

Vagueness and Mechanistic Explanation in Neuroscience

PHILIPP HAUEIS
*Institute of Philosophy
Humboldt University of Berlin*

The problem of fuzzy boundaries when delineating cortical areas is widely known in human brain mapping and its adjacent subdisciplines (anatomy, physiology and functional neuroimaging). Yet, a conceptual framework for understanding indeterminacy in neuroscience is missing, and there has been no discussion in the philosophy of neuroscience whether indeterminacy poses an issue for good neuroscientific explanations. My paper addresses both these issues by applying philosophical theories of vagueness to three levels of neuroscientific research, namely to (i) cytoarchitectonic studies at the neuron level (ii) intra-areal neuronal interaction measured by the BOLD-signal of functional magnetic resonance imaging (fMRI) and (iii) inter-areal connectivity between different cortical areas. The rest of the paper explores how this framework can be extended to mechanistic explanations in neuroscience. I discuss a semantic and an ontic interpretation of vagueness in mechanistic explanations and argue how both become scientifically interesting from the perspective of a philosophy of scientific practice.

Key words: Vagueness, mechanisms, neuroscience, explanation, fuzzy boundaries.

Introduction

In a variety of studies concerning the anatomy of the cerebral cortex, neuroscientific researchers report different forms of indeterminacy when describing the boundaries between different cortical areas. In a study that investigated how the folding of the cortical surface predicts structural features of different brain areas, Fischl et. al (2007) notice that the “changes that define the borders between adjacent association cortices (such as 44/45) are considerably more subtle than in primary areas, [...], making the precise and repeatable localization of higher

areas considerably more difficult” (ibid., 1978). In a similar vein, but concerned with the “activation maps” of functional magnetic resonance imaging (fMRI) studies, Saxe, Brett and Kanwisher (2010) write that these maps are not “unambiguous and sharp [...]. Instead, as most researchers are aware, the apparent sharp boundaries are subject to the choice of threshold applied to the statistical tests that generate the image” (ibid., 39) The same phenomenon is described by Passingham, Stephan and Kötter (2002) in a different research area, where neuro-anatomical boundaries are detected by classifying the structural and functional connections between different spatial locations in the cortex. Although “it is possible to detect clusters of areas with a similar, although not identical pattern of connections”, these authors also stress that “there is no objective criterion for defining the size of a family [...]. The threshold for defining ‘families’ of areas is arbitrary” (ibid., 609).

These quotations show that neuroscientists acknowledge that studying brain anatomy involves dealing with *fuzzy boundaries* between different cortical areas. What is missing, however, is a conceptual framework to understand and distinguish this issue from other methodological and technical problems in neuroscientific research. Saxe and colleagues for instance believe that the problem of fuzzy boundaries is merely an empirical issue which requires a consensual solution, and they do not distinguish it from ambiguity or uncertainty due to technological shortcomings. But disentangling these issues is crucial to decide which mode of analysis is the appropriate problem solution (e.g. empirical, conceptual or computational analysis), which in turn has fundamental implications for practical applications of neuro-anatomical knowledge (e.g. localization studies, invasive experiments or surgery). Therefore, the first part of this paper applies philosophical theories of vagueness to three levels of neuroscientific research, namely to cytoarchitectonics, where the identity conditions for brain areas are identified the neuron level (section 1.1), voxel activation at the level of intra-region neuronal interaction measured by the BOLD-signal of functional magnetic resonance imaging (fMRI, section 1.2) and community membership of nodes at the level of functional connectivity between brain regions (section 1.3). My analysis of vagueness in neuroanatomy serves a twofold purpose: first, to develop an account of vagueness that adequately captures indeterminacy in neuroscience in general (section 1.4), and second to explore what implications indeterminacy has for neuroscientific explanations in particular. Thus, the second part of the paper investigates how the conceptual framework developed above applies to Craver’s (2007) model of mechanistic explanation. To date, the mechanistic account is the most prominent explanatory model in philosophy of neuroscience and its multi-level nature suggests that vagueness at one level has implications the other ones as well. But since mechanisms are held to be individuated by function and not spatiotemporal location, I will first discuss an objection to the thesis that mechanisms or their explanations are vague at all (section

2.1). My rebuttal attempts to show that the objection is inconsistent with Craver's account of mechanistic explanation. Then, the vagueness of mechanistic explanation will be explicated by the use of "fuzziness" in the recent literature on philosophy of psychiatry (section 2.2). Based on this explication, a semantic and an ontic version of the vagueness of mechanisms can be distinguished (section 2.3). Since I claim that both versions are unsatisfying for the practicing neuroscientist, my discussion will end by arguing how vagueness becomes a scientifically interesting problem from the perspective of philosophy of scientific practice (section 2.4).

1. *Vagueness in Neuroanatomy*

Before describing how and at which levels of research vagueness is to be found in neuroanatomy, some general remarks about the philosophical concept of vagueness are necessary. The paradigmatic form of vagueness leads to the so-called *Sorites*-paradox (from *soros*, Greek for 'heap'), where the scope of a predicate is gradually increased until absurdity. The paradox arises when an uncontroversially true statement like "10,000 grains of sand make up a heap" is placed in an inductively generated series of statements that stand in a transitive relationship to one another, at the end of which stands a statement that is obviously false, such as "three grains of sand make up a heap". In between there are statements referring to borderline cases of heaps, i.e. cases for which it is indeterminate whether the statement "X grains of sand make up a heap" is true or false. Vagueness, therefore, poses a challenge to classical logic, as the borderline case sentences seem to express propositions that are neither true nor false (they have truth-value gaps). Distinct from this form of vagueness concerning the extension of predicates is the "problem of the many" (Unger 1980), concerning the individuation of objects. Take for example a single cloud in the sky, for which there are different aggregates of water droplets in the sky that are all candidates of being the cloud. Hence, for the water droplets that are contained in some, but not all cloud candidates, it is indeterminate whether they are part of the cloud or not. These types of vagueness share certain features, namely that these predicates/objects a) allow for borderline cases (of application/membership), b) have fuzzy boundaries (extensional/spatiotemporal) and c) are tolerant to small changes. Both forms of vagueness can be interpreted as either semantic or ontic vagueness: in the first reading, vagueness is a feature of language and concerns the indeterminacy of applying a concept with a vague rather than a precise extension (Russell 1923). In the second reading the world itself is such that there is no matter of fact whether a certain object is a part of the vague object (Tye 1990, 535f.), or whether an object has a certain property. Ontic vagueness can be based on both categorical and gradable properties (e.g. number of water droplets or shades of color). Furthermore, vagueness is to be distinguished from

epistemic uncertainty, where indeterminacy arises through incomplete knowledge. These general distinctions provide a heuristic tool to describe the different forms of indeterminacy neuroanatomists face in their research.

1.1. *Cytoarchitectonics*. On the level of cytoarchitectonics, researchers study the cortex at the microscopic level of histological sections of post mortem brains. The structural properties of neuron cells (e.g. shape, size, density) can be used to parcellate the cortex into various cortical areas. Although in some cases, the distribution and structure of neurons changes abruptly (e.g. in the primary visual cortex), in other cases the transition between two areas changes gradually. Consider the following example: A characteristic feature of Brodmann area 4 (BA 4) is the presence of giant pyramidal cells (Betz cells) in layer V, but their size varies considerably between different individuals, both in height (60–120 μm) and width (30–60 μm) (Amunts, Schleicher, and Zilles, 2002). Furthermore in a single individual, Betz cells can be found outside of area 4, and the distance between them increases towards BA 6 (Economo and von Koskinas, 1925; Zilles et al., 1995). Thus, there are multiple instances of vagueness: the variability of height and width introduces imprecision into the definition of what a Betz cell is, although for every individual, these cells can be identified. But there is also vagueness in every single individual, because gradually increasing distance gives rise to vagueness of individuation between BA 4 and 6: At the periphery there are parts $X_{1, \dots, n}$ such that it is indeterminate whether the proposition expressed by the sentence “ $X_{1, \dots, n}$ is part of BA 4” is true or false.¹ Notice that contrary to the observation of Fischl and colleagues quoted in the introduction, the vagueness of individuating brain areas is not identical to the problem of inter-individual variability. Even if in the case of higher cortical areas it is more difficult to *repeat* localization, it is already impossible to localize an area precisely in the singular case, insofar as there is a gradual transition of a cytoarchitectonic property between two areas. Albeit being illustrative example, however, the vagueness of BA 4 does not quite reflect neuroanatomical practice. The reason is that in the given example, the individuation of brain areas is made on the basis of just one feature, while usually cortical areas are delineated by using many dimensions, such as myelination (the insulation of axonal connections between neurons), or chemical receptor binding sites. But even in the case of multi-dimensional characterizations of brain anatomy, vagueness appears in the form of *transitional zones*, that share some, but not all cytoarchitectonic properties of its adjacent neighbors. Area 9/46 in the human prefrontal cortex represents such a

¹ The same applies for sentences using “is part of BA 4 or BA 6” where “or” is an exclusive disjunction. Notice that this case can be also construed as a Sorites-series: By constructing a series of statements of the form “This part belongs to BA 4” one moves from uncontroversially true statements to obviously false ones, leaving borderline cases in between.

case (Rajkowska and Goldman-Rakic, 1995a, b). It shares with area 9 a pale sublayer Vb, and with area 46 a distinct layer IV and uniformly sized cells in layers III and V. Additionally, 9/46 possesses combined features at the periphery towards other areas (classified as 9–8, 9–45, 46–10, and 46–45) and cannot be separated distinctly into areas 9 or 46 based on myeloarchitecture. The case of 9/46 most likely presents a case of absolute vagueness, where all known anatomical properties of the brain fail to clearly demarcate a spatiotemporal patch of the cortex into distinct areas.² It would therefore meet the philosophical working definition of vagueness as presenting speakers with an uncertainty that is not due to the “ignorance of the facts” (Grice 1989, 177).³ Notice that the introduction of a “transitional zone” does not represent more than the collection of borderline cases of the neighboring regions. Thus, it creates an artificially sharp line between all the borderline cases and the clear cases of membership to cortical regions where there is none (this phenomenon is called higher-order vagueness, see section 1.3 for further discussion).

1.2 Task-based fMRI and statistical thresholding. Since neuroscientists hold that cortical anatomy is structured by the functions that different subparts of the brain fulfill, the cytoarchitectonic study of structural boundaries is accompanied by the investigation of functional boundaries through cognitive neuroscience. In fMRI studies, functional anatomy is examined by inferring neural activity from the contrast between oxygenated and deoxygenated blood (blood oxygen level dependent, or BOLD). In experiments where subjects have to execute a cognitive task (e.g. rhythmic finger tapping), the contrast is generated by comparing the BOLD-response during the task to a baseline condition. The contrast is only regarded as an activity of the respective brain area when it is statistically significant, i.e. when it passes a statistical null-hypothesis (H_0) test. In functional neuroimaging, these tests generate a value for each voxel that indicates whether the BOLD-data would have occurred if H_0 had been true (p -value). Since H_0 holds that the data would have occurred even if the subject did not execute any task, neuroscientists infer that the neurons in a voxel were causally involved (i.e. active) during the experiment in case the p -value is below a pre-determined significance level α . But setting α at any level is arbitrary

² This is partly an empirical conjecture, because there has been no chemoarchitectonic study of this area I know of. There are good empirical reasons, however, that such a study will not result in a precise demarcation of 9/46 either: not all receptors show areal boundaries equally (Zilles et al., 1995), and some are also heterogeneously distributed within an area (e.g. GABA_a receptors within area V1, Zilles and Schleicher, 1993).

³ Such a demand seems problematic insofar the phenomenon of vagueness gives rise to empirical issues in the sciences. Although I cannot explore here further in what way empirical evidence *could* alter how researchers deal with vagueness (see Weiner 2007), I will discuss the issue with regard to the semantic and ontic vagueness of mechanisms in section 2.3.

since the supposition that data above such a level occurred by mere chance is unwarranted. The reason is that the brain is a *causally dense system*, where every neuron is (at least weakly) causally connected to every other neuron (cf. Savoy, 2001, 30). Since for such systems, H_0 is always strictly speaking false, there is no rationally compelling reason to distinguish between activity and non-activity despite the current technological ability to detect BOLD-responses at a certain level (Klein 2010). To put it in different words: the distinction between “activity” and “non-activity” is vague, because the setting of a significance level is susceptible to Sorites-reasoning: If we regard any data at $p < 0.05$ as statistically significant, why should we not regard data at $p < 0.051$ as statistically significant too? And why not for $p < 0.052$? If not established by convention, the choice of a statistical threshold would present the cognitive neuroscientist with the uncertainty typical for the phenomenon of vagueness.⁴

In empirical research, the vagueness of “neuronal activity” is overlapping with several issues that can be analytically distinguished from one another. In practice, fMRI researchers face various measurement uncertainties such as downstream effects, where increased oxygenation comes from other areas (Tehovnik et al., 2006) or the blurring of the signal across anatomical landmarks (gyri) due to the magnetization of brain tissue (Ojemann et al., 1997). Indeterminacy that is based on these uncertainties is merely epistemic and can probably be resolved by technological innovation. A question that cannot be answered by improved technology however, is how the fMRI signal is to be interpreted physiologically. Logothetis (2008) described the interpretation of the BOLD-contrast as ambiguous, since different levels of oxygenation do not allow for a distinction between cortical networks that are largely inhibitory, excitatory or both. What is meant by “ambiguous” here is not the overdetermination of conventionally established semantic meanings by one term, but the underspecification of the causal source through the displaying sign. Just like smoke can indicate fire, the BOLD contrast can but does not decisively indicate one of the three above possibilities. The causal ambiguity is to be distinguished from vagueness of “neural activity”, since the former concerns the question what kind of neural activity is measured by fMRI studies and not how, based on that research, functional boundaries of cortical areas are to be distinguished. That these questions may be ultimately related is no reason to take them to be identical.

1.3 *Inter-Areal Connectivity and Networks.* The neuroscientifically informed skeptic could object at this point that vagueness in cytoarchitec-

⁴ Notice that the conventional setting of a threshold is also dependent on the context of research: while neuroimagers regard false positives as an issue for good research (Simmons et. al, 2011), neurosurgeons include them to prevent the injury of important parts around the lesion that has to be operated (Gorgolewski et. al, 2011).

tonics and task-based fMRI do not pose a real problem to brain mapping, since both make unwarranted assumptions about cortical organization that are not upheld anymore in current research practice. Neuroscientists nowadays recognize that the brain consists of large-scale networks including many cortical areas, and that the majority of physiological activity is not explained by task-based activation. Therefore, anatomical studies that search for the connectional patterns based on fMRI measurements in the absence of a cognitive task (so-called resting-state connectivity studies) are immune to the cases of vagueness described above. But such an argument would misconstrue the relationship between connectivity studies and the other two levels of research. Firstly, the aim of connectivity studies is to confirm existing cytoarchitectonic delineations of cortical areas and eventually to discover hitherto unknown ones. Therefore it makes sense to ask how precisely or vaguely the brain can be parcellated based on this research method. Secondly, as indicated by the quote of Passingham and colleagues, connectivity studies employ statistical thresholds to fMRI data to decide how similar connectional patterns within an area have to be to count as one “family”. Considering these interdependencies with the previous levels, it has to be explored how vagueness does get imported into connectivity studies. Two examples will serve as an illustration.

The first example is a study by Wu et al. (2010) that investigated the overlapping community structure of a structural brain network. Describing cortical connections as networks is based on the assumption that communication is the most efficient when the connection between any two nodes of the network is the shortest possible (high integration), while the whole network is clustered in maximally distinct communities (high segregation). Wu and colleagues constructed such a “small-world” network by correlating the volumetric changes of grey matter (indicating past anatomical connection between areas cf. Mechelli et al., 2005) to construct an inter-areal correlation matrix between all subjects. As the network analysis based on that matrix showed the highest mutual information at a subgraph size where 15 nodes belonged to more than one community, the researchers concluded that their analysis revealed “fuzzy boundaries” between communities in the structural brain network (a result that was robust across various thresholds). What remains unanswered by the scientists is how their finding can be understood as vagueness of cortical networks.⁵ First of all, the vagueness described here is intensional, i.e. only showing possible borderline cases (Fine, 1975), since “grey matter volume” is not a sortal concept that allows scientists to count cortical areas, but only an indirect measure of cortical connectivity. Second—and because of their intensionality—these borderline cases can either represent gradual transitions of cytoarchitectonic properties or the fact that it is

⁵ I argue in more detail against understanding the overlapping community structure as a family resemblance relation in Haueis (*forthcoming*).

indeterminate to which community a node belongs to is independent of cytoarchitectonics because the overlapping community structure is an organizational principle of the brain. In case the first of these two alternatives is correct, the fuzzy community boundaries would be analogous to higher-order vagueness, where the vagueness of the object language gets imported into the meta-language (Varzi, 2001).

The second example concerns the study of functional connectivity by Cohen et al. (2008), who sought to delineate functional cortical boundaries by applying an edge detection algorithm to their collected resting-state fMRI data. These data were allocated to seed regions in a grid system that was based on a structural MRI scan of the subject. To each of these seed regions, the researchers assigned a coefficient η^2 which expresses how similar the connectivity profile is to every other seed. To determine which seed regions represent putative functional boundaries, the researchers used two thresholds (t) to sort the η^2 values into three classes: (1_C) ($t_{\text{low}} < \eta^2$): no edge; (2_C) ($t_{\text{high}} < \eta^2 < t_{\text{low}}$): edge if neighbor $\eta^2 > t_{\text{high}}$, no edge if $\eta^2 < t_{\text{low}}$; (3_C) ($\eta^2 > t_{\text{high}}$): edge. The idea behind introducing the intermediate class (2_C) is to consider η^2 values only as edges if they are part of a consistent series of points which are above the high threshold. By ordering these continuous values into three distinct classes, the method of Cohen and colleagues is analogous to the philosophical theory of supervaluationism, which attempts to eliminate truth-value gaps in propositions that are expressed by sentences containing vague predicates (Fine 1975). These sentences are also ordered in three classes: (1_S) clear cases (2_S) borderline cases and (3_S) non-cases. Sentences in (1_S) and (3_S) are true or false under all admissible precisifications respectively, while sentences (2_S) are true under some, and false under other admissible precisifications. Thus, every borderline case has a definite truth value after supervaluation. Now considering the analogy of dealing with borderline cases, both supervaluationism and the method of Cohen and colleagues suffer from higher-order vagueness. In case of the former, the term “admissible precisification” in the supervaluationist semantic is itself vague, and therefore the boundaries between (1_S) – (3_S) are fuzzy (cf. Williamson, 1999). In case of the latter, higher-order vagueness is imported through the use of arbitrary thresholds which serve the purpose to ensure spatial stability across short stretches of the η^2 profiles, because what counts as “short” is vague. These fuzzy boundaries between (1_C) – (3_C) were also confirmed empirically in the study of Cohen and colleagues, since they report three seed regions for which the connectivity profile between the angular and supramarginal gyrus did not show an abrupt change.

1.4 *Interlude: Theories of Vagueness and Neuroscience.* Before turning to mechanistic explanations, let me sum up the previous analysis and point out what its implications are for both neuroscience and philosophy. Regarding the neuroanatomical task of delineating cortical areas,

the conceptual framework proposed above allows disentangling different issues that were contained in the introductory quotations of different neuroscientific researchers. With respect to cytoarchitectonics, vagueness of individuation is not identical to inter-individual variability because it is based on one-dimensional gradual transitions between two cortical areas in one brain. Concerning functional neuroimaging, the vagueness of “neural activity” can be distinguished from “ambiguity” of the fMRI signal by means of differentiating causal underspecification from causal density. And in the case of connectivity studies, different forms of vagueness can be described depending from which level the borderline cases get imported. But *vice versa*, the above analysis was supposed to show that empirical research puts different theories of vagueness to a practical test. Regarding such a test, the criterion for a good theory of vagueness should not be whether it is able to solve the *Sorites*-paradox, but whether it allows the speaker to *reasonably deal* with fuzzy boundaries.⁶ As exemplified by the second example in section 1.3, supervaluationism fails to meet such a criterion because it reiterates the problem rather than providing a sensible solution.

A full account of what a theory of vagueness would have to deliver for neuroscience cannot be given here. But there are at least two demands such an account would have to include, namely that (i) parcellations of the brain are context-dependent, i.e. determined by the purpose of research (e.g. functional localization, detecting brain lesions, exploring network structure); and that (ii) different levels of research have different degrees of precision, depending on what instrumentation they use. As far as I can see, a form of non-indexical contextualism that includes more parameters than just place, time and speaker is the only option that can meet these demands (see Åkerman and Greenough, 2010). Such an account assumes that the propositions are individuated by the degrees of precision, i.e. the ability of a speaker to discriminate details of an object by using a concept which applies to it (cf. Keil, 2010, 68). Degrees of precision are themselves preset by the institutional context (not by the individual interest of the speaker). Thus, non-indexical contextualism captures that neuroscience is organized in different levels of research, and that besides instrumentation, the various demands of basic research or practical application will determine how the individual neuroscientist deals with the problems of fuzzy boundaries in the cases described above.

⁶ That may imply that the philosophically more fruitful task would be not to develop a *general* theory of vagueness but rather to inquire how well different theoretical options apply to issues generated by the phenomenon of vagueness in different contexts (for such a project concerning law, see www.unscharfe-grenzen.de).

2. *Vagueness and Mechanisms*

2.1 *Can Vagueness in Neuroanatomy be extended to Mechanisms?* Even if the above conceptual framework were to be held convincing, there is a serious objection against extending vagueness to the level of mechanisms. It holds that mechanisms cannot be vague because neuroanatomical *descriptions* are fundamentally different from neuroscientific *explanations*.⁷ Thus, the brain is treated as a different kind of object, depending on whether its anatomical properties are described or its involvement in the execution of the abilities of an organism is explained. Arguments for the vagueness of predicates or individuation, however, are only applicable when the objects under question are treated as *aggregates*. According to Craver's (2007) account of neuroscientific explanations, aggregates exhibit the following features:

The mass of a pile of sand is an aggregate of the masses of the individual grains. When wholes are sums of their parts, the wholes change continuously with the addition and removal of parts. Intersubstitution of parts makes no difference to the property of the whole. The parts do not interact in ways that are relevant to the aggregate property. The pile gets heavier continuously as one adds new grains of sand, and moving them about has no effect on the weight. Replacing individual grains with equally weighted replicas has no effect on the weight of the pile, and the grains do not interact with one another in ways that influence the weight of the pile (ibid., 186).

It is not only interesting that Craver discusses here the favorite object of philosophers of vagueness, but also that the relevant characteristics of aggregates seem to apply to the structural properties picked out by between cytoarchitectonic studies. If cortical areas are wholes consisting of parts (neuronal cells) which are defined by size, shape and mass, these parts are a) intersubstitutable b) show no interaction (because they are obviously dead) and c) increase the properties of the whole continuously. Therefore it can be inferred that cytoarchitectonic descriptions carry no information about the *function* of cortical areas.

Precisely the absence of functional specification can now be used to argue that a cortical area is not analogous to a cloud, which would be the relevant object of comparison since here, vagueness of individuation is discussed.⁸ The argument for the disanalogy comes from the difference between levels of aggregates and levels of mechanisms. For levels of the latter kind, the part-whole relationship is not one of homogeneity, but one of composition. Furthermore, the characteristic of the composing parts is not that they have a smaller spatial dimension, but that they are *acting entities*, which are organized together such that

⁷ In my discussion of explanations I follow the account of mechanistic explanations as developed for the biological sciences by Machamer, Darden and Craver (2000) and which has been applied to neuroscience by Craver (2002, 2007).

⁸ Understood as a type of description, aggregativity also holds for vaguely individuated non-aggregates such as Tibbles the cat, because plucking hairs does not change the constitution of the cat as a living being.

they constitute a higher-level phenomenon. The phenomenon of spatial memory of rats navigating through a maze, for instance, is constituted by the function of cells in the hippocampus as “spatial maps”, whose generation is explained by the mechanism of Long Term Potentiation at the synaptic level, which is in turn composed of the release and binding of magnesium and calcium ions by NMDA and AMPA receptors (cf. Craver 2002, S89). Since the crucial factor is the higher-level phenomenon, a mechanism is individuated, i.e. picked out amongst other entities of the same kind, via its function, and not via its spatiotemporal structure.⁹ Similarly, the parts of the mechanism are individuated through their role in the causal relation at the level in question. It follows that once cortical areas (such as the hippocampus in the example above) are recruited by a mechanistic explanation, these areas are individuated according to their function and thus, the argument for the (spatiotemporal) vagueness of individuation is blocked.

Although this objection may appear powerful at first glance, I will try to show that of its several assumptions are problematic or cannot be consistently upheld upon Craver’s own account. First, the inference from cytoarchitectonic cell properties to the absence of information about function is challengeable. It was already argued by Korbinian Brodmann, the author of one of the most well-known and widely used parcellation schemes of the human cortex, that anatomical differentiation corresponds to functional differentiation (Brodmann 1909, 289). Thus, how cells are shaped or what size they have is not irrelevant to what function they execute. In present neuroscience, it is also acknowledged that the development of new axonal connections between neurons is itself influenced by repeated utilization (“neurons that fire together, wire together”). Therefore the distinction between function and structure becomes itself blurred (Fingelkurts, et al., 2005, 828). Second, the distinction between neuroanatomical description and mechanistic explanation reflects an unresolved tension between the normative and descriptive demands of the concept of “level” in Craver’s account. On the one hand, *good* neuroscientific explanations are supposed to reveal the causal mechanical structure of the world, which is why neuroscience ought not to describe the brain at the levels of mere spatiotemporal aggregativity or pure psychological function. On the other hand, classifying the brain into different levels does not follow *one* ultimate structure of the world (pace Wimsatt 1976 or Churchland and Sejnowski 2000) but is rather determined by the pragmatic purpose of different research programs (cf. Craver 2002 S89; 2007, 190). Since cytoarchitectonics is regarded as a distinct level of microanatomy

⁹ The primacy of function over spatiotemporal structure does not imply, in my opinion, that mechanisms are abstract entities without spatiotemporal identity conditions. Since the relationship between the higher-level phenomena and the mechanistic parts is one of constitution, and these parts in turn are spatiotemporally organized acting entities, the function of a mechanism (and therefore the mechanism) would not exist without the spatiotemporal structure being in place.

(Devlin and Poldrack 2007) and is acknowledged as its own research program by the neuroscientific community, it would be too much of an abstraction not to regard it as a “level” in Craver’s sense. And since the cytoarchitectonic definition of “cortical area” is a homogenous and architectonically distinct region in the brain (Amunts et al. 2002), it is reasonable to regard the spatiotemporal vagueness at this level as a possible obstacle for good explanations, once cytoarchitectonically defined areas are parts of mechanisms. Furthermore the objection above only holds for the first of the three levels previously outlined. Hence it is not applicable once the vagueness of tasked-based fMRI or resting-state connectivity studies is considered (with regard to mechanisms, I will return to the problem of causal density in section 2.3).

My critique of Craver’s account is not intended to show that his conception does not leave room for indeterminacy in neuroscience. Quite on the contrary, when writing about the spatiotemporal localization of mechanisms in the brain, he writes:

In many cases, the components picked out in a mechanistic decomposition fail to correspond to paradigmatic entities with clear spatial boundaries. The synapse, for instance, is composed of part of a pre-synaptic cell (the axon terminal), part of a post-synaptic cell (the dendrite or bouton), and the gap between them. What unifies these items into a component is their organized behavior: the pre-synaptic cell releases transmitters that traverse the cleft and act on the post-synaptic cell. Synapses are not cells or parts of cells (Craver 2007, 190).

Here, it appears to be that the mechanism of chemical transmission at the synapse itself introduces some indeterminacy into neuroscientific explanations of the brain. Hence it is surprising that in a more recent paper, Craver describes the imprecise localization to be only a feature of our limited epistemic access, and not a characteristic of the mechanism in question: “Indeed a structural component might be so distributed and diffuse as to defy tidy structural description, though it no doubt has one if we had time, knowledge and patience to formulate it” (Picinini and Craver 2011, 291). The apparent tension between both quotations seems to reflect a gap within the mechanistic account of explanation which has not been filled in with a consistent interpretation. Before closing this gap by interpreting the imprecise localization of mechanisms in the brain as part of the vagueness of neuroscientific explanations (section 2.3), let me discuss how the recent discussion in the philosophy of psychiatry may elucidate such an interpretation.

2.2 *‘Fuzziness’ in the recent literature on psychiatric kinds.* At first glance, it does not appear to be obvious how the debates about classifying psychiatric disorders may help to understand the nature of vagueness in mechanistic explanations of neuroscience. Since the introduction of the third *Diagnostic and Statistical Manual of Mental Disorders (DSM III)*, psychiatric classifications are intended to be atheroetical and do not attempt to *explain* how the classified disorders came into

being (American Psychiatric Association 1980). Based on their atheoretical character, it may be assumed that the diagnostic criteria do not form parts of a mechanism for the psychiatric kinds in question, because their presence does not imply a causal relevance for the genesis of the disorder—at least as long as the symptom is not part of the etiology of a disease. Despite these purported differences, however, mechanistic explanations in neuroscience and psychiatric classifications share important characteristics. Both are essentially multi-level (Craver 2002; Kendler 2012a), albeit that the levels can differ in their grades of resolution. While neuroscientific explanations usually stop at the behavioral level (as in the case of the rat example above), psychiatric diagnoses often comprise supra-individual factors such as environmental and/or social influences. The contribution of various levels to the completeness of the mechanistic explanation or psychiatric diagnosis in question may also support the claim that there is no fundamental level the higher-level phenomena can be reduced to (Craver 2007, Kendler 2012b). Taking these similarities together with the influence of both biological and cultural factors upon psychiatric disorders, there is a question equivalent to whether the imprecise spatial location of mechanisms results from our epistemic practices or the entities themselves: It is whether the overlapping character of current diagnostic categories is a limitation of our classificatory practices or results from the disorders themselves.

In order to tackle the latter question, Kendler, and Zachar and Craver (2011) introduced the notion of “mechanistic property cluster kinds” (MPC) to characterize what kinds of things psychiatric disorders are. Their notion is supposed to capture the intuitions of two contrasting kind conceptions without suffering from their disadvantages. From the *essentialist* kind conception they adopt the assumption that kinds are real, i.e. existing independently of our classifications, without insisting that one etiologic agent is directly and causally responsible for all symptoms of the disorder. Thus allowing for multi-level characterizations, the MPC conception shares with the *social constructivist* kind conception that social factors influence the manifestation of disorders, but the disorders themselves do not disappear when the *concepts* of the classifications change. The stability of the categories over time is guaranteed through the same *causal mechanisms* which are underlying specific disorders. But since these mechanisms span multiple levels and the symptoms they produce causally influence each other, the authors characterize the “boundaries between MPC kinds to be fuzzier than with essentialist kinds” (ibid., 1148). How one disorder is distinguished from another is, therefore, often influenced by the practical goals of the classifying scientists (such as reliable diagnosis, prognosis etc.).

Similarly to the example of the overlapping community boundaries in section 1.3, however, these authors leave open the question whether

the fuzzy boundaries of MPC kinds can be understood in the technical sense of vagueness, i.e. as admitting for borderline cases. Contrary to Wu and colleagues, however, the notion of ‘fuzziness’ is imported into the MPC conception through the adaptation of Richard Boyd’s “homeostatic property cluster” (HPC) conception of kinds (Boyd, 1999). According to Boyd, kinds are defined by properties whose co-occurrence is temporarily stable. But since it is not required for an individual to exhibit all properties of the property cluster that define the kind, there is an extensional indeterminacy with regard to the membership of certain individuals that share some, but not all features of the kind in question (comp. *ibid.*, 144). Although Boyd is alluding to vagueness here he leaves open which theoretical option is best suited to interpret his conception. If one follows the characterization of biological species as ‘fuzzy sets’ by Kendler (2012a), the HPC (and subsequently the MPC) kinds may be interpreted by the fuzzy logic approach developed by Zadeh (1965). Fuzzy logic applies different degrees of truth to the propositions expressed by sentences that contain vague predicates. Applied to kinds, every individual would receive a value of membership depending on how many properties of the relevant cluster it realizes. Thus, any individual with a value of membership around 0.5 would be considered a borderline case for which it is indeterminate whether it is a member of the kind or not. There are several general and particular reasons, however, why the fuzzy logic interpretation is problematic. Strictly speaking, biological species cannot *be* sets, because sets have extensional identity conditions: the identity of a set depends on the number of members while the identity of a kind does not. With respect to vagueness, general utility of fuzzy logic to deal with vague predicates has been questioned because it is over over-precisifying statements whose communicative use is generated by their fuzzy referential boundaries (cf. Keil 2010). The point can be illustrated by the example of a psychiatrist using a diagnostic manual. It does not help to know that a patient fulfills the diagnosis of schizophrenia by 0.47 or 0.53 because these values are too precise to decide the binary question of applying the diagnostic category or not. Moreover, stipulating two cut-off points that define the class of borderline cases is arbitrary (comp. section 1.2 and 1.3) and eliminates the practical utility of the classification system to allow for treatment or non-treatment in individual cases (false positives and negatives). The fuzzy logic approach moreover models continuity and not imprecision (comp. Pinkal 1995, 166). But neither do the properties of HPC/MPC kinds have to be gradable, nor does every combination of properties have to be realized in nature (cf. Hauswald 2012).

A theoretical alternative that can accommodate for the lacks of the fuzzy logic approach is combinatory vagueness (Alston, 1967). Combinatorially vague concepts (e.g. ‘religion’) are a subclass of family resemblance concepts (Wittgenstein 1953/2001) that allow for borderline cases of application (e.g. ‘religion’ to the Quaker movement or ideologies

such as communism). Although there are clear examples (e.g. Catholicism), the list of features characterizing such a concept does neither comprise a list of necessary and together sufficient conditions, nor does it require the shared characteristics of the ‘family members’ to stand in the same logical relation to each other. Now if HPC kinds are considered as combinatorially vague, then there exist individuals that realize a number of properties such that it is indeterminate whether they are a member of the kind or not. Notice that in contrast to the fuzzy logic approach, the identification of borderline cases does not additionally require the assignment of membership values but only the manifestation of properties in the individuals themselves (plus somebody who classifies them into kinds).¹⁰ With regard to the MPC conception of psychiatric disorders, the fuzzy boundaries between different diagnostic categories, two kinds of vagueness can arise. First, it may be indeterminate whether or not an individual should be diagnosed as having a psychiatric disorder when he exhibits a borderline number of diagnostic criteria (e.g. having two of the five symptoms of schizophrenia). Second, it may be indeterminate whether or not scientists ought to split a category into two, based on the discovery that a particular mechanism causally contributes to the manifestation of a disorder (comp. Kendler, Zachar and Craver 2011, 1149).

2.3 Two interpretations of the vagueness of mechanistic explanations. The combinatory vagueness approach can now be applied to mechanistic explanations in neuroscience. What remained open from section 2.1 is whether the mechanisms themselves or our descriptions of them are vague. With regard to the MPC kind conception, the combinatory vagueness appears to be ontic, because it arises from the distribution of properties among the individuals themselves. It is unclear, however, whether the ontic interpretation is upheld by the authors discussed above, since they write that “in our current state of relative ignorance about the mind/brain [...] the boundaries between these mechanisms may be as confusing as those between the disorders themselves, and the mapping of these mechanisms onto our current clinical syndromes may be anything but pretty” (ibid., 1148). The position expressed here resembles Picinini and Craver who assert that in the end, mechanisms *do* have a determinate spatiotemporal localization regardless of whether we are able to uncover it. Thus there is an apparent tension: how can mechanisms be real and ontologically determinate and at the same time combinatorially vague as outlined above?

¹⁰ Hauswald (2012) furthermore argues that the HPC conception is compatible with semantic externalism and therefore the kind terms do not—in contrast to combinatorially vague or family resemblance concepts—possess any intension. Although I am employing the same interpretation of combinatory vagueness, I am neutral with regard to the question which semantic theory is most appropriate to understand Boyd’s approach.

The straightforward answer to this question is to interpret vagueness as semantic, i.e. as a feature of our language. Most philosophers regard the semantic version of vagueness to be the only intelligible one (Russell 1923, Dummett 1975, Lewis 1986, Varzi 2001). On the semantic reading, vagueness is the kind of mismatch between language and the world where speakers cannot delineate precisely which objects fall under the concepts they are using. And indeed, the semantic version does capture an important aspect of the MPC kind conception, where the disorder concepts are variable, thus creating potential mismatches between diagnostic categories and the mechanisms underlying the disorders. Similarly, the incompleteness of our *descriptions* may lead to a mismatch between our explanatory sketch and the brain mechanism in question. But why should our incomplete knowledge count as semantic vagueness? According to the conceptual framework introduced above, it seems like the indeterminacy here is one of epistemic uncertainty and not of spatiotemporal fuzziness. I believe that the sharp distinction between these two kinds of indeterminacy comes from the assumption that no additional empirical evidence can resolve conceptual vagueness. But it is an assumption untenable for most issues of indeterminacy in the empirical sciences, and it makes the philosophical concept of vagueness uninteresting for most scientific practitioners. Scientific knowledge is naturally incomplete, and the fuzzy concepts scientists utilize change according to the demands of empirical discovery.¹¹ Thus if the strict demarcation from epistemic uncertainty is loosened, semantic vagueness can be understood as epistemic because it reflects our inability to precisely grasp the ultimate structure of the world (I will argue in section 2.4 why I find the epistemic/ontic divide problematic as well).

Conceived as a semantic problem, vagueness in neuroscientific explanation arises when a scientist has to decide which and how many causal factors have to be included in the description of a mechanism. Ontologically, “the boundaries of mechanisms—what is in the mechanism and what is not—are fixed by reference to the phenomenon that the mechanism explains” (Craver 2007, 123). But although the identity conditions of the mechanisms may be set by the phenomenon, identifying the causal components is only possible through experimental manipulation. Now consider that in four different experiments, the mech-

¹¹ The fact that there is almost no debate on vagueness and scientific concepts can be partially explained by the history of philosophy itself. Under the vein of ideal language approaches, philosophers thought that vagueness is a defect of natural languages that can be eliminated by logical or mathematical formalization (Frege 1893/1903; Russell 1923). In 1939, a short debate arose between Max Black and C.G. Hempel about the role of vagueness in science. But after the rise of ordinary language philosophy (Wittgenstein 1953/2001; Austin 1975), the conviction that vagueness is ineliminable from any language led philosophers again to develop general theories of vagueness. Only recently, writers attempted to to apply vagueness to issues in the special sciences, such as microphysics (Chibeni 2004).

anism A which is individuated by the phenomenon P and composed by the components XYZ E is experimentally investigated in four studies of the individuals $I_1 - I_4$. In I_4 however, only component E is detected, and thus the question arises whether the neuroscientists in the fourth experiment investigated the same phenomenon. Thus I_4 is an instance of a borderline case of P_A because it is indeterminate whether or not P_4 has been produced by A.¹² The vagueness here described is combinatory because it arises from comparing different contexts of research. In typical neuroscientific experiments (such as studying the dorsal and ventral pathway of visual processing), phenomena can be investigated through different why-questions (e.g. ‘why does the dissociation between dorsal and ventral pathway result in visual but not grasping illusions?’) such that only parts of the components appear to be causally relevant for the particular context of research. In most cases, there is no decision rule to determine which or how many components have to be included in order for an experiment to count as a class-member of why-questions which investigates the same phenomenon.

According to Sirtes (2010), the problem of causal (or explanatory) relevance arises in Craver’s account because he maintains two versions of the explanandum phenomenon. In the *ideal* version quoted above, the components of the mechanism are potentially infinite and thus lead to the actual combinatory vagueness of various descriptions of the mechanism. In the *partial* version, every particular why-question would pick out a different mechanism and thus result in a “very messy ontology” (ibid., 16). Leaving aside the problem of ontology for the moment (I will return to it in section 2.4), let me point out how Sirtes’ pragmatic-ontic account of mechanisms elucidates their spatiotemporal vagueness. He argues that even the mutual manipulability criteria introduced later in Craver’s book do not solve the problem of explanatory relevance, because they do not determine what *degree of precision* the mechanistic explanation will have. These criteria hold that an acting entity is part of a mechanism if its activity changes once the system is manipulated and the activity system changes once the entity is manipulated (comp. Craver 2007, 155). Sirtes’ point is now that in biological species, the type-token relationship of mechanisms is one of family resemblance, and therefore the manipulation of an entity may result in the activity of one system but not the other. Thus his criticism supports the characterization of Craver’s account of mechanisms as combinatorially vague, when he writes that there is “no objective way of deciding where one family of mechanisms stops and another begins” (ibid., 19). Depending on the degree of precision in the description, the scope of tokens one type of mechanism subsumes will vary. But the introduction of degrees of precision makes it also possible to account for the vagueness of cortical areas described in the first part of the paper. Since parcelations of the brain are context-dependent, investigating mechanisms

¹² This is also called the „subtyping problem“. I am thankful to Beate Krickel for pointing out the connection to my interpretation of combinatory vagueness.

under different experimental conditions also implies that the spatial boundaries picked out by the mechanistic description can be vague. For example: using a cytoarchitectonic atlas to allocate the mechanism of spatial guidance of movement (Taira 1990) in the premotor cortex suffers from the vague delineation of BA 6. If the spatial boundaries of the same mechanism would be allocated differently by another study using a connectivity-based parcellation (Wise et al., 1997), the spatial boundary can additionally be combinatorially vague.

If non-indexical contextualism (section 1.4) is now taken to characterize the vagueness of mechanistic explanations, it becomes apparent that vague mechanistic descriptions may not only disadvantageous. Leaving out details at lower levels can be now understood as a change in the degree of precision, which allows researchers to express mechanistic explanations when investigating higher-level phenomena. To take Sirtes' example: A neuroimaging scientist can talk about the action potential using the coarse vocabulary of a textbook (hyper- and depolarization, sodium and potassium channels etc.), while an electrophysiologist can talk about the sodium-voltage gated channels Nav1.1, Nav1.6 in human Purkinje cells. Here, the vaguely defined mechanism of the action-potential ensures communication across different neuroscientific subdisciplines. Both researchers can say something true about the *same* mechanism because the disciplinary context sets the degree of precision that is utilized in their description. The drawback of the semantic interpretation of vagueness is that it remains unclear how scientists find out about the identity conditions of the 'real' mechanisms. Craver's anti-reductive position prohibits a decision at which degree of precision the description of the mechanism is the most accurate. The mechanistic account thus runs the risk of becoming a "frictionless spinning in the void", to utilize John McDowell's phrase (1994, 67).

In order to avoid the dilemma, mechanisms could be interpreted as ontically vague. Notice that such an interpretation does not hold that vagueness is an additional *property*, such as the degree of membership in the fuzzy logic interpretation of MPC kinds. Everything that is needed is a non-linguistic *source* of vagueness (Tye 1990) which in this case is that the mechanisms are realized by a causally dense system, namely the brain.¹³ A causally relevant factor is now understood to be part of the causal mechanistic relation itself, not only as part of its description. Because of the weak causal connection between neurons in the brain, the mechanisms are tolerant to small changes and thus become susceptible to Sorites-reasoning: 'if this particular neuron is not causally relevant, this one cannot be causally relevant either', etc.¹⁴

¹³ Nothing that will be said in the following hangs on the term 'brain', as the mechanism is realized by a causally dense system even if it spans through the whole organism and possibly even parts of the environment.

¹⁴ The Sorites-series outlined above only works for probabilistic conceptions of causation, which is allowed by Craver (2007), 103 and Kendler, Zachar and Craver (2011), 1148.

The result of the procedure should by now be clear: A mechanism is a vaguely individuated entity because there are acting neurons such that there is no matter of fact whether or not these neurons are part of the mechanism. Notice that everything outlined for the semantic version also holds for the ontic version: mechanisms are a) combinatorially vague since they can be realized in individually varying brains with different subsets of neurons and b) neutral with respect to the degree of precision with which the acting entities are individuated (e.g. spiking outputs of single neurons, phase synchronization of neuronal clusters, cortical microcircuits, large-scale cortical networks). An implication of a) is that they furthermore leave the neuroscientist with the problem of c) deciding whether a borderline case of a phenomenon may be constituted by a different mechanism. The only difference to the semantic version is that the scientists themselves discover the vagueness of the 'real' mechanisms. I do not want to delve into the question whether ontic vagueness is a coherent philosophical position or not. It suffices to point out that interpreting mechanism with the standard version is unsatisfying in the same respect as semantic vagueness if it is sharply demarcated from epistemic uncertainty. Firstly, it follows from b) that mechanisms have also borderline causal factors well below and above anything neuroscientists are able to investigate, such as quarks or whole biospheres. Secondly, and more importantly, ontic vagueness would hardly worry any scientific practitioner if there is *no* matter of fact to be discovered about the vague boundary of a mechanism. The question which thus remains is how the vagueness of mechanisms becomes a (neuro-)scientifically interesting problem.

2.4 Reconciling ontic and semantic vagueness of mechanisms through philosophy of scientific practice. The framework of Joseph Rouse's pragmatic naturalism is in my opinion suited to explain why vagueness of mechanisms should matter to neuroscientists. In his extensive discussion, which I can only sketch here, Rouse (2002) attempts to show that the philosophical debate about naturalism and anti-naturalism suffers from a false dualism between nature and normativity (comp. *ibid.*, 11). His attempt to undermine the dualism is to argue that nature becomes itself normatively accountable through its manifestation in scientific practice. Thus the picture of science is not one of finding out about the world from a God's eye perspective, but one where there is "something at stake" by choosing between alternative ways of engaging in scientific practice (*ibid.*, 260). The implication of this view for the philosophy of neuroscience is that mechanisms cannot be described from a "view from nowhere" either. The material-discursive settings within which they are situated need to be taken into account. More recently Rouse (2011) has taken up the problem alluded to by McDowell's phrase in section 2.3. His point is that scientific concepts do not gain their experiential content by perceptual receptivity but through the isolation and stabilization of phenomena in experimental systems (comp. *ibid.*, 245;

furthermore Hacking 1983, 220ff. and Rheinberger 1997). The result is that entire domains of the world become only accessible through the establishment of an experimental practice investigating it. Rouse's conception therefore accommodates for an important feature of vagueness in neuroscience, namely that concepts like 'cortical area' of 'action potential' deal with a domain where we do not know the boundaries, "for non [have] been so far drawn" (Wittgenstein 1953/2001, §68, 33). Since vagueness is modeled on everyday concepts which usually possess extensional boundaries—however vague—the semantic version suggests a split between our descriptions and the mechanisms themselves that is not present in scientific practice (and arguably, not even in any other domain of human activity).

Placing vagueness in scientific practice also repudiates Sirtes' ontological worries and makes the ontic version scientifically interesting. If the problem of explanatory relevance in Craver's account is interpreted as combinatory vagueness, it is false that *every* why-question describes a different mechanism, because there are always clear instances, and not only borderline cases. The clear cases may not be necessary because they depend on the contingent distribution of causal factors in the world. But they are real insofar as they explain the phenomena that can be isolated and stabilized under laboratory conditions (cf. Rouse 2011, 251). Furthermore, I believe that Sirtes' refusal to grant mechanisms' existence stems from a too conservative picture of ontology. It is by no means clear whether science always attempts to explain everything with as "few entities as possible", (Sirtes 2010, 16; the sixteen elementary particles of the standard model may be a point in case). Furthermore, the exclusion of pragmatic factors such as the degree of precision or the causal relevance relation from ontology results from the wrong dichotomy of concepts and the world. Once the experimental interventions of the scientists themselves are understood as constitutive for conceptual articulation, mechanisms isolated in the laboratory and their descriptions become co-substantive:

Practices are not something *else* that accounts for or accomplishes the formation and location of such patterns and boundaries in within the world. They *are* the establishment and maintenance of boundaries and the accountability that sustains them. Practices do not preexist the objects, boundaries and stakes they constitute (Rouse 2002, 289).

Rouse's quotation also makes possible to see what implications ontically vague mechanisms have for neuroscientific practice. Through its embeddedness in experimental systems, the class of borderline cases of a mechanism is restricted to the entities which a scientific practitioner can manipulate. Once more, non-indexical contextualism is of help here, because the degrees of precision in neuroscience are themselves something *material*. They are set by the ability to discriminate between two points of measurement (e.g. pixel, voxel) in the brain with the corresponding imaging modality (e.g. EEG, PET, fMRI) or invasive instrument (e.g. scalpel, electrode). These degrees are furthermore

normative insofar as they are part of a research context within which a specific set of instruments, programs and routines of action is available. Only through these sets, the boundaries that appear in cytoarchitectonic atlases, fMRI activation maps, or the diagram of a mechanism like Long Term Potentiation can be articulated. But looking at these visualizations does not reveal the degree of precision with which they have been generated (cf. Keil, 2010, 65). Thus, there is potential for conflicting normative demands in different research contexts: a neurosurgeon may not be able to discriminate the degree of precision of a cytoarchitectonic map generated by a neuroanatomist. But acting upon such a map in a surgery *creates* matters of fact. *These* are the facts why vagueness *matters*, and not that there is no matter of fact to where a spatial or functional boundary lies in the cortex. The example can be extended to mechanisms, because their spatiotemporal location is important to know where to intervene for their manipulation, be it experimentally or neurosurgically. To sum up: the vagueness of mechanisms is ontic because how neuroscientists deal with it has *real* consequences (comp. Rouse 2002, 156). But it is equally semantic because the concepts that articulate mechanistic boundaries provide the possibility for further action.

Conclusion

My paper tried to develop a conceptual framework which is adequate to capture the phenomenon of fuzzy boundaries between cortical areas in neuroanatomy. I furthermore argued that vagueness in neuroanatomy can be extended to mechanistic explanations, because they invoke the knowledge of the three levels of research described in the first part of the paper. The recent literature on psychiatric kinds showed that the mechanistic framework leaves conceptual room for vagueness, but a consistent interpretation was missing so far. Again, vagueness in neuroanatomy and mechanistic explanations does not imply that functional or spatial localization is *impossible*, because there are always clear cases. But factors such as degrees of precision and the context of research have to be taken into account to understand how boundaries are determined in neuroscientific practice, as shown in section 2.4. To be adequate, however, an account of vagueness also has to show why indeterminacy matters for the particular scientist dealing with it. A correct description alone is insufficient. To illustrate my point, think about whether it matters to any geographer who studies mountains if a certain atom is part of Mt. Everest or not. Examples such as these are the reason why philosophical discussions about vagueness of individuation remain without implications for the sciences.

My refutation of the ontic version of vague mechanisms in section 2.3 was not meant to decide the *metaphysical* question what kinds of things mechanisms are. I agree with Sirtes that Craver's picture of mechanisms having smaller parts which are themselves also mechanisms may

be too tidy. What I do not share is his conclusion that if mechanisms are no acting entities, they do not exist at all.¹⁵ The pragmatic-naturalistic conception of vague mechanisms in section 2.4 is agnostic with regard to these questions, because it points towards a different focus of further research. Considering degrees of precision in different research contexts becomes especially interesting once brain maps or mechanistic explanations travel as “brain facts” (Choudhury and Slaby 2012, 30) outside the laboratory into clinical applications or other areas of society (such as law or public media, see also Dumit 2004). Here, philosophy of scientific practice can play a constructively critical role in understanding how people deal with vagueness in these areas.

References

- Åkerman, J., and Greenough, P. 2010. “Vagueness and Non-Indexical Contextualism.” In *New Waves in Philosophy of Language*, ed. S. Sawyer. Basingstoke New York: Palgrave Macmillan. 8–23.
- Alston, W. 1967. “Vagueness”. In *The Encyclopedia of Philosophy*, vol. 7, ed. P. Edwards. London and New York: MacMillan. 218–21.
- American Psychiatric Association 1980. *Third Diagnostic and Statistical Manual of Mental Disorders*. Washington, DC: American Psychiatric Association.
- Brodmann, K. 1909. *Vergleichende Lokalisationslehre der Großhirnrinde*. Leipzig: Barth.
- Chibeni, S. S. 2004. “Ontic Vagueness in Microphysics”. *Sorites* (15):29–41.
- Choudhury, S., & Slaby, J. (eds.) 2012. *Critical Neuroscience: A Handbook of the Social and Cultural Contexts of Neuroscience*. Chichester: Wiley-Blackwell.
- Churchland, P., and Sejnowski, T. J. 2000. “Perspectives on Neuroscience.” In *Cognitive Neuroscience*, ed. M. S. O. Gazzaniga. Oxford: Blackwell. 14–24.
- Craver, C.F. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Oxford University Press.
- 2002. “Interlevel Experiments and Multilevel Mechanisms in the Neuroscience of Memory”. *Philosophy of Science*, Supplemental 69, S83–S97.
- Devlin, J., and Poldrack, R.A. 2007. “In praise of tedious anatomy”. *Neuroimage* 37 (4):1033–1041.
- Dumit, J. 2004. *Picturing Personhood. Brain Scans and Biomedical Identity*. Princeton: Princeton University Press.
- Dummett, Michael 1975. “Wang’s paradox”. *Synthese* (30):301–324.
- Ettlinger G. 1990. “Object vision’ and ‘spatial vision’: the neuropsychological evidence for the distinction”. *Cortex* 26 (3):319–41.
- Gorgolewski, K., Bastin, M. Rigolo, L. Soleiman, H.A., Pernet, C., Storkey,

¹⁵ An alternative metaphysical view which conceives of mechanisms as processes consisting of temporally extended events has been proposed by Krickel (2012). Her position has the interesting consequence that inter-level causation is possible.

- A., Golby, A.J.. 2011. "Pitfalls of Thresholding Statistical Maps in Presurgical fMRI Mapping". In *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine*. Montreal, 2430.
- Economo, C., and Koskinas, G. von 1925. *Die Cytoarchitektur der Hirnrinde des erwachsenen Menschen*. Berlin: Springer.
- Fine, K. 1975. "Vagueness, Truth and Logic". *Synthese* 30 (3-4):265–300. doi: 10.1007/BF00485047.
- Fingelkurts, A., Fingelkurts, A., and Kahkonen, S. 2005. "Functional connectivity in the brain—is it an elusive concept?". *Neurosci Biobehav Rev* 28 (8). doi: 10.1016/j.neubiorev.2004.10.009.
- Hacking, I. 1983. *Representing and Intervening*. Cambridge: Cambridge University Press.
- Haueis, P. (forthcoming). "The Fuzzy Brain. Vagueness and Mapping Connectivity" in the Human Cortex. *Frontiers in Neuroanatomy*, Special Issue on "Mapping Connectivity of the Human Cerebral Cortex".
- Hauswald, R. 2012. "Was sind vage natürliche Arten?", Conference Presentation at GAP.8, September 17–20 2012, University of Konstanz.
- Keil, G. 2010. "Halbglätzen statt Halbwahrheiten. Über Vagheit, Wahrheits- und Auflösungsgrade". In *Wahrheit - Bedeutung - Existenz*, ed. M. Grajner and A. Rami. Frankfurt, M. [i.e. Heusenstamm], Paris, Lancaster, New Brunswick, NJ: Ontos-Verl), 57–86.
- Kendler, K. S. 2012a. "Levels of Explanation in Psychiatric and Substance Use Disorders: Implications for the Development of An Etiologically Based Nosology". *Molecular Psychiatry* 17 (1):1–12.
- _____. 2012b. "The Dappled Nature of Causes of Psychiatric Illness: Replacing the Organic-functional/hardware-software Dichotomy with Empirically Based Pluralism". *Molecular Psychiatry* 17 (1):11–21.
- Kendler, K.S., Zachar, P. and Craver, C.F. 2011. "What kinds of things are psychiatric disorders *Psychological Medicine* 41:1143–1150.
- Krickel, B. 2012. "Making sense of Interlevel Causation". Conference Presentation at GAP.8, September 17–20 2012, University of Konstanz.
- Lewis, D. 1986. *On the plurality of worlds*. Oxford: Blackwell.
- Machamer, P., Darden, L., Craver, C.F. 2000. "Thinking about Mechanisms". *Philosophy of Science* 67:1–25.
- McDowell, J. 1994. *Mind and World*. Cambridge, MA: Harvard University Press.
- Mechelli, A., Friston, K., Frackowiak, R., and Price, C. 2005. "Structural covariance in the human cortex". *J Neurosci* 25 (36). doi: 10.1523/JNEUROSCI.0357-05.2005.
- Ojemann, J. G., Akbudak, E., Snyder, A. Z., McKinstry, R. C., Raichle, M. E., and Conturo, T. E. 1997. "Anatomic localization and quantitative analysis of gradient refocused echo-planar fMRI susceptibility artifacts". *Neuroimage* 6 (3). doi: 10.1006/nimg.1997.0289.
- Piccinini, G. and Craver, C.F. 2011. "Integrating Psychology and Neuroscience: Functional Analyses as Mechanism Sketches". *Synthese* 183(3):283-311.
- Pinkal, M. 1995. *Logic and lexicon: The Semantics of the Indefinite*. Dordrecht: Kluwer Academic Press.

- Rajkowska, G., and Goldman-Rakic, P. S. 1995a. "Cytoarchitectonic Definition of Prefrontal Areas in the Normal Human Cortex: I. Remapping of Areas 9 and 46 using Quantitative Criteria". *Cerebral Cortex* 5 (4):307–22.
- Rajkowska, G., and Goldman-Rakic, P. S. 1995b. "Cytoarchitectonic Definition of Prefrontal Areas in the Normal Human Cortex: II. Variability in Locations of Areas 9 and 46 and Relationship to the Talairach Coordinate System". *Cerebral Cortex* 5 (4):323–37.
- Rheinberger, H.J. 1997. *Toward a History of Epistemic things*. Stanford, CA: Stanford University Press.
- Russell, B. 1923. "Vagueness". *The Australasian Journal of Psychology and Philosophy* 1 (2):84–92.
- Rouse, J. 2002. *How Scientific Practices Matter*. Chicago, IL: University of Chicago Press.
- _____. 2011. "Articulating the World: Experimental Systems and Conceptual Understanding". *International Studies in the Philosophy of Science* 25(3):243–254.
- Savoy, R. 2001. "History and future directions of human brain mapping and functional imaging". *Act. Psychol.* 107:9–42.
- Sirtes, Daniel 2010. "A Pragmatic-Ontic Account of Mechanistic Explanation". [Preprint] Accessed online at <http://philsci-archiv.pitt.edu/5181/07/31/2012>, 1:26 pm.
- Simmons, J.P., Nelson, L.D. and Simonsohn, U. 2011. "False Positive Psychology. Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant". *Psychological Science* 22 (11):1359–1366.
- Taira, M., Mine, S., Georgopoulos, A.P., Murata, A., Sakata, A. 1990. "Parietal Cortex neurons of the monkey related to visual guidance of hand movement". *Experimental Brain Research* (83):29–36.
- Tehovnik, E. J., Tolias, A. S., Sultan, F., Slocum, W. M., and Logothetis, N. K. 2006. "Direct and Indirect Activation of Cortical Neurons by Electrical Microstimulation". *Journal of Neurophysiology* 96 (2):512–21.
- Unger, P. 1980. "The Problem of the Many". *Midwest Studies in Philosophy* 5:411–67.
- Varzi, A. 2001. "Vagueness in Geography". *Philosophy and Geography* 4 (1):49–65.
- Wimsatt, W. 1976. "Reductionism, Levels of Organization, and the Mind–Body Problem". In G. Globus, I. Savodnik, and G. Maxwell (eds.), *Consciousness and the Brain*. New York: Plenum Press. 199–267.
- Wise, S. P., Boussaoud D., Johnson P. B., Caminiti R. 1997. "Premotor and parietal cortex: corticocortical connectivity and combinatorial computations". *Annu. Rev. Neurosci.* 20:25–42.
- Weiner, J. 2007. "Science and Semantics: The Case of Vagueness and Supervaluation". *Pacific Philosophical Quarterly* 88:355–74.
- Wu, K., Taki, Y., Sato, K., Sassa, Y., Inoue, K., Goto, R., Okada, K. et al. 2011. "The overlapping community structure of structural brain network in young healthy individuals". *PLoS One* 6 (5). doi: 10.1371/journal.pone.0019608.
- Williamson, T. 1999. "On the structure of higher-order vagueness". *Mind* 108 (429):127–43.

- Wittgenstein, L. 1953/2001. *Philosophical investigations*. The German text, with a revised English translation. 3rd ed. Oxford and Malden, Mass: Blackwell.
- Zadeh, L. 1965. "Fuzzy sets". *Information and Control* 8:338–53.
- Zilles, K., and Schleicher, A. 1993. "Cyto- and myeloarchitecture of human visual cortex and the periodical GABA-A receptor distribution." In *Functional organization of the human visual cortex*, ed. B. Gulyas, D. Ottoson, and P. Roland. Oxford: Pergamon Press. 111–20.
- Zilles, K., Schlaug, G., Matelli, M., Luppino, G., Schleicher, A., Qu, M., Dabringhaus, A., Seitz, R., and Roland, P. E. 1995. "Mapping of human and macaque sensorimotor areas by integrating architectonic, transmitter receptor, MRI and PET data". *J Anat* 187 (Pt 3):515-37.