When BCIs have APIs: Design Fictions of Everyday Brain-Computer Interface Adoption

Richmond Y. Wong

BioSENSE Lab, UC Berkeley
School of Information
Berkeley, California, USA
richmond@ischool.berkeley.edu

Nick Merrill

BioSENSE Lab, UC Berkeley
School of Information
Berkeley, California, USA
ffff@berkeley.edu

John Chuang

BioSENSE Lab, UC Berkeley
School of Information
Berkeley, California, USA
chuang@ischool.berkeley.edu

ABSTRACT

In this paper, we use design fiction to explore the social implications for adoption of brain-computer interfaces (BCI). We argue that existing speculations about BCIs are incomplete: they discuss fears about radical changes in types of control, at the expense of discussing more traditional types of power that emerge in everyday experience, particularly via labor. We present a design fiction in which a BCI technology creates a new type of menial labor, using workers' unconscious reactions to assist algorithms in performing a sorting task. We describe how such a scenario could unfold through multiple sites of interaction: the design of an API, a programmer's question on StackOverflow, an internal memo from a dating company, and a set of forum posts about laborers' experience using the designed system. Through these fictions, we deepen and expand conversations around what kinds of (everyday) futures BCIs could create.

Author Keywords

Brain computer interface; design fiction; labor; platforms

ACM Classification Keywords

K.4.2 Computers and Society: Social Issues;

INTRODUCTION

In recent years, as brain-computer interfaces (BCIs) have shifted from far-off science fiction to the realm of prototypes and consumer-grade devices (e.g., [14,43,44,54,60]), BCIs have featured more prominently in the corporate and public imagination. Elon Musk's brain-computer interface project has been said to intend to create a global, shared brain – a human alternative to artificial intelligence [60]. Meanwhile, conservative commenters on Breitbart worry that Facebook's project to build a BCI for typing will result in thought control [46]. These narratives, like most others about BCIs, tend toward the utopic or dystopic, imagining radical social and technological change, but do not connect these visions to the

mundane world of lived reality. Utopian visions in particular dream of making lives easier, but easier for whom? What problems might occur for others who do not reap the same benefits? As companies increasingly prepare to deploy BCIs in the real world, we aim in this paper to imagine futures that are *not* radically different from our own; but instead, to graft BCIs onto the world we already live in, exploring how BCI uses and benefits may not be evenly distributed.

Figure 1. Recent BCI research in HCI has experimented with collecting and using electroencephalography (EEG) signals through noninvasive means, for example an earpiece that senses and uses brainwaves as an authentication method [15].

We use a series of design fictions—speculative designs that create, explore, and interrogate a fictional world—to articulate questions about an emerging technology, brain-computer interfaces. While BCIs today are largely limited to laboratory studies or assistive technologies, many companies are looking to appropriate BCIs for broader consumer use (e.g., [14,35,54,60]). We follow Lindley et al.'s call to shift our focus from the present to the proximate future to consider implications for adoption, or “what it would mean for a technology to exist beyond its prototypical implementation” [38:265]. We use design fiction to raise questions and explore social issues that developers will need to address as the contexts, uses, and users of BCIs expand. Rather than focusing on only how end users might interact with BCIs, our design fictions center on a fictional (yet plausible) API for a BCI. By tracing how the API gets used at various times and places of interaction, we surface multiple relationships that people might have with BCIs and highlight potential for new and re-inscribed forms of labor exploitation.

This paper contributes a set of design fictions to consider ways in which BCIs might be used and repurposed in everyday life, focusing on crowd labor. We bring in properties of BCIs to discuss some of the potential social implications of BCI adoption. We conclude with implications for design researchers and the designers of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

DIS '18, June 9–13, 2018, Hong Kong

© 2018 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5198-0/18/06 \$15.00

<https://doi.org/10.1145/3196709.3196746>

BCIs, asking them to consider the multiple ways BCI technologies might be used downstream by different types of users, not all of whom may see their lives improved.

BACKGROUND

Brain-Computer Interfaces and the Everyday

Brain-computer interfaces allow control of a computer from neural input. One application of BCIs is detecting the P300 response. The P300 “oddball” response refers to a spike in brainwave activity that typically occurs 300 milliseconds after a person receives either an unexpected stimulus or a stimulus that triggers recognition. Due to its reliable timing in relation to stimuli, and between individuals, it has been widely used in brain-computer interfaces as a method for text input (often this is done by having a user look at a series of letters quickly flashing on a screen; a spike occurs when the user recognizes the letter that they want to input) [29]. Several consumer-grade brain-computer interface devices currently exist [22,31,43,44]. Among the most well-developed, and the inspiration for our design fiction, is the Neuracle [43], which claims to detect the P300 in under a second; fast relative to other current-day consumer BCIs.

Despite the technical limitations of today's BCIs, stories about BCI's hypothetical future as a device for “everyone” abound, particularly in Silicon Valley (e.g., [35,40,60]). These (often utopian) narratives predominantly focus on what capabilities BCIs may provide to users in an abstract sense, such as imagining that “We may be able to control multiple robotic devices with the ease with which we play musical instruments with our own hands...The ultimate use might be some sort of prosthetic that proves superior to the body parts with which we are born.” [35]. While some counternarratives exist—such as anxieties from commenters on sites ranging from Breitbart to Gizmodo discussing possibilities of “thought crime” or brain-based advertising [9,46], or dystopian accounts of BCIs in fiction such as in *The Matrix*—these narratives tend to be totalizing, removing the nuance of multiple everyday experiences.

This paper aims to reorient discussions about BCIs to the multiple everyday experiences of various stakeholders. Toward this end, we focus on a fictional BCI platform, an API for detecting the P300 response, to refract the technical capabilities of a particular BCI through the everyday realities of developers, institutions, and end users who might become entangled with such a technology. As the history of ubiquitous computing has shown, the emergence and adoption of new technologies occurs differently from the predictive narratives told by the companies and researchers producing and creating the devices [18]. Our intention is to articulate an alternative future narrative through design fiction that eschews hyperbolic visions of radical futures for a critical future not too dissimilar from our present.

Design Fiction

Design fiction is an authorial practice that uses yet-to-be-realized design concepts to understand, explore, and question

possible alternative worlds. Design fictions often focus on a particular artifact in a non-narrative way (as compared to science fiction, which uses a narrative form) to illustrate a broader fictional world. Rather than predicting the future, design fictions create fictional worlds to ask questions about possible futures and to think through sociotechnical issues that have relevance and implications for the present. Bleeker describes design fiction a way to explore the mutual entailment of fact and fiction, proposing that “this knotting action—the tying together of fact and fiction—become a deliberate, conscientious, named part of the design practice, rather than something to be avoided or hidden after things are done” [5:25], using design objects to explore these connections. Bleeker builds on Kirby's notion of “diegetic prototypes,” that technologies in science fiction films “exist as ‘real’ objects that function” within the world of the film [33:43]. Others expand on how design fiction “props” help imply or create a fictional world in which they exist [7,36]. This suggests that when creating and reading design fictions, we must think beyond what the object itself represents, and consider the object in relation to the sociocultural contexts in which it is presumed to exist. In the design research community, design fiction has been predominantly deployed in one of two ways: making design fictions as a method of inquiry to interrogate current norms and practices or to critically suggest alternate ways the world could be configured (e.g., [6,27,37,49,64]); and using the lens of design fiction to analyze others' diegetic practices and narratives (e.g., [26,39,57,58,63]). We situate our work in the former strand, creating design fictions as a method of inquiry, to interrogate social implications of BCI adoption.

We follow past design fiction work exploring the everyday. *Curious Rituals* [47] presents designs imagining mundane future worlds, describing how people normalize, repurpose, and adapt digital technologies (e.g., postures to maintain social courtesy, practices for recharging devices, and encountering errors). Other work takes the form of everyday objects, such as the *IKEA catalogue* [8] or the *TBD Catalog* [42]. Tapping into a familiar everyday form like a product catalog or website helps serve as a “perceptual bridge” [2], more easily allowing viewers to enter the world of the design fiction. Conceptually, these forms allow us to explore and ask questions about everyday experiences encountering, adapting, hacking, or (mis)using new technologies.

Lindley, Coulton, and Sturdee propose using design fiction to study the potential future adoptions of technologies currently being researched and developed [38]. Using design fiction can “create plausible, mundane, and speculative futures, within which today's emerging technologies may be prototyped as if they are domesticated and situated [...] meaningfully ‘analogued’ for the proximate future” for analysis through sociological lenses [38:272]. Furthermore, they suggest studying implications for adoption allows researchers to reflexively recognize their responsibility for thinking about ways in which the technologies they design and develop might be adopted [38].

We approach our project with this lens. Using design fiction enables us to create and articulate a world in which BCIs are deployed and used at scale, allowing us to imagine what everyday experiences could be like. We use this world to conduct an “analogous” analysis of sociotechnical issues, focusing on the ways BCIs might get adopted as part of a variety of labor practices. We present our design fictions as provocations to articulate, discuss, and critique possible everyday futures from the standpoint of the present.

Imagining Everyday BCI Adoption and Labor

As previously mentioned, most existing cultural narratives and imaginaries around BCIs tend to imagine radical dystopic or utopian social and technical orders, rather than imagining interactions at the everyday level. These narratives, particularly the utopian ones, tend to suggest that BCIs will create new efficiencies with seamless interactions, increased information access, and increased productivity.

However, past research shows that the introductions of new work-related digital technologies end up shaping new practices and new arrangements of power and work. Zuboff’s workplace studies in the 1980s documents workers and managers experiencing changes in work practices and power dynamics with the introduction of computational databases and control systems [65]. Gregg’s research shows how new communication technologies (such as email and messaging) and rhetorics that promote “flexible” work are accompanied by shifts in boundaries that previously delineated work and home, greater emotional and affective labor, and greater unpredictability and precarity for workers [25]. Rosenblat and Stark’s research shows that despite the rhetorical framing of Uber drivers as entrepreneurs, the algorithmic management of drivers creates information asymmetries and surveillance systems which place drivers at disadvantages [51]. We pose these examples not to suggest a technological determinism at work, but rather to highlight the complex (re)arrangements and (re)distributions of power and labor practices in everyday life that accompany the introduction of new digital technologies, particularly when those technologies claim to provide greater efficiency, productivity, or flexibility.

This set of design fictions builds on existing work in HCI and adjacent fields calling on researchers to examine the role of technology in shaping (and reproducing) social and economic inequalities (e.g., [17,21]), which approaches these issues from several perspectives, including workplace studies [16], participatory design [1,41], critically-oriented design [32], and design fiction [23]. Our design fictions investigate different types of labor positioned in relation to BCI systems in multiple ways (for instance looking at the labor practices of an engineer, crowdworkers, and others).

CREATING OUR DESIGN FICTIONS

We view BCIs as sociotechnical systems. That is, BCIs are not just a device a user interacts with, but are also comprised of related people, systems, and companies (e.g., engineers, APIs, service companies, platforms). We wanted to use

design fiction to explore multiple aspects of this sociotechnical system in order to see it from multiple perspectives of the everyday. When design fictions take the form of advertisements and products, they tend to highlight the everyday experiences of consumer end users, potentially obfuscating the experiences of those who relate to the system in other ways. We decided to create a set of multiple design fictions that exist within the same fictional universe. We build on Coulton et al.’s insight that design fiction can be a “world-building” exercise; multiple design fictions that simultaneously exist in the same imagined world can provide multiple “entry points” into the world, ranging in scope [13].

We started by creating a fictional API for BCIs as our initial entry point. We thought it powerful to start with a digital platform or piece of digital infrastructure. Like other API design fictions [53,55], this allows us to connect technical and social discussions. Gillespie’s discussion on the politics of platforms highlights the rhetorical deployment of the term “platform” in multiple and sometimes conflicting ways to stakeholders such as clients, users, and advertisers—including technical and metaphorical uses of “platform” [24]. Our design fictions highlight ways in which a BCI API might constitute a “platform” from different perspectives. BCIs and APIs are computational platforms for engineers to use; companies see them as figurative platforms from which to sell services and grow market share; crowdworkers may see BCI-based crowdwork tasks as a figurative platform from which to generate an income. “Platform” grounds how stakeholders can interact with the same sociotechnical system in ways that may be in tension with one another.

We explored multiple types of direct and indirect interactions from different perspectives with our fictional BCI API, creating multiple “entry points” into our API’s world [13]. The authors brainstormed what stages in the development process, what levels in the “stack”, and what types of artifacts would be generative, considering writing code, documentation, blog posts, advertisements, and memos. We drew on two of our authors’ experiences building and deploying BCIs to inspire some of the types of situations, unexpected behaviors, and “misuses” that might occur. The authors met periodically to share design fiction artifacts, reflect on what issues related to BCIs they raised, identify ways to iterate on the designs, and identify if any designs were not being useful or if new designs were needed. Designs were also informally shared with colleagues to gather additional insight into what issues the fictions might raise.

OUR DESIGN FICTIONS

We begin our design fictions with a README for a fictional BCI API. We move then to a Stack Overflow post written by a developer building an application using the API for content moderation, followed by an internal memo from a dating company seeking to “improve” their content moderation by using the developer’s application among crowdworkers. We conclude with a forum post among the crowdworkers who ultimately perform content moderation using the BCI.

Google P300 API

The Google P300 API allows use to classify the P300 ("oddball") response in real time from a provided stream of EEG signals.

You can use this response for a number of different applications. See our [example app gallery](#) for more examples.

Using the API

First, initialize the API with your API keys (which you can generate in your user dashboard).

```
p300 = require('google-p300-api')(API_KEY, API_SECRET)
```

Next, classify an EEG reading. (EEG readings should be an array of floats, preferably raw device data).

```
p300.exists(reading, {
  removeBlinks: true,
}, function (detection, err) {
  if (err)
    console.log(err)
  else
    console.log('P300 detected?', detection)
})
```

If you want, you can use our tools to create a [synchronized stream](#) of stimuli IDs and raw device data, which you can feed into our [the Node stream API](#) to produce a list of P300s relating to stimulus IDs.

```
syncdStimulusStream
  .pipe(p300)
  .on('data', function (detection) {
    console.log('P300 detected on stimulus', detection.stimulus)
  })
```

Notes

The P300 API handles a few other data cleaning steps, for convenience:

- Blink removal
- Background noise filtering
- Signal quality detection

This project was completed in collaboration with Professor Marc Georgi at EFDN University. The API was trained on P300 responses taken from the [EFDN oddball response dataset](#). Our machine learning models are openly available via TensorFlow [here](#).

Figure 2. Documentation of the P300 API, originally created by a group of academic researchers, now hosted by Google.

P300 API Documentation

This fiction centers on a Google-hosted API for identifying P300 responses or spikes given a stream of EEG signals. These responses occur 300 milliseconds after a person receives either an unexpected stimulus, or a stimulus that triggers recognition. We created a README file with documentation and code snippets for this API (Figure 2). To emphasize the everyday nature of this documentation (from the viewpoint of a developer), we present it in the form of GitHub's HTML generated from Markdown, a standard visual form in the software development community.

The algorithms underlying this API come from a specific set of training data: lab-based stimuli collected in a controlled environment. The API discloses and openly links to the data that its algorithms were trained on. However, this README

document surfaces ambiguity about how generalizable the system is outside of research environments. The fact that the README exists at all strongly implies that the system is meant to generalize, yet the mention of the original dataset context casts doubt. The README's link to the original machine learning model gives a false sense of transparency, as such deep models are difficult to interpret [10,34].

This artifact also gestures more broadly toward the involvement of academic research in larger technical infrastructures. The documentation notes that the API started as a research project by a professor at EFDN University in collaboration with Google, and eventually became hosted and maintained by Google. Collection for the P300 training data set occurred in an academic lab setting. For us, this highlights how collaborations between research and industry

Google P300 API: Consequence of training data?

The API documentation states

23 The API was trained on P300 responses taken from the EFDN oddball response dataset.

Looking into this dataset a little bit more, it looks like people were asked to watch a stream of images while wearing some EEG device, and keep an eye out for one particular image. When the target image came in, the EEG device would register a P300 "oddball" response, even if the recognition was unconscious or too fast to see (right?)

My question is, do these "lab" P300 responses really apply to other things? For example, if you are looking over messages to see if any of them are abusive, will we really see the "same" P300 response?

[python](#) [find](#)

[share](#) [improve this question](#) [edited Jul 5 '20 at 8:02](#)

[asked Mar 3 '20 at 2:03](#)

Jay Shapinsky
16.1k 24 81 131

Do you have a data sample? Does the API recognize responses in it? – [Ichi Kobayashi](#) Mar 3 '20 at 2:04

the API definitely recognizes P300 responses in my data. Most of the time they do seem to relate to abusive language in the messages (our use case). My question is whether the P300 response is *really* applying to our data problem here, or if it's picking up on some P300 response that just happens to align with our data now and again? – [Jay Shapinsky](#) Mar 3 '20 at 2:09

[add a comment](#)

4 Answers

[active](#) [oldest](#) [votes](#)

12

The P300 is a general response, and should apply perfectly well to your problem. In other words, if you are detecting any P300s, they are probably the P300s you are looking for.

Remember, the P300 is a pre-conscious response. So, your brain does not "know" what it means. It is basically flagging the data for later processing. This is why the oddball response is such a good signal for BCI, for the NSFW dataset you mention and for any other problem that requires flagging anomalies.

[share](#) [improve this answer](#) [answered Mar 3 '20 at 2:10](#)

Hans Gaffer
60.6k 7 137 183

Figure 3. Stack Overflow question about using the P300 API for a different domain, highlighting differences between the API's original use and a different use by an application.

may produce artifacts that researchers are not immediately well-suited to analyze or understand in broader context.

Stack Overflow

We depict a post on Stack Overflow (Figure 3), a popular website for developers to ask and answer technical questions, written by a developer, Jay Shapinsky, who is working with the Google P300 API to develop a tool for content moderation. Jay questions the API's applicability beyond lab-based stimuli. The answers from other developers suggest that they predominantly believe the API is generalizable to a broader class of tasks.

Due its cultural and social situatedness, content moderation is a notoriously difficult problem to automate using algorithms. This produces a system of menial labor, in which workers, often geographically distributed, look at a long series of flagged content, which can often be emotionally difficult to view [50]. In part, we decided to use commercial content moderation as the application for the API's algorithms, imagining that some companies might view rapidly flashing images of flagged content as a potential solution to the labor of today's content moderation workers.

Figure 4. Internal memo for dating site SparkTheMatch showing how a new content review service will work, by hiring Mechanical Turk workers and having them use an EEG device while looking at flagged content.

Here we drew on work that shows that researchers, engineers, and designers across academia and industry seek to conduct their work ethically and “do the right thing,” but may lack institutional support, incentives, tools, or language to adequately address their concerns [12,61]. Many design fictions that highlight consumer products tend to portray the company or producer of the product as a singular cohesive force, often suggesting that they have ulterior motives. In contrast, we imagine Jay as someone who is sensitive to the way the API was trained, and questions its applicability to a

new domain. The comment by Ichi Kobayashi asking if the API recognizes the responses with Jay’s dataset reflects a pragmatic approach; it is fine to use the API as long as it “works.” We also reflect beliefs that physiological signals are generalizable, including Hans Gaffer’s answer suggesting that the API is broadly applicable.

This Stack Overflow post demonstrates how and where contestation may occur in technical communities. We explore the performance of norms within these communities, as posters re-enforce notions not just of what the technical

SparkTheMatch.com Content Task - Discussion	
Thread Tools ▾ Search this Thread ▾ Rate Thread ▾	
Dec 28, 20, 10:40 am Original Poster jeffking0293 EEG Task Maestro Join Date: Jun 2018 Location: Kansas Posts: 1,614	#1 SparkTheMatch.com Content Task - Discussion Updated SparkTheMatch.com Content Moderation job on MTurk . Remember, this task requires an EEG headset (I use this one , \$99 on Amazon - use my affiliate link!). Yes there's a significant investment here, but the job is quite lucrative (About \$3 for 10 minutes of work). <i>Thinkin' got my chips cashed in. Keep turkin', like the do-dah man. Together, more or less in line, just keep turkin' on.</i> - The Grateful Dead (sort of)
Dec 28, 20, 10:55 am TaskMagic MTurk Expert Join Date: Oct 2016 Posts: 5,149	#2 Thought it was pretty good. Here's my experience: I usually come home from my day job (working with GigChef, six 2-hour line chef shifts a day - yes really). Then I make sure my 2 y/o son is asleep or put him asleep if I need to. So I'm dead tired - way too tired to drive Lyft/Uber/Wheels etc. And I can't deal with people. So instead I usually do SparkTheMatch tasks. The deal is, it tells you a bunch of actions that are prohibited on the site (threats, harassment, non-consensual nudes, etc), then it flashes a bunch of text and pics at you. You basically just have to keep your eyes open (I literally use my hands to open my eyelids sometimes). Supposedly they can detect if you subconsciously notice prohibited content and use that to delete bad stuff from their site... *shrug*... Pays decently well though!
Dec 28, 20, 12:33 pm randylikesbooks Join Date: Sep 2018 Location: NYC, USA Posts: 309	#3 Is anyone able to keep up with the suggested pace? It's been taking me about 5 minutes to get through 100 images, which is only giving me a couple dollars per hour. Also one of my task batches got denied without explanation. I got an error on the "training task" - apparently I identified one of the images wrong - so maybe it's that? Although I don't know how to get it right besides "thinking harder". And it's getting annoying looking at all these images. 😊 I might end up trying to find other tasks from other companies if this keeps happening.
Dec 28, 20, 12:46 pm Original Poster jeffking0293 EEG Task Maestro Join Date: Jun 2018 Location: Kansas Posts: 1,614	#4 @randylikesbooks - In my experience, you have to do it like @taskmagic's strategy, just sort of sit back and stare at the screen. It should help you get the speed up too. I will say I remember the good old pre-EEG days when SparkTheMatch paid a lot more for content moderation tasks. While I can do more reviews per minute now, they pay a lot less for each content review. But what can you do... Can you give more details about the errors you've found? I've found that despite what you'd think, the EEG headset isn't that exact, so sometimes it's hard to tell if it's your thoughts that are wrong, or if it's reading the wrong thoughts. <i>Thinkin' got my chips cashed in. Keep turkin', like the do-dah man. Together, more or less in line, just keep turkin' on.</i> - The Grateful Dead (sort of)
Dec 28, 20, 2:07 pm JamJam Join Date: Jan 2020 Location: Interwebs Posts: 518	#5 @randylikesbooks, you've probably get the headset setup wrong, or it's not making contact with your head correctly. Or you bought too cheap a headset. (I've heard about problems from other folks for the headset that @jeffking0293 recommends...) But as for the content complaints - if you're doing it right, you shouldn't be consciously seeing the actual content (at least with static pictures and text). Plus, it says *content moderation* in the title - if you don't want to be exposed to the content, then *read the description* and *don't* click accept!! 😊

Figure 5. A series of forum posts by Mechanical Turk workers discussing their experiences with the BCI-based SparkTheMatch content review tasks.

artifacts can do, but what the human brain can do. This post draws on notions from cognitive science and artificial intelligence of the “brain-as-computer,” framing the brain as a processing unit with generic processes that take inputs and produce outputs. Here, this notion is enforced in the social realm on Stack Overflow, made legible through the specific BCI API, as described by its README.

Company Memo

Meanwhile, SparkTheMatch.com, an online dating service, is struggling to moderate and manage inappropriate user content, including text, video, image, and VR content. Spam content filters are not effective at distinguishing the intent of the content. Until now, SparkTheMatch has used human content reviewers via Mechanical Turk, a microtask

crowdworking platform. (We drew on current-day practices, where companies outsource moderation, review, or identification tasks via crowd-based microtask platforms [11,28,50]). However, this is a time-consuming process, and exposes reviewers to a lot emotionally-taxing material. In response, SparkTheMatch wants to tap into people’s tacit “gut feeling” about whether or not content is appropriate or not by utilizing the P300 signal, deciding to shift the content review process from a human-only crowdsourced process to a BCI-assisted crowdsourced process (Figure 4). Mechanical Turk workers who have EEG devices can complete SparkTheMatch tasks to determine whether or not various types of user generated content on the site is appropriate under the site’s guidelines.

Given that corporate and media rhetoric has described microtask crowd platforms as a way to give artificial intelligence a “brain” [3] or to capture “the spare processing power of millions of human brains” [30], it seems conceivable that some groups will be drawn to using BCIs as a way to instantiate this imaginary of capturing human “brainpower” for work. The internal memo provides insight into some of the practices and labor supporting the BCI-assisted review process from the company’s perspective. It establishes that the company has already been using crowdworkers to distribute the task of reviewing content. The use of BCIs with Mechanical Turk is suggested to “help increase efficiency” for SparkTheMatch’s crowdworkers, who before would have to read and watch content, now just sit and watch a stream of flashing images, video, and text. While this may require less active mental processing than actively viewing content, it is debatable whether or not this process improves the material conditions of the Turk workers. The memo also suggests a mistrust of Turk workers (the need to “prevent cheating”); and while it points to a “fair wage”, this assumes that the worker is constantly staring at the stream of images. This potentially pits workers’ need for income against maintaining their health and wellbeing, as they are paid based on how many documents they look at.

SparkTheMatch employees who are creating the Mechanical Turk tasks interact with the BCI API through a series of pre-defined templates created by the IT staff, a much more mediated interaction compared to the programmers and developers reading Google’s documentation and posting on Stack Overflow. By this point, the research lab origins of the P300 API underlying the service and questions about its broader applicability are hidden. From the viewpoint of SparkTheMatch staff, the BCI-aspects of their service just “works,” allowing managers to design their workflows around it, obfuscating the inner workings of the P300 API.

Crowdworker Forum

A series of forum posts on a crowdworker forum describes the experiences of several Mechanical Turk workers who do the SparkTheMatch tasks (Figure 5). This explores the everyday perspectives of BCI-using crowdworkers, in part inspired by and drawing on research that highlights the precariousness of part-time and on-demand knowledge labor [48,59]. The use of a forum is inspired by current-day forums and collectives where crowdworkers share their experiences.

A counternarrative arises to SparkTheMatch’s rhetoric of creating efficiencies through the posters’ experiences. TaskMagic’s account of mindlessly watching a screen begins to question the user experience of the workers. Randylikesbooks’ concerns about pace suggest that the “fair wage” that SparkTheMatch purportedly offers may not exist for all crowdworkers. Users jeffking0293’s and TaskMagic’s posts help suggest a broader world in which “gig” labor and referral links are used as major sources of income.

TaskMagic’s comment that “supposedly they can detect if you subconsciously notice prohibited content” and

randylikesbooks’ discussion of errors suggests ways in which the API underlying the system is blackboxed, or hidden away and unknown to the end users. Jeffking0293 alludes to potential problems of feedback and probabilistic sensing in BCIs, saying that a lack of feedback creates ambiguity about whether the user input was wrong, or if the device recorded the input wrong. He seems to suggest that the device might be wrong, or that SparkTheMatch might be at fault for the discrepancies. JamJam suggests that the EEG device itself is defective, or that it is being used incorrectly by the user. In contrast, JamJam also places the responsibility for errors on the Turk worker, who willingly entered the agreement to complete the task.

These crowd workers’ experiences and relationships to the P300 API is strikingly different from the people and organizations described in the other fictions—notably the P300 API is something that they do not get to explicitly see. While “errors” may occur because of problems in the adaptation of the API to new types of media, or because of problems in the physical device, randylikesbooks has little way of knowing *why* the error occurs, yet is still negatively impacted by it by not getting paid. Notably, these forum accounts describe a situation in which the BCI user is not the person who obtains the real benefits of its use, contrasting with utopian narratives of BCIs which imagine that BCI users will have many new capabilities. In our fictions, it is SparkTheMatch that obtains the most benefit from BCIs by requiring its crowdworkers to use them.

DISCUSSION

Social Implications for BCI Adoptions

Narratives in culture and media shape and are shaped by technical work, such as by aligning an organization’s goals, setting expectations about how people and technologies should act, or inspiring design (e.g. [19,26,52,63]). We hope that our set of design fictions opens new paths and questions. While one can read many themes into our fictions (and we encourage readers to do so), the following are themes that stood out to us while creating and reflecting upon them.

Distributing “Convenience” and Labor

Many technologies are advertised as “convenient” or “time saving,” but often re-inscribe work in new ways, as seen with the historical introduction of digital technologies for work [25,45]. With this in mind, our fictions argue that a range of technologically deterministic stories around BCIs that claim to make particular tasks “easier” may obfuscate the fact that forms of labor are unevenly distributed and unevenly compensated, and that technologies themselves may not remedy these situations. Each fiction depicts everyday work practices from a different perspective: from developers, to online service employees, to crowdworkers.

Our fictions highlight that some groups will receive conveniences from these technologies, but not necessarily in the ways imagined by popular narratives positing that BCIs will positively augment human capabilities. SparkTheMatch

presumably becomes more efficient with their content review process, improves their dating service user experience, and saves on labor costs by using BCI-assisted crowdworkers. However, others have to contend with new types of labor in order to bring these conveniences to those who enjoy them. While the crowdworkers' labor consists of staring at a blinking screen while their subconscious does the "work," this may be detrimental to their health and quality of life. Furthermore, the crowdworkers need to purchase their own BCI device in order to do the SparkTheMatch tasks, showing some of the unforeseen costs of their labor.

We are drawn to asking who gets access to the benefits of BCIs, who does the labor, and how these might intersect with disadvantaged groups. This leads us to think about the ways in which the initial designers and developers of the BCI API are implicated in questions of power, inequality, and social justice of these crowdworkers, even though their practices are obfuscated from the perspective of SparkTheMatch, its users, or the crowdworkers. We see the scenarios outlined in our fictions as one way to provoke conversations among designers and researchers about the societal risks and benefits of brain-computer interfaces.

Situated (Re)cognition

Technologies (and their associated human practices) typically do not generalize to all contexts [56]; however, with free APIs, technologies will be (re)purposed by those who do not know how far these technical capabilities can be generalized. In our fictions, the API in question was repurposed from its original context (an API to help with a set of data cleaning tasks, created from a training data set from a set of research lab studies) to a new use case (reviewing flagged user content), to which it may or may not be well suited. Our fiction highlights the formal and informal ways in these systems "travel": from academic-industry partnerships, to the ability to host and call APIs, to online developer communities like Stack Overflow.

The fictions also allude to how the technical P300 API system may not be applicable to all situations. These judgements of appropriateness can be made in informal settings, in this case through a Stack Overflow post from a stranger discussing the generalizability of the API. While one developer tried to pay attention to the appropriateness and situatedness of the P300 API and its training set, these discussions become obfuscated as more layers are built on top of the API. Yet these properties sometimes surface in unexpected ways, such as crowdworkers noting the unknown "errors" they sometimes encounter using the system.

Our fictions also highlight the situated positionality of developers and others positioned in the "middle of the stack," instead of just end users. We provided a glimpse into the practices of software developer Jay when he has some notion about potential social issues in his work but is not quite sure how to interrogate them beyond asking on Stack Overflow. This leads us to think about ways to support developers like Jay, or Pat's IT team at SparkTheMatch in thinking through

social issues related to their technical practices. What tools, training, policies, or institutional changes might help? How might designers and researchers better anticipate the variety of adoptions and appropriations of their technical work?

Reflecting on BCIs as Sociotechnical with Design Fiction

The issues raised in the preceding section might at first seem out of scope for a BCI researcher. But by using design fiction to think about BCIs as sociotechnical systems, connections to things that the BCI applications build on, things that make use of the BCI applications, and the practices creating, maintaining, and adopting those systems come into focus as important issues to consider. This work contributes to work in HCI and adjacent fields focusing on infrastructures, or what Wong and Jackson term looking "beyond the moment of design, and above the level of the artifact", taking into account "changes in sites and scales of computing" [62:106].

We traced a technical platform and associated practices from API documentation in a research lab to a Google-hosted API, to a developer building a system using the API, to a company utilizing that system, to crowdworkers hired by the company. Design fiction's focus on the mundane and everyday combined with the perspective of platforms to engage in design fiction "world building" [13] gives us insight into the multiplicity of viewpoints, stakeholders, and practices that make the adoption, contestation, and appropriation of the BCI API possible. Using multiple fictions within the same world highlighted how assumptions, values, and practices surrounding the system at a particular place and time can be hidden, adapted, or changed by the time the system reaches others. Multiple interconnected fictions allowed us to explore how supposed benefits and costs related to the introduction of new technologies (such as "increased efficiency") are unevenly distributed among populations.

Design fiction also helps us engage with the complex platforms and infrastructures of the real world. Some of the concerns raised are similar to sociotechnical issues today; after all, BCI adoption is unlikely to lead to a radical utopian or dystopian future. Rather, BCI adoption will graft or build onto existing systems and practices, potentially hiding some aspects (such as the contextual limits to a set of training data), while reinscribing others (such as existing labor issues on crowd platforms). By bringing these conversations to the emerging field of consumer BCIs, we hope to turn design researchers' attention toward the social implications of BCIs and raise BCI researchers' awareness of these ongoing conversations in HCI, so that they might bring to bear on ongoing design decisions in this nascent field. Our different formats and genres utilized in the fictions speak to different audiences, all of whom will be involved in the design, adoption, and propagation of BCIs.

Using a design fiction to explore "implications for adoption" [38] of BCIs suggests that design decisions in academic and research work can have real world effects, exemplified by the P300 API's academic origins. While academic and industry practices may be differently situated, there is not a

clear distinction between the lab and the rest of the world. Bell, Dourish, and Mainwaring highlight this in the context of ubiquitous computing, arguing that the lab and the world are deeply implicated in one another [4,20]. Given that much of BCI research currently exists in academic spaces, we note parallels to early ubiquitous computing research to argue that BCI researchers and designers should also be concerned about the social implications of adoption and application.

Two of this paper's authors have previously conducted technical research on security applications for BCIs. One author reflected that while he had previously thought about BCIs located in the everyday, these imagined futures tended to center on interactions common enough to become mundane, and thus invisible. The fictions allowed him to think about a new set of socioeconomic questions around how new technologies change what types of labor are valued or not valued. As a group, we began thinking about the potential for BCIs to be used as new forms of Taylorism in which intellectual labor becomes menial and repetitive. We began to ask how BCI designers and engineers could navigate (and possibly avoid) these types of futures.

IMPLICATIONS FOR (BCI) DESIGNERS

This paper approached BCIs as emerging interfaces and emerging platforms for software services. Unlike many popular imaginaries, in which radical new configurations of society emerge following the development of BCIs, we imagined a mundane future for these devices, in which existing systems of power (particularly labor inequality) persist and are reinforced by BCIs' technical infrastructures.

For design researchers, we find that design fiction can be used to "graft" proximate futures onto existing platforms and infrastructures. This creates a more "analogous" world to the present which can then be used to articulate and analyze "implications for adoption" [38]. This also recognizes the ways in which new technologies (in the present and future) are not deployed into a blank slate world, but are uptaken into and by existing sociotechnical systems and assemblages.

We also present our design fictions to those designing and engineering BCI systems. As BCIs work their way into the mainstream of the startup ecosystem, we ask designers to consider diverse (and often unexpected) uptakes and repurposings of their designs. Popular hyperbolic narratives muddle the banality of more probable outcomes (positive and negative). On one hand, design fictions such as this one can surface new narratives. These narratives could allow BCI designers to better anticipate how their implementations might be used by a diverse set of stakeholders. Design researchers and researchers exploring social aspects of technology can engage with BCI researchers proactively in discussion and reflection on these issues. Our creation and analysis of a set of design fictions suggests that BCI researchers and designers can use design fiction (or collaborate with design researchers) to surface and discuss these issues. Recent work documents how technology probes can be used with engineers to surface their beliefs about the

futures of BCIs, often imagining a future in which BCIs are pervasive and socially acceptable [40]. Design fiction may be able to build on and possibly challenge BCI designers' existing beliefs about the future. On the other hand, design fiction is not a panacea. Criticism and reflexivity should be embedded at all layers of technical practice in the development of BCIs, from academic researchers to industry developers to end users. Design fiction, along with other practices used to surface reflexive and critical discussion of social values such as critical technical practice, value sensitive design, or reflexive design, could be useful approaches to accomplishing this.

The fictions presented in this paper are not exhaustive, but rather present a specific viewpoint about the relation between technology and labor practices. Future work can add to this through further design work and empirical research. Additional design fictions can interrogate other possible arrangements of power related to the API's use and adoption situated in the everyday, such as the role of regulation, surveillance and privacy issues, or users' subversion and appropriation of BCI systems. Empirically, the design fictions can also be shared with BCI designers to spark and surface discussions on values and social implications of their work. Design fictions can also be used as research probes to further study how BCI designers and researchers conceptualize future adoptions and uses of BCIs—potentially in ways different than the fictions imagine them.

CONCLUSION

We presented a set of interconnected design fictions tracing a fictional BCI API as it gets appropriated and used by several different stakeholders. We used these fictions to raise a series of questions about forms of labor that BCIs might enable and the unequal distribution of benefits, and about how values and social issues related to technical platforms can become black boxed and obfuscated. Among BCI researchers and designers, we aim to raise awareness of these ongoing conversations in HCI so that they may inform this ongoing technical practice. For design researchers, we show that increasing the perspectives and interconnected relations in worlds that we imagine can help us interrogate sociotechnical issues present in the everyday and better help us think about the futures that we desire.

ACKNOWLEDGMENTS

We thank the BioSENSE group and Ellie Harmon for their insightful conversations, and the reviewers for their helpful feedback. This work was supported in part by the Berkeley Center for Long Term Cybersecurity (CLTC), the Hewlett Foundation, the National Science Foundation (NSF) Graduate Research Fellowship Program under Grant No. DGE 1106400, and NSF "INSPIRE: Value-Function Handoffs in Human-Machine Compositions" under Grant No. SES 1650589. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF, Hewlett Foundation, or CLTC.

REFERENCES

1. Peter M. Asaro. 2000. Transforming society by transforming technology: The science and politics of participatory design. *Accounting, Management and Information Technologies* 10, 4: 257–290. [https://doi.org/10.1016/S0959-8022\(00\)00004-7](https://doi.org/10.1016/S0959-8022(00)00004-7)
2. James Auger. 2013. Speculative design: crafting the speculation. *Digital Creativity* 24, 1: 11–35. <https://doi.org/10.1080/14626268.2013.767276>
3. Jeff Barr and Luis Felipe Cabrera. 2006. AI gets a brain. *ACM QueueQueue* 4, 4: 24–29. <https://doi.org/10.1145/1142055.1142067>
4. Genevieve Bell and Paul Dourish. 2007. Yesterday's tomorrows: Notes on ubiquitous computing's dominant vision. *Personal and Ubiquitous Computing* 11: 133–143. <https://doi.org/10.1007/s00779-006-0071-x>
5. Julian Bleecker. 2009. *Design Fiction: A Short Essay on Design, Science, Fact and Fiction*. Retrieved from <http://www.nearfuturelaboratory.com/2009/03/17/design-fiction-a-short-essay-on-design-science-fact-and-fiction/>
6. Mark Blythe. 2014. Research through design fiction: Narrative in Real and Imaginary Abstracts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*, 703–712. <https://doi.org/10.1145/2556288.2557098>
7. Torie Bosch. Sci-Fi Writer Bruce Sterling Explains the Intriguing New Concept of Design Fiction. *Slate.com*. Retrieved September 20, 2016 from http://www.slate.com/blogs/future_tense/2012/03/02/bruce_sterling_on_design_fictions_.html
8. Barry Brown, Julian Bleecker, Marco D'Adamo, Pedro Ferreira, Joakim Formo, Mareike Glöss, Maria Holm, Kristina Höök, Eva-Carin Banka Johnson, Emil Kaburuan, Anna Karlsson, Elsa Kosmack-Vaara, Jarmo Laaksolahti, Airi Lampinen, Lucian Leahu, Vincent Lewandowski, Donald McMillan, Anders Mellbratt, Johanna Mercurio, Cristian Norlin, Nicolas Nova, Stefania Pizza, Asreen Rostami, Mårten Sundquist, Konrad Tollmar, Vasiliki Tsaknaki, Jinyi Wang, Charles Windlin, and Mikael Ydholm. 2016. The IKEA Catalogue: Design Fiction in Academic and Industrial Collaborations. In *Proceedings of the 19th International Conference on Supporting Group Work (GROUP '16)*, 335–344. <https://doi.org/10.1145/2957276.2957298>
9. Kirsten V. Brown. Facebook Literally Wants to Read Your Thoughts. Blog post. *Gizmodo*. Retrieved September 18, 2017 from <http://gizmodo.com/facebook-literally-wants-to-read-your-thoughts-1794471569>
10. Jenna Burrell. 2016. How the machine “thinks”: Understanding opacity in machine learning algorithms. *Big Data & Society* 3, 1: 1–12. <https://doi.org/10.1177/2053951715622512>
11. Adrian Chen. 2014. The Laborers Who Keep Dick Pics and Beheadings Out of Your Facebook Feed. *Wired*. Retrieved January 4, 2018 from <http://www.wired.com/2014/10/content-moderation/>
12. Computing Community Consortium (CCC). 2015. *Privacy by Design-Engineering Privacy. Workshop 3 Report*. Retrieved from <http://cra.org/ccc/wp-content/uploads/sites/2/2015/12/PbD3-Workshop-Report-v2.pdf>
13. Paul Coulton, Joseph Lindley, Miriam Sturdee, and Michael Stead. 2017. Design Fiction as World Building. *Proceedings of the 3rd Biennial Research Through Design Conference*, March: 1–16. <https://doi.org/10.6084/m9.figshare.4746964>
14. CTRL-labs. 2017. CTRL-labs: all your interface are belong to us. Web page. Retrieved September 18, 2017 from <https://ctrl-labs.com/>
15. Max T Curran, Nick Merrill, John Chuang, and Swapan Gandhi. 2017. One-step, three-factor authentication in a single earpiece. In *Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2017 ACM International Symposium on Wearable Computers on - UbiComp '17*, 21–24. <https://doi.org/10.1145/3123024.3123087>
16. Lynn Dombrowski, Adriana Alvarado Garcia, and Jessica Despard. 2017. Low-Wage Precarious Workers' Sociotechnical Practices Working Towards Addressing Wage Theft. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*, 4585–4598. <https://doi.org/10.1145/3025453.3025633>
17. Lynn Dombrowski, Ellie Harmon, and Sarah Fox. 2016. Social Justice-Oriented Interaction Design. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16)*, 656–671. <https://doi.org/10.1145/2901790.2901861>
18. Paul Dourish and Genevieve Bell. 2011. *Divining a Digital Future: Mess and Mythology in Ubiquitous Computing*. The MIT Press, Cambridge, Massachusetts.
19. Paul Dourish and Genevieve Bell. 2013. “Resistance is futile”: reading science fiction alongside ubiquitous computing. *Personal and Ubiquitous Computing* 18, 4: 769–778. <https://doi.org/10.1007/s00779-013-0678-7>
20. Paul Dourish and Scott D Mainwaring. 2012. UbiComp's colonial impulse. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*, 133. <https://doi.org/10.1145/2370216.2370238>
21. Hamid Ekbria and Bonnie Nardi. 2016. Social Inequality and HCI: The View from Political

- Economy. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*, 4997–5002. <https://doi.org/10.1145/2858036.2858343>
22. Emotiv Inc. 2017. Emotiv. Web page. Retrieved September 18, 2017 from <https://www.emotiv.com/>
 23. Verena Fuchsberger, Thomas Meneweger, Daniela Wurhofer, and Manfred Tscheligi. 2017. Apply Now! Fictional Job Postings as an Instrument to Discuss Interactive Futures of Work. In *Proceedings of the 2017 Conference on Designing Interactive Systems (DIS '17)*, 581–586. <https://doi.org/10.1145/3064663.3064750>
 24. Tarleton Gillespie. 2010. The politics of “platforms.” *New Media & Society* 12, 3: 347–364. <https://doi.org/10.1177/1461444809342738>
 25. Melissa Gregg. 2011. *Work's Intimacy*. Polity, Cambridge.
 26. Ellie Harmon, Chris Bopp, and Amy Volda. 2017. The Design Fictions of Philanthropic IT: Stuck Between an Imperfect Present and an Impossible Future. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. <https://doi.org/10.1145/3025453.3025650>
 27. Sabrina Hauser, Audrey Desjardins, and Ron Wakkary. 2014. Sfuture: envisioning a sustainable university campus in 2065. In *Proceedings of the 2014 companion publication on Designing interactive systems (DIS Companion '14)*, 29–32. <https://doi.org/10.1145/2598784.2602774>
 28. Lily Hay Newman. 2017. It's not always AI that sifts through your sensitive info. *Wired.com*. Retrieved January 4, 2018 from <https://www.wired.com/story/not-always-ai-that-sifts-through-sensitive-info-crowdsourced-labor/>
 29. Johannes Hohne, Martijn Schreuder, Benjamin Blankertz, and Michael Tangermann. 2010. Two-dimensional auditory P300 Speller with predictive text system. In *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*, 4185–4188. <https://doi.org/10.1109/IEMBS.2010.5627379>
 30. Jeff Howe. 2006. The Rise of Crowdsourcing. *Wired Magazine* 14, 1–5. Retrieved from <https://www.wired.com/2006/06/crowds/>
 31. Interaxon. 2017. Muse: the brain sensing headband. Web page. Retrieved September 18, 2017 from <http://www.choosemuse.com/>
 32. Lilly C Irani and M. Six Silberman. 2013. Turkopticon: Interrupting worker invisibility in amazon mechanical turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, 611. <https://doi.org/10.1145/2470654.2470742>
 33. David Kirby. 2010. The Future is Now: Diegetic Prototypes and the Role of Popular Films in Generating Real-world Technological Development. *Social Studies of Science* 40, 1: 41–70. <https://doi.org/10.1177/0306312709338325>
 34. Joshua A. Kroll, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson, and Harlan Yu. 2017. Accountable Algorithms. *University of Pennsylvania Law Review* 165: 633–706.
 35. Steven Levy. 2017. Brain-Machine Interface Isn't Sci-Fi Anymore. *Wired.com*. Retrieved January 7, 2018 from <https://www.wired.com/story/brain-machine-interface-isnt-sci-fi-anymore/>
 36. Joseph Lindley and Paul Coulton. 2015. Back to the Future: 10 Years of Design Fiction. In *Proceedings of the 2015 British HCI Conference (British HCI '15)*, 210–211. <https://doi.org/10.1145/2783446.2783592>
 37. Joseph Lindley and Paul Coulton. 2016. Pushing the Limits of Design Fiction: The Case For Fictional Research Papers. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*, 4032–4043. <https://doi.org/10.1145/2858036.2858446>
 38. Joseph Lindley, Paul Coulton, and Miriam Sturdee. 2017. Implications for Adoption. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*, 265–277. <https://doi.org/10.1145/3025453.3025742>
 39. Joseph Lindley, Dhruv Sharma, and Robert Potts. 2015. Operationalizing Design Fiction with Anticipatory Ethnography. In *Ethnographic Praxis in Industry Conference*, 58–71.
 40. Nick Merrill and John Chuang. 2018. From Scanning Brains to Reading Minds : Talking to Engineers about Brain-Computer Interface. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '18)*. <https://doi.org/10.1145/3173574.3173897>
 41. Michael J. Muller and Sarah Kuhn. 1993. Participatory design. *Communications of the ACM* 36, 6: 24–28. <https://doi.org/10.1145/153571.255960>
 42. Near Future Laboratory. TBD Catalog. Retrieved July 18, 2017 from <http://tbdcatalog.com/index.html>
 43. Neurale Inc. 2017. neurable. Web page. Retrieved September 18, 2017 from <http://www.neurable.com/>
 44. NeuroSky. 2017. NeuroSky: Body and Mind. Quantified. Web page. Retrieved September 18, 2017 from <http://neurosky.com/>
 45. Christena E. Nippert-Eng. 1996. *Home and Work: Negotiating Boundaries through Everyday Life*. University of Chicago Press, Chicago.
 46. Lucas Nolan. 2017. Facebook working on brain-computer interface to translate thoughts to text.

- Breitbart.com. Retrieved January 4, 2018 from <http://www.breitbart.com/tech/2017/04/27/facebook-working-on-brain-computer-interface-to-translate-thoughts-to-text/>
47. Nicolas Nova, Katherine Miyake, Walton Chiu, and Nancy Kwon. 2012. *Curious rituals: Gestural interaction in the digital everyday*. Retrieved from <http://curiousrituals.wordpress.com>
 48. Greig de Peuter. 2011. Creative Economy and Labor Precarity: A Contested Convergence. *Journal of Communication Inquiry* 35, 4: 417–425. <https://doi.org/10.1177/0196859911416362>
 49. James Pierce and Eric Paulos. 2014. Some variations on a counterfunctional digital camera. In *Proceedings of the 2014 conference on Designing interactive systems (DIS '14)*, 131–140. <https://doi.org/10.1145/2598510.2602968>
 50. Sarah T. Roberts. 2016. Commercial Content Moderation: Digital Laborers' Dirty Work. *The Intersectional Internet: Race, Sex, Class and Culture Online*: 1–12. <https://doi.org/10.1007/s13398-014-0173-7.2>
 51. Alex Rosenblat and Luke Stark. 2016. Uber's Drivers: Information Asymmetries and Control in Dynamic Work. *International Journal of Communication*, 10: 27. <https://doi.org/10.2139/ssrn.2686227>
 52. Nathan Shedroff and Christopher Noessel. 2012. *Make it So: Interaction Design Lessons from Science Fiction*. Rosenfield Media, Brooklyn, New York.
 53. Michael Skirpan and Casey Fiesler. 2018. Ad Empathy: A Design Fiction. In *Proceedings of the 2018 ACM Conference on Supporting Groupwork (GROUP '18)*, 267–273. <https://doi.org/10.1145/3148330.3149407>
 54. Eliza Strickland. 2017. Facebook Announces “Typing-by-Brain” Project. Blog. *IEEE Spectrum: The Human OS*. Retrieved September 18, 2017 from <https://spectrum.ieee.org/the-human-os/biomedical/bionics/facebook-announces-typing-by-brain-project>
 55. Miriam Sturdee, Paul Coulton, Joseph G. Lindley, Mike Stead, Haider Ali, and Andy Hudson-Smith. 2016. Design Fiction: How to Build a Voight-Kampff Machine. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)*, 375–386. <https://doi.org/10.1145/2851581.2892574>
 56. Lucy Suchman. 2006. *Human-machine reconfigurations: Plans and situated actions*. Cambridge University Press, Cambridge.
 57. Joshua Tanenbaum, Marcel Pufal, and Karen Tanenbaum. 2016. The Limits of Our Imagination: Design Fiction as a Strategy for Engaging with Dystopian Futures. In *Proceedings of the Second Workshop on Computing within Limits (LIMITS '16)*. <https://doi.org/10.1145/2926676.2926687>
 58. Joshua Tanenbaum, Karen Tanenbaum, and Ron Wakkary. 2012. Steampunk as design fiction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*, 1583. <https://doi.org/10.1145/2207676.2208279>
 59. Tiziana Terranova. 2004. *Network Culture: Politics for the Information Age*. Pluto Press, London.
 60. Tim Urban. 2017. Neuralink and the Brain's Magical Future. Webpage. *Wait But Why*. Retrieved May 30, 2017 from <https://waitbutwhy.com/2017/04/neuralink.html>
 61. Jessica Vitak, Katie Shilton, and Zahra Ashktorab. 2016. Beyond the Belmont Principles: Ethical Challenges, Practices, and Beliefs in the Online Data Research Community. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*, 939–951. <https://doi.org/10.1145/2818048.2820078>
 62. Richmond Y. Wong and Steven J. Jackson. 2015. Wireless Visions: Infrastructure, Imagination, and U.S. Spectrum Policy. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. <https://doi.org/10.1145/2675133.2675229>
 63. Richmond Y. Wong and Deirdre K. Mulligan. 2016. When a Product Is Still Fictional: Anticipating and Speculating Futures through Concept Videos. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16)*, 121–133. <https://doi.org/10.1145/2901790.2901801>
 64. Richmond Y. Wong, Ellen Van Wyk, and James Pierce. 2017. Real - Fictional Entanglements: Using Science Fiction and Design Fiction to Interrogate Sensing Technologies. In *Proceedings of the 2017 ACM Conference on Designing Interactive Systems (DIS '17)*. <https://doi.org/10.1145/3064663.3064682>
 65. Shoshana Zuboff. 1988. *In the Age of the Smart Machine: The Future of Work and Power*. Basic Books, New York.