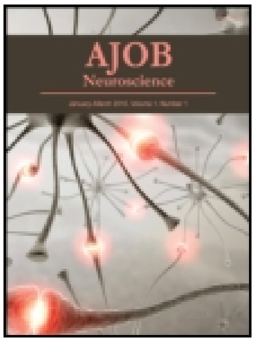

AJOB Neuroscience

ISSN: 2150-7740 (Print) 2150-7759 (Online) Journal homepage: http://www.tandfonline.com/loi/uabn20

Ethical Challenges of Simulation-Driven Big Neuroscience

Markus Christen, Nikola Biller-Andorno, Berit Bringedal, Kevin Grimes, Julian Savulescu & Henrik Walter

To cite this article: Markus Christen, Nikola Biller-Andorno, Berit Bringedal, Kevin Grimes, Julian Savulescu & Henrik Walter (2016) Ethical Challenges of Simulation-Driven Big Neuroscience, AJOB Neuroscience, 7:1, 5-17

To link to this article: http://dx.doi.org/10.1080/21507740.2015.1135831

Published online: 04 Apr 2016.

AJOB Neuroscience, 7(1): 5–17, 2016

ISSN: 2150-7740 print / 2150-7759 online
DOI: 10.1080/21507740.2015.1135831

Target Article

Ethical Challenges of Simulation-Driven Big Neuroscience

Markus Christen, University of Zurich
Nikola Biller-Andorno, University of Zurich
Berit Bringedal, Institute for Studies of the Medical Profession
Kevin Grimes, Karolinksa Institutet
Julian Savulescu, University of Oxford
Henrik Walter, Charité-Universitätsmedizin Berlin

Research in neuroscience traditionally relies on rather small groups that deal with different questions on all levels of neuronal organization. Recent funding initiatives—notably the European "Human Brain Project" (HBP)—aim to promote Big Neuroscience for integrating research and unifying knowledge. This approach is characterized by two aspects: first, by many interacting researchers from various disciplines that deal with heterogeneous data and are accountable to a large public funding source; and second, by a decisive role of information and communication technology (ICT) as an instrument not only to perform but also to structure and guide scientific activities, for example, through simulations in the case of the HBP. We argue that Big Neuroscience comes along with specific ethical challenges. By examining the justification of Big Neuroscience and the role and effects of ICT on social interaction of researchers and knowledge production, we provide suggestions to address these challenges.

Keywords: Big Neuroscience, brain simulation, ethics in organizations, Human Brain Project, knowledge production, neuroethics

In January 2013, the Human Brain Project (HBP) was announced as one of two flagship projects funded by the European Commission's Future and Emerging Technologies Programme (FET). The matched funding for the HBP of about 1.16 billion Euros over 10 years provided by the European Union (EU) and partner shall enable a concerted effort to "lay the technical foundations for a new model of ICT-based brain research, driving integration between data and knowledge from different disciplines, and catalyzing a community effort to achieve a new understanding of the brain, new treatments for brain disease and new brain-like computing technologies" (HBP Report 2012, 3). Only a few weeks later, U.S. President Barack Obama announced the "Brain Research through Advancing Innovative Neurotechnologies (BRAIN) Initiative" in the United States, aiming to invest several billion dollars to examine the workings of the human brain (Markoff 2013). Also, Australia, China, and Japan have recently announced large-scale projects in neuroscience (Grillner 2014).

Such large-scale initiatives aim to involve close collaboration among hundreds of scientists from widely diverse disciplines, highly organized into several scientifically and technically distinct work subdomains (e.g., the HBP has 12 research areas and more than 100 partner institutions) for many years. Such initiatives answer to large public funding organizations (the HBP to the European Commission). Moreover, these consortia must address specific challenges to accountability and governance of the projects. In terms of duration, investment, and number of people involved, these projects are comparable in size to, up to this time, the largest Big Science project in biology, the Human Genome Project (1990–2003, US$2.7 billion, consortium of 20 institutions; see http://www.genome.gov/11006943). Furthermore, such large-scale initiatives must emphasize methodology. For example, the BRAIN Initiative focuses—as the name indicates—on neurotechnologies (e.g., imaging), and the HBP aims for an information and communications technology (ICT) infrastructure for setting up an experimental facility to study and to simulate the structure and functions of the human brain. The latter aligns with the huge transformation of various scientific fields that rely increasingly on computer power not only to organize

Address correspondence to Markus Christen, University Research Priority Program Ethics, University of Zurich, Zollikerstrasse 117, Zurich CH-8008, Switzerland. E-mail: christen@ethik.uzh.ch

data, but to generate new knowledge (Winsberg 2010). These converging technologies are often referred with the acronym NBIC: nanotechnology, biotechnology, information technology, and cognitive science.

Such large initiatives are usually coined as "Big Science" projects that include—to our understanding, as we outline in the following—two main features: first, a challenge to organize, coordinate, and manage such a large number of researchers and research data in a way that enables accountability toward the funding organizations; and, second, an important role for advanced ICT serving as a structuring principle for the research carried out. In two words, Big Science is about structure and technology. Our aim is to outline potential consequences with ethical significances of these two features of Big Science in the field of neuroscience.

In doing so, we focus on the Human Brain Project for two reasons. The first is because the HBP has already faced opposition in the scientific community while setting up this initiative (Waldrop 2012), which culminated in an "open letter" sent in July 2014 to the European Commission signed by more than 800 scientists (http://www.neurofuture.eu). These scientists harshly criticized governance within the HBP and its technological focus—but they also called for "redirect[ing] the HBP funding to smaller investigator-driven neuroscience grants," "strongly support[ing] the mechanism of individual investigator-driven grants," which points to a general skepticism about Big Science approaches used in neuroscience. Nobel laureate Edvard Moser recently emphasized that "The brain is too complex, and neuroscience is too young, for all funding to be put into a single-aimed project" (in Requarth 2015). According to this opinion, neuroscience by its theory is not sufficiently developed for "Big Science"; rather, one should concentrate research efforts on relatively small grants to diverse laboratories around the world (Requarth 2015). This criticism is in line with studies suggesting that impact per dollar remains constant or decreases with grant size, which favors funding strategies that target diversity (Fortin and Currie 2013). An effort at mediation was undertaken before arbitration (Marquardt 2015), before the HBP Board of Directors changed the HBP scientific program and reframed its governance structure, which included removing the HBP Executive Committee and requiring representation on the board of each of the subprojects.

Second, we focus on the HBP, indeed, because of our roles. We are regular members of the Ethics Advisory Board (EAB) of the HBP,[1] except for one of us. We regular EAB members are independent advisors to HBP management on ethics issues and have no duties with the HBP to perform scientific or technical work outside of our ethics work. No regular EAB member receives funding from the HBP or is formally involved in research activities of the Science and Society program of the HBP. The one exception is an ex officio member, paid for work in the HBP to establish and coordinate the ethical committee(s) toward accomplishing the committees' mission. Thus, this contribution is well within the prescribed roles and responsibility of the EAB to independently and critically comment, from our unique perspective, on the developing HBP. Our views expressed here we own, and are not intended to reflect either the view of the committee as a whole, or the views of the leaders of the HBP.

Furthermore, we clarify our understanding of ethics, as some readers may question whether issues related to the organization of scientific work or public policy really fall into the domain of ethics. We see two reasons why the ethical analysis should go beyond "classical" ethical topics that already are discussed in the HBP context, such as the domain specific potential misuse of findings, or problems such as whether brain simulation might once create entities with moral status (Choudhury et al. 2014; Lim 2014; Rose 2014). First, research carried out in Big Science contexts can accentuate well-known ethical problems like data and privacy protection. Second, such contexts affect how neuroscience is understood, pursued, and also perceived in the public. This is ethically relevant. For, with respect to the ethical climate of internal collaboration (Martin et al. 2014), or where misallocation of public funding may possibly exist, policy decisions must involve moral judgments. Also, the European Commission's first HBP ethical, as well as technical, reviews, in January 2015, substantially addressed ethical issues raised by organizational and governance issues of the HBP (Technical Review Report 2015).

This contribution is structured as follows. In the second section, we investigate whether Big Neuroscience is actually needed as claimed to promote progress in this field. In the third and fourth sections, we analyze in detail the ethical issues related to structural aspects of Big Neuroscience and issues related to the role of ICT (in particular, simulations) to integrate project activities with knowledge. Finally, we offer recommendations on how the ethical challenges of Big Neuroscience might be addressed and which competences are needed within the neuroethics community that would support implementation of such infrastructure.

THE JUSTIFICATION OF BIG NEUROSCIENCE

Large, structured and formalized collaborations between scientists are a rather recent phenomenon in the history of science. This is related partly to the increasing importance of science for the military (Mendelsohn et al. 1988), as exemplified by the Manhattan Project in World War II, but also to the nature of certain research questions (e.g., in particle physics) that require large research infrastructures, as exemplified by the CERN facilities of the European

1. Originally, two independent ethics bodies advised the HBP Board of Directors: the Ethics, Legal and Social Aspects Committee (ELSA) and the Research Ethics Committee (REC). In September 2015, both committees merged into a single ethics advisory board in conjunction with structural reorganization of the HBP. All co-authors are currently EAB members and are unpaid except for one (KG), who is an ex officio member.

Organization for Nuclear Research in Geneva, Switzerland. The basis of these formal collaborations is that technology, broadly conceived, exists and the consequent collaborative activities can be characterized according to how different types of data-generating technologies are combined with levels of formalization in the modes of organizing access to those data (Shrum et al. 2007). "Big Science" is defined as a substantial expansion of scientific collaboration along several axes (Galison 1992): geographic (in the concentration of scientific expertise and technological capacities within cities or regions), economic (in the monetary sponsorship of major research endeavors, on the order of several billion dollars), multidisciplinary (in the necessary coordination of teams from previously distinct fields), and multinational (in the coordination of groups with very different research styles and cultures). Furthermore, more recent initiatives, like CERN's Large Hadron Collider, the Sloan Digital Sky Survey, or the NASA Climate Simulation Center, are strongly related to information technology infrastructure development allowing these programs to create and/or to handle large amounts of data (Markram et al. 2014).

Large-scale collaboration and a focus on technological infrastructure are also characteristic of Big Neuroscience initiatives (Grillner 2014), and Big Data efforts purportedly have become modus operandi in neuroscience, replacing smaller scale, hypothesis-driven science (*Nature Neuroscience* 2014). A major motivation is that the current research activities in neuroscience generate huge amounts of data and scientific papers on all levels of neuronal organization[2]: on gene expression in neurons, neuronal connectivity, brain activity patterns captured with neuroimaging, and human and animal behavior, just to name a few. And all these data are collected by distinct research traditions, structural biology at one end of the spectrum (ion channel identification) and behavioral research at the other end. However, neuroscientists face numerous problems when trying to integrate diverse data sets into a coherent understanding of brain functioning, which will require a cultural shift in not only sharing data across laboratories, but also making theorists central players in its analysis and interpretation (Sejnowski et al. 2014). One example concerns the reuse of so-called "long-tail" data, small data sets whose reuse is often stymied by a lack of community data-sharing standards (Ferguson et al. 2014)—and creating these standards is one aim of Big Science approaches. Representing huge data sets related to brain health and disease as descriptive clinical data, laboratory results or brain images are other examples that assist diagnosis and medical decision making. A large domain of ICT infrastructure must be set up and numerous technological and ethical problems will have to be solved (Christen et al. 2016)—an endeavor that is tackled by the Human Brain Project.

In summary, models, simulations, and large-scale ICT infrastructures acquire new functions within neuroscience as instruments to integrate systemic biological knowledge, gained on all levels of neuronal organization. Indeed, it is hard to imagine how neuroscience would be able to address the most difficult questions of its field—such as understanding how changes on all levels of neuronal organization affect behavior or how to address neurodegenerative diseases—without an integrative perspective that focuses not only on the processes in each level or organization, but also on the interplay of levels—an insight that in social neuroscience (and other fields) has led to the notion of "multi-level analysis" (Cacioppo and Decety 2011). Therefore, models and simulations are not only used as methods to observe specific brain processes, but also as a strategic tool to (re)organize knowledge. An example of success using this approach was recently described, namely, semiautomated text mining of a considerable amount of neuronal biophysical data from the vast electrophysiological research literature. Tripathy and colleagues (2015) found that experimental conditions (e.g., electrode types, recording temperatures, or animal age) can explain to a substantial degree the biophysical variability observed within a neuron type; that is, electrophysiological data are far more reproducible across labs than previously appreciated. Furthermore, a novel class of cortical and olfactory bulb interneurons that exhibit persistent activity at theta band frequencies has been identified using this approach.

From the onset, the HBP has been considered a project that aims to tackle the hard problem of integrating data with knowledge. As the former HBP director Henry Markram has pointed out: "We are not building a model; we are building a data integration strategy to render biologically realistic models" (personal communication, February 21, 2013; Christen 2013). To pursue this data integration goal, the HBP is developing six ICT-based platforms dedicated, respectively, as Neuroinformatics, Brain Simulation, High Performance Computing, Medical Informatics, Neuromorphic Computing, and Neurorobotics (for more detailed information compare https://www.humanbrainproject.eu). In that sense, the initiative not only aims to unify knowledge on several levels, but also to guide empirical research (see fourth section). Because it is practically impossible to determine the connectivity pattern of neurons for individual human brains in the same way we can

2. Grillner (2014) estimates the number of neuroscientific publication to be on the order of 100,000 papers a year. To put this in context, the number of annual publications in the Web of Science Core Collection subject categories reveals the following numbers for 2014: chemistry (physical and multdisciplinary) ~140,000, material science ~115,000, electrical/electronic engineering ~105,000, applied physics ~75,000, biochemistry and molecular biology ~70,000, oncology ~70,000, multidisciplinary sciences ~65,000, surgery ~60,000, and neuroscience ~55,000—i.e. the fields that are ranked before neuroscience cover a much broader domain of application (remember that the subject categories only provide a rough estimate of the number of papers published in a discipline).

now sequence the genome of a person, neuroscientists will need tools that tell them what to look for in real brains. According to the promotors of the HBP, simulations and brain atlases hold the promise to be "integrators" of knowledge and "lenses" through which scientists look to tackle the complexity of the brain.

Although there are several caveats, as last year's controversy around the HBP has demonstrated (Frégnac and Laurent 2014) and as we will outline in the fourth section, we consider this scientific strategy appropriate to narrow the gap between brain data on the one hand, and brain theories and real-world applications of neuroscience on the other hand. We think that the fundamental critique toward the HBP—namely, that the sheer complexity of the human brain involving the interplay of various organizational levels makes it impossible to "simulate the brain" (see first section)—fails to appreciate the main intended purpose of the HBP brain simulations. As explained earlier, they are not simply a tool to investigate a specific brain process or function. Rather, they serve as mediators to structure existing data such that they can be used for simulations, thereby shaping the processes on how new data is generated. Thus, the uncertainty about whether the brain as a whole can or cannot be "simulated" is not the main issue, nor is it a question concerning the likelihood of success; rather, apart from delivering insights into the mammalian brain in general, this strategy would generate new methodologies and techniques for handling large amounts of data, and would cause scientists to understand more deeply the feasibility of transforming "brain sciences" into Big Science (Dudai and Evers 2014).

Therefore, an ethical analysis should not focus solely on potential results of brain simulations as such, but also on the methods, because the social process of designing and implementing simulations and the corresponding reorganization of knowledge have normative implications even if certain results will never be obtained. "Simulating the brain" is not only a matter of technical know-how or practices involving creating simulation code, building interfaces, and using the models; it also involves complex processes to organize empirical and theoretical knowledge that aim to inform the models. The HBP promoters themselves indirectly point to this when writing: "We propose that the HBP use these techniques to generate a scaffold of *strategically selected data* [our emphasis] on the structure and functional organization of the human brain at different ages and at different levels of biological organization" (HBP Report 2012, 30). This indicates the necessity to choose among potentially conflicting data, which involves important but hidden normativity in model generation. We discuss this point in more detail in the fourth section when outlining the role of simulations as knowledge integrators.

In summary, the current Big Neuroscience initiatives, like the HBP, address a relevant goal, the achievement of which would be unlikely in a looser cooperative setting, but that requires structured collaboration across a large range and number of researchers and institutions. Like many other Big Science projects, the HBP initiative is not just a "container" in which research groups continue to work on their own agendas. This includes a technology-driven strategy to integrate data and to guide research—and doing Big Science also has significant effects on the way scientific collaboration itself is generated, structured, managed, and promoted. Taking this into account, to our opinion, a fundamental rejection of this approach misses the point. Rather, the focus of the analysis should be on how Big Neuroscience features influencing neuroscience are to be implemented and realized.

STRUCTURE: RISK–BENEFIT ANALYSIS OF STRUCTURAL FEATURES OF BIG NEUROSCIENCE

We analyze Big Neuroscience with respect to three basic structural features: first, the size of the pool of scientists who are involved in the collaborative structure ("Big Number"); second, data, that is, the amount and heterogeneity of the involved data ("Big Data"); and third, money, that is, the decision of a large funding body to invest a significant amount of research money into a single project ("Big Money"). We now discuss all three features in detail in order to inform a risk–benefit analysis regarding ethical aspects related to them. Table 1 provides a summary of this analysis.

Big Number

In general, Big Science means that a large number of scientists, often from many institutions across many countries, are involved in a rather formal mode of collaboration that creates a trade-off between organizational coordination and individual scientific freedom. In the "classic" bottom-up discovery ethos of many scientific fields, the interaction of research groups is a form of cooperation, where the interests of the various groups are more or less equally taken into account in decision making and everyday work. But in large research consortia, the interests of others may be overridden in the pursuit of the collective goal (Shrum et al. 2007). If a researcher or a research group within a large consortium has a better idea about what the consortium should be doing, the group may not be allowed to do it. Certainly, "small neuroscience" is also embedded into institutional structures that impose constraints—but not in the same way as in Big Neuroscience. This itself is not an ethical problem, but may actually be a normative issue, that is, an ethically justified requirement called for given the increased accountability due to the large amount of funding involved (discussed later). This comes with an ethical advantage due to the directive coordination needed in multiple projects, which can enable effective ethically motivated oversight, including protective measures against publication bias and efforts ensuring adequate systematic review prior to commencing research and detection of research fraud. Such advantage is more possible, at least, when such oversight structures are embedded as an integral part of the organization. An example is to place

Table 1. Risk–benefit analysis of structural features of Big Neuroscience.

	Ethical risks	Ethical benefits
Size: One big organizational unit with top-down defined patterns of collaboration instead of many small groups as independent units, organized mainly in a "bottom-up" way.	• Conflict with "bottom-up" work ethics of scientific cooperation • Suboptimal investment of research money in terms of productivity	• More effective ethical oversight to prevent publication bias and research fraud
Data: Large and heterogeneous data sets emerging from coordinated research actions instead of data mining and pattern recognition in existing data bases ("classical Big Data").	• Data security and privacy issues across countries • Poor understanding of role responsibility in managing and protecting data • Informed consent procedures across different cultures • Failure because of complexity and/or regulatory issues	• Larger database will yield more power and reliability • Data sharing, with clear understanding of each party's role, incorporated into science work as a virtue or even duty • Maximization of data collection to contribute to human welfare
Money: Connection to a large public funding body that commits itself to the project, creating a special kind of interdependence.	• Big promises that undermine research credibility • "Too big to fail" problem	• Anticipating and responding to calls for accountability regarding use of public money for research

ethics advisory functions in position to observe and react around major decisions concerning the infrastructure. Seen from that perspective, Big Neuroscience could seize the opportunity to strengthen principles of responsible research and innovation.

However, there is another important aspect related to size that could turn into an ethical risk. As the collaboration among scientists is, by necessity, more formalized, the size of the research groups will likely be larger compared to "small neuroscience"[3]—at least in the sense that the genuine groups (many of them existed already before the collaboration started) are more constrained in their work and are embedded in larger planning schemes after having entered the collaboration. This may affect group climate and research productivity. The latter, in particular, has been investigated empirically within scientometry and sociology of science. The majority of studies do not find a positive relation between research productivity and group size (Hemlin and Ollson 2013). Data rather indicate that research groups tend to be most productive if a certain threshold in size is not exceeded (Kenna and Berche 2011), a finding that is also supported by general theoretical considerations (Kao and Couzin 2014). The problem seems to be further aggravated when the disciplinary diversity is increased (Cummings et al. 2013), which is a characteristic of Big Neuroscience due to the involvement of information technology (for a further discussion of this point see the fourth section). "Scaling up" the number of interacting researchers could be a bad investment of research money. Big Neuroscience binds into one project a quite large part of available research money in neuroscience, and no existing data inform proponents of small neuroscience and those of Big Science which will, in the end, be the most efficient use of taxpayer money. The moral choice must be made. The HBP reply toward the "open letter" (available at https://www.humanbrainproject.eu/documents/10180/17646/HBP-Statement.090614.pdf) stating that HBP funding is taken from an ICT arm of the European Commission (EC) budget does not fully appreciate this point, as the funding scheme demands matching the commission's investment by third parties, which may have set different priorities for relative distributions to neuroscience versus ICT. Thus, disadvantaged researchers may be competing from outside the HBP for some of the same research money.

However, authors have argued that the usual way to measure research productivity—in terms of publications and citations (Kenna and Berche 2011; Fortin and Currie 2013)—is not adequate to measure the performance of Big Science, as its success will include additional factors such as infrastructure-built-up (like the CERN accelerator), the generation of more efficient methodologies (e.g., cheap genome sequencing in case of the Human Genome Project), or public education due to a more coherent presentation of the object under investigation (the brain in our case).

3. Although this is often claimed, for example, also in the "open letter" (http://www.neurofuture.eu), we found no recent published data on neuroscience research group mean sizes. A count performed in December 2014 in all 123 research groups of the Neuroscience Center Zurich that cover the whole disciplinary spectrum of neuroscience and on which data was available revealed a mean group size of 8.4 members (professors, postdocs, PhD students, technicians).

Big Data

It is often claimed that Big Neuroscience and Big Data go hand in hand. Big Data has become a large subject not only in biomedicine in general—as well as in many other fields like consumer research—but also in neuroscience (see the special issue of *Nature Neuroscience* in November 2014 dedicated to this subject). Roughly, Big Data can be defined as the availability of large amounts of heterogeneous data requiring sophisticated data-processing algorithms including data mining and pattern recognition. Although it is true that Big Science generally involves the generation and/or handling of large amounts of data, Big Data does not by itself imply Big Science; that is, it does not need a formal collaboration of a large number of scientists involving Big Money. Already smaller research projects might include Big Data (e.g., analyzing connectivity patterns in neuroimaging). Only when Big Data require new and highly expensive technologies (e.g., the development of new high-performance computing architectures, which is one of the goals of the HBP), they imply Big Science (see the next paragraph). More importantly, ethical issues related to Big Data in general are not specific to Big Neuroscience. What is, however, of particular relevance for Big Neuroscience is that data size is combined with data heterogeneity including sensitive personal data—and this involves not only practical but also ethical issues that go beyond traditional research; this is also because neuroscientific data may contain, and therefore leave vulnerable, much more personalizing information in a way that genetics or other Big Data fields do not.

For example, if the research aim is to integrate heterogeneous data across a large numbers of institutions and countries, the data flow, transfer, and identification would require new ways to secure data privacy, and to deidentify and tag data. Data are subject to differing jurisdictional regulations governing informed consent, or—in the case of laboratory animals—animal protection laws (e.g., compared to data emerging from the United States or China). Although the principle of subsidiarity offers guidance within the EU, integrating data from non-EU countries might be problematic—an issue that was also identified as being ethically problematic by the European Commission itself in information written on the current "Horizon 2020" funding scheme. Heterogeneous data, due to its interdisciplinary nature, will also require coordinating procedures across different disciplinary cultures. Animal data use may encounter transnational welfare governance conflicts due to different legislations. Human brain data are, by nature, sensitive: even if they do not contain health care information, because they contain information about the organ of the mind. These issues will multiply in relevance whenever the medical areas of neurology and psychiatry are concerned.

However, these potential pitfalls of Big Neuroscience, in particular with relation to health care, should carefully be weighed against the potential positive effects on providing sounder knowledge, reliable databases, and sufficient large samples with longitudinal data that are a necessary basis for evidence-based medicine. "Codes of conduct" must be developed to promote responsible research and innovation (RRI) and help create benefits for society (von Schomberg 2013). A more difficult dilemma would emerge if evidence were to be interpreted such that standards of data protection would have to be lowered in order that the Big Data approach applied to neuroscience would indeed create benefit.

Big Money

Big Neuroscience starts with the decision by a large funding body to invest ample research money into a single project. Clearly, whenever a big chunk of money is conspicuously (and correctly) visible to the public eye, its purpose should be publicly explained and its responsible use also accounted for publicly. But this may lead to unrealistic justifications on the part of researchers and unrealistic expectations in the public, in research agencies, and in politicians who have to justify their budget allocation in the political arena. Thus, justifying large investments of public money in research is more and more accompanied by "big promises" that may in the long run undermine the credibility of science. High-volume funding creates considerable pressure for success and tangible outcomes that are in line with the expectations and priorities of the sponsor. "Big promises" also affect the way project goals and achievements are communicated to the public. An example is the notion of "social bubbles," that is, human expectations being inflated beyond reason. A recent study showed that the Human Genome Project shared several characteristics of financial market bubbles (Gisler et al. 2010). If the picture is biased toward success, impact, and "return on investment," it is difficult to convey a realistic picture of progress as well as obstacles and problems to the public. This may in turn create unrealistic expectations by the public of what science can realistically achieve, leading to more pressure on the researchers, and increasing the risk of disappointment if the anticipated results cannot be delivered. Furthermore, justification and expectations are usually inflated by simplified, often exciting and enticing explanations in mass media. Excessive societal investment and expectations may have contributed to the fabrication of data by Woo Suk Hwang (Saunders and Savulescu 2008) in the early stages of therapeutic cloning. Thus, expectation management on side of the researchers should explicitly be addressed by leaders of Big Neuroscience projects to downmodulate the amplitude in the well-known cycle of hype and criticism (Caulfield et al. 2010).

Setting aside the expectation problem, Big Science, Big Money, and public attention often combine in ways that generate pressure to create strict governance and oversight structure. More is "at stake," and the pressure to deliver practical applications often increases, which may inevitably call for "narrowing" the scientific work program, leaving out relevant subfields and collaborators. This may counteract the ethical benefits mentioned earlier, namely, that Big Neuroscience could make it easier to implement ethical governance within the organization.

Related to this, the funding body is in a special way committed to the project. Large-scale funding in science involving hundreds of millions of dollars or Euros may become "not falsifiable" in a practical sense, as all involved stakeholders have incentives to make the project a success. This is the "too big to fail" issue. Certainly, however, even large projects can fail, as the recent cancelation of the so-called National Children's Study shows. This longitudinal study involving 100,000 children that consumed more than $1.2 billion within the last 14 years was stopped on December 12, 2014, despite this investment, the buildup of 40 recruitment centers, and the enrollment of already 5,700 children (Reardon 2014).

Each of these risks could be turned on its head. Greater integration could lead to better international oversight, ensuring coordinated and expert vetting of proposed research, requiring systematic literature review, an undertaking to make results publicly accessible (so preventing publications bias) (Savulescu et al. 1996), and greater oversight to detect fraud or misconduct internally. By establishing expert scientific and ethical review with full funding, such science could be conducted more efficiently and more ethically.

TECHNOLOGY—ETHICAL ISSUES OF THE SIMULATION APPROACH

A second characteristic of Big Science, beside structural features, is the decisive role of ICT as an instrument to structure and guide scientific activities. It's unthinkable that organizations like the CERN, which generates terabytes of data, or the Human Genome Project could ever have functioned without such technology. As large-scale simulations[4] have become important in many scientific fields, such as cosmology and climate research, supercomputer infrastructures have become an indispensable tool. This is certainly also the case in neuroscience (Gerstner et al. 2012), and a broad spectrum of large-scale models and simulations is found in neuroscience (De Garis et al. 2010). Examples include the "Blue Brain" (Markram 2006), the "SyNAPSE project" (Systems of Neuromorphic Adaptive Plastic Scalable Electronics; Ananthanarayanan et al. 2009), a large-scale model of the mammalian thalamocortical systems (Izhikevich and Edelman 2008), and the SPAUN model (Semantic Pointer Architecture Unified Network; Eliasmith et al. 2012).

The interconnection of simulation tools and simulation objects and the importance of the information metaphor in neuroscience raise various questions, some of which already are discussed in the literature (e.g., Bennett and Hacker 2003; Garson 2003; Falkenburg 2012). This includes epistemic issues like the meaning of "information" or the epistemic status of simulation results. These are indeed important and should be investigated further in the discourses about *in silico* experiments in research—this is, however, beyond the scope of this article (see for further reading, e.g., DeLanda 2010; Dudai and Evers 2014; Gramelsberger 2010).

In the following, we focus on the effects of using simulation technologies in neuroscience. We suggest that the modes of interactions of four basic units that characterize simulation-driven neuroscience—the real brain, the knowledge generated out of researching the brain, the brain simulations, and the public observer of neuroscience—are affected. Of concern are basic values that are also of ethical concern: trust, community, truth, and credibility. We discuss each of these points separately—partly by including experiences made in climate modeling, a field with profound experiences in setting up large-scale simulations in a highly interdisciplinary setting to demonstrate that the risks identified are not merely theoretical but could lead to consequences with ethical significance (most of the examples referring to climate modeling emerge from Lahsen [2005]).

Interrelation Between Real Brains and Brain Knowledge

The first point relates to trust. Within the Human Brain Project, simulations are intended to be used as predictive tools; that is, they should guide which experimental measures should be obtained in order to test theories. Predictive models, however, are an important extension of the classical hypothetic-deductive approach in science. Deducible from complex theory, explaining the experimental variables and measures may be difficult; indeed, neuronal dynamics may not any longer be understandable to the experimenter conducting the empirical tests. Simulation code represents the theory and it can be used to generate hypotheses (e.g., a specific distribution of synapses) that experimenters may then test. Simulations are intended to guide experiments in this way, which enlarges the usual scope of simulations in science—namely, to solve problems that are intractable to current analytical theory and to gain insight into physical phenomena where the accuracy and scope of experimental results are limited.

If simulations "guide" experiments, several aspects have to be considered. First, the better the laws are that govern the system, the more the predictive power of simulations increases. This is why simulations in solid-state physics are a common tool used, for example, to validate the integration of new chemical elements into computer chips (Pignedoli et al. 2007) and in this way informing the

4. To clarify the terminology used in this contribution: Models are abstractions of real-world structures and/or processes mostly in the form of mathematical equations or algorithms (although some models are physical, e.g., in hydrology). Simulations refer to the behavior of the model in time, whereas the equations or algorithms are usually implemented on a computer, requiring in most cases numerical approximations. Simulations may specify inputs, information handling mechanisms, or outputs in order to allow for prediction, retrodiction, explanation, or exploration. Due to the numerical nature of most simulation calculations, simulations can be understood as approximations of models.

actual chip design process. Neuroscience, however, is not physics in terms of the precision and comprehensiveness of available first principles (e.g., the laws of quantum physics). This may undermine trust regarding the guiding role of simulations. Second, experiments in neuroscience can be more difficult to conduct in terms of time and other resources (e.g., animals) compared to solid-state physics. This means that there may be a gap in the speed of how the simulations proceed and the time needed to actually perform the experiments the simulations propose. Maintaining insight and awareness of the mechanisms and basis for predictions therefore is essential to enable close and intense collaboration between modelers and those who use them.

These issues concern moral choices about the way to organize research. Setting aside solid-state physics, simulation has (to our knowledge) not yet been used in a way to guide experiments with complex systems like ecosystems or the climate system—so experience is lacking regarding the effects of a "reversed command structure" on the social environment within research groups. By "reversed command structure" we mean that simulation technologists suggest to empiricists what experiments could be performed, and not vice versa. From an ethical point of view, such a strategy benefits when it reduces the number of "unnecessary" animal experiments. Nevertheless, we would like to note that this specific way of using simulations as guides to the research process may be in some tension with the ethos of science—namely, to be open to the unexpected (i.e., experiments that are not suggested from simulations). Recall that a substantial part of initial scientific skepticism regarding the HBP refers to mistrust in the usefulness of simulations as "research guidance," given the vast complexity of the brain (Waldrop 2012) and a perceived lack of corrective loops between hypotheses and experimental tests within the HBP (Frégnac and Laurent 2014; see also Dudai and Evers 2014). Simulations in other biomedical fields (protein folding and molecular dynamical simulations) demonstrate the important role of coarse graining (i.e., finding physically sound simplifications) in order to handle computational complexity—but this coarse graining is based on a well-elaborated theoretical grounding of the underlying processes (Kamerlin et al. 2011). Critics question that neuroscience has achieved theoretical understanding to make decisions for what levels of detail should be used for specific simulations (Requarth 2015).

Interrelation Between Real Brains and Brain Simulations

The second point relates to community. The creation of brain simulation requires the integration of expertise that is not traditionally present in neuroscience—building a neuroscience/ICT modeling community. This goes along with the integration of different scientific cultures that see the object under investigation from different perspectives—which by itself might be a virtue. However, historically, it is interesting to observe that the first attempt to grasp neurobiological processes in terms of information theory and computation in the 1950s and 1960s failed to generate a productive research orientation in a similar way as in molecular biology. In the latter field, the notion of a "genetic code" became a powerful unifying metaphor for guiding research, whereas the search for a "neural code" provided a perplexity of approaches in neuroscience (Christen 2006). This may partly explain why there is still a large skepticism within neuroscience about the role of simulation as a method—in particular when the intention is to simulate the whole brain—as exemplified in the "open letter" and discussed in the HBP mediation report (Marquardt 2015). Therefore, the day-to-day collaboration among scientists and technologists representing quite different scientific cultures must be considered in design of the interactions.

Experiences in climate modeling demonstrate the pitfalls of interactions between different scientific cultures. First, climate simulations that intend to model a complex phenomenon are by themselves complex structures. Often, model developers build only parts of a model, integrating submodels and representational schemes (parameters) developed by other modeling groups. Even scientists (model users) who are not primarily model developers typically modify the models they have obtained from elsewhere. The difficulty of distinguishing properly between developers and users complicates a clear identification of the source of a model. Second, this increased specialization has reduced the amount of time model developers have to study the atmosphere using empirical data. This has contributed to an alienation of the empiricists from the real world whose role is checking models against empirical knowledge. Empiricists live in a culture that also involves humility about the accuracy of forecasts of atmospheric conditions, which is supported by the common experience of seeing synoptic and numerical weather forecasts turning out wrong. They complain that model developers often freeze others out and tend to be resistant to critical input, living in a "fortress mentality" (Lahsen 2005). Therefore, an ethical risk emerges on the level of the climate of research groups that are intended to interact in an open fashion, for example, due to failures to acknowledge contributions in large-scale collaborations or general distrust and noncooperation that harm the execution of the project goals.

Interaction Between Brain Simulation and Brain Knowledge

The third point relates to truth: Simulations as intended by the HBP assume a certain model of knowledge. That is, structured access to data and integration and interpretation of data across all levels will, due to the enormous number of publications in neuroscience,[5] depend on

5. The promoters of the Human Brain Project estimate that the "publication body" relevant for the project consists of at least 30 million papers (HBP 2012, 37).

Table 2. Risk–benefit analysis of technological features of Big Neuroscience that substantially relies on simulations

	Ethical risks	Ethical benefits
Trust: Simulations as "guiders" of empirical research.	• Interference with the ethos of free and open science (also due to pushing efficiency)	• More efficient empirical research that minimizes the use of resources (in particular, animals)
Community: Integration of an ICT-based "simulation culture" in neuroscience.	• "Fortress mentality" and alienation phenomena	• Increased incentive for interdisciplinary collaboration
Truth: Restructuring the knowledge base of neuroscience to make it more compatible for modeling.	• Ignorance of conflicts in data/knowledge • Missing peer-review culture	• Better awareness for conflicting findings due to systematic mining of available knowledge
Credibility: Using simulations as tools for communicating research results.	• Lacking standards regarding visualization blur the boundary between reality and simulation	• Novel ways to communicate complex phenomena increase public understanding of neuroscience

automated procedures of text mining and the like. This buildup of knowledge models, that is, structured access to data and publications referring to the phenomena one wants to model, is guided by normative decisions—what should be included in these knowledge libraries and what should not? This selection procedure differs from traditional strategies relying on peer review, as the latter allows for contradictory information and multiple disagreements (Rose 2014). But when creating models, at some point one has to choose the framework of the model. Models utilizing only one or very few mechanisms may be unproblematic—the model would be a tool to assess which mechanism leads to better predictions (e.g., which distribution of synaptic weights onto cortical neurons shape their spiking activity; Iyer et al. 2013). But more complex models that use techniques like parameterization make more decisions necessary. Thus, there may be an incentive to resolve conflicting data to establish a "coherent" knowledge base by ignoring these conflicts. This problem is aggravated when knowledge once considered "confirmed" is questioned again. When this knowledge has already been embedded in simulation code, much greater effort will be needed to resolve the conflict, tempting one to neglect this discrepancy. Those who manage the model would need to recognize that it needs to be modified, but they may be heavily invested in maintaining the model because of their effort in building the model and their reputation.

As an illustration, the functional role of columns in the cortex of higher mammals, according to the designers of the "Blue Brain" project—the forerunner of the HBP—was assumed to be "building blocks" in the sense that simulating a column in the rat barrel cortex would be a first step toward a whole cortex simulation. However, whether columns have any such functional role within neuroscience is controversial (Horton and Adams 2005). In that particular case, that would, perhaps, not be relevant for validating the model (Srikanth Ramaswamy, personal communication during the workshop "Future Neuroscience and the HBP," June 11–13, 2015, Brocher Foundation, Geneva, Switzerland), but it illustrates the problem of "strategic selection" of data. Conflicting knowledge becomes critical when creating models using data or generating hypotheses based on the model's assumptions. Dealing with this issue requires careful governance during the buildup of large-scale simulations. It also requires a commitment to levels of confidence rather than knowledge, and a preparedness and flexibility to use alternative models.

Again, experiences in climate modeling concerning the effect upon the structure of knowledge demonstrate that this risk is not merely theoretical. This results from the fact that simulation codes are tricky to develop and work with (Sundberg 2010). It often takes a long time to develop a code that does not crash during calculations and, when sufficiently stable, maintains sensitivity calibration, specifications, and parameters (see, e.g., Winsberg 2003). Therefore, the psychological and social investment in models, and the social worlds of which the modelers are a part, can reduce the critical distance one has from one's own creations. Although such personal and professional investments are not unique to the field of modeling, the problem is aggravated by the finding that, already in a time when codes were simpler, model codes were seldom subjected to peer review (Bankes 1993) and large-scale model studies are never replicated in their entirety by other scientists. Doing so would require reimplementing an identical conceptual model. Replication in science is generally difficult (Collins and Pinch 1993), and in the field of climate modeling the exact reproduction of a climate model outcome will never happen due to the internal model variability

Table 3. Recommendations

	Issue class	Recommendations
Structural features	Size	• Ensure the HBP has scientific and ethical expertise to enable systematic review, publication of all data, coordination between groups, and that it provides oversight for fraud and misconduct, etc., that is beyond the usual research ethics review processes. • Scan the patterns of collaboration for potential ethical issues that affect the climate in research groups. • Embed ethics contact persons (ethics rapporteur) within each of the research divisions. • Protect "whistleblowing" within the organization.
	Data	• Harmonization of procedures across workgroups that may have different political persuasions. • Cooperation agreement should include also standards for research ethics, including codes of conduct. • Do not demand matching of every detail regulated by law ("appropriate granularity of regulations").
	Money	• Enable project responsiveness to anticipate calls for accountability regarding use of public money for research. • Engage project internal public relations office to develop and maintain awareness of the need, and to design appropriate response, to avoid "overselling" results.
Technological features	Trust	• Ensure that the experimenters performing the verification/falsification tests meet core competency requirements regarding how simulations are generating predictions, including system strengths and weaknesses. • Allow some openness for explorative neuroscientific experiments outside of the predictive range of simulations.
	Community	• Define interfaces and modes of collaboration between modelers and empirical scientists to allow for knowledge and experience transfer and to avoid "fortress mentalities." • Ensure through education some understanding of the scientific culture of the counterparts.
	Truth	• Define procedures such that working with or adapting of simulation code is reproducible. • Communicate openly when models involve choices among conflicting data/theories. • Use programming strategies to avoid having empirical knowledge embedded in code that cannot be revised due to prohibitive investments when revising the code. • Determine protocols that allow for quality control of simulation code in a similar way as peer review of scientific contributions. • Support varieties of models/simulations that deal with similar problems. • Analyze the effect of model creation on structuring and selecting the data that provides the foundation of the models.
	Credibility	• Communicate openly that simulation visualizations involve design decisions that are not present in the real object (e.g., regarding color choices). Develop some standards of different types of visualization that emerge out of simulations (e.g., regarding the appearance of neurons) and ensure distinctness from other types of visualizations (e.g., brain imaging).

that results in chaotic dynamic perturbations. The sucomponents of the optimal climate models are closely scrutinized and compared in international peer-reviewed studies, among various models, aiming to find convergence in their findings. This kind of lack of reproducibility could seriously threaten Big Neuroscience—if spectacular predictions from climate models or, in our case, brain models don't come true, this may jeopardize funding for the whole field.

Interaction Between Brain Simulation/Knowledge and the Public

The final point concerns credibility. Simulations become communication tools; that is, they generate a new type of evidence (visualizations, movies) relevant to communication between scientists working on the various levels of brain organization, as well as for informing the public. It is well known that representational formats generate a certain authority and strength of persuasiveness, which grows out of its analytical power, its power to suggest and to communicate (Giere 1988)—a claim that has been empirically confirmed by Keehner et al. (2011) for different types of brain images. Visualizations guide scientific perception and interpretation of underlying data (Huber 2011). However, using simulations as communication tools requires conventions in information visualization of the output generated by the simulation. Whether these conventions are present in brain simulations can be questioned, given the lack of standards in neuroimaging (Christen et al. 2013). Thus, it will be crucial that as one goes from displaying a stream of numbers (referring, e.g., to structural positions of synapses, current flow within dendrites, and the like) to visualizing rich and dynamic data, for example, in animated graphics and movies, simulation results must be presented in a visual language that distinguishes them from empirical results.

A potential confusion of simulation with naturalistic data may be part of a more general phenomenon. Conflation of simulations with "observations," "samples," and "data" has been identified in studies of scientists in several fields of research (Dowling 1999). Simulation techniques may especially encourage such conflation, however. For example, Stefan Helmreich's (1998) ethnographic study of artificial life simulators revealed the powerful effect of simulations on the imagination of their creators and users.

In summary, although the use of technology within Big Neuroscience primarily relates to epistemological and sociological (i.e., on how the involved scientists interact) issues, it also can have relevant ethical effects—both positive and negative ones. Table 2 provides an overview in that respect.

CONCLUSION AND RECOMMENDATIONS

Here, we have outlined that Big Neuroscience, due to its structural and technological features, is associated with specific ethical challenges that should be a focus in neuroethical discourse. The intention of this article is to raise awareness regarding these issues that—on first sight—appear to be more of a methodological or organizational kind. Ethical reflection should not only concern the possible products of Big Neuroscience, but also the specific way these products are made. Certainly, the issues raised deserve a much broader analysis than our overview can provide. Nevertheless, we would like to point to some general recommendations that could help to address them. We structure them along the seven points we have outlined in the article (see Table 3).

The analogy to climate modeling shows that ethical risks associated with Big Science gain relevance when Big Science is intended to be used to support solving global problems. It is thus not surprising that some of the critique raised against the political implications of climate modeling relied on issues related to methodology or modes of interaction, for example, with respect to the "climate-gate" controversy of 2009 resulting from an illegal attack of the servers at the University of East Anglia in Britain.[6] There is admittedly much more to say on this debate—but the point here is merely that the practice of modeling involves various critical issues that have the potential to undermine the function of models and simulations, in particular when they obtain political relevance. Given the enormous burden brain-related diseases have, it is possible that also brain simulations will obtain such a political role, for example, with respect to guiding allocation of resources for research in neurodegenerative diseases. Possibly, brain simulations, in combination with personalized medical strategies, may at some time in the future guide therapeutic decisions for individual patients, making the ethical impact immediate. Projects like the virtual brain (www.virtualbrain.org) have as an explicit goal to customize models to individuals for precisely this purpose. It will thus be important that the neuroethics community increases its knowledge on issues related to information technology and brain modeling in order to be able to critically accompany these new developments. ■

REFERENCES

Ananthanarayanan, R., S. K. Esser, H. D. Simon, and D. S. Modha. 2009. The cat is out of the bag: Cortical simulations with 109 neurons, 1013 synapses. *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis 2009*, 1–12. Portland, OR: ACM.

Bankes, S. 1993. Exploratory modeling for policy analysis. *Operations Research* 41(3): 435–49.

6. Unknown hackers stole several thousand e-mails and other text files that then were considered to demonstrate that climate researchers manipulated data in favor or climate change and attempted to suppress critics. Later, committees investigated the allegations and published reports and found no evidence of fraud or scientific misconduct.

Bennett, M. R., and P. M. S. Hacker. 2003. *Philosophical foundations in neuroscience*. Malden, MA: Blackwell.

Cacioppo, J. T., and J. Decety. 2011. Social neuroscience: Challenges and opportunities in the study of complex behavior. *Annals of the N.Y. Academy of Science* 1224:162–73.

Caulfield, T., C. Rahul, A. Zarzeczny, and H. Walter. 2010. Mapping the coverage of neuroimaging research *SCRIPTed* 7(3): 421–28.

Choudhury, S., J. R. Fishman, M. L. McGowan, and E. T. Juengst. 2014. Big data, open science and the brain: Lessons learned from genomics. *Frontiers in Human Neuroscience* 16(8): 239.

Christen, M. 2006. *The role of spike patterns in neuronal information processing. A historically embedded conceptual clarification*. ETH-Diss No. 16464. Reprinted 2012, Südwestdeutschen Verlag für Hochschulschriften.

Christen, M. 2013. Gehirn-Simulationen—ein hindernisreicher Erkenntnisweg, 60. *Neue Zürcher Zeitung*, March 27.

Christen, M., J. Domingo-Ferrer, B. Draganski, T. Spranger, and H. Walter. 2016. On the compatibility of Big Data driven research and informed consent—The example of the Human Brain Project. In *Ethics of biomedical big data*, ed. L. Floridi and B. Mittelstadt. Springer.

Christen, M., D. A. Vitacco, L. Huber, et al. 2013. Colorful brains: 14 Years of display practice in functional neuroimaging. *Neuroimage* 73:30–39.

Collins, H., and T. Pinch. 1993. *The golem: What everyone should know about science*. Cambridge, UK: Cambridge University Press

Cummings, J. N., S. Kiesler, R. Bosagh Zadeh, and A. D. Balakrishnan. 2013. Group heterogeneity increases the risks of large group size: A longitudinal study of productivity in research groups. *Psychological Science* 24(6): 880–90.

De Garis, H., C. Shuo, B. Goertzel, and L. Ruiting. 2010. A world survey of artificial brain projects, Part I: Large-scale brain simulations. *Neurocomputing* 74:3–29.

DeLanda, M. 2010. *Philosophy and simulation. The emergence of synthetic reason*. London, UK: Continuum.

Dowling, D. 1999. Experimenting on theories. *Science in Context* 12 (2): 261–73.

Dudai, Y., and K. Evers. 2014. To simulate or not to simulate: What are the questions? *Neuron* 84:254–61.

Eliasmith, C., T. C. Stewart, X. Choo, et al. 2012. A large-scale model of the functioning brain. *Science* 338:1202–5.

Falkenburg, B. 2012. *Mythos Determinismus. Wieviel erklärt uns die Hirnforschung?* Heidelberg, Germany: Springer.

Ferguson, A. R., J. L. Nielson, M. H. Cragin, A. E. Bandrowski, and M. E. Martone. 2014. Big data from small data: Data-sharing in the 'long tail' of neuroscience. *Nature Neuroscience* 17(11): 1442–47.

Fortin, J.-M., and D. J. Currie. 2013. Big science vs. little science: How scientific impact scales with funding. *PLoS ONE* 8(6): e65263.

Frégnac, Y., and G. Laurent. 2014. Where is the brain in the Human Brain Project? *Nature* 513:27–29

Galison, P. 1992. The many facets of Big Science. In *Big science: The growth of large-scale research*, ed. P. Galison and B. W. Hevly, 1–17. Stanford, CA: Stanford University Press.

Garson, J. (2003). The introduction of information into neurobiology. *Philosophy of Science* 70:926–36.

Gerstner, W., H. Sprekeler, and G. Deco. 2012. Theory and simulation in neuroscience. *Science* 338:60–65.

Giere, R. 1988. *Explaining science: A cognitive approach*. Chicago, IL: University of Chicago Press.

Gisler, M., D. Sornette, and R. Woodard. 2010. Exuberant innovation: The human genome project. *ArXiv Physics and Society*. Available at: http://arxiv.org/abs/1003.2882 (accessed January 22, 2015).

Gramelsberger, G. 2010. *Computerexperimente. Zum Wandel der Wissenschaft im Zeitalter des Computers*. Bielefeld, Germany: Transcript Verlag.

Grillner, S. 2014. Megascience efforts and the brain. *Neuron* 82(6): 1209–11.

HBP Report. 2012. *The Human Brain Project. A report to the European Commission*. Available at: https://www.humanbrainproject.eu/documents/10180/17648/TheHBPReport_LR.pdf/18e5747e-10af-4bec-9806-d03aead57655 (accessed January 22, 2015).

Helmreich, S. 1998. *Silicon second nature: Culturing artificial life in a digital world*. Berkeley, CA: University of California Press

Hemlin, S., and L. Olsson. 2013. The psychology of research groups: Creativity and performance. In *Handbook of the psychology of science*, ed. G. J. Feist and M. E. Gorman, 393–415. New York, NY: Springer.

Horton, J. C., and D. L. Adams. 2005. The cortical column: A structure without a function. *Philosophical Transactions of the Royal Society B* 360:837–62.

Huber, L. 2011. Norming normality. On scientific fictions and canonical visualisations. *Medicine Studies* 3:41–52.

Iyer, R., V. Menon, M. Buice, C. Koch, and S. Mihalas. 2013. The influence of synaptic weight distribution on neuronal population dynamics. *PLoS Computational Biology* 9(10): e1003248.

Izhikevich, E. M., and G. M. Edelman. 2008. Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Science of the United States of America* 105(9): 3593–98.

Kamerlin, S. C., S. Vicatos, A. Dryga, and A. Warshel. (2011). Coarse-grained (multiscale) simulations in studies of biophysical and chemical systems. *Annual Reviews in Physical Chemistry* 62: 41–64.

Kao, A. B., and I. D. Couzin. 2014. Decision accuracy in complex environments is often maximized by small group sizes. *Proceedings of the Royal Society B* 281(1784): 20133305. doi:10.1098/rspb.2013.3305

Keehner, M., L. Mayberry, and M. H. Fischer. 2011. Different clues from different views: The role of image format in public perceptions of neuroimaging results. *Psychonomic Bulletin & Reviews* 18 (2): 422–28.

Kenna, R., and B. Berche. 2011. Critical mass and the dependency of research quality on group size. *Scientometrics* 86:527–40.

Lahsen, M. 2005. Seductive simulations? Uncertainty distribution around climate models. *Social Studies of Science* 35:895–922.

Lim, D. 2014. Brain simulation and personhood: A concern with the human brain project. *Ethics and Information Technology* 16:77–89.

Markoff, J. 2013. Obama seeking to boost study of human brain. *New York Times*, February 17. Available at: http://www.nytimes.com/2013/02/18/science/project-seeks-to-build-map-of-human-brain.html?pagewanted=all&_r=0 (accessed January 22, 2015).

Markram, H., R. Frackowiak, and K. Meier. 2014. Big Digital Science—A roadmap for the brain. Available at: http://ec.europa.eu/archives/commission_2010-2014/kroes/en/content/digital-minds-new-europe.html

Markram, H. 2006. The Blue Brain Project. *Nature Reviews: Neuroscience* 7:153–60.

Marquardt, W. 2015. Human Brain Project mediation report. Available at: http://www.fz-juelich.de/SharedDocs/Pressemitteilungen/UK/DE/2015/15-03-09hbp-mediation.html;jsessionid=BE4F5917ECDF2E5F8CC0ED6380219726 (accessed July 19, 2015).

Martin, S. R., J. J. Kish-Gephart, and J. R. Detert. 2014. Blind forces: Ethical infrastructures and moral disengagement in organizations. *Organizational Psychology Review* 4(4): 295–325.

Mendelsohn, E., M. R. Smith, and P. Weingart, eds. 1988. *Science, technology, and the military*. Dordrecht, The Netherlands: Kluver.

Nature Neuroscience. 2014. Editorial. Focus on big data. *Nature Neuroscience* 17(11): 1429.

Pignedoli, C. A., A. Curioni, and W. Andreoni. 2007. The anomalous behavior of the dielectric constant of hafnium silicates: A first principles study. *Physical Review Letters* 98(3): article 037602.

Reardon, S. 2014. NIH ends longitudinal children's study. *Nature News*, doi:10.1038/nature.2014.16556. Available at: http://www.nature.com/news/nih-ends-longitudinal-children-s-study-1.16556 (accessed February 4, 2015).

Requarth, T. 2015. The big problem with "big science" ventures—Like the Human Brain Project. Available at: http://nautil.us/blog/the-big-problem-with-big-science-ventureslike-the-human-brain-project (accessed September 9, 2015).

Rose, N. 2014. The Human Brain Project: Social and ethical challenges. *Neuron* 82(6): 1212–15.

Saunders, R., and J. Savulescu. 2008. Research ethics and lessons from Hwanggate: What can we learn from the Korean cloning fraud? *Journal of Medical Ethics* 34(3): 214–2 21.

Savulescu J., I. Chalmers, and J. Blunt. 1996. Are research ethics committees behaving unethically? some suggestions for improving performance and accountability. *British Medical Journal* 313(7069): 1390.

Sejnowski, T. J., P. S. Churchland, and J. A. Movshon. 2014. Putting big data to good use in neuroscience. *Nature Neuroscience* 17(11): 1440–41

Shrum, W., J. Genuth, and I. Chompalov. 2007. *Structures of scientific collaboration*. Cambridge, MA: MIT Press.

Sundberg, M. 2010. Organizing simulation code collectives. *Science Studies* 23(1): 37–57.

Technical Review Report of the HBP. 2015. Available at: https://www.humanbrainproject.eu/-/hbp-technical-review-report-now-available (accessed July 21, 2015).

Tripathy, S. J., S. D. Burton, M. Geramita, R. C. Gerkin, and N. N. Urban. 2015. Brain-wide analysis of electrophysiological diversity yields novel categorization of mammalian neuron types. *Journal of Neurophysiology* 113(10): 3474–3489. doi:10.1152/jn.00237.2015

Von Schomberg, R. 2013. A vision of responsible innovation. In *Responsible innovation*, ed. R. Owen, M. Heintz, and J. Bessant, 51–74. London, UK: John Wiley.

Waldrop, M. M. 2012. Computer modelling: Brain in a box. *Nature* 482:456–58.

Winsberg, E. 2010. *Science in the age of computer simulation*. Chicago, IL: University of Chicago Press.

Winsberg, E. 2003. Simulated experiments: Methodology for a virtual world. *Philosophy of Science* 70:105–25.