ADVANCES IN NEURAL ENGINEERING

BRAIN-COMPUTER INTERFACES

Edited by

AYMAN S. EL-BAZ

JASJIT S. SURI

ELSEVIER

ACADEMIC PRESS
An imprint of Elsevier

Academic Press is an imprint of Elsevier
125 London Wall, London EC2Y 5AS, United Kingdom
525 B Street, Suite 1650, San Diego, CA 92101, United States
50 Hampshire Street, 5th Floor, Cambridge, MA 02139, United States

Copyright © 2025 Elsevier Inc. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

Publisher's note: Elsevier takes a neutral position with respect to territorial disputes or jurisdictional claims in its published content, including in maps and institutional affiliations.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or any information storage and retrieval system, without permission in writing from the publisher. Details on how to seek permission, further information about the Publisher's permissions policies and our arrangements with organizations such as the Copyright Clearance Center and the Copyright Licensing Agency, can be found at our website: www.elsevier.com/permissions.

This book and the individual contributions contained in it are protected under copyright by the Publisher (other than as may be noted herein).

Notices

Knowledge and best practice in this field are constantly changing. As new research and experience broaden our understanding, changes in research methods, professional practices, or medical treatment may become necessary.

Practitioners and researchers must always rely on their own experience and knowledge in evaluating and using any information, methods, compounds, or experiments described herein. In using such information or methods they should be mindful of their own safety and the safety of others, including parties for whom they have a professional responsibility.

To the fullest extent of the law, neither the Publisher nor the authors, contributors, or editors, assume any liability for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions, or ideas contained in the material herein.

ISBN: 978-0-323-95439-6

For information on all Academic Press publications visit our website at
https://www.elsevier.com/books-and-journals

Publisher: Mara Conner
Acquisitions Editor: Chris Katsaropoulos
Editorial Project Manager: Emily Thomson
Production Project Manager: Sujithkumar Chandran
Cover Designer: Vicky Pearson Esser

Typeset by TNQ Technologies

ELSEVIER
Book Aid International
Working together
to grow libraries in
developing countries
www.elsevier.com • www.bookaid.org

CHAPTER

20

Assessing systemic benefit and risk in the development of BCI neurotechnology

Prashant Desai[1], *John R. Shook*[2] *and James Giordano*[3]

[1]**Neuroethics Studies Program, Pellegrino Center for Clinical Bioethics, Georgetown University Medical Center, Washington, DC, United States;** [2]**Philosophy, Politics, and Economics Program, Bowie State University, Bowie, MD, United States;** [3]**Department of Neurology and Neuroethics Studies Program, Pellegrino Center for Clinical Bioethics, Georgetown University Medical Center, Washington, DC, United States**

1. BCI innovation and ethics on a global stage

The rapid pace and expanding scope of neurosciences and technologies (neuroS/T) in general, and BCIs in particular, have prompted both concerns about and calls for increased oversight and guidance to promote responsible innovation and use of such technologies in national security and defense contexts, as well as diffused and perhaps direct effects in civilian spheres. Many of the methods of implementation research and science can be employed to effectively and efficiently address and bridge expanding gaps between technological capability and ethical approvability. While a novel neuroS/T such as BCIs proceeds through design and developmental phases, anticipating issues and addressing problems can supply a vital feedback loop. In such development and implementation, ethics should be a fundamental process rather than an unwelcome pest.

Minimum standards for ethical use of current and emerging neurotechnology should not be dismissed as weak or ineffective. Rather, perceived impotence is due to the short-sighted way that ethics is often applied retroactively, after new markets have widely impacted consumers across many ecologies and jurisdictions. Proactively, we believe that good ethics is a smart preventative measure for the scientific and industrial community, minimizing the risk of reactive overregulation; and we also point out how ethical practices are always more trustworthy from the user/consumer perspective. Multinational investments and efforts in

Brain-Computer Interfaces
https://doi.org/10.1016/B978-0-323-95439-6.00018-1

neuroS/T are establishing a global neurobioeconomy and are influencing the foci, scope, and tenor of both neurotechnological research and its applications in medicine, public life, and military operations.

Given this international momentum, it becomes evermore important—if not essential—to acknowledge differing socio-cultural needs, values, philosophies, and ethics when engaging discourse and deliberations about right and good engagement of neurotechnology in various contexts on the 21st-century world stage [1]. Thus, ethics must be focal to the enterprise, and this dictates frameworks and implementation methods to guide, support, and govern and such particular uses-in-practice.

2. BCI beyond the brain

Medical and clinical settings easily come to mind when contemplating BCI neuroS/T and its fruitful applications to improve the human condition. Of distinct note are defense industry stakeholders who have emerged as active investors in neuromodulation technology development. Neurotechnology research and development, along with particular implementations have been undertaken by the United States Army, Air Force, Navy, and Marine Corps units, with some of the most ardent translational research efforts conducted by the Defense Advanced Research Projects Agency (DARPA). Key DARPA programs in this space are Restoring Active Memory (RAM), with aims of producing closed loop neuromodulatory systems that could mitigate cognitive impairment following brain injury or illness; the Systems-Based Neurotechnology for Emerging Technologies (SUBNETS) program, established to develop closed-loop neuromodulation systems to treat neuropsychiatric conditions; and most currently (and perhaps far-reaching), Next Generation Noninvasive Neuromodulation (N3), which seeks to develop tools and methods assess and affect neural node and network activity more precisely, less invasively, and remotely. Of further note is the way that defense-related engagements of BCI are not limited to the militaries of the United States and its allies [2,3].

We should also draw attention to that reciprocity and complementarity between military and civilian effort and engagement in neuroS/T. A number of CRADAs (military-to-civilian cooperative commercialization agreements), as well as commercially engaged performers (i.e., civilian-to-military contracts), have been executed over the past five to 10 years in the U.S. alone. Academic discussions are ongoing about the multinationalization of neurotech R&D toward utilization, and the establishment of both collaborative opportunities (e.g., research and medical tourism) and competitive markets. Each and all of these sectors of innovation display "systemic characteristics and dynamics" that could be considered singularly within particular systems, as well as collectively among the interactive effects of systems-engaged-with-systems.

Vast implications for benefit, burden, and risk occurrence and analyses necessarily ensue across the whole range of actual and potential neuroS/T. Prior efforts and recommendations to best regulate neurotechnology have brought into stark relief both the importance of cooperation between the private, public, and defense sectors, and the necessity and value of developing realistic standards for research and application(s) in practice. Essential to any such enterprise is the assessment, mitigation, and/or prevention of risks in pursuit of worthy rewards.

We certainly urge the careful consideration of benefit, burden, and risk analyses in advance of further BCI design and development. To promote both impactful and sustainable technological advances and advantages, it becomes important to design implementation tactics and strategies that are explicitly directed at overcoming or side-stepping existing and probable barriers to use-in-practice in specified settings. Implementation science research provides methods to evaluate benefits, burdens, risks and threats within a system, and in these ways, contributes to flexibility in developing and instituting approaches that enable sustainable use of a neurotechnology in a particular real-world ecology, over time. Responsible and sustainable innovation must therefore be examined in its systemic as well as individualistic dimensions.

Individualism tends to prevail when ethical concerns receive attention. Discussions of BCI technology usually presume that a modulatory—and modular—technology is raising decomposable issues. As first-generation BCIs are attaining remarkable results, their design still demarcates the brain “over here” and the computer “over there”, only contingently connected and temporarily cooperating. An “interface” at most implies two facing informants, able to exchange messaging and influence each other. Nothing essential about either party gets transformed; how one or the other party processes information is unaffected during or after an interface. Furthermore, BCIs have been conceived as little more than prosthetic controllers with linear causal powers: the brain has thoughts, thoughts prompt computations, and systemic computing executes the thought in motor action (bodily or robotically). Distal, detachable, and directable BCI designs suggest a measure of safety and security, providing for reversibility (for the brain) and upgradability (for the computer), with a single cyber-neuro connector to safeguard. Traditional medical criteria continue to dominate the debates. Evaluations of BCI and its potential accordingly revolve around its proximity and invasiveness into the brain, its intrusiveness into subjective thought, and its prioritization for therapeutic goals rather than enhancement [4,5].

Lists of ethical issues for BCIs in experimental and implemental phases reflect that modular approach to BCI operationalization. The subject of BCI research deserves protections from idiosyncratic (individual-level) risks and hazards. As a mainly therapeutic and restorative technology, to benefit the particular patient under treatment, key priorities include obtaining informed consent, avoiding brain damage, minimizing unintended cognitive problems, monitoring personality and behavioral changes, and preserving privacy and confidentiality [6].

More philosophically intriguing concerns, phenomenologically and ethically, revolve around the effects of BCI on agency and autonomy. BCI functioning may introduce unwanted variability to plans and actions because it might misinterpret proceed hastily from a fleeting neural signal. Guidance from BCI computing may translate into a sensed loss of agency as much as enabled agency. Reliance on the BCI performance may enlarge one’s felt autonomy, although a perception of dependency can emerge as well, especially when the BCI is disconnected or removed entirely. Nothing about the brain or the person is entirely insulated from the workings of the BCI [7,8]. Integrable technologies are indeed integrated into the human users, and the sense of “self” extends to the sphere of effective action, with the BCI becomin a dimension of the human organism.

Such concerns are magnified with the development and use of closed-loop bio-neuro-feedback BCI. The interfusion of co-cognitive processing blurs a line between brain and computational system activity. Thus, matters are getting far more systematic than individualistic, as BCI designs permit closer integrations between minds and machines. Proactive modulations capable of entraining neuronal activity faster than ordinary self-conscious thinking will be needed to prevent the brain from slowing BCI performance. Future BCIs will be empowered by AI capable of anticipating, channeling, nudging, and effectively encouraging (as well as discouraging) some kinds of cognitive patterns and choices, much like it (currently) completes queries and fulfills prompts. Unless brains are tightly yoked with AI, human awareness will tend to be merely "on the loop" instead of "in the loop." To be surfe, the desired value is for mind-controlled BCI neurotechnology, not BCI mental distortion or mind control.

Such quasi-autonomous AI merged with BCI—which we here refer to as *brAIn* interfusion—would not only mold one's habitual conduct, but could (and likely would) also modify the social field around the individual. Knowing where the impact of BCI halts therefore becomes somedwhat pointless. Little about BCI implementation and modulatory influence could remain modular and discrete, much less isolatable for discerning the "individual" moral issues.

3. BCI as social and systemic

Societies will be pragmatically interested in synthetic (human-machine) agents, not merely "what is the human part doing?" or "what is the programming doing?" but rather "what is being done here?" Systemic evaluations will look to those systems employing neuroS/T—not humans in "isolation" apart from neurotechnology. Behaviorally, the whole person is still in sight, with the BCI functioning as an internal "organ". However, a behavioristic assignment of agency would not translate neatly into responsibility, either morally or legally. When an action looks good, the subject happily takes the credit; when an action looks bad, faulting the BCI suddenly seems fairer, despite getting held to inhumanly higher standards.

BCI engineering and programming will fall under increasingly intense technical and ethical scrutiny as mistakes and mishaps (repeatedly) occur. BCI-enabled hybrids will serve well enough as the subject of evaluations for capability, agency, responsibility, competency, culpability, and merit. Asking what the person did is an inapt question—personhood is entirely extensible just as mentality and capability is extendable. That extensibility of personhood does not halt upon encountering another person. Any participatory cooperation or structured practice that includes a BCI-enabled person eventually hybridizes the group, more or less tacitly, but nonetheless quite effectively. That hybridization therefore unavoidably extends into collective agency and group responsibility as well.

Law and policy lag behind neuroethical insights to the breadth of BCI-enacted agency and responsibility. The notable gap between clinical application and commercial availability further illustrates the entrenched dichotomy between idiosyncratic and systemic concerns. The remedy for idiosyncratic harms might be litigation naming particular parties; the remedy for systemic harms may be legislative or administrative, with vocal lobbying on behalf of at-risk groups. Ignoring prolonged harms to many people eventually heads toward class-action litigation, long after damage was done. Now that BCIs are arriving, what shall be the course of action taken about them? If any perceptible harms are revealed to be genuinely systemic rather than device-specific, no identifiable party could be more than incidentally responsible. Even as attendant risks (as well as rewards) of BCIs can already be envisioned, if not cautiously predicted, in advance, specifiable damages are still speculative at most. However, BCI innovators and device manufacturers that feel immune from liability should not lean toward complacency. A public provoked by real worries inspires precautionary approaches that generate short-term overregulation because industry would not assume long-term accountability.

Hence, regulatory considerations must be sensitive to evolving social conditions. The growing use of BCI and brAIn innovation(s) indeed present systemic benefits as well as burdens. On one hand, neurotechnology will incur a spectrum of (both idiosyncratic and systemic) effects. The adjustability of BCIs to enable a variety of actions and activities to be more precisely evoked does indeed contrast the less precise effects of pharmaceuticals. On the other hand, the relative potential for precision through BCIs makes them more powerful interventions with broader implications far beyond any individual utilizing a BCI. No individual could be reasonably expected to understand those implications or know where BCI utility reaches a limit. No matter how well a given BCI is adjusted for practical operability, the desired and ideal benefits will only become apparent after users discern how to achieve previously unimagined effects and results. What a BCI "does" will eventually be whatever its users determine what can actually be done with it. Past design will not define future performance.

Going further, as the intended and desired benefits of BCI-based neurotechnology are iteratively achieved, it is also likely that demand would drive continued and expanded use. In the event, idiosyncratic and systemically relevant caveats and warnings are doubly warranted. The broadening adoption of neurotechnologies should be regarded as an inherently opportune yet risky enterprise for both individuals and societies. As already noted, the performance of social groups will most surely be affected by the actions and influence of hybrid-BCI constituents. All participants will become "second-hand interfacers" through habituated interactions and engagements. Although benevolent motives and goals may prompt the development and initially intended use(s) of any neurotechnology, latent effects occurring in a broadening population of users will, as matter of fact, elicit diversity of effects that are both conspicuously apparent as well as more subtle, if not initially unrecognized.

Parallels are not hard to find across the history of novel technology. Writing was the first great leap in technical communication, extending cognition beyond the present moment and truly beyond the grave. A society with even 1% of its members utilizing writing was profoundly different from a society without any writing at all. For another example, let us simply ask which if any society proceeded entirely unaffected after a single printing press went into operation within its milieu? A semi-literate society promptly becomes a literary society (even as most people still aren't literate!) in most every institutional, economic, political, educational, and even religious sphere of life. Similarly, current and near-term future society will necessarily become a BCI-hybridized society around the time that about 5% of inhabitants are so equipped. Dare we deny that the first 5% of personal computer users, or the first 5% of smartphone adopters, were indeed what all the rest of us were to become, but became so just a little sooner? Human cognition is already computationally hybridized and enveloping the planetary semiosphere of signaling and communication.

4. Systemic implementations and their implications

For both neurotechnology as an engineering field and neuroethics as an academic discipline, it must be imperative to explore and address the wide implications of BCI use (and nonuse) at the aggregate level as well as the individual level. Deliberations should analyze and elucidate both idiosyncratic and systemic effects, in order to better develop balanced approaches to their consideration, occurrence, and consequences in real-world settings.

Longer-term consequences from alterations of neurocognitive function and neural phenotypes, along with modifications to their behavioral manifestations (and social influences) may remain unknown at first, only to manifest (and spread further) as a consequence of their diffusive effects over time. In addition, there is an appreciable risk that once a workable intervention for a particular problem has been devised, it might mask factors that regularly contribute to that problem. For example, where adolescent behavioral issues have causes from structural or local failures of the educational system, those deleterious conditions are less noticeable while a potent intervention masks those issues.

Effects incurred by the use of BCIs can be sufficiently potent to alter psycho-behavioral kinds and phenotypes. Also, the pace of diffusion of these effects is influenced—if not directly determined—by systemic (viz. social) pressures in and between the groups affected by technology adoption. Moreover, incentives and disincentives that are typically considered in the use or nonuse of BCIs from the idiosyncratic perspective are not necessarily accurate reflections of systemic impact. There are two reasons for this.

First, individual utility does not always mirror aggregate welfare. Examples of situations in which individual utility differs from systemic benefit abound in social scientific literature (the Prisoner’s Dilemma being prototypic). Herein, we define such effects, wherein adoption of a technology affects the individual one way, but affects the system in a diametrically opposite way as *strictly systemic effects*. Second, even in cases where individual utility affords some systemic benefit, certain selective pressures are not revealed in the short(er) timeframe in which idiosyncratically based decisions to adopt neurotechnologies are made. In this way, the adoption of a technology may appear to confer advantage(s) for the individual but can produce (idiosyncratic) risks that are unrecognized and therefore unaccounted in the decision-process. When these effects are compounded, they can incur multilevel and multiscalar systemic effects; we define these as *weakly systemic effects*.

Many of the individualizable risks inherent to the medical use of BCIs have been identified; these include those issues focal to the technology, such as hardware malfunction, surgical injury, explantation harm, and other complications that pose serious problems for patients and must be minimized. These risks can be considered to be *idiosyncratic* in that harm accrues to the individual, rather than society. To be sure, the realization of widespread idiosyncratic risks can impart considerable burden upon the systems in which individuals are constituent. And so, while it is crucial that idiosyncratic risks be identified and managed, *systemic* risks can be incurred even when idiosyncratic risks are managed.

Any technology available to a sufficient number of people can promptly generate a systemic social issue, without an arbitrary way to know in advance when that line will be crossed. Moreover, any sufficiently novel BCI will be thoroughly social from the start (in its import and effects) unless there is a priori assumption that changing one brain can only modify one brain. That presumption does not comport with the way that human mentality works, since all mentality *is* sociality. For BCI neuroS/T, individual enablement *is* social implementation.

Aggregating individual preferences and choices will not capture or fully explain the consequences of BCI adoption and integration into society. For a start, individual and social incentives are not necessarily aligned in all contexts. In *theory*, as neoclassical economics avers, rational individuals are axiomatically modeled as expected utility maximizers. All the while, actual individuals, in *practice*, are pursuing their own values and lives as constituent participants of society. All else staying equal (which never really happens), increasing individual utility should strongly correlate with increasing social (i.e., systemic) welfare. However, to reiterate, all else is rarely if ever equal: individual utility easily arouses and establishes unpredictable emergent effects that may *decrease* systemic welfare or at least diminish net social utility.

The converse case, where restrictions upon individual behavior (such as theft) can increase social/systemic welfare, makes good sense. And, thus far, discussions of ethical concerns relating to neuroS/T have focused on discerning and regulating sources of individual disutility: concerns about autonomy, changes in personal identity, unwanted side effects, and the like, all matter because they pose identifiable risks to persons taken individually. Those risks cannot be disregarded, especially by any neuroethical discourse. Nevertheless, as indicated above, the use of neurotechniques and neurotechnology has implications for intra- and inter*systemic* risk(s), as well.

Indubitably, there are—and will be—economic and socio-ecologic pressures that are likely to lead to increased adoption of BCIs if and as these approaches become accepted for use in developing physical and cognitive resilience, resistance to detrimental occupational effects, and/or performance-enhancement. Within a preventive medical framework, the use and (systemic) adoption of neurotechnology could be incentivized, especially in competitive settings, wherein organizations seek improvements for sustained efficiency, effectiveness, and productivity.

While the (at least perceived) risk of using BCI (or any neurotechnology) to "correct" certain conditions has been generally regarded to be somewhat less if such approaches are (1) only used in patients diagnosed with a condition, and (2) only used to treat the condition, such considerations beg the question of what condition(s) would be treatable. This is especially pertinent given that (3) neurotechnology can be—and often is—used to define and diagnose neurocognitive conditions (and establish diagnostic standards and norms); and (4) preventive medicine explicitly focuses upon fortifying health and preventing injury and/or disease, and therefore preventive interventions would need only define "potential" conditions that are to be mitigated or averted. This latter predicate re-establishes the propagating conditions for interest in, if not use of "preventive neurotechnologic interventions."

It is not obvious that potential macro-level benefits should be evaluated as normative imperatives: the extent to which concerns over systemic risk should constrain adoption of a particular neurotechnology is a question to be resolved through ongoing address, discourse and dialectic. We posit that important to any such discourse is cognizance of the ways in which adoption (or rejection) of a technology could have latently adverse effects. Given the complexity of the systems affected, it is admittedly difficult to precisely model these outcomes ex ante, and toward development of such effort, we have previously proposed a risk assessment and mitigation paradigm for operationalizing current and emerging neurotechnologies [9,10]. This paradigm of "Prepared Consequentialism" establishes groundwork responsibilities, questions, and framing parameters, which can be applied within a generalized framework for considering technology adoption, and an implementation protocol for articulation and ongoing review of systemic benefit, burden, and risk in practice.

5. Systematic ethics

At both idiosyncratic and systemic levels, a principled approach borrowed from bioethics asks key ethical questions about neuroS/T [11]. (1) Beneficence asks has the neurotechnology provided measurable and reliable effectiveness and value, in and across contexts of use, when compared to alternatives? (2) Nonmaleficence queries if the availability of a neurotechnology has been shown to be sufficiently safe, both in terms of consequences upon the recipients/users, as well as the local and nonlocal systems of their ecologies? (3) Autonomy inquires if and how the implementation of this neurotechnology can gain users' understanding, approval, and consent? (4) Justice evaluates to whom—and which groups—will the neurotechnology be available, and provided, and what are the contingencies for such provision?

Justice looms large, and rightly so, for any systemic approach to a neurotechnology of broad social impact and implications. The realities of what constitutes just distribution and use fosters consideration of whether more refined and focused principles can ensure that neuroS/T maintains legitimacy and merit within and across levels of application (i.e., among and between various systems).

To this latter point, multisystem (e.g., cross-cultural or international) considerations merit economic and ethical consideration. The increasing use of neurotechnology (of all sorts) will likely generate considerable reconfiguration and disruption of extant social, economic and power balances on the 21st-century world stage. These will inevitably include instabilities due to moral and justice inequalities and inequities, unless market and other hegemonic domain considerations and precautions are taken in advance.

Evaluation of investments, benefits, and risks should be inherent to a balanced process of both commercial cooperation and competition (i.e., "coopetition"), and achieving this balance should be a regulatory priority. Medical ethics, bioethics, and neuroethics can concur on certain precepts that could be important to, useful, and thus of value in informing and formulating such regulatory social policy. Specifically, we recommend that short-term evaluations of an emerging neuroS/T, as established by the aforementioned process of questions, context(s), and contingencies, should seek a fair balance among the following "6-P" priorities:

1. Protecting the autonomy and liberties of persons;
2. Promoting public health and general welfare;
3. Presenting economical ways to fairly distribute resources;
4. Preventing neglect of the vulnerable and disadvantaged;
5. Preserving the justice of the legal system;
6. Publicizing sound science for better public understanding [12].

Legal, political, economic, multicultural, and educational factors would not be left out of an adequate systematic analysis of neuroS/T innovation and adoption. Yet they all have ethical character, too. The future of neuroS/T will be *socialized* for systemic adoption and *commercialized* for market consumption, all without the least contradiction. That harmonization in theory lends all the more reason to view ethics as a fundamental factor in realizing and sustaining cooperative competition, and competitive cooperation, in real-world practice.

This sort of ethical integration is particularly important where emerging markets and new stake/shareholders are involved. Just as economic issues (such as asymmetries of commutative and distributive justice) continue to foster ethical issues, ethical hazards pose economic risks, even within long-established socio-economic sectors. New technologies and their markets display significant lability, liability and fragility, so their consumers-qua-recipients are more vulnerable. Society is never interested in building-in greater vulnerabilities, and markets take on risks only to suffer the consequences.

6. A framework for risk management

Building moral safeguards into the development and applications of neurotechnology is sound wisdom. Thus, establishing ethical standards for market networks and international trade and use is essential to smart policy. We urge the viewpoint that fiscal analysis should be ethically informed, and ethical analyses must address economics, especially given the widening niches of neuroS/T research, and its use(s) in a variety of practices—from the biomedical to the bellicose.

Here we invoke Owen Flanagan's construct of ethics as "human ecology" [13], which reinforces that broader contexts, namely, those ecological domains of human society in which relative goods, resources, and services and exchanged, must be acknowledged in any authentic approach to ethics. Aligned with Flanagan, we have posited, and reiterate here, the fundamental and dynamical relationship between a biopsychosocial approach to both economics [14,15] and ethics, which is entailed in and obtained by neuroethics [16]. Given that social reality, where brains are embodied in individuals living in groups that are embedded in a temporospatial environment, then philosopher Alasdair MacIntyre's query "Whose justice; which rationality?") [17] can be re-cast as: "Which ecology; whence good?"

Of course, defined "goods" differ, and increasingly may diverge, as the capabilities and use-propositions of neurotechnology expand in and across particular biopsychosocial domains and ecologies in which they are employed. The wisest recourse must openly assess the interactive nature and effect(s) of both idiosyncratic and systemic benefits, burdens and risks, and honestly ground evaluations in terms of "cui bono; *cui malo*?"—to whose benefit and whose harm? Resulting proposals in the form of regulations and principles need not appear simply as "top-down" dictates of dubious provenance, but rather as democratically accumulated wisdom garnered from hard trial and experience.

Guidance and policy should be open to perspectives, needs, and values of both innovators and their clienteles, but should also be informed by, and responsive to, systemic effects of the uses of neurotechnology on both local and nonlocal scales. Legal analyses abound; yet the regulatory realm is a poor place for first hearing of worries and grievances. Why should the economic sphere be deaf to articulable concerns, while it listens so carefully to the markets themselves? To be sure, not all benefits and risks are purely financial. Even if the economic sphere can only notice what gets represented as fiscal matters, we argue for a "minimax" calculation, akin to the difference principle of John Rawls [18] to translate ethical sense to market sense.

We opine that any framework of guidelines and policy for neurotechnology implementation must distinguish between (to borrow from the microeconomics community) "microprudential" policy, which relates to the management of idiosyncratic risk, and macroprudential policy, which limits the accumulation of systemic risk. In the neurotechnologic context, idiosyncratic risk is that incurred by an individual. In the United States, for example, oversight and development of such risk identification and reduction is currently regulated by the Food and Drug Administration's (FDA) Center for Devices and Radiological Control. There is, however, no equivalent body to ensure oversight and management of systemic (social and economic) risks fostered by and/or associated with neurotechnology.

Macroprudential regulation involves three key elements: a system to monitor risks, instruments to operate, and a strong institutional basis. The first of those elements, risk monitoring, should be comprised of a strong research program and active assessment of user populations. A coordinated research program is particularly important because of the speculative nature of risk prevention in the context of emerging neurotechnology; namely, it is not clear how extensive systemic effects might be, and in what domains they may arise. More precise identification of those domains and dimensions in which neurotechnology will most probabilistically incur and evoke impact (e.g., organizational-sector capability and productivity), will allow for more accurate modeling, monitoring, and evaluation.

Ensuring that whatever body is tasked with regulating neurotechnology has capacity for control of the appropriate policy instruments will be essential. Specifically, this involves oversight of device approval and distribution, while holding providers responsible for employing and maintaining the devices. At present, pending acquisition of further information about actual risks, it seems that the most viable form of oversight would be some type of nonbinding guidelines for use in particular (preventive and therapeutic) clinical contexts, and stricter regulations defining parameters of use in nonclinical contexts, given the relative ambiguity of clear medical need in/under such circumstances.

Ensuring that the regulating agency has clear mandate and appropriate input from other relevant government actors (such as the US National Institutes of Health, the Food and Drug Administration, and Departments of Labor and Commerce) will be crucial. Further, we note the importance of insulating the regulatory body from potential diversion or capture by partisan political influences. Instead, we advocate that any and all regulatory direction should reflect common standards of best interest in and across those systems in which the technology will be utilized and exert effect(s).

7. Implementing neurotechnology guidelines

As previously stated, the pace and extent of development(s) in neurotechnology are ever-increasing. The biomedical use of such advances unquestionably undergirds the importance of establishing evidenced-based practices, so as to safely and effectively guide various applications and improve outcomes. Yet, there is an extant lag of at least 15 years (if not longer) between research evidence and integration into practice, which can directly influence both resource allocation and the time to incur intended benefit(s) in user populations [19]. What's more, it is important to acknowledge what has been called the Collingridge dilemma: namely, that actual benefits, burdens, risks, and harms of a particular development cannot be fully known until after its widespread utilization across a variety of real-world settings. Despite best intentions, outcomes always vary, since the adoption and use of any emerging neurotechnology can, and likely will, evoke some burden, risk, and harm within and between systems over time.

A due appreciation for the complexities inherent at the macro-utilization scale does not mean that guidelines and principles must lose salience or force. Quite the opposite: strict general tenets are all the more reasonable and responsible for helping to parse and highlight ethical norms from the background of merely prudential values. That view of ethical guidelines (or protocols, regulations, principles, etc. by whatever label) moves away from simplistic ethical relativism, as relevant moral priorities should lose none of their force as they extend from idiosyncratic to systemic foci and dimensions.

Principled guidance is the intelligent response to systemic problems caused by interpersonal and intercultural diversity. Therefore, it remains important to appreciate—and remain responsive to—dynamic interactions within and between systems—inclusive of state-level systemic effects. To wit, it is important to note that neuroS/T is a multinational enterprise (and is becoming ever more so), and formalized regulations are not yet universally harmonized on the global stage, and may never be (the example of pharmaceuticals serves well here). It is easy to imagine and will be noteworthy to recognize the current realities of widespread regulatory arbitrage—developers and providers of neurotechnology exploiting variabilities across national (viz.- cultural) values, norms, and regulations, seeking competitive advantage in and across a variety of leverageable markets and settings.

Ethical principles are not deterred by a spectacle of prudential calculation and commercialization, as the global context around BCIs shifts and sways. We have herein described multinational similarities and distinctions in neuroS/T enterprises and defined key ethico-legal and social issues generated by these developments. To align with these realities, we propose a cosmopolitan approach—and methods—for neuroethics to address and potentially resolve such issues.

Inherent to this approach are four core guidelines for transnational appreciation and implementation, to address neuroS/T in general and especially innovations such as BCI [20].

Self-creativity: The right of persons to re-create themselves for enriching their lives.
Non-obsolescence: The duty to avoid the creation of obsolete or single-use people.
Empowerment: The right of persons to maintain their capabilities to live autonomous and fulfilling lives.
Citizenship: The duty to promote free, equal, law-abiding, and participatory citizenship.

These guidelines exemplify cross-cultural considerations and liberties, rather than ignoring them. What non-oppressive culture or country anywhere in the world wishes the contrary for its own people?

Regardless of the particular regulatory guidelines that may be eventually adopted by convention, regulation, and treaty, they should respond to all people invested in, and affected by, the social future of neuroS/T. Neuroethical recommendations should especially heed this counsel.

We argue that no impactful neurotechnology could be used in a "neutral manner" by generic consumers apart from their social surroundings and cultural traditions. As Latour [21] has emphasized, new scientific discoveries do not merely answer existing questions, but instead foster other emergent queries. In this spirit of inquiry, we [22] have raised issues and posed questions about enhancing and transformative neurotechnologies such as BCIs. How will people directly experience the application of neurotechnology during the ongoing course of their lives, given their own values and goals, and living within their own collectives and groups? How will the application of neurotechnology be evaluated within each group for responsible innovation, genuine need, and justice for all members? How will a group generally assess a neurotechnology's relationship to customary values, ideals, norms, and established laws? And, how will such evaluations and practices comport with (i.e., both affect, and be affected by) multi-systemic needs, values, economics, hegemonies, and ethics?

Answers to these—and evolving—questions will require considerable diligence, discourse, and dialectic within and between the systems (groups, communities, nations) involved. Opportunities for improving individual health along with social welfare are beckoning [23]. Their ongoing review may necessitate further revisitation, revision, and/or replacement of ongoing risk assessment methods: inclusive of regnant management, ethical, and implementation principles, protocols, and paradigms. Recognition, readiness, and responsivity to the changing realities of neurotechnology, and the systems that are affected by—and affect—its development and use should serve as cornerstones upon which to construct and iteratively remodel approaches to neuroS/T research and implementation across its dynamic, systemic ecology.

Acknowledgments

JG is supported by federal funds from Award UL1TR001409 from the National Center for Advancing Translational Sciences (NCATS), National Institutes of Health, through the Clinical and Translational Science Awards Program (CTSA), a trademark of the Department of Health and Human Services, part of the Roadmap Initiative, "Re-Engineering the Clinical Research Enterprise"; National Sciences Foundation Award 2113811 - Amendment ID 001; Henry Jackson Foundation for Military Medicine; the Institute for Biodefense Research; and Leadership Initiatives.

References

[1] Lanzilao E, Shook JR, Benedikter R, Giordano J. Advancing neuroscience on the 21st century world stage: the need for—and structure of—an internationally-relevant neuroethics. Ethics Bio Engineer & Med 2013;4(3):211—29.
[2] Kosal M, Putney J. Neurotechnology and international security: predicting commercial and military adoption of brain-computer interfaces (BCIs) in the United States and China. Politics Life Sci 2023;42(1):81—103.
[3] DeFranco J, Rhemann M, Giordano J. The emerging neurobioeconomy: implications for national security. Health Security 2020;18(4):267—77.
[4] Sattler S, Pietralla D. Public attitudes towards neurotechnology: findings from two experiments concerning brain stimulation devices (BSDs) and brain-computer interfaces (BCIs). PLoS One 2022;17(11):e0275454.
[5] Saha S, Khondaker A, Mamun K, Mostafa R, Naik G, Darvishi S, Khandoker A, Baumert M. Progress in brain computer interface: challenges and opportunities. Frontiers Systems Neurosci 2021;15:578875.
[6] King B, Read G, Salmon P. The risks associated with the use of brain-computer interfaces: a systematic review. Int J Human-Computer Interaction 2022:1—18. https://doi.org/10.1080/10447318.2022.2111041.
[7] Gregg M. Brain-computer interfaces: taking thoughts out of the human body. In: Wickramasinghe N, Chalasani S, Sloane E, editors. Digital disruption in healthcare. Cham, Switz: Springer; 2022. p. 17—26.
[8] Gilbert F, Ienca M, Cook M. How I became myself after merging with a computer: does human-machine symbiosis raise human rights issues. Brain Stim 2023;16(3):783—9.

[9] Giordano J. Toward an operational neuroethical risk analysis and mitigation paradigm for emerging neuroscience and technology (neuroS/T). Exp Neurol January 2017;287(4):492–5.

[10] Shook JR, Giordano G. Brains and bioethics. Cambridge, Mass: MIT Press; 2024.

[11] Beauchamp T, Childress F. Principles of biomedical ethics. 5th ed.. Oxford: Oxford University Press; 2001.

[12] Shook JR, Giordano J. Toward a new neuroethics in a multipolar and multicultural world. Global-E 2020;13(56):2–10.

[13] Flanagan O, Ancell A, Martin S, Steenbergen G. Empiricism and normative ethics: what do the biology and the psychology of morality have to do with ethics? Behaviour 2014;151(2–3):209–28.

[14] Cañadas A, Giordano J. A philosophically-based bio-psychosocial model of economics: evolutionary perspectives of human resource utilization and the need for an integrative, multi-disciplinary approach to economics. Int J Interdisciplinary Soc Sci 2010;5(8):53–68.

[15] McIntosh T, DuBois J, Perlmutter J. Ethical challenges in the commercialization of neurotechnology: contending with competing priorities. AJOB Neurosci 2022;13(1):60–2.

[16] Desai P, Shook JR, Giordano J. Addressing and managing systemic benefit, burden, and risk of emerging neurotechnology. AJOB Neurosci 2021;12(4):8–11.

[17] MacIntyre A. Whose justice? Which rationality. Notre Dame, Ind.: University of Notre Dame Press; 1988.

[18] Van Parijs P. In: Fleurbaey M, Gravel N, Laslier J, Trannoy A, editors. Justice as the fair distribution of freedom: fetishism or stoicism? In freedom in economics: new perspectives in normative analysis. London and New York: Routledge; 2002. p. 207–15.

[19] Morris Z, Wooding S, Grant J. The answer is 17 years, what is the question: understanding time lags in translational research. J Royal Society Med 2011;104(12):510–20.

[20] Shook JR, Giordano J. A principled and cosmopolitan neuroethics: considerations for international relevance. Phil Ethics Humanities in Med 2014;9(1):1–13. https://doi.org/10.1186/1747-5341-9-1.

[21] Latour B. Science in action: how to follow scientists and engineers through society. Cambridge, Mass: Harvard University Press; 1987.

[22] Shook JR, Giordano J. Defining contexts of neurocognitive (performance) enhancements. In: Jotterand F, Dubljevic V, editors. Cognitive enhancement: ethical and policy implications in international perspectives. Oxford: Oxford University Press; 2016. p. 76–98.

[23] Lillywhite B, Wolbring G. Having the ability to have a good life: what might be the impact of BCIs? In: Dubljević V, Coin A, editors. Policy, identity, and neurotechnology: the neuroethics of brain-computer interfaces. Cham, Switz: Springer; 2023. p. 117–50.