Review Article

The Mind-Brain Relationship as a Mathematical Problem

Giorgio A. Ascoli

Krasnow Institute for Advanced Study, George Mason University, Fairfax, 22030-4444 VA, USA

Correspondence should be addressed to Giorgio A. Ascoli; ascoli@gmu.edu

Received 11 February 2013; Accepted 7 March 2013

Academic Editors: J. A. Hinojosa and S. Rampp

Copyright © 2013 Giorgio A. Ascoli. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

This paper aims to frame certain fundamental aspects of the human mind (content and meaning of mental states) and foundational elements of brain computation (spatial and temporal patterns of neural activity) so as to enable at least in principle their integration within one and the same quantitative representation. Through the history of science, similar approaches have been instrumental to bridge other seemingly mysterious scientific phenomena, such as thermodynamics and statistical mechanics, optics and electromagnetism, or chemistry and quantum physics, among several other examples. Identifying the relevant levels of analysis is important to define proper mathematical formalisms for describing the brain and the mind, such that they could be mapped onto each other in order to explain their equivalence. Based on these premises, we overview the potential of neural connectivity to provide highly informative constraints on brain computational process. Moreover, we outline approaches for representing cognitive and emotional states geometrically with semantic maps. Next, we summarize leading theoretical framework that might serve as an explanatory bridge between neural connectivity and mental space. Furthermore, we discuss the implications of this framework for human communication and our view of reality. We conclude by analyzing the practical requirements to manage the necessary data for solving the mind-brain problem from this perspective.

1. Introduction

The relationship between mind and matter has been a fundamental topic of investigation in many if not all cultures and traditions since the most ancient records of human thought, from the Hindu orthodox school of Sankhya nearly 27 centuries ago [1] to the classic Greek philosophy of Plato (e.g., in the dialogue *Phaedo*) 300 years later [2]. With few exceptions (most noticeably that of panpsychism: [3]) most theories of the mind throughout history related it to the body and its various parts, including the heart in Aristotle's view [4] and the endocrine pineal gland in the work of Descartes [5]. Early physicians Hippocrates [6] and Galen [7] were among the first influential proponents of the central role of the brain in the operation of the mind based on anatomical and physiological observations. The development of modern neuroscience led to the (still ongoing) accumulation of massive evidence that irreversibly linked the mind to the brain [8].

The goal of this spotlight paper is emphatically *not* to provide an extensive review or even a balanced perspective of

the enormous body of work on the brain-mind relationship in cognitive philosophy. To appreciate the breadth and depth of this topic, we refer the reader to a sample collection of over 200 articles on consciousness and neuroscience available online [9]. Instead, we selectively review a set of specific topics in neuroscience and cognitive science together supporting the notion that, within a particular interpretation, certain aspects of the mind-brain relationship can be framed as a rigorously defined and in principle soluble mathematical problem. In order to build this argument, we first describe in the next two sections the somewhat delimited facets of the mind and the brain we aim to bridge together. Next, we explicate what in our view would count as "solving" the mind-brain problem. Then we elaborate on relevant general topics in the investigation of the brain and of the mind that will likely play a central role in a satisfactory explanation of the mind-brain relationship. Finally, we overview a sample of more specific available frameworks in neuroscience and cognitive science that appear particularly promising among the known existing candidates to help crack the mystery of the link between mental and neural activities.

2. Declarative Mental States: Content and Meaning

The term “mind” is commonly employed to signify a broad variety of connotations even within the scientific discourse [10, 11]. When referring to the mind-brain relationship, mind is most often taken to stand for human consciousness [12]. Consciousness is itself challenging to define, which may be viewed as a puzzling paradox, considering that it constitutes perhaps the most immediately and intimately accessible characteristic of the life of every person. Consciousness cannot even be taken as a trivial or necessary feature that is always present in our existence, because we all experience the transition of fading into dreamless sleep and awaking from it tens of thousands of times in the course of a typical life span. The distinction between inner mental content and outer behavioral observables is equally apparent when considering the obvious difference between dreamless sleep and a vivid dream [13]. Judging from the paralyzed body (except for eye movement), an external spectator could not even begin to guess the mental state of a dreaming person [14]. In contrast, we often toss, turn, kick, twitch, mumble, and moan during dreamless sleep [15].

Here and in the rest of this paper we do not consider the legal or medical definitions of being conscious as opposed to unconscious, because the shared goal of physicians and lawmakers is to evaluate other people rather than themselves. Thus, they must necessarily take an objective, third-person perspective, that is, one entirely based on behavioral evidence. In contrast, what makes mental states so hard to relate to the brain, the body, and matter in general is precisely their being felt from “within”, subjectively, or from a first-person perspective [16]. Similarly, there are clear resemblances between the properties of consciousness and those of attention [17]. Neuroscience has advanced considerably in the elucidation of the neural mechanisms underlying attentive processes (e.g., [18]). Nevertheless, recent evidence clearly indicates a double dissociation between attention and consciousness [19].

Excluding the immediately cyclical definitions (such as equating consciousness to being aware and awareness to being conscious), one of the simplest approaches to describe consciousness is to define it as “what it feels like being” [20]. Thus, a particular conscious sensation (such as the taste of chocolate) or thought (such as remembering someone’s name) would just correspond to what it feels like to taste chocolate or what it feels like to remember that person’s name. The challenge of the mind-matter relation, then, is to understand how particular configurations of matter (e.g., a person’s body or brain) or material processes (the dynamic activity of that body or brain) systematically coexist or cooccur with feeling particular mental states [21]. The ultimate “hard problem” is to explain why any material entity should feel like anything at all.

Many aspects of conscious experience can be reported explicitly by the subject [22, 23] either internally, by inner dialogue or memory rehearsal, or externally, such as by a verbal description or numerical rating. A defining characteristic of such *declarative content* of the conscious mind is that it can be

communicated. The most common form of communication among human beings is through spoken or written language, but declarative content can also be communicated by other means, such as body gesture, drawings, movies, and many others. Independent of the signs or “signifiers” used (e.g., words and sentences in language), the declarative content they communicate is what we call their *meaning*. The field of semantics studies meaning and its communication by signifiers.

Consider a subject experiencing a certain declarative mental state (e.g., anger) and communicating it to another person “I’m mad at you!”. This message informs the recipient of the sender’s conscious content. A fascinating facet of declarative communication is that the content intended by the sender is generally not identical to that understood by the recipient (e.g., why and how seriously is s/he angry?). We will return to this mismatch at the end of the paper.

3. The Brain: Spatial and Temporal Patterns of Neural Activity

Overwhelming scientific data unequivocally link the conscious mind to the brain, as opposed to matter at large. The general consensus on this interpretation in the research community boosted the popularity of the materialist credo that “the mind is what the brain does” [24]. Available evidence, however, is more specific than generically pointing to the brain as the most relevant organ of the body in relation to the mind. Particular regions of the brain are more directly and intimately involved with mental activity, including the thalamocortical system [25] and the hippocampal complex [26]. Other parts of the central nervous system, such as the cerebellum and the spinal cord, appear less pertinent to the declarative content of the mind.

The majority opinion is that what renders the parts of the brain relating to the mind “special” is not the intrinsic substance they are made of but their functional organization. In fact, those same thalamocortical and hippocampal regions are definitely active even during dreamless sleep, but their activity is different from that during conscious states [27]. Thus, it is the information process, or computation, carried out by these brain structures, that most closely corresponds to consciousness [28, 29]. Implementing that same computation in different mediums, in artificial machines, or in virtual environments, would then in principle also result in conscious experience.

Unfortunately, however, what might in fact constitute the essence of this computation, or its neural implementation, is far less clear. Proposals range from involvement of broadly distributed oscillatory rhythms ([30]; see, however, [31]) to activation of highly specific “concept neurons” [32]. For the purpose of this paper, we do not need to dwell into specific hypotheses, except noting a general consensus supporting spatial-temporal patterns of neural activity as key material correlates of conscious experience. The carriers of neural activity are neurons, biological objects each occupying a physical volume in the brain. Thus, the complete electric and chemical record of every small portion of each neuron and

their connections in the entirety of the appropriate brain networks throughout the course of an arbitrarily large number of declarative mental states would in principle include all the necessary information about “the brain” that is relevant to the content of the mind [33].

Within the particular framework we have chosen to adopt, the key players of the mind-brain relationship can be summarized as the following. On the one hand, each of us has access to the immediate experience of what it feels like to be, in any one moment of our conscious existence and throughout our life span. We can communicate that experience to each other, and our mental states are altered by receiving those messages. On the other hand, the brain is a dynamical information processing system, consisting of a large number of neurons, each characterized at any one time by the spatial distribution of its electrochemical state. Within a brain, neurons continuously exchange signals with each other, mutually affecting their spatially distributed electrochemical state.

In this framework, the mind-brain problem can be formulated as a simple question. Why and how do certain spatio-temporal patterns of neural dynamics relate to declarative mental states? In other words, why do those “neural dynamics that feel like something” feel like those (or any!) experiences?

4. What Would Count as an Explanation?

Many relationships are routinely observed in the course of everyday life, such as how the weight of bodies relates to their size or how satisfying it feels to drink when thirsty. Why is the mind-brain relationship so different as to constitute a *problem*? The main challenge seems to be that typically related observables are somehow of the “same kind.” For example, size and weight are both physical attributes of material objects, while thirst and satisfaction are both subjective sensations of minds. Brains and minds are so different from each other as to inspire several kinds of dualist philosophies in the course of time. This fundamental difference has been considered challenging by scientists and philosophers alike, from confronting possible inconsistencies in temporal dynamics [34] through a retreat to the acceptance of limited explanations [35], to overt calls for giving up scientific accounts altogether [36].

The more extreme positions maintaining that there is no solution to the mind-brain problem have been refuted both on philosophical [37] and scientific grounds [38]. The key issue is in fact to understand what would count as a solution. In other words, if the mind-brain relationship can be explained, what type of explanation is being sought? To answer this central question, it is useful to consider previous scientific breakthroughs that resulted in satisfactory understanding of relationships that had been considered difficult to explain before. The most illustrative examples would concern discoveries relating (sets of) observations that appeared completely disparate from each other, linking phenomena seeming of “different kinds.”

A prime case in point is the explanation of thermodynamics in terms of statistical mechanics. The 1687 publication of

Newton’s *Principia* [39] consolidated a century of gains from the scientific method introduced by Galileo [40] into a coherent and complete theory of mechanics. Newton’s famous laws relating force to acceleration and defining the mutual attraction of masses by gravity provided an excellent description of planetary motion as well as accurate predictions of interactions among material objects. More than one hundred years later the first comprehensive treaty on thermodynamics was published [41], describing the relationship between heat, energy, temperature, and what was later called entropy.

Thermodynamics appeared completely independent of Newton’s mechanics which explained instead how bodies move and respond to forces. While both mechanics and thermodynamics were (and still are) recognized as landmark advancements of scientific progress, they seemed to describe phenomena of different kinds. Yet, some relationships, discovered earlier by the likes of Gay-Lussac, Avogadro, and Boyle [42], consistently crossed that divide, such as the proportionality of gas pressure (a mechanical attribute) and temperature (a thermodynamic one) in a container of fixed volume. Moreover, simple observations remained challenging to reconcile on both sides of the divide, such as why matter can apparently be heated up to arbitrarily high temperatures but not cooled off below a certain limit (“absolute zero” or -273.15°C) and the continuously jittery (“Brownian”) motion of particles visualized under a microscope [43].

Between 1856 and 1871, Kroenig, Clausius, Maxwell, and Boltzmann developed a “kinetic theory” explaining thermodynamics in terms of mechanics, based on pioneering ideas originally proposed more than a century earlier [44]. Specifically, temperature relates to the average quadratic velocity of a large number of microscopic particles and pressure to their momentum. These relationships intuitively match the observation that, when warming up one’s hands by rubbing them together, the faster the movement, the higher the generated heat. Entropy is related to the number of possible states a system can be in, which clarifies why disorder tends to increase in the absence of other constraints. The demonstration of the equivalence between thermodynamics and mechanics is as convincing and direct as the derivation of the corresponding thermodynamic and mechanical laws. In particular, the law of gas can be derived from Newton’s formulas. Furthermore, this link solved the mystery of the lower bound of temperature: there is no practical limit to the maximum speed that particles can have, but they cannot go any slower than being perfectly still (corresponding to absolute zero). Notably, the kinetic theory also led to the famous 1905 Einsteinian explanation of Brownian motion (recounted in [45]).

Another striking example of scientific advancement linking phenomena “of different kinds” is the explanation of optics in terms of electromagnetism. The phenomena of refraction and diffraction are fully explained by Maxwell’s equations of electrodynamics once light is understood as electromagnetic wave. Yet the properties of mirrors and the passage of light through various media (such as air and water) seem so different from the phenomena of electric current and magnetic dipoles. Other illustrations of the same relationships are the quantum physics foundation of chemistry, the

genomics bases of genetics, and the explanation of neuronal firing in terms of voltage-dependent sodium and potassium channels [46].

When Mendeleev compiled the first draft of the Periodic Table of the Elements, there was no physical justification of the observed proportions of chemical reactions: the truth and correctness of the Table relied on its predictive power and descriptive elegance. Mendeleev's few initial "mistakes," most of which he himself corrected in subsequent versions, were due to using the atomic mass (the observed "chemical weight" of a substance) rather than the atomic number (the number of protons in the nucleus) as the organizing principle. The atomic nucleus was discovered 42 years later, just 4 years after Mendeleev's death. Thus his corrections, required to describe then-available data parsimoniously, de facto predicted the atomic number nearly half a century prior to its actual discovery. The subsequent development of quantum mechanics showed that the atomic number emerged as a mathematical necessity from more fundamental assumptions, such as Schrödinger's wave equation or Heisenberg's uncertainty principle. Today we readily accept that many chemical phenomena (e.g., the tendency of metals and halogens to combine, as well as their "valence" ratios) are ultimately "explained" by quantum mechanics, even if these two phenomena would appear to relate to different and independent aspects of reality.

Is there a common thread in these seemingly disparate, if illustrious, precedents? We purport that the illusion of mystery at one level (e.g., the law of gases, light diffraction, and the combining ratios of chemical elements) evaporated as soon as those same sets of properties were discovered or demonstrated at a different level (particle kinematics, electromagnetic waves, and quantum orbitals). Breakthroughs of these sorts are fundamentally distinguishable from hand-waving "just-so" speculations because once the parallel is drawn, it systematically explains and predicts an entire body of observations rather than isolated phenomena. Here we surmise that the content and meaning of mental states, the most inescapable yet ineffable puzzle of human cognition, will eventually be understood as a direct reflection, if not simply an aspect, of brain computation, much like thermodynamics is statistical mechanics. Given the complexity of brain dynamics and of our inner experience, we expect the mathematical constructs connecting neuroscience and mind to be far less simple than Newton's, Maxwell's, or Schrödinger's equations. The principle of what would count as an explanation, however, remains the same: to solve the mind-brain conundrum one "only" needs to show that the math of the mind and the math of the brain are equivalent.

We are not simply proposing mental properties to be probabilistically supervenient on brain properties, that is, that they can be inferred statistically from brain measures within any given error rate [47]. On the contrary, we are asserting the possibility of a formal equivalence between the two, through all temporal scales and plastic changes [48]. The explanatory power of mathematical theory in neuroscience is recognized in principle [49], but the extent of its reach has not yet been fully realized, and the path forward has never been chartered before. This is in stark contrast to the third simulation-based

leg of scientific progress (complementary to experiments and theory), which is blossoming into maturity in computational neuroscience and cognitive modeling alike [50] and in the study of consciousness in particular [51, 52].

5. Neural Connectivity as the Most Informative Constraint in the Brain

To explain the equivalence of brain and mind by mapping them onto each other, it is essential to identify the relevant levels of analysis in order to define proper mathematical formalisms for their quantitative description. We start from the brain in this section and tackle the mind in the next. Nervous systems are gigantic networks of intercommunicating neurons. From the computational point of view, it matters relatively little that neurons are electrical devices. Instead, brain signal processing is fundamentally dependent on circuit connectivity [53]. Specifically, how neurons are connected to each other constrains network dynamics [54] and therefore determines the possible flow of information transmission [55].

A human brain has an estimated $\sim 10^{11}$ neurons and 10^{15} synapses [56], which could in theory form a humongous number ($\sim 10^{10^{16}}$) of distinct connectivity patterns, between a googol (10^{10}) and a googolplex ($10^{10^{100}}$). Because real brains are wired to a certain degree according to stereotypical principles [57], the actual number of connectivity patterns that could be found in any one human brain is undoubtedly lower. However, brain circuitry is neither random nor regular, and the information content of a single human brain remains far greater than the number of fundamental particles in the whole universe, let alone just the complete biochemical specification of that individual brain. Thus network connectivity is necessarily more informative than the entire molecular profile of each of all of its neurons, including the expression of every gene and protein constituting the biophysical machinery at the basis of neuronal electrophysiology.

While neuroanatomy provides the foundational roadmap of information transmission in nervous systems, neural activity is itself characterized by chaotic dynamics [58] typical of complex systems [59, 60]. As these aspects are particularly relevant to conscious brain function [61], a full understanding of the brain as it relates to mental content will have to integrate adequate accounts of both neural dynamics and connectivity [62–64]. Nevertheless, the network architecture specification is absolutely central to the assumed correspondence between spatial-temporal patterns of neural spiking and mental states.

The notion of "connectomics," characterizing the circuit blueprint of the nervous system [65], has progressively grown close to practical feasibility with the recent dramatic advancements in genetic manipulations allowing for multicolor microscopic visualization [66, 67]. Two distinct connectome concepts refer, respectively, to the cellular level, or "synaptome" [68, 69], and the regional level, or "projectome" [70]. The former is further distinguished in the dense reconstruction of the entire synaptic matrix and the statistical potential of synaptic connectivity, both highly

relevant to computational processing [71–73]. In contrast, the much coarser description of regional connectivity has less direct implications for a mechanistic understanding of brain cognition. However, this latter approach is also substantially more realistic to achieve in the near future, using existing histological techniques in animal models [74, 75] or noninvasive imaging in humans [76, 77].

A number of “big-science” as well as grass-root data acquisition efforts are underway in both cellular and regional connectomics. These include the Human Connectome Project [78, 79], the 1000 Connectomes Project [80, 81], the Mouse Brain Architecture Project [82], and the FlyCircuit Database [83], among many others. This flurry of developments along with nonconventional approaches (e.g., [84]) is generating a wave of optimistic expectation in the research community that massive neural connectivity data will become increasingly available in the foreseeable future.

The branch of mathematics dealing with connectivity is graph theory. In light of the previous considerations, it is not surprising that graph theory has become a considerably popular topic in neuroscience (e.g., [85, 86]). It is remarkable that important properties of general graphs that have been found to apply to many types of networks, including random connections [87], small-world attributes [88], scale invariance [89], and motif distributions [90], are prominently relevant to neural circuits [91–94].

The application of graph theoretic analysis to neural circuit has already revealed a number of features, including network communities [95] and rich clubs [96], but also general principles of wiring economy [97] and network organization [98] as well as potential implications of circuit structure on signal communication [99, 100]. It is important to stress that, while two cells are never exactly alike, neurons can be organized in distinct classes such that neurons within each class are much more similar to each other than across classes [101]. Thus the statistical properties of brain connectivity are likely to be strongly determined at the level of connection probability among neuron classes. Initial progress is being made in the application of the relevant field of mathematics, stochastic block modeling, to this problem [102, 103].

Two further facets are worth considering in the characterization of the brain in terms of its network connectivity. The first is the all-important issue of intersubject diversity. While in invertebrates it is sometimes possible to recognize the same individual neurons across subjects, in mammals it is not even possible to match the same types of neurons bilaterally within subject, such as in motor neurons innervating symmetric muscles [104]. In humans, intersubject variability is already very considerable at the regional level [105] and can be expected to be extraordinary large at the level of individual neurons across subjects.

The second vital element of brain circuitry is structural plasticity, that is, dynamical changes in the synaptic connectivity not just during development but throughout adulthood. In the cortex of normally behaving mice, for example, 4% of axonal boutons change over the course of a few months [106], with similar proportions reported in dendritic spines [107]. Abundant experimental evidence suggests that this

form of plasticity is activity- and experience-dependent [108–111]. This is just one of many mechanisms underlying neural plasticity across spatial and temporal scales, from short- or long-term alterations in synaptic strengths to neurogenesis [112], which are believed to support memory storage [113, 114].

Much as the brain is in constant flux and its functional connectivity continuously changes with every spike and synaptic discharge, so is the mind never exactly the same before and after instantiating each and every subjective representation. These aspects of both brains and minds are strongly resonant with Heraclitus’ “panta rhei” [115].

6. Quantifying Declarative Mental States: Semantic Maps

One may find it disturbing that despite knowing so much about the structure and activity of the brain, we have a hard time “guessing” what it is thinking. But it is somehow even more peculiar that in spite of direct, detailed, continuous, and complete access to each and all conscious mental states we experience, we find it difficult to describe them comprehensively, let alone quantitatively. Indeed, it seems absurd that we can measure the concentration of Substance P in single neurons to the fifth significant digit; yet we can only measure the resulting sensation of pain semiquantitatively on a 7-point discrete scale. In order to bring the study of conscious content into the realm of hard science, we need to devise a quantitative measurement system for subjective states [116].

Language has often been considered a convenient proxy to access mental states, if not the most direct tool to describe them. The scientific characterization of the meaning of language, or semantic analysis, has a long history and remains one of the most active research areas in (computational) linguistics. Here we do not aim to review or even to provide a balanced commentary on the state of the art of semantic analysis techniques. Instead, we introduce and explain a very specific, nonconventional approach to this problem that is particularly pertinent to the topic of this spotlight paper.

Most if not all of the best known computational methods of semantic analysis are based on (variations of) the common principle that the meaning of words relates to the contextual occurrence of their use in language [117]. For example apples, oranges, and grapes tend to be used in similar contexts as reflected by their cooccurrence with similar words in the same sentence (e.g., eat, ripe, juice, and vitamin). Thus, they share similar semantics (they are all types of fruit). The notion that word meaning relates to the (relative) frequency of their cooccurrence is shared by many broadly adopted approaches, including Latent Semantic Indexing [118], Latent Dirichlet Allocation [119], Hyperspace Analogue to Language [120], and many others [121]. In practice, these techniques rely on the identification of statistical patterns of word usage in large-scale text corpora by computational parameter extraction.

Although the details vary among types of computational semantic analysis, words (or more generally, concepts) are often allocated to a multidimensional abstract space such that the location of each concept reflects its meaning. These spaces are sometime referred to as “semantic maps.” Meaningful

semantic dimensions of these spaces can be associated with their geometrical shape, a bit like the location on a globe can be described with the polar and azimuth angles, or the size of a cigar can be described by its length and thickness. Alternatively, meanings can be identified with clusters of words in this space. For instance, all fruit words in the previous example would be located in the same region of the space. By nature of its own principle, latent semantic analysis and its variations generate results that are highly context dependent. In other words, the semantics extracted from a cookbook are typically quite different from those detected in movie reviews or obituaries. In fact, use of nonhomogeneous collections of corpora from different domains typically fails to yield meaningful semantics. Moreover, this general class of methodologies tends to produce a large number of highly specific dimensions.

A rather complementary (and historically precedent) goal of lexical semantics has been to seek the fundamental (or at least context independent) dimensions of word meaning. Perhaps the most seminal study in this sense has been that of Osgood's "semantic differentials" [122]. In that work, subjects were asked to rate a large number of words in various hand-picked dimensions defined by two opposite extremes (e.g., soft-hard, fast-slow, clean-dirty, valuable-worthless, and fair-unfair), using a Likert (discrete) scoring scale. Subsequent analysis identified three principal dimensions that were robust to cultural and geographical differences, namely, evaluation (also known as valence: good-bad), potency (strong-weak), and agency (active-passive). A limitation of these studies and other similar psychometric approaches [123] is that they involve human subjects and arbitrary choices of starting terms. Thus, they are not amenable to automated, high-throughput computational extraction.

Word meaning has of course also been characterized for thousands of years in many languages and cultures through the creation of dictionaries. Here the beginning of modern times can be considered to correspond to the systematic (but again, ultimately arbitrary) classification introduced by Roget's Thesaurus of English words and phrases, now accessible online 160 years after its original publication [124]. Among contemporary efforts, the most comprehensive academic resource is Princeton's WordNet [125]. Researchers in computational linguistics are vigorously pursuing the topic of conceptual ontologies [126]. Yet, it remains to be established if and how formal ontological theories could map semantic spaces such as those generated by latent semantic analyses.

We have recently introduced a novel paradigm for semantic mapping that allows systematic construction of a low-dimensional metric system for the context-independent ("fundamental") meaning of words [127]. Our approach combines certain elements of Osgood's work (use of antithetical meanings) with the scalable use of "objective" dictionaries. Specifically, using a novel self-organization process, we constructed a semantic map of natural language that simultaneously represents synonymy and antonymy. Synonyms and antonyms are commonly listed in dictionaries for most terms. We extracted these relationships from digitally accessible dictionaries (Microsoft Word and Princeton's WordNet) in each of several languages (English, French, German, and

Spanish). For each dictionary and language, we initially allocated words at random locations in a finite, multidimensional spherical space. Then we started moving the position of every word following a simple rule: every word would "attract" its synonyms and "repel" its antonyms. Thus, pairs of synonyms would tend to move closer to each other, and pairs of antonyms would move farther apart (within the bounds of the multidimensional sphere).

This process "converges" in the sense that all words reach a stable position that could not be further improved in terms of proximity to synonyms and distance from antonyms. The resulting space only had a limited number (~4) of statistically significant dimensions. This means that, even if the starting space is a homogeneous sphere of many (~100) dimensions, the resulting emergent shape can be quite completely described with just four numbers. Most importantly, the emergent semantics of the map's principal components are clearly identifiable: the first three correspond to the meanings of "good/bad" (valence), "calm/excited" (arousal), and "open/closed" (freedom), respectively. The semantic map is sufficiently robust to allow the automated extraction of synonyms and antonyms not originally present in the dictionaries used to construct the map, as well as to predict connotation from their coordinates.

The map's geometric characteristics include a bimodal distribution of the first component, increasing kurtosis of subsequent (unimodal) components, and a U-shaped maximum-spread planar projection. Both the semantic content and the main geometric features of the map are consistent between dictionaries, among tested Western languages, and with previously established psychometric measures. Some of the mathematical formalism and speculative interpretations are elaborated in a second follow-up paper [128]. Interestingly, the main emerging dimensions of this semantic map loosely correspond to the primary modulatory neurotransmitter systems in the mammalian brain [129].

The previous paradigm can be expanded with appropriate adaptations to extract additional, independent dimensions of word meaning by considering other linguistic relations besides synonyms and antonyms. In particular, the relationship "is-a" (hypernyms and hyponyms) captures the abstractness (or more precisely the ontological generality) of words. For example, the statements "Mickey is a mouse", "the mouse is a rodent", and "a rodent is an animal" reflect a hierarchy of concepts from the more concrete, or rather specific (Mickey), to the more abstract/general (animal), whereas rodent is a hypernym of mouse and a hyponym of animal. In this paper we refer to this property as "abstractness" because "generality" might be confused to indicate how common a term is (usage frequency).

Is-a relationships are commonly used in dictionary definitions following the classic recipe of Aristotle's Logic "A is a type of B with property C" [130]. However, hypernyms and hyponyms are seldom listed in immediately machine-readable form in digital collections, the way synonyms and antonyms are. One exception is provided once again by WordNet, which explicitly provides is-a relationships among many of its terms. Unlike synonyms and antonyms, which are symmetric relations (if A is synonym of B, B is synonym of

A), hypernyms and hyponyms are directional and mutually antisymmetric (if A is hypernym of B, B is hyponym of A). We thus changed the form of the energy functional in the previously described optimization procedure [131].

The resulting allocation of words in space yielded a ranking of all terms along a single dimension, that is, a simple scalar measure of their abstractness (ontological generality). For example, the ten “top” scoring terms (out of ~124,000) are entity, physical entity, psychological feature, auditory communication, unmake, cognition, knowledge, noesis, natural phenomenon, and ability. The bottom (11 ex aequo) of the list reads *Edmontonia*, *Coelophys*, *Deinocheirus*, *Struthiomimus*, *Deinonychus*, *dromaeosaur*, *Mononykus olecranus*, *oviraptorid*, *superslasher*, *Utahraptor*, and *Velociraptor*.

One advantage of a metric system is that words can be compared in terms of their abstractness even if they are unrelated, for example, the term “governance” (whose abstractness value is 1.726) can be determined to be more abstract than “newspaper” (abstractness of 0.541) even if there are no (sequences of) is-a relationships connecting them. Moreover, because the measure is quantitative, it allows evaluation of relative comparisons. For instance, “power” (1.805) is more abstract than “revolutionary” (1.070). However, the generality of “governance” relative to “newspaper” (abstractness difference: $1.726 - 0.541 = 1.185$, abstractness ratio: $1.726/0.541 = 3.190$) is greater compared to that of “power” relative to “revolutionary” (abstractness difference: $1.805 - 1.070 = 0.735$, abstractness ratio: $1.805/1.070 = 1.687$). This opens the possibility to establish a probabilistic estimate of whether a word is more abstract than another.

The metrics of context-independent word meaning along the principal dimensions described previously can be applied to characterize declarative mental states. The most straightforward application is to quantify the content of verbal examples along the main dimensions of the map. This can help in relating semantic content to neural signals. It should be noted that the semantic map described here represents a complementary, rather than alternative, tool to more established latent semantic analyses. While maps produced by the latter are corpus (and context) dependent, this space adds general dimensions that are applicable to all corpora and context. For example, we successfully used the map to rank online collections of movie reviews and biomedical abstract based on the average measure of their words (the very primitive “bag of words” approach) along our dimensions. We indeed found that the first dimension (good-bad) was an excellent quantitative predictor not only of the movie critique score but also (together with the second dimension) of its genre (high valence and arousal: action, low valence and high arousal: drama, high valence and low arousal: romantic comedy, and low valence and arousal: documentaries).

It is also possible to extend the same approach to the has-a relationship (holonyms/meronyms), which is also explicitly included in WordNet. For example, in the sentence “a mouse has a whisker,” the mouse is holonym of whisker and whisker is meronym of mouse. This relationship is (like is-a) also antisymmetric, and we thus expect to be able to extract a “partonomy” scale much with the same approach described previously.

7. A Radical View of Reality, Information, Consciousness, and Remaining Challenges

Semantic mapping provides a possible approach to quantifying mental states that can be expressed declaratively. Mental states, or more practically words, sentences, paragraphs, and text in general are allocated in a multidimensional space such that their distances and relative positions reflect particular semantic relationships, including hypernymy/hyponymy (is-a), holonymy/meronymy (has-a), synonymy/antonymy (is similar/opposite to), and in principle many others, such as causation and cooccurrence [132]. In this framework, mental states and their relationships can be themselves represented as graphs of nodes and edges, respectively. If one believes that (at least some) mental states reflect properties of outer reality, it is possible to conceive reality itself as occurring in a giant graph in which any possible observable is a node, and edges correspond to probabilities that two observables would cooccur. We call this conceptual construct the Universal Reality Graph.

In this view, reality would unfold in time as a sequence of events constituting patterns of activation of subsets of nodes and all edges among them within the Universal Reality Graph. Any agent capable of observation will witness a subset of these activation patterns, that is, a sequence of partial events, each consisting of a collection of active nodes and edges. At each moment of experience, every agent would learn some (but not all) of those associations s/he witnesses. In this process, the agent would progressively form a mental graph representing part of his/her experienced history, which is itself part of the general occurrence of reality, sampled at each instant from the Universal Reality Graph.

While this notion of reality as a graph is purely speculative, even more extreme theories have proposed that physical reality is ultimately a product of information (“It from Bit”), rather than the other way around [133]. Most importantly, the possibility to conceive reality as a graph offers interesting vistas on the solution of the mind-brain problem. If agents form graph-like minds to represent (and therefore predict) their experience of graph-like reality, it stands to reason that the fittest physical substrates selected by evolution to encode these representations be themselves graph like, namely, brain networks. *The relationship between minds and brains could then be resolved as a mapping between their respective graphs and their embeddings.*

In this framework, the fundamental operation to grow a mind is pairwise association between observables [134], that is, establishment of edges between nodes in the mental graph based on corresponding experiences in the reality graph. An interesting aspect of mental representation is that we only learn a small fraction of associations from those observed in reality. In particular, our ability to learn is gated by previously acquired background information. We have recently proposed that this constraint may be a consequence of the spatial relationship among the tree-like shaped neuronal axons and dendrites that underlie brain connectivity [135, 136]. Specifically, in order for new synapses to be formed, the axon of the presynaptic neuron must be sufficiently close to a dendrite of the postsynaptic neuron, arguably because

of preexisting connectivity with other neurons encoding for related knowledge.

These ideas are also consonant with the Information Integration Theory (IIT) of consciousness [137], which is emerging as a leading candidate among the fundamental theories of mental content. The underlying assumption of IIT is that consciousness is fundamentally a property of information processing. Specifically, according to the IIT, when a brain (or in principle any other computing device) is in a particular state, its amount of consciousness, called Φ , depends not only on the actual content represented in that state but also on the absence of all contents represented in the states that are not being (but could be) instantiated. Thus silent neurons contribute to the conscious state as much as the active neurons, because consciousness depends as much on the content that *could* be represented by the network as on the content that is actually being represented. Therefore, consciousness is a product of the integrated activity in the network and is measured by information integration, a property that has been defined in graph structures [138].

While the IIT profoundly links consciousness to information [139], its cognitive underpinning is shared by other theories (e.g., [140]) and experimental approaches [141]. A crucial and unique outcome of IIT, however, is that the definition of integrated information enables a geometric characterization of mental states or qualia [142]. This can in principle provide a neurally based bottom-up correlate to the spaces that emerge from top-down semantic mapping of natural language. If the information processing product of neural network activity can be shown to correspond mathematically to a quantitative description of subjective mental content, the brain-mind problem would be effectively resolved.

Information is not only an essential element of consciousness, reality, and brain activity, but also of communication among conscious agents. Consider a dialogue between two individuals, in which one tells the other: "It's almost midnight, I'm tired. I had such a day." What does that *mean*? More precisely, what does it mean to the speaking individual, and what does it mean to the listener? Shannon's notion of information [143] captures the reduction of uncertainty of the listener's mind upon receiving the message. Assuming that the second individual had no idea of what time it was, whether the first person was rested or fatigued, and so forth, communication is indeed informative. Mapping brain and mental states, however, opens another perspective on the meaning of communication. Consider the brain/mind state X of the first individual (being tired, etc.), and suppose that the goal of the message was to instantiate X in the second individual. As a result of communication, however, the second individual's brain/mental state is Y rather than X . If Y equaled X , communication would be 100% perfect, but that is never the case.

The listener's understanding of "it's almost midnight" may be fairly close to the speaker's meaning of those words, but even in that case, "midnight" could be associated with different feelings and memories in the two individuals, and the term "almost" might be interpreted as 20 minutes by the first individual and as 2 minutes by the second. When

we analyze the second portion of the message, "I'm tired", we realize that the listener will think of his/her notion of tiredness, which is at best a coarse approximation of the actual state expressed by the speaker, for example, in terms of physical versus mental, chronic versus acute, concerned versus conversational, and so forth. The connotation of "I had such a day" is even more prone to subjective interpretation. The listener may think of his/her days s/he would describe as "such a day", but those days and associated sentiments are likely quite different from the events and related mental state the speaker was referring to.

This simple example can be generalized to all of human communications. How much of the intended meaning is effectively transmitted between communicating conscious individuals on average? Even in the most favorable cases of colleagues discussing their joint work or spouses talking about their family, we speculate that communication hovers between 50% and 85%. If this is the case, the mean human communication between any two individuals (including casual interactions) might be 10% or less. Even reminiscence and planning (retrieval of autobiographic and prospective memories, resp.), as well as other forms of internal dialogue, can be viewed as special cases of communication (with oneself). Such a type of communication between two instances of the same individual at different points in time can be expected to be much more effective than between different human beings, but even in those situations it will not be perfectly effective, as the mind is in constant flux. In all cases, mental state quantification by semantic mapping and its corresponding neural correlates in brain activity spaces could dramatically enhance communication effectiveness, deeply altering human relationship.

8. Concluding Remarks: A Vision Forward

The relationship between mind and matter has perhaps been, in one form or another, the most debated issue in the history of human thought, and it still constitutes, in the modern "mind-brain" incarnation, an open scientific and philosophical problem. Specifically, why do certain brain states "feel" like something, and why specific brain states feel the way they do? We have proposed that a satisfactory answer can ultimately come from mathematics, if the abstract spaces of brain activity and mental content can be quantitatively characterized and geometrically mapped onto each other. Such a "solution" will connect the conscious mind and the relevant aspect of brain states by demonstrating the equivalence of their properties, much like statistical mechanics and thermodynamics are nowadays accepted as one and the same phenomenon, even if they are practically treated as distinct for every day purposes.

We argued that semantic maps constitute a useful initial framework to establish a rigorous description of the mind and that network connectivity provides the most informative constraint on brain dynamics. However, defining the proper mathematical states to effectively bridge brain and mind still constitutes a formidable challenge. State-of-the-art semantic maps only scratch the surface of the necessary quantification

of the human mind. Next-generation voice recognition and optical character recognition software programs might soon enable real-time acquisition and analysis of the complete life-long natural language corpus experienced by an individual. Exhaustive compendia of semantic relationships could be extracted from such a resource, enabling the creation of a comprehensive semantic map for that individual. Such a resource could then be used to systematically report subjective mental states.

While the main obstacle in quantifying mental content appears to be the required paradigm shift towards a science of first-person perspective, neuroscience faces mostly a technological hurdle in creating brain-wide neuron-level maps of network connectivity and activity. Specifically, existing techniques can indeed map all of synaptic circuitries, but only in a very small volume (a fraction of cubic millimeter) of nervous tissue [144], only in animal models, and not *in vivo*. Other techniques to analyze neuronal anatomy, only hinting at the *potential* connectivity [75, 145], are possibly scalable to entire brains of live animals [146, 147], but again not human beings, let alone in normal behavioral conditions. The only noninvasive imaging techniques available to investigate the human brain (e.g., [148]) are by many orders of magnitude too coarse to probe the level of neurons and spike.

Several “futuristic” scenarios have been proposed to solve the technological gap between small-scale animal-model neuron-level analysis and a full activity map of all neurons and synapses of a sentient human brain [149], such as the eventual adoption of nanotechnology [150]. One present-day partial solution is to use molecular homology to identify existing correspondences between neuron types in rodents and humans by comprehensive genetic mapping [151] and single-neuron sequencing [152]. The subsequent extension of rodent brain connectivity to human cognitive architecture would only be tentative, requiring extensive computational testing and refinement by multiscale simulation [153]. An initial pilot project in this regard might tackle a suitable brain region (and related computational functions), such as the mammalian hippocampus [154].

Assuming that, at least in principle, technological advancements enabled accumulation of sufficient datasets to adequately map the neuronal activity of the human brain, such a feat would likely involve massive automation. High-throughput, machine-acquired, and large-scale data poses the outstanding matter of human interpretation [155–157]. This issue has recently promoted considerable growth in the field of neuroinformatics [158], that is, the establishment of an information framework for neuroscience (e.g., databases and other electronic resources), which is especially needed in computational neuroanatomy [159]. Recent initiatives have proposed a formalism to represent connectivity structure in neuronal network models [160] and seeded web-based multimodal connectivity databases [161]. A parallel informatics effort is required to enable storage, manipulation, and analysis of machine- and human-readable empirical data on cognitive functions, behaviors, and introspection [162, 163]. The neuroinformatics of language might provide a useful bridge between neural and cognitive frameworks [164].

The proposed vision offers a path towards a mathematical solution of the relationship between brain and mind that is consistent with contemporary philosophical positions [165]. Tremendous advancements in physics, chemistry, and biology provided an increasingly unified understanding of the material world. Bridging neuroscience with a quantitative description of inner subjective life may provide a fundamental closure to human scientific inquiry.

Acknowledgments

The author is grateful to Professor Aurelio Ascoli (the author's father) for his useful feedback on this paper, with particular attention to historical details. He also thanks Professor James L. Olds for his encouragement and suggestion to write this material for publication after its presentation at the 2012 Krasnow Institute retreat. The author's work is supported in part by Grants R01-NS39600 from the National Institutes of Health and MURI-N00014-10-1-0198 from the Office of Naval Research.

References

- [1] S. G. Weerasinghe, *The Sankhya Philosophy; A Critical Evaluation of Its Origins and Development*, South Asia Books, New Delhi, India, 1993.
- [2] B. Jowett, *The Dialogues of Plato Translated into English with Analyses and Introductions in Five Volumes*, Oxford University Press, 3rd edition, 1892, Chapter: PHAEDO, <http://oll.libertyfund.org/title/766/93700>.
- [3] W. Seager and S. Allen-Hermanson, “Panpsychism,” in *The Stanford Encyclopedia of Philosophy*, E. N. Zalta, Ed., 2012, <http://plato.stanford.edu/archives/spr2012/entries/panpsychism>.
- [4] M. Shuttleworth, “Aristotle's Psychology,” Experiment Resources, 2012, <http://www.experiment-resources.com/aristotles-psychology.html>.
- [5] G. J. C. Lokhorst and T. T. Kaitaro, “The originality of Descartes' theory about the pineal gland,” *Journal of the History of the Neurosciences*, vol. 10, no. 1, pp. 6–18, 2001.
- [6] F. Adams, “Translation of Hippocrates' ‘On the Sacred Disease’ (University of Adelaide Library, retrieved from Internet Classics Archive,” 2006, <http://web.archive.org/web/20070926213032/http://etext.library.adelaide.edu.au/mirror/classics.mit.edu/Hippocrates/sacred.html>.
- [7] V. Nutton, *Galen of Pergamum (Encyclopædia Britannica)*, 2012, <http://www.britannica.com/EBchecked/topic/223895/Galen-of-Pergamum>.
- [8] E. R. Kandel and L. R. Squire, “Neuroscience: breaking down scientific barriers to the study of brain and mind,” *Science*, vol. 290, no. 5494, pp. 1113–1120, 2000.
- [9] D. J. Chalmers and D. Bourget, “Consciousness and Neuroscience. In Online Papers on Consciousness,” 2009, <http://consc.net/online/8.1>.
- [10] A. Zeman, “Consciousness,” *Brain*, vol. 124, no. 7, pp. 1263–1289, 2001.
- [11] G. A. Ascoli and J. Grafman, Eds., *Consciousness, Mind and Brain*, Massom Publisher, Milan, Italy, 2005.
- [12] G. A. Ascoli, “Brain and mind at the crossroad of time,” *Cortex*, vol. 41, no. 5, pp. 619–620, 2005.

- [13] J. A. Hobson, "REM sleep and dreaming: towards a theory of protoconsciousness," *Nature Reviews Neuroscience*, vol. 10, no. 11, pp. 803–814, 2009.
- [14] D. Kahn and T. Gover, "Consciousness in dreams," *International Review of Neurobiology*, vol. 92, pp. 181–195, 2010.
- [15] R. E. Brown, R. Basheer, J. T. McKenna, R. E. Strecker, and R. W. McCarley, "Control of sleep and wakefulness," *Physiological Reviews*, vol. 92, no. 3, pp. 1087–1187, 2012.
- [16] A. V. Samsonovich and G. A. Ascoli, "The conscious self: ontology, epistemology and the mirror quest," *Cortex*, vol. 41, no. 5, pp. 621–636, 2005.
- [17] J. G. Taylor, "Mind and consciousness: towards a final answer?" *Physics of Life Reviews*, vol. 2, no. 1, pp. 1–45, 2005.
- [18] J. H. Reynolds and R. Desimone, "The role of neural mechanisms of attention in solving the binding problem," *Neuron*, vol. 24, no. 1, pp. 19–29, 1999.
- [19] J. J. Van Boxtel, N. Tsuchiya, and C. Koch, "Consciousness and attention: on sufficiency and necessity," *Frontiers in Psychology*, vol. 1, article 217, 2010.
- [20] T. Nagel, "What is it like to be a bat?" *The Philosophical Review*, vol. 83, no. 4, pp. 435–450, 1974.
- [21] A. Gierer, "Brain, mind and limitations of a scientific theory of human consciousness," *BioEssays*, vol. 30, no. 5, pp. 499–505, 2008.
- [22] N. Block, "Consciousness, accessibility, and the mesh between psychology and neuroscience," *Behavioral and Brain Sciences*, vol. 30, no. 5-6, pp. 481–548, 2007.
- [23] O. J. Hulme, K. F. Friston, and S. Zeki, "Neural correlates of stimulus reportability," *Journal of Cognitive Neuroscience*, vol. 21, no. 8, pp. 1602–1610, 2009.
- [24] B. Peeters, "Language and the mind. On concepts and values," *Pragmatics and Cognition*, vol. 4, pp. 139–152, 1996.
- [25] A. B. Butler, "Hallmarks of consciousness," *Advances in Experimental Medicine and Biology*, vol. 739, pp. 291–309, 2012.
- [26] R. P. Behrendt, "Contribution of hippocampal region CA3 to consciousness and schizophrenic hallucinations," *Neuroscience and Biobehavioral Reviews*, vol. 34, no. 8, pp. 1121–1136, 2010.
- [27] L. De Gennaro, C. Marzano, C. Cipolli, and M. Ferrara, "How we remember the stuff that dreams are made of: neurobiological approaches to the brain mechanisms of dream recall," *Behavioural Brain Research*, vol. 226, no. 2, pp. 592–596, 2012.
- [28] G. Tononi and G. M. Edelman, "Consciousness and complexity," *Science*, vol. 282, no. 5395, pp. 1846–1851, 1998.
- [29] A. Cleeremans, "Computational correlates of consciousness," *Progress in Brain Research*, vol. 150, pp. 81–98, 2005.
- [30] V. G. Hardcastle, "Consciousness and the neurobiology of perceptual binding," *Seminars in Neurology*, vol. 17, no. 2, pp. 163–170, 1997.
- [31] J. Aru, N. Axmacher, A. T. Do Lam et al., "Local category-specific gamma band responses in the visual cortex do not reflect conscious perception," *The Journal of Neuroscience*, vol. 32, pp. 14909–14914, 2012.
- [32] R. Q. Quiroga, "Concept cells: the building blocks of declarative memory functions," *Nature Reviews Neuroscience*, vol. 13, no. 8, pp. 587–597, 2012.
- [33] V. A. F. Lamme, "Towards a true neural stance on consciousness," *Trends in Cognitive Sciences*, vol. 10, no. 11, pp. 494–501, 2006.
- [34] B. Libet, "Reflections on the interaction of the mind and brain," *Progress in Neurobiology*, vol. 78, no. 3–5, pp. 322–326, 2006.
- [35] G. Hesslow, "Will neuroscience explain consciousness?" *Journal of Theoretical Biology*, vol. 171, no. 1, pp. 29–39, 1994.
- [36] C. McGinn, "Can we solve the mind-body problem?" *Mind*, vol. 98, no. 391, pp. 349–366, 1989.
- [37] J. R. Searle, "How to study consciousness scientifically," *Philosophical Transactions of the Royal Society B*, vol. 353, no. 1377, pp. 1935–1942, 1998.
- [38] G. A. Ascoli, "Is it already time to give up on a science of consciousness? A commentary on mysterianism," *Complexity*, vol. 5, no. 1, pp. 25–34, 1999.
- [39] I. Newton, "Philosophiæ Naturalis Principia Mathematica," 1687, English translation by A. Motte 1729, http://en.wikisource.org/wiki/The_Mathematical_Principles_of_Natural_Philosophy.
- [40] M. Clavelin, *The Natural Philosophy of Galileo*, MIT Press, Cambridge, Mass, USA, 1974.
- [41] S. Carnot, *Réflexions Sur la Puissance Motrice du Feu et Sur les Machines Propres À Développer cette Puissance*, Bachelier, Paris, France, 1824, English translation of R. Thurston, 1890, <http://books.google.com/books?id=tgdJAAAAIAAJ>.
- [42] D. Cardwell, *From Watt to Clausius: The Rise of Thermodynamics in the Early Industrial Age*, Cornell University Press, Ithaca, NY, USA, 1971.
- [43] R. Brown, "A brief account of microscopical observations made in the months of June, July and August, 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies," *Philosophical Magazine*, vol. 4, pp. 161–173, 1828.
- [44] D. Bernoulli, "Hydrodynamica, sive De viribus et motibus fluidorum commentarii," 1738, <http://echo.mpiwg-berlin.mpg.de/ECHODocuViewfull?mode=imagepath&url=/mpiwg/online/permanent/library/AZ870BWE/pageimg&view-Mode=images>.
- [45] A. Einstein, *Investigations on the Theory of Brownian Movement*, Dover, New York, NY, USA, 1956.
- [46] G. A. Ascoli, "From data to knowledge," *Neuroinformatics*, vol. 1, no. 2, pp. 145–147, 2003.
- [47] J. T. Vogelstein, R. J. Vogelstein, and C. E. Priebe, "Are mental properties supervenient on brain properties?" *Scientific Reports*, vol. 1, article 100, 2011.
- [48] M. Allen and G. Williams, "Consciousness, plasticity, and connectomics: the role of intersubjectivity in human cognition," *Frontiers in Psychology*, vol. 2, article 20, 2011.
- [49] G. A. Silva, "The need for the emergence of mathematical neuroscience: beyond computation and simulation," *Frontiers in Computational Neuroscience*, vol. 5, article 51, 2011.
- [50] W. Gerstner, H. Sprekeler, and G. Deco, "Theory and simulation in neuroscience," *Science*, vol. 338, no. 6103, pp. 60–65, 2012.
- [51] G. O'Brien and J. Opie, "A connectionist theory of phenomenal experience," *Behavioral and Brain Sciences*, vol. 22, no. 1, pp. 127–148, 1999.
- [52] G. Leisman and P. Koch, "Networks of conscious experience: computational neuroscience in understanding life, death, and consciousness," *Reviews in the Neurosciences*, vol. 20, no. 3-4, pp. 151–176, 2009.
- [53] G. Buzsáki, "The structure of consciousness," *Nature*, vol. 446, no. 7133, article 267, 2007.
- [54] O. Sporns, G. Tononi, and G. M. Edelman, "Connectivity and complexity: the relationship between neuroanatomy and brain dynamics," *Neural Networks*, vol. 13, no. 8-9, pp. 909–922, 2000.
- [55] G. A. Ascoli, "The complex link between neuroanatomy and consciousness," *Complexity*, vol. 6, pp. 20–26, 2000.

- [56] S. Herculano-Houzel, "The human brain in numbers: a linearly scaled-up primate brain," *Frontiers in Human Neuroscience*, vol. 3, article 31, 2009.
- [57] R. J. Douglas and K. A. Martin, "Mapping the matrix: the ways of neocortex," *Neuron*, vol. 56, no. 2, pp. 226–238, 2007.
- [58] C. C. King, "Fractal and chaotic dynamics in nervous systems," *Progress in Neurobiology*, vol. 36, no. 4, pp. 279–308, 1991.
- [59] P. So, J. T. Francis, T. I. Netoff, B. J. Gluckman, and S. J. Schiff, "Periodic orbits: a new language for neuronal dynamics," *Biophysical Journal*, vol. 74, no. 6, pp. 2776–2785, 1998.
- [60] A. A. Fingelkurts and A. A. Fingelkurts, "Making complexity simpler: multivariability and metastability in the brain," *International Journal of Neuroscience*, vol. 114, no. 7, pp. 843–862, 2004.
- [61] S. Roy and R. Llinás, "Dynamic geometry, brain function modeling, and consciousness," *Progress in Brain Research*, vol. 168, pp. 133–144, 2008.
- [62] G. Buzsáki, *Rhythms of the Brain*, Oxford University Press, Oxford, UK, 2006.
- [63] T. Klausberger and P. Somogyi, "Neuronal diversity and temporal dynamics: the unity of hippocampal circuit operations," *Science*, vol. 321, no. 5885, pp. 53–57, 2008.
- [64] O. Sporns, *Networks of the Brain*, The MIT Press, Cambridge, Mass, USA, 2011.
- [65] O. Sporns, G. Tononi, and R. Kötter, "The human connectome: a structural description of the human brain," *PLoS Computational Biology*, vol. 1, no. 4, article e42, 2005.
- [66] J. W. Lichtman, J. Livet, and J. R. Sanes, "A technicolour approach to the connectome," *Nature Reviews. Neuroscience*, vol. 9, no. 6, pp. 417–422, 2008.
- [67] J. Kim, T. Zhao, R. S. Petralia et al., "mGRASP enables mapping mammalian synaptic connectivity with light microscopy," *Nature Methods*, vol. 9, no. 1, pp. 96–102, 2011.
- [68] J. DeFelipe, "From the connectome to the synaptome: an epic love story," *Science*, vol. 330, no. 6008, pp. 1198–1201, 2010.
- [69] J. Lu, "Neuronal tracing for connectomic studies," *Neuroinformatics*, vol. 9, no. 2-3, pp. 159–166, 2011.
- [70] N. Kasthuri and J. W. Lichtman, "The rise of the 'projectome'," *Nature Methods*, vol. 4, no. 4, pp. 307–308, 2007.
- [71] G. A. Ascoli, "Potential connectomics complements the endeavour of 'no synapse left behind' in the cortex," *Journal of Physiology*, vol. 590, no. 4, pp. 651–652, 2012.
- [72] S. L. Hill, Y. Wang, I. Riachi, F. Schürmann, and H. Markram, "Statistical connectivity provides a sufficient foundation for specific functional connectivity in neocortical neural microcircuits," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 109, no. 42, pp. E2885–E2894, 2012.
- [73] A. M. Packer, D. J. McConnell, E. Fino, and R. Yuste, "Axo-dendritic overlap and laminar projection can explain interneuron connectivity to pyramidal cells," *Cereb Cortex*, 2012.
- [74] D. Kleinfeld, A. Bharioke, P. Blinder et al., "Large-scale automated histology in the pursuit of connectomes," *Journal of Neuroscience*, vol. 31, no. 45, pp. 16125–16138, 2011.
- [75] M. Helmstaedter and P. P. Mitra, "Computational methods and challenges for large-scale circuit mapping," *Current Opinion in Neurobiology*, vol. 22, no. 1, pp. 162–169, 2012.
- [76] O. Sporns, "The human connectome: a complex network," *Annals of the New York Academy of Sciences*, vol. 1224, pp. 109–125, 2011.
- [77] J. Sepulcre, M. R. Sabuncu, and K. A. Johnson, "Network assemblies in the functional brain," *Current Opinion in Neurology*, vol. 25, no. 4, pp. 384–391, 2012.
- [78] A. W. Toga, K. A. Clark, P. M. Thompson, D. W. Shattuck, and J. D. Van Horn, "Mapping the human connectome," *Neurosurgery*, vol. 71, no. 1, pp. 1–5, 2012.
- [79] D. C. Van Essen, K. Ugurbil, E. Auerbach et al., "The human connectome project: a data acquisition perspective," *Neuroimage*, vol. 62, no. 4, pp. 2222–2231, 2012.
- [80] B. B. Biswal, M. Mennes, X. N. Zuo et al., "Toward discovery science of human brain function," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 107, no. 10, pp. 4734–4739, 2010.
- [81] D. N. Kennedy, "Making connections in the connectome era," *Neuroinformatics*, vol. 8, no. 2, pp. 61–62, 2010.
- [82] J. W. Bohland, C. Wu, H. Barbas et al., "A proposal for a coordinated effort for the determination of brainwide neuroanatomical connectivity in model organisms at a mesoscopic scale," *PLoS Computational Biology*, vol. 5, no. 3, Article ID e1000334, 2009.
- [83] A. S. Chiang, C. Y. Lin, C. C. Chuang et al., "Three-dimensional reconstruction of brain-wide wiring networks in *Drosophila* at single-cell resolution," *Current Biology*, vol. 21, no. 1, pp. 1–11, 2011.
- [84] A. M. Zador, J. Dubnau, H. K. Oyibo, H. Zhan, G. Cao, and I. D. Peikon, "Sequencing the connectome," *PLoS Biology*, vol. 10, no. 10, Article ID e1001411, 2012.
- [85] E. T. Bullmore and D. S. Bassett, "Brain graphs: graphical models of the human brain connectome," *Annual Review of Clinical Psychology*, vol. 7, pp. 113–140, 2011.
- [86] M. Kaiser, "A tutorial in connectome analysis: topological and spatial features of brain networks," *NeuroImage*, vol. 57, no. 3, pp. 892–907, 2011.
- [87] P. Erdős and A. Rényi, "On random graphs," *Publicationes Mathematica*, vol. 6, pp. 290–297, 1959.
- [88] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, pp. 440–442, 1998.
- [89] A. L. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, pp. 509–512, 1999.
- [90] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon, "Network motifs: simple building blocks of complex networks," *Science*, vol. 298, no. 5594, pp. 824–827, 2002.
- [91] O. Sporns and J. D. Zwi, "The small world of the cerebral cortex," *Neuroinformatics*, vol. 2, no. 2, pp. 145–162, 2004.
- [92] S. Song, P. J. Sjöström, M. Reigl, S. Nelson, and D. B. Chklovskii, "Highly nonrandom features of synaptic connectivity in local cortical circuits," *PLoS Biology*, vol. 3, no. 10, article e350, 2005.
- [93] W. J. Freeman and M. Breakspear, "Scale-free neocortical dynamics," *Scholarpedia*, vol. 2, article 1357, 2007.
- [94] B. J. Pettejohn, M. J. Berryman, and M. D. McDonnell, "Methods for generating complex networks with selected structural properties for simulations: a review and tutorial for neuroscientists," *Frontiers in Computational Neuroscience*, vol. 5, article 11, 2011.
- [95] M. A. Porter, J. P. Onnela, and P. J. Mucha, "Communities in networks," *Notices of the American Mathematical Society*, vol. 56, no. 9, pp. 1082–1097, 2009.
- [96] M. P. van den Heuvel and O. Sporns, "Rich-club organization of the human connectome," *Journal of Neuroscience*, vol. 31, no. 44, pp. 15775–15786, 2011.

- [97] M. Rivera-Alba, S. N. Vitaladevuni, Y. Mishchenko et al., "Wiring economy and volume exclusion determine neuronal placement in the *Drosophila* brain," *Current Biology*, vol. 21, no. 23, pp. 2000–2005, 2011, Erratum in: *Current Biology*, vol. 22, no. 2, article 172, 2012.
- [98] E. Bullmore and O. Sporns, "The economy of brain network organization," *Nature Reviews Neuroscience*, vol. 13, no. 5, pp. 336–349, 2012.
- [99] C. Assisi, M. Stopfer, and M. Bazhenov, "Using the structure of inhibitory networks to unravel mechanisms of spatiotemporal patterning," *Neuron*, vol. 69, no. 2, pp. 373–386, 2011.
- [100] L. da Fontoura Costa, J. L. Batista, and G. A. Ascoli, "Communication structure of cortical networks," *Frontiers in Computational Neuroscience*, vol. 5, article 6, 2011.
- [101] G. A. Ascoli, L. Alonso-Nanclares, S. A. Anderson et al., "Petilla terminology: nomenclature of features of GABAergic interneurons of the cerebral cortex," *Nature Reviews Neuroscience*, vol. 9, no. 7, pp. 557–568, 2008.
- [102] D. E. Fishkind, D. L. Sussman, M. Tang, J. T. Vogelstein, and C. E. Priebe, "Consistent adjacency-spectral partitioning for the stochastic block model when the model parameters are unknown," *SIAM Journal on Matrix Analysis and Applications*, 2012.
- [103] C. E. Priebe, J. T. Vogelstein, and D. Bock, "Optimizing the quantity/quality trade-off in connectome inference," *Communications in Statistics*, 2011.
- [104] J. Lu, J. C. Tapia, O. L. White, and J. W. Lichtman, "The interscutularis muscle connectome," *PLoS Biology*, vol. 7, no. 2, article e32, 2009.
- [105] C. Kelly, B. B. Biswal, R. C. Craddock, F. X. Castellanos, and M. P. Milham, "Characterizing variation in the functional connectome: promise and pitfalls," *Trends in Cognitive Sciences*, vol. 16, no. 3, pp. 181–188, 2012.
- [106] V. De Paola, A. Holtmaat, G. Knott et al., "Cell type-specific structural plasticity of axonal branches and boutons in the adult neocortex," *Neuron*, vol. 49, no. 6, pp. 861–875, 2006.
- [107] M. Park, J. M. Salgado, L. Ostroff et al., "Plasticity-induced growth of dendritic spines by exocytic trafficking from recycling endosomes," *Neuron*, vol. 52, no. 5, pp. 817–830, 2006.
- [108] K. M. Harris, "Structure, development, and plasticity of dendritic spines," *Current Opinion in Neurobiology*, vol. 9, no. 3, pp. 343–348, 1999.
- [109] N. Gogolla, I. Galimberti, and P. Caroni, "Structural plasticity of axon terminals in the adult," *Current Opinion in Neurobiology*, vol. 17, no. 5, pp. 516–524, 2007.
- [110] A. Holtmaat and K. Svoboda, "Experience-dependent structural synaptic plasticity in the mammalian brain," *Nature Reviews Neuroscience*, vol. 10, no. 9, pp. 647–658, 2009, Erratum in: *Nature Reviews Neuroscience*, vol. 10, no. 10, article 759, 2009.
- [111] M. Fu and Y. Zuo, "Experience-dependent structural plasticity in the cortex," *Trends in Neurosciences*, vol. 34, no. 4, pp. 177–187, 2011.
- [112] E. Bruel-Jungerman, S. Davis, and S. Laroche, "Brain plasticity mechanisms and memory: a party of four," *Neuroscientist*, vol. 13, no. 5, pp. 492–505, 2007.
- [113] D. B. Chklovskii, B. W. Mel, and K. Svoboda, "Cortical rewiring and information storage," *Nature*, vol. 431, no. 7010, pp. 782–788, 2004.
- [114] S. Fusi, P. J. Drew, and L. F. Abbott, "Cascade models of synaptically stored memories," *Neuron*, vol. 45, no. 4, pp. 599–611, 2005.
- [115] C. Kahn, *The Art and Thought of Heraclitus: Fragments with Translation and Commentary*, Cambridge University Press, London, UK, 1979.
- [116] G. A. Ascoli and A. V. Samsonovich, "Science of the conscious mind," *Biological Bulletin*, vol. 215, no. 3, pp. 204–215, 2008.
- [117] T. K. Landauer and S. T. Dumais, "A solution to Plato's problem: the latent semantic analysis theory of acquisition, induction, and representation of knowledge," *Psychological Review*, vol. 104, no. 2, pp. 211–240, 1997.
- [118] C. H. Papadimitriou, P. Raghavan, H. Tamaki, and S. Vempala, "Latent semantic indexing: a probabilistic analysis," *Journal of Computer and System Sciences*, vol. 61, no. 2, pp. 217–235, 2000.
- [119] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent Dirichlet allocation," *Journal of Machine Learning Research*, vol. 3, no. 4–5, pp. 993–1022, 2003.
- [120] C. Burgess and K. Lund, "Modelling parsing constraints with high-dimensional context space," *Language and Cognitive Processes*, vol. 12, no. 2–3, pp. 177–210, 1997.
- [121] T. K. Landauer, D. S. McNamara, S. Dennis, and W. Kintsch, Eds., *Handbook of Latent Semantic Analysis*, Lawrence Erlbaum Associates, Mahwah, NJ, USA, 2007.
- [122] C. E. Osgood, G. Suci, and P. Tannenbaum, *The Measurement of Meaning*, University of Illinois Press, Urbana, Ill, USA, 1957.
- [123] D. C. Rubin, "51 properties of 125 words: a unit analysis of verbal behavior," *Journal of Verbal Learning and Verbal Behavior*, vol. 19, no. 6, pp. 736–755, 1980.
- [124] P. Roget, "Roget's Thesaurus of English words and phrases," 1852, <http://www.roget.org>.
- [125] C. Fellbaum, *WordNet: An Electronic Lexical Database*, MIT Press, Cambridge, Mass, USA, 1998.
- [126] T. Gruber, "Ontology," in *Encyclopedia of Database Systems*, L. Liu and M. Tamer Özsu, Eds., Springer, 2009, <http://tomgruber.org/writing/ontology-definition-2007.htm>.
- [127] A. V. Samsonovich and G. A. Ascoli, "Principal semantic components of language and the measurement of meaning," *PLoS ONE*, vol. 5, no. 6, Article ID e10921, 2010.
- [128] A. V. Samsonovich, R. F. Goldin, and G. A. Ascoli, "Toward a semantic general theory of everything," *Complexity*, vol. 15, no. 4, pp. 12–18, 2010.
- [129] H. Lövhheim, "A new three-dimensional model for emotions and monoamine neurotransmitters," *Medical Hypotheses*, vol. 78, no. 2, pp. 341–348, 2012.
- [130] G. Kemerling, "Aristotle: logical methods," in *The Philosophy Pages*, 2011, <http://www.philosophypages.com/hy/2n.htm>.
- [131] A. V. Samsonovich and G. A. Ascoli, "Augmenting weak semantic cognitive maps with an "abstractness" dimension," *Computational Intelligence and Neuroscience*, 2013.
- [132] G. Recchia and M. N. Jones, "The semantic richness of abstract concepts," *Frontiers in Human Neuroscience*, vol. 6, article 315, 2012.
- [133] J. A. Wheeler, "Information, physics, quantum: the search for links. Ch. 19," in *Complexity, Entropy, and the Physics of Information*, W. Zurek, Ed., Addison-Wesley, Redwood City, Calif, USA, 1990.
- [134] G. A. Ascoli, "Association, abstraction, and the emergence of the self," *Noetic Journal*, vol. 2, pp. 9–20, 1999.
- [135] G. A. Ascoli and M. Mainetti, *Gated Learning: Much ADO About Background Information. Towards A Science of Consciousness*, Stockholm, Sweden, 2011.
- [136] M. Mainetti and G. A. Ascoli, *Much ADO about BIG Learning*, APS Mtg, Washington, DC, USA, 2011.

- [137] G. Tononi, "Consciousness as integrated information: a provisional manifesto," *Biological Bulletin*, vol. 215, no. 3, pp. 216–242, 2008.
- [138] G. Tononi and O. Sporns, "Measuring information integration," *BMC Neuroscience*, vol. 4, article 31, 2003.
- [139] D. Balduzzi and G. Tononi, "Integrated information in discrete dynamical systems: motivation and theoretical framework," *PLoS Computational Biology*, vol. 4, no. 6, Article ID e1000091, 2008.
- [140] B. J. Baars, "Global workspace theory of consciousness: toward a cognitive neuroscience of human experience," *Progress in Brain Research*, vol. 150, pp. 45–53, 2005.
- [141] S. Dehaene and J. P. Changeux, "Experimental and theoretical approaches to conscious processing," *Neuron*, vol. 70, no. 2, pp. 200–227, 2011.
- [142] D. Balduzzi and G. Tononi, "Qualia: the geometry of integrated information," *PLoS Computational Biology*, vol. 5, no. 8, Article ID e1000462, 2009.
- [143] C. E. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, pp. 623–656, 1948.
- [144] D. B. Chklovskii, S. Vitaladevuni, and L. K. Scheffer, "Semi-automated reconstruction of neural circuits using electron microscopy," *Current Opinion in Neurobiology*, vol. 20, no. 5, pp. 667–675, 2010.
- [145] E. G. Jones, J. M. Stone, and H. J. Karten, "High-resolution digital brain atlases: a Hubble telescope for the brain," *Annals of the New York Academy of Sciences*, vol. 1225, no. 1, pp. E147–E159, 2011.
- [146] H. Hama, H. Kurokawa, H. Kawano et al., "Scale: a chemical approach for fluorescence imaging and reconstruction of transparent mouse brain," *Nature Neuroscience*, vol. 14, no. 11, pp. 1481–1488, 2011.
- [147] Y. Mishchenko and L. Paninski, "A Bayesian compressed-sensing approach for reconstructing neural connectivity from subsampled anatomical data," *Journal of Computational Neuroscience*, vol. 33, no. 2, pp. 371–388, 2012.
- [148] A. Irimia, M. C. Chambers, C. M. Torgerson, and J. D. Horn, "Circular representation of human cortical networks for subject and population-level connectomic visualization," *Neuroimage*, vol. 60, no. 2, pp. 1340–1351, 2012.
- [149] R. A. Koene, "Experimental research in whole brain emulation: the need for innovative in vivo measurement techniques," *International Journal of Machine Consciousness*, vol. 4, no. 1, pp. 35–65, 2012.
- [150] R. Kurzweil and T. Grossman, "Fantastic voyage: live long enough to live forever. The science behind radical life extension questions and answers," *Studies in Health Technology and Informatics*, vol. 149, pp. 187–194, 2009.
- [151] M. J. Hawrylycz, E. S. Lein, A. L. Guillozet-Bongaarts et al., "An anatomically comprehensive atlas of the adult human brain transcriptome," *Nature*, vol. 489, no. 7416, pp. 391–399, 2012.
- [152] G. D. Evrony, X. Cai, E. Lee et al., "Single-neuron sequencing analysis of l1 retrotransposition and somatic mutation in the human brain," *Cell*, vol. 151, no. 3, pp. 483–496, 2012.
- [153] O. Schmitt and P. Eipert, "neuroVIISAS: approaching multiscale simulation of the rat connectome," *Neuroinformatics*, vol. 10, no. 3, pp. 243–267, 2012.
- [154] G. A. Ascoli, "The coming of age of the hippocampome," *Neuroinformatics*, vol. 8, no. 1, pp. 1–3, 2010.
- [155] K. H. Ambert and A. M. Cohen, "Text-mining and neuroscience," *International Review of Neurobiology*, vol. 103, pp. 109–132, 2012.
- [156] L. French, S. Lane, L. Xu et al., "Application and evaluation of automated methods to extract neuroanatomical connectivity statements from free text," *Bioinformatics*, vol. 28, no. 22, pp. 2963–2970, 2012.
- [157] G. A. Ascoli, "Global neuroscience: distributing the management of brain knowledge worldwide," *Neuroinformatics*, vol. 11, no. 1, pp. 1–3, 2013.
- [158] H. Akil, M. E. Martone, and D. C. Van Essen, "Challenges and opportunities in mining neuroscience data," *Science*, vol. 331, no. 6018, pp. 708–712, 2011.
- [159] G. A. Ascoli, Ed., *Computational Neuroanatomy: Principles and Methods*, Springer, Totowa, NJ, USA, 2002.
- [160] M. Djurfeldt, "The connection-set algebra—a novel formalism for the representation of connectivity structure in neuronal network models," *Neuroinformatics*, vol. 10, no. 3, pp. 287–304, 2012.
- [161] J. A. Brown, J. D. Rudie, A. Bandrowski, J. D. Van Horn, and S. Y. Bookheimer, "The UCLA multimodal connectivity database: a web-based platform for brain connectivity matrix sharing and analysis," *Frontiers in Neuroinformatics*, vol. 6, article 28, 2012.
- [162] G. V. Gkoutos, P. N. Schofield, and R. Hoehndorf, "The neurobehavior ontology: an ontology for annotation and integration of behavior and behavioral phenotypes," *International Review of Neurobiology*, vol. 3, pp. 69–87, 2012.
- [163] J. A. Turner and A. R. Laird, "The cognitive paradigm ontology: design and application," *Neuroinformatics*, vol. 10, no. 1, pp. 57–66, 2012.
- [164] M. A. Arbib, J. J. Bonaiuto, I. Bornkessel-Schlesewsky et al., "Action and language mechanisms in the brain: data, models and neuroinformatics," *Neuroinformatics*, 2013.
- [165] T. E. Feinberg, "Neuroontology, neurobiological naturalism, and consciousness: a challenge to scientific reduction and a solution," *Physics of Life Reviews*, vol. 9, no. 1, pp. 13–34, 2012.