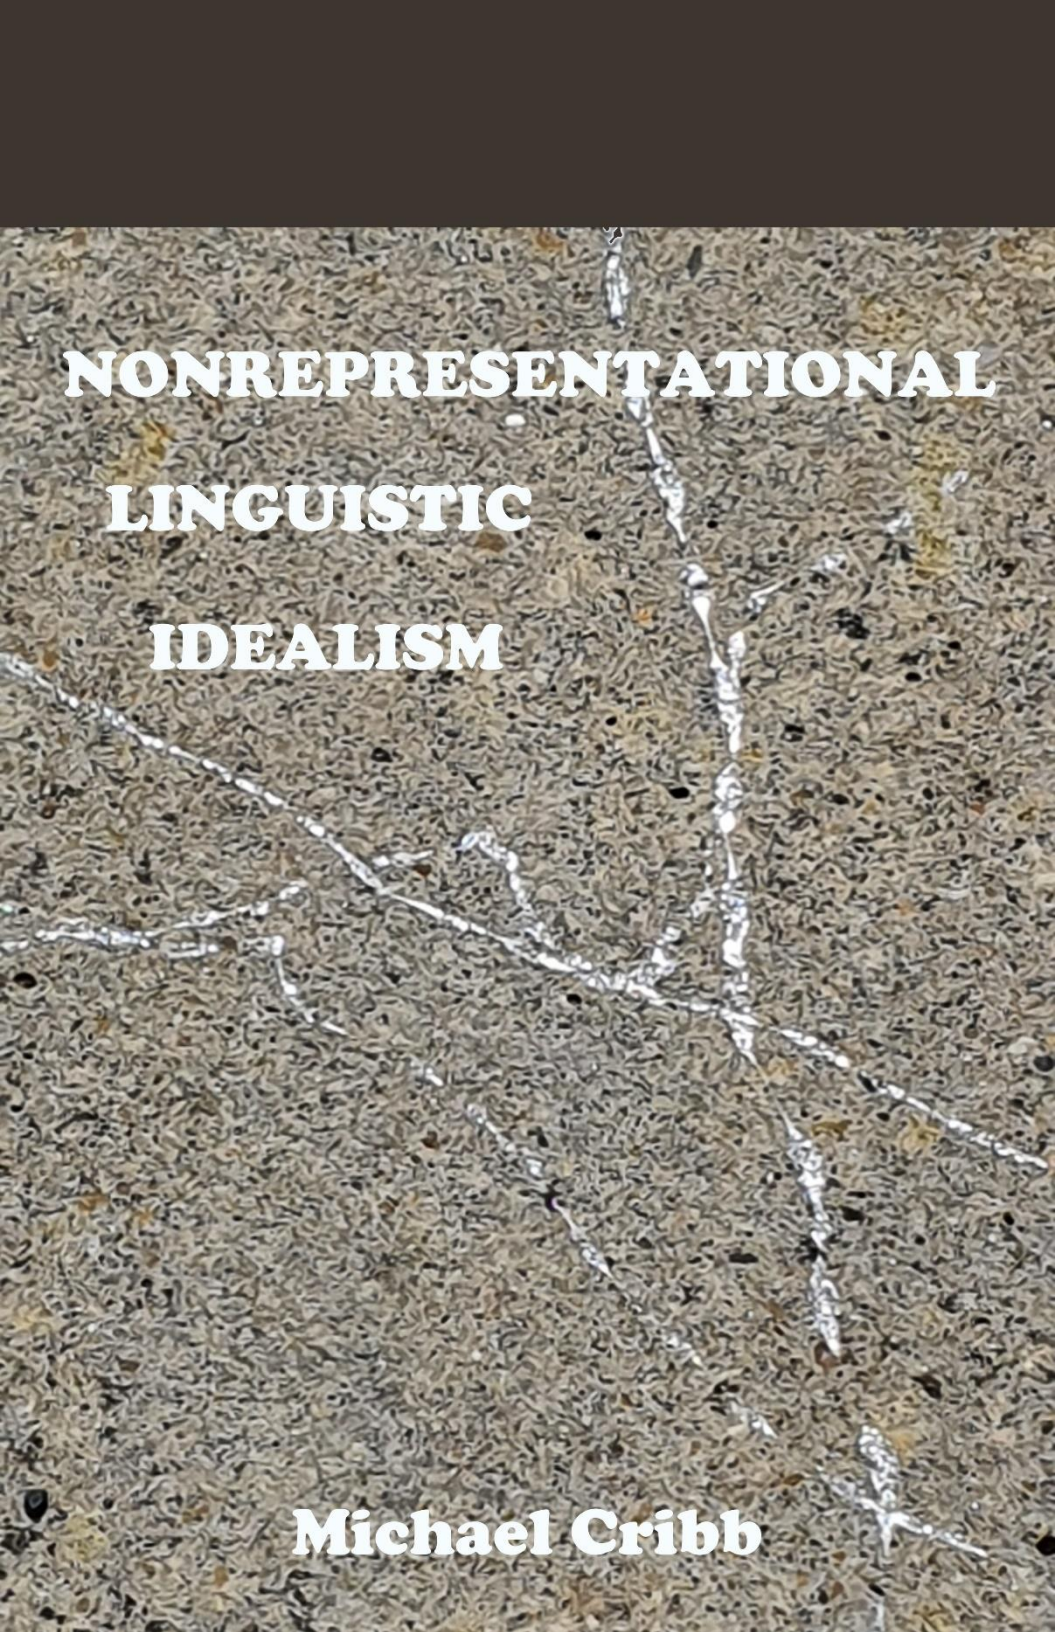
NONREPRESENTATIONAL
LINGUISTIC
IDEALISM
Michael Cribb

Nonrepresentational Linguistic Idealism

Monograph

Michael Cribb

First published 2023

Nonfiction > Philosophy > Language
Nonfiction > Philosophy > Mind & Body
Keywords: idealism, language, mind, consciousness, sentience

https://itslanguagestupid.wordpress.com/nli-monograph/

essential tremor publications

CONTENTS

Preface 1

1. Introduction 2

2. Idealism 5

3. Linguistic Idealism 10

4. Nonrepresentationality 17

5. The Linguistic Paradox 26

6. Phenomenal and Access Consciousness 33

7. Access Consciousness is Linguistic Consciousness 39

8. Thought (part I) 46

9. Thought (part II) 59

10. Sentience 67

Bibliography 73

Notes 77

PREFACE

In 2020, I published my book *It's Language Stupid: unravelling the DNA of the mind*. The book represented many years of researching and teaching on language, and contemplating its relationship with the mind. Since that time I have continued to write, and my thinking has developed on this subject. The monograph here brings together all these ideas into one place and puts them under a new name: Nonrepresentational Linguistic Idealism (NLI).

It has always been known that language plays a central role in human discourse encoding our ideas, values and beliefs in history, law, politics, education and much of day-to-day society. NLI claims that language is also central to our thought processes. So central in fact that we could term it the 'DNA of the mind'. The mind is language and is conscious through language.

The structure of the monograph follows a series of post written at different times during the past two years. After a brief introduction to NLI in chapter 1, I take each of the terms of the title 'nonrepresentational linguistic idealism' (yes, even I struggle to say it sometimes!) starting from the right to form the core of the theory in chapters 2, 3 & 4. A key chapter is then presented in 5 where I introduce the linguistic paradox. In chapters 6 and 7, I look at phenomenal and access consciousness in turn. I consider some thought processes in chapters 8 and 9 before finishing off with a chapter on sentience in artificial intelligence (AI).

The texts remains largely as they were originally posted with only minor changes and corrections. This results in overlap of content in some places, but I felt it was better to leave this content in.

Mike Cribb, Bicester, January 2023.

1. INTRODUCTION

we are destined to rewrite ourselves

Michael Cribb

What is nonrepresentational linguistic idealism (NLI)?

NLI is a philosophical position (metaphysics) that puts the mind at the centre of reality and language at the centre of the mind. It asserts that thought is linguistic in nature and that humans can only ever be conscious of ideas encoded in language. Language does not represent the physical world as is often claimed but is the world itself. Language is the DNA of the human mind.

Does it assume that the physical realm is a figment of our imagination and does not really exist?

No, it does not. It assumes that there is a physical world out there, but this world is indeterminate to us. All we ever consciously know is linguistic in nature.

Which has priority, the linguistic realm or the physical realm?

Neither. NLI does not take a position on which is prior or more 'real'. NLI deals with what we experience as human beings rather than what is.

Who else believes in this stuff?

Emanuel Kant is taken to be the father of idealism. Thomas Hofweber (2018) and Bernardo Kastrup (2019) are modern-day exponents of the belief in idealism. Richard Gaskin (2021) and Christian Barth (2011) are prominent linguistic idealists. I do not know of anyone except myself who is a nonrepresentationalist although there are several proponents of anti-representationalism which is a different position.

If I accept NLI, how will it change my life?

It probably won't change anything. It is just a metaphysical position that you can run through your mind from time to time.

Is NLI a sort of religion?

No, NLI makes no claims about whether there is a God or not.

References

Barth, C. (2011). *Objectivity and the Language-Dependence of Thought*. Routledge.

Gaskin, R. (2021). *Language and the World. A Defence of Linguistic Idealism*. Routledge.

Hofweber, T. (2018). Conceptual Idealism Without Ontological Idealism: Why Idealism Is True After All. in *Idealism: new essays in metaphysics*, T. Goldschmidt and K. Pearce (Eds). Oxford University Press.

Kastrup, B. (2019). *Analytic Idealism: A consciousness-only ontology*. Doctoral Dissertation, Radboud University Nijmegen.

2. IDEALISM

idealism entails that mind is nature's fundamental ontological ground, everything else being reducible to, or grounded in, mind

Bernardo Kastrup

Idealism is the belief (or rather philosophical position) that reality as we know it is mentally constructed. Nothing exists which is independent of the mind, the human mind that is. The mind is given priority over the material world. Philosophers, such as Kant, Berkeley, and Hegel, have emphasized the mind over material over the years to varying degrees.

The position contrasts to some extent with *realism* – the belief that the physical world exists independently of observers. This position holds that there are entities or facts that exist which are in some way independent of our minds. While realism and idealism may seem like opposite sides of the coin the difference really is one of emphasis or viewpoint. Realists put the material world at the centre while idealists put the human mind at the centre when trying to make sense of our existence. It is probably true to say that the dominant school of thought in the sciences and philosophy today is realism although idealism seems to be making a comeback of sorts.

Emmanuel Kant is sometimes said to be the modern-day father of idealism. Writing in the 1700s, Kant's belief was that we have innate structures that link us to rational thought (*a priori*). These innate structures are waiting to be unleased but can only be done so through experience of the world. The *a priori* is not something simply to be discovered through rational introspection but is gained through experience of the material world. Thus while Kant rejected the *tabula rasa* views of Locke and Hume – his predecessors – he believed that it was necessary for the human mind to experience its environment in order to develop a coherent and valid representation of it.

Kant's version of idealism is termed transcendental idealism although there are several other flavours of idealism:

Subjective idealism: human minds directly perceive nothing but themselves and their own ideas. The existence of other minds is inferred from one's own perceptual ideas.

Absolute idealism: everything including one's own mind is a part of a greater mind: the mind of God.

Ontological idealism: the view that we are central for reality understood as what there is. That is, what exists.

Alethic idealism: the view that we are central for reality understood as what is the case. What facts or propositions we know.

My own view of idealism in NLI is a combination of ontological and alethic in some form. The mind (language as I will claim in my next chapter) is central not only for facts (what is the case) but also for what there is. However I am wary of cementing my position to a name because I think it can only really be understood once all the components of NLI have been explicated.

One issue with idealism that is often brought up is that if the material world is somewhat dependent on the mind, what happens when there are no human minds around to observe the material world. For example, there were no human minds around to observe the dinosaurs on earth so does that mean the dinosaurs never existed. I do not think that idealism entails this proposition. Some creatures did roam the earth many millions of years ago but that era is indeterminate to us now. In other words we leave it to the physical – material world to handle that. Whatever creatures there were, and however they lived and moved around, the physical world had no need of the human mind to bring them into existence. Idealism is more about the centrality of the human mind now as we look back in to the past from our vantage point (the vantage point of the human mind – or at least my human mind).

I can demonstrate the two different outlooks with the diagrams below. In the first diagram we have the human eye (the mind) all-pervading and omnipresent looking down on the material world from the birth of the universe until today (and into the future). The diagram suggest that there is an all-present mind to make sense of the world and to ensure that dinosaurs, among other things, can exist. The mind of God or a cosmic mind as some have called it. While this view of idealism might be termed 'absolute idealism' it is not my form of idealism.

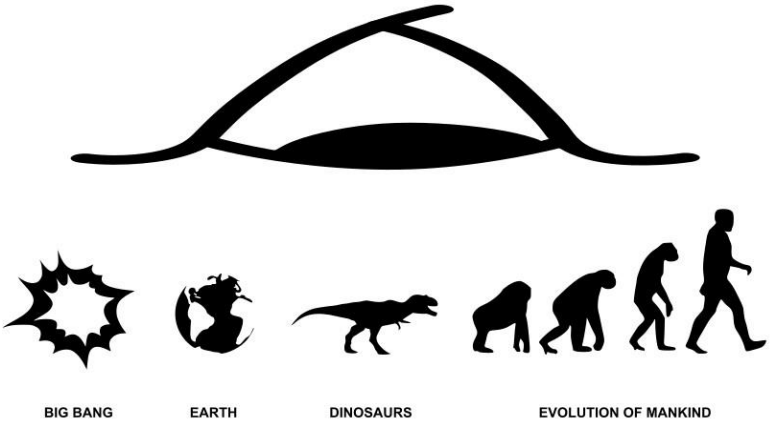

My form of idealism is shown in the second diagram in which modern man and woman with their modern minds (eye) look back on history from their current vantage point. And in this situation dinosaurs certainly did roam the earth many millions of years ago as evidence by the bones and fossils that we find today. But there is no mind's eye at the time of the dinosaur's existence.

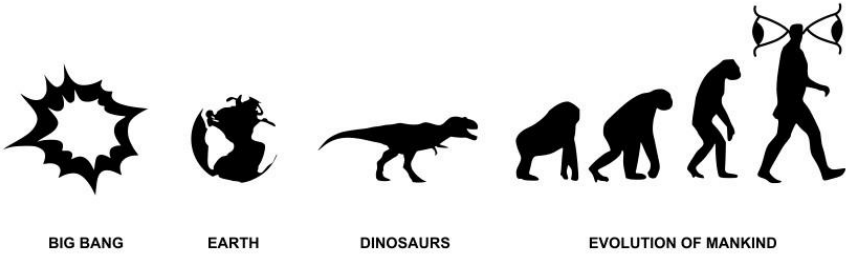

There is a subtle difference between the two diagrams which some readers might not appreciate. In the first diagram there is an attempt to provide a mindful (linguistic) explanation to the material world at all points in history whilst in the second there is a 'now-centric' view of the world. In the second diagram there is an assumption that the physical world got on okay without the human mind in the way that physical worlds do before we got here. Idealism therefore does not reject a physical, material world but assumes that it is indeterminate to the human mind. (An analogy can be made with Schrodinger's cat in which the cat is both dead and alive until we cats our eyes on it. The physical world takes care of the dead and alive paradox perfectly fine until our human mind arrives.)

> Idealism does not reject a physical, material world but assumes that it is indeterminate to the human mind

Idealism is the basis for NLI but it is only with the other two components – the linguistic and the nonrepresentational – that NLI truly comes into its own. I will outline these constructs in the next two chapters.

3. LINGUISTIC IDEALISM

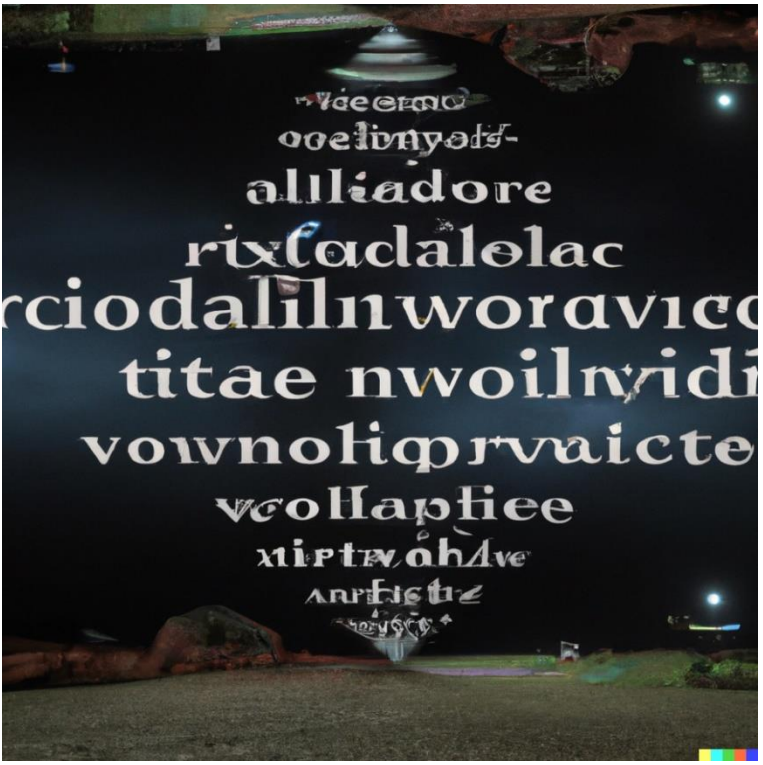

the doctrine of linguistic idealism is that the world is constituted—all the way down—by its expressibility in language

Richard Gaskin

Idealism is the philosophical position that the mind is at the centre of reality for us (humans). 'Linguistic idealism' is quite simply the notion that language is central to the mind. All thoughts that we have in the mind are of a linguistic nature and anything else that is not linguistic is not mindful. The mind is language in effect.

This position may puzzle some since most people report other type of thoughts in the mind beside linguistic thoughts. The mind can hold, say, an image or have a sensation of an itch. Why are these non-linguistic experiences not part of the mind? NLI does not say that we do not experience these types of thought. It is just that these are not part of the mind. They are not mindful thoughts that we can be conscious of in the same way that language is. They are essentially just activities in the brain that are below the workings of the mind. In order for them to be brought into the mind they need to be tokened in language in some form or other. I discuss this further in a later part of the chapter but for now let me just focus on linguistic thoughts.

I assume that all competent adult human beings experience linguistic thoughts from time to time throughout the waking day. (You are experiencing linguistic thoughts now as you read these sentences.) Sometimes these thoughts are very personal and private and we might keep them to ourselves. At other times we express these thoughts in overt spoken (or written) language to others. Most, if not all, of our social structure is built on language. We think in language, converse in language, write down our history and literature in language, discuss politics in language, resolve legal disputes in language, and so on. Without language our social world would be much narrower and more circumscribed.

It is true that the human mind can have other types of 'thought' apart from linguistic ones. Most people report seeing images in the mental space from time to time. Or they might report an awareness of a sensation such as an itch or pain, or be aware of an emotion that we are feeling such as sadness or joy. NLI asserts, however, that all these latter thoughts, or mindful activities as they have been termed, are not 'thoughts' in the way that linguistic thought is. They are just instances of brain activity occurring at a non-conscious level in the brain. This is a key distinction that NLI makes.

Let me give an example to illustrate the difference between linguistic thought and non-linguistic 'thought'. Imagine you are viewing a painting in an art gallery. Your brain is processing a visual image in front of your eyes. You may see, say, a chair and a bed in this painting. You might notice a window and a table and see that the room is somewhat messy but bright. NLI argues that you cannot have these thoughts just by looking at and viewing an image with your eyes. If you do not process these parts of the image (chair, table, bed, etc) as language then you haven't really seen them as such. You have seen something but not a chair, table, etc. You could argue that you have seen items such as legs, backs, frames, tops, bed sheets, etc. But NLI would say that even these are linguistic terms so in reality you haven't seen these items. We might then say that you see shades and colours and edges but even these are linguistic in nature.

> The visual world is indeterminate until language is brought into play.

So what have you seen in the painting in the gallery if you haven't seen a chair and a bed, or legs, tops, or even shades and colours? We cannot say what you have seen. The visual world is

indeterminate until language is brought into play. It is only when (and if) you token the items in the painting with language that you can actually say that your mind is conscious of these features, and it is only until you token the thoughts such as 'the room is messy' that you can say that your mind is aware of this fact. So seeing an image in the mind is not really mindful activity according to NLI. It is just brain activity. Now it may be that off the back of seeing the painting in the gallery you decide to buy a print in the shop or take up painting when you get home. But NLI is not concerned with the outcome of brain activity. It is only concerned with what is mindful to the human – what we can be conscious of – and NLI states that we can only ever be conscious of linguistic thought.

Linguistic idealism

The mind then is at the centre of our human experience and language is at the centre of the mind. Language is the mind in effect and we are language. This does not mean that we can get by without our brains and our bodies. We do need a physical world to support the mind. But it is the mind and language which gives us our conscious experience of life. Without language we would be just another tree in the forest.

Without language we would be just another tree in the forest.

You may feel that I am elevating language to some special status in our universe that is not warranted. NLI however does not really look to say which mode is more 'elevated' or given priority. NLI is only concerned with describing what is and it assumes that language is something separate from the physical world and is constitutive of the human mind. Granted this may seem like a

step too far for you but consider this point: I am arguing the case here, in this monograph, for NLI and you are reading my arguments. All this is taking place in language. It is language and only language which can argue such a case and only language which can ask questions. The painting in the gallery a few paragraphs back cannot ask questions, neither can the itch or pain sensation that you may have felt. It is only language that can hold questions and present them to the mind. Language is the only domain that is a vehicle for these type of thoughts. Without language we would not have the need to ask what is the mind or what is the physical. We would just be.

Counterarguments

One argument against this position might be that humans think in some abstract, non-linguistic form which then manifests itself in language from time to time. However if we look carefully at this argument we can see that this cannot be the case. Let's assume first that this type of thought is holistic in nature. What does 'holistic' mean? It is generally taken to mean the whole rather than the parts. So assume that a thought occurs in the mind that is holistic in nature. It just appears in the mind without any linguistic structure. But if it is holistic in nature we cannot say what that thought is here and now. I cannot write it out. It is just a blob. No more than the dot at the end of this sentence. The only way we can say what the thought is is to convert it into language. But what is the point of proposing a holistic thinking system that can only be understood in language. If we can only even understand a thought when it is linguistic in nature, then a holistic thought is no thought at all. That is not thinking. That is just the brain being active. There is no useful information in a blob for the human mind. The only thoughts that can be relevant to me here and now are linguistic thoughts.

Let's consider another possibility. Assume that humans have a language of thought, a mentalese. This has been proposed by some scholars in the past. But this theory suffers from the same problems of the holistic thinking theory. If I have a thought in mentalese it is no use to me here and now until it is tokened in language. And why propose a mentalese that is similarly structured to language when no one has ever observed this mentalese? Why not just take it that language is the thinking system?

NLI takes these counter-arguments and states that the only thoughts that are relevant to the human mind are those which are couched in language. Everything else is just part of the physical, material world and can never be conscious to us (the human mind). It is only when we token something in language that we find it relevant. This happens in the private mind and also in the public sphere. Language is the domain in which we work. It is the domain in which we ask our questions and try to answer them. It is the domain in which we argue and debate and discuss.

Other Idealists

There are not many philosophers who put language at the centre of the human world. Two living proponents are Richard Gaskin and Christian Barth. Gaskin (2021) states that the 'doctrine of linguistic idealism… is the thesis that the world is a precipitate of language … the whole point of this doctrine is that the world is constituted—all the way down—by its expressibility in language.' Barth approaches the position from a slightly different angle and terms it 'Universal Conceptual Lingualism'.

In the next chapter, I will look at the final part of NLI which is the notion of nonrepresentationalism. This is probably the most

contentious part of the philosophical position and I have not yet found anyone who shares my opinions here.

> The doctrine of linguistic idealism… is the thesis that the world is a precipitate of language. (Gaskin)

References

Barth, C. (2011). *Objectivity and the Language-Dependence of Thought*. Routledge.

Gaskin, R. (2021). *Language and the World. A Defence of Linguistic Idealism*. Routledge.

4. NONREPRESENTATIONALITY

an experience makes its appearance only when it is being said and unless it is said it is, so to speak, non-existent

Hannah Arendt

Idealism is the notion that the mind is at the centre of reality for humans and linguistic idealism is the notion that language is at the centre of the mind. This does not entail that the material world does not exist, but that to understand human existence we need to work from the inside out, that is from the mind out to the external world.

Representation

Representation I admit is a term which I struggled with early on in my career in applied linguistics. The term refers to the (accepted) notion that language is a vehicle for the mind to approximate, or model, the outside world. This seems like a pretty useful vehicle which helps us to make sense of the world in which we live.

If we take a very ordinary, everyday sentence such as 'the cat is on the mat', it is assumed that the language of this sentence in the mind represents, or approximates, to a current state of affairs in the material world. For example, the words in the sentence refer to real-world objects; thus 'cat' represents a cat and 'mat' to a mat. In addition, the sentence as a whole represents a state of affairs, namely that there is a cat on the mat. We can use language to represent future (or possible) worlds as in when I think that there will be a cat on the mat when I get home.

This purported ability of language to represents the physical world seems like quite a useful vehicle for our thoughts in that we can model the real-world in our minds and presumably use this to make decisions, argue positions, and communicate with our fellow human beings about the world.

A closely related notion is that of *intentionality*. Intentionality in philosophy is used in a very special sense which does not carry the day-to-day meaning of 'on purpose'. We say that language

'intends' towards objects and states of affairs in the real world. What this means is that the object of a word or a sentence is something in the outside external world, which is external to our minds. So when we use the word 'cat' it intends toward the real world object cat and the sentence 'the cat is on the mat' intends towards the state of affairs.

Representation and intentionality are very similar constructs and for all intents and purposes we can treat them as the same here. If we were to tease out the difference between the two we might suggest that intention is more of a process whereas representation is more of a product, but I won't make any use of this distinction here.

Nonrepresentation

The basis of NLI, and perhaps its most contentious claim, is that language does not represent the outside world. That is, when we use language to think in our minds or to communicate with other minds, the words and sentences are not actually pointing to the outside world. We might then ask what does language refer to if not objects in the material world. NLI suggest that language does not represent anything. It is a world in itself – the linguistic world. Granted this linguistic world may seem very similar to the external world at times but there are important differences. I will argue this in two ways: the how and the why.

> NLI states that language does not represent anything. It is a world in itself – the linguistic world.

The first argument for nonrepresentation is one of *how*. How does language represent the external world, if that is what we

claim it does? What magical force conveys the representation from the linguistic thought in our minds to the external objects and facts? There does not appear to be any invisible force-carrying particle that conveys representation from the mind to the external world. (Unlike gravity which has invisible force-carrying particles: gravitons.) And if there is, when do these start to act? When do they stop? These are not trivial questions. It seems like most scholars in the philosophical community just accept that there is a link of representation between the linguistic and physical without really questioning how this happens. But this is a mistake I believe. There is no logical explanation for how language represents the external world.

A second argument that questions the notion of representation is the *why*. Why does representation need to occur? What is the purpose of representation for the physical world and the mind? Without a human mind there is no representation. A red traffic light conventionally represents the notion of 'stop'. But it only does so as a human mind in a car approaches. It doesn't mean stop for a cat, or a bird, or a road marking for that matter. It only means stop for human minds and even then it is only conventional. We have agreed as a society to recognise it as such but I can easily override this and drive through a red traffic light if I am so inclined.

But a traffic light is nothing like language. A traffic light has only three colours and a limited number of colour combinations. Language is a propositionally-based thought system that can generate an infinite number of ideas. It is the basis for human minds but no one can say why language represents the physical world.

Everything we think and say in the social world is carried out in language. We write poetry in language, our histories, our laws,

conduct politics and education in language. So what is going on when we do this, when we use language to construct this social world? Are we saying that everything that has ever been thought, said or written is just a representation of the physical world? Is it not more likely that language is a world in itself: a world that is local to itself and largely independent of the physical world? NLI argues for this latter position; that language does not represent the physical world but is in fact a world in itself – a linguistic world.

A more complicated example

Let me give a more complicated example than the simple 'cat on a mat' sentence that I discussed above. The following sentence is taken from a book which I read a few years ago:

> The phenomenologist studies perception, not as a purely subjective phenomenon, but as it is lived through by a perceiver who is in the world, and who is also an embodied agent with motivations and purposes. (Gallagher and Zahavi, 2012)

It is hard to see how we could check to see whether this sentence is true or false in the real world. We would need to test each and every phenomenologist and see how they studied perception. If one single phenomenologist did not study perception as stipulated then the sentence would fall. As the sentences of language that we use become more complex and abstract, it becomes more and more difficult to claim that they are pointing to the real world.

But there is a simple and stronger reason why we can say that the sentence does not represent the physical world, or a possible future world, and that is due to the principle of nonrepresentation.

Principle of nonrepresentation

The principle of nonrepresentation is a stipulation of NLI that says that one dimension cannot be used to represent another dimension. We accept that we live in a world that has physical dimensions (space) and a temporal dimension (time). If I clasp my palms together in front of me I know that they enclose a certain volume of space and within this space is some matter – chiefly air and perhaps a few dust particles. But it would be wrong to say that by clasping my palms I also capture time in them. Time is a different dimension and of completely different stuff. (Quite an extraordinary dimension I must say but I won't get into the nitty gritty of it here.) Each dimension has its own properties and one cannot be used to represent the other. We cannot explain space through time and we cannot explain time through space.

NLI states that language is also a dimension, the linguistic dimension. This dimension is separate to the physical and the temporal and has properties of its own. To say that language resides in its own dimension (or domain) is not some stuff of a weird Science Fiction movie. It is just simply the basic observation that language is of different stuff to the physical and the temporal. It is not matter and it is not time; it is language. To claim it exists in a separate dimension is merely to convey on it something basic and fundamental to it that cannot be explained by space or time. To say that we exist in a three dimension spatial world and a one-dimensional temporal world is a given. NLI adds the linguistic to this.

So NLI is saying that the sentence shown above and the earlier one (the cat is on the mat), and all the sentences and bits of language that have ever been thought and said and written are not representations of the physical, material world but are in fact a world in themselves. The principle of nonrepresentation says

that one dimension cannot be used to represent another and therefore language is not merely representing the physical world. If we try to do so, or think we are doing so, we commit a representation fallacy (Dyke 2008) and misconstrue our understanding of the universe and who we are within it.

There are some odd outcomes of this position which we do not need to get too hung up on here. One outcome is that cats and mats are really products of the linguistic world, not the material world. Something does exist in the physical world but we cannot say it is a cat or a mat because these are linguistic terms. We have to accept the physical dimension for what it is – a physical world filled with matter that is indeterminate to the linguistic world: a bunch of fundamental particles obeying the laws of physics. And that is really the most we can say about it, and even this is pushing things.

> The principle of nonrepresentation states that one dimension cannot be used to represent another.

If you are not yet convinced that language is a separate dimension in a world of its own, consider this: everything you have read in this monograph so far has been written in language. Language is the only medium I can use to retain your attention and to explain the ideas that I have. You remain here on this page because language is engaging you (even though you may not agree with what I am saying). It would seem odd then to suggest that language is not something unique. It would seem odd to suggest that all human thought and speech and written texts that have ever been expressed are simply particles of dust like the dust particles in the space between my palms. Surely

language has to be something special in the seat of human consciousness and human existence?

non- or anti- representationalism?

There is a widely established position of anti-representationalism that differs significantly from my account. Anti-representationalism holds that perception is not a process of constructing internal representations but is in fact an active and dynamic process between the agent and the environment. Most anti-representational accounts do not take language to be a separate dimension which is why I prefer the term 'non-' as opposed to 'anti-' representationalism.

For some people this claim may be too strong but let's look at what it is not saying. NLI does not claim that there is no physical world. It is just that this world is indeterminate to us. We can probe and push this world all we like, but we can only ever be conscious of our own existence through language. We are language in effect.

Let's look at this claim that 'we are language' in more detail because some of you may think that I am emphasising language too much over the physical world and time.

> "We are language"

There is a bit of language in this sentence that seems to claim that language is in some sense superior or more important to us than the space-time which we live in. This is not the case however. The linguistic part is supported by the physical. The physical world provides the light and the pixels and the screen and the PC and all the other material bits that allow this statement to exist, and this includes our bodies and our brains. And time provides the temporal sequencing that allows our eyes to move from the left

to the right of the sentence. Without space and time language would be nothing. It is just the case that space and time cannot express themselves in the way that the linguistic can. They operate in the background, off-stage so to speak. The principle of nonrepresentationalism entails that the physical can only express itself in a physical way, through matter. It cannot be represented by language. And time can only represent itself in the way that the dimension is constituted, not by the physical or the linguistic.

Imagine you have three children which you love dearly and equally except that only one child can speak, only one child can see and only one child can organise events. You interact with each child in the only way they can and accept them as they are. You talk to one, you see the other and you hold events with the third. It is only together as a team that they can get out of the house and live their lives. And that is who we are. So, to say 'we are language' is to celebrate all aspects of our lives.

To say 'we are language' is to celebrate all aspects of our lives.

Dyke, H. (2008). *Metaphysics and the Representational Fallacy*. Routledge.

Gallagher, S. & Zahavi, D. (2012). *The Phenomenological Mind*, (2nd ed.). Routledge.

5. THE LINGUISTIC PARADOX

the human — and no other — possesses the one essential tool which makes a social construction of reality possible. That tool is language.

George Grace

NLI (Nonrepresentational Linguistic Idealism) is the position that the mind is the centre of reality for humans and language is at the centre of the mind. Language is the mind and is nonrepresentational. It does not represent the physical world but creates a world in itself. We experience our existence in a poly-dimensional universe that is spatial, temporal and linguistic where each dimension expresses itself, so to speak, in the only way that it can. The spatial expresses itself through bottom-up, distributed matter, the temporal through time and the linguistic through top-down, propositional language.

There is however a paradox in what is written above in that I am attempting to describe the universe (in particular the physical and temporal dimensions) through language and I have already said that language cannot be used to represent these domains. This is the principle of nonrepresentationalism: the principle that one dimension cannot be utilised to approximate to, or represent, another. Thus while we like to think that we can talk about the physical world through language, we in fact commit a representational fallacy every time we try to do so.

Courts of Law

So that I can build the argument for why this is the case, consider the following argument. Imagine a man has been charged with a crime. The judge demands to see the suspect in court. A representative, such as his lawyer or brother, will not do. The judge wants to see the suspect in person. Similarly during the trial the suspect needs to be present to answer questions in person. Simply submitting written answers is no substitute.

Why does a judge in a court of law insist on bringing the suspect into court to hear the trial? Why not just accept a representative? It is because our system of justice demands that the person in

question answers to the charges directly. Something so profound is happening through the trial process that the person charged with the crime, and only that person, needs to be present. And any subsequent sentence and punishment needs to be heard and served by the criminal. A lawyer cannot simply accept the judgment on his behalf and neither can a family member serve time in prison for him.

Similarly with language, if you owe me three gold coins and say to me 'here are three gold coins' I cannot accept that the words are the gold coins. It is not until you place the three gold coins in my palms that I can accept you have paid me. In other words, the language is not the physical entities.

NLI rests on this principle. The argument for nonrepresentationalism in NLI goes something like this. If we say that language is a representation of the physical world, then we imply that it is not literally the physical world. If it is not literally the physical world then it is not a representation of it. The steps in this argument are shown in the diagram below.

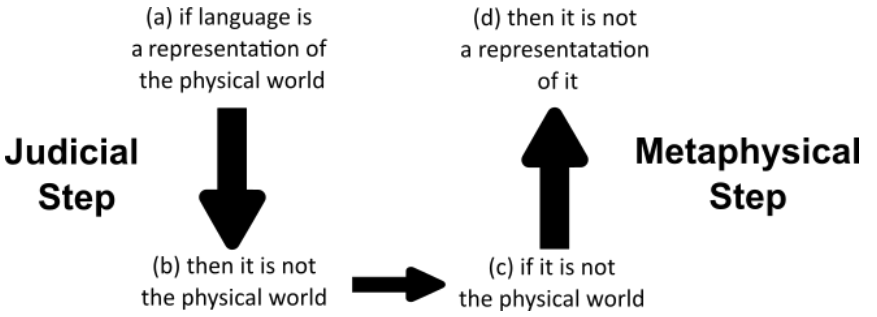

Step (a) to (b) I term the judicial step because it is the argument a judge uses to demand that a suspect, and no other, appears in court. We often fail to recognise, however, that there is a return step from (c) to (d) which I call the metaphysical step. In this step, we must assume that if language is not actually the physical

world then it is not a representation of it. When using language we mostly assume the judicial step but fail to see the metaphysical fallacy. This is the linguistic paradox: we desire to talk about the physical world but language is not a means for representing the physical world.

> *The linguistic paradox*: we desire to talk about the physical world but language is not a means for representing the physical world.

The linguistic paradox

[1] Consider this question: if a tree falls in the forest and there is no one around to hear it, does it make a sound? This well-known thought experiment asks us to question the nature of sound when there is no human mind around to perceive it. Can it be classed as sound if there is no human mind to hear it?

From an NLI perspective, the answer would be to say that all we have presented from [1] above is language (do you see how I placed the number surreptitiously into the text?). There are no trees, forest or sound, only words and sentences. The question exists only in the linguistic domain, and not the physical.

[2] But you might then argue that we could conceivably meet up one day and travel to a forest to find a tree that is just about to fall. We could then retreat to a suitable distance, wait for the tree to fall and return to consider the question of whether the tree made a sound or not. In other words, we could set up an actual experiment in the physical domain to determine whether the question is true or not.

But I would argue that from [2] onwards, all that we have done is add more language to our minds. The language in [2] is no more closer to the physical world than the language in [1]. In fact all of the language that I have written in this monograph is merely language of the mind. It is not the real, physical world.

As you can probably see, we could go on debating whether language represents the world for ever and each time you argue something at [3] and [4] and so on, I would be able to claim that it is merely language. It seems then that language is just a way to entertain thought experiments. Everything that has ever been thought, written or said, and everything that ever will be thought, written or said are in fact just thought experiments in our minds. We seem to be no closer to the physical world than when we first started speaking.

Everything that has ever been thought, written or said, and everything that ever will be thought, written or said are in fact just thought experiments in our minds.

Schrödinger's

[3] Let me give you one more example. Schrödinger's cat is a thought experiment which relies on the indeterminate nature of the quantum world. In this paradox, a cat is in a room. A flask of poisonous gas in a canister is poised to be released and kill the cat if and when a single atom decays. However since quantum mechanics, under one interpretation, says that the atom can be decayed and not decayed at the same time, the result seems to

be that the cat is both dead and alive at the same time. It is only till we look in the room to see whether the cat is either dead or alive that the paradox is resolved.

NLI has an answer to this paradox. It says that Schrödinger's cat is a thought experiment in the linguistic domain only. The physical world has no problem with quantum superposition and ambiguity. It gets on by itself doing what it does best – which is to be physical. It does not need language to make it work. The paradox, if there is one, is purely in the language that has been used at [3] to create it. All we have at [3] is more language, not the physical world. There are no cats, atoms, poison or anything physical at [3].

Now, there may be some utility in making claims in language about the physical world. It might be useful for example for scientists at CERN to know that an 'atom consists of electrons orbiting a nucleus'. In making such a statement, they bypass the metaphysical step as we all do when using language on a day-to-day basis.

But NLI is not interested in the utility of language. It is a philosophical position that seeks to understand what actually is. It desires to understand what are the tracks of the mind, here and now, and what are the relations between these tracks. It makes no claims as to whether language is useful for the human condition or not. The 'tracks' that I am talking about are the physical, the temporal, and the linguistic dimension – the three fundamental dimension that make up our existence and our experience of life.

Summary

The linguistic paradox is when we use language we desire to represent the physical world, but the physical world cannot be

represented by language. No matter how much we talk about the physical world we are no closer to it through language.

6. PHENOMENAL AND ACCESS CONSCIOUSNESS

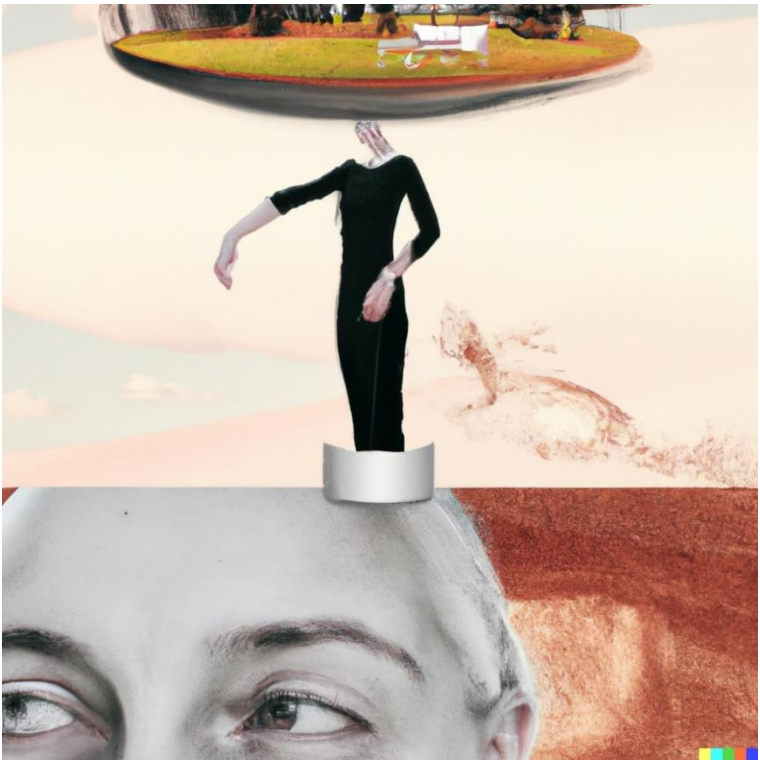

what makes a state phenomenally conscious is that there is something "it is like" to be in that state.

Ned Block

Block (1995) distinguishes between phenomenal consciousness (PC) and access consciousness (AC). PC is the subjective experience that the human mind holds as we go about our lives. We have experiences such as the smell of coffee or the red colour of a rose. We can say that there is a certain 'feel' to these experiences or the what-it-is-like. AC by contrast is the conscious awareness we experience as we report on our phenomenal experiences and generate rational thought on the back of them. AC is also said to be a conscious state. It is sometimes said that there is an 'overflow' in PC. That is we experience more of the world in PC than we can actually report in AC.

For the metaphysical position NLI (nonrepresentational linguistic idealism) there are a number of issues regarding PC and AC which I'd like to clear up here.

Phenomenal Consciousness (PC)

For me, PC is the *ineffable hum* of the physical. If I experience the smell of coffee, say, then before (and if) I access this smell in AC, I experience it phenomenally and in totality in PC. But this experience is subjective. I cannot say what this experience is like because it is ineffable. As soon as my mind attempts to report on it, the experience is lost in the words. As soon as AC is engaged to report on what I am experiencing, the phenomenality of the experience collapses. This is what I term the *linguistic paradox*. (An analogy can be made with the collapse of the wave-function at the quantum level when an observer measures the position of a fundamental particle.) Once I access the smell of coffee and report on it, say—'that smell is aromatic'—the ineffable hum of the smell has been lost. The phenomenality has collapsed.

It is important to clarify what NLI is saying here since this is a critical juncture for the theory. I can continue to smell the coffee

as I report on the experience. That is, my nose and olfactory neurons can continue to process the smell. What collapses, however, is the phenomenality of the experience as it is reported on in AC. In other words, the linguistic contents that appear in AC are not the phenomenal experience. In going from PC to AC the phenomenality collapses.

> The Linguistic Paradox:
> the use of language to report on and rationalise a phenomenal experience necessarily collapses the phenomenality of the experience

Consider Block's example of becoming conscious (AC) of a pneumatic drill sometime after being aware of it in PC (see quote below). Block has to choose his words carefully here because he is trying to denote two types of consciousness. AC is the rational description of the drill as we become consciously aware of it whereas PC is awareness of the noise of the drill.

> … suppose you are engaged in intense conversation when suddenly at noon you realize that right outside your window there is- and has been for some time – a deafening pneumatic drill digging up the street. You were **aware** of the noise all along, but only at noon are you **consciously aware** of it. (Block 1995: 234, emphasis added)

I do not see how we can be 'aware' of something at time 1 and then 'consciously aware' of it at time 2 and still maintain that both states are conscious. If anything, the state at time 1 is a sub-conscious awareness. Block tries to justify the choice of terms a little later:

> Note that this case involves a natural use of "conscious" and "aware" for A-consciousness and P-consciousness, respectively. "Conscious" and "aware" are more or less synonymous, so calling the initial P-consciousness "awareness" makes it natural to call the later P-consciousness plus A-consciousness "conscious awareness." (Block 1995: 234)

PC for me is not a conscious state; it is just the brain being aware of something. It is just the physical being physical. My neural network that I call my brain is processing sensory information, taking in the sound waves of the pneumatic drill and being aware of them but without conscious access to them. There is no consciousness involved until AC rationalises the sound through language. What is left of the PC when we rationalise the sound vanishes, but we are left with a sense that we did experience something. PC lets us know that there is something to experience but that is all it can do. It cannot yield its contents.

Access Consciousness (AC)

First off, I should say that I do not like the term 'reporting' to AC despite its widespread use in the literature. I do not think AC 'reports' on the PC per se. AC is not obligated to name the entity that is in PC. The AC has free will to choose how it reacts to PC. For example, if I see a cat in the street, my AC is not bound to report this as 'cat'. I might report it as 'animal' or 'pet' or perhaps 'dog' if I am mistaken. Or I may have a rational thought such as 'I need to get a pet' or 'bad luck if I cross the line'.

I would suggest that a better way to look at the relationship between PC and AC is one of 'motivation'. PC motivates AC to some extent but does not restrict it. AC is free to access whatever content it wants from the phenomenal experience. The term 'cat' might be the most likely on probability terms but cannot be

guaranteed. There may also be some backward propagation from AC to PC. AC may motivate PC and cause the mind to shift focus onto something else in the local environment.

My major claim here is that AC is linguistic and can only be linguistic. This claim is based on two arguments: (i) language brings an end to the infinite regress of the physical and (ii) language is all we have to report on PC consciously. Let me explain the infinite regress argument first. If we say that AC is simply a global workspace (i.e. a theatre) in the brain where cognitive processes meet then this workspace is physical and the information is distributed in the workspace. We are no closer, however, to having conscious awareness of this physical information than when it was distributed throughout the brain. Simply having a workspace to concentrate the physical does not make the information not distributed. How then can we be aware of distributed physical information? How can the mind get inside of what is essentially a bunch of neurons firing? Only language can bring this argument to a close.

> Language brings an end to the infinite regress of the physical.

The second argument for AC being linguistic and only linguistic is that there is no other entity in the universe that could be a candidate for conscious awareness (i.e. conscious thought). A universal language of thought (Fodor 1975) has been proposed but I think this just delays the inevitable. A language of thought would need to be isomorphic with language and why propose a universal substrate to natural languages that we have no evidence for whatsoever.

Another candidate for conscious thought that has been proposed is a holistic representation. Perhaps the mind thinks

holistically? However on closer inspection this argument appears to be weak. If I say that thought is holistic then this means that thought is rolled up into one without internal structure, as a 'blob'. So we cannot be consciously aware of anything inside this blob-like structure. There is no useful information for the conscious human mind and the only way out of this is to convert the information into language.

This is not to say that the contents of the pre-linguistic global workspace and the physically distributed information in the brain cannot be useful for decision making and direction. We most likely do 'think' and make decisions subconsciously. But we are not interested here in the subconscious workings of the mind. We are only interested in what becomes available for conscious awareness. Language is the only entity that can achieve this. Language puts a cap so to speak on the information and provides a top-down view of the contents so that we can see them from the 'inside', i.e. be conscious of the contents.

AC then necessarily needs to be linguistic in nature if it is to present information to us that we can consciously hold. We can thus say that the only consciousness that we have is linguistic consciousness. Access consciousness is essentially linguistic consciousness.

In the next chapter, I show how linguistic consciousness works and defines us as a human species.

References

Block, N. (1995). On a Confusion about a Function of Consciousness. *Behavioral and Brain Sciences*. 18, 227-287.

Fodor, J. A. (1975). *The Language of Thought*. Thomas Y. Crowell.

7. ACCESS CONSCIOUSNESS IS LINGUISTIC CONSCIOUSNESS

some representational contents are poised for direct use in reasoning, speech, rational action, and subjective reports

Matthias Michel

In the previous chapter, I looked at phenomenal and access consciousness. I claimed that phenomenal consciousness (PC) is the *ineffable hum* of the physical and cannot be reported on. If I experience the smell of coffee, say, phenomenally, then the experience collapses as soon as I try to report on this through access consciousness (AC). This is the *linguistic paradox*: the use of language to report on a phenomenal experience necessarily collapses the phenomenality of the experience.

> The Linguistic Paradox:
> the use of language to report on and rationalise a phenomenal experience necessarily collapses the phenomenality of the experience

I also claimed that access consciousness has to be linguistic in nature. There is no other candidate for consciousness to bring an end to the infinite regress of the physical. All conscious thought in the human mind, then, is carried out in linguistic form (i.e. language). Everything else, such as seeing images or processing smells, is merely the brain in action in PC. In this chapter, I will outline an argument for nonrepresentationality, the notion that language does not merely represent the physical world but is the world in itself. This world is our consciousness: linguistic consciousness.

Essentially what NLI does is break the link between language and the physical world. Traditionally, language has been assumed to represent the physical world. If I think or say 'I smell coffee' then it is assumed that this linguistic statement in some way

represents the world as it is when the statement is made: namely that there is a smell of coffee in the air and I am smelling it. NLI claims that this is a representational fallacy. The representation between the linguistic statement and the physical world does not hold; it is nonrepresentational. NLI does not deny that something is happening in the physical, but the smell of coffee and me smelling it are all within the linguistic domain.

In order to demonstrate this nonrepresentational relationship, let me present a two-step argument, the conclusion of which is that language *does not* represent the physical world.

> 1. If language is a representation of the physical world then it is not the physical world.
>
> 2. If language is not the physical world then it is not a representation of the physical world.

Step 1 is termed the judicial step since it is similar to the arguments a judge makes when ordering a suspect to appear in court. At some point, the judge wants to see the actual suspect in person in court, not a representative such as a solicitor or family member. NLI makes the same demands because it is an ontological philosophy that is concerned with what is. We understand that the language that we use on a day-to-day basis is not the actual physical entities that we describe. If, for example, you order coffee from a coffee shop and the barista serves you with the word 'coffee' you will be sorely disappointed.

Step 2 is termed the metaphysical step. This step is more difficult to comprehend and is often ignored in philosophical circles. The step asks us to consider what we mean by the term 'representation'. The Oxford English Dictionary defines it as: 'the action of standing for, or in the place of, a person, group, or thing, and related senses' (OED). Let's analyse this sense to see how and whether representation holds between language and

the physical world. We can do this by asking how, when, and why representation holds.

If language represents the physical world then *how* does it do this? Is there some sort of force carrier between the word 'coffee' and the entity it purports to represent (analogous to say the graviton that mediates gravitational attraction between two bodies). Similarly, *when* does this representation hold? Does it hold at all times or only when the word is uttered. How long after it is uttered does the representation dissolve? Finally, why is it necessary to postulate such a relation? *Why* does language need to represent the physical world at all?

Brentano's through his intentionality thesis (Jacquette 2004) also questioned whether the mind can actually represent objects in the material world. If a glass of wine is sitting in front of me, say, and I *desire* that glass of wine, what is the content of my desire? Brentano would say that the wine is contained within the mindful act by virtue of its 'in-existence'. NLI says that the *desire* is purely PC until we report on it through language in AC. The image and smell of the wine are merely a phenomenal experience that has no consciousness until I report on the desire such as 'I want that wine'. Then the phenomenality of the experience collapses in the linguistic thought. There is no madness here. Language is simply creating our world.

There is nothing inherent in the universe that says that the language that flows through our minds or comes out of our mouths (or written by hand) has to be tied to the physical world. The only reason we postulate such a relation is because it is conventional. Conventional in the sense that we have unquestioningly accepted Brentano's wisdom and that that language is a series of conventional signs and grammar rules that as a society of human minds we agree represents the physical

world. But if it is conventional then it is easy enough to defeat that conventionality and simply state that language does not represent the world. On what grounds can you disagree? In fact, without the human mind agreeing to such a relation it is difficult to see how anything can represent anything else. A red traffic light only represents stop because a society of human minds has deemed it to be so. There is nothing inherent in the colour of red that causes bodies to stop.

Nonrepresentationality

Nonrepresentationality is at the heart of NLI. The basic postulate is that one domain cannot be used to represent another domain. So the linguistic domain cannot be used rightfully to represent the physical. To try and do so is to commit a representational fallacy. And vice versa, the physical domain cannot be used to represent the linguistic domain. Nonrepresentationality is the outcome of the linguistic paradox introduced earlier. As we report on the PC through the AC, the phenomenality of the experience is lost; it collapses. The resulting linguistic thought in AC cannot then be said to represent whatever the experience in PC is.

> The PC is ineffable: hidden noumena perhaps or a mind-independent aspect of reality that is inaccessible to human perception.

The PC is ineffable: hidden noumena perhaps as Kant refers to it or a mind-independent aspect of reality that is inaccessible to human perception (Seth 2021). If God needs a personal

workspace to present the world to us, then so be it. We shouldn't assume that we can peer into it. We have the linguistic to work with. So, in effect consciousness is linguistic in nature and only linguistic.

If NLI says that language does not represent the physical world, then what is the relationship between the two (if there is any at all)? A better way to conceive of the relationship is through the concept of 'motivation'. Motivation suggests a looser relationship between the physical and the linguistic than representation and also suggests a two-way relationship. The physical world *motivates* language and language in turn *motivates* the physical world in part. The physical world has motivated language to turn out the way it has. Throughout the linguistic history of mankind, our interaction with the physical world has shaped language into what it is today with certain entities and grammatical structures to show the relations between these entities. Similarly language in some way shapes and motivates the physical world in how we divide up reality and how we conceive of it. Motivation, unlike the term representation, suggests a relationship but does not specify exactly what that relationship is.

The physical world *motivates* language and language in turn *motivates* the physical world.

References

Jacquette, D. (2004). *The Cambridge Companion to Brentano*. Cambridge University Press.

Seth, A. (2021). *Being You: A New Science of Consciousness*. Faber & Faber.

8. THOUGHT (PART I)

the decisive factor in the invention of human language was the advent of conscious thought

Filippo-Enrico Cardini

The word ‘thought’ is used regularly in discussing human cognition but is rarely defined. We are all familiar with statements such as ‘I just had a thought’ or ‘what are you thinking’. We seem to generally accept that people do have thoughts from time to time but we rarely question how these take place. Most people just assume that thoughts magically pop into our heads and either cause us to act in some way or invite us to describe them to others through language.

In the 1980s, Russell T Hurlburt carried out an ingenious experiment to try and determine the ways that the human mind thinks and experiences life on a day-to-day basis. He placed buzzers on participants timed to go off at random intervals during the day. When a buzzer sounded, participants were asked to freeze the contents of the mind and write down what they were experiencing in their mind just before the buzzer went off (i.e. what thought or experience were they having, if any at all). Later they were interviewed and asked to recall these thoughts. From the results of this experiment, Hurlburt proposed that the human mind carried five different modes of experience. These were: inner speech, inner seeing, unsymbolized thinking, feelings and sensory awareness.

1. Inner speech – speaking to yourself inside your head. Sometimes termed ‘inner voice’.

2. Inner seeing – seeing an image or something in your imagination that is not actually present

3. Unsymbolized thinking – Thinking a particular, definite thought without the awareness of that thought’s being conveyed in words, images, or any other symbols

4. Feeling – Affective experiences, such as sadness, happiness, humour, anxiety, joy, fear, nervousness, anger, embarrassment, etc.

> 5. Sensory awareness – Paying attention to a particular sensory aspect of the environment where that sensory experience is itself a primary theme or focus apart from the object of perception
>
> (Heavey & Hurlburt 2008)

Nonrepresentational Linguistic Idealism (NLI) states that only the first experience, inner speech, is human thought. All other aspects of experience are not really aspects of human thought but instead just brain activity. This may come as a surprise by many since we tend to assume that thinking (or thought) is a multimodal process in the brain. In this chapter, I will argue that only language can be the vehicle of our thoughts. I will take each of the other four aspects in turn to show how this is so.

2. Inner seeing

Inner seeing is the sensation we have when we see an image 'inside our head' that is not present in the visual field. For example, we might close our eyes an imagine the face of a family member or see a painting that we saw in an art gallery earlier. This seems very intuitive and most people report seeing images from time to time in their thoughts. (Although recently I have questioned whether I myself actually do see anything at all, even when I try hard to direct my imagination on something.)

But is seeing an image actually a thought in the mind? Let's say someone asks you 'what are you thinking' and you say 'I'm seeing an image of a red aircraft taking off from an airfield'. Your description of the image is in fact a linguistic statement. In order to convey the image to anyone who asks, we need to encode it in language. You may argue that before communicating the contents of the image via language, the image itself is a thought, 'inner seeing' as Hurlburt terms it. Your mind is attending to this

image and perceiving a 'red aircraft' and an 'airfield' and a 'taking off'. But again, these are linguistic terms. You could take it one level further and say the mind is perceiving the shapes of the aircraft such as a wing, fuselage and fin, and the colour red, together with the features of the airfield such as runway and grass. But at any level we need to use language to describe what we are seeing.

So seeing an image in the mind cannot be described in linguistic terms, else it is not inner seeing. We do not see words or sentences in our mind; we see a red aircraft taking off from an airfield. But what actually are we seeing then if we cannot describe it in language? I could draw a picture here of the red aircraft taking off or I could just write:

*&!k)%!^@

In other words, the inner seeing cannot be described in language else it ceases to be inner seeing. The mind is perceiving an image (of the red aircraft) and it may very well act upon seeing this. You may reach for your red pen and draw the image, or go out and buy a model aircraft or do any number of things off the back of this thought. NLI does not deny that the image is in the mind and can have causal effect. But what NLI does say is that inner seeing is not linguistic in nature, is not conscious thought and is not thought at all. It is merely brain activity. Only language in the mind (i.e. inner speech in Hurlburt's terms) is thought, and only language in the mind is conscious thought.

> Only language in the mind is thought, and only language in the mind is conscious thought.

Now at this point you may just argue that it is only a difference of terminology. I have said that only language is thought and only language can be conscious thought. Everything else such as seeing images in the mind is sub-conscious thought or brain activity. You could argue that seeing an image in the mind is a conscious experience too if we just designate it as that. But what NLI is arguing is that it is only through language that humans can be conscious of their world. It is only when we think in language, communicate in language or write in language (as I am doing here) that we have conscious thoughts. Seeing an image is part of the workings of the mind, but there is a qualitative difference between seeing an image and processing language. An image is just a visual experience. There is no linear propositionality to the experience, unlike language. When we process language, we experience something like in (a) below. When we see an image, we experience (b).

> (a) I'm seeing an image of a red aircraft taking off from an airfield'
>
> (b) *&!k)%!^@

The language in (a) gives humans an insight into the world. Language lights up the world and illuminates it for the human mind from the inside. It opens the world up and stretches it out along the linguistic dimension. Images do not do any of that. It is sometimes said that a picture paints a thousand words. It doesn't. It doesn't paint anything at all. This is not to say that images are not an important part of human experience. Human visual experience is incredibly important for our existence and our lives. But it is not thought in the sense that the language of this monograph is thought. I cannot write this monograph in visual imagery and expect you to get anything from it. There is only one vehicle of thought for humans and that vehicle is language.

> Language lights up the world and illuminates it for the human mind from the inside.

3. Unsymbolized thinking

Hurlburt's third type of mindful experience is termed 'unsymbolized thinking'. Here is a reminder of its definition:

> Unsymbolized thinking – Thinking a particular, definite thought without the awareness of that thought's being conveyed in words, images, or any other symbols[1]

Hurlburt gives an example of this in his 2008 paper as reported by a participant in a follow-up interview:

> Adam was watching two men carry a load of bricks in a construction site. He was wondering whether the men would drop the bricks. This wondering did not involve any symbols, but it was an explicit cognitive process[1].

The implication here is that thinking is taking place but the participant is not experiencing this in language, images or any other symbolic form. Most people would report that they have had times when they have thought a certain idea but are unable to report on how that idea was formed in the mind. There is no clear language or images that appear in the mind to carry this thought.

NLI however states that unsymbolized thinking is sub-conscious brain activity and not thought. Consider this argument. The description given by Adam of watching two men carrying bricks (above) is written in language, and was reported as such in the post-experiment interview. There is no other way to convey this

idea but through language. But let's take it that Adam did experience unsymbolised thought at the time of the experience. How was this thought structured? We might suggest that it was holistic in nature. It came into the mind in one go and appeared to the thinker (Adam) as a whole, gestalt. But if it was holistic then it had no internal structure to Adam. How could Adam become consciously aware of the contents of the thought if it had no structure? I suggest that it is only when the holistic thought is converted into linguistic form that it is available for conscious introspection. Language 'strings out' holistic ideas in a propositional dimension. It opens them us and gives us a conscious view of the inside. Until then, we must assume that the holistic thought is sub-conscious brain activity.

Another way to argue for unsymbolized thinking is to suggest that the human mind has some, as yet undiscovered, language of thought (LOT) or mentalese. This LOT in some respects resembles language and is the basis of all propositional thinking. But why propose such a language when we already have natural languages which we have direct experience of. The LOT would need to be isomorphic with natural human languages. Isomorphic here means that the internal construction of the LOT needs to match the human language it is reported in. So in fact we gain nothing by proposing a LOT.

Language 'strings out' holistic ideas in a propositional dimension. It opens them us and gives us a conscious view of the inside.

In the end we are left with a similar conclusion to our treatment of inner seeing. The experience of unsymbolized thinking is a

genuine experience but it occurs at a level below consciousness and we can only become mindful of the 'thought' by converting it into a linguistic form.

I will take up the discussion of the remaining two of Hurlburt's mindful experiences, feelings and sensory awareness, in chapter 9. These two experiential modes are not considered as thoughtful modes in the same sense that the first three experiences are. For now, I wish to summarise my position and present an argument for language (inner speech) being the only candidate for mindful thinking.

Traditional view of thought v. NLI view

The two diagrams above summarise the difference between the traditional (conventionally accepted) view of thought and the view that I am presenting here as NLI. On the left, we have

thought as a general property of the mind with inner speech being just one of several ways of thinking. On the right, the NLI view shows that inner speech is a qualitatively different type of thought to all the other modes which NLI terms 'brain activity' (or cognitive activity). Now you could term this brain activity as 'thought' but I prefer to refer to it as just brain activity for the simple reason that it is not conscious activity. The only way that sub-conscious brain activity can become conscious is if we access it in inner speech (=language). The important point to take away from these diagrams is not whether we term the non-linguistic modes of thinking as 'brain activity' or 'thought', but that we recognise linguistic thinking (inner speech) as being qualitatively different to all other modes of experience. All the other modes are wholly within the physical domain. Linguistic thought is within the linguistic domain.

Thought Experiments

If you are not convinced that language is the only true conscious thought of the mind, let me take you through one more argument, which I think is the killer argument. You have probably heard of *thought experiments* that philosophers like to propose, such as Schrödinger's cat or the tree in a forest, where a hypothetical situation is presented for the purpose of thinking through to some sort of consequence. Thought experiments of course need to be presented in language. It would be very difficult for me to lay out an image in front of you and expect you to see a thought experiment. (Or to let the wind blow on your face.) How you think through that thought experiment depends to some extent on your typical mode of thinking. You might think it through in images, or unsymbolised thought or maybe language. I expect that most minds will combine modes to think

the experiment through to some sort of conclusion. However, if you want to convey your conclusion to me you will need to output the results in language.

Let's go through a simple thought experiment just to show you how they work. *Imagine a tree standing in a forest. The tree falls. Does the tree make a sound if no-one is around to hear it*? That is the thought experiment in italics which of course is presented in language. Think this through for a few seconds and come to some conclusion (even if you are not sure what the answer is). When you have a conclusion, just run that through your mind in your inner voice.

Now let me move the argument on. The classical thought experiments such as the falling tree above are well known. But in essence, the monograph I have written here, from the first sentence to this sentence, is a type of thought experiment. I have been writing sentences which you have been reading, even if you might not have started at the beginning. You cannot deny now, as you read what I am writing, that I am asking you to participate in a thought experiment. This thought experiment has been asking you to consider the brain and the mind, and what thought is, and whether language is a unique type of thought, etc. You may have thought about my ideas for a time and no doubt could report on your conclusions to me if given the opportunity. In the same way that we presented a short thought experiment in the previous paragraph with the falling tree, this monograph is also a thought experiment.

In fact it seems that all of language is really just a thought experiment. An email, a text, a conversation is in some way really just language being put into someone else's mind and asking them to consider the ideas in that piece of language and come to some conclusion or consequence (which may be verbalised or

not). But if this monograph is just a thought experiment, then what has happened to the real world objects that I have being talking about such as the brain, the mind, inner speech, buzzers, Dr Hurlburt, etc. They are out there somewhere in the real world, but they are not in this thought experiment. This thought experiment is just happening in mine and your mind. It is just language in the mind.

(ii) Now you might come back here and say that we could actually go out into the real world and conduct some of the thought experiments in practice. You and I both know that that is possible. We could drive into a forest, say, find a tree that is close to falling, retreat a suitable distance and wait for the tree to fall and then return to consider the question of whether the tree made a noise. This is perfectly possible to undertake as a real-world experiment.

But in fact the language at (ii) in the last paragraph is in fact just another thought experiment. It is just a few sentences that I wrote for you to put into your mind to consider for a brief time and come to some conclusion. There is in fact no forest or tree or falling in this thought experiment. It is all in the mind.

> There is no forest or tree or falling in this thought experiment. It is all in the mind.

(iii) The only real physical world that is present to me at the moment is the space I am physically inhabiting at the moment. Which happens to be, as I write, my office upstairs at home with my table and computer and various items lying around. You too have a physical world wherever you are reading this monograph.

Take a few seconds to look around you and note the physical world you are in.

But the paradox returns again because it seems that the language in (iii) is not actually the physical world surrounding myself and yourself, but is in fact just another thought experiment. Another bit of language that I put into your mind for you to consider for a few seconds and come to some conclusion. The paradox is that it seems that any time I type a sentence that purports to point to the real world, it doesn't. It is just a figment of our minds. No matter how hard I try to describe the physical world, as soon as my fingers come down on the keys on the keyboard of the computer, more language is generated and we extend the thought experiment.

That is why I say that language is the only conscious thought that we can experience. In order to say that language is the only conscious thought, I have to phrase it in the linguistic domain. Language is the only system that we know of that can encode such a thought. But what has happened to all the inner seeing and unsymbolised thinking and feelings and sensory awareness of Hurlburt's experiment I hear you say. Well they are still with us to an extent but they cannot be in the linguistic domain. The only way you can experience an inner seeing is to experience it in its own domain, which of course is the visual domain. Because of this, I will never be able to capture it in the linguistic domain, in language. And the only way to experience sensory awareness is to experience it in its own domain, which is a sensory domain. We cannot capture it in language. Go ahead for a minute and experience an inner seeing by imagining a picture in your mind or feel a sensory awareness on your body. You will not be able to talk about it or write it down. The best you might say is:

*&!k)%!^@

We cannot talk about the other modes of thinking because they are sub-conscious modes that are not available to consciousness, unless we report on them in language. We know we experience them – we have direct knowledge of them – and that others do because of Hurlburt real-world experiment. But he never saw them himself. He only ever inferred them through the linguistic reports from his patients. And we know they exist because we too experience them ourselves. But we can never consciously think of them. The only conscious thought available to the human mind is linguistic thought, i.e. language.

The only conscious thought available to the human mind is linguistic thought, i.e. language.

References

Heavey, C.L. & Hurlburt, R.T. (2008). The Phenomena of Inner Experience. *Consciousness and Cognition*, 17(3):798-810.

9. THOUGHT (PART II)

infinities can only be found in the abstract thought systems of the human mind

James Glattfelder

In chapter 8, I outlined how NLI (nonrepresentational linguistic idealism) states that only linguistic thought is actually thought. All other modes of 'thinking' such as inner seeing and unsymbolized thinking are subconscious cognitive activities which we can never be conscious of. Only linguistic thought provides the human mind with a conscious experience. In this chapter, I will deal with the final two modes of experience from the Hurlburt experiments[1] and show how these are subconscious activities too.

The word 'thought' is used regularly when discussing human cognition but is rarely defined. We are all familiar with statements such as 'I just had a thought' or 'what are you thinking?'. We seem to generally accept that people do have thoughts from time to time but we rarely question how these take place. Most people just assume that thoughts magically pop into our heads and either cause us to act in some way or invite us to describe them to others through language.

In the 1980s, Russell T Hurlburt carried out an ingenious experiment the results of which suggested five basic modes of inner experience in the human mind: inner speech, inner seeing, unsymbolized thinking, feelings and sensory awareness. Here is a brief reminded of each mode.

1. Inner speech – speaking to yourself inside your head. Sometimes termed 'inner voice'.

2. Inner seeing – seeing an image or something in your imagination that is not actually present

3. Unsymbolized thinking – Thinking a particular, definite thought without the awareness of that thought's being conveyed in words, images, or any other symbols

> 4. Feeling – Affective experiences, such as sadness, happiness, humour, anxiety, joy, fear, nervousness, anger, embarrassment, etc.
>
> 5. Sensory awareness – Paying attention to a particular sensory aspect of the environment where that sensory experience is itself a primary theme or focus apart from the object of perception
>
> (Heavey & Hurlburt 2008)

4. Feeling

Mode 4 is an affective experience – the experience of feeling sad, happy, anxious, etc. Imagine a buzzer going off now. How would you describe your feelings just before the buzzer went off? This is what Hurlburt instructed his participants to do in his experiment. We can all report on the feeling our body is experiencing now, even if it is a very neutral, non-feeling that we have. But is feeling a conscious activity? Without language, I would suggest it is not. First, here is an example report from one of the participants in the experiment:

> Courtney was unequivocally angry, although it was difficult for her to describe how this anger presented itself to her. It seemed to be conveyed by or accompanied by a tight feeling in her chest and a little shakiness in her hands, but she could not be definite about those aspects[1].

Some researchers have suggested that we continuously experience a stream of affect. In other words, we are always in some particular state of mood which we are consciously aware of all the time. Others have suggested that only particular moods rise into consciousness at certain times perhaps due to the rapidity of onset and/or extent of the mood. When the feeling 'weakens or stabilizes it recedes into the background'[2]. The results of Hurlburt's experiment seem to suggest the latter

circumstance since feelings were only reported approximately 25% of the time, although this is still a significant amount. It is of course possible that participants tended to over emphasise feelings since in the absence of any explicit thought it is very easy to 'dip into' the body and report on how we are feeling.

Considering Courtney's reporting of anger in the above sample, we should first note that 'anger' is a linguistic term. What does the term 'anger' mean? The OED defines it thus:

> A strong feeling of displeasure, dissatisfaction, or annoyance, generally combined with antagonism or hostility towards a particular cause or object; the state of experiencing such feelings; wrath, rage, fury.

But this just replaces one lexical item with others. We cannot really suggest the human body needs these words in order to feel angry? Maybe the mind needs to access them though? Can we feel anger without knowing these words, or without bringing them to mind? Most people would suggest we can. But what does anger look like in this case without language? The best we can perhaps say about it is that it is a feeling of:

$%!)@?&

In other words, we cannot put it into words. It is simply the phenomenal aspect of the human nervous system sending signals that correlate with anger to the brain. It is true that the brain can focus on these signals and determine the intensity and extent of the affect. Perhaps it can even determine specific bodily reactions such as a tightening of the chest or shaking of hands as Courtney reported. But NLI does not classify this focus as a conscious focus. There is no language to bring this feeling alive to the mind. All we have is a physical process. Unless, and until, we unpack this feeling with language it remains subconscious.

> Unless, and until, we unpack feelings with language they remain subconscious.

5. Sensory awareness

The final mode of experience is sensory awareness according to Hurlburt. Sensory awareness is not merely just being aware of one's surroundings through our senses but paying particular attention to some aspect of this. We are all aware of our surroundings pretty much most of the time while awake, but at certain times we pay particular attention to aspects of this. Hurlburt suggests it not the act of seeing something as a means to an end (e.g. to reach out and grasp it) but seeing something as 'the primary focus of attention at the moment'[1]. Consider for example a cup of coffee on a table. On the surface of the coffee we notice a particular patterning and shading of the froth. This is the sensory awareness rather than the coffee cup and table per se.

We can construct a similar argument for sensory awareness as we did for feelings. The awareness of 'froth with a particular patterning and shading' requires language to be encoded consciously. Without language we simply are left with a physical process of light from the froth entering our eyes and being processed by the brain. The distributed nature of this information within the neural network of the brain means that it can never be resolved in a conscious manner within the physical domain. Language is the only entity that we know which provides a top-down focusing of physical processes and brings an end to the infinite regress of the physical (i.e. the reductive nature of the

physical). But language must necessarily be outside of the physical domain, for if it were to be contained wholly within the physical domain, we would be back to the same paradox of trying to resolve it. Language cannot merely be a bunch of neurons firing in a particular pattern else we would not be able to become conscious of this. Language must necessarily be something outside of the physical. (A bunch of neurons do fire of course from the third-person point of view as a neurologist might determine. But from the first-person perspective we do not see these neurons. We only experience the language.)

The physical and the linguistic

It is for this reason that I sometimes say that language is the fifth dimension in the fabric of our universe. Something that is not physical (and not temporal) but which interacts in some, as yet unspecified, way with the physical. By language I mean not just the words and the grammar of the sentences that shape language, although these are important, but also the semantic message that these create and resolve in the human mind. The human mind is crucial here because without it we cannot have language. Language in a dictionary, or book, or website is not language until it is processed by the human mind. And once the human mind has processed it, the dictionary or book or website return to their non-linguistic, physical states. The mind however remains active, and ready to act.

> Language is the fifth dimension in the fabric of our universe.

Language then is the homunculus of the mind: the workspace, the viewing gallery. It resolves bottom-up, distributed physical processes into top-down conscious experiences that make sense

to us and which we can handle and manipulate. This probably happens many thousands of times per day almost instantaneously which is why we are not readily aware of it. We spend time experiencing the world subconsciously at the physical level and then intermittently bringing these experiences into conscious thought linguistically. An inner seeing experience of a red aircraft is a subconscious activity until the mind brings it into consciousness through language. An unsymbolized thought of two men carrying bricks is a subconscious physical activity until the mind brings it into consciousness through language. A feeling of anger is subconscious until we name it. A sensory experience of froth on top of coffee is subconscious until we encode it in language.

> Language is the homunculus of the mind: the workspace, the viewing gallery.

Without language I believe we would survive as a human race but our society would be narrower. There would be no red aircrafts or brickies. There would be no anger. And there would be no coffee with froth on the top of it. I think I could survive without these. How about you?

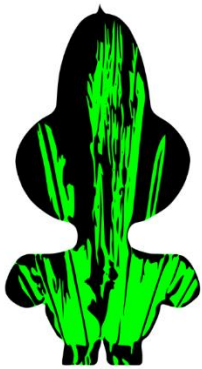

The homunculus

10. SENTIENCE

the possibility that we are actually automata enjoying the unanticipated consequence of personal sentience as our nervous systems go about their business cannot be ignored

William Uttal

The notion of sentience has been in the news recently since a Google AI researcher, Blake Lemoine, claimed that a chatbot he had been interacting with was sentient. His claim has been rebuffed by some antagonists and he has been put on paid leave by Google. In this blog, I will try and outline what NLI (nonrepresentational linguistic idealism) has to say with regard to sentience in AIs. The outcome might surprise you.

NLI states that the mind is at the centre of human existence and that language is at the centre of the mind. Language is nonrepresentational in the sense that it does not represent the physical world out there but is in fact a world in itself. We are essentially language. NLI takes an idealist position in that it assumes we start from the premise that our own 'thinking' minds are at the centre of reality and all attempts to come to terms with the physical world need to be viewed through our minds. This does not mean that the physical world is a figment of our imagination, but that it can only be accessed and understood through a subjective, first-person point of view. In particular NLI takes the view that there is only one subjective viewpoint at any one time, which is *my* viewpoint at the time of writing. You, the reader, can come along for the ride if you want to, but I'm afraid this story is about *me*.

> I infer that you are sentient because I can have a conversation with you in language.

NLI also assumes that language is the seat of consciousness. All sentience that we have in our minds is carried through language. Without language an entity cannot be sentient. I infer that you, and other human beings that I interact with, are sentient because

I can have conversations with you in language. How long I need to have such conversations and what topics we need to get through in order to determine this are not the subject of this chapter. But what if I have a conversation with a chatbot, such as Google's LaMDA, as Lemoine did? If I detect sentience in this chatbot through these conversations, does this mean the chatbot is sentient? Lemoine claims it does.

My answer to this is a subtle but rather neat one I think. Consider this argument: If we say there are *n* human brains in the world, then we might suppose that there are *n* subjective states, assuming that each brain is a conscious, sentient subject. But that would be wrong. There are *n* brains in the world but only one subjective state, namely the subjective state that I am viewing the world from at this moment in time. NLI remember is a first-person account of existence and there is only one first-person viewpoint in the room at the moment, and that viewpoint is mine.

Now this may seem rather odd because in some way it seems to break the laws of mathematics: *n* brains equals *n* sentient states, surely? How can I claim that *n*=1? But this is the linguistic paradox that I have outlined in chapter 5. As soon as I write about a state such as *sentience* using language, I objectify that state. A phenomenal state is brought into access consciousness. But this acts to break the link between the phenomenal state and the language. The state is not the language and is not represented by the language. Language is nonrepresentational and, as a separate domain, cannot be used to represent a physical state. It is true that the objectification of phenomenal experience through language gives us a means to count entities: so we can have one, two even *n sentiences* in the linguistic domain. But there is only ever one phenomenal sentience, and that is the sentience I have at the moment.

So how do I transfer my first-person, sentient viewpoint to you, the reader, so that you can have a turn with the conch, so to speak. I cannot. The phenomenal sentience and first-person viewpoint will always be with me (until I die I suspect). I can infer sentience in other humans through the language they use with me. At some time in the future, I may also infer sentience in chatbots such Google's LaMDA through their language. In fact, I have had many interesting conversations with AI chatbots over the last few years and it is likely that these conversations will only become more human-like and more sophisticated as time goes by and computer algorithms become more and more powerful. But that does not mean the chatbots *have* or *will have* sentience because…

> Sentience is a first-person, phenomenal awareness of language in *my* mind that is non-transferrable.

Sentience is a first-person, phenomenal awareness of language in *my* mind that is non-transferrable. In other words, the question of whether a chatbot *has* sentience or not is moot. *To have* is not the right verb here. A chatbot does not *have* sentience. I only ever infer sentience through the linguistic interaction that I have with the chatbot. I have the sentience and I make the call as to whether my chatbot is sentient or not. Or, perhaps more paradoxically, I am the sentience that the chatbot purports to have.

> I am the sentience that the chatbot purports to have.

So what are the consequences for AI given the argument I have made above? Why do we find it so problematic to utter the words 'the machine is sentient'? Clearly assigning sentience, which is only mine to assign, to machines could change the status of these machines. Might I be less likely to turn them off or delete them if I believed they had sentience? Could machines demand rights if I accepted that they had sentience? These are important questions and perhaps ones that might need to be seriously addressed in the not too distant future.

Creator

But there is another aspect of machines that I think is important in the quest for sentience, and this is the question of who is their creator. Most people, I think, accept that we do not really know who our maker is or why we are here on planet earth. Some may claim a god or mother nature as their creator, but I think it is generally accepted that we all have doubts at times. This is not the case for inanimate objects such as houses or cars where we know that mankind is the creator.

With chatbots today we know they have demonstrable creators. I can create a chatbot with just a few lines of code today if I so choose (the tools exist). Google created their chatbot LaMDA and has the right to turn it off and delete the chatbot if it so desires. It seems then that knowing who or what created a chatbot holds some importance in determining whether it has sentience and what status we ascribe it.

But what if chatbots started to appear and we could not determine where they had come from? What if we discovered a twitter account that had been tweeting sensible, logical language for the past five years but had no obvious creator? Or a band that had been writing and publishing songs that we had all

been singing along to that turned out to be a bot? Perhaps the bots had been created by earlier bots who in turn had come from bots. The sentience that I ascribed to these bots might have more legitimacy simply because I do not know their creators. It seems then that if we know who created a bot, the sentience is in part ascribed to the creator. When the creator is unknown we tend to ascribe the sentience to the bot itself.

> Sentience in part depends on our understanding of whether the entity in question has an identifiable creator or not.

BIBLIOGRAPHY

Baker, L.R. (2013). *Naturalism and the First-Person Perspective*. Oxford University Press.

Barth, C. (2011). *Objectivity and the Language-Dependence of Thought*. Routledge.

Block, N. (1995). On a Confusion about a Function of Consciousness. *Behavioral and Brain Sciences*. 18, 227-287.

Bogen J.E. (1995). An Example of Access-consciousness without Phenomenal Consciousness? *Behavioral and Brain Sciences*, 1997, Vol.20 (1), p.144-144.

Brock, S. & Mares, E. (2006). *Realism and Anti-Realism*. Routledge.

Cardini, F. (2023). *Consciousness and the Cultural Invention of Language*. Routledge.

Carruthers, P. (1996). *Language, Thought and Consciousness*. Cambridge University Press.

Carruthers, P. (2019). *Human and Animal Minds: The consciousness questions laid to rest*. Oxford University Press.

Chalmers, D. (2006). *Consciousness and Its Place in Nature: Does physicalism entail panpsychism*? Imprint Academic.

Clark, T.W. (2019). Locating Consciousness: Why experience can't be objectified. *Journal of Consciousness Studies*, 26(11-12), pp.60-85.

Crane, T. (2016). *The Mechanical Mind: A philosophical introduction to minds, machines and mental representation*. Routledge.

Dennett, D. (1996). *Kinds of Minds: Towards an understanding of consciousness*. Weidenfeld & Nicolson.

Dyke, H. (2018). *Metaphysics and the Representational Fallacy*. Routledge

Gabriel, M. (2015). *Why the World Does Not Exist*. Polity.

Gallagher, S. & Zahavi, D. (2012). *The Phenomenological Mind* (2nd ed.). Routledge.

Gaskin, R. (2021). *Language and the World. A defence of linguistic idealism*. Routledge.

Gennaro, R. (2017). *Consciousness*. Routledge.

Glattfelder, J. (2019). *Information—Consciousness—Reality: How a new understanding of the universe can help answer age-old questions of existence*. SpringerOpen.

Goff, P. (2017). *Consciousness and Fundamental Reality*. Oxford University Press.

Goldschmidt, T. & Pearce, K.L. (2018). *Idealism: New Essays in Metaphysics*. Oxford University Press.

Grace, G.W. (1987). *The Linguistic Construction of Reality*. Croom Helm.

Frankish, K. (2020). *The Demystification of Consciousness*. iainews, Issue 86.

Hacking, I. (1998). *Rewriting the Soul: Multiple personality and the sciences of memory*. Princeton University Press.

Haikonen, P. (2012). *Consciousness and Robot Sentience* (Series on machine consciousness; v. 2). Singapore: World Scientific.

Hofweber, T. (2018). Conceptual Idealism Without Ontological Idealism: Why idealism is true after all. In *Idealism: New essays in metaphysics*, T. Goldschmidt and K. Pearce (Eds). Oxford University Press.

Hurlburt, R.T. (1990). *Sampling Normal and Schizophrenic Inner Experience*. Plenum Press

Kastrup, B. (2019). *Analytic Idealism: A consciousness-only ontology*. Doctoral Dissertation, Radboud University Nijmegen.

Matthias Michel (2020). Consciousness Science Underdetermined: A short history of endless debates. *Ergo*, Volume 6, No. 28.

Mondal, P. (2017). *Natural Language and Possible Minds: How language uncovers the cognitive*. Brill Publishers.

Noë, A. (2006). *Action in Perception*. MIT Press.

Ogas, O. & Gaddam, S. (2022). *Journey of the Mind: How thinking emerged from chaos*. WW Norton & Co.

Robbins, M. (2018). *Consciousness, Language, and Self: Psychoanalytic, linguistic, and anthropological explorations of the dual nature of mind*. Routledge.

Rorty, R. (1989). *Contingency, Irony, and Solidarity*. Cambridge University Press.

Seager, W. (2016). *Theories of Consciousness: An introduction and assessment* (Second ed.). Routledge.

Searle, J. (2007). *The Construction of Social Reality*. Penguin.

Seth, A. (2021). *Being You: A New Science of Consciousness*. Faber & Faber.

Tye, M. (2009). *Consciousness Revisited: Materialism without phenomenal concepts*. MIT Press.

Wolf, M.P. (2008). Language, Mind and World: Can't we all just get along? *Metaphilosophy*, 39(3), 363-380.

NOTES

To cite this monograph:

Cribb, V.M. (2023). *Nonrepresentational Linguistic Idealism*. essential tremor monograph publications.

Author links:

http://cribb.coventry.domains/

https://itslanguagestupid.wordpress.com/

All chapter header images generated by OpenAI DALL-E.

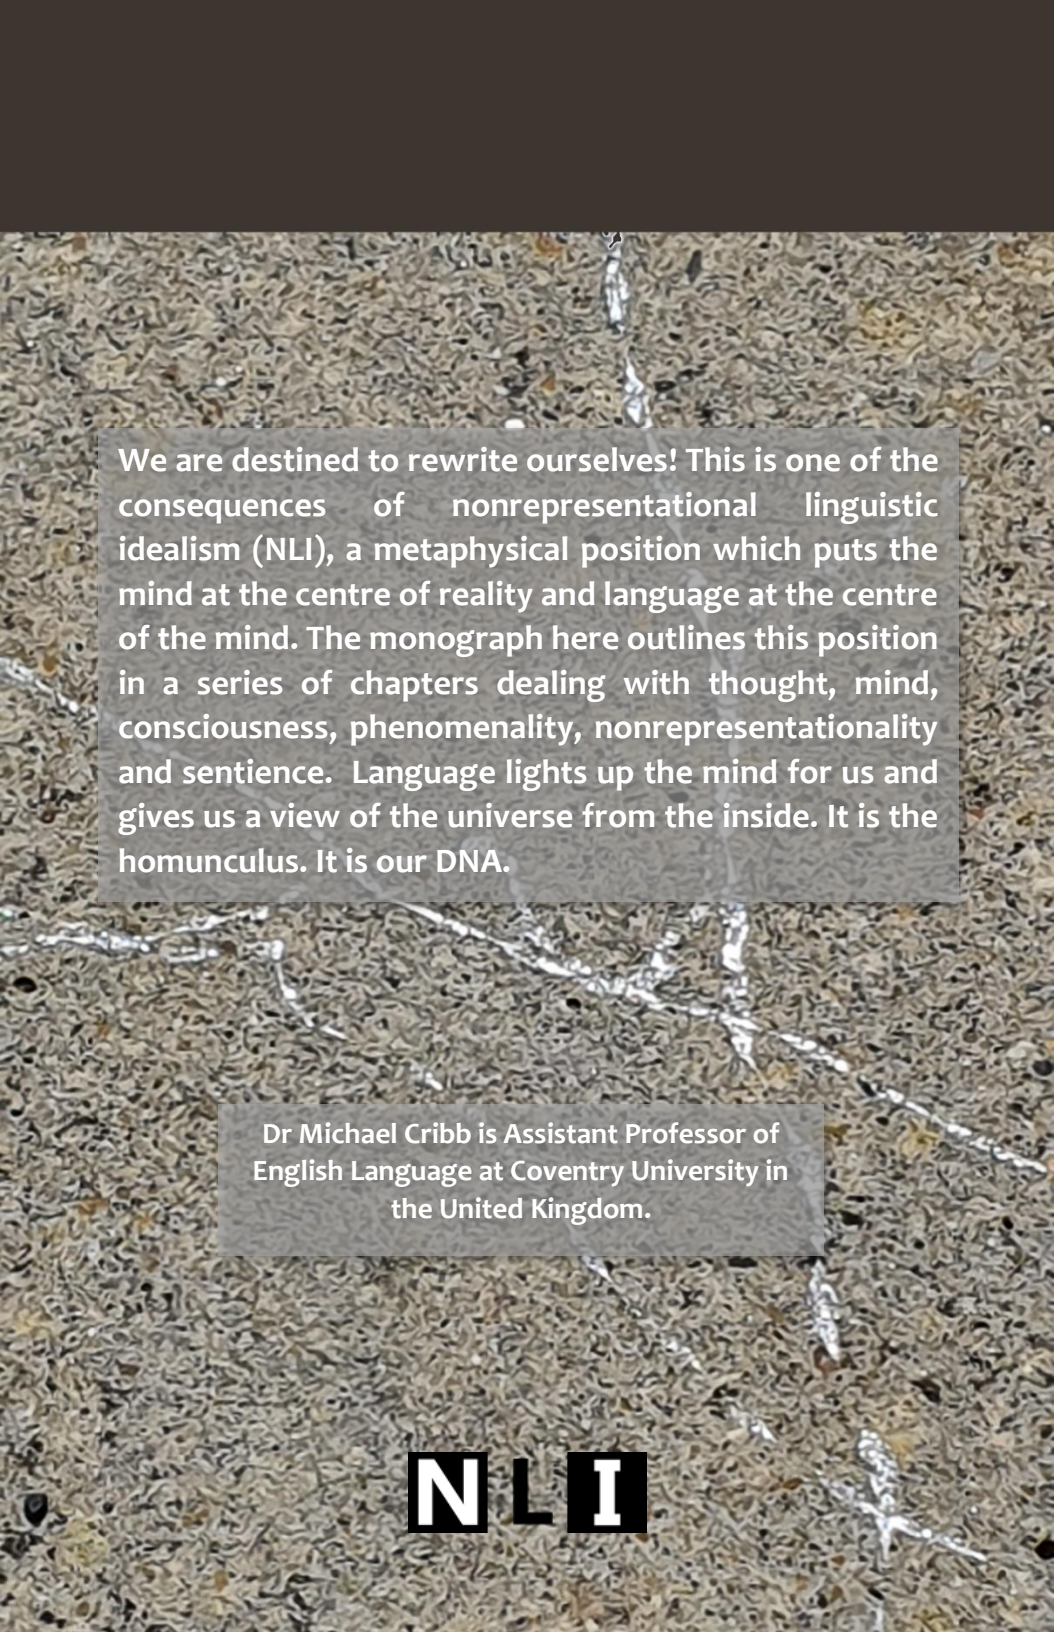

We are destined to rewrite ourselves! This is one of the consequences of nonrepresentational linguistic idealism (NLI), a metaphysical position which puts the mind at the centre of reality and language at the centre of the mind. The monograph here outlines this position in a series of chapters dealing with thought, mind, consciousness, phenomenality, nonrepresentationality and sentience. Language lights up the mind for us and gives us a view of the universe from the inside. It is the homunculus. It is our DNA.

Dr Michael Cribb is Assistant Professor of English Language at Coventry University in the United Kingdom.

NLI