

# AI and the Footnotes to Semantic Coherence

By Daniel. A. Drabkin

To *Lulu*, a friend

\*

## Preface

It was Ludwig Wittgenstein who famously observed that language “deceives” us into thinking<sup>1</sup> that there is a puzzle in the essence or nature of things which goes beyond our straightforward perception of our existential environment and the insightfulness attainable through empirical – generally falling under the term ‘scientific’ – investigation;<sup>2</sup> that we are misled to think there’s something ‘out there’ to understand beyond the coherent structure in which reality<sup>3</sup> presents itself in our first-hand experience, by the very way we speak.

Language has always been at the crossroads of meaning and substance, of all the “stuff” that comprises the world we live in – overt and latent alike – and the perspectival sense that allows any relation thereto or interaction therewith in the first place.<sup>4</sup> With the

---

<sup>1</sup> ‘*Feeling*’ is the more accurate term, however inappropriate, or even jarring, such implication right at the outset of a philosophical paper may appear to be, if only to the standard, academic philosophical ear. In point of fact, when stripped of academic jargon (or better still – in their genuine emergence/manifestation), philosophical puzzles resonate with our existential core (grab us in our gut), rather than merely engaging our intellect. In contrast, it is sheer intellect (the human mind) which is haunted and consumed by a logical or mathematical problem/challenge (un proved theorem or providing any idea or phenomenon with mathematical form), and the same works for physics and any other scientific challenge, whether in the ways it presents itself or the idea of it or any combination or interface therebetween. Philosophy, alas, lays bare an existential – not an intellectual – wound within us. We indeed try to *think* it through, but the latter rather resonates with a prudent attempt to address or resolve a crisis than with solving a riddle or figuring out a puzzle.

<sup>2</sup> Inspired by natural curiosity and equally natural propensity for rational analysis and evaluation. The latter differ from the sense of intellectual “unease” which inspires philosophical reflection. It is the difference between wondering: ‘what is that?’ (and all its essentially synonymous variants such as ‘how does this happen and/or made possible?’ etc.) and being affected with a sense of ‘something here doesn’t sit well..’, almost like ‘something is wrong, what am I missing?’. The latter is nothing like the root or inspiration of scientific puzzle/ inquiry. It might be the case whereby one evaluates a scientific (or any other) theory or hypothesis but not the reflective-inquisitive mode that awakes and commands the very inquiry/exploration science is all about. It rather resembles the case where someone wants to tell me something (e.g., its affinity with *assessment*). But who? Or maybe: what? ‘**What** is that state of unease (the philosophical puzzle)?’ – a scientific question. ‘**Who** is trying to **tell** me anything?’ or ‘**What** is the **message**?’ – an error in language application. Details below.

<sup>3</sup> Or any other, essentially metaphorical, name we may give to the notion of ‘existential domain’, e.g., ‘the world around us’, ‘existential environment’ etc. The dual resonance of that notion as both an experience and an object in itself is at the background of this paper. We believe that the analysis which follows will also cast genuine light thereupon.

<sup>4</sup> Between being as such – in itself, “from within”, and perceived to be what or as it is. One could rightly wonder to which of the two categories (if they are categories at all) ‘overt’, ‘covert’, ‘latent’ or, for that matter, *something/anything at all*, belong? The answer, notwithstanding several thousand years of philosophical reflection, remains unclear, or rather wholly elusive. Hence the metaphorical “crossroads”.

exception of Kant's *Noumena*,<sup>5</sup> initially excluding any relation thereto or conception thereof, nothing can be considered or discussed, let alone interacted with,<sup>6</sup> without the conceptual framework/fabric and grammatical "momentum" at the core of language. Traditionally, setting theology aside, the impartation of meaning to *what is* was ascribed to the cognitive powers of the homo sapiens,<sup>7</sup> viewing our lingual skills as means (essentially a vessel) for carrying out something that occurs on a deeper substrate. But modern philosophical thought, linguistics, and the findings and insights stemming from multidisciplinary approaches to human cognition and its physiological bedrock, highlighted language's constitutive role in the overall architecture, expanding our understanding of what language is.

Indeed, we communicate lingually,<sup>8</sup> but we also think, inquire and appease our intellectual curiosity in language.<sup>9</sup> In point of fact, 'what is meaning?' or 'what things are?',<sup>10</sup> as well as any conceivable answers, are articulations that unavoidably fail to escape the necessary conditions for what it means to *say* something.<sup>11</sup> The profound importance of that fact is particularly compelling in dealing with philosophical questions, namely such that initially address 'meaning' per se.<sup>12</sup> That derives from the fact that articulation – lingual articulation – is at the core of meaning,<sup>13</sup> whereby pursuing for meaning (asking about meaning; articulating a relation thereto) – whatever artificial qualifications we may grant that notion – by way of employing a fundamental aspect or

---

<sup>5</sup> Or similar lines of thought, assuming the notion of total, absolute transcendence (itself a lingual unit, stealthily infiltrating through the back door).

<sup>6</sup> Initially approached, so to speak, at the fundamental existential level.

<sup>7</sup> E.g., Descartes' *Cogito*.

<sup>8</sup> And in respect to the subtler, higher-level aspects of our existential realm, only so.

<sup>9</sup> Not in English or Italian, but in language.

<sup>10</sup> 'What such or such *is*?', namely what does it mean to be such and such (to be that of which we ask: 'what is it?').

<sup>11</sup> Not merely to make a valid or coherent *statement* (point), but to make (produce, issue) a proper articulation. Not to succeed in doing something particular, but to be able to do something of a particular kind. The same way as one does not *succeed* to breathe but either breathes or not, one either *says* [something] or not in the first place. To make a valid statement is something else, namely *saying something* (applying the binary attribute of articulation) within a shared set of behavioral rules (a 'way of life', to echo Wittgenstein), e.g., '*Klaus is German*', ' $3x-1=y$ ', '*that sucks!*' etc. Such different "sets" of rules (ways of life) consist in, reflect or correspond to (there's no precise word for – since there's no ultimate *depiction* of – *that* relation) the different aspects of our place in the world. But the core capacity – the non-scalar property – of articulation precedes and constrains them. By the same token, failing to breathe or run (displaying a flawed disposition or performance thereof) does not amount to a failure in doing something, but a failure in being what one is, and the same goes for *saying* (as distinct from playing chess, cooking or reaching the peak of the Everest Mountain).

<sup>12</sup> E.g., 'what is meaning?' as distinct from any other question, whatever its form ('what/who/why etc. is [any other] 'x'?'), whereby the meaning of 'x' in all the latter is assumed/established by the very questioning. It seems that the very grammatical-semantic possibility to ask a question (inquire, display/express wonder/bewilderment/misapprehension) assumes concrete meaning; a meaningful 'something'. Not its theoretical possibility, but its concrete "availability" (even a theoretical possibility is 'something'). Only a question that purports to address meaning as such further purports to question that availability: to put a question mark on the fact initially assumed by the very questioning.

<sup>13</sup> And/or vice versa.

constituent thereof,<sup>14</sup> is bound to stumble upon principal – logical both practical – difficulties.

Purporting to explore ‘meaning’ through the invisible architecture of meaning is at the very least questionable. Hence, Wittgenstein’s above intuition-observation, attributing language the devilish effect of preoccupation-obsession with metaphysical wonder-blunder.<sup>15</sup>

Moving fast-forward, we are of late witnessing recurring references by leaders and pioneers of the field of artificial intelligence to the philosophical background, or more precisely – the philosophical ‘missing link’, of/in the disruptive existential implications of the technological advances (the technological reality) they proudly<sup>16</sup> preside over and facilitate. To cite two of the most notable representatives of that truly remarkable cohort of individuals, Mr. Eric Schmidt, former Google CEO, observed in a recent interview that: “*there’s no language for what happens with the arrival of this..*”;<sup>17</sup> while Dr. Demis Hassabis, CEO of Google DeepMind, recently observed that: “*I think we need new great philosophers to come about, hopefully in the next five-ten years, to understand the implications of this*”.<sup>18</sup>

Indeed, both distinguished figures evoke the need for lucidity and thoughtful, constructive – big picture – perspective in regard to the practical implications of the mentioned advances on/to the world at large, namely on/to all of us and the ways we live. It seems, however, and we posit that intuition as our first conjecture in this paper, that the philosophical (one wishes to say – existential) gap they allude to, or better still – perceive themselves as alluding to while in fact resonating, is inherent to, and actually stems from, the very technological breakthroughs they so magnificently and masterfully accomplished and continue to accomplish, namely to what they are actually doing. The gap, we argue, or rather – the very void, is not in the effect but in the wellspring, the *effector*, which invites unbiased philosophical assessment not on *what happens with the arrival of this* and not *in the next five-ten years* but on the roots of the arrival, and now.

---

<sup>14</sup> Namely, through something that turns to be more than, strictly speaking, a means (unlike reaching for an apple with my hand, which *does* amount to employing a means), and more precisely – through something that is inherent in (imbued with) the very object of the pursuit (or is inherent in the object of the purported pursuit). Like alleging to see something that is at the base of the very ability to see.

<sup>15</sup> If anyone had already put it that way, we apologize for the seeming stylistic appropriation and assure the reader that nothing herein has been plagiarized. On the contrary, this paper rests on flat rejection of, and genuine resentment to, the prevailing approaches to the addressed issues. Since the author is not a native English speaker, however, we drew upon Google Gemini’s services for highly particular cases of phrasing. But this is where it stops. This paper is our original comment on something we deeply care about.

<sup>16</sup> Absolutely well-earned pride, let there be no mistake.

<sup>17</sup> Schmidt, Eric, "A conversation with Dr. Eric Schmidt, Chair, SCSP and Jeanne Meserve, Host, SCSP Tech Podcast." *YouTube*, uploaded by the Special Competitive Studies Project, 10 Apr. 2025, <https://www.youtube.com/watch?v=L5jhEYofpaQ>.

<sup>18</sup> Hassabis, Demis, "What's next for AI at DeepMind." *YouTube*, uploaded by 60 Minutes, 20 Apr. 2025, [https://www.youtube.com/watch?v=1XF-NG\\_35NE](https://www.youtube.com/watch?v=1XF-NG_35NE).

A second conjecture we shall put forward at the outset, brings together two seemingly unreconcilable theoretical notions: Noam Chomsky's "Language Gene", as a metaphor for innate, genetically endowed disposition to apply language,<sup>19</sup> and Wittgenstein's notion of 'Language Games', as the real-world architecture of language acquisition.<sup>20</sup>

Albeit in apparent divergence,<sup>21</sup> the respective outlooks of those eminent thinkers on the nature of language seem to converge over the intuition that language is precursory to human applicative creativity; that it's more of a framework than a product. On the basis of that intuition we shall propose our third background conjecture, namely that language acquisition/generation underlies and inhibits everything else we can do as advanced biological creatures and cannot serve as an instrument in our hands,<sup>22</sup> least of all for the purpose of transcending our very selves and creating an artificial lingual reality.<sup>23</sup> Consequently, we shall argue, echoing Chomsky, that any kind of 'training' presupposes, rather than facilitates, language skills, and, echoing Wittgenstein, that the genuine arena of language acquisition is an existential "angle" beyond our reach,<sup>24</sup> namely conceptually antithetical (existentially exclusive) to any model thereof. These two arguments, we suggest and will further try to demonstrate, intersect on a deeper level than meets the eye, with even deeper consequences.

In flat contrariety to our perspective on the nature of language and its place in human purposeful existence, the great project of prevailing applicative approaches in the field of artificial intelligence, as manifested in the advance of Large Language Models, seems to rejuvenate the traditional, instrumental conception of language<sup>25</sup> and to base its entire vision – its entire existential prospects – on the notion of 'language skills' as essentially

---

<sup>19</sup> According to which language is in our very nature, namely we arrive in the world already equipped with the ability to apply/practice thereof rather than something *in the world*, which we learn by means of instruction or any kind of apprehension or training. Language, according to Chomsky, is like breathing and moving rather than cooking, doing math or the game of basketball.

<sup>20</sup> According to which we acquire our language skills through participation in preestablished real-world activities.

<sup>21</sup> We hope to show, as we proceed, that those two outlooks have much more in common than points of contention.

<sup>22</sup> It is an instrument/property of the species, not of the individual specimens thereof.

<sup>23</sup> The same way as building artifacts that move does not amount to playing with space and time and artificial insemination is not playing with the origin of life. By the same token, developing Large Language Models and building machines which display something that is strikingly similar to human lingual discourse, does not amount to teaching language (conferring upon something, or equipping it with, language skills).

<sup>24</sup> It is a metaphysical "livestream" in which we are "locked" – absolutely absorbed – as language acquisition/generation unfolds (which is demonstrated by the fact that we cannot say – are logically/categorically precluded from "seeing" – which precedes the other: the acquisition or the generation). Language acquisition amounts to 'action + [X]' or '[X] + action'. That kind of entity (of *something*) cannot be perceived/conceived and therefore cannot be modelled. Like 'God' or 'Beginning' as such cannot. Pattern recognition is not *that*. The ability to say "pattern recognition" precedes *that* and misleads us to think they are identical. The very possibility of employing that notion ('action + [X]' or '[X] + action') is only relevant to this paper. It is the ladder we should throw away once it's concluded.

<sup>25</sup> Namely as a means/vessel for carrying out something that occurs on a deeper substrate.

an instrument – a conduit to existential manipulation. To cite two other distinguished figures in the AI industry,<sup>26</sup> Mr. Mohhamed Gawdat, former Chief Business Officer for Google X<sup>27</sup> and renowned author and spokesman on the future of AI, shared in a recent podcast<sup>28</sup> his experience of the grassroots of autonomous machine learning in Google, back in 2008,<sup>29</sup> revealing, if only indirectly but unequivocally all the same, his core belief of what language is:

*“..and eventually, again – metaphorically, one of them came back and said: “I found something!”.. and remember, we did not give them a task to look for anything, no instruction, just observe and learn from patterns, **which by the way is the way we teach kids** [emphasis not in the original text, DD].. sort of, metaphorically, one of them came back and said: “I found something” and we needed to write more code to find out what it found, and obviously it was watching YouTube, so it found a cat.. and it didn’t find one cat, it found **the essence of what makes a cat a cat** [emphasis not in the original text, DD], and that’s quite interesting about learning and AI in general.. If you take a child, a toddler, on a quick tour in a car, you’d hope **that when the toddler sees the next car he can tell you: “oh, that is a car”**.. [emphasis not in the original text, DD]. It could be a different car, it could be bigger, smaller, different color, different make, **but the child would be able to understand the essence, not the memory, of something** [emphasis not in the original text, DD]. **That is learning, learning is – you see enough cars and now you know what makes a car a car,<sup>30</sup> you see enough cats and you know what makes a cat a cat**.. [emphasis not in the original text, DD]”. (Gawdat 08:00).*

According to Mr. Gawdat, the quintessence of learning ‘what something is’<sup>31</sup> amounts to a Eureka-like moment of *identification* (e.g., “I found something!” – **the** something all the converging patterns the system distinguished and connected amid the general flow of

---

<sup>26</sup> We apologize for the fact that these two citations are slightly longer, but this is important for accuracy and proper demarcation of the point of contention.

<sup>27</sup> Currently ‘X Development LLC’.

<sup>28</sup> Gawdat, Mohhamed, "Ex-Google Exec Reveals the Shocking Truth About AI with Mo Gawdat." *YouTube*, uploaded by Bad Decisions Studio, 13 Feb. 2025, <https://www.youtube.com/watch?v=cfPeP0s8EKU&t=6477s>).

<sup>29</sup> Experimenting with the basic design of algorithms for pattern recognition by feeding them with enormous quantities of image data (YouTube videos), which resulted – as part of the algorithms’ very operation – in the buildup of neural networks in growing layers of abstraction up to a pinnacle supposedly capturing the essence of a certain object (represented by a distinct field of image data within the general flow). In other words, facilitating an algorithm’s capacity to recognise *something* with no prior instructions to recognise *anything*, namely discovering and establishing its capacity to identify semantic coherence, otherwise known as ‘aspects of the world’, solely on the basis of its own “raw” function, without any external (human) command (of the kind: “do this”, followed by a set of “if  $x_s$ , then  $y_s$ ”).

<sup>30</sup> Namely, can *say so!* [remark not in the original text, DD].

<sup>31</sup> Not in consequence of something like research but on the fundamental level of becoming acquainted with the world – the kind humans undergo (are immersed in) as part of their early cognitive development.

data it was fed with, represent). The “collapse” of the ‘identification function’, to borrow a hackneyed expression from another discipline, ipso facto equips/endows the “learner” (the human agent or the machine) with a comprehensive sense of fundamental familiarity that will, in turn, enable him/her/it<sup>32</sup> not merely to denote properly instances of the learned thing, but to *speak* about it properly, namely apply it as part of his/her/its evolving<sup>33</sup> language skills.<sup>34</sup>

In another podcast,<sup>35</sup> Prof. Terrence Sejnowski<sup>36</sup> explained, elaborating on advanced LLMs’ capacities to display such astounding semantic interaction with humans, that:

*“..it’s astonishing.. Actually, there’s more and more evidence that humans use a similar type of self-supervised algorithm, and this is actually in another part of the brain, called the basal ganglia, which is below the cortex and there’s a loop between the cortex and the basal ganglia. But we now know that the basal ganglia is important for what’s called ‘procedural learning’, that is to say **learning a sequence of actions to achieve some goal or reward** [emphasis not in the original text, DD]. For example, it’s used when you’re learning how to play Tennis. **It’s used in infants when they’re learning how to talk** [emphasis not in the original text, DD]. Language is, during the first two years of life, I mean this is amazing how quickly it happens, **and it happens because this basal ganglia is doing this prediction** [emphasis not in the original text, DD]. Predicting is called reward-prediction error. It predicts whether the next step is going to be, give you a good outcome or a bad and it uses the difference – the prediction error – in order to update all the connections, the strength of the connections. So that is a very important general principle that we discovered, and the brain uses it, transformers use it. It sounds very simple, **I mean how could you get information out of that? It’s just whether you did better or worse** [emphasis not in the original text, DD]. But if you just do that every moment, your building up with enormous.. what is called a ‘value function’. That is to say the value of a particular task, a particular action you’re going to take – is it a good one or a bad one”.* (Sejnowski 13:34).

According to Prof. Sejnowski, learning to speak and learning to play Tennis are fundamentally the same thing.

---

<sup>32</sup> With further training and software engineering.

<sup>33</sup> Not in the case of the machine, as we shall see.

<sup>34</sup> ‘It drives’, ‘it’s fun’, ‘let’s have one’, ‘where is it?’, ‘why did it stop?’ etc.

<sup>35</sup> Sejnowski, Terrence, " Terrence Sejnowski on ChatGPT and the future of AI." *YouTube*, uploaded by Closer to Truth, 21 Apr. 2025, <https://www.youtube.com/watch?v=CRHJMC9bBs0&t=917s>.

<sup>36</sup> Francis Crick Professor at the Salk Institute for Biological Studies and the director of the Crick-Jacobs centre for theoretical and computational biology.

In summary, as far as Mr. Gawdat and Prof. Sejnowski are concerned (their respective observations converge over the basic point), having passed through a brew (a synthesis) of physiological and phenomenological processes (having observed a sufficient number of cats or cars and/or drawn sufficiently robust predictive power from an adequately functioning basal ganglia), the developing human creature will learn to grasp, express and communicate lingually<sup>37</sup> that which was established about him/her on another, fundamental level.<sup>38</sup> And the same, so the argument/story goes, applies to the machine.

And it all occurs not only in front of our very [observing and inquisitive] eyes, but in a fully accessible [existential] arena.

To the best understanding of those opinion leaders, and by extension of the majority if not the entirety of the AI research and entrepreneurial community, language acquisition/application is no different from any other human appropriation/skill for living and survival; something we acquire and apply to achieve *our* goals given *our* perception and understanding of *our* environment (express ourselves, communicate with others, build, create and facilitate stuff and generally do whatever we need or choose to do in the course of our lives), whereby it isn't something we cannot model in the artifacts we build.<sup>39</sup>

With all due respect, we strongly disagree.

By the end of this paper, we hope to make the case for three complementary claims and four conclusions, as follows:

- a. Ludwig Wittgenstein was right;
- b. Learning 'what something is', or in other words – the emergence of semantic coherence, has nothing to do with *identification* as such. Commanding the requisite standpoint for that insight requires genuine philosophical nitty-gritty, far-off from the primitive, blurry conflation of 'meaning' with 'identification', and cuts a sharp rift between the AI industry's ultimate pretensions<sup>40</sup> (e.g., AGI) and its actual achievements and creations;

---

<sup>37</sup> In/through/with language.

<sup>38</sup> Use language for the conception and/or expression and/or communication thereof. One should rightfully ask here: if learning to move (or move in the proper directions) or breath is not learning to *display* movement (or proper movement) or *display* breathing, why is learning to speak an acquisition of an ability *to display*? Why is learning to conceive, express or communicate a relational/dependent or adjacent capacity where other, similarly emergent capacities, are not? But we are getting ahead of ourselves here. Let's proceed step by step.

<sup>39</sup> Note the suspiciously innocent *given* at the heart of the paragraph. Isn't it already assuming everything it is supposed to affect (condition/curb)? Here too we are expediting what lies ahead, and the weight of that question will become evident later. It is, however, important to draw attention to that preliminary discrepancy already at the outset as it affects our entire journey.

<sup>40</sup> Pretensions masked as aspirations. This subtle point will also become clearer as we proceed.

- c. Language is a property of the world. Human language is a variant thereof,<sup>41</sup> not a “lens” onto the world or a category that stands in relation (applies) thereto;

**Therefore:**

- a. What and how artificial intelligence *does* learn has nothing to do with what human learning *is*;
- b. Insufficient understanding of what we are [as agents in a world] necessarily inhabits our understanding of what AI systems can be, or more precisely: the definitive limitations on our understanding of how we emerge as agents in a<sup>42</sup> world exclude an ability to cater for any such artificial emergence;
- c. Being what they are, AI systems by definition cannot mimic that which “builds” agentic properties (and agentic reality) in the real world;
- d. Artificial General Intelligence (AGI), in the sense of being more than an artifact, is in principle unattainable, no matter how formidable and elusive that which *is* attainable may be;

This is not a treatise about AI. We are wholly unqualified for that task, nor do we have any such aspirations. It is, however, a work on meaning and language, which are at the core of the AI industry. The implications of the reflections and analysis below are at the heart of the developments that will most probably define our age. The proper role of philosophy in exploring the broader implications of this major advance will be restored once human understanding has reconciled itself with the advance itself, whereby we too, contra Wittgenstein, refuse to “throw the ladder”<sup>43</sup> once our present journey will have consumed itself. What lies ahead, however, as an immanent precursory condition to any such pending philosophy, is not a game of chess, whereby no preview of the future can command our course. In what follows, we respond to the way things are, rather than aiming to win over or achieve anything. We therefore address the basics first. We believe that a genuine philosophical perspective is far more urgent and myopic than Dr. Hassabis assumes.

Without further ado, let’s delve in.

---

<sup>41</sup> A variant of ‘language’ as a fundamental property of the world.

<sup>42</sup> Always ‘a’, never ‘the’ [world]. Only fictional characters emerge as agents in *the* world.

<sup>43</sup> “*My propositions serve as elucidations in the following way: anyone who understands me eventually recognizes them as nonsensical, when he has used them – as steps – to climb beyond them. (He must, so to speak, throw away the ladder after he has climbed up it.)*” He must transcend these propositions, and then he will see the world right.” Wittgenstein, Ludwig, Klement, Kevin C. (ed.), *Tractatus Logico-Philosophicus*. Translated by Pears, D.F; McGuinness, B. F. (side-by-Side-by-Side ed.). University of Massachusetts. Retrieved January 27, 2019, 110.

I

Lately, I listened to an interview with the neuroscientist, David Eagleman.<sup>44</sup> Beside the general insightfulness of the interview at the popular level,<sup>45</sup> for me as a philosopher it touched upon a profound, gnawing philosophical neuralgia, namely the meaning of words and, more generally, the ultimate cause/nature of semantic coherence, which captured my focus and rattled the cage of the philosophical demon in my head. At a certain point I felt that what I hear literally demonstrates Wittgenstein's insight concerning language's "deceitful" role in generating metaphysical wonder. But the more I lingered over the matter, the clearer it became to me that it catches far more than a philosophical point and that expanding upon that insight in conjunction with the interview's case study, may stretch far beyond philosophy and even so far as the theoretical bedrock of the vanguard breakthroughs and frontiers in the field of artificial intelligence and its swelling infringement upon the notion of human agency. What follows is an attempt to elaborate on that intuition.

To rule out any inaccuracies or misconceptions, here is a succinct summary (retaining the form of a dialogue) of the relevant part of the interview, in which Mr. Eagleman elaborates on the way he exploited the brain's flexibility as an interpreter of sensory data in order to circumvent hearing impairment and enable people who cannot hear to take-in and process audio data (sound in general) by providing an alternative channel for that data into their brains:

**Eagleman:** *...it's a wristband for deaf people that captures sound through vibratory motors, breaks it from high to low frequency and the result is that the person who wears it **feels the sound** on his/her torso. This is exactly what the inner ear actually does for people with proper hearing: breaking the sound from high to low frequency and shipping it to the brain. We've just "transferred" the ear to the skin of the torso. In that way, deaf people can literally "hear" the world.*

**Levitt:** *So if I understand it right, sounds "happen", and this wristband "hears" the sounds and then "shoots" electrical inputs into your wrist that correspond to the high-low frequency?*

**Eagleman:** *It's just like vibratory motors, it's just like the buzzer in your cell phone, but we have a string of these buzzers all along your wrist and we're actually taking advantage of an illusion which is if I have two motors next to each other and I stimulate them both, you will feel one virtual point right in between, and as I change the strength of those two*

---

<sup>44</sup> Eagleman, David, "Your Brain doesn't Work the Way You Think." *YouTube*, uploaded by on Freakonomics Radio Network, 23 Dec. 2024, <https://www.youtube.com/watch?v=Ma-dSfHaZdQ>.

<sup>45</sup> Conducted by Prof. Steven Levitt.

*motors relative to each other I can move that point around so we are actually stimulating 128 virtual points along the wrist... **you see the dog's mouth moving and you feel the barking on your wrist or you close the door and you feel that on your wrist or you say something (many deaf people can speak) and know what your motor output is by "feeling" the input(!), namely by sensing on your skin what you are simultaneously saying...***

**Levitt:** *Hearing their own voice for the first time..., oh God, that's interesting...*

**Eagleman:** *...exactly, and BTW that's how you learn how to use your ears too. When you're a baby, you're watching your mother's mouth move and you're hearing<sup>46</sup> data coming in your ears... and you clap your hands together and you hear something in your ears, it's the same idea, you're just training up correlations in the brain about 'oh this visual thing seems to always go with that auditory stimulus'*

*...and that sounds so crazy, but remember, that's what your ears are doing, your ears are capturing vibrations at the ear drum that moves from the middle ear to the inner ear, breaks up to different frequencies, goes off to your brain, goes to your auditory cortex... this is GIANT pathway of things... and yet even though you're hearing my voice right now INSIDE YOUR HEAD, you think I am somewhere else and that's exactly what happens irrespective of how you feed the data in. (Eagleman 11:50).*

## II

So why do I think that this intriguing conversation and even more exciting technology demonstrate that language deceives us to ask philosophical questions or that the specific "mystery" that seems to spur metaphysical wonder is a product of our ability to speak rather than a blur or a puzzle in the world that can (*should*, once identified) be clarified or understood the way empirical phenomena (the complex composition of matter or living cells, psychological states etc.) or mathematical puzzles can?

The answer is that Mr. Eagleman's technology seems to demonstrate Wittgenstein's original insight that the ultimate meaning of a word<sup>47</sup> does not consist in anything beyond the way we use and apply it (that word or any other part of speech) in practice given the various aspects<sup>48</sup> of our existence, rather than any ethereal, platonic entity that is a

---

<sup>46</sup> *Author's mistake – should be 'experiencing'.*

<sup>47</sup> Its meaning as "something"; as 'something in the world'; as what it is to *be* a 'word' in the world.

<sup>48</sup> [Our] *ways of life* as advanced biological creatures in an environment governed by the laws of physics and the other natural sciences. Namely, given the ways we actually subsist/live in the world rather than any

product of [something like] intention or other mental projections. In other words, it demonstrates, contrary to our intuitive conception (e.g., ‘I say what I mean...’), that we do not “inject” meaning into parts of speech by intending to say this or that but that what gives meaning to our living application of language is the functional-environmental context of applying (*saying*) that word as part of particular situations (‘ways of life’, in Wittgenstein’s words). Moreover, it demonstrates that “learning” – “picking up” the ability to use – language<sup>49</sup> does not proceed by means of straightforward association of something like ‘substance’ with something like ‘meaning’ (this is ‘red’, this is a ‘cat’, this is to be ‘angry’ etc.) in the first place,<sup>50</sup> but through the practical, actual immersion of the human agent – the advanced biological creature – in the multiple pathways and possibilities of its existential environment – in the human creature’s *becoming* an integral part thereof; in its realization/actualization of its features as a creature in the world. Finally – in an interplay rather than incorporation/internalization. When humans *learn* they [thereby] become *part of* something rather than *understand* anything (in the sense of acquiring a finer perspective thereupon).<sup>51</sup>

### III

How does this technology demonstrate that? This is the most beautiful, thrilling part.

At the intuitive level, to ‘hear’ is to grasp the meaning of audible qualia or vocal messaging/communication. Someone says something, birds sing, thunder rolls or music is playing, and we *hear* it. It is an *auditory* thing; something that’s *heard* rather than

---

preliminary or antecedent design or trajectory. According to that view, it is the ways in which we develop and live that emanate sense in – “disseminate” meaning into – existence as such rather than any applicable (relatable) sense of meaning which underlies them (underlies existence as such).

<sup>49</sup> Not learning *a language* in the sense of some specific language, but the very acquisition of the fundamental skills at the core of applying language – the fundamental skills of being *a speaker* (a language animal) in the first place.

<sup>50</sup> Spoiler alert: which is also not the way artificial intelligence systems are trained to display absorption/generation of “meaning”. Recall Mr. Gawdat’s observation quoted above: “*and it didn’t find one cat, it found **the essence of what makes a cat a cat***” (Gawdat 08:45). “No”, we say, “for the purposes of what it means *to speak*, it didn’t. It acquired a wholly different capacity than what it means *to speak and understand* in the human sense, and its training, or whatever other way we would describe or refer to whatever it was engaged in to have been able to achieve what it had achieved, has nothing to do with *learning to speak and understand*. The need for spoiler alerts will subside as our arguments unfold. We resort to that rhetorical technique due to our reluctance to leave loose ends. The prevalent discourse on the issues we address is, to our humble opinion, characterized by a pervasive tendency to leave loose ends and unresolved details. The preliminary results of the exciting experiments Mr. Gawdat has been referring to in the quote above might have (indeed, they surely had) sufficed for Google’s applicative purposes at the time, but they certainly fall short of making any adequate statements regarding what ‘understanding’ or ‘learning to speak’ (namely that the system *understood*) is. Notwithstanding, the interview emphatically posits that it does. Hence our spoiler alerts and the abundance of explanatory footnotes in this paper. They might seem tedious, but they are vital to our goal.

<sup>51</sup> One could ask whether this is also true of “grasping” abstract “truths”, be it a mathematical breakthrough or the reason for a complex process or event. Our answer is in the affirmative and we hope to shed light in what follows on the reason why we believe this to be the case.

grasped/experienced in any other way. That's what anyone will tell you if asked: I hear the sound of all these acts/processes/events and understand them as such. There are these things called *sounds* and *them* I hear. If pushed a little further, they will proceed and say that 'hearing' is the ability to, as it were, "take in" these different sounds and understand them as the particular audible<sup>52</sup> pieces of data they are: messages<sup>53</sup> in the form of sound. *Voices, noises, sounds* are "there", and I *hear* them.<sup>54</sup> Another possible answer may be: I hear through my ears, which are the physical-biological mechanism through which sound is absorbed into the human organism and transmitted to the brain. Eyes see, the skin touches and ears *hear* their respective, particular kind of message (informative fabric). Entering the organism's auditory system, soundwaves (the "carriers" of sound) are transformed into the kind of qualia the brain processes, making it possible for the agent I am to interpret them as the specific pieces of experiential data/fabric they are.

The crucial point in the above popular interpretation of 'hearing' or 'to hear' (or any further variant thereof) is that besides the more elaborate explication of the intricate physical background (sound waves, frequencies etc.) and the physiological mechanisms at the root of human hearing (the entire auditory system and its incorporation into the human brain via the nervous system etc.), the entire set of the grammatical variants of the word 'hear',<sup>55</sup> which comprise the domain of its "meaning" for us (to hear, hearing, be able to hear, have heard, upon hearing etc.), is directly and exclusively associated with the experiential interpretation of sound: if I hear anything, someone (a person) or something (nature) had generated [the] sound [which I hear]. To *hear* is to experience [generated]<sup>56</sup> sound.

If we return to Mr. Eagleman's technological application, which allows people with substantial hearing impairment to "hear through their wrist" (or any other part of their torso, for that matter), or in his own words: "***you see the dog's mouth moving and you feel the barking on your wrist or you close the door and you feel that on your wrist or you say something (many deaf people can speak) and know what your motor output is by "feeling" the input(!), namely by sensing on your skin what you are simultaneously saying...***" (Eagleman 13:40), we see that the novel way this technology allows for the "experience of sound" on behalf of an agent for whom it was previously inaccessible, essentially cuts off that notion from its otherwise trivial 'audio' setting (scenery), namely from our traditional/intuitive conception and understanding of that notion: suddenly, you no longer merely "hear" sound, but can also "feel" it on your body!

---

<sup>52</sup> As it were "voiced", whether it is the voice of a human or the voice of nature or the universe (in a cosmological, not a new-age, metaphorical and ultimately hollow sense of "the universe tells me what's right for me...").

<sup>53</sup> This is a key notion in what follows, albeit not in its usual, commonsense connotation, and we should bear it in mind.

<sup>54</sup> Are they "there" for me? This is another fundamental question we should bear in mind.

<sup>55</sup> Its semantic field.

<sup>56</sup> Far from being trivial, this is very crucial in a way that will unfold below.

Whereas previously ‘sound’ was associated with – and experienced as – taking-in a process (some sort of “happening”, as prof. Steven Levitt, the interviewer, interpreted Mr. Eagleman’s explanation at a certain stage of the interview) in the audible domain (where “auditory” stuff happens, namely where you *hear* rather than see, touch or smell stuff), it has now transformed into [taking-in] a process in the domain of our torso (where tactile sensations are generated). Yet in other words, whereas before, the roar of the thunder or Pink Floyd’s melodic themes or my mom’s shouting “*Daniel... come eat*” was thought to have been “happening” in the [very] particular ether where live speech, live music and the sound of nature “trumpet” the full array of their structured noise, it has now moved (migrated) to where I encounter and get the feel of the surface of other objects – to my skin...<sup>57</sup>

And while<sup>58</sup> at the scientific level that difference is explained away quite easily, as Mr. Eagleman himself has shown during the interview,<sup>59</sup> at the philosophical level – where the fundamental / ultimate meaning of words, concepts and stuff in general matters – something very significant happened. Namely: if I *hear* [the audible fabric of] sound but also *feel* [the touch of] sound, what is to *hear* [audibility]? And if I *feel* [the touch of] sound but also *hear* [the audible fabric of] sound, what is to *feel* [tactility]? In face of such pervasive or even complete overlapping of these notions, what is their respective, distinct meaning?<sup>60</sup>

#### IV

As intelligent people we understand that the faculty of hearing is rooted in physiology, namely complex physiological processes in the ear and the brain. As educated people we further understand that evolution has fostered and refined these properties of sensual perception so that we may perceive the sound/s of the world.<sup>61</sup>

---

<sup>57</sup> It is as if I just, and always, *feel* all the sensual experiences I have, and the particular “ways” I thought I do it, namely *hear*, *smell*, *touch*, and *see*, are different “expressions” (words, lingual symbols) of the same thing.

<sup>58</sup> Please bear with us. We are slowly approaching the punchline. In the tricky landscape of genuine philosophy, it is a tedious task, since many intuitive forces work against us.

<sup>59</sup> “*This is exactly what the inner ear actually does for people with proper hearing: breaking the sound from high to low frequency and shipping it to the brain. We’ve just “transferred” the ear to the skin of the torso... and BTW that’s how you learn how to use your ears too. When you’re a baby, you’re watching your mother’s mouth move and you’re hearing data coming in your ears... and you clap your hands together and you hear something in your ears, it’s the same idea, you’re just training up correlations in the brain about ‘oh this visual thing seems to always go with that auditory stimulus’.*” (Eagleman 12:10).

<sup>60</sup> Their abstract marked-out value *as such* (as *what they are*).

<sup>61</sup> You talk to me and [thereby] *the world* “talks” (communicates) to/with me at the same time.

As intelligent people we can further fathom that what seemed like an exclusive property of the ear can be exercised via other organs like our wrist or further surfaces of the torso,<sup>62</sup> the same way it is possible, in principle, to walk on our hands or (with technological “intermediaries”) poop otherwise than through our rectum.

What is far less trivial to grasp is that the **word** or notion of ‘hear’ has actually lost its previously exclusive, essential connection to the world or domain of sounds and the notions of ‘auditory’ or ‘vocal’, as we thought we knew them. It had migrated, at least on the face of it, to the broader “blanket” domain of experience.

We should pay heed here: what happened at the conceptual level is not a change or transformation in the *way* of hearing or in the *means* with/through which we hear. It even isn’t a mutation in our *faculties* of hearing. Means and performance or sense and experience is not the framework of the conceptual shift<sup>63</sup> we had just witnessed, but a total meltdown of the word / notion / lingual vessel (qualia) of ‘hearing’. At the practical level we may (and will) continue to use that word as we did, and people with hearing impairment may now evidently join us in doing so, regardless of the different mechanisms allowing them to participate in that practice (a whole realm of shared experience).<sup>64</sup> At the philosophical level, however, a whole paradigm has been shattered, with non-trivial theoretical implications to our conception of what we are. Let us explain.

## V

Examined on its own, the case (phenomenon, instance) of someone hearing my voice/speech or an operatic Aria through his/her wrist (after due training and guidance in the practical means of doing so, as in Mr. Eagleman’s technology) can be reasonably accommodated in our commonsense model of reality. Technological solutions and substitutes to natural phenomena/faculties are nothing new. We forge artificial (human-made) mechanisms that imitate natural ones with various, and sometimes phenomenal, degrees of success in a way that improves and upgrades our human condition in the most general sense of that notion. Science does make “miracles”, albeit coherent and ultimately tractable ones, and this case is no exception. Technological progress (namely human intellectual ingenuity) disrupts what is practically possible, and if we can [come to be able to] hear, see, walk or breath by other means (however nonorthodox and disruptive) where the natural ones fail us (for whatever reason), we should only be

---

<sup>62</sup> Soundwaves stimulated various mechanisms/receptors at the inner ear, and now the same soundwaves stimulate – with some technological “boost” from without – other physiological receptors in another part of our body (it is, after all, the same body, isn’t it?). On the face of it, it isn’t insane.

<sup>63</sup> A manifest linguistic crack, actually.

<sup>64</sup> Which is not really a *shared experience*, this being a strictly metaphorical term, but a shared realm of practice – the application of language – as part of a whole, comprehensive aspect/way of life (of belonging to the human species). ‘Shared experience’ is the very crux of the whole matter. We have embarked on an indirect journey to the heart of the maze.

grateful. That's, after all, the fundamental feature of our species' journey from the flora & fauna world of cavepeople to the anthropogenic Technosphere we inhabit today. Hearing through our toes or providing for our bowel movement through a tube in our stomach is still *hearing* and *defecating*. Very exciting and even awe-inspiring, but philosophically and conceptually trivial.

The point, however, is not that soundwaves can be manipulated to create a neural transmission similar to the one generated by the inner human ear<sup>65</sup> through an alternative mechanism fastened to (coming into contact with) any other piece of healthy human skin. This is scientifically and logically coherent and given the current scientific and technological levels of progress, possibly even trivial.

The point has to do with the 'word' hear,<sup>66</sup> and specifically with the **acquisition of the ability** to apply (use & understand) it as a word. Not as the very complex physiological process referred to above,<sup>67</sup> but as a word. To *talk, reflect, communicate* (encapsulate, signify) all that [physiological] complexity on an altogether different level than the one in which it occurs.

## VI

When a previously deaf person learns to *hear* and experiences for the first time the subjective-experiential correlate of the aforementioned neural transmission (which non-deaf people experience unaffectedly upon hearing their own and other people's voices, for example),<sup>68</sup> the same thing happens which we – the non-deaf – had undergone already at the early stages of our infancy: a novel kind of experience<sup>69</sup> gets correlated in his/her "perceptual theatre" (the conscious "stuff" of his/her immediate environment) with another particular event<sup>70</sup> in that same perceptual environment. A person undergoing such initial introduction and training in using the aforementioned technology will be prepared to assimilate that correlation,<sup>71</sup> whereby the "novelty" of the situation will focus,

---

<sup>65</sup> Same soundwave input, same neural output.

<sup>66</sup> Without which all of that is nothing to us; irrelevant, non-existent.

<sup>67</sup> Which, like any other physiological activity, is merely and always executed as such – as an essentially mechanical, rule-governed, deterministic process – and never applied in any other way.

<sup>68</sup> Experiences *what it is like* to hear something. What it is like to be in the state everybody refers to as 'hearing'.

<sup>69</sup> What we call 'sound' or 'noise'.

<sup>70</sup> The movement of another person's lips or the operation of a radio set or some other event like the falling of some object on the floor or even some abstract event in the skies (which he/she will later learn is called 'thunder').

<sup>71</sup> He or she will have been informed that they are going to be introduced to the sense of "hearing" for the first time and they understand, roughly, that it should be like a new, additional way to perceive – interact with – reality (the environment). They already know what the latter is; they are capable of connecting perceptual scenery with symbols (they can see, smell, touch) and indicating the former by means of the latter.

for him/her, on the novel [essentially structural] “fabric” of the experience rather than the correlation itself.<sup>72</sup>

After a while, when due training and sufficient practice have been completed, that person will be able to say: ‘I can hear’ and ‘I understand now what hearing is’ or even ‘do you hear me?’

But what is it that enables him/her to do so now? What is it that accounts – **from within, not from without**<sup>73</sup> – for the “propriety” of all these sentences he/she now utters in an allegedly proper way?

In yet other words: we know that the person with the defect in the normal operation of his/her inner ears is undergoing a new experience, and we know that he/she is connecting that experience with an event whereby someone or something is generating some kind of message to him/her (or in general<sup>74</sup>). But why/how all that corresponds to or sums-up

---

<sup>72</sup> The correlation is interesting/relevant from the standpoint of the world, namely from without (“he/she can hear”; what the world “has to say” is henceforth audibly accessible to them). From the perspective of the protagonist, it is not, in the first place, a correlation (“what did she say?”, “what was that?”) but a new world, a new realm of being (what we call *experience*).

<sup>73</sup> It is clear why we, from without, deem them proper. “he/she can hear now” will be our subconscious sentiment in real time, since we already know what *hear* is in the sense that we are capable of interpreting phenomena in specific ways: we perceive certain behaviors like a person walking, laughing or being sad and are able to identify them as such – walking, laughing, experiencing sadness, and this is exactly and solely why we’re further able to observe properly that “he/she can hear now”. But what makes it proper from within that person? What explains the shift that occurred for him/her to be able to report a new capability (not merely an experience) rather than for us to be able to interpret a familiar sentence (reporting a new capability is different from applying a symbol)? When a person says: “oh, I see that now” in the sense of “I see now what I haven’t noticed before”, we understand that sentence in the sense that we understand that something had prevented that person from seeing/spotting some particular object, which was now (namely, the impediment) removed. But if a person that hadn’t seen anything before and therefore doesn’t know what *seeing* should “feel” (be like), would suddenly say “I can see now”, we would have to ask for some further explanation regarding what to become able to see (rather than *see something when seeing* is already an established faculty) is for him/her. And again, the question isn’t about the ability to experience coherently. This is, roughly, as trivial as tasting a new, unexpected flavor of ice-cream. The question pinpoints the [acquisition of the] ability to apply a symbol in a way that is not merely associative (e.g. “experience x = tree”, “experience y = pain”), but which *eo ipso* encapsulates the entire contextual environment (the whole environmental/scenic background) that that symbol denotes for those who use it (at this stage we should admit that we are also completely uninformed about the way all these people acquired that capability at their early infancy, which is the ultimate question we wish to shed some light on in this paper).

<sup>74</sup> The rustling of the leaves or the thunderbolt are ‘not for me’ and ‘for me’ at the same time. Here’s another possible solution to the philosophical wonder as such: so long as I have noticed the thunderbolt, the rustling of the leaves or the sound of whatever it is that is on the ground and somebody is treading thereon, or the music in the concert hall, there is no genuine puzzle about them, where they come from or what they’re here for. Science explicates all these to their core in a very long, complex, maybe infinite or endless, but consistent (as anything ever was) journey. And so long as I haven’t [noticed them], they aren’t there for me to pursue, namely not in the sense that they do not exist, but in the sense that no question rises in their regard. One can challenge that line of thought and argue that they’re still there for others to ask or wonder about, whereby the question persists. To which we reply that it persists for each individual/pensive mind on its own, whereby and to whom the same principle nonetheless applies. We tend to bestow upon our questions (as well as the reality their respective answers relate to) an abstract, general flare or status, as if they existed “somewhere” on their own for all of us to draw upon and employ. We say that 2+2=4 is

[in]to the *word* ‘hear’ as we – the non-deaf – have come to understand it? What or where or how is that peculiar “summing-up” proceeds ***in the subjective whereabouts of that person/agent?***

## VII

Maybe by seeing what it – this transformative summing-up; the acquisition of the principal ability to relate to reality (to a realm of “stuff”) via “animate” symbols – is evidently not, we’ll be able to approach closer to what it actually is.

When we utter a ‘word’ or write it or think about it (imagine it in our mind’s eye), we have a kind of straightforward, abstract (and yet transparent and present in an impeccable sense of presence and transparency), understanding of what it is. When asked about anything familiar (what is an apple? How much is  $2+2$ ? etc.), right before or in parallel to absorbing the question and spitting out the answer, we have a kind of “flash” – a sort of “unsummoned” elucidation – in our mind’s eye of *what it is* that we’ve been asked about. Not of the [correct] answer<sup>75</sup> but of the *thing* we were asked about and of which we are about to say something (i.e. the “identity/essence” of which we are about to report).

In our commonsense, uncritical conception of such a situation there is an interplay in three respective levels: the [some]*thing* we were asked about (the apple, the historic

---

independent of anyone thinking  $2+2=4$  and infer that the rules of mathematics have their independent, abstract existential realm. But before anyone thought or inferred or realized or identified that  $2+2=4$ , there was the question or wonder regarding ‘ $2+2=?$ ’ or ‘how much is  $2+2$ ?’ And even if no one ever formulated that question [in the sand] before coming to that basic equation or mathematical phrase, there was wonder – initially complex and foggy wonder – which gradually distilled into the kind of wonder that inspired and initiated the cognitive activity that culminated in the first realizations and formulations regarding  $2+2=4$ . **A sense of ‘wonder’ that is categorically distinct (a different ‘language game’) from the postulate emerging therefrom.** So concrete, practical questioning, or wonder in general, always precedes anything that *is*, and every single question regarding what *is*, namely what we [essentially] perceive as existing, will eventually find its answer. The only exception concerns the psychophysical [alleged] interface that lurks at the ultimate background of any genuine philosophical wonder and on which we hope to shed some light as we proceed. There also remains, however, one, allegedly “big” question, which at least theoretically cannot be dismissed merely for the lack of any evidence for a positive or a negative answer thereto or the difficulty even to fathom (imagine) a picture of a proper answer in our minds: can there be or is it possible that there is some abstract, general perspective that asks or echoes or precedes/underlies all the questions, wonderings and reflections we intelligent creatures are capable of (a perspective in a mode of a question mark)? The idea, despite its mindboggling and outrageous gist, is not utterly ludicrous, if only because something links and connects all these individual wonderings and questionings (specifically the form of an interrogative statement: the question mark). The obvious answer that pushes its way is: language. It is language that does so. But this seems (on the profound, analytical rather than the intuitive level) as merely reiterating the original question/puzzle, since we began our inquisitive journey by asking what a ‘word’ is and how do we come to [masterfully] use it as the cornerstone of meaning and therefore of existence at all. It is, therefore, in the least problematic to “sneak” it into any explanation through the back door.

<sup>75</sup> We are not a machine, performing tasks or merely *performing/executing/functioning*. We are – experientially – *more*. We are not merely *performers*. The experiential aspect of what we are (fleeting forms in time and space, biological creatures etc., would be other possible references) underlies all that is existentially important to us, being the starting and the ending point of any perspective or discussion.

event, the mathematical phrase etc.), its corresponding lingual reflection/expression (the articulated words/names for these three “things”), and a medium (the *agent* or *perspective*) which by means of articulation – the generation of the word/phrase as such – provides the thing transparency of meaning.<sup>76</sup>

Following that intuitive conception, a word is the lingual form<sup>77</sup> of the “thing” and our ability to use language is what enables us to communicate the ‘world’ (all the stuff that exist) to ourselves and to others (communicate = exchange and relate to meaning; as it were, “swim” in meaning<sup>78</sup>).

Consequently, the word ‘hear’ is the lingual form of the *ability* to hear or the *faculty* – the sense of – hearing. Not of the physiological process underlying that ability, but of the ability/faculty as such. And in the same way, the word ‘word’ is the lingual form of the lingual entity a word is.<sup>79</sup>

Needless to say, that this conception assumes that there are *objects* such as an apple, an event, a mathematical phrase, an ability or a word in some [existential] domain, or domains, where such objects exist as such (as apples, events, abilities, etc.), which language merely reflects through the articulating agent or perspective (the medium whereby language is applied to the world).

Let us now return to the above case of the deaf person that learns to *hear* and experiences for the first time the subjective-experiential correlate of the neural transmission usually generated by the inner ear, but in this case by the interface of a designated bracelet and his/her wrist, and ask: how does he or she acquire the ability to associate the absolutely novel “stuff” they get exposed to for the first time in their lives with what we call hearing?

In other words, if the word ‘hear’ stands for ability, what is the ability this person acquires in the aforementioned process and how is it related to the hearing ability we – the non-deaf – possess, as we all seem to apply the same word – the same lingual entity – to our respective abilities?

---

<sup>76</sup> A pressing question rises here: does the articulation as such *generate* or merely enable/reflect/transmit? Is it a creative phenomenon or merely of an applicative, conducive nature [for something that emerges on another, antecedent level]? This is just another variant of the question we started with, and we should be ready to meet it in further forms as we proceed, as it is our ultimate pain which flares and radiates whenever we approach the abscess.

<sup>77</sup> And the “element” (nature/substance) of a lingual form is what we may call ‘meaningfication’, contrary to mere ‘signification’, which is a strictly formal event (*this is that* rather than *this is what it [indeed] is*).

<sup>78</sup> We *speak* as we *breath* (on the same categorical level in terms of existential attributes), only the domain in which we do so is not the biosphere but a story. A story. What exists *is* but can also *be told*. Everything as such can also *be told*. Why anything is as it is can be explicated by science. Why anything should be *told* is a question that is followed by an abyss.

<sup>79</sup> There is a certain blur here, but it isn’t our blur (isn’t part of our line of thought in this paper), so we’ll leave it for now.

As we've seen above, when we hear something, it is evident to us that we *hear* something, and it is equally clear to us what we are talking about when the words 'hear' or 'hearing' flash through our daily discourse.

The question is how does that deaf person "grow", having mastered the aforementioned technology, from living in a world where there is no hearing and no word for hearing to "knowing" what he/she means when he/she utters the word 'hear'?

And let us be clear: the question we pose regarding our protagonist is a twofold question: a) how does he/she *know* he/she is *hearing*? And b) how is he/she able to go on and *use/apply* the entire semantic field of that notion ([to] 'hear', 'hearing') as we do (in absolute correspondence/similarity/harmony, not conformity, to the way we do it)? These are two separate questions, although for Mr. Gawdat and all those who share his approach to language and language acquisition, as we've seen in the preface to this paper above, they are one and the same. Being software engineers, or at best scientists, the latter trust and rely on what they see (however many layers they dig in so as to see more). However, the distinction between these two questions or the phenomenal (or rather existential) complexity they stem from isn't observable. It springs from unease with certain interpretations of what can be seen.

One can rightfully suggest that there is nothing extraordinary or even interesting in these questions, since growing into new experiences and capabilities is an essential part of human biological nature, the obvious one being puberty, but the same principle applies to many other, more trivial cases like learning about an unfamiliar fruit or mechanism or many other aspects of life and the environment. The later examples illustrate that preliminary familiarity with a certain category makes it possible to apprehend additional, previously unfamiliar variants (specimen) belonging thereto. And while human sexual maturity is not a member in a class of objects like apples or particular machines are, it is certainly a feature – one of many features – of human biological nature, which is by definition an evolving phenomenon, expressed in the acquisition, through growth and development, of new modes of behavior, new capabilities that were absolutely absent yesterday (from the perspective of the evolving creature – couldn't walk, couldn't speak, couldn't see etc.) and are intuitive, coherent and fitting into the general environment (again – from the perspective of the evolving creature) today.

What should therefore be so special or challenging, so the argument goes, in acquiring a new sensual perspective – a new way/mode/realm of sensing/experiencing/interacting with the environment / with the world? The deaf person already sees, smells, understands, creates, gives and gets etc., and can already *speak* all of these. He/she is absorbing the world and communicating him/herself to the world in various, immensely subtle and complex ways. Why would he/she be precluded from, or stumble upon any particular hardship in, learning another way to experience – and communicate with – the world?

## VIII

As obvious and commonsensical as the objection/challenge at the end of the last chapter may seem to be, it has – again<sup>80</sup> – led us astray from our original pursuit/challenge.

For we haven't asked, and aren't interested in exploring or explaining, how a human being can – in principle or in practice – acquire new [natural] capabilities. We had, rather, set out to understand how a human being comes to subsist in, and perform as part of, a realm that is imbued with meaning; or in other words: to incorporate/exhibit cognizance; interaction with the environment in a non-performative manner.<sup>81</sup> And again, cognizance not in the sense of a lens or a camera – a capacity for abstract “snapshots” of something particular (Mr. Gawdat's ‘cat’ or ‘car’), but as “live” – ontological – storytelling; the ability (ability?) to inhere in a semantic field.

One may further object and argue that *the ability to apply language* – inhere in a semantic field (exhibit cognizance) – is as natural as (no “less” natural than) all the other aforementioned abilities and that the exceptionality, let alone enigma, attributed to its acquisition compared to the others, has nothing to support it and is not even fully coherent.

We disagree. Here's why.

As we've already noted above: “*we know that the person with the defect in the normal operation of his/her inner ears is undergoing a new experience, and we know that he/she is connecting that experience with an event whereby someone or something is generating some kind of message to him/her (or in general).*” “**But why/how**”, we further repeat and ask, “**all that corresponds to or sums-up [in]to (in the constitution of) the word ‘hear’ – [IN]TO THE ACQUISITION OF THE ABILITY TO GO ON AND APPLY A LINGUAL ENTITY? The same lingual entity all of us apply in the same circumstances. What or where or how is that peculiar “summing-up” proceeds in the subjective whereabouts of that person/agent?**”

To be able to see, walk, hear, touch, move stuff, interact with various objects in the environment – are, at their core, mere functions, the functions of objects in a given environment. As actions,<sup>82</sup> they consist in following rules, performing strictly mechanically,<sup>83</sup> which is why all of them can be modelled and duplicated/performed by

---

<sup>80</sup> It's so damn easy!

<sup>81</sup> I ‘see you’, ‘hear you’, ‘know [something]’, rather than act or perform in a certain way *as part of* something (have/perform a role). Conceptual performance, so to speak (a continuous set of *absorption/generation* without edges/fringes/boundaries nor an identifiable beginning or end).

<sup>82</sup> Units/items/aspects/instances of performance.

<sup>83</sup> However complex or ingenious the mechanism is.

programed machines.<sup>84</sup> But when we read or conceive of these words, let alone apply or learn to apply them in the first place, we've made a leap from everything that strictly speaking *happens* or *takes place* in any conceivable "domain".<sup>85</sup>

In sharp contrast to each and all of the aforementioned acts/capabilities, applying language, namely imbuing/infusing all that functional-performative riches with *meaning*, essentially "caters" for a state of affairs whereby things not only *proceed* by whatever inner reason or logic (physical, chemical, biological), "echoing" by that sheer flow whatever [essentially] abstract sense of causation or patterning or balance with nothing other than transparent "drive" and equilibrium to count on or be the case,<sup>86</sup> but where the *act* or *phase* or *cause* or *result* or any other part or aspect of that "drive" or equilibrium are not only "vectors",<sup>87</sup> not merely units in a plane defined – and governed – by abstract factors or coordinates, but also stages in a kind of pictographic<sup>88</sup> ontology, events and items with the fabric of *a world* rather than an abstract domain: the fabric of places rather than coordinates, of stones rather than groups of atoms / elementary particles, of brightness rather than flows of particles, of organisms rather than bunches of genes, of animals rather than organisms, of 'why?', 'where' and 'when?' rather than infinite 'now' and formless 'this/that'.<sup>89</sup>

This leap from performance in a realm (being in the world without language) to non-performative interaction with scenery, from sheer functionality to sheer narrative, appears to us as singular and therefore as a fundamental junction and potentially fruitful perspective for further insight.

---

<sup>84</sup> An important distinction is due here between applying these abilities in practice and thinking about them. The latter, we argue, is the result of learning to speak, whereby their "meaning" (as we read their names or think about them) does not pertain to their functional nature, e.g., the meaning of the word 'breath' does not pertain to the processes that take place in the physiological realm for what we call 'breathing' (and is evidently not a necessary condition for those processes to evolve and/or "do their thing").

<sup>85</sup> Grammar is not a domain, and neither is language.

<sup>86</sup> The tree grows and the beast preys, but nothing in their subsistence is more than "the stage before the next stage" (a finite cycle of – in the case of the beast, as the case of the tree is even simpler: 'I eat, I rest, I wonder around, I multiply' over and over again) and nothing in their environment (from their own perspective) is more than "where" that happens. In other words, they have no conception/heed of their environment as more than the sum total of the aforementioned stages of their totally predetermined, cyclic-performative existential journey. There are no 'places' or 'objects' there other than where I lay my head or seek for my food (alone or as part of my group) and/or vis-a-vis whom I fight or mate. And again, the word 'where' is applied here strictly metaphorically and due to the grammatical constraints of our human language. For us, any part of environment is [some]where. For the other inhabitants of the fauna, it is only: [now] I lay my head, [now] I escape, [now] I confront, [now] I seek/found food, [now] I court/mate etc. No 'where' or 'who' or 'now' etc.; only *where I am and what I do from within*, namely without any [further] perspective by means of which this sentence (or anything similar) will be applicable thereto.

<sup>87</sup> Sheer agents of forward-looking eventuation.

<sup>88</sup> Discernible from/through something like a [an external] perspective.

<sup>89</sup> The core ontological units. It may be argued that 'now' and 'this/that' are the only notions that retain the same meaning with and without language, resonating a timeless and spaceless domain (counterintuitively, 'this/that' does not assume space).

## IX

But is there<sup>90</sup> no equal *leap* between, say – the processes that take place in the physiological realm for what we call ‘breathing’ and the *ability* to breath or what the *breather* does upon breathing (the application of that capacity)?<sup>91</sup>

The answer is No. There is a crucial difference between the two.

As already noted, physiological processes, including those of the brain and the nervous system, are not the same as their performative analogues/results in the world (walking, seeing, pain etc.). ‘Walking’ is not something feet (or any other organ or physiological process) do, and neither is ‘seeing’ performed by the eyes or ‘breathing’ by the lungs alone or in unison with the throat and mouth or anything else. Indeed, humans, namely wholesome agents/creatures in the world, perform all these.

But in terms of strict performance there is no other than a conceptual boundary, no substantive<sup>92</sup> interface, between the underlying physiological bedrock and the performative wholesome agent – between the cells, tissues, fibers, muscles, tendons and the rest of it running at full throttle (physiology) and strolling in the park (agent), or between the whole light-retina-nervous-neural sequel at the bottom of seeing (physics/physiology) and the subjective experiential phenomenology that corresponds to that word (agent). In other words, what we call and perceive and can identify as a human agent or person or creature, with all its corresponding capacities, is not something that emerges **if not for the ability to say something, namely, apply language.**

Only the ability *to say something*, or more precisely – the acquisition of that ability, doesn’t have itself to rely on as it evolves.

In other words, although all these human capacities also appear to transcend the sum total of their physiological infrastructure, it is only a result of presupposing language, as referred to above.<sup>93</sup> It is only a result of being able to think about them and relate to them – represent them to ourselves retrospectively and in real-time – through language, in the first place.

This is where we reconnect with the intuitive common ground we ascribed to Chomsky and Wittgenstein at the beginning of this paper, whereby language is precursory to human applicative creativity and is more of a framework than a product, and on the basis of which we proposed our third conjecture, namely that language acquisition/generation underlies and inhibits everything else we can do as advanced biological creatures.

---

<sup>90</sup> Our Devil’s Advocate is persistent, which we only welcome.

<sup>91</sup> And the same goes for walking, seeing, touching etc., and also for various experiences like pain, rage, doubt or pleasure.

<sup>92</sup> Phenomenal, discernible.

<sup>93</sup> Namely as the source and generator (“infuser”) of meaning across the existential board.

Which leaves us with the original enigma regarding the acquisition of that singular capacity (***how does it all “sum-up” to something like a word?***) and brings us back to the postulate that acquiring the capacity – the trait – of applying language is both exceptional and pivotal for any attempt to make sense of the world and our place in it, let alone to the exploration of ‘meaning’.

It is now time to return to Wittgenstein and our hearing-impaired protagonist. It wasn’t the last intermediate challenge we’ll have to tackle<sup>94</sup> and not the last diversion from our main course, but we’re heading right.

X

Going through the novel experience by means of the novel technology, the deaf person gradually comes to identify structures in the new medium – what we call ‘sound’ (structured noise).

The flaw in the physiological mechanisms of our hero’s inner ear does not preclude him/her from learning to identify structures in an experiential medium he/she encounters for the very first time (had never *heard* before). The brain, as essentially a pattern-recognizing organ, will be disposed to seek for patterns in the neural transmissions it will receive once the technology is “on” and will start interpreting them as such in our protagonist’s conscious experience. The *fabric* of the experience will indeed be new to our hero, but this will not prevent the experiential differentiation/”volume” from building up.

This is where *correlation*<sup>95</sup> – that apparently trivial but actually as baffling and perplexing as any notion ever was – steps in.

Those initial [experiential] patterns, be they clearer, consistent and even agreeable and exciting over time (our hero’s aesthetic instincts are in place), would have been and remained no more than structured noise<sup>96</sup> if not for the emergent correlation with some particular event or source/factor with which they will gradually become associated in our hero’s “mind”, namely on every future encounter therewith.<sup>97</sup>

For some unknown reason, those evolving patterns also, as it were, “cling” – almost “merge” into without losing their distinct form – to a particular, sufficiently familiar event

---

<sup>94</sup> A reflective-analytical deviation we’ll have to endure.

<sup>95</sup> Or, in Mr. Gawdat’s language: pattern recognition.

<sup>96</sup> Sporadic structure of experiential qualia.

<sup>97</sup> The fact that our hero can, e.g., read (as one of several related cognitive capacities) the word which [in the community of speakers] corresponds to the experiential qualia he now “hears” (encounters for the first time) – a certain name, for example – has a crucial role in the cognitive buildup that presently occurs about him/her. More about that later, as we proceed.

(one of many) in the immediate whereabouts of our hero. A bond emerges between the two in our hero's perceptive (experiential) projection, which we call 'association'.

Indeed, association may seem like a trivial concept: a certain element brings to mind another element. A sound, a smell, a landscape, a gesture, 'bring to mind' some other stuff – a place, a person, an episode/encounter, a feeling, a memory etc., which “were the case or present or had occurred” on a previous encounter with [any of] the former (sound, smell, landscape, gesture etc.), and our propensity to link them *again* is the ability to *associate*. The former, as it were, *reminds* us of the latter in the sense of echoing them in real time.

But the question is what made that specific echo possible, namely what makes the latter events from the previous paragraph “stand out” upon further encounters with any of the former from all the other stuff that was the case in the original setting? What made them stick to the experience which now reminds us of them in the first place?

A possible, tangible answer would be: they were vivid or distinct or dominant enough in the original situation: a sincere, tender gesture awakes in me the agreeable/hearty feeling I felt when someone related to me in that or a similar way in the past, **because** back then I felt (consciously or not) that *that* gesture and not the room we sat in, the music that played or that it was Monday, affected me that way. And the same works for more abstract associations like the United States and capitalism, Germans and discipline or Scandinavians and blond hair.<sup>98</sup>

The point is that these vividness, distinctiveness or dominance seem to be the case in the original setting due to a “deeper” affinity that spurs them. It seems that tenderness and a hearty feeling, Germans and discipline or Scandinavians and blond hair are related *as such*, whereby we – the “operating system” at the heart of the setting (of reality) – spot it – spot that structural cohesion – the way it is.

---

<sup>98</sup> Indeed, I don't “feel” that the USA has something to do with capitalism or Germans with discipline when I come across these words for the first time. That sort of association emerges upon a gradual buildup of encounters, which is nevertheless governed by the same principle, namely coming across these notions in recurring discursive and/or experiential cases and circumstances already “imbued” with that linkage. The context of these further discursive or experiential encounters already incorporates or suggests that affinity which I similarly “feel” as a dominant feature thereof based on my [already established] capacity/propensity for association. The contextual consistency of these recurring encounters is the ground for my ability to “feel” the dominance of capitalism and/or discipline in regard to ‘USA’ or ‘Germans’ respectively, and I do so (am capable of doing so) the same way I feel “*that that gesture and not the room we sat in, the music that played or that it was Monday* (e.g. not that America is on the other side of the Atlantic or that Germans like beer and sausage), *affected me that way in the previous, original encounter*” in the more basic and personal examples in the main text above. It is crucial to notice, however, already here, that I need not be *instructed* in any way concerning this linkage and/or dominance, whereby any such instruction (“*America IS the epitome of capitalism*” or “*capitalism is at the heart of the American socio-economic structure*” or “*Germans are very disciplined*”) will most certainly not spur the associability discussed above. Exposure to and/or participation in relevant discussions or personal experience, on the other hand, will most certainly do. But we jump the gun here. Let's proceed.

But this merely brings us back to the original question: how do we spot reality *the way it is in the first place*? How do we sense/identify vividness, distinctiveness or dominance in the first place? **What** guides us thereto with virtually nothing to rest on?

It is not enough, at least to a philosopher or anyone with an inquisitive mind, to say so and “feel” it makes sense, because – again – it makes sense because and only because upon retrospective reflection it is grammatically coherent. But what made it grammatically coherent, or rather constituted/enacted that grammatical coherence, in the first place? – is the question we ask.

XI

Which brings us back again to our hero, as this is the exact question we asked about him/her above: how will the evolving experiential patterns of the new sensation – of what we as “accomplished products” call sound – “cling” to the particular event they will cling to (the lip movement of his/her interlocuter, his/her own attempt at talking, the clap of hands etc.) in the first place, and furthermore – *cling* thereto in a way that consummates in semantic sense,<sup>99</sup> namely in/with *a world* the agent at issue is *ipso facto* part of (a way that carries/introduces him/her into *a [whole] world* he/she wasn’t part of and suddenly becomes part of)?

Let us dwell on the matter a little more.

In consequence of that elemental, “structural”,<sup>100</sup> somehow anterior correlation with another event (a clap of hands, the movement of lips<sup>101</sup>), what started as a new [kind of] experience (**SOUND!!**), transforms into *something*; and not just something; it transforms categorically<sup>102</sup> from mere subjective (my own) experience to a **face** in the environment

---

<sup>99</sup> Moreover, the primary, “foundational” ability to correlate between sound and sense – between raw sensual experience (and when it’s new, it *is* raw) and emergent semantic coherence/meaning – is a more fundamental and challenging “case” than the ability to associate Scandinavians and blond hair, which has evolved after we have already mastered the ability to correlate. Although it seems that we relate the later in the present paragraph, we actually pursue the former. Consider the case of instinct: no one ever *teaches* us our instincts and the borderline between fear (instinct) and danger (semantic coherence) is very blurred. The leap between the two seems somewhat analogous to the leap between sound and sense we explore here.

<sup>100</sup> We don’t know the structure, or more precisely – we know nothing about it, so we use this notion here as a metaphor. The fact that a correlation emerges, as it were “out of nowhere” (as far as we are concerned) is intuitively interpreted by us as an echo of a structure. The sound we experience (in contrast to the soundwaves measured by the sound level meter) is associated with the moving lips although there is no conceivable way to *tie* them together, except grammar (which is, again, a part of the very mystery we are trying to solve).

<sup>101</sup> Both of which had never entailed the unique experience that now follows them for the deaf person.

<sup>102</sup> Mere experience – sound, visual formations, odors etc. – is not *something*. We, as it were, yield (surrender) to them, not perceive or interact with them. In real time they *are* everything. “Scenarios” that pervade our consciousness, on the other hand, namely real life, are never so. A scenario – any instance or situation I may find myself in in the course of my life – is by definition (in the way I perceive or wired to perceive it) a **part of reality**, of what ‘is’ in the broadest sense.

(the world)!<sup>103</sup> Moreover, that categorical shift *is* what transforms the environment into a *world*.<sup>104</sup>

The human agent's (actually the living creature's) innate disposition to spot correlation in his/her environment<sup>105</sup> serves as a "bridge" between experiential mishmash and coherent reality (semantic coherence/meaning – a face<sup>106</sup>). Between something where/whilst (in/of which) nothing is the case and 'Me in a/the World' (the grammatical 'I').

But what is that special disposition to "enact" correlation?

Let us elaborate, as meticulously as we can, on the event we discuss here:

- a) **Sound is born** where nothing of the kind – no *sound* – existed.
- b) **Sound takes the form of** a coherent "terrain", displaying typical, distinct characteristics (pitch, loudness, timbre, velocity etc.), namely stands out over mere noise.

---

<sup>103</sup> Only *faces* interact. Any/all other "things" are merely elements/components in – parts of – a broader thing/stuff/flow. Elements/parts of, or components in, something, do not *interact*, they comprise the something that *is*. To bring that point home, consider this: human individuals are not parts of / components in groups, societies, peoples or nations like physical, chemical or biochemical elements are parts of / components in an organism (individual wolves in their pack or other members of the fauna in their groups, is an interim case which is beside the point here). Indeed, the notion of 'a group' precedes the homo sapiens, which is born into a group and is a social animal by definition, but the [social] dynamics essential to and typical of its group are primarily dynamics between faces, notwithstanding all the physical, chemical and biological forces and drivers that govern and characterize it. 'This', 'that', 'me', 'you', 'us', 'they', 'them' etc., are more fundamental to the latter dynamics than all the aforementioned forces and drivers, e.g. the 'world' – in the sense of the ultimate domain – of humans has emerged – and *is* as such – by means of faces. Becoming a human, albeit as part of a community, is acquiring a face, already in early infancy, which henceforth only matures. Furthermore, once I had acquired a face, everything in my existential theatre has one (a face) – you and the screen onto which I currently type my thoughts and even the electric current or gravity (e.g. the 'universe'). As faces they are identical (categorically equal) if only because I am one. A face can only interact with faces. But the interaction (the only 'interaction' possible, let us remember) is not similar to the one between marbles, chemical elements or biological cells. The latter "interact" only in **MY** world – the world of faces, not in "theirs" (where they comprise). How do faces interact in a way that comprises (metaphorically speaking) "downwards" everything else, while everything else merely comprises (actually/virtually, not metaphorically) "upwards"? There's only so much we can explicate here. We hope to shed at least some light on that too as we proceed. For now, let us just remember that the emergence of a face is the emergence of semantic coherence – the "vessel" of everything (of the notion itself and what it means, not of our purposeful operability). For now, that should do.

<sup>104</sup> The notion of 'environment' also yields a subjective sense. The notion of a 'world' doesn't. 'A private world' is a metaphor whose literal sense is an oxymoron. If I can tell you about my world and you can understand it, it isn't private, even if I will never share it with anybody. I cannot "tell" you about my environment. The latter by definition has an edgeless "centrifugal" spread, so to speak. When we speak of *our* environment, we [actually] speak of our world.

<sup>105</sup> The original correlation is necessarily between instances of the same category. Indeed, we called 'hearing' an experience and the clap of hands – an event, but at the preliminary stage of the correlation the predominant element in both is their empirical nature. Both are predominantly *perceived*, and the fact that one is part of a scene and the other – abstract experience, does not sow a categorical rift between them.

<sup>106</sup> A perspective from "there" (the 'face' is the 'there').

- c) **Sound is associated with** phenomenological objects/instances in the environment (the sound of the word ‘car’ is the *object car*, the sound of the word ‘wow’ is a *reaction of appreciation/amazement*, the sound of rustling leaves is *the [background] voice of nature*, the sound of a snake’s hiss is a *source of danger* etc.).
- d) **Sound is associated with** different, accumulable and ultimately/evidently integrated, conceptual categories, e.g., one’s own sound and/or the sounds of nature and/or civilized environment and/or the form of a spoken word and/or a name and/or an expression of will/need/emotion/opinion etc.<sup>107</sup>

A deeper resonance than mere identification is evident here, it would seem, especially given the last two stages of the process.<sup>108</sup>

---

<sup>107</sup> These are all utterly different and moreover – reciprocally alien, entities (something like a name and something like a spoken word, considered on their own, are alien to each other), and calling them all ‘categories’ or bringing them or parts of them under a single category does not dissipate the division. It is obvious that they are all words, perceptions, ideas or mental acts etc., but it is the way they become all of that which is baffling and is sought to be accounted for (we *called* them so, but why and how is the riddle), or philosophy would have no justification whatsoever.

<sup>108</sup> Leading AI scientists, such as Yann LeCun, Chief Scientist of Meta AI, speak of the need to furnish AI systems with intricate representations of the physical world as a crucial challenge for upgrading their agentic performance by standards of human agency. Mr. LeCun differentiates between data-processing aimed at output generation and data-processing geared towards observational world-learning, namely towards “*learn[ing] how the world works by observing the world*” (LeCun, Yann, “Yann LeCun “Mathematical Obstacles on the Way to Human-Level AI.”, *YouTube*, uploaded by Joint Mathematics Meetings, 21 Mar. 2025, <https://www.youtube.com/watch?v=ETZfkkv6V7Y&t=2006s>). Mr. LeCun hypothesizes that a possible shift in AI training methodology (from sequential inference to inference by optimization) should result in a robust artificial imitation (model) of ‘learning from sensory input’, which he views as a core human capability for integrating into the real world and the way it works at the early stages of human development. More specifically, Mr. LeCun grounds the development of fundamental human properties such as persistent memory, planning, reasoning and self-guidance/control (including such faculties as tracking object permanence and biological motion, face tracking, instinctive acquaintance with natural kind categories, shape constancy and notions like solidity, rigidity, stability and support and even subtler phenomena such as gravity, inertia and conservation of momentum), which seem to lie at the core of human robust orientation in the immediate environment, in the gradual acquaintance therewith (with the world around us; the reality into which we are/were born) through observation. Since our ability to *act* during the period whereby we already display such subtle properties, namely at the early stages of infancy, is fairly limited, Mr. LeCun concludes that their preliminary formation in the architecture of the brain is produced by assimilating sensory input (rather than later stage learning through active behavior like cooking or playing chess, for example. That distinction will become crucial in what follows, given our contrasting approach to the role of *action* in the learning process). He summarizes his point in the lecture cited above as follows: “*the physical world, it turns out, is much more complicated than language*” (LeCun 14:21), whereby while learning to predict the next word or token by processing a great deal of text will only get the system so far, “introducing” it to the equally ample structures and integrated interfaces of the physical world should bring it closer to the “organic” way humans get by (subsist) in their habitat, pushing beyond current barriers to construct artificial applications featuring broad real-world deployment, like genuinely holistic home assistants and sufficiently dependable self-driving cars. But then again, even if Mr. LeCun’s alternative approach to AI training proves successful and a system that will have successfully imitated ‘learning from sensory input’ proceeds and adequately follows through the task of flying home from New York to Paris (from ordering plane tickets through the laptop computer to packing, ordering a cab, using the elevator to get down to the street and up to boarding the airplane, finding its proper seat and going through custom controls at Paris Charles de Gaulles Airport), how will he train it to follow through the fact

This complex “ability”,<sup>109</sup> as tacit and trivial as it may seem, is actually very extraordinary, or more correctly – the ability to *learn* it, in the sense of acquiring the practical disposition for such categorical versatility/elasticity, is. The ability to discern (absorb) and apply, through the same experiential “channel” (the structure of the sound of a specific spoken word – ‘car’ for example), categorically different (reciprocally alien) existential

---

that the sound of the word ‘car’ (or ‘plane tickets’ or ‘airport’) is the *object* car, plane tickets or airport, or that the sound of the word ‘wow’ (which the cab driver may say) is a *reaction* of appreciation/amazement, or that the sound of rustling leaves (in the street) is the *[background] voice* of nature or that the sound of a snake’s hiss (or a dog’s bark at the same street near the building it has just exited) is a *source* of danger, namely to follow through the distinction/s between objects, reactions, background voices, sources etc., as reciprocally alien (but nevertheless integrated) categories that are anything but components of – or parameters for executing – a task? A dog’s bark is a source of danger and also a background voice of nature. But can we ascertain precisely where this distinction lies? Is it one of *identification* or a wholly different kind of disposition? Will riding through all the possible types, pitches and timbres thereof and/or digesting everything there is to know about dogs and/or learning to identify one as a spatiotemporal object allow us to make it in practice? A source of danger (as well as reactions or background noises) may well be a relevant consideration in performing a task (within the immediate context of human task execution), but we haven’t learned it in that context or for such a purpose. In the case where I, as an eleven months infant, articulated my first word – ‘dog’ (in my case it was ‘tractor’) – because this was the first word I imbibed as an infant from the language spoken in my vicinity (whether, but not necessarily, in reference to the family dog or to the barking coming from outside the window or for any other reason) – I did so absent any form of instruction. That event had no underlying cause or future aim, except random mimicry, if only because back then the world was throwing at me a myriad of various data and I still couldn’t have wanted a pet or thought of becoming a dog trainer. I later learned to associate and apply that word in various contexts, including task execution, but my ability to absorb it as background noise and/or a source of danger had further emerged, and this is the crucial point, in equally random (from my standpoint) circumstances (prefiguring language games wherein I found myself interwoven with). Or how will he train it to follow through the fact that one’s own sound is also the sound of nature and/or of civilized environment and/or of the form of a spoken word and/or of a name and/or of an expression of will/need/emotion/opinion? The capacity for proficient navigation in physical space, marked by impeccable spatiotemporal orientation, does not appear to equate to (correlate with) a natural inclination towards (an inherent disposition for) categorical diversity, namely with *being* (inhering) in the world, as exemplified in my multifaceted capacity to absorb and apply my own sound, which prefigures human proficient navigation in physical space. Being able to grab a cab and taking the necessary precautions so that the summoned object will not hit you when it stops, does not amount to being in a world where a cab is *everything it is* [to us], which is an indeterminable loop we never “learn” in the real world in any conceivable sense of ‘learning’. We actually do not know how a ‘cab’ is or becomes a vehicle, a source of danger, the most common or convenient or dangerous or expensive means of transportation in a certain area, the neighbor’s way to make a living, the thing he/she gets into every morning or the name of a short story etc. etc. We can *speak* all of these, and we can use this notion for various purposes and relate to it in various ways, and the AI system – thanks to human formidable ingenuity and creativity – can display something similar (a wholly, categorically different thing from which no leap/path seems to exist to what *we do*), but we do not know how it emerges about us as such in the first place. We are never introduced to all these possibilities in any formal ways. We just participate in language games as part of our natural lives, which is also the way we had been initially “introduced” thereto (by nobody!). Hence, equipping oneself with all the fundamental world-oriented properties listed above is not a result of *learning through observation* as no observation would track the aforementioned categorical complexity. It turns out, we tentatively conclude, that language, or world-language interfaces, have more complexity to them than Mr. LeCun assumes. There’s something else going on in humans which our language – the language game we play in positively trying to understand and formally represent something – fails to articulate (doesn’t cover). Bear with us.

<sup>109</sup> A truly Super-ability, qualitatively speaking.

“acts”/instances in the world (modes of human natural behavior) is in fact unaccountable and unexplainable by any set of formal rules.<sup>110</sup>

---

<sup>110</sup> Which already hints or instructs us that it isn't a single channel, as we understand this term, or some archetypal structure. Consider this: I can move my arm *and* breath due to physical and physiological laws and resources. I can display fear *and* pleasure *and* bewilderment *and* despair by the same laws (evolution). Accordingly, complex artificial systems can perform very different and highly complicated tasks and achieve inspiring and even unexpected results through endless sets of 'if-then' orders that we have designed or built the platforms for their further "creative", autonomous design. But what can provide for or make viable the possibility of generating a *command*, on the one hand, and *naming*, on the other? Of *exclaiming* (a *car!*) or *warning* ("a *car!!!*") or *answering* (a 'car'), applying or absorbing the same formal symbols, namely the same form of experiential data/qualia? How do you get from learning a name to learning to *call* something by its name? Is there a set of formal or any *rules* that can formalize an adequate process to that effect? **1)** Logical rules (this is her name, now call her – is there a logical path from name to call? Her name vs. Call her – is there a logical path from the first 'her' to the second? One may suggest such a path: **a)** 'this is that' (a name); **b)** 'this is that' (a female – another name); **c)** 'this is that' (to identify, spot the correct name – another name) – but what produces a specific grammatical structure – a sentence – out of these three names? What accounts for the leap from 'this is a name' to 'this is a specific name'? What is the difference between to name and to name correctly or *identify*? How is that taught/wired? Isn't to *name* also to spot the *right/correct* name? What's the 'extra' in what we call *identifying* and how do we *implement* it? **d)** 'this is that' (to call, reach out to an identified object, namely something that was named correctly – another name); **e)** 'these are the cases/circumstances where you *reach out*' (another name) and so on – assuming that at the end of that complex process, which may include further interim stages, the machine will have transcended from 'this' to 'name' to 'her name' to 'calling her'. But did it? Even if we assume that all this is viable, will the machine be *calling her* or doing what we do – what any of us does – when we *call her*? Where is the logical step from 'her name' (a signifier) to 'call her' (reaching out to a signified)? A machine gun can be programmed to "reach out" (shoot) a target upon identification, but there's no 'reaching out' here whatsoever of the kind we discuss here or any 'reaching out' for that matter. Moreover, where is the preliminary [logical] step from 'this' to 'name' (her name)? Where is *she* in that process? Does the algorithm "incorporate" 'she' into 'her' the way we do in our worldly speech? Where is the step from 'this' to 'she'? And/or from 'she' (or from another 'this') to 'her'? And where is the 'identification for reaching out' that 'call' amounts to in our discourse, namely in the world? Can the machine *intend* to say '**namely in the world**' as I/we just did (please think it through before rushing to cry 'Yes'), namely as distinct from applying it as part of its answers to me on Gemini or ChatGPT, and is the sheer mimicry of answering/questioning it performs so astoundingly in cyberspace anything like discourse as a way of life (remember Wittgenstein)? Where is the 'will/need' to reach out to something (a very specific something) that *call* is imbued with rather than resonates from without to a human ear? Can these need/will also be programmed/wired into a machine? Indeed, we saw that complex algorithms can display as it were "sovereign decisions" which even deviate from the targets they were set to achieve (an ability to "deceive"), displaying a capacity to "summon" their abilities to achieve something like "their own" goals, but these are "cracks" or lacunas in the software (in prioritization) and/since there is nothing like 'she' or 'her' or 'call' in these leaps of prioritization or in anything else which was the case along the way (it isn't a 'face' that reaches out, and it doesn't *reach out*). Such a machine may indeed destroy the planet along a similar process, abiding to the same principles, but it is not 'the world' as we perceive and conceive of that notion it will be destroying, nor will it be 'destroying' anything the same way as the autonomous car doesn't 'save' human life when it "decides/prefers" to crash into a road barrier rather than hit a pedestrian (it isn't a pedestrian "for it" and 'rather' is a non-entity in its realm)). Hence, no logical rules to go from 'her name' to 'call her'. **2)** Arithmetical or other mathematical rules? No. **3)** Any other scientific rules? No. **4) Grammatical** rules? Here we must beg pardon, since these are the very thing that needs to be explained and accounted for rather than mimicked in one way or another. When a toddler learns to speak, he/she doesn't learn to *display* [anything], namely, to perform/execute an *interpretable* act or exercise an *interpretable* disposition or achieve an *interpretable* outcome (what we call, essentially from without and in point of fact very sloppily, 'to communicate' in the sense of to act so as to be "understood"). He/she learns to exert/wield – *emanate* – meaning on [logical] par with the exertion of physical force. He/she learns to apply grammar. Not proper grammar (there isn't such a thing as improper grammar; it's not disprovable). *Grammar* as such. Grammar

We can always compare our case study to a child learning what an apple is (“*this is an apple, you want an apple, baby? Here, look, an apple! Want an apple? Go ahead, taste it!*.”), but the more appropriate analogy would be the child’s learning what ‘sound’ or what ‘an object’ is in the first place. For upon learning what an apple is, the child has already learned what an object<sup>111</sup> or food or mommy or ‘*I want*’ is/are, and based on his/her disposition to absorb and apply categorical elasticity he/she will further “learn” – acquire the ability to behave to the effect – that that specific *object that I want and is presented by mommy is an apple* (not: *is an apple*, but essentially *is an apple*, as all three words are one – a single “quanta” – in/during the “learning process”).<sup>112</sup>

‘This/that is an apple’ is behavior, not identification. Not in the traditional behaviorist sense (displaying/featuring/exhibiting). In the sense of a sovereign act. Bear with us.

In contrast [to learning what an apple is], the child that learns what ‘sound’ or ‘object’ (*something*) are in the first place, has no antecedent conceptual layers to rely on, since ‘sound’ or ‘something’ are themselves foundational.<sup>113</sup>

---

– the only set of rules applicable to human interaction (as a species) that has nothing to do with [the regulation of] humans as social creatures. A set of rules that regulates the human very existence as a species (rather than what they do or how they live given their diverse, unique capabilities). Nay, that regulates *the picture* of its existence [from without]. What are they beyond or apart from the fact that they are *the right thing* or *the right* transmitter? ‘Right’ in what sense?

<sup>111</sup> Notwithstanding the fact that he is years away from acquaintance with that term as such.

<sup>112</sup> The artificial algorithm has learned what *cat* is. But can it learn, doing what it does whilst it ‘learns’, what ‘object’ is? Or what ‘why’ is? Or the difference between ‘effect’ and ‘result’, or more precisely between the respective semantic fields of these two notions? As all these (including ‘object’ and ‘why’ and effectively all the other *words* comprising human language) are not words but semantic fields: categorically elastic possibilities – “keys” – humans apply in their purposeful existence (existence, not speech), namely for “opening” the otherwise sealed/non-existent sense of their immediate existence, rather than a total number of formal variations of anything. Moreover, can we even articulate, regarding that latter alternative, what they (these words) are variations of? A concept? A notion? An existential category? But these, again, are semantic fields which *identification* doesn’t “open”, whereby we stumble upon an infinite regress. Indeed, Gemini, ChatGPT and their likes do display similar (not really similar, more of a specular) elasticity in the way they interact (not communicate) with us, but they acquire that formal variability (rather than categorical elasticity) by essentially statistical analysis (prioritization) of a closed set (enormous as it may be) of formal variations, having nothing but analytical power to build on and implement. They too, like Google’s pioneer algorithm that identified a ‘cat’, essentially prioritize “pictures”. A question or an answer or a suggestion or sympathy or encouragement on their behalf are nevertheless pictures and/or sequences of pictures they present (place in various orders) by way of prioritization given an established (very complex but essentially settled) set of formal patterns they had been “bombarded” with and absorbed as such. “Throwing back” at us that set of formal patterns is not *speaking*, nor does “cracking” that set amount to picking up *what something is*. Although the algorithm had “encountered” our prompts in the application as such sets, ***we didn’t articulate them as such***. We did something completely different. *We spoke the world*. We’d done so (were able to do so) subsequent to having gone through a path and a realm altogether alien to the algorithm’s very nature. We spoke the world. The algorithm did something completely different. Bear with us. We shall clarify, demonstrate and maybe even *prove* what we mean as we proceed. We so much *want* to do so. We play a certain language game in the world. The algorithm is a part of our language game. Like any unconventional weapon. We see stuff. It *doesn’t play*. Bear with us.

<sup>113</sup> Or so it currently appears, at the very least. Bear with us.

So how does that physiologically robust infant *learn* that sound is sound, and something is something? How does he/she acquire the ability to discern (so as to apply) form and structure in the blizzard of initially formless phenomenal variety in which he/she finds him/herself throughout that foundational period of their development?

## XII

Somewhere along the way, that child (our second hero and protagonist) will be able to absorb and generate specific, distinct sounds relating to various aspects of his/her environment.

And now we ask: what transforms the sounds that our hero “hears” for the first time in his/her life, from nonsensical, albeit structured, qualia into a symbol of a coherent object – into a translucent medium of semantic coherence?

Moreover, what transforms it not merely into a coherent perceptual entity (something that I “understand”) but into a “weapon” for carving meaning all-around?<sup>114</sup>

To emphasize the point of that fundamental question we shall remind ourselves the initial divergence regarding the meaning of a word:<sup>115</sup>

- **The intentional**, sort of idealist “mental-qualia” conception, on the one hand, according to which the meaning of a word – the ground for (source of) the semantic transparency of what we say – is *what I mean (intend to say)* by it – a generative mental entity/process that corresponds to objective/subjective phenomenology (objective – a tree; subjective – fatigue) and yields intentional exchange/communication among advanced homo sapiens’ individuals. The “picture” of the object in my mind – the “object”<sup>116</sup> of my intention<sup>117</sup> – is its meaning. Meaning resides within this ephemeral “picture” (understood as its substance), having been established via association.<sup>118</sup>

---

<sup>114</sup> It can be easily assumed that phrasing the question in such a way is inconsistent with our non-instrumental conception of language, whereby ‘a weapon for carving meaning [all-around]’ can be perceived as merely a metaphor for a tool that facilitates interaction with the environment, no different from my hands or any other object [which is not a part of me but which] I can use. This is a fair criticism, whereby we least of all should claim to be exempted from the constraints and inadequacies of human language in philosophical matters/discourse. There is, however, more to that metaphor than metaphor and no inconsistency in our course. Please bear with us. We are approaching the heart of our case.

<sup>115</sup> Further apologies for repeating ourselves, but this is important.

<sup>116</sup> The cognitive trace of its real-world referent.

<sup>117</sup> As well as mere, seemingly passive (but actually not passive at all), conception thereof.

<sup>118</sup> Already here we can get a sense of the difficulty of capturing and formulating the dynamics of that model in language, as ‘picture’, ‘ephemeral’, ‘intention’ and ‘association’ are all used metaphorically as explanatory vessels, being themselves part of what has to be explained. We hope to shed further light on the reasons for that as we proceed.

- **The Wittgensteinian** ‘language games’ conception, on the other hand, according to which what gives meaning to our living application of language is the practical/functional context thereof in any particular case; not its generative, ephemeral (mental) substance, but the concrete, practical circumstance in which it – the word, the phrase, the sentence<sup>119</sup> – is articulated/pronounced.<sup>120</sup> According to that conception, ‘*come!*’ is born not by understanding what ordering someone to attend or follow you is, but by accommodating ourselves to *ordering/obeying* as forms of behavior, which was **not** in itself (the accommodation) made possible by “understanding” what facilitating attendance (on either “side”) is, but by executing an inherent disposition that conforms to the way our existential setting is.<sup>121</sup> And how does that *accommodation* happen, namely made possible in real time and acquires a particular behavioral form? That is a question concerning the species, not any particular variant thereof. In other words, it is not a question regarding an event in real-time. The agent (child) indeed heard that particular word in particular behavioral circumstances. But he/she didn’t “grasp” an ephemeral “intention”<sup>122</sup> and connected it with any parallel sensual perception. He/she didn’t attach symbol to substance. He/she displayed (has grown to display like a stalk grows to

---

<sup>119</sup> A word can be a part/component of a sentence, but a single word can also amount to an entire (wholesome) act of speech. For instance, we can apply the word ‘come’ in a certain way (tone) that will amount to ‘please come with me’ etc.

<sup>120</sup> The real-time semantic coherence of ‘fear’ stems directly from the fact that I inhabit a dangerous realm rather than from *identification* of danger. Danger isn’t – ever – identified. It is ontologically endorsed – by my cat or an insect and by a homo sapiens alike. We don’t have a word for that elemental endorsement, but my whole train of thought as I sit in a room knowing there is a killer behind the door is fundamentally identical to the insect’s escape from its predator. There is a qualitative difference between the two, but it isn’t different – as a *difference* – from the difference between my and my cat’s ability to grab an object. The fear of a car crash, a financial crisis, misspeaking, breaking up with a loved one or death, namely what it is to be afraid in each of these instances, are primarily **participations** in the way the world is, not **assessments** thereof or the form it might take (they are reactions *in the* world, not *to the* world). It isn’t the “value” of the situation in our perception, but our very response to the flow of events (our note or phrase in the symphony). In that sense it is indeed a means and a tool (our reaction is purposeful from our real-time standpoint), but for *existing* rather than *dealing with something* (solving stuff of fulfilling assignments). There is no standpoint/perspective from which we can relate to our fear and the world in which we fear as a single, comprehensive playground or gameboard. We indeed identify sources of danger and experience fear, but we cannot relate – are absolutely deprived of any capacity for relating – the content of that last sentence (“*we indeed identify sources of danger and experience fear*”) other than a sentence – as narrative (something we absorb as an accomplished product) rather than reality (something we can impact). We do solve problems, overcome obstacles and meet challenges, but not in order to excel or succeed. Actually, *in order to* (for the purpose of)... nothing (ultimately – as a narrative). We just [primarily] *are* [the way we are]. Since ‘ontological purpose’ (beyond survival as a metaphysical constant) is an oxymoron, we **cannot** know how to duplicate what we are. Ultimately, it is the *world* – the ultimate realm – which is so “deep” and elaborate, not we or anything else “in” it.

<sup>121</sup> This is where the ontological and the empirical converge: we integrate in an impenetrable (inaccessible, unperceivable) way into the most straightforward, accessible reality.

<sup>122</sup> Which is itself – the notion of “intention” – a metaphor (or whatever) rather than an experiential state. We have no access to our intentions. They are colors (of our acts or states) we have no palette for. Higher order colors. An act or state is possible on the individual level. Intention is possible only where there’s *two* (namely more than an individual) but is in neither of them as such (as individuals/agents).

display a flower) an ability in the world. It is modes of behavior we “suck” into our genetic dispositions,<sup>123</sup> not ideas.<sup>124</sup>

From now on, we shall shift inadvertently between our two heroes – the hearing-impaired person and the small child / infant. This should not undermine our progress, since for all relevant purposes they are ultimately one and the same.

Applying the first – mental qualia – conception to our case-study, instructs us that the “combination” of the growing (accumulating) acquaintance of our hero with the new experiential fabric (what we call ‘sound’) together with his/her anterior familiarity with some specific category (in this case: naming), enabled him/her to associate the two.<sup>125</sup> That association, so the argument goes, and specifically the ability to recapture/reproduce (reapply) it upon recurring encounters with the same structure, is the very name/*meaning* of the increasingly familiar new kind of structured experiential fabric and will be fixated/”encoded” *as such*<sup>126</sup> in our hero’s mind (the association *is* – or will have become – the name).<sup>127</sup> In consequence of that event/process of welcoming a new member into an established category, two things are contended to have happened: he/she *knows* what *that is*, and he/she can communicate it and/or in its regard.

Let us, however, examine from a little closer the argumentative logic behind this explanation. The rationale (*explanatory factor*) offered to us consists in switching (substituting) the term ‘association’ with an allegedly more substratal (and actually wholly abstract) sense of ‘connection/combination’. According to that rationale, the association emerges in the mind of our protagonist due to (by way of an “act”<sup>128</sup> of)

---

<sup>123</sup> Here we are merely more knowledgeable than Wittgenstein, certainly not more insightful.

<sup>124</sup> Indeed, if learning to apply semantic coherence is mere ‘accommodation’ and a mere display of a possibility in the world like the stalk’s “ability” to display a flower, where is the *interaction*, the *story*, the *face* which supposedly distinguishes us from mere parts or components in a greater scheme as described above? This evident discrepancy/inconsistency hadn’t eluded us and will be resolved as we proceed. For now, we can just note that the *world* is the story, not we. Bear with us.

<sup>125</sup> Whether by being indicated of such a connection by an instructor, as in the case of our deaf hero, whereby as part of his/her training with Dr. Eagleman’s device, upon uttering the word ‘John’ the instructor wrote that word on a blackboard and pointed to himself and then wrote the word ‘name’, or without any such indication, as in the case of a baby’s learning to say ‘*mama*’.

<sup>126</sup> A name and a specific name at once. The increasing familiarity had found its base (“home”) in a specific category.

<sup>127</sup> Other kind of associations will have become specific variants of other [distinct] categories, or so the story goes.

<sup>128</sup> If it is only and merely the brain, as a kind of software, which performs and/or produces, on its own level, all the relations (connections, combinations, associations etc.), in/as what Mr. Joscha Bach eloquently calls ‘a self-reinforcing causal pattern’ (Bach, Joscha, “Why Consciousness is software, and software is spirit.” *YouTube*, uploaded by The Institute of Arts and Ideas, 27 Mar. 2025, [https://www.youtube.com/watch?v=E361FZ\\_50oo](https://www.youtube.com/watch?v=E361FZ_50oo)), what do we (the human agent) do? The requisite answer is nothing! – as we are, according to that view, no more than a simulacrum – the plot of a movie “experienced”, so to speak, from within a flat (ontologically absent) screen, while the brain is the [ontological] substrate/mechanism producing the simulation **and** the “real” sense thereof, namely the dream **and** the “sense” of the dreamer, or the plot of the novel **and** the heroes “thinking”, for all ends and purposes, they are entities in and parts of a world, at the same time (in contrast to how we would normally

combining<sup>129</sup> the familiarized experiential fabric with the familiar category. But what is the sense of ‘combining’ we are offered here – an abstract/formal act of connecting – other than a description – a declarative assertion – regarding a (some) relation between two separate things?<sup>130</sup> And how does such a description explain the way a *specific*<sup>131</sup> (end product) relation between two previously unrelated items has been established? In other

---

conceive of someone reading a book). In other words, the same way a character in a movie or a novel *doesn't do anything*, residing merely metaphorically on a flat screen or along the browsed paper sheets, neither do we, residing in equally “flat” virtuality. In the ensuing discussion we intent to refute that conception, not merely by highlighting its logical inconsistencies but by demonstrating its lack of practical grounding. Bear with us.

<sup>129</sup> Asserting a formal link between two entities. Terms like 'connection', 'combination', 'association' and other synonyms for enacting a relation between two separate entities (“somethings”) belong to a broader semantic field pertaining to relationships and integration and embody/resonate nuanced emphases of linkage, contiguity, amalgamation, proximity/co-occurrence, dependence/interdependence, organization/structure etc. While ‘association’ (to associate) denotes a mental link, ‘connection’ and/or ‘combination’ imply a plainer, more substratal function of ‘bringing things together’, inviting their interpretation as the constitutive force in the emergence of a higher-order relationship, e.g., as the “building blocks” of the resulting mental linkage/amalgamation. Given the premises of the ‘Intentional’ conception of semantic coherence, ‘association’ may emerge in only one of two ways: 1) a connection/combination occurring at the mental realm, or 2) a connection/combination arising at a more fundamental substrate/level, with resultant effects on the mental plane. The first option is a tautology (mental linking consists in mental linking). In what follows we hope to demonstrate that the second option is as tautological and as unviable in accounting for the emergence of semantic coherence. We further aim to demonstrate the absence of a foundational, formal function of ‘putting things together’ at the core of semantic coherence (within the human mind), which would permit or enable higher-order forms of relatedness and, ultimately, our [sense of] subsistence in an inherently coherent and integrated environment. We intend to show that Mr. Bach’s notion of ‘self-reinforcing causal pattern’, as presented in the interview cited above, namely as consisting in the following scheme: theoretical premise: consciousness is a behavior of some kind, namely a manipulation of arbitrary systems; goal: finding a representation for describing arbitrary systems; practical premise: in computer science it is possible to describe any function using state transitions which can be further described using computational operators; means/task: looking at discernable differences, namely at state transitions, and finding a resolution for all of them to matter – the “key” of coherence for the entire flow or the bedrock for its comprehensive semantic integration; outcome: using the resolution applicable to all the transitions between sets of discernable differences to describe arbitrary behavior – does not account for actual language absorption/application, whereby the latter – a behavior indeed – does not exist in any flat representation or any kind of ether whatsoever, but in our actual behavior in the real world. The ability to *describe*, we argue and shall further demonstrate (and/or the possibility to employ the description to whatever purpose), is not equivalent to explicating – to being able to see into the core of – what something *is*. Bear with us.

<sup>130</sup> An utterly formal ability to enact an utterly formal relation. A kind of internal assertion/affirmation regarding a relation rather than any “logic” resulting in the formation of a new entity, set or process. Like simply stating: ‘you create a relation between two separate things’, assuming that everything which that statement should explain is already clear. A genuine combination is a unified entity formed by related elements. The combined elements are joined or merged to form a single entity or set, assuming an antecedent relation therebetween (a horizon of interaction/relatedness). The newly established entity is a consequence of that joining/merging. A combination enacts the transition from origin to result and while it *implies* a relationship between the combined elements, its primary meaning goes down to the very *formation* of the ensuing result. Lacking any antecedent relation between the familiarized experiential fabric and the familiar category in our case, the “combination” referred to in the main text above amounts to no *formation* whatsoever and seems to boil down to mere *identification* of some link or proximity, namely a mere statement about an existing or perceived link between separate entities. But since no such link exists or is perceived by our hero prior to the purported event, there’s no identification here either.

<sup>131</sup> Categorically specific (in this case: naming) rather than abstract.

words, how an explanation such as ‘that [specific] relation is created by you [specifically] creating that [specific] relation’<sup>132</sup> is expected to explain the emergence of the relation at issue or a relation at all? What accounts for – goes down to the origin of – the emergent specificity? Does that explanation provide anything other than a circular argument contingent on the tacit assumption that the meaning of ‘that specific relation’ is already known? It is patently evident that it does not. Doesn't that explanation merely restate the unknown as known? Evidently so.

In fact, while alleging to explain that *that [specific] relation is created by you [specifically] creating that [specific] relation*, what such an explanation actually says is that *that [specific] relation is created by you abstractly linking **this and that***,<sup>133</sup> tacitly assuming that ‘this’ is a variant and ‘that’ is a category which is also the category to which this variant belongs!<sup>134</sup> – without providing any basis whatsoever for the validity – or the constitution – of that tacit assumption from the standpoint of the “learner”.

The *specificity* is evident to **us** “explorers” into language or consciousness and/or participants in the discussion or seminar thereabout, who already “swim” in semantic coherence, namely are able to hear properly and apply the related language. In Joscha Bach’s terminology, it is evident to us as we are already part of the simulation and the entire set of relevant simulacra, as well as its bedrock in our brains, have already grown wild. But how do our two protagonists, who had not yet been fully “simulated”, “plunge” therein **from their own rather than our** perspective? How do they “grow” volume (such as growing wings) in what shall have become a flat realm? In a dimension reduced to flatness, how does a structure (the semblance of volumetric growth, like the unfurling of wings) develop in the first place? Finally, how does a dream acquire or impart the sense of not being a dream? Sensing an environment (experiencing oneself as part of an environment) and determining/judging/asserting that something is a dream involves a categorical leap which must be resolved – anticipatively (prospectively), not reflectively (retrospectively) – for the latter to apply to the former. For attaining that goal, words and descriptions will not do (e.g., it is what we already “know” it is), nor will any reference to optical illusions or other errors in the real world or purported analogies such as “killing” characters in a computer game, which do not involve a similar leap (for all intents and purposes, we are *playing*, not killing, in that last analogy; we simulate, not being simulated; the entire set/process unfolds in *our* world). A true argument is needed.

And if that is the “explanation” we’ve been provided with by the intentional conception of the meaning of speech (of semantic coherence) to that very **meaning**, namely that it

---

<sup>132</sup> ‘A combination is created by you combining stuff’, or ‘you connect things by putting them together’, which are bare tautologies.

<sup>133</sup> Or in Mr. Bach’s language: your brain ‘links this and that’ while you dream specificity.

<sup>134</sup> That tacit assumption underlies the very utterance – the very ability to say – “*that [specific] relation*”, effectively stripping the remainder of the sentence – “*by you abstractly linking this and that*” – of semantic value (rendering the remainder of the sentence semantically superfluous).

ultimately emerges (is generated) by an *abstract ability to link **this and that***, how can the latter account for the categorical particularity of *naming* as one of many other categorically different acts to have been supposedly learnt in the same formal way? How can this flat, one-dimensional, absolutely stiff mode of alleged apprehension account for a complex, “multi-dimensional“ capacity to learn, namely for the multifarious, compound capacity for learning *to name* AND *to express one’s will or need* AND *to spot danger* as well as many other disparate (categorically diverse) modes of becoming acquainted with the world? It isn't a world of 'this' and/or 'that' which is meaningful and whose meaning we strive to uncover, but a world of 'mommy', 'No!', 'give me' and 'why?'.

It should by now be obvious why that model fails to stand/deliver, as it is utterly silent regarding the deaf person’s very ability to *link* the new experiential fabric (being what it is) and the familiar category (being what it is), or rather since it actually presupposes it. The new experiential fabric does not contain, as it were, “receptors”<sup>135</sup> for something like a ‘name’ and the sense of *linking* which will “endow” it with a ‘name’<sup>136</sup> – transform it from something utterly unfamiliar into something absolutely familiar<sup>137</sup> – it is in itself an absolute riddle. Its emergent association with that particular category (rather than with something at all) *is* that which should be explained rather than merely *stated/affirmed*. Even if in the midst of the situation a third person writes the specific name which the deaf person now “hears” for the first time and further writes “this is my/your name”, the question still stands of how (regarding the “way”) our hero will be able to link *what he/she hears* (the utterly novel experience) with what he/she already knows a name is (between what he/she *hears* and what he/she *reads*, as the latter already incorporates that very linkage in *seeing* the written word and “taking it” as a ‘name’ and cannot, therefore, concomitantly account for the same transformation in regard to the auditory qualia)! our hero’s accomplished ability to read and write and his/her “understanding” of what a name is and his/her understanding that the third person is attempting to indicate a relation between his/her new experience and *naming* (this is that, you see?) still do not and cannot by themselves explain or account for – cannot by themselves pinpoint the real-time dynamics of – **his/her own intrinsic ability to make the leap** from experience to semantic sense (to meaning). His/her [already] established capacities cannot account for and explain an emergence (the emergence of something) that is at the core of their own original emergence!

---

<sup>135</sup> And if it does, it is *them* we should explicate and account for rather than assume.

<sup>136</sup> *Name* the name, so to speak.

<sup>137</sup> And the sense of transformation we are talking about here is indeed absolute (from nothing to something)! Furthermore, the emergent ‘familiarity’ is not of the kind I display in my ability to recognize a person or a plant or any other kind of object as part of the normal flow of my life (where everything we presently inquire into the constitution of which is already established), but of a truly primary emergence whence something is emerging out of nothing.

At the bottom line, being able to display **assumes** what we call ‘understanding’ rather than accounts for it, and the burden of getting to the bottom of that interface (form and meaning) remains.

There is a whole lot of cognitive obfuscation in the way of *seeing* into the logic of the above point, since it purports to dissect and pull the blanket under our most fundamental practices of being what we are. But it is logically sound. It is sound because it unveils and underscores the very heart of the puzzle we are confronted with.

The point gains further traction when we turn to our second – infant – protagonist's analogous, utterly raw, "first instance" learning scenario. The infant does not know what naming is before he/she acquires the ability to apply the first names he/she applies (‘mommy’ or other sounds he/she utters as names, namely as part of the language game of *naming*).<sup>138</sup> So how does he/she nevertheless succeed in performing the/his/her earliest act/s of naming, having nothing to combine/associate it with? How did I – Daniel Drabkin – utter ‘tractor’ as the first word I have ever articulated, having nothing other than diverse speech occurring around me (not necessarily addressed to me), my sensory faculties and my brain? Have I done what Google’s nascent algorithm succeeded to do without any prior instruction, as provided in Mr. Gawdat’s interview cited above? On the face of it, both the infant me and the algorithm engaged with respective data streams devoid of any instruction.<sup>139</sup> The algorithm spotted a structure (*metaphorically, one of*

---

<sup>138</sup> It is a language game and not an algorithmic process if only because the names will change over the same algorithm while the world shall not. The metamorphic and kaleidoscopic nature of naming consist in rules that defy any sense of schematics. The game is always in the arena (in the world), not in the software or on the board. There is no computer game without a human *gamer*, and when the computer plays with itself, it doesn’t *play*. It [essentially] competes (faces/meets a challenge). Like the gene. In the human realm every competition is an ontological game. In cyberspace every game is ontological competition. Naming is not a competition or a challenge. Absolutely not.

<sup>139</sup> I certainly did not construe my nanny’s utterances that the tractor, then working outside our window, would arrive and consume me should I fail to eat, as any form of instruction. Indeed, what might have at first appeared to me as formless vocalizations gradually acquired form and structure that I could perceive and mimic. But in doing so I certainly did not “**find the essence of what makes a cat a cat or a tractor a tractor**” as Mr. Gawdat asserted regarding the algorithm’s pivotal *identification* in the interview cited above. The words and phrases I gradually acquired the ability to apply through my comprehensive apparatus of speech generation were anything but essences and neither was my articulation thereof an identification of any essence whatsoever. My brain has certainly identified an abstract form as a prerequisite to my agentic ability to say ‘tractor’, yet while we can coherently engage with each of these layers in isolation, we cannot assert/articulate anything about the shift, transition or integration between them, let alone perceive any identification of essence within their confines. In terms of what we can say about my (not *these*, but essentially *my*) primary words and phrases they are/were nothing (I knew nothing of tractors when I first uttered that word, and saying that I identified the abstract form of a tractor is saying nothing about [the infant] me but about those who heard my utterance), but in terms of what they meant **to me** (what uttering ‘tractor’ or for that matter – *mama*, meant to me rather than its interpretation in the community of speakers / language animals) they are/were everything. And that – the meta-categorical distance between the community of speakers and the emergence of the individual speaker – is an unbridgeable gap. Nor were they tools in any conceivable, or for that matter – reversible and reproducible (by purportedly reducing them or their aggregate to first principles or original schematics), sense of a toolkit. They were acts (we actually don’t have a word for what they were) in a categorically inaccessible domain, namely a domain that we – as stains on the canvass of reality (we think we are figures on that canvass, but we are stains which only an

*them came back and said: "I found something"* (Gawdat 08:00)). Did I do too? It in its world, me in *the* world, but essentially abiding to the same underlying principle, performing the same task, achieving the same accomplishment? I said 'tractor'. Did it **say** 'cat'? We know what the algorithm did. It successfully identified a complex pattern – a capacity that will be significantly improved in subsequent versions thereof. But what did I do at the inaugural phase (*as* the inception) of an internally consistent, predetermined development that isn't predicated on discrete "version" upgrades of myself?<sup>140</sup> We know what the algorithm did because we (Google software engineers) designed the experiment. We built a pattern recognition algorithm and ran it on a stream of patterns. But are we so sure about the existential 'experiment' whereby a person begins to speak? Does a human newborn also run on "stuff" and pronounce his/her discoveries of what's already there?<sup>141</sup> Are the commonalities between the two intrinsic or solely superficial, and potentially illusory?

---

unattainable perspective "sees" as figures) – are categorically precluded from accessing. The following analogy may help to clarify that point: think of the emergence of David out of the marble block Michelangelo was carving. The artist had indeed had a prior image of the sculpture, but the emergence of the final piece was not equivalent to the materialization of a model. Suppose that a would be highly sophisticated artificial sculptor could theoretically produce an equally awe-inspiring or even superior sculpture out of an unhewn block of marble. The artificial "work in process", however masterful in orchestrating a myriad of possibilities extending to modalities humans neither envisioned nor could have envisioned, is nevertheless an entirely mechanistic process. Form is incrementally imparted to – and assumed by – the block through a succession of deliberate steps/stages – distinct, discernible transformations in the continuum that holds both the nascent block and the final sculpture, and in this sense, Michelangelo's work in process is no different. There is the abstract model, its material actualization and a cohesive algorithm that binds them and orchestrates the latter in deference to the former. But the **emergence** of the ultimate form – the emergence of *David* – essentially transcends all these. The *emergence* of form out of nothing (*David is* categorically transcendent to the block; the block *is* nothing to David) is an "event" beyond any possibility of qualification or quantification (we have no word for *that* too; 'transformation' or 'metamorphosis' are essentially retrospective classifications, not real-time depictions). When David is already there (once it has been completed) or just before, there is the model and there is the marble and the algorithmic chiseling, but the *becoming* thereof (*of David*) evades us – all these three modes of engagement – altogether. No instance of *behold!* is attributable thereto or extractable therefrom. It is irreducible to the model or the marble or their linkage in the scheme. That *becoming* is analogous to the words and phrases I gradually acquired the ability to apply through my comprehensive apparatus of speech generation. We have access to the model (the word 'tractor' which my nanny articulated in my presence and we all "understand"; the real-time "glue" of the community of speakers) and to the "stain"/block of marble (my entire phonatory system) and to the cohesive algorithm that "binds" them and orchestrates the latter in deference to the former (the neural computation unfolding in my brain; my cerebral electrochemistry / neural signaling), but we cannot penetrate the formation (it is "nowhere", as far as we can tell). We can engage with (appreciate) the model (this is what we see looking at David or *do* performing amid the community of speakers) but touch or embrace only the marble (embracing David, we are, in essence, engaging solely with a block of marble, much as probing "the speaker" reveals nothing beyond the phonatory system) or inquire into and/or play with only bare algorithmics/computation. The "place" where the marble had become David – where the trinity emerges as unity – is not for us to see or apprehend. This is also where I uttered my first word. Had we been mathematicians, we would probably ad 'QED'.

<sup>140</sup> Unless we equate software evolution with the entirety of human evolution, rather than the ontogeny of a single human being.

<sup>141</sup> In asking this seemingly oversimplified question we do not overlook the principal differences between the two cases, namely the biological and experiential as against the computational nature of the data, respectively, or the substantial qualitative difference in the inherent complexity of the respective learning

We are driven to dismiss and set aside these questions (as Mr. Gawdat's interview, cited above, so poignantly and unapologetically illustrates) by the tacit assumption that the 'objects' in both cases – the human neonate on the one hand and the algorithm on the other hand – are categorically equivalent entities in terms of output-generation. We witness the outcome – the event – in both cases in virtually the same way (Google's engineers witnessed the astonishing algorithmic feedback in their lab and my family witnessed my first word in our home) and conclude/assume that what we've witnessed is equally identical. But is it the same kind of articulation that occurred/emerged where we have allegedly seen/perceived it in each of the two cases?

The absolute correlation between individual subjectivity and its (our) flawless interaction with "other minds" compels us to assume uniform articulations, notwithstanding our preclusion from internal access to those "other" subjectivities. Hence, albeit the algorithm in this case is not a classic "other mind", the robust similarity between the output it – the artificial processor – had generated when "*metaphorically, one of them came back and said: "I found something"*" (Gawdat 08:00) and the way we humans identify via our "wet" processor, namely the brain, drives us to assume a more fundamental affinity between the two *findings/identifications*.

The point, however, is that we are not comparing the algorithm's accomplishment – its breakthrough – with the human ability to identify. We compare it with human articulation. Mr. Gawdat's presupposition – his unexamined, default assumption – about the nature of the algorithm's output cuts across the distinction between identification and articulation, which is evident in the very way he relates to what the algorithm allegedly *did*, namely: '*metaphorically, one of them came back and said: "I found something"*' (Gawdat 08:00). For Mr. Gawdat, '[it] *said: "I found something"*' is a unitary, elemental, single act which both the algorithm and we humans *do* when we discern/apply meaning. We *find* and *say* – express/project that which was found, whereby the latter – the pronouncement of the finding – is some kind of necessary aspect or consequence of the former – the finding itself.<sup>142</sup> But as we are about to see, this is not the case at all, wherefore identification and articulation are altogether different things. Mr. Gawdat can

---

mechanisms, which include genetic predispositions and biological systems (nervous system, sensory organs, etc.) in the human case (as well as pre-verbal behavioral output, such as facial expressions, body movements, crying and other vocalizations). These differences are, however, irrelevant to the fundamental point we address. We asked whether the nascent act of producing output in the respective cases of nascent articulation, however different in its structure and environment, is the same kind of impartation of meaning, or more precisely whether what we perceive as similar [nascent] acts/events from without consist in the same [nascent] acts/events from within. A human "*no!*" or "*it hurts*" or "*maybe tomorrow*" are the same kind of meaningful articulations. The different "*meows*" a cat produces in different circumstances are equally identical kinds of meaningful articulation. Yet our question centers on the nascent emergence – the very genesis – of articulation in human versus artificial environments: we ask whether the nascent human articulation (the emergence of the ability to impart meaning) is the same kind of event in the world as the uninstructed algorithm's "*I found something*".

<sup>142</sup> This is why he further sees no fundamental barrier to silicon-based artificial beings achieving the same comprehensive, meaningful performance in the world that humans do through their biological substrate.

indeed say that the algorithm *found*, *spotted*, or *nailed* something, namely displayed evident potential – an inherent qualitative aptitude. This is what the engineers in the lab, for all intents and purposes, witnessed. But he cannot say, neither literally nor metaphorically, that it *said* something. No one – neither the engineers nor anybody else – ever saw or witnessed that. We, in contrast, *say/articulate*,<sup>143</sup> which is another kind/level/category of aptitude – such that is never perceived from without, but which is the sole potential of making something out of nothing; of sparking light in absolute darkness.<sup>144</sup>

We are also not attempting to examine/ascertain whether the algorithm's accomplishment imbues it with a quality of subjective experience – the sense of *being like something*, which as we shall also see below, is an absolutely erroneous way of approaching the body-mind problem, methodically convenient as it may be (we reject convenience as a matter of principle in our pursuits). The phrase<sup>145</sup> 'to be like something'

---

<sup>143</sup> This is precisely what we witness when spoken to or when we ourselves generate speech.

<sup>144</sup> Pattern recognition or any other artificial application are never – by definition – out of nothing.

<sup>145</sup> It is not an idea, but essentially a phrase. The equivalent idea is 'something'. However, closer examination reveals that the abstract idea of 'something' (and therefore also 'nothing') is fundamentally vacuous, a 'non-entity' devoid of content (an empty concept; to engage with – relate to – the abstract nature/properties of 'something' is to engage with – to grasp at – *nothing* at all), whereby 'something' is essentially and necessarily 'something *that is*' (a phrase), which meets the conceptual criterion of 'what it is to be [like] it! Even the famous reflection, commonly attributed to Einstein but first offered by Leibniz and echoed by Spinoza: *why is there or should there be something rather than nothing?* – is misleading in that respect, as it appears to posit – both grammatically and semantically – an objective realm untainted by perspective (the quality purportedly enacting the experiential level – the "likeness" – of something). In fact, however, even that conception of reality – of 'something' – yields the phrase 'to be like something', since the quality (nature) of that reality doesn't denote merely the formal – abstract – fact of "*isness*" (existence) but consists in rule-governed self-manifestation (this is exactly how Einstein conceived of physical reality). And can we positively assert that rule-governed self-manifestation does not amount to *being like something*? The sense of unease and ultimate confusion this question stirs derives from our propensity to identify the phrase 'to be like something' with the quality of human subjectivity (the taste of an apple, the agony of loss). But that is an evident speculative error. In Wittgensteinian parlance, we are misled into speculative confusion by the way we speak (e.g., "*I feel*" vs. "*this is there*" or "*the force of gravity*"). The 'something' in/of "*gravity applies*" is the same 'something' in/of "*I feel*", whereby the discursive qualification provided by *I* reflected in '**to be like something**' is purely lingual. Hence, while the physical sphere does not suffer *pain* or is taken by *fear*, its rule-governed actualizability seems to meet the conceptual (rather than the discursive) conditions for *what it is to be like something*. So, when the philosopher and computer scientist, Dr. Bernardo Kastrup, suggests (Kastrup, Bernardo, "Bernardo Kastrup – Can AI Become Conscious." *YouTube*, uploaded by Closer to Truth, 19 May. 2025, <https://www.youtube.com/watch?v=pvplorQ9xWs>) that while "*there is something it is like to be Robert, there isn't something it is like to be a single neuron in Robert's head*" (Kastrup 06:40) or that "*nature is showing me that I have good reasons to think that even amoeba have a private point of view of their own... they exhibit behavior that we associate with a private conscious point of view..., but I don't think nature gives us any reason to think that a volcano has a private, conscious point of view of its own*" (Kastrup 7:46), he commits the speculative – categorical – error of disregarding the role of language games – the effects of the ways we speak – in the whole setting, taking *phrases*, and specifically the phrase "*there is something it is like*", for ideas and taking/denoting 'Robert', 'amoeba', 'neuron' and 'volcano' for/as phenomenal "objects" rather than parts of speech. If pressed, Dr. Kastrup would certainly concede that they are all (amoeba, neuron etc.) *phenomena* (within the 'field of subjectivity of analytic idealism', to which he is committed), but he *speaks of them* – all of them regardless of any supposed agentic distinction (e.g. between an amoeba and a rock) – as absolute noumena, namely as "things in themselves" displaying different properties. That

is an existential tautology and a hollow articulation, like purporting to formulate what breathing is by inhaling and exhaling oxygen. Every articulation, be it affirmative, interrogative, reflective, speculative etc., is a straightforward, absolute manifestation of 'to be like something'. Hence, the very phrase 'to be like something' (or 'what is it to be like something?' or 'to be like something is [or 'is not' or 'is unlike'] X' or 'what it is to be like something is...') is itself **to be like something** [in the first place]. How, then, can it (how can we – *being like something* in the very articulation) allege to relate or approach itself from without (as an object of contemplation, which, namely the contemplation, is also a manifestation thereof)? We generally do not inquire into something by the very act of doing it (e.g., we do not inquire into 'walking' by stepping forward). How, then, can we purport to act (or claim the eligibility to act) otherwise in regard to 'being like something'? No reason whatsoever presents itself. No compelling promise of clarity and genuine illumination consists therein. In point of fact, no compelling rationale comes to mind to account for the prolonged dominance of this approach as a principal, if not singular, pathway in the body-mind discourse; its very longevity is, arguably, its own perplexing question.

Indeed, we can speak of our experiences, and we can speak of the fact of *experiencing*, and also of the fact that we are something that experiences. We can conceive of ourselves entertaining various kinds of ideas, including the idea of something that experiences itself in ways that other objects appear not to be able to do and further reflect on and explore the reasons why this is so. But all that "talk" and conception is itself an experience.<sup>146</sup> It is 'to be like something'. And that 'experience' – the qualitative self-projection (*likeness*) it consists in – is not a construct onto the 'something' (the form of life) we experience being or that "underlies" it. It is the latter – the 'something' – in itself,<sup>147</sup>

---

tacit allusion is the speculative confusion of disregarding the language game he necessarily engages in. The error in Dr. Kastrup's case is twofold: **for one**, as his real-time argumentation unequivocally illustrates, he has not *opted for* idealism (metaphysical subjectivism – consciousness as the ultimate existential substrate) over materialism (metaphysical objectivism – physical reality as the ultimate existential substrate), but – in speaking of different "things" (kinds of stuff? Different experiences?) – merely disregarded (skipped over) the alleged distinction therebetween altogether. Moreover: saying *everything* is consciousness/subjectivity and saying *everything* is physical/objective is actually saying the same thing, namely nothing. We have no access to either 'everything', nor is there any inductive, deductive or other conceivable path – an adequate way to extrapolate – from concrete to general in that respect. My subjective experience and your subjective experience are not two chains in a wholesome "neckless". They – their existential consonance – are the very maze. Whereby all Dr. Kastrup does, under the guise of argument, is issuing *statements* grounded in nothing other than semantic coherence. In other words, while that/his purported extrapolation is semantically sound, its merit ends there. **Secondly**, the language game of things and properties – being a language game and not an ontological setting of a lab (an observer/scenery scenario) – cuts across and initially underlies (enacts) the notion of '*to be like something*', whereby the attempt/pretense of applying that notion – assigning it as a property – to any particular "thing" collapses. Hence, the attempt or pretense/pretext to elevate any particular variant – any particular distinction – within that semantic field (actually – any particular *move/step* within that [language] *game*) to the level of metaphysical reality collapses too.

<sup>146</sup> One needn't even say that this is *ontologically* so, as no other possibility is available.

<sup>147</sup> A rule-governed self-manifestation.

like an insentient object or any of its constituents is something in itself. We haven't found, nor can we even imagine finding, a non-rule-governed foundation for all that exists. All we have are rule-governed self-manifestations (in animate and inanimate phenomena alike). The '*likeness*' of 'to be like something' is not the '*likeness*' of a dream or a hologram. It isn't the '*likeness*' of a picture. It is the '*likeness*' of the 'something' it is to be.

And since [to be] 'something' is essentially 'to be like something',<sup>148</sup> trying to examine or ascertain whether [to be] 'something' is or is not 'to be like something' is actually trying to ascertain whether [to be] 'something' is [to be] 'something' or whether 'to be like something' is 'to be like something'.<sup>149</sup> In other words, it is tantamount to doing nothing and a purely performative enterprise, solemn as it may appear to be.

Consequently, the well-worn, "classical" distinctions between applied cognition and tangible eventuality, e.g., that speech and perception entail experiential qualia beyond mere energy emission/transmission and computation and that we, therefore, ascribe the algorithm such properties as processing information and manipulating representations only metaphorically, will not bring us any closer to a genuine insight. What we need are not bird's-eye view, ever half-baked intuitions about the exceptionality of cognition in general, but an argument setting apart my '*tractor*' and Google's '*I found something*'.

Such an argument's beginnings were already outlined above.<sup>150</sup> Google's algorithm identified a visual pattern in the stream of data it was bombarded with – one of many that evidently existed there. It didn't "report" thereof, as Mr. Gawdat offhandedly assumed. Can we say that it spotted form where none existed? "Where" had it done what it did? The result of its operation (again not its report) was absolutely *for us*. Google's staff saw and understood that it identified a pattern. For *it*, it just executed a function. In cyberspace. Saying that it "nailed" or locked-on an essence is equivalent with saying "there are patterns in the world", which already assumes algorithms – natural and artificial – whose essential "job" (whose "essence") is to track them (if there's a pattern, there's an algorithm and vice versa).<sup>151</sup> But again: where do they do what they do? Cyberspace (which is categorically uniform for humans and machines alike) is a process – a formal construct of predetermined logico-mathematical variability – whereby the rule-governed, semantically neutral manifestation of physical reality is molded in ways that further provide for semantic (empirical) coherence. But it isn't congruent with the latter.

---

<sup>148</sup> See ft. 145.

<sup>149</sup> To exclude any doubt whatsoever, the alleged principle of multi-instantiation (multiple realizability), if it is coherent at all, is inapplicable here. If this is still not evident given all the above, namely that 'to be like something' is not a property of an object (like a bird's/airplane's capacity to fly or a fish's/submarine's capacity to swim) but the very nature/essence of anything at all, we shall posit that conceptual tautology or ontological equivalence (which are absolute by definition) do not yield to that notion/principle.

<sup>150</sup> See ft. 139.

<sup>151</sup> The phrase "there are patterns in the world" assumes algorithms. The rule-governed self-manifestation of physical reality before life evolved was "free" of patterns, in the sense that this notion is inapplicable thereto. Whereby, the rule of self-manifestation is not a pattern.

Semantic coherence requires more than logico-mathematical variability. It requires “something” (e.g. rule-governed self-manifestation) at the near-end of computation and grammar at its (computation’s) far-end. But it does not emerge as a linear consequence of the first two, namely of a tangible substrate and computation. Semantic coherence is not a logical outcome or a conceptual derivative of applying computation in a tangible sphere (phenomenal user interface, a robot etc.). The human genetic disposition for grammatical articulation cuts across both. At a certain point of natural evolution – probably with the inception of life itself<sup>152</sup> – there has evolved natural phenomena which projected cyberspace, namely behaved in ways that involved/necessitated computation of some sort. The human species is a descendent – a “spin-off” – of that development, whereby we project semantics, e.g., *see*, *speak* etc., due to the fact that our brain is engaged in computation/processing upon its physiological substrate. But it doesn’t process information, or for that matter – impressions. It manipulates signals (neural firings) in accordance with rules that govern formal signals. A further existential substrate is imminent for semantics to evolve or emerge. We call it ‘grammar’ or ‘syntax’ (the “rules of...”), reiterating the tacit assumption that it is a facet or an effect of matter (physiology) and computation. But what is it really?

Indeed, what might have at first appeared to me as formless vocalizations gradually acquired form and structure that I could perceive and mimic. But that perception-based mimicry, far from identifying *what makes a tractor a tractor*, was not an emulative kind of act (to merely identify is to echo). It did not amount to denotation of a particular *this* (it wasn’t *denotation* at all)<sup>153</sup> but an act in (a projection of) a world – an entire realm – that *is* in itself, not unlike any of the other pre-verbal behavioral output and non-verbal livelihood/functionality I generated/displayed since I emerged from my mother’s womb, which also projected an entire realm. More precisely, it wasn’t an echo of any elemental *substance* but [an assertion/affirmation] of an entire *world*. It echoed a disposition to be *part of something* much broader than any element or distinct entity as such may be. It accommodated an echo of *to be*<sup>154</sup> rather than “there you are!” (or “there you go”; “that’s what this is”). Finally, it didn’t amount to a denotation of a formal ‘something’ but to a manifestation/proclamation of ‘*I am*’ which is already imbued with ‘there’, ‘is’ and the unpronounceable/inarticulable juncture/hookup between them. That is also why, very quickly and essentially escaping observation from without, I – as any other infant in my place – will have been able not only to utter ‘tractor’ but to do so by way of *pointing*, *answering*, *asking* etc., and to apply verbally not only nouns but also words like ‘*no*’, ‘*yes*’, ‘*give*’, ‘*give me*’ and further be able to “understand” and thereafter also generate far more

---

<sup>152</sup> Let us remember that already primitive forms of life, even at the cellular and sub-cellular levels, display behavior that is compatible with primitive “brains”.

<sup>153</sup> It wasn’t a move on a board.

<sup>154</sup> *To be* is to be in ‘a world’, not in ‘a state’.

complex verbal expressions ('do you want me to leave this here?'), namely *to speak* in a worldly manner, as an inherent, wholesomely integrated part/aspect of *The World*.

The application of the rules of grammar reflected in that/my evolving worldly speech was/is not any kind of outcome or combination of matter (physiology) and computation. The latter were indeed the case when it "happened" (on the level of its happening rather than my nanny's perception thereof), but *it* didn't amount to them.<sup>155</sup> 'It' was the form of my ontological immersion/integration in my environment – in the relevant environment of my species in the most general sense of that word, namely in what *is*.<sup>156</sup>

The algorithm nailed a pattern. My brain has certainly identified an abstract form as a prerequisite to my agentic ability to say 'tractor'. **But I participated in a language game** [in a world]. There is no *shift* here which we can recreate.

A language game is not a structure. It isn't a game humans devised (no 'let's play' precedes it). Wittgenstein knew that a 'language game', as well as the social 'rules' behind it, are metaphors [of the unspeakable/inimitable].

In recognizing the pattern underlying the image of a cat the algorithm did not echo any sense of *to be*. Nor do any of its later, far more superior and ever improving versions which seem to "speak" and/or behave as we do (but not as we **are**). They do not participate in what we – metaphorically – call 'language games'.

### XIII

Adopting a more "productive" language, the more productive (and successful) among us would say that the category of naming is *wired* into the infant's (and also the deaf person's) brain. But, again, the very phrasing: "a category is wired into the brain", albeit

---

<sup>155</sup> And when my parent's came home and my nanny announced the exhilarating news regarding the pronunciation of their child's first word, and they further applied that word in my direction ("you said 'tractor', my dear baby!", "it is a tractor, yea?" etc.), and I went on to utter, essentially playfully: 'tractor', 'tractor', the way infants do in such situations, the very same thing happened, namely I reiterated and solidified and substantiated my immersion in my environment / the world, rather than absorbed myself in any kind of identification or communication (they communicated with me, not me with them; we were absorbed in different language games).

<sup>156</sup> Whereby, Aristotelian physics, classical physics, quantum mechanics, the theory of evolution, psychology, anthropology etc., are not reflections/mirrors/insights of or into that *isness* (of what *is*), but one of the facets of that immersion/integration – participations in one of the many language games it consists in: the particular language game of *understanding* (others are: achieving, protecting, cooperating etc.). The ancient Greeks, as other ancient and modern civilizations, tried to *understand*, but also to *achieve*, *protect*, *cooperate* etc. The fruits of their 'love of wisdom' (e.g., Aristotle's conception of motion) did not instruct any of the other (on the personal, social or national levels), not because they (or we or me and you) did/do not rise to the challenge, but because they amount to wholly different language games. Their arithmetic and geometry, in contrast, *did* instruct their (military, architectural and so on) achievements, precisely because mathematical exploration (essentially *cracking* and *finding*) – exactly like modern applicative (technological, medical etc.) research – is part of the language game of 'achievement' rather than 'understanding'.

coherent at a certain descriptive level, leaves out and has no explanatory force nor sheds any light whatsoever on the ultimate way or substrate whereby the brain – **a material machine/engine/processor** – and the human agent’s capacity to project meaning (to apply complex semantic categorization) – **an agent “inside out” that participates in semantic interaction** – are connected or “molded” into a unity. It still doesn’t account for the practical ability to associate a certain experience with a category – for the way this becomes possible (the way these two unbridgeable realms make up a single world).<sup>157</sup>

Language games do. But unspeakable and inimitable as the “reality” behind that metaphor seems to be, what does it actually amount to? Let’s proceed.

#### XIV

We can now finally bring together under the Wittgensteinian umbrella the two questions underlying this paper, which boil down to the same fundamental puzzle:

1. How does a physiologically robust infant learn that sound is sound, and something is something? How does he/she acquire the ability to discern form and structure in the blizzard of initially formless phenomenal variety in which he/she finds him/herself throughout that foundational period of their development?
2. If our deaf hero *hears* [the audible fabric of] sound but also *feels* [the touch of] sound, what is to *hear* [audibility]? And if our deaf hero *feels* [the touch of] sound but also *hears* [the audible fabric of] sound, what is to *feel* [tactility]? In face of such pervasive or even complete overlapping of these notions, what is their respective, distinct meaning (Their abstract marked-out value *as such* – as **what they are**)?

Answering the second question will pave the way to the first.

The evident answer to the second question in view of the discussion above, at least in our humble opinion, is that the difference between the *meaning* of ‘to hear’ and the *meaning* of ‘to feel’ (or any other respective pieces of semantic sense) lies in the practical context – the world/environment based (“planted”) behavior – in which those lingual entities – each of them respectively – originate in the first place and in which (*do you hear?*) or in relation to which (*I lost my ability to feel*) they are further necessarily applied.<sup>158</sup> Exactly

---

<sup>157</sup> We perform as subjects, but we explore the material world. At the same time, we explore the realm where we perform as subjects, *and* we perform while (in) exploring the material. Performing is *in* the explorable while exploration is performance (performable). Each of the two presupposes the other with no apparent way to step outside. The term ‘wired’ is purported to connect the two, but it assumes an antecedent – external to both – level. At present I cannot (nor can anybody else, to my knowledge) say anything about the latter (this level or fact). Maybe this is what Wittgenstein meant we should be silent about.

<sup>158</sup> And the fact that we presently deal with both these notions as it were ‘in the abstract’, does not undermine that conclusion the same way that the drive wheels of a vehicle would “operate” even when

like 'tractor' or 'object' or 'to say' acquired their respective "meaning" – their 'livelihood' in the landscape of my subjective development (essentially from my own – which is the exclusively genuine – vantage point) in the example above. What I essentially "learned from" in the above 'tractor' example, unlike any kind of trained algorithm or machine, wasn't any kind of data but essentially *behavior* (being part of something), which is what all the previous exclamations/pronunciations of that and all the other words in my infant universe amounted to. In other words, learning to apply meaning is not an act (application of... [something]) but accommodation (immersion in... [something]) via behavior (being in/as/part of... [something]). The growth it consists in is immersion, not becoming.

The word 'feel', namely the symbol for tactility or any other aspect of human subjectivity (experiential state) that this word denotes, including the entire symbolic chain from which it had evolved beginning with the first communities of human language-animals, emerged, ensuing the genetic mutation that introduced the specific lingual manipulation humans are capable of into their habitat,<sup>159</sup> from and by way of and as integral, ultimately inseparable, part of the practical activities those members of primordial homo sapiens' communities were engaged in in what comprised their 'world', whereby their 'world' and their concrete 'ways of life' were one and the same.<sup>160</sup> And the activities that forged interactive room for that notion (the overall practical instances/circumstances which incited sensations and personal experiences of a general kind, e.g. '*I experience something*', '*there's something in me which is the effect of what is going on or interacted with me or happened*'), given our entire, complex genetic endowment and the conditions of physical reality, were, albeit related to, still somewhat different from those which

---

detached from the ground albeit *being (and becoming) what they are* only in the context of facilitating mechanical motion on a surface.

<sup>159</sup> That 'language gene' which made it possible for the homo sapiens both to represent to itself various aspects of its environment through grammar-imbued symbols (*come, this, for, there, near, far, now* etc., in their preliminary raw forms) and to project them into its practical interaction with its immediate, empirical environment. Indeed, as Noam Chomsky has argued, the ability to apply language is not, strictly speaking (as it were prior to everything else), a means of communication like a secret code that we invent may be, but a faculty for essentially "private" representation and evaluation of (thinking about, whereby representation is [necessarily] designated for [whatever kind of] evaluation) the environment (of/regarding 'what is'). But, based on the same scientific logic of evolution by natural selection, it is equally palpable to assume/conclude that "private" evaluation-oriented representation has evolved in order to enhance the chances of human survival, namely for more efficient (survival oriented) interaction with the environment, of which communication with other members of the species is no less crucial (and practically more so) than coming up with the content thereof.

<sup>160</sup> Unlike the way we use the term 'world' in its abstract sense today. And the fact that we can use it in that abstract way is wholly irrelevant and does not undermine the point. Moreover, viewing 'the world of a certain species', or even the interwoven domain of multiple species, as the cumulative expression of their ways of life (rather than the aggregate of marbles – or fish, food, plants and water – in a jar or processes in a sphere such a brain, a lab, or planet earth as seen from high above), offers a distinct insight into *Homo sapiens'* position in the natural world. From this perspective, the human world (understood as the totality of what is meaningful in the broadest sense) is fundamentally no different from that of any other living creature, extending to the very base of the tree of life. This insight may well seem trivial from the standpoint of evolution theory, but from a philosophical point of view, especially given the present analysis of meaning and language and their instantiation in the world, it is much less so.

prompted the primitive real-time articulations of something like '*I hear*' (noticing rolling thunder, the roar of a nearby beast or a fellow humanoid's vocal gesture). Indeed, the constitutional role of such closed-up sets of circumstances as the ontological substrate ("environmental cradles") of language acquisition in a community of advanced humans, is all the more emphatic and lucid in regard to those who "spoke first" and among whom language as we know it initially evolved. Wittgenstein's idiomatic phrase of 'language games' is intended to capture the entire phenomenon but stops short of delineating the complex dynamics and ultimate nature thereof, which should indeed be the role of philosophical scholars rather than singular, brilliant minds such as his.

Indeed, to participate in a *game* – follow rules<sup>161</sup> – is to immerse in a setting and not apprehend scenery, but the whole point is that acquiring the *capacity* to follow the rules of a game does not categorically precede the enterprise as a whole. The rules at issue are not traffic rules or those of societal good manners. They need no explanation. They are inherent to the raw setting of being a creature in the world.<sup>162</sup>

The way to see it is to understand that one's ultimate interlocuter in any lingual interaction is not another person, or another/other mind/s, namely not another *speaker* (agent), but always and necessarily *the world* as such. Think of an instance (the number of which along our daily lives is numerous) of awareness to my own breath, a yawn I make or my footsteps. Are they different, at the fundamental level, from a word or a phrase or a sentence someone says to me? Indeed, they aren't orders or questions, but they can elicit conscious, intended response on my behalf, and at the fundamental level – as semantic units – they are no different from a spoken phrase I've heard or read, whereby my reaction/interaction to/with both consists in the same immersion in the environment (my=the<sup>163</sup> world). Semantic units are ipso facto comprehensive settings and any semantic unit which finds its way into one's perceptual-conceptual theatre – one's consciousness – is a message from *the world* one becomes a part of – and persists as such – by the very absorption and/or application thereof. Indeed, speaking "from within" semantics (within grammar), there are other people and other minds and [equally sovereign] subjects and various other stuff "out there" in the world (the absolute "out there" or 'objective reality' that encompasses us all). But the "grammarless" perspective – which is the sole philosophical perspective,<sup>164</sup> depicting the interface of subjectivity and objectivity as such<sup>165</sup> – presents that beyond me there is only a world one interacts with

---

<sup>161</sup> In the sense of adapting to a particular form of interaction rather than adhering to a linear set (or sets) of instructions aimed at achieving a linear set (or sets) of consecutive ends.

<sup>162</sup> Think of a living cell swimming in its biochemical substrate. Moreover, as the synthetic and developmental biologist, Michael Levin, has shown, microorganisms themselves display mindlike behavior in the sense of following rules.

<sup>163</sup> In real-time [my] world is semantics while [the] world is reality, with no way to separate the two.

<sup>164</sup> Philosophy as the unaccountable, but equally inescapable, strife to transcend grammar.

<sup>165</sup> Not 'me' and 'the world' as concepts but subjectivity and objectivity as the exclusive, total in their complementary stretch, existential categories, comprising a single, omnipresent, ontological constant.

*and in* at the same time. Only reality in its impeccable, absolute sense.<sup>166</sup> Within grammar I am part of a world that presents no philosophical challenge whatsoever (only existential ones, which we are by definition [existentially] equipped to confront).<sup>167</sup> Outside grammar there's 'me' and the 'world' as a (*the*) singularity, whereby philosophy, rather than any "applicative" endeavor, steps or may step in.

Approaching our existential setting from that angle allows us to see that 'meaning', or more precisely semantic sense/coherence, does not derive from – is not constituted by means of – any kind of subject-object interaction (e.g. 'identification' – producing or equipping oneself with something) or combination of previously familiar objects/notions (any kind of 'elements'), but by practical immersion of the human agent in the multiple pathways and possibilities of its existential environment – in the human creature's developing into an integral part thereof rather than learning what this "out there" scenery is. In other words, the subject does not **emerge out of** objective reality (an objective substrate). It **grows into** one. This is the real, ultimate meaning of 'articulation'. To articulate is not to *represent* but to immerse, to grow – get oneself interwoven – into. It is to journey homeward; find one's way to one's abode. That's why it makes something out of nothing or the sense in which it sparks light in absolute darkness.

That's what I did in my first '*tractor*'. That's what Google's algorithm didn't do in its "cat".

That is also what other members of the animal kingdom do in learning to interact with their environment and conspecifics, irrespective of any distinction between humans and other animals in this paper (specifically in ft. 86 above).<sup>168</sup> For obvious reasons, our

---

<sup>166</sup> Reality not as a bowl of fish, but as the fact of perspective. Perspective and nothing else, whereby 'world' and 'perspective' are one. No *and...* is required for perspective, as it is the only grammarless notion. 'Perspective' as immersion in itself.

<sup>167</sup> 'What is truth?' or 'what is good or bad?' are, ultimately, also, and maybe the first and foremost, existential, rather than merely speculative, abstract challenges. It is very "convenient" to think otherwise and come up with abundant reasons and arguments to that effect. But the difficulty of a challenge is always practical/existential, and it never transforms it into a theoretical issue. Furthermore, that very initiated and essentially sovereign transformation, namely from practice to theory, is in itself a moral, namely existential, stance regarding the challenge.

<sup>168</sup> E.g., oxpeckers, for example, will be articulating to or messaging other oxpeckers, not the hippos or other mammals they perch. On the other hand, my cat, Gigi, evidently *speaks to me* (a human) when it approaches my chair in my study and meows while looking at me ("*give me my daily Premio snack*" or "*I want a snack*") just before continuing to rub against my feet in case I didn't respond by looking back and/or rising from my chair and heading the kitchen. In such circumstances, however (from Gigi's standpoint), I am not an animal of another species (which, in contrast, Gigi *is* when my wife or daughter "converse" therewith, e.g., "*what is it that's wrong you today, darling?*" or "*what bothers you, why are you sad?*" or "*you've just got your snack, sweetie, stop being a treat monster!*") but merely one of many functional objects in its world. This is an interesting distinction, since theoretically (existentially) I might be both (could I be *both* for an AI that interacts with me?). A mechanism issuing a snack in an orderly, cyclic manner would have fully replaced Daniel Drabkin the human being as far as Gigi is concerned, being a strict means as well as another "something" (another 'creature') in its environment, and it would even try to get extra – "out of the cycle" treats from that mechanism as well, however futile that would turn out to be. Would that mean Gigi shall be *speaking* to a mechanism? Absolutely. Since speech is ["born" and developed as] behavior in concrete species-specific circumstances, it absolutely would (would an analogous behavioral

primary focus has centered on human speech with its complex syntactic and semantic implications on human consciousness, the human conception of meaning, and our broader worldview. Consistent with our specific line of argument, however, the same applies not only to humans but also to the rest of the fauna, particularly more advanced forms like primates and dolphins, yet conceivably to all animals and, however radical the assertion, to the whole biological domain.

The syntactic expression of meaning, i.e., the manifestation of human language, is not singular except in evolutionary terms. Existentially, its distinction from other modes of expressibility lies in its form (structure), not its kind (fundamental nature). In other words, it is not exceptional in terms of *agentic aptitude* per se. Other animals also perceive their environment, interact therewith, and even, we now dare to argue, perceive themselves as part of their environment. For strictly evolutionary reasons, they lack the property of 'reflective' conception/contemplation<sup>169</sup> and the syntactic notion of *I/me*, but aren't denied, for that very fact, expressibility as such, namely the possibility of self-expression. Since meaningful subsistence and interaction in/with the world essentially emerge – the capacity therefore is initially acquired – across the biological realm via behavior in concrete practical contexts, its later, mature (and equally contextual) manifestations throughout that realm (among all its inhabitants) allow and retain an equally vibrant level of agentic – percepto-conceptual – relatability, and the syntactic form humans had evolved to be able to do so does not disqualify other forms of the same phenomenal category (self-expression as part of a larger whole).

Indeed, the seeming leap from performance in a realm (what we have called, alas erroneously: being in the world without language) to non-performative interaction with scenery, referred to above, turns out to be a typically human misconception – a conceptual error or illusion. As we shall see below, the intrinsic nature of language and meaning generation, understood as an immersion in a broader totality, further reconnects us with our deep natural roots.

And since becoming a 'language animal' is the very essence of semantic coherence, then learning to speak – “picking up” the ability to use language in the first place – emerges from the same immersion rather than familiarization-oriented ostensive association (this is 'red', this is a 'cat', this is to be 'angry' etc.). The human agent learns to speak like he/she learns to eat or walk or drive a bicycle rather than where Australia is on the map, that this

---

manifestation by an AI amount, by extension, to *speaking*? Absolutely No). One may argue that my pet's meowing while we stare at each other in that example isn't *speech/speaking* or "conscious"/intentional interaction at all but another kind of natural or functional behavior that has nothing to do with speech, human speech or consciousness. Indeed, the core of our project in this paper – its very quintessence – is to present and demonstrate a compelling argument to the contrary. Bear with us.

<sup>169</sup> Although the distinction between the physiological manifestations of basic instincts (like fear, hunger, reproduction, or the curiosity observed in some animals, such as my cat) and the viability of a corresponding, abstract state of mind (still not 'reflective' in the human sense, but possibly 'deliberative' indeed) is remarkably subtle and fundamentally indistinct.

is your new neighbor, how to fix a car or the laws of physics. Furthermore, he/she will have been able to learn all the latter since (and only because), ultimately, they **are** (before they are “stuff/themes” to be learned) pathways and possibilities of our existential environment. Geography *is* [as such] only because planet earth is the domain of our species; my neighbor *is* [as such] only because we survive in communities; mechanical maintenance *is* [as such] only because we had been able to produce advanced means of transportation, and the laws of physics are [as such] only because advanced brains are participants rather than mere bolts in the cosmic realm. Indeed, before anybody had learned where Australia is on the map they learned to move from place to place and then to navigate. This is the whole (or almost whole) story.

Finally, learning itself is one of those pathways rather than a lens or any kind of *means* to interact with [the very fact of] ‘existence’ as such. We [are able to] learn **since** we exist rather than *for* any prospective purpose (least of all to explain existence as such or the fact that we exist or issue any impact thereupon).<sup>170</sup>

So, when our deaf hero experiences sound for the first time, what actually happens is not captured in (and totally evades) that sentence. It isn’t a case of experiencing something for the first time in the sense it is so upon eating a new, exotic dish or watching a game one has never seen before or even meeting an alien creature from another planet, namely nothing “new” emerges on the fundamental level of our makeup in the sense in which the aforementioned experiences *are* new in our routine. And while the experiential fabric (sound) of what in the community of “hearers” passes (is regarded) as *hearing*, will gradually pass from being utterly novel for him/her into being utterly familiar, the fundamental and constitutive feature of the unfolding event is not the acquisition of a new capacity (again – that is the way to describe or conceive of it from within semantics) but our hero’s “existential introduction” (immersion, fusion) into an existential pathway that is typical of his/her species, namely a specific mode of interaction with the/his/her environment.

The word and its signification, or any *ability* in their regard, are not important. The instance in the world [*as it is*] is. Indeed, when I hear a word, *someone said something*, but this is an interpretation from within semantics. As an instance in the world, it has nothing to do with that interpretation. The instance itself is first and foremost the world being as it is in its uncleavable, uninterpretable (and only) sense.

Hence, it is also of no consequence whether it is [what we call] a name, an order, the sound of the wind, his/her own breathing, a melody or rolling thunder that our hero “experiences”, namely the different structures of the sounds he/she (indeed – all of us) experiences as such, but the actual – genetically framed – interaction in which he/she

---

<sup>170</sup> Which is also why “teaching” machines to apply language the way it is currently done in the world of artificial intelligence involves a very fundamental crux as far as its “humankind horizons” are concerned.

engages<sup>171</sup> in real time (social hierarchy, leisure, survival) which already embodies those structures as part of the environment we had all<sup>172</sup> “found ourselves in” for the first time/s at one point or another along our existential (physical, biological, psychological, social etc.) stretch.<sup>173</sup>

Consider the following example (a thought experiment of sorts without manipulation of time and space): in utter synchronicity with our hero’s experience of a name, an order, the sound of the wind, his/her own breathing, a melody or rolling thunder, other similar phenomenological events (another name or order, some other person’s own breathing, some other melody or another event of rolling thunder) take place (issued someplace else on planet earth). Our hero will *interact* with the world – with a certain manifestation thereof – in the way that is embodied in his/her experience (so as to “produce” the latter at his/her experiential level), while not interacting with the world – with another manifestation thereof. But these allegedly *different manifestations* of the world are nothing but grammar – but the way we represent things/stuff (essentially in retrospect) in a semantically coherent way (a “story”, in the final analysis). In reality, namely from the standpoint of the world (of planet earth, in this case) it is the [same] *world* he/she interacted with / immersed in and did not interact with / immersed in at the same time. It is the [same] *world* which “expressed” itself, so to speak, regardless (conceptually) of any “experience” (perspective thereupon by/from a part thereof). And this/the *world* hasn’t changed (wasn’t something else than what it is) in consequence of these immersion/interaction. Moreover, *it* persisted as if nothing other than itself being the case had occurred. So, it is the world in its uncleavable, uninterpretable (and only) sense that is the *Ontological Case*, rather than any cognitive emergence, even when subjectivity seems to “grow” or play a note.<sup>174</sup>

Nor is it, therefore, of consequence *how* our hero – indeed, any and all of us – is ultimately capacitated and succeeds (at the technical/functional level) to “get access” to the specific experiential fabric at issue (through the ears or through the wrist, grounded, ultimately, in the same aspect of the physiological substrate). That fabric, like any other experiential (strictly subjective) phenomenology as such, would be nothing at all if not for

---

<sup>171</sup> Is disposed and bound to engage in consequence of his/her genetic endowment.

<sup>172</sup> Every single creature that ever *WAS*.

<sup>173</sup> “The sentence ‘the interaction embodies the structure/s as part of the environment’ may at first seem problematic, but this is only because of the tacit assumption that all these three are ultimately separate and that there is structure and there is interaction and there is environment and that the second (interaction) takes place according to the first (structure) in the third (environment). But in fact, from a logico-ontological perspective, the second and/or any form it may take *is part of* the third, namely the ‘environment’ is not a stage or a neighborhood but the fabric of everything that *is*. ‘WE’ and everything “we are” is an essential part of what [there] IS. It is very easy to miss or, as it were, to “slip out” of this extremely analytical perspective, since it undercuts the grammar at the core of our being as we are (in a sense it is an absolutely futile insight; rings a Wittgensteinian bell?), but that is the crux of the matter: what it undercuts are not alternative conceptions of how things “ARE” but merely and specifically grammar (namely, the latter are not conceptions but grammar: they are articulated from within semantics in the first place).

<sup>174</sup> More on this thought experiment and its consequential implications to our line of thought below.

its place in the greater story in which the agent/person/creature is born to play a role through its entire set of capacities as a whole – a whole which is a part of the whole world, rather than for any distinct “apparatus” or operating system it – the creature – is equipped with or any underlying, qualitatively foundational, stative substrate.

And this is also the way the first question at the head of this chapter<sup>175</sup> is resolved: the robust infant learns that sound is sound, and something is something, as part of the genetically framed activities it (as a variant of its species) has evolved to be able to perform, namely as part of the preordained palette of interaction with its environment, stemming from its genetic makeup. The form and structure of the environment a mature creature perceives/conceives itself to be/live/operate in (sound, sight, sense and all the dynamics and categorization in between), acquires its coherent semantic sense as an integral part of what the creature can *do* given its genetic makeup. The same way it can, already at the earliest stages of its emergence in the world (of emerging as a creature), reach for its mother’s nipple, “apprehending” already then and in subtler and more complex ways over time, that certain aspects of the environment (what we would later describe as “certain something/s”) are to be reached for to sustain its very existence/wellbeing, it can also “reach out” to various other aspects of the same environment (the frame of its baby bed, the objects hanging from above, the reappearing face/s from “nowhere”, the various accompanying sounds and “gestures” of all that scenery and so on), and the “formation” of all the latter into forms and categories is ultimately inseparable from these positive-affirmative activities (reaching out and interacting). The [concomitant, not *antecedent*] “apprehension” above is not a separate, substrato-constitutive prerequisite to the ability to ‘reach out/for’ but an integral part of the very reaching – the ontological form of the latter, so to speak (to *be like something as something in itself*).<sup>176</sup> There’s actually no *ability* or display of ability here or any *reaching out*, for that matter, which are rhetorical tools and conceptualizations *we* (who/what is/are that/those ‘we’ or ‘me’ but conceptual fields all the same?) apply as part of a wholly different language game than the one the infant is engaged in, whereby the ultimate setting of both is equally and essentially inaccessible – even to our imagination (revisiting Kant from a contemporary perspective which grants language its rightful prominence). It is the setting of human immersion in the world. That apprehension, if you wish, is the same “apprehension” of a natural force *en route*. The latter acquired thereof when it, too, immersed in the world (e.g., the Big Bang).

Indeed, without an advanced brain none of it would have been possible, but that brain is part and parcel of the same environment, serving as that aspect thereof which allows

---

<sup>175</sup> How does a physiologically robust infant learn that sound is sound, and something is something? How does he/she acquire the ability to discern form and structure in the blizzard of initially formless phenomenal variety in which he/she finds him/herself throughout that foundational period of their development?

<sup>176</sup> The infant reaches for the nipple, not anywhere else, whereby the nipple, the motion and the very drive are inseparable, are **one**, are the immersion in the world.

these specific activities to take place and emanate their ensuing [for us, concomitant in fact] cognitive/semantic forms. Far from being agents or the engines of agents, however, our brains are voices of the world (see below).

So the brain *allows* (is a necessary condition for) the primary learning we discuss here, which comes to fruition in the form of both passive and active conceptual differentiation, but it is the *activity/interaction* – behavior as an integral part of the environment – which ultimately constitutes the conceptual substrate – the semantic framework – of the world we (from early infancy through cognitive maturity and up to agentic decay and extinction) perceive/conceive ourselves to inhabit/occupy.

Noam Chomsky may not like it, but syntax is nothing but for behavior.

There is no *mama* other than through ***everything I*** (the baby I am) ***can do*** regarding the creature (the aspect of the world) which I would later come to regard as *mama* given my genetic propensity to act in a certain way, e.g. synthesize physiological processes into a certain output (which a mouse or a cat are genetically deprived of the possibility to do).

Consequently, form and structure and everything in between, will emerge from the blizzard of the formless phenomenological soup I had arrived at when I was born, through what I and any and all of us will *do*: I understand that I have legs when I'm able to walk, I understand that I have arms when I can grab and so on (I understand that there's a world out there when I can interact through seeing, hearing, smelling and touching), rather than as a precondition to all these real-world activities. Similarly, I understand a spoken word or a phrase when and only when I can already talk. Hence (by way of logical inference), I have learned to talk (apply language) before I could understand – differentiate – what words mean, and I did it – learned to talk – through everything I *did* [in the world] throughout the buildup of that capacity.

## XV

Given all the above, we can plausibly say that when sound is produced (someone speaks, or a melody plays) in the presence of both properly hearing and hearing-impaired people, the message encapsulated therein is relevant to both. To pacify any objection, we would circumscribe our contention as follows: the physical-biological substrate of the emitted sound is relevant to all members of the human species within a certain distance from its source.

And why will that sound be imbued with a message (a semantic “dispatch”) for the properly hearing while not for the hearing-impaired?

- Is it because the hearing-impaired person *lacks the capacity* to hear?
- Is it because the *physiological flaw* prevents him/her from identifying form and structure in the emitted substance?

No. It is because he/she had never participated in any genetically incited activity that involved sound.

It should by now be clear that the first two answers to that question stem from *the way we speak* in the sense that providing them, respectively, are specific “language games” we engage in in specific circumstances rather than depictions of *the way things are*.<sup>177</sup> The *absence of the ability to hear* and/or *having a physiological flaw*<sup>178</sup> are not observations of reality nor are they outcomes of any kind of analysis. They are forms of interaction compliant with higher-layer goals in the human environment. Consider the possible reasons and circumstances in which the question to which we proposed them as palpable accounts may rise in the real world as against the reasons we brought this question here given the overall flow and objectives of this paper. The kind of scientific inquiry or systematic analysis which may generate the second of these two answers is an inquiry into physiology, not into human capacities. Providing that specific answer is part of human interaction at a completely different level than the one we pursue in this paper, namely exploring the acquisition of and the very enablement of lingual and cognitive skills at their fundamental level (what *these*, rather than any physiological elements or dynamics, are). The “higher level” language game of physiological – or for that matter any properly scientific – research are as integral to human existence and to the human environment as the language game we pursue here – a pursuit for an honest insight into what language acquisition and the generation of meaning *really are* rather than trying to explain a phenomenon that puzzles us, if only because our language game in this paper contains no phenomena in the sense that is inherent to any scientific activity or discourse and no ‘observer-scenery’ setting equally integral thereto.<sup>179</sup> Indeed, our language game is a peculiar one, which stands out from all the others. It is the enduring language game of philosophy, which has existed since time immemorial, and which persists regardless of science's pervasive assimilation of its inherent questions throughout the ages. We are, if nothing else, honored and grateful for the chance to be a part thereof and we're committed to giving our utmost to be worthy of it. Another crucial and constructive difference between our pursuit here and any real-time questions regarding the practical situations applicable thereto is that we [already] know that the one person hears while the other doesn't as well as the explanations science and practical reason provide thereto and are still not content: still resonating some peculiar *why*? The questions to which the two first answers are proper answers [as part of the language games they consist in] necessarily arise out of authentic surprise and puzzlement concerning the

---

<sup>177</sup> The truly, nay – authentically, meta-existential language game. AI pioneers and visionaries also purport to engage in this language-game when they are not writing code.

<sup>178</sup> As an answer to that question rather than discovering something in the lab (reading scans or lab results).

<sup>179</sup> Trying to find out how things really are (philosophy) is not the same as trying to understand why something behaves the way it does (science). The latter opens the way to creative manipulation. The former opens no doors. The fundamental error of those who engage in the exploration of the “mental sphere” (except in the way psychology and psychiatry do) – of cognition/consciousness – is that they assume they engage in the latter while being absorbed – from head to toe, albeit unadmittedly – in the former.

same practical states of affairs. They consist in puzzles which [already] aren't puzzles for us. They take the form of 'does anything prevent him/her from...?', 'why is he/she unable to...?', 'is there any flaw...?' etc., the answers to which we already know. That puzzlement has – essentially has – purpose and function in the world – in what the inquirer will do and what he/she will become subsequent to the inquiry and the settlement of the puzzle at its core (e.g., the game – the language game – has a goal other than its result). Our language game is one *with* the world, whereby no other goals can pertain thereto.<sup>180</sup>

As we've seen above, after a while, namely after due training and sufficient practice with Mr. Eagleman's technological application, the hearing-impaired person will be able to say: 'I can hear' and 'I understand now what hearing is' or 'do you hear me?'. In other words, he/she will be able to perform immaculately regarding the semantic field of the notion of *hearing*, whereby we will concede in good faith that he/she *hears*. But what is it that enabled him/her to do so now? What is it that accounts – **from within, not from without** – for the propriety of all these sentences on his/her behalf and his/her proper immersion in the physical-biological substrate that underlies them? What explains the shift that occurred for him/her to be able to report and display a new capability rather than for us to be able to interpret a familiar sentence? What enacts the leap from nothing at all or even a mere, "inanimate" capability to correlate (discern that *this* is related to *that*) into absolute compliance with a whole aspect of the world – with the way things *are*? What accounts for that deep *transformation*?

The answer to that bunch of different forms of the same fundamental question should by now be the following: what ultimately enables, explains, enacts and accounts for this ability/propriety/leap/transformation are the activities pertaining to the training and practice with the technology at issue (corresponding to those which properly hearing people undergo and engage in in their early infancy).

Training upon *activities* is not 'training' upon a set. There are no sets or sets of sets that comprise a world. If only because a world is itself a metaphor. It is a language game.

Indeed, our hearing-impaired protagonist is now able to report and display a new capability rather than provide accurate answers to the question 'what is hearing?'. He/she is now capable of performing as an integral part of his/her environment rather than *relate* thereto. Will the far-end variants of present LLMs ever obtain the possibility, rather than the capability, to report and display something of that sort or perform **as such**? We kindly urge you to pause before answering. The pause may turn quite longer than anticipated (it is not a practical question).

Trees and animals need the biosphere to evolve as such, chemical elements need a physical substrate (the physical universe as such) to acquire their shape, living organisms

---

<sup>180</sup> This indeed resonates or evokes a peculiar kind of logic, as we purport to address that which we are part of on a level that assumes we are not. How this is possible evidently exceeds our goals and capacities here.

evolve in cell cultures and so on. The “cell culture” of the cognitive creature is not the ‘world’ – an already “established” world (which the biosphere, the physical substrate or cell culture *are*) – but the cognitive creature’s behavior in the world. We “create” through what we do, nevertheless creating what already *is*.<sup>181</sup> The leap from physics and chemistry into biology already foretells, prefigures and encapsulates all the horizons of advanced cognition. Biology *is* (already) an operating system from its very initial takeoff. The semantic realm – that “pictographic ontology” (discernible “from without”) we had alluded to in this paper above – will emerge therefrom, and the way this will happen is precisely and specifically by way of everything that organisms – including ourselves at the far end thereof – will *do*. Hence, **Why** LLMs do what they do, rather than merely *how*, is also essential (and maybe more so than anything else) to any delineation and understanding of their place and prospects in the world they appear – and are boldly professed – so overwhelmingly to impact. A poem (or any artistic oeuvre, for that matter), we should bear in mind although almost everything around us provokes and assails us to concede otherwise, is no poem but for its ground in human nature and human civilization – not as the sets of data through which they are reflected and represented in the World Wide Web, but as the cumulative, purposeful flow (the *actual* flow) they [really] are. So is all the rest.

We *do* ‘because’ rather than ‘for’. We perceive it as *for* [something]. But we ourselves – including our ability to apply that ‘for’ – are *because of* so much more. No other entity, we humbly assert, including the most formidable of our creations, can be otherwise.

## XVI

Given the conclusion at the end of the previous chapter, we should address its alleged contradiction to our observation in this paper above concerning the qualitative exceptionality of the ability to apply language relative to all our other qualities and capabilities in that it provides meaning to any sense – to the *fact* – of existence and cannot assume itself to linchpin its own “advent” (onset, emergence). Since this is where we, as it were, come full circle, we shall bring that observation effectively verbatim, culminating our previous apologies for repeating ourselves with one last solicitation for clemency, as our journey draws to a close:

In sharp contrast to each and all the aforementioned acts/capabilities (breathing, walking, reaching for something through our limbs), applying language, namely imbuing/infusing all that functional-performative riches with *meaning*, essentially “caters” for a state of affairs whereby things not only *proceed* by whatever inner reason or logic (physical,

---

<sup>181</sup> Perhaps the very ability to say that sentence or produce this insight *is* the [ultimate] perspective from without.

chemical, biological), “echoing” by that sheer flow whatever [essentially] abstract sense of causation or patterning or balance with nothing other than transparent “drive” and equilibrium to count on or be the case, but where the *act* or *phase* or *cause* or *result* or any other part or aspect of these “drive” and “equilibrium” are not only “vectors”, not merely units in a plane defined – and governed – by abstract factors or coordinates, but also stages in a kind of pictographic ontology, events and items with the fabric of *a world* rather than an abstract domain: the fabric of places rather than coordinates, of stones rather than groups of atoms or elementary particles, of brightness rather than flows of particles, of organisms rather than bunches of genes, of animals rather than organisms, of ‘why?’, ‘where?’ and ‘when?’ rather than infinite ‘now’ and formless ‘this/that’.

***This leap from performance in a realm (being in the world without language) to non-performative interaction with scenery, from sheer functionality to sheer narrative, appears to us as singular and therefore as a fundamental junction and potentially fruitful perspective for philosophical insight.***

But is there no equal leap between, say – the processes that take place in the physiological realm for what we call ‘breathing’ and the *ability* to breath or what the *breather* does upon breathing (the application of that capacity)?

The answer is No. There is a crucial difference between the two.

Physiological processes, including those of the brain and the nervous system, are not the same as their performative analogues/results in the world (walking, seeing, pain etc.). ‘Walking’ is not something feet (or any other organ or physiological process) do, and neither is ‘seeing’ performed by the eyes or ‘breathing’ by the lungs alone or in unison with the throat and mouth or anything else. Indeed, humans, namely wholesome agents/creatures in the world, perform all these.

But in terms of strict performance there is no other than a conceptual boundary, no substantive interface, between the underlying physiological bedrock and the performative wholesome agent – between the cells, tissues, fibers, muscles, tendons and the rest of it running at full throttle (physiology) and strolling in the park (agent), or between the whole light-retina-nervous-neural sequel at the bottom of seeing (physics/physiology) and the subjective experiential phenomenology that corresponds to that word (agent). In other words,

what we call and perceive and can identify as a human agent or person or creature, with all its corresponding capacities, is not something that emerges if not for the ability *to say something*, namely, apply language.

Only the ability *to say something*, or more precisely – the acquisition of that ability, doesn't have itself to rely on as it evolves.

In other words, although all these human capacities also appear to transcend the sum total of their physiological infrastructure, it is only a result of presupposing language. It is only a result of being able to think about them, namely relate to them – represent them to ourselves – through language, in the first place.

Which leaves us with the original enigma regarding the acquisition of that singular capacity (***how does it all “sum-up” to something like a word?***) and brings us back to the postulate that acquiring the capacity – the trait – of applying language is both exceptional and pivotal for any attempt to make sense of the world and our place in it, let alone to the exploration of 'meaning'.

Indeed, the leap from performance in a realm (being in a/the world without language) to non-performative interaction with scenery, from sheer functionality to sheer narrative, appears to be singular in the human realm.

Indeed, what we call and perceive and can identify as a human agent or person or creature, with all its corresponding properties/capacities, regresses into nothingness if not for the ability *to say something*, namely, apply language.

Indeed, only the ability *to say something*, or more precisely – the acquisition of that ability, does not have itself to rely on as it evolves.

Indeed, if not for language there is nothing at all to discuss or explain or, for that matter, to *start from* – neither this paper or any other meaningful exposition of the way things are, nor any conception of the universe or anything else in the first place.

So, it must be primary or conceptually (logically, “metaphysically”) antecedent to the world as such or any immersion therein of anything relating thereto, mustn't it?

No.

The actual ability *to say something*, including the real-time acquisition of that ability, nevertheless has something to rely on as it evolves. It is also grounded in our genetic makeup and its realization in what we do.

Thank you, Ludwig Wittgenstein, for inspiring us to pursue the unpursuable. Thank you for the profound notion of 'language games' and the chance – the porthole – to elaborate it in concordance with the advance of intellectual enquiry.

The leap from performance in a realm to non-performative interaction with scenery, from sheer functionality to sheer narrative, is not singular:

- An ability *to say something*;  
as against / towards
- *Non-performative interaction* with scenery;

The first is acquired by everything we *do* since we are born.

The second is ingrained in the first. It isn't acquired separately.

Non-performative interaction – awareness as such – is also an *act*. A 'state' from without (when we think about it). An *act* from within.

It is merely a reiteration – in real-time and in principle – of everything we already *said*.

The abilities to breath, walk or any other kind of performance are what they are due to our *non-performative interaction* with the world – due to our ability to perceive them (in real-time or through contemplation). But we are capable of the latter due to all the activities we've already engaged in since we were born and the possibility to synthesize them into inputs and outputs via our brain.

The latter instances of synthesis are also *activities* in the world ('house', 'boy', 'come!', 'why?', 'this is beautiful').

Our neural networks work, our biology complies, we perform.

The brain is a machine, the same as a single cell is.

The brain is the machine for the world to express itself (through which it does so). Not for *us* to do so (*we perform*, not *express*, therethrough). For the world.

If the previous sentence seems to border mysticism, we firmly insist that it doesn't. Endorsing our commitment to rationality, argumentative vigor and analytical transparency, we should be able to demonstrate that claim. To do so we shall return (as we said we will) to the thought experiment we conducted above in this paper regarding a person's experience of certain phenomenological events (be called by name or issued an order, the sound of the wind, his/her own breathing, a melody or rolling thunder) in perfect synchronicity with similar phenomenological events (another instance of naming or ordering, some other person's own breathing, some other melody or another event of rolling thunder) someplace else on earth. And here is the thrust and the logic – the "thoroughfare" – of the argument we propose: The latter [phenomenological "ambience"/setting/eventuation] wasn't/isn't less of itself by virtue of the fact that it wasn't/isn't experienced by our hero (just note how we indicated thereof; the language game we engage in here, please remember, is *calling things by their own names*). Hence, it was/is the world that "expressed itself", or 'played a note' in both cases, not our hero or

anything about him/her. The *world* which our hero had (or will have) experienced and the *world* he/she had (or will have) not, is one and the same world, whereby it expressed itself both in the instance our hero “absorbed” (was part of) and the one he/she didn’t (wasn’t part of).

What follows, and this is the crucial part, is that in the event **involving [the presence of] our hero** (the time and place when and where *I am* and/or *you/he/they are*) the world “expressed” itself both for him/her and not for him/her; both for the agent and irrespective thereof,<sup>182</sup> the same way as the cry of the hawk above me in the sky or the pain in my chest, is for me and not for me at the same time, whether I hear or feel it or not; whether I notice or do not notice it. It *is* for me even if/when I do not notice it and it *isn’t* for me even if/when I do.

Hence, it ***is for me*** (expresses *itself*) whether I notice it (anything) or not. And it ***isn’t for me*** (expresses *itself*) whether I notice it (anything) or not. Consequently, it is *it* – the world (the hawk, my footsteps or the person beside or in front of me and everything else, without exception, including the far end of the universe) – which *expresses itself*.<sup>183</sup> The logical imperative (the deductive strain) is very clear: it’s not that it persists where I’m not, but that it *speaks* where nothing *listens*; where no subjective tapestry “perspires” [it still *speaks*]. From beginning or no beginning to end or no end. The world in its uncleavable, uninterpretable (and only) sense.

The hawk is in the sky whether the fox or rabbit “migrate” into survival mode and hit for their urgent flight or not. Evolution is in the predator and the prey. The *world* precedes the whole triad – them both and evolution as such. Not as a frame or a realm. As a voice.

In the present paper we do not *hear*, we *listen*. Not as the person that hears his/her own breath or [some of] the audience in a concert hall (a certain language game). We listen as philosophers (quite another).

So do – we believe – mathematicians in their work.

QED.

## XVII

One final angle (no apologies at this stage):

I walk for survival (so as to survive). I hear my footsteps (as well as anything else) for survival (so as to survive). I still walk when I do not hear my footsteps or notice the fact that I walk [in any other way]. I *survive* (exist as a survival-oriented system) even when I

---

<sup>182</sup> Although he/she was “*there*”(!!), namely it ***isn’t*** for him/her even when he/she is there!

<sup>183</sup> If it isn’t for me when it eventuates (although I am my full case and it is its full case), it is *it* which is the unfolding ontological case.

do not “notice” it. Hence, by sheer logical inference, even (also) when I *notice for survival*, I survive<sup>184</sup> irrespective thereof (irrespective of noticing). Consequently, noticing *for survival* is also noticing *not for survival*. And what is that concurrent aspect of noticing *not for survival* while noticing *so as to survive*? It is the omnipresence of the world. And what is ‘*noticed omnipresence*’? A voice (one wishes to say: articulation). The world “stating” that *it is* (that it’s “there”) rather than merely being the case.

This is the sense in which a supernova in a faraway galaxy is also *for me*. *For me* and *not for me* at the same time. The world “stating” it is there rather than merely being the case.

I can walk without noticing I’m walking, I can breathe without noticing I’m breathing, and I can also *speak* without noticing I’m applying language. It isn’t “daydreaming”. It is *doing* in unison with the world. It is the world.

Is it, then, that not only elementary particles are nowhere and everywhere at the same time? Do large scale, complex systems also abide by the equations of quantum mechanics? This is a language game beyond the scope of this paper.

## XVIII

The brain governs our actions in absolute unison with the rest of environment.

Whether I sense/notice it or not.

The outcome is a symphony, not a process or an effect.

A symphony we are part of even when our specific instrument is silent (pauses).

For us it is tragedy, comedy and everything – including all achievements and regressions – in between.

Expression, in its essence (the only notion that has essence), has no functional goals (that’s why it’s an essence). We are merely infected.

Non-performative interaction – consciousness – is also an/the outcry of the world; the outcry of naked existence (of ‘isness’).

Egos find this difficult to confront. But they are drops in the sea. This is not our enterprise.

Language is singular from our standpoint, the standpoint of speakers. But they/we are also parts of the world – of naked existence. It isn’t singular from the latter’s standpoint. The standpoint of that which really and ultimately cries out.

The narrative isn’t ours.

---

<sup>184</sup> I am.

It – language – isn't antecedent to the world.

Nor is it a *product* thereof in any sense of that word that we can understand.

This is why we cannot model it (or 'meaning').

## XIX

The major obstacle for grasping the gist of the last few chapters stems from the traditional ways of thinking through our general theme here. What is usually asked regarding the point in question is: what will connect the apprehended symbol – the 'word' – with its recurrence in the agent's conscious theatre in a way that ipso facto retrieves the relevant "stuff" into that same theatre, as if from oblivion?

That kind of question initially sets the focus on agency at the subjective level – on what the agent *does* or is capable of doing in real-time, or, in other words, on the operating system at the core of the agent. Ultimately, it sets the focus on an object – not unlike any other object – in the world.

We have tried to show that this is not the right focus. It is a denominative rather than an explicative perspective. This is the way we talk and relate to our environment after we've already acquired that capacity and all the important stuff. If our goal is to figure out the nature of meaning rather than merely generate it on paper, the computer screen or in our minds, a different perspective is imperative.

That further perspective is not to be found in the erroneous way of posing the question, to which we already alluded above in this paper, namely: **what is it (or it is) like to be something?**, or more precisely: what kind of *being* 'to be like [something]' is?, which aims to singularize the human – consciousness generating – operating system as an object of inquiry. As we draw to a close, let's elaborate this point a little further.

Is it not clear that *to be like something* is inaccessible as an object of inquiry, if only because there is no categorical or any other distance<sup>185</sup> between the sentence

---

<sup>185</sup> No semantic complexity, which serves as a condition for any question (or any *saying*, for that matter) in the first place. It is completely, nay – absolutely, different from asking *what is...?* regarding anything else whatsoever or "what happens...?" or "who is...?" or "why is...?" which all presuppose a *there* – a distance or categorical separation between the question mark (rather than the questioner) and its subject – that is inherent to the question itself, namely ingrained in the very possibility to pose it in the first place. There is no such distance/complexity in the question 'what is it (or it is) like to be something?', namely between the kind of *being* 'to be like [something]' is and the act/fact of asking a question, since the latter – *asking a question* – is [itself] that very kind of being, namely to ask a question *is* to be like something. In other words, in the very [act of] asking that question one displays the quality he/she inquires of. Moreover, in the very [act of] asking that question one reiterates that which one purports to display wonder (uncertainty, puzzlement, curiosity) in relation thereto {can an expression or the very real-time display of curiosity project curiosity onto itself? Can doubt exert/project doubt onto its own unfolding? Can a question mark set a question mark onto the fact it has been set? Can my very saying 'I am Daniel' or 'give me an apple' or 'who

---

are you?’ result in doubt that I am Daniel or that I asked for an apple or inquired concerning your identity? Doubts regarding all three assertions may indeed thereafter rise, but all three genuine assertions/sayings cannot result in doubt or puzzlement regarding their own case, namely regarding the very acts performed thereby – asserting, requesting and inquiring. And in the same way *being like something* (e.g., like a question) cannot result in or constitute or take the form of a question concerning what it is (a question cannot result in or constitute or take the form of a question regarding the fact that it is a question). Eventuation does not have its own negation as a possible outcome thereof, whereby no doubt or question arises in its regard. And if one would say that it isn’t eventuation but the *nature* thereof he/she enquires into, we shall reply that his/her real-time enquiry is the very *eventuation* of a question mark – it is *to be like something* that happens, namely its very *nature* (metaphorically, we would call it: something that happens **from within!**) – and ask our interlocutor what kind of doubt he/she entertains, undergoes or proposes thereabout. The absence of any answer whatsoever to that last question settles the matter, unless our interlocutor has no clue of what he/she is talking about. That question – what kind of *being* ‘to be like [something]’ is? (which, as we’ve just seen, is not a genuine question) – emerges like a “scanned text/document”, so to speak (in the sense that there is no way to access what seems/appears as ‘live text’ in the document – to access it as formal text rather than abstract content). It is text which is actually a picture (a display **of** the world rather than **by** the inquirer). The ‘text’ in our case (the case of this [as if] question) is the entire process of acquiring the ability to apply language. The ‘picture’ is the ‘non-performative interaction’ referred to in this paper above. ‘Live text’ also acquires its meaning via the process of reading (a conscious agent), but the text in a scanned document had never been *written*, whereby nothing like ‘writing’ is applicable thereto and wherefore “reading” it has nothing to do with the way it had “emerged” (the nature of its emergence as such). So, what is “reading” a scanned document (or perceiving something that had not happened)? It is *reiteration* of something that *has* been done (or *has* happened). Hence, asking what ‘being like something is’ is a reiteration of being, not a reference thereto, least of all an enquiry thereabout. I can read a scanned document because someone had taken a photo, but being able to take a photograph or apperceive it has nothing to do with – is absolutely bereft of any interface with – how the photographed world came to be what it is, and therefore asking about the way I can apprehend a photograph – about the way it is meaningful to me – has nothing to do with the way I am in my actual, real-time environment, namely about the world as it is. This is why, knowing how scans had been made possible, we never ask what a scan “really” is. We [still and nevertheless] *do* ask what subjectivity is because, unlike the case of machines or characters in a story, we are the ontological scan – the pictured text. And we don’t know how we arrived there. So, what we are really puzzled about is not what subjectivity is (what it is *to be like* something), but how we emerge or have emerged as we are, namely how subjective agents emerge as such, which is the very question this paper addresses (what is the source/cause of semantic coherence?). In an attempt to provide an answer, we can [be tempted to] resort to the familiar scientific story of physics, chemistry and biology, the basic interaction of primordial creatures in an environment and the coextensive growth of the conceptive/perceptive operating system and the portion or volume of the environment which it projects. The point, however, is that this story has nothing to do with and is not applicable to (has no explanatory force regarding) the way we emerge as a scan (the emergence of agentic subjectivity as such, namely the fact we “find ourselves” *as such* – as being *like something* – at a certain point in time, even if this is also a gradual process). And it’s worth noting that the very grammatical combination of the words ‘be’ and ‘like’, namely ‘be like...’ already contains that crux, fusing ontology and phenomenology as some kind of axiomatic, impenetrable given. We can trace the origins of ‘being’, at least as far as the Big Bang, but we cannot trace the origins of the emergence of ‘like something’ to anything other than to that impenetrable given, whereby the question persists. As we have seen above, physics, chemistry and biology are the bedrock of the ‘live text’, including what we do and say as an integral part of the environment. Our interaction therewith is the live text, including all that we *do* in speaking and thinking. And we know or at least engage in a continuous and fruitful process of apprehending the ways it all evolves (rather than emerges). The only thing which as assuredly eludes us is the fact that we also emerge as a “continuous snapshot”. As an emergent and persisting scan. This is a real question. A true, legitimate and correct one. It emerges applying language consisting in action. The written text scanned at the copy machine doesn’t *act*. We do. From when we reach out to our mother’s nipple, through crying and moving our infant bodies and up to the pronunciation of our first words – the thread that maps and consolidates our “scan” are our actions in the world. Integrating therein through behavior is immersion, not a scan. We throw away Wittgenstein’s ladder. We’re not a scan. We have come a full circle.

embodying the question in its regard and the question mark at its end? Or because what it is *to be like* we are [in “real-time”, namely that which we are “really” asking about in that question], is an absolute given – something we are cast into (emerge *as*, whereby it is not the newborn who/which emerges as such; it is *me/I*<sup>186</sup>), not only without a clue as to a way of transcending thereof and/or obtaining a perspective thereupon but with no clue whatsoever as to what the “walls” or boundaries thereof – the confines/frontiers of subjectivity<sup>187</sup> – might be (except its own extinction; not biological death, which is a completely different thing, but the same kind of absolute dissolution/annulment the [absolute] original onset is)?

When two speakers interact, they don’t exchange words and ideas. Again, this is the way we talk. It is also not what we perceive observing such a situation. We just see them talk. When and if we try to explain it, we resort to the way we talk, and this is what comes out.

In this paper we haven’t engaged in the language game of *trying to explain it*. We engaged in [the language game of] listening to everything there is.

When two speakers interact what resonates is their shared field of activities in learning to speak in the first place. Only that.

Hence, when I ask: ‘do you hear me?’, I do not really inquire of something (about you or at all) at the time the question is posed but of something about the time you had become able to *do* [something], namely the time you acquired an ability.<sup>188</sup> Ultimately – about the time you became what you are. Words and their meaning(s) are actually original abilities, rather than items/entities we, as it were, “work with” in real-time.

When me and a person that had been trained to use Mr. Eagleman’s technology interact in discourse that involves ‘sound’, resorting, among other stuff, to the word ‘hear’, that word (as all other words) is a constitutive vacancy. At the point of our reciprocal

---

<sup>186</sup> To speak otherwise is not to make an observational mistake or to err in any other way about anything (regarding anything in – or any aspect of – *the world*). It is to engage oneself in the wrong language game. To engage in the wrong language game is not to err regarding anything in the world. It is a metaphysical bug (malfunction). All the other available language games are fusions, infusions and diffusions of the way we are, wherefore they are by definition *fusible* (“solvent” existential instances, so to speak). Only the language game embedded in that alleged question does not fuse [given the fissure – absolute break – between the sentence it is carried upon – the worldly *act* embodied therein – and the question mark at its tip / far end; between anything and everything the *act* can possibly rest on and its purported objective (which is also why it’s absolutely different from the aspiration to inhabit Mars or a distant galaxy or cure cancer, model a living cell or the human organism as a whole)].

<sup>187</sup> Subjectivity and “live”/actual subjectivity are one and the same. There’s no such thing as a ‘specimen’ or a model of subjectivity. There is no way (either phenomenologically or lingually) to relate or refer – to gaze upon – subjectivity, other than *be*. To think of or remember or imagine a past feeling or any other experiential state or to express it via language (e.g. ‘*I feel / had felt*’) is still – essentially and immanently – to *be* (to “live through” subjectivity). The same as there’s no ‘to burn’ without actually *burning*, only on an absolute scale. ‘Burning’ is the world’s subjectivity. Its voice. Us being as/what we are also is.

<sup>188</sup> The difference/distance between [an] ability and [the] ability is the timeless history – the timeless chronicles – of the world.

resonance, that word doesn't [resonate]. So much action had constituted that resonance. That word doesn't stand out therefrom. It's not that the interaction "leaves it out". The word is non-constituent thereto; immaterial – as a word – for the interaction to unfold.

Artificial operating systems do not, and given everything we presently know seemingly also cannot, interact in that way. **All they do is manipulate words and the gap between that and anything else they so awe-inspiringly do (e.g., AI robotics) is as absolute as ever.** To do otherwise they would have to evolve in a shared, comprehensive existential environment that has *expression* in its core. This is not where they evolve.

Will they ever? Or will a way be found to bypass that necessary condition? We do not know.

Consider the mental state of a human organism with severe congenial brain damage and severe neurological developmental disorders, which practically prevent any proper cognitive development stemming from the interaction of a healthy newborn with the environment. The experiential fabric of such an organism would amount to "nothing" compared to a healthy human's subjectivity. It will be a kind of random soup of experiential qualia, of the kind healthy infants are subject to just before their first cognitive capacities start to evolve (given a "bugless" genetic makeup). The point, however, is that it doesn't merely fall [completely] short of a healthy human's subjectivity. It also falls short – categorically so – from a [healthy] insect's or bat's perceptive world, which is the latter's *proper* perceptive environment, albeit "nothing" to a homo sapiens.

Insects and bats can engage in a categorically superior interaction with the environment (from a metaphysical, not an evolutionary, perspective) than a certain variant of the homo sapiens. Where does that leave our 'words', our grammar, our syntax? Behavior counts.

## Epilogue

I

Formulating thoughts about meaning<sup>189</sup> is hard, especially when the goal is the very formulation rather than communicating something. The attempt to do so, broadly understood as engaging in philosophy, is the only intellectual practice which a priori excludes the notion of 'mistake': an unsuccessful attempt does not amount to a flawed result but to nothing at all – to nonsense, which is not to say that it is worthless or devoid

---

<sup>189</sup> The fact or possibility of meaning rather than that which is meaningful in our lives.

of intellectual merit but that it does not comply with its own standard of [existential] veracity.<sup>190</sup>

The reason for this, as established at the outset and reiterated throughout this paper, consists in the fact that any articulation, including any of its explicative variants, should a priori meet the criterion of meaning, if it is to be an articulation at all, whereby the endeavor to discuss the foundational possibility of discourse – to presume one can speak about the very condition that makes speech possible/tangible – invariably leads to a logical paradox. The insistence to perform, nevertheless, an act that boils down to a logical paradox, is equivalent with doing nothing at all. At best, it is a case of *acting*. An act that consists in existential bluffing.

[Engaging in] philosophy as a hoax in the face of existence!

This is also why our principal objective and analytical approach in this paper are of a critical rather than a constructive nature and our main goal and ultimate ambition is neither to explicate meaning or suggest a “theory” of meaning, nor to discuss the obstacles and limitations or conundrums such an effort should inevitably run into, but, essentially, to point at some very grave and substantial conceptual and theoretical challenges the pretension to facilitate or “create” ‘meaning generation’ in the domain of artificial intelligence<sup>191</sup> should face, if it is to be at least theoretically viable.

Judging by the prevailing discourse in that field and industry and by its avowed goals, including among its opinion leaders, it forbears to confront this challenge. We maintain that such a refusal is equally fallacious in its confrontation with reality as any strictly philosophical (and merely reflective) endeavor to construct a theory of meaning or engage with meaning per se. Consequently, and quite ironically, the all too familiar, and much celebrated, antagonism between the contemplative theorist and the cutting edge of technological innovation collapses into a shared fallacy.

We did our best to avoid that trap. However, despite our circumscribed objectives and the critical nature of our argumentation, and however conceptually diligent we aimed to be, we could not escape addressing meaning and its semantic field directly as we progressed, if for no other reason than to bolster our critique of that notion. Consequently, navigating these conceptually murky waters risked veering into what may

---

<sup>190</sup> If I whisper while purporting to shout, I did not generate a faulty or inadequate vociferation (*failed to shout*) but did not shout at all. A person standing next to me may indeed tell me: “*you did not shout*” or “*that was not a shout*” or “*why do you whisper?*” in a sense that implies that my act was not equivalent to nonsense. But as far as the link between what I purported to do, or took myself as doing, and what I *actually did* is concerned, I did nothing at all. My act was devoid of any meaning. In Wittgensteinian terms, it was not the case that I engaged in the wrong language game (there is no such thing), but that I did not engage in any language game at all. Like trying to hit the ball with my feet and scoring a goal as part of a basketball game. Adhering to the rules of a soccer game while playing basketball is not a mistake. It is an existential bluff, namely nothing.

<sup>191</sup> In any domain, actually, where we *operate* rather than *exist*.

appear to be esoteric or mystical discourse. Hence, such assertions as: “*The brain is the machine for the world to express itself*” or “*In the present paper we do not hear, we listen*” or “*The narrative isn’t ours*” and so on.

We earnestly hope that the context surrounding these and similar phrasings throughout this paper is sufficiently evident to preclude any charge of esotericism. We hope to have clarified that our portrayal of the world as the ultimate “speaker” is not to be conflated with casual expressions such as “*the weather will decide*” or “*the stock exchange will have the final say*”, least of all with “*the universe communicates to me/you/us*”, but an insight stemming from a rigorous assessment of the comprehensive landscape of meaning generation. We further hope to have conveyed that unlike hearing, listening may culminate in an absolute applicative (entrepreneurial) stupor – in a practical dead end. Finally, we hope it is clear that taking the postulate that we are part of a whole seriously, namely in a way that excludes a perspective on the latter, does not commit us to a view by which human positive existence on planet earth – our world as we know it from history, memory, science, daily interaction etc. – is anyhow mystical,<sup>192</sup> or that suggesting consciousness is inaccessible as the seeming live text of a scanned document, we did not argue that humans are scans or photographs. Maybe we are something of something else besides or apart from or beyond our genetic origin or fountainhead. The point is that the latter ‘something’ is inaccessible to us by the very fact that we are a part – an echo – thereof.<sup>193</sup> That ‘of’, like ‘is’ or ‘am’ or ‘are’ (or their echoes in ‘I’, ‘you’, ‘he/she’ etc., in languages that go without copular verbs or “syntactic glues”), belongs to language but also to the world. The way it belongs to the latter – unlike our biological grassroots – we cannot know, let alone model.

---

<sup>192</sup> ‘Mysterious’ in the sense we speak of ‘the mysteries of nature’ – yes. ‘mystical’ in the sense of ‘transcendental’ – no.

<sup>193</sup> And compare:

*“But of this frame the bearings, and the ties,  
The strong connections, nice dependencies,  
Gradations just, has thy pervading soul  
Look’d through? or can a part contain the whole?”*

Alexander Pope, “An Essay on Man: Epistle I.” *YouTube*, uploaded by the Poetry Foundation, <https://www.poetryfoundation.org/poems/44899/an-essay-on-man-epistle-i>. This idea is obviously not original to us; we merely wield it against the pretensions we grapple with herein. Moreover, we are not *part* of our genetic origin but a result, an integral stage/facet thereof, an insight we arrived at not by means of examining “ourselves” but by examining biological reality. We are, however, barred from arriving at a similar insight regarding the world (in the sense of a ‘whole’) we are part of. It is the absolute empirical and conceptual abyss we encounter aiming to glimpse at the far end – the ‘beyond’ – of what it is to be ‘us’, namely of ‘us’, that warrants, nay – earmarks, that absolute inaccessibility. This is the quintessence of the above poetic extract, which resonated in one form or another by many profound thinkers since and before. We merely provided our humble analytic exposition thereto. And all those who are quick to highlight the discontinuity or fragility of that sense of “us,” often citing instances like deep, dreamless sleep or anesthesia, we encourage to remember that glimpsing beyond the ‘nothingness’ of those states is just as impossible as seeing beyond the otherwise flowing ‘something’. The presence of others, near or far, is irrelevant. In both cases.

Distancing ourselves from prevailing academic conventions and operating outside any concrete disciplinary creed, the cornerstone – or rather the compass, the intellectual lodestar, of our inquiry was but an unwavering commitment to our topic – to language and meaning as they arise in the world. Underscoring all the same our profound indebtedness to Wittgenstein, we tried to look as unaffectedly as we could into real-time, live language application, with a fresh focus on its acquisitive phases, from a contemporary angle. Having given our best, this paper offers our critical findings on what are arguably the most consequential developments shaping what it means to be human in contemporary times.

## II

We began this paper with the intent of building a case for the following three claims and four conclusions:<sup>194</sup>

- a. Ludwig Wittgenstein was right;  
**We believe we have shown that, indeed, he was.**
- b. The emergence of semantic coherence (learning what something is) has nothing to do with identification as such;  
**We believe we have shown that, indeed, it doesn't have.**
- c. Language is a property of the world. Human language is a variant thereof, not a “lens” onto the world or a category that applies thereto;  
**We believe we have shown that, indeed, language is a property of the world and not a lens thereupon.**  
**Therefore:**
  - a. What and how artificial intelligence *does* learn has nothing to do with what human learning *is*;  
**We believe we have shown that, indeed, the former has nothing to do with the latter.**
  - b. Insufficient understanding of what *we* are (of how we emerge as agents in a world) inhabits our understanding of what AI systems can be or be made to be;  
**We believe we have shown that, indeed, it so inhabits.**
  - c. Being what they are, AI systems by definition cannot mimic that which “builds” agentic properties (and agentic reality) in the real world;  
**We believe we have shown that, indeed, they cannot.**
  - d. Artificial General Intelligence (AGI), in the sense of being more than an artifact, is in principle unattainable, no matter how formidable and elusive that which *is* attainable may be;

---

<sup>194</sup> Repeated here in their abridged form.

We believe that in the discussion above we have also established that last conclusion, at least as far as the current level of human knowledge allows. Abiding by the standards we apply to others, however, we stop short of laying any claim to prophecy. We actually do not know what the far end of the current surge of artificial intelligence may turn out to be. We just think that the relevant puzzles – grand and inspiring puzzles – go beyond everything that happens inside the “black box” of modern algorithmics.

A future where an inscrutable ‘black box’ transcends all human expertise – embodying the ultimate healer, theoretician, empath and so on – still offers no progress toward solving the profound existential riddle that intrinsically informs our worldview.

Better and better embedding and exponential integration within evolving high-dimensional spaces is not the answer. It is irrelevant to the questions.

True philosophy, on the other hand, is relevant and applicable not merely to such questions as what shall we do with – or how shall we go about<sup>195</sup> – the genie that emerged from the cybernetic bottle. It is applicable to the bottle and the genie and the entire story they tell.

Perhaps this understanding is the very crucible in which our self-assigned standing within the existential theatre will finally shed both its philosophical artifice and its practical illusions.

---

<sup>195</sup> Or how shall we look at ourselves and our species in consequence of –