

Machine Learning Uncertainty as a Design Material: A Post-Phenomenological Inquiry

Jesse Josua Benjamin
j.j.benjamin@utwente.nl
Department of Philosophy
University of Twente
Enschede, Netherlands

Nick Merrill
ffff@berkeley.edu
Center for Long-Term Cybersecurity
University of California, Berkeley
Berkeley, United States

Arne Berger
arne.berger@hs-anhalt.de
Computer Science and Languages
Anhalt University of Applied Sciences
Koethen, Germany

James Pierce
jppierce@uw.edu
School of Art + Art History + Design
University of Washington
Seattle, United States

ABSTRACT

Design research is important for understanding and interrogating how emerging technologies shape human experience. However, design research with Machine Learning (ML) is relatively underdeveloped. Crucially, designers have not found a grasp on ML uncertainty as a design opportunity rather than an obstacle. The technical literature points to data and model uncertainties as two main properties of ML. Through post-phenomenology, we position uncertainty as one defining material attribute of ML processes which mediate human experience. To understand ML uncertainty as a design material, we investigate four design research case studies involving ML. We derive three provocative concepts: *thingly uncertainty*: ML-driven artefacts have uncertain, variable relations to their environments; *pattern leakage*: ML uncertainty can lead to patterns shaping the world they are meant to represent; and *futures creep*: ML technologies texture human relations to time with uncertainty. Finally, we outline design research trajectories and sketch a post-phenomenological approach to human-ML relations.

CCS CONCEPTS

• **Human-centered computing** → **HCI theory, concepts and models.**

KEYWORDS

post-phenomenology, machine learning, design research, thingly uncertainty, horizontal relations

ACM Reference Format:

Jesse Josua Benjamin, Arne Berger, Nick Merrill, and James Pierce. 2021. Machine Learning Uncertainty as a Design Material: A Post-Phenomenological Inquiry. In *CHI Conference on Human Factors in Computing Systems (CHI '21)*, May 8–13, 2021, Yokohama, Japan. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3411764.3445481>



This work is licensed under a Creative Commons
Attribution-NonCommercial-ShareAlike International 4.0 License.

CHI '21, May 8–13, 2021, Yokohama, Japan
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8096-6/21/05.
<https://doi.org/10.1145/3411764.3445481>

1 INTRODUCTION

Just as prior advances in computing technology have led to design researchers approaching algorithms, software interfaces, sensors, and actuators as *design materials* [44, 60], similar developments are happening with Machine Learning (ML) technologies (cf. [10, 27, 41, 80]). Recent advances in this regard have resulted in guidelines from computing corporations for curating ML results, improving user experience design and Explainable Artificial Intelligence (XAI) (e.g., [2, 19, 43, 81]). In light of the increasing ubiquity and technical opacity of ML, design research methodologies such as Research-through-Design, Speculative Design or Design Fiction are urgently needed in this space to develop better tools grounded in a rich and socially-situated understanding of how ML shapes everyday life.

We argue that design research is well-posed to describe, explore and reflect on how the “statistical intelligence” [10] of ML decision-making bleeds into the intimate human experience of lifeworlds; and to productively engage with emerging personal, societal and ethical issues. At the same time, we observe that design research struggles to engage ML as a material for design, as the probabilistic inference of models from data patterns withdraws from being present-at-hand. A major focus of prior work is the technical opacity of ML. However, just as ML-driven systems are difficult to interpret and exhibit emergent and oftentimes unpredictable behavior [8], the outputs they generate are inherently characterized by uncertainty from data noise and model variance [35].

ML uncertainty is a problem for HCI design research, because the field has not yet framed it as a design material to assess, design, and reflect novel applications, objects and services that build on ML. On the contrary, prior work, for example in the area of XAI, is often concerned with using design methods to explain, rather than utilize, ML uncertainty. We propose that approaching ML uncertainty as a specific material property of ML can help design research consider novel spaces for design intervention and interaction.

In this paper, we investigate ML uncertainty from a post-phenomenological perspective to develop a conceptual vocabulary for engaging ML uncertainty in designerly ways. The principal argument of post-phenomenology is that technologies actively *mediate* human relations to the world [33]: technologies are neither completely deterministic (technological determinism) but neither are

they neutral tools (technological instrumentalism). Instead, human-technology relations are co-constitutive, with artefacts shaping how humans as specific subjectivities relate to the world [67]. In this view, design is *doing* philosophy “by other means” [69], and design materials are those involved in shaping technological mediation. We follow Hauser et al.’s framing of Research-through-Design (RtD) as post-phenomenological practice [25], and expand this view to further design research methodologies and the specific topic of ML. We thereby also continue prior work from the HCI community on using and extending the post-phenomenological framework (cf. [26, 58, 72]). Specifically, we investigate four design research case studies using a post-phenomenological lens, and develop provocative concepts for ML uncertainty similar to “strong concepts” [32] or robust “annotations” [14, 25] in HCI. In this, we take a modest step forward in making ML uncertainty tangible as a design material for HCI design researchers. We argue that uncertainty is the material expression (cf. [60]) of ML *decision-making*. While there are reasonable engineering incentives to minimize uncertainty in many use cases, uncertainty constitutes a fundamental attribute of any ML-driven system. Uncertainty offers a representation of the ‘fault lines’ of ML decision-making, not only as a negative attribute of solutions but as simply part-and-parcel of ML technologies. As such, designerly making with uncertainty offers an opportunity to design artefacts and scenarios that attribute or more fully exploit the characteristics of ML decision-making. Our contributions consist of provocative, conceptual shorthands for ML uncertainty. We encourage designers to not see ML uncertainty as “to be explained away,” but rather as generative of particular relations that can be designed for. At the same time, we also show how emerging ML applications are not readily accounted for with prior human-technology relation concepts.

We provide three provocative concepts that build on ML uncertainty as a source for future research in HCI, design research and philosophy of technology communities. As a general concept, we propose *thingly uncertainty* to capture the capacity of ML-driven artefacts to be uncertain about the world, and thereby generating and adapting to a wide continuum of relations to other things, their datafied environment and people. We furthermore distinguish two specific concepts. First, we propose *pattern leakage* for describing how ML models become generative of patterns which enter and alter the everyday. Pattern leakage describes how ML models alter the world they seek to represent. Second, we propose *futures creep* to denote the often subtle transformations of the present by ML predictions; changing human relations to time by injecting probabilistic events such as climate change predictions into the direct perception of the present. Based on our derived concepts, we propose how design research can engage with ML uncertainty, and furthermore suggest *horizontal relations* as a post-phenomenological research trajectory into the specifics of how ML capacities and human experience co-extend and overlap. Among the contributions of this paper, we (1) construct a post-phenomenological lens on ML based on related work; (2) analyze four case studies from different design research methodologies through this lens to discern ML uncertainty as a design material; (3) present *thingly uncertainty*, *pattern leakage* and *futures creep* as provocative concepts for future work; and (4) propose *horizontal relations* as a distinct human-technology relation, and lay out research trajectories for investigation.

2 BACKGROUND

In the following section we (1) provide a brief introduction to ML uncertainty; (2) frame how design research can pragmatically and critically engage with emerging technologies; (3) argue that post-phenomenology is a promising lens to make ML uncertainty graspable in design research.

2.1 ML Technologies between Opacity, Interpretability and Uncertainty

Many of the diverse algorithmic techniques common in ML today stem from cybernetics. McCulloch and Pitt’s notion of neuronal activity [49] and Rosenblatt’s *Perceptron* as a probabilistic model for learning and remembering information about the environment [63] are in direct connection to today’s advances in deep learning, while Wiener’s concept of negative feedback [74] can still be seen as the general principle which make ostensibly novel technologies tick. And like cybernetics, today’s ML technologies are reliant on probabilistic techniques as a computational means for the “taming of chance” [20]. The deployment of increasingly powerful probabilistic techniques (e.g., expectation-maximization, gradient descent, backpropagation) capable of adapting models to large datasets in real-world settings has also come at the cost of opacity, and become generative of types of uncertainty originating from *within* technological deployments themselves.

In general, ML algorithms operate according to the principle of “insight through opacity” [50], using specific probabilistic techniques to infer a model that describes statistical patterns in datasets [6], which approximates some assumed real-world functional relationship [18]. The insight, i.e. a model that can detect significant patterns in relevant data (e.g., does this image show tumors?), depends on opacity: as each variable such as a pixel is computed as a vector in relation to all other variables, the combined dimensionality of vectors exceeds those that humans can intuit. In short, the often “unreasonable effectiveness” [21] of current ML implementations comes at the cost of understanding how exactly insights were arrived at. From within ML research, the challenge of opacity is one of the most urgent topics of research, showing in fields such as XAI or ML interpretability research. A common approach for ML research is to deploy algorithmic methods, so-called interpretability techniques [45, 53], to extract information from ML pipelines in order to explain outputs via textual or visual explanations. While this research area is fundamentally oriented at experts, its focus on actual ML technologies offers up a catalogue of properties potentially of interest to design research. Given ML’s epistemological as well as technical origins in cybernetics, we are specifically interested in how uncertainty is dealt with in light of the contemporary insight-through-opacity approach.

In ML research, the aim of dealing with uncertainty takes on a distinctly *material* grounding: given the reliance on probabilistic techniques, uncertainty is not only a human disposition but part and parcel of ML technologies. ML researchers frequently distinguish between two types of uncertainty which are related to data (both training and input) and models respectively [12, 35, 37]. The former, “aleatoric” uncertainty, can be framed as ‘noise’: incoming signals in training or real-world deployments are inevitably impure and may affect the performance of ML algorithms. The latter,

“epistemic” uncertainty, refers to the complex questions surrounding the ‘fit’ of a generated ML model. Inferred models embody, firstly, only one way of describing patterns from a given dataset, and given the insight-through-opacity approach the relationship to unknown models is uncertain. Secondly, an inferred model may also *generate* additional uncertainty when deployed “out-of-data”, i.e. in settings different to training environments. Data and model uncertainty frequently feature in ML interpretability research. It is important to note that both are *computable* depending on the given ML deployment. Hohman et al., for example, deploy a technology probe for ML experts featuring various data visualizations in one interface, which includes a “regions-of-error” technique showing the model uncertainty of predictions [30]. Similarly, Kinkeldey et al. use a landscape metaphor in a cluster visualization, indicating through a grey-scale topography how certain the clustering model is about the membership of each individual point by their location in “peaks or slopes” [36]. Concerning data uncertainty, Kendall and Gal note that in image segmentation using deep learning, noise affects the boundaries surrounding objects [35]. Similarly, Kwon et al. note that data uncertainty in brain lesion detection with neural networks manifests around affected brain regions [39].

On a general level, we therefore consider data uncertainty to manifest with the *objects* of ML decision-making (e.g., data as images, strings, vectors), whereas model uncertainty concerns the *mode* of ML decision-making (e.g., clustering, classifying, predicting). While the technical fields take an understandably solutionist stance on ML uncertainty by attempting to either minimize or explain it, we argue that the notion of computational uncertainty as a part-and-parcel property of ML promises opening actual ML technologies to designerly research.

2.2 Design Research in ML and Emerging Technologies

Uncertainty, understood in an everyday sense as ambiguity or chance, is a well-known resource in design research (cf. [15]). As an alternative to quantitative studies, design research focuses on unearthing, exploring and understanding individual and situated encounters of people with technology that are often based on uncertainty and indeterminacy. Particularly ‘third-wave HCI’ [24] methodologies such as Research-through-Design, Design Fiction or Speculative Design have engaged this trajectory, foregrounding how technological artefacts are not merely solutions to discrete problems but rather embody open-ended, contextually dependent questions. Yet, design research has predominantly focused on either using ML *for* design, or using design *for* ML. With regards to the former, for example, Yang et al. have discussed how ML is engaged by design practitioners and researchers [80–82] to improve user experience through adaptation or personalization. They found that while practitioners and researchers are enthusiastic about using ML, design research so far lacks distinct methods that engage ML as a design material. Similarly, Dove et al., conclude that integrative prototyping methods reflecting both “ML statistical intelligence and human common sense intelligence” [10] are missing in the field. However, in contrast to this research focus on ML *for* design, we argue that more fundamental research is needed into how, and to what end, ML technologies can be a design material in their own right.

Similarly, in HCI approaches to XAI, researchers envision design research methods for improving a given system’s explainability; by e.g. developing explainability scenarios [3, 76] or conceptualizing contextually sensitive questions [43, 73] for various stakeholders. We see a similar, if inverted, limitation in this use of design *for* ML: in XAI, ML technologies tend to become the target (for e.g. designing explanations), not the material, of design research methods. We therefore next consider critical design research projects that, while not always directly addressing ML technologies, deal with uncertainties of complex design materials.

For example, Merrill et al. found that lay people believe different types of biosensors can reveal much more, or much less, than they actually can [52]. Others have explored how people speculate on smart things for the home, based on the fluidity and diversity of the values they attribute to their individual homes. This unraveled the idiosyncrasies and situatedness of potential future smart things, and highlights that these are bound to individual experiences with uncertainty [4, 56]. Pierce’s design-led inquiry into smart camera systems [58] investigates the ways in which smart camera systems opaquely embody specific relations to the world—that is, functions which may be concealed due to undesired effects (e.g., distrust, failure). Redström and Wiltse further interrogate how user interactions are tied to infrastructural functions, and outline how “surface-level simplicity” of interactions such as pressing play in Spotify belie “dynamic, sophisticated, and hidden backend complexity” [61].

This brief overview of critical design research projects shows that in general, situated yet concealed design materials are nothing new, and opacity and its often uncertain effects are a recurring theme across application areas. However, we observe that few critical design research projects deal with specific rather than generic ML technologies. We hypothesize that design research currently lacks a conceptual grasp on material properties that characterize ML technologies’ inference of models from data and decision-making (e.g., prediction, classification). Due to their relationship with input, inference and output of ML technologies, we propose that computational model and data uncertainty from the technical fields discussed above are promising candidates. That is, rather than a human-centered notion of uncertainty, we propose that the technical framing of uncertainty promises a stronger foothold on the ML design space. However, it is unclear how ML uncertainty can be framed specifically for design research. For example, Hemmentt et al. call for artistic, designerly practices of revealing the “distortions in the ways in which algorithms make sense of the world” [29]; yet do not outline how this may relate to existing design research methodologies as well as actual ML technologies. In the following, we discuss how the philosophical framework of post-phenomenology offers a starting point to expand on this gap.

2.3 Post-Phenomenology and Sketching Human-ML Relations

Post-phenomenology is an empirical-analytical framework in philosophy of technology. In this view, technological artefacts shape human perception and action by mediating the world in specific ways [33]. The framework has become particularly influential in HCI design research, as it foregrounds the responsibilities and ethico-political stakes of designing technological artefacts—surfacing how

technologies co-constitute our relations to specific ‘slices’ of the world, and how this relation makes us who we are [13, 25, 69].

Below, we outline established concepts of post-phenomenology in the form of their basic relational schemata, which allow empirical-analytical research to probe their objects of study for the roles that human, technology and world play.

The schemata use established notations for simple *connections* between entities (-); *interpretation* of one by the other (\rightarrow); being experienced together (\leftrightarrow); being in the *background* (/) of another entity; or being already *thematized* (i.e., meaningfully focused [28]) in some way (\square) before being experienced. We then argue that post-phenomenology is a well-suited framework to describe, analyze, and interpret ML as a technology with active yet concealed relations to the world.

2.3.1 Technological Mediation: Human ways of perceiving and acting in the world are shaped by technological artefacts, co-constituting who relates to what/whom in which way. Depending on the technology (e.g., an ultrasound scanner), ways of perceiving phenomena in the world are shaped (e.g., the fetus-as-patient, the womb-as-potentially-dangerous-enclosure), and actions are invited or inhibited (e.g., decisions on fetal care), for a specific human subjectivity (e.g., the non-pregnant parent-as-caretaker) [67]. The general schema for technological mediation is:

$$\text{Human} \text{ -- } \text{Technology} \text{ -- } \text{World}$$

2.3.2 Human-Technology Relations: Post-phenomenology studies technological mediation through human-technology relations (Table 1), the structures of the empirical settings in which humans and technologies encounter each other, and how such encounters shape how humans relate to the world [62].

Ihde initially proposed four types of human-technology relations [33]. In *embodiment relations*, technologies become an inseparable part of human bodily-perceptual experience. Wearing glasses shapes our experience of the world, not of the glasses themselves. In *hermeneutic relations*, technologies make the world legible in a specific way. Reading a map mediates the world as a grid; a thermometer translates heat phenomena onto a legible scale. More overt, in *alterity relations* technologies are seen as a quasi-Other. For example, an ATM, toy or robot are interacted with as if they have human-like intentions. In *background relations*, technologies merge with the background of experience, yet “texture” (i.e., generate an atmosphere for) that experience. Central heating or thermostats are absent presences in the experience of our home, but are rarely interacted with directly. Verbeek further expanded these relations to reflect the more subtle and also radical mediation by emerging technologies. In *immersion relations*, for instance, technologies such as augmented reality glasses or ambient intelligences merge with the environment, shaping how social relations can be enacted within them [70]. More radically, in *cyborg relations*, technologies such as microchips or pacemakers merge with the body, becoming indistinguishable in direct experience [68]. And lastly, in *composite relations*, technologies such as extremely long-exposure photography or computational imaging make things experienceable that have no direct correlation to ordinary human modes of perceiving space and time. Instead, composite relations mediate a “reality that can only be experienced by technologies” [68].

Relation	Schema	Examples
Embodiment	(I-Technology)World	Glasses, Cane
Hermeneutic	$I \rightarrow (\text{Technology} - \text{World})$	Thermometers, Maps
Alterity	$I \rightarrow \text{Technology} (- \text{World})$	ATMs, Robots
Background	$I (- \text{Technology/World})$	Heating, Thermostat
Immersion	$I \leftrightarrow \text{Technology/World}$	Virtual/Augmented Reality
Cyborg	$(I/\text{Technology}) \leftrightarrow \text{World}$	Implants, Pacemakers
Composite	$I \rightarrow (\text{Technology} \rightarrow \text{World})$	Computational Imaging

Table 1: Human-technology relations defined by Ihde and Verbeek, with associated schemata and examples.

2.3.3 Sketch for Human-ML Relations: In all the above human-technology relations, post-phenomenology attributes intentionality to technological artefacts, a material ‘directedness’ that these artefacts exhibit in relating to the world [67]: through intentionality, such as the thermometer’s combination of quicksilver and a scale, specific aspects of the world become legible in a specific way (e.g., ‘reading’ temperature). ML processes of inferring models from patterns found in data arguably strongly evidence this characteristic. At the same time, ML processes occur outside the phenomenological “horizon” of experience [33] in everyday life: we engage newsfeed-interfaces as technological artefacts, not the ML sorting algorithms. This leads us to two considerations for grasping human-ML relations post-phenomenologically. First, ML technologies shape our perception of the world via artefacts (e.g. devices and interfaces) which display, or are composed according to, their outputs. Second, outputs are not pre-configured, but depend on how data patterns and the parameters of the specific ML algorithm converge in a model. Accordingly, though the “model-world relations” [22] of ML algorithm and data affect how we experience the world, they do so in a way that is not directly empirically present. Given the established post-phenomenological concepts of *texturing* from the background of experience, *thematization* and the schema of *composite relations*, we can sketch a preliminary schema for technological mediation in human-ML relations as follows:

$$\text{Human} - \text{Technology} / (\text{Model} \rightarrow [\text{World}]) - \text{World}$$

This schema represents how an ML model, though active in the background of technological artefacts (/) that we experience within our phenomenological horizon, nevertheless *textures* that experience through its interpretive (\rightarrow), data-driven (\square) relations with the world. It can thus serve similar to the general schema of human-technology-world relations shown above in describing the technological mediation of ML. However, this schema alone does not yet provide concrete guidance on how precisely design research methodologies may gain a firmer grasp on ML through uncertainty, or whether there are particular human-technology relations beyond those in Table 1. So far, all it does is indicate that design researchers may probe or actively pursue ML outputs for higher variance (e.g., more or less uncertain), but not whether there may be ML-specific phenomena which should be paid special attention to. This forms our rationale for conducting an investigation of design research projects that engage ML technologies, as such an investigation will point to specific phenomena which then also feed back into our preliminary schema.

Case Study	Methodology	ML Technology	Context	Design Output
Pierce’s <i>Shifting Lines of Creepiness</i> [58]	Research-Through-Design	Image Recognition	Corporate / Home	Artefact Scenarios
Wong et al.’s <i>When BCIs have APIs</i> [79]	Design Fiction	Classification	Corporate / Labor	Infrastructure Scenarios
Wakkary et al.’s <i>Morse Things</i> [72]	Material Speculation	Reinforcement Learning	Home	Counterfactual Artefact
Biggs and Desjardins’ <i>Highwater Pants</i> [5]	Speculative Design	Linear Regression	Climate Futures	Speculative Artefact

Table 2: An overview of the selected design research projects as case studies for our analyses.

3 METHODOLOGICAL APPROACH

In this section we describe our rationale for selecting our case studies, and how we use them to frame ML uncertainty as a design material through post-phenomenological analyses. Post-phenomenology considers design as a practice of shaping technological mediation, as a way of *doing* philosophy “by other means” [69]. We build our approach on Hauser et al.’s work on the relationship between Research-through-Design in HCI and post-phenomenology [25], outlining the former as an experimental variant of the “interpretive empiricism” of the latter.

The hypothesis for our approach is twofold. Firstly, design research unfolds and shapes specific relations between humans and the world by designing technological artefacts. Secondly, as philosophy-in-practice, design research may hold latent propositions on how to think ML uncertainty in a post-phenomenological, designerly way. Therefore, the goal of the remainder of this paper is to use post-phenomenology to explicate and conceptualize the role of ML uncertainty in specific design research projects involving ML technologies.

3.1 Case Study Selection

Our selected case studies are design research projects from diverse methodological strands and application domains, to reflect both research and real-world concerns relating to ML technologies. We first gathered potential case studies from the corpus of CHI and DIS from 2015 to 2020. For the final selection, all authors met and discussed candidates (which also included e.g. [31, 40, 51]). We were especially interested in investigating a selection of case studies with a methodological, technological, contextual and designerly diversity. The presented case studies (cf. Table 2) were selected to cover a wide range of human-technology relations with different ML applications, while at the same time representing the diversity of HCI design research methodologies and contexts.

This diversity covers a wide spectrum of established human-technology relations (cf. Table 1), offering an empirical-analytical grounding to our own analysis (cf. [25]). Specifically, we selected Pierce’s design-led inquiry into smart home security cameras to reflect issues of leaking surveillance in the home [58], Wong et al.’s design fiction on Brain-Computer Interfaces (BCI) to reflect infrastructural relations [77], Wakkary et al.’s material speculation into more-than-human design of artefacts [72]; and Biggs and Desjardins’ speculative design of artefacts for relating to climate change predictions [5]. Our case study selection enables us to consider how, within the case studies, ML uncertainty becomes an explicit or implicit facet in the process of the designerly shaping of technological mediation, and thereby how people relate to the world.

3.2 Analytic Procedure

Analyses were led by the first author in an iterative process, subsequent to the selection of case studies. Our high-level analytic process begins with human-technology relations articulated in the work of Ihde and Verbeek, and cited and extended within HCI by e.g. [26, 58, 72]. We focus on the human-technology relations given through each case’s design output (e.g., artefact, scenario) and context (e.g., home), and then expand to the particular ML technology employed in the case studies to reconsider how the former are affected. We do so through our initial theoretical schema for human-ML relations (cf. subsection 2.3). In our analyses, we search for designerly “intermediate-level knowledge” [32] that ties in human-technology and model-world relations; which also allows us to investigate in how far established post-phenomenological notions can be advanced. By analyzing the entanglement of the specific ML technology (e.g., goal-driven reinforcement learning) with the given human-technology relation (e.g., embodiment relation), we discern phenomena unaddressed within the design research projects themselves. We then generalize these phenomena under provocative concepts, which we propose be used in design processes to engage ML uncertainty as a design material.

4 ML UNCERTAINTY AS A DESIGN MATERIAL: FOUR ANALYSES

In the following, we investigate four case studies from distinct design research methodologies and unfold phenomena and research questions related to ML uncertainty as a design material.

4.1 Pierce’s *Shifting Lines of Creepiness*

4.1.1 Description. Pierce undertakes a Research-through-Design (RtD) inquiry into smart camera systems, focusing on notions of ‘creepiness’ so as to interrogate how design artefacts navigate creepiness and acceptability [58] of smart cameras in the home. Pierce investigates through RtD how smart home cameras may opaquely or involuntarily “leak data” into system ecologies; house more directly opaque “hole-and-corner” applications that exploit data; and lay the groundwork for future smart services by acting as a “foot-in-the-door.” Pierce proposes speculative design scenarios which transpose the described phenomena into future application domains. The artefacts, smart home security cameras, are predominantly driven by ML image recognition algorithms, such as Convolutional Neural Nets or deep learning variants (e.g., [54]). Pierce’s design considerations construct specific human-technology relations, which we now probe for the (implicit or explicit) presence of ML uncertainty.

4.1.2 Analysis. We focus specifically on the notion of ‘data leakage’ in Pierce’s framing, as it is most distinctly tied to the physical artefact of the camera and the ML algorithm. Depending on how the camera is oriented, recognized objects and people can become processed as data unbeknownst to or without the explicit intention of the owner of the smart camera (e.g., the neighbor’s visitor). This implies one of Ihde’s established human-technology relations, a background relation:

I (- Camera / Home)
Background Relation

In this human-technology relation, the home is shaped as a distinct zone-of-observation, and humans or animals become detections-in-waiting. However, Pierce is explicitly concerned with the agential capacities of smart cameras, i.e. the capacity of ML recognition models to extract patterns from incoming data. The smart camera is less passive than heating appliances or a thermostat, it has a dynamic relation to the surrounding world in its interpretation of incoming data. Therefore, we add this capacity as a distinctly composite, but non-experiential side to the background relation:

I (- Camera / (Recognition → World) / Home)
Composite Background Relation

Yet, this does not quite cover what an ML algorithm within a smart camera *does*. Crucially, it is inferred from a particular relationship between patterns of pixels, e.g. the presence of patterns x and y in an input image indicates the detection of person-is-entering-a-camera-frame (in non-conceptual terms, the model does not ‘know’ this). As these relationships between patterns are learned from data, the relation can be refined by saying that the ‘world’ which the smart camera relates to is already technologically *thematized* to a certain extent: there’s a propensity to recognize specific things in the world, which we indicate with square brackets ([]).

I (- Camera / (Recognition → [World]) / Home)
Composite Background Relation

This notation now brings the specifics of ML in the smart camera home security system to the fore: it is not only the camera’s physical lurking which textures people’s experience of the home. Like a thermostat, people are not constantly aware of the camera, and also like a thermostat, sometimes people are, through notifications or alarms. Yet, people do not have explicit access to the layer of texturing constituted by the direction (i.e., intentionality) of the model towards *latent* recognitions in every frame. Depending on the trained model, some patterns are more probable to be recognized: it is not only data that leaks from the outside-in, as Pierce’s data leakage covers, but furthermore *patterns that leak from the inside-out* (cf. Figure 1). That means, the ML model active within the smart camera has been trained to recognize faces of people within its field of view. It is not capable to distinguish the neighbor’s porch from the porch of its owner. As such, there is a likelihood that a person on the neighbors porch or even a portrait on a delivery truck decal are categorized as a person, depending on how the ML model is inferred.

This is where we find a first manifestation of ML uncertainty as a design material: patterns that leak are also always probabilistic, e.g., there’s an inherent variance in the detection due to model uncertainty. Kendall and Gal note that in their deep learning image

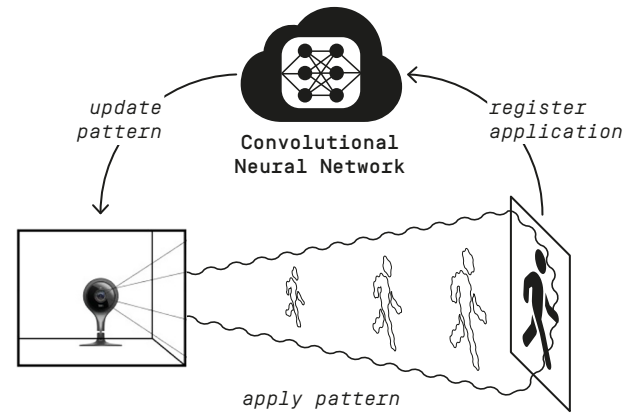


Figure 1: ML uncertainty may not only lead to data leaking from the outside-in, but *patterns leaking from the inside-out*. Camera Illustration © Pierce [58]

segmentation use case, model uncertainty leads to footpaths becoming part of roads [35]. This kind of model uncertainty suggests that ML recognition through smart cameras may *generate* phenomena, rather than merely register them: when ML-driven artefacts recognize or process observations, model uncertainties can lead to patterns being projected into the world, which in turn may solidify the propensity for that pattern to be recognized. Or, falsely associated patterns may lead to new entities becoming significant. While Pierce proposes speculative artefacts to counteract data leakage, we consider the phenomena of *pattern leakage* as a prompt for design research to actively, productively, and also critically engage with. Possible research questions for design research into smart camera pattern leakage may be the following: what novel hybrid entities, such as detections compressing human bodies and suburban surroundings, could become ‘self-evident’ through ML feedback loops? What unforeseen consequences could such uncertain entities have? In what way may they have an effect on end-users? If pattern leakage leads to smart camera patterns that leak from the inside out into the world in a way either nonsensical or odd to humans, what type of domains, products, services could embrace such an uncertain mode of access to the world? Future design research could therefore specifically design for pattern leakage as an exemplary phenomenon of the “ontological surprises” [41] of ML, for example, by engaging in “ludic design” [16] of artefacts interpreting the world by allowing for or exaggerating model uncertainty.

4.2 Wong et al.’s *When BCIs have APIs*

4.2.1 Description. In their design fiction on brain-computer interface (BCI) applications, Wong et al. speculate about the creation of an API for a Google service which lets developers tap into the reading of P300 occurrences in brainwaves. These P300 readings indicate recognition by measuring spikes in brainwave activity, which can be used for e.g. character recognition for input devices. Wong et al. employ design fiction to surface new forms of labor and asymmetries such technologies may engender [79]. Additionally, the researchers also outline distinct human-technology relations

which allow us to probe the surfacing of ML uncertainty. Wong et al. design their fictional API on the basis that the algorithm used to infer the P300 signal was trained on “lab-based stimuli from a controlled environment;” which is representative of various real-world applications and the tensions between training and general application. While the researchers did not specify the algorithm they had in mind, from state of the art research in the field of ML-based P300 recognition (e.g., [59]) we can reasonably assume that a Support Vector Machine (SVM) was involved. SVMs follow the insight-through-opacity approach of ML, as they attempt to form ‘hyperplanes’ in high-dimensional spaces in order to classify (i.e., separate) data [7].

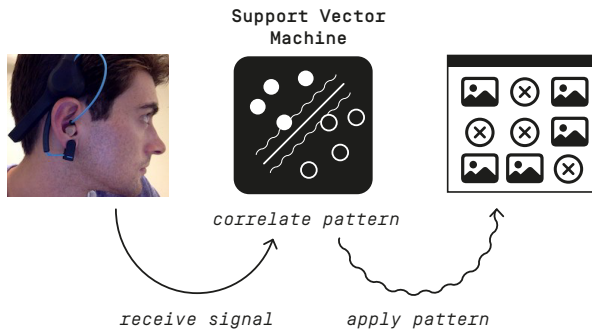


Figure 2: Data uncertainty can lead to pattern leakage of P300-labelled phenomena into everyday experience, affecting human-technology relations beyond the original domain. Photograph © Wong et al. [79]

4.2.2 Analysis. It is within an application based on the P300 API where we encounter ML uncertainty as a potential design material. In this scenario, Wong et al. present the following situation: Crowdworkers are employed with Human Intelligence Tasks (HITs) of content moderation that is using BCIs. This task does not necessitate direct human action from the crowdworkers, but builds on a technical process which registers crowdworker’s brainwaves in relation to potentially harmful images presented on the crowdworking platform (e.g., Amazon’s Mechanical Turk). In the scenario, crowdworkers voice insecurity over these HITs, debating possible P300 recognition errors and doubts over the fitness of their cognition to the task. Nevertheless, in order to successfully complete their HITs, workers employ drastic measures such as holding open their eyelids so as to moderate as much content as possible. At this point, we already encounter a complex, two-fold human-technology relation: Crowdworkers are firstly wearing their BCI, rely on it passively to work correctly and, secondly, looking at the HIT interface which presents them with content moderation tasks. As such, an initial human-technology relation in this scenario may be depicted as follows:

$$(I - BCI) \rightarrow (HITInterface - World)$$

Embodiment Hermeneutic Relation

However, this schema does not yet show how the SVM employed to detect P300 signals is involved. As the P300 signal *model* interprets (\rightarrow) incoming brainwave-data about the world ($[\]$), we may schematize this process as follows:

$$\text{Signal} \rightarrow [\text{World}]$$

We may therefore extend the above scenario that integrates an embodiment relation (crowdworkers are wearing BCIs) and a hermeneutic relation (crowdworkers are observing images for content moderation on an interface) with a further composite relation:

$$(I - BCI / (\text{Signal} \rightarrow [\text{World}])) \rightarrow (\text{HITInterface} - \text{World})$$

Composite Embodiment Hermeneutic Relation

Here, ML uncertainty surfaces in Wong et al.’s research as an implicit design material. The model for P300 stems from the SVM algorithm constructing a hyperplane to separate other signals from the P300 pattern. The latter was, as Wong et al. present, learned from lab-stimuli in controlled environments. Hence, the data uncertainty of the P300 pattern recognition pipeline provides another dimension to the above concept of *pattern leakage*. Through data uncertainty, the SVM hyperplane may be inclusive of non-P300 events, and the P300 pattern (i.e., not the ‘actual’ P300 occurrence, but its model) can thereby ‘leak’. When the trained model is transposed into other domains (e.g., content moderation), overlap between real-world stimuli in phenomenological experience with data uncertainty, then, can lead to P300 patterns leaking into human-technology relations (cf. Figure 2). For example, we can imagine automated filtering of social media newsfeeds based on BCI-driven content moderation. This would make the world legible in particular ways building on the emergence of leaked patterns. These could moderate content based on how a particular BCI algorithm interprets users cognitive functions. Pattern leakage thus describes a phenomenon in Wong et al.’s project, which the researchers have considered as an error generative of labor exploitation. In the spirit of Wong et al.’s research, of using design fiction to unfold the “banality of more probable outcomes” [79] of emerging technology applications, ML uncertainty beyond the notion of error in the form of pattern leakage can be engaged explicitly as a design fiction material. Design fiction can probe for a range of ‘real-world’ events that due to pattern leakage become associated with a particular service or artefact in unplanned, yet *generative* processes. For example, P300 pattern leakage into automated social media filtering can, due to emergent networked media effects, become generative of events such as novel trends, memes, or socio-material practices that are irreducible to human decision-making, but rather are distinct effects of ML data uncertainty in the interpretation of human cognition by a specific algorithm. Design fiction research can therefore use this concept to attend to non-linear, unplanned yet generative effects of future systems.

4.3 Wakkary et al.’s *Morse Things*

4.3.1 Description. Wakkary et al. have developed ‘Morse Things’ as a set of computationally enhanced bowls that can function as everyday household objects while simultaneously communicating among themselves in a human-excluding fashion [72]. Following the material speculation methodology, Morse Things are actual artefacts that are nonetheless counterfactual to expected ways of interaction. They yield a ‘possible world’ in which such artefacts

can exist in their own right [71]. Sets of three (from small to large) were distributed among households by the researchers, and after living with the artefacts for six weeks, a workshop was held in which participants shared stories and constructed scenarios for future technologies. Wakkary et al. thereby sought to probe how thing-centered design methodologies can elucidate design spaces for living with things that are not exclusive to human utility. Specifically, the Morse Things do not register their interactions with humans, but rather are designed to communicate amongst themselves when awake, only sonically emitting Morse code (i.e., short and long beeps) into their surroundings:

“The Morse Things mostly sleep (computationally speaking) and wake at random intervals during the day at least once every eight hours. Upon waking a Morse Thing will send and receive messages to and from other Morse Things in its set. The messages sent by each Morse Thing are in Morse code and simultaneously expressed sonically and broadcasted on Twitter [as cryptic encodings].” [72]

In a subsequent reflection on the project, Oogjes et al. outline further details of the design process, particularly the use of ML to promote a “thing-centered logic”¹ in the Morse Things’ decision on active periods based on how many other Things were successfully communicated with previously [57]. Specifically, a reinforcement learning algorithm was used, taking incoming data to update a plan of action over time (cf. [42, 75]).

4.3.2 Analysis. Through this description we can first sketch the human-technology relation with Morse Things as a background relation:

$I(-\text{Morse Things} / \text{Home})$
Background Relation

While human domestic dwellers experience their home through their daily routines, Morse Things lurk in the background, doing what they do. Morse Things cannot be urged to do what they do, nonetheless their activities form a backdrop to the domestic experience, texturing the dweller’s perceptions of the home, which becomes layered with latent technological activity and opaque beeping. This background relation was often made explicit by participants, who become quite involved with the (real or imagined) activity and purpose of the Morse Things. A typical example is how one of Wakkary et al.’s participants would “continue to keep trying to grab the bowls while they are ‘tweeting’ [...] Maybe I’ll be able to tell them apart eventually.” Therefore, a more accurate rendering of the human-technology relation in this mode is an alterity, a quasi-Other:

$I \rightarrow \text{Morse Thing}(-\text{Home})$
Alterity Relation

While the communication among Morse Things is ‘in itself’ beyond the horizon of our experience, it clearly is a focus even in its opacity, and becomes a feature of the ‘objecthood’ of the Morse Thing for Wakkary et al.’s participants, who actively probe the opaque communication. This tension *constitutes* the ‘gap’ between things and us that Wakkary et al. find to be fruitful as a design space for ambiguity and reflection. Participants of the study echoed

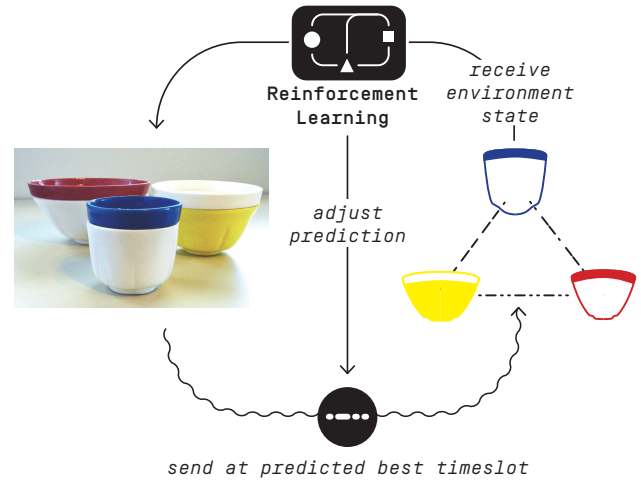


Figure 3: The ideal ‘waking-up’ time is predicted by each Morse Thing based on the amount of communication during prior cycles; with adjustments after every activity. These probabilistic ‘futures’ of the Morse Things creep into human experience. Photograph © Wakkary et al. [72]

this sentiment through statements such as “that’s why I like the idea of something else, let them be themselves. Other stuff is going on that we’re just totally unaware of and it doesn’t matter.”

For our purposes, the crucial aspect is that while each Morse Thing is randomly initiated into ‘waking up’, over time they learn from the most ‘successful’ phases of waking by logging the times when communication with many other Morse Things occurred, and predicting the optimal ‘timeslot’ for waking using ML [57]. Through reinforcement learning algorithms, each artefact updates an internal model based on the prior ‘success’ of its actions in the overall environment of other Morse Things’ activity (Figure 3). With the constantly updating model of an opportune timeslot, ML uncertainty becomes a particularly rich resource that allows for potentially broadening the design research inquiry. First, we can schematize the timeslot model as a machine interpretation of how other Morse Things have previously acted in the world:

$\text{Timeslot} \rightarrow \text{Morse Things}$

The model for a timeslot is based on a prediction of Morse Things waking and communication at a specific time, relating to all Morse Things’ specific interpretation of the ‘world’ as represented by all Morse Things’ activity. Next, we can use this schema to more precisely outline how ML figures in the previously presented human-MorseThings relations:

$I(-\text{Morse Things} / (\text{Timeslot} \rightarrow \text{Morse Things}) / \text{Home})$
Composite Background Relation

$I \rightarrow \text{Morse Thing} / (\text{Timeslot} \rightarrow \text{Morse Things})(-\text{Home})$
Composite Alterity Relation

Our denotation allows us to characterize how ML processes affect the given human-technology relation: it is explicitly related to *time*. Each Morse Thing learns about ‘ideal’ times individually, with the overall assumption that learning will converge in specific

¹<https://doenjaogjes.com/portfolio/morse-things/>, accessed 09/15/2020.

times across all Morse Things. This constant attempt towards convergence invariably affects situated human-technology relations: participants wonder when a Morse Thing will act, and this textures not only the spatial surroundings of a home but also its temporal characteristics. Considering model uncertainty in this scenario, we term the phenomenon of Morse Things' involvement in human experience of time as a *futures creep*: the impact of predictions on situated human experience. Put differently, ML uncertainty allows us to explicitly denote how the *thingly* model of time impacts participants' experience of Morse Things. There is not a singular prediction, but multiple entangled predictions of varying degrees of probability. Various models of a timeslot coalesce around the human experience of beeps and tweets; 'thingly futurings' that characterize and at the same time are irreducible to human-artefact interaction. Expanding on this design space could, for instance, see an expansion of the "animistic" [48] tendencies of Morse Things into even more individualistic thingly futurings. For example, by designing different confidence thresholds for a Morse Thing to decide on a timeslot, or by representing uncertainty in modulating the frequency of beeps based on uncertainty. Futures creep, then, could be used to characterize a *thingly uncertainty* of actual artefacts, further interrogating questions such as: How could material speculation artefacts be more explicitly designed around notions of time and uncertainty? For example, investigating how ML-driven products or services mediate relations to time in specific settings, design researchers could purposely pursue material speculation on intermediary, 'time-keeping' artefacts for futures creep.

4.4 Biggs and Desjardins's Highwater Pants

4.4.1 Description. In their speculative design project, Biggs and Desjardins design and deploy the artefact 'Highwater Pants'—a pair of pants whose legs lengthen or shorten based on whether the wearer is in an area threatened by predicted sea-level rise in the future [5]. The researchers mobilize speculative design to probe human relations to possible futures of climate change, and argue that the Highwater Pants make such futures tangible as they "bend time." The artefact was deployed with cyclists, as these have "embodied and sensorial" knowledge about their environment and specific garments are part of the cycling culture. The Highwater Pants are equipped with a fabric micro-controlling unit, which combines a variety of operations. First, it controls actuation of the pant legs into an up or down state. Second, it houses a GPS module which retrieves the current geographical position of the wearer. Third, it compares the position to a set of polygons on a map, which have been curated by the researchers. Within the polygons, sea-level rise in "30 to 50 years" was predicted using ML linear regression algorithms by the National Oceanic and Atmospheric Administration (cf. [65]). Such algorithms extrapolate tendencies for values to increase linearly based on previous data (cf. [9]); in this case, decades worth of temperature and longitude/latitude data.

4.4.2 Analysis. Riding a bike is, classically, an embodiment relation. As we navigate the world, the bike itself withdraws, becoming part of our extended bodily relationship to our environment. The Highwater Pants, initially, also withdraw into this relation:

$$(I - \text{HighwaterPants} - \text{Bike})\text{World}$$

Embodiment Relation

Based on our prior analyses, the prediction model of the linear regression algorithms can be schematized as an interpretation of an already thematized 'slice' of the world:

$$\text{Prediction} \rightarrow [\text{World}]$$

As Biggs and Desjardins were mindful of the uncertainty of ML predictions, they "padded" their polygons so as to increase the zones of prediction and to give participants more opportunities for reflection. When comparing the wearer's GPS location with the "geofencing" of the polygons, the Highwater Pants pants leg is rolled up or down. Both the Highwater Pants and the bike remain intimately tied to bodily-perceptual experience, but the embodied perception of the environment is textured by likely future states of the same environment. Biggs and Desjardins' empirical findings from deploying pairs of Highwater Pants with experienced local cyclists reflect this intimately. For instance, one participant muses on how "it'd be a totally different experience living here without [the waterfront parks]" after the Highwater Pants indicated that these may disappear. Based on the schemata so far used, we may note the entangling of an embodiment relation (with the bike *and* the Highwater Pants) with a probabilistic prediction of future sea-levels (i.e., a prediction of an already 'thematized' world) as follows:

$$(I - \text{HighwaterPants} / (\text{Prediction} \rightarrow [\text{World}]) - \text{Bike})\text{World}$$

Composite Embodiment Relation

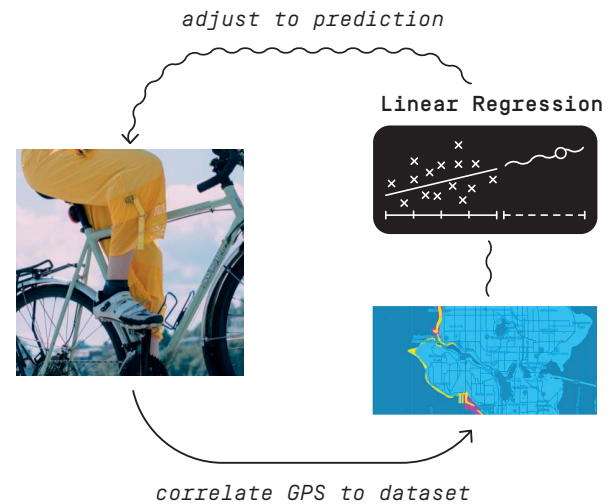


Figure 4: The Highwater Pants correlates the GPS location of the wearer with the dataset for sealevel-rise prediction, and adapts its activity to the relevant prediction. The futures creep mediated by the Highwater Pants affects how bikers relate to possible future states of their environment. Photograph © Ioan Butiu / Biggs and Desjardins [5]

Initially, the above denotation resurfaces the phenomenon of pattern leakage which we discussed with regards to smart cameras and BCIs above: both training data and predicted data are intrinsically uncertain, and the patterns of likely sea-level rise are projected onto the 'real' world. However, the researchers' empirical findings also suggest a further nuance to the Highwater Pants: participants

actively tried to find the boundaries of polygons in order to discover where predictions lose their “time-bending” hold on the present. In their experience, participants noted how the Highwater Pants would be “oscillating between up and down at geofence boundaries, creating a kind of anticipatory sensation in a liminal zone.” Therefore, the initial relation to the prediction model becomes more overt in the form of an alterity relation:

(I – Bike) → HighwaterPants / (Prediction → [World]) (World)
Composite Embodiment Alterity Relation

This complex multi-relationality of human-HighwaterPants interaction shows how the predicted post-sea-level rise world is probed via the Highwater Pants against the backdrop of one’s own bodily-perceptual relations with the environment, while *simultaneously* being part of the human-bike embodiment relation. The Highwater Pants as a quasi-Other becomes, in the words of Biggs and Desjardins, an intermediary “oracle, or translator, speaking for/from an ecology-to-be,” which directly affected participant’s immediate bodily perception of their own future. This adds a further variant to our concept of futures creep: the ecology-to-be is not only an ‘image’ of the future. Rather, the “future present” [11] of inferred models, fluctuating with the data uncertainty of past measurements, becomes an implication of human perception and action in *that* future. The time-bending phenomenon which Biggs and Desjardins explicate is thus not a technological operation triggering a human response (as had been the case with the Morse Things’ futures creep), but rather the injection of prediction into the human phenomenological experience of time (cf. Figure 4). The way participants’ think of the future is co-shaped by the Highwater Pants’ mediation of a prediction, and takes on a particular *shape* as a potentially endangered lifeworld—additionally reflecting questions on whether *my* present is the one leading up to the realization of *this* prediction. The uncertainty of this prediction has a direct effect on how someone *imagines* relations between their present and future, and their capacity in shaping this specific, technologically thematized time. This aspect of futures creep, then, can be used by design researchers to further investigate how design artefacts shape and transform relations between present and future. As futures creep is a phenomenon of ML uncertainty, design researchers can also question: What kind of positionality towards such temporal shapes does the variance of model and data uncertainty provoke, for whom, and towards which ends? Design research methodologies such as speculative design or RtD can use this facet of futures creep for an inquiry into the specifics of time-bending by artefacts, for example investigating how model uncertainty could be deliberately calibrated by users to explore non-anthropocentric notions of time.

5 DISCUSSION

Through our analyses, we derived three provocative concepts for ML uncertainty: thingly uncertainty, pattern leakage and futures creep. We summarize and define these for future research below. While we have separated out case studies into the phenomena we tried to make referable, it should be noted that they are not mutually exclusive and simply offer lenses that bend toward the specific phenomena generated by ML models. Furthermore, our analyses have also surfaced potential research trajectories for more general post-phenomenological ML studies, which we address separately.

5.1 Three Provocative Concepts for Designing with ML Uncertainty

Data and model uncertainty are intrinsic and defining properties of ML. From an engineering or XAI perspective, uncertainty may be seen as a phenomenon to be curbed or explained to frame outputs more unambiguously. But, as we have argued, ML uncertainty is also a promising design research material: as an inherent attribute of contemporary ML, data and model uncertainty speak to the material involvement of ML decision-making with the world. In what follows, we take a further modest step forward and articulate working definitions of our concepts as provocative shorthands for future design research to productively engage with ML uncertainty as a design material. We conclude this section by summarizing their utility for design, and connecting our concepts to the emerging discourse of more-than-human design in HCI.

5.1.1 Thingly Uncertainty. With thingly uncertainty, design researchers can go beyond human uncertainty *about* an artefact and engage how the uncertainty *of* an artefact can become generative of specific, technologically mediated phenomena in the world. As a general concept, we posit that this a particularly powerful shorthand for ML-enabled artefacts. Again, uncertainty in this regard is not a negative attribute, but simply part-and-parcel of the use of probabilistic techniques in ML. Ihde and Verbeek have previously elaborated on a “thingly” [34] or “material hermeneutics” [33, 67]: technological artefacts, through their material properties (e.g., affordances, representations, sensors, actuators), shape how the world becomes legible in specific ways. Thingly uncertainty, however, attributes more precisely the kind of agency that ML-driven artefacts exhibit. Rather than fixed, “scripted” (e.g., [1]) readings, ML-driven artefacts can be much looser and uncertain *about* the legibility of the world. As such they act and adapt within a *continuum* of relations to their environment and the humans that experience them. Thingly uncertainty can help design researchers to more directly explicate the variance of both ML and human, and point to non-normative ways of how the human sees and is seen in human-ML relations: How do the entities, assets and attributes that define humans shift across types (i.e., data and model) and amplitudes (e.g., low or high thresholds) of ML uncertainty? What could a speculative set of norms based on such ML-mediated variances look like?

For example, Wong et al.’s proposal of “infrastructural speculations” [78] can be extended through thingly uncertainty. Considering speculative design and design fiction, Wong et al. propose this concept to more closely define the role of ‘actual’ infrastructures, like the socio-cultural, economic or operational micro- and macro-infrastructures, that need to be in place for an artefact to exist. Often, such infrastructures will involve ML or advanced AI capacities, and thingly uncertainty offers a means to consider the quantitative and qualitative particularities at the ‘joints’ of infrastructural speculations. On the one hand, researchers can investigate how thingly uncertainty will have to be explained, minimized or ignored for the object of research (e.g., a speculative scenario, practice, artefact) to exist. On the other hand, infrastructural speculations also gain a concrete technological dimension of *variance* within lifeworlds: Where, for instance, can pattern leakage or futures creep occur? Who or what needs to perform care, or is affected by, these computational phenomena? Thingly uncertainty can thereby offer

a basis to consider actual phenomena generated and mediated by ML for design research.

5.1.2 Pattern Leakage. This concept describes how ML uncertainty affects the ways in which objects, entities, events and people are recognized in ML-driven systems; noting the propensity of probabilistic patterns to shape the world they are deployed to represent. Through our case studies, we found instances of pattern leakage related both to data and model uncertainty. Design research can appropriate this concept to probe both types of ML uncertainty in the following ways.

When investigating Wong et al.'s *When BCIs have APIs*, we proposed that pattern leakage names how phenomenological experience becomes affected by data uncertainty. Brainwave signals become registered as P300 instances, yet due to data uncertainty it is likely that the world becomes populated with 'surplus' P300 instances. Future design research could therefore pay attention to how data uncertainty leads to a data-driven 'inclusivity' of classifications or predictions that due to intent, oversight, subtlety or opaqueness bleed into specific socio-material constellations. Focussing on pattern leakage due to data uncertainty, design research can more precisely reflect on the thresholds for participation in ML-driven data ecologies. For instance, design fiction could use this concept to avoid assumed linearities in future technological settings, paying close attention to how 'slippage' in classification or prediction can not only lead to breakdowns, but rather be generative in its own right.

In analyzing Pierce's *Shifting Lines of Creepiness*, we noted that smart cameras, through model uncertainty, may actively generate phenomena rather than passively register them. Learned patterns (e.g., a person in a restricted area) may leak onto events in the wild (e.g., a movie poster), and affect how things and humans see the world. We propose that future design research take on this concept to investigate how the intentionality of diverse ML algorithms (e.g., artificial neural networks, SVMs) to read the world in a specific way is generative of distinct patterns leaking into situated human-technology relations. For example, an ML-driven IoT artefact could have various options for its algorithmic functionality. Researchers may actively provoke and design for patterns to leak under different algorithmic choices. Deploying various artefacts, investigations can then begin into whether and how distinct pattern leakage phenomena differentially affect, or become generative of unanticipated, human-technology relations with the ML-driven artefact.

5.1.3 Futures Creep. Futures creep denotes how ML-driven artefacts affect human relations to time through probabilistic, uncertain predictions. The impact of technologies on human conceptions of time is a complex issue of investigation, particularly when considering situated manifestations of meta-concepts such as the "ontologies of times" of specific eras (cf., [47]). However, as ML is fundamentally a technological approach to making predictions (e.g., of classification, recognition, translation) about data, the concept of futures creep provides an opening into this more subtle side of ML for design research.

With regards to model uncertainty, in our analysis of Wakkary et al.'s *Morse Things* we found that futures creep in ML-driven artefacts denotes a specific, human-excluding side: *when* such artefacts do what they do may be not directly correlated with human

experience, yet that is precisely how the artefacts are humanly interpreted to have a specific *character*. In line with Marenko and van Allen's proposal for "animistic design" [48] research, this facet of futures creep can be used to inquire into how different kinds of data uncertainty thresholds for ML-driven artefacts mediate the characteristics that humans attribute to them. For instance, artefacts could be purposely designed to exhibit animistic tendencies by allowing for higher variance in data uncertainty for activity, and researchers may investigate whether such technological decisions translate into mediations of particular artefactual 'characters.'

Model and data uncertainty combined showed a yet more subtle facet of futures creep in Biggs and Desjardins' *Highwater Pants*: the prediction of future sea-level rise was mediated by their speculative design artefact in such a way that it affected participant relations to the present and possible futures, with participants actively probing the range of predictions. The futures creep related to time-bending, i.e. directly affecting the shape of present and future, can be used to actively interrogate how and whether distinct ML algorithms and respective forms of model uncertainty generate specific relations to 'temporal shapes.' Specifically, futures creep may be used in "attending to temporal representations" [38] in design fiction. For example, researchers may interrogate whether a higher degree in variance concerning ML-driven predictions may bring about novel political or civic norms in speculative scenarios, e.g. people choosing to be in loose, variable relations to the future.

Our provocative concepts can serve as a novel conceptual vocabulary to build design artifacts that provocatively engage with ML uncertainty; which allows for in-depth investigations of human-ML relations throughout the design process. Designing for thingly uncertainty with futures creep and pattern leakage can shed light on how human subjectivities become entangled with ML-driven artefacts. This is not only a symbolic or aesthetic exercise, but rather a potentially powerful way of investigating how standard thinking on human-ML relations rely on normative assumptions (e.g., anthropocentric, capitalist, hetero-normative) about the technological as much as the human side of those relations. In this light, and echoing recent calls for more-than-human design (cf., [13, 17, 46, 48, 55]), we propose that our concepts are readymade for research that takes seriously the role of non-human entities within design processes and products.

5.2 Refining ML-driven Technological Mediation: Horizontal Relations

Post-phenomenology's strengths lie in its "methodological post-humanism" [64], in interrogating how technological mediation affects how humans perceive and act in the world. However, ML's thingly uncertainty (i.e. agential capacities for perception, prediction and adaptation) and technical opacity seem to require further in-depth consideration. In our inquiry, the human-technology relation schemata grew ever more complicated and convoluted as we investigated the relationship of model-world and human-technology relations in our case studies; and the phenomenological difference between present artefacts and absent ML technologies was not entirely resolved. Specifically, we can see this in the use of the backslash (/) to indicate both the background of experience *as well* as the workings of ML in the background of an experienced artefact.

Whereas post-phenomenology is mostly focused on technological mediation in the here and now of the phenomenological horizon, the presented phenomena of ML uncertainty trouble this selective focus. As ML algorithms infer models from data representations of the world, they ‘populate’ the world that humans experience with ready-made yet ultimately uncertain entities (e.g., music recommendations, likely traffic jams, people to follow, coasts to disappear). And as everyday phenomenological experience becomes textured by probabilistic models, our capacities for perceiving and acting in such ‘probable’ worlds are shaped accordingly. Thus, while we may not be aware of ML models’ involvement in technological mediation, our ways of relating to the world nonetheless become “imbricated” [23] (i.e., overlapping and co-extensive) with ML technologies. Accordingly, ‘our’ phenomenological horizon, the matters and modes of perceiving, acting and sense-making, is textured by ML’s thingly uncertainty. Investigating such *horizontal relations* can become a promising trajectory for post-phenomenological ML studies, which we briefly sketch as follows. Our preliminary schema mirrors Goodfellow et al.’s use of the tilde operator (\sim) for ML inference [18]:

$$\begin{aligned} I &\sim (\text{Technology} - [\text{World}]) \\ &\quad \text{ML} \sim \text{World} \\ &\quad \text{Horizontal Relations} \end{aligned}$$

Similar to background relations, horizontal relations recede from a specific interface or device that intentional human-technology relations are formed with (e.g., in a hermeneutic relation). And similar to composite relations, horizontal relations feature technologically-exclusive interpretations of the world. But more than that, in horizontal relations, human-technology relations are *embedded* within ML technologies’ specific capacity of being uncertain about the world. Here, model-world relations infer a model (\sim) from the world-as-data ($[\]$), which ‘converges’ with a particular human-technology relation. For example, the Facebook newsfeed algorithm is not only an operation on graph data, but implies a specific human way of relating to particular arrangements of data patterns. The hermeneutic relation that I take up with the Facebook newsfeed interface is therefore ‘textured’ by the uncertainties in predicting that data pattern. Horizontal relations may therefore be characterized as a human-technology *vector* associated with a given ML model, pointing towards a specifically ‘thematized-by-data’ world. The ways in which model and data uncertainties of such relational ‘pointing’ manifest (as e.g. pattern leakage and/or futures creep) in what we perceive, then, are the fundamental concerns of studying horizontal relations. Further investigations involving diverse ML implementations, can use the preliminary schema and concepts we have developed to investigate the established human-technology relations (e.g., Table 1); refining our schema and deriving associated phenomena.

5.3 Limitations of the Inquiry

With this paper, we aimed to make ML more tangible for design research through a post-phenomenological perspective on ML uncertainty. Our analyses have generated concepts to be used by researchers, yet there are important areas of research that we have not touched upon. A challenge to consider is how designing with ML uncertainty may become an ethico-politically reflective practice

beyond a possibly detached aestheticization of ‘glitchy’ ML technologies. Verbeek stresses that post-phenomenology is particularly suited to the anticipation of ethical issues. Future work deploying our developed concepts should therefore be especially attentive to how ML uncertainty may play a role in the “hybrid moral agency” [66] constituted in the relationships of technology and people.

6 CONCLUSION

In this paper, we took a post-phenomenological lens to investigate design research projects for phenomena related to ML uncertainty. From our analyses, we generated three main provocative concepts. *Thingly uncertainty* denotes a general characteristic of ML-driven artefacts: the capacity for relating to the world along a variable continuum. *Pattern leakage* describes the propensity for the learned patterns of ML models to be projected into the world. *Futures creep* names the mediation of particular relations to the present and future of ML-driven artefacts. All concepts offer distinct opportunities for design research to engage ML-driven technological mediation. We argue that these concepts offer a promising foothold for design research of ML technologies, which has been a difficulty for the field. Additionally, we noted that the concepts derived from our case studies can also feed back into post-phenomenological ML studies, adding a more precise description of how human intentionality co-extends and overlaps with ML capacities in the form of *horizontal relations*. As such, we offer a modest step forward for design research and post-phenomenology to engage with ML.

ACKNOWLEDGMENTS

We thank Peter-Paul Verbeek, Michael Nagenborg and Christoph Kinkeldey for input and inspiration; Richmond Wong and Doenja Oogjes for clarification regarding their work; and our reviewers for their invaluable feedback. The authors of our case studies retain copyright for image material used as stated.

REFERENCES

- [1] Madeleine Akrich. 1992. The De-Description of Technical Objects. In *Shaping technology/Building Society: Studies in Sociotechnical Change*, Wiebe E. Bijker and John Law (Eds.).
- [2] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for Human-AI Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, 3:1–3:13. <https://doi.org/10.1145/3290605.3300233> event-place: Glasgow, Scotland Uk.
- [3] Josh Andres, Christine T. Wolf, Sergio Cabrero Barros, Erick Oduor, Rahul Nair, Alexander Kjærsum, Anders Bech Tharsgaard, and Bo Schwartz Madsen. 2020. Scenario-based XAI for Humanitarian Aid Forecasting. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (CHI EA '20)*. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3334480.3382903>
- [4] Arne Berger, William Odom, Michael Storz, Andreas Bischof, Albrecht Kurze, and Eva Hornecker. 2019. The Inflatable Cat: Idiosyncratic Ideation of Smart Objects for the Home. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, Glasgow, Scotland Uk, 1–12. <https://doi.org/10.1145/3290605.3300631>
- [5] Heidi R. Biggs and Audrey Desjardins. 2020. High Water Pants: Designing Embodied Environmental Speculation. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, Honolulu, HI, USA, 1–13. <https://doi.org/10.1145/3313831.3376429>
- [6] Christopher M. Bishop. 2006. *Pattern recognition and machine learning*. Springer, New York.
- [7] Christopher J C Burges. 1998. A Tutorial on Support Vector Machines for Pattern Recognition. *Data Mining and Knowledge Discovery* 2 (1998), 121–167.

- [8] Jenna Burrell. 2016. How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society* 3, 1 (Jan. 2016), 205395171562251. <https://doi.org/10.1177/2053951715622512>
- [9] T. Tony Cai and Peter Hall. 2006. Prediction in functional linear regression. *The Annals of Statistics* 34, 5 (2006), 2159–2179. Publisher: Institute of Mathematical Statistics.
- [10] Graham Dove, Kim Halskov, Jodi Forlizzi, and John Zimmerman. 2017. UX Design Innovation: Challenges for Working with Machine Learning As a Design Material. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 278–288. <https://doi.org/10.1145/3025453.3025739> event-place: Denver, Colorado, USA.
- [11] Elena Esposito. 2013. The structures of uncertainty: performativity and unpredictability in economic operations. *Economy and Society* 42, 1 (Feb. 2013), 102–129. <https://doi.org/10.1080/03085147.2012.687908>
- [12] Craig R. Fox and Gülden Ülkümen. 2011. Distinguishing two dimensions of uncertainty. *Perspectives on thinking, judging, and decision making* (2011), 21–35.
- [13] Christopher Frauenberger. 2019. Entanglement HCI The Next Wave? *ACM Transactions on Computer-Human Interaction (TOCHI)* (Nov. 2019). <https://dl.acm.org/doi/abs/10.1145/3364998>
- [14] Bill Gaver and John Bowers. 2012. Annotated portfolios. *Interactions* 19, 4 (July 2012), 40–49. <https://doi.org/10.1145/2212877.2212889>
- [15] William W. Gaver, Andrew Boucher, Sarah Pennington, and Brendan Walker. 2004. Cultural Probes and the Value of Uncertainty. *Interactions* 11, 5 (Sept. 2004), 53–56. <https://doi.org/10.1145/1015530.1015555>
- [16] William W. Gaver, John Bowers, Andrew Boucher, Hans Gellerson, Sarah Pennington, Albrecht Schmidt, Anthony Steed, Nicholas Villars, and Brendan Walker. 2004. The Drift Table: Designing for Ludic Engagement. In *CHI '04 Extended Abstracts on Human Factors in Computing Systems (CHI EA '04)*. ACM, New York, NY, USA, 885–900. <https://doi.org/10.1145/985921.985947> event-place: Vienna, Austria.
- [17] Elisa Giaccardi, Nazli Cila, Chris Speed, and Melissa Caldwell. 2016. Thing Ethnography: Doing Design Research with Non-Humans. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16)*. Association for Computing Machinery, Brisbane, QLD, Australia, 377–387. <https://doi.org/10.1145/2901790.2901905>
- [18] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>
- [19] Google PAIR. 2019. People + AI Research. <https://pair.withgoogle.com>
- [20] Ian Hacking. 1990. *The Taming of Chance* (1 edition ed.). Cambridge University Press, Cambridge England ; New York.
- [21] Alon Halevy, Peter Norvig, and Fernando Pereira. 2009. The Unreasonable Effectiveness of Data. *IEEE Intelligent Systems* 24, 2 (March 2009), 8–12. <https://doi.org/10.1109/MIS.2009.36>
- [22] Leif Hancox-Li. 2020. Robustness in machine learning explanations: does it matter?. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. Association for Computing Machinery, Barcelona, Spain, 640–647. <https://doi.org/10.1145/3351095.3372836>
- [23] Mark B. N. Hansen. 2015. *Feed-forward : on the future of twenty-first-century media*. University of Chicago Press, Chicago ; London.
- [24] Steve Harrison, Phoebe Sengers, and Deborah Tatar. 2011. Making epistemological trouble: Third-paradigm HCI as successor science. *Interacting with Computers* 23, 5 (Sept. 2011), 385–392. <https://doi.org/10.1016/j.intcom.2011.03.005>
- [25] Sabrina Hauser, Doenja Oogjes, Ron Wakkary, and Peter-Paul Verbeek. 2018. An Annotated Portfolio on Doing Postphenomenology Through Research Products. In *Proceedings of the 2018 Designing Interactive Systems Conference (DIS '18)*. ACM, New York, NY, USA, 459–471. <https://doi.org/10.1145/3196709.3196745>
- [26] Sabrina Hauser, Ron Wakkary, William Odom, Peter-Paul Verbeek, Audrey Desjardins, Henry Lin, Matthew Dalton, Markus Schilling, and Gijs de Boer. 2018. Deployments of the Table-non-table: A Reflection on the Relation Between Theory and Things in the Practice of Design Research. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, 201:1–201:13. <https://doi.org/10.1145/3173574.3173775>
- [27] Patrick Hebron. 2016. *Machine Learning for Designers*. Technical Report. O'Reilly Media, Inc. <https://www.oreilly.com/library/view/machine-learning-for/9781491971444/>
- [28] Martin Heidegger. 2010. *Being and time*. State University of New York Press, Albany. OCLC: ocn608297834.
- [29] Drew Hemment, Ruth Aylett, Vaishak Belle, Dave Murray-Rust, Ewa Luger, Jane Hillston, Michael Rovatos, and Frank Broz. 2019. Experiential AI. *AI Matters* 5, 1 (April 2019), 25–31. <https://doi.org/10.1145/3320254.3320264>
- [30] Fred Hohman, Andrew Head, Rich Caruana, Robert DeLine, and Steven M. Drucker. 2019. Gamut: A Design Probe to Understand How Data Scientists Understand Machine Learning Models. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, 579:1–579:13. <https://doi.org/10.1145/3290605.3300809> event-place: Glasgow, Scotland Uk.
- [31] Stacy Hsueh, Sarah Fdili Alaoui, and Wendy E. Mackay. 2019. Understanding Kinaesthetic Creativity in Dance. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, 511:1–511:12. <https://doi.org/10.1145/3290605.3300741> event-place: Glasgow, Scotland Uk.
- [32] Kristina Höök and Jonas Löwgren. 2012. Strong concepts: Intermediate-level knowledge in interaction design research. *ACM Transactions on Computer-Human Interaction* 19, 3 (Oct. 2012), 23:1–23:18. <https://doi.org/10.1145/2362364.2362371>
- [33] Don Ihde. 1990. *Technology and the Lifeworld: From Garden to Earth*. Indiana University Press, Bloomington and Indianapolis.
- [34] Don Ihde. 1997. Thingly hermeneutics/Technoconstructions. In *Hermeneutics and the Natural Sciences*, Robert P. Crease (Ed.). Springer Netherlands, Dordrecht, 111–123. https://doi.org/10.1007/978-94-009-0049-3_7
- [35] Alex Kendall and Yarin Gal. 2017. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? *arXiv:1703.04977 [cs]* (Oct. 2017). arXiv: 1703.04977.
- [36] Christoph Kinkeldey, Tim Korjakow, and Jesse Josua Benjamin. 2019. Towards Supporting Interpretability of Clustering Results with Uncertainty Visualization. The Eurographics Association. <https://doi.org/10.2312/trvis.20191183>
- [37] Armen Der Kiureghian and Ove Ditlevsen. 2009. Aleatory or epistemic? Does it matter? *Structural Safety* 31, 2 (March 2009), 105–112. <https://doi.org/10.1016/j.strusafe.2008.06.020>
- [38] Sandjar Kozubaev, Chris Elsdon, Noura Howell, Marie Louise Juul Søndergaard, Nick Merrill, Britta Schulte, and Richmond Y. Wong. 2020. Expanding Modes of Reflection in Design Futuring. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3313831.3376526>
- [39] Yongchan Kwon, Joong-Ho Won, Beom Joon Kim, and Myunghee Cho Paik. 2018. Uncertainty quantification using Bayesian neural networks in classification: Application to ischemic stroke lesion segmentation. Amsterdam, The Netherlands, 13.
- [40] Gierad Laput, Yang Zhang, and Chris Harrison. 2017. Synthetic Sensors: Towards General-Purpose Sensing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, Denver, Colorado, USA, 3986–3999. <https://doi.org/10.1145/3025453.3025773>
- [41] Lucian Leahu. 2016. Ontological Surprises: A Relational Perspective on Machine Learning. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16)*. ACM, New York, NY, USA, 182–186. <https://doi.org/10.1145/2901790.2901840>
- [42] Jae Won Lee. 2001. Stock price prediction using reinforcement learning. In *ISIE 2001. 2001 IEEE International Symposium on Industrial Electronics Proceedings (Cat. No. 01TH8570)*, Vol. 1. 690–695 vol.1. <https://doi.org/10.1109/ISIE.2001.931880>
- [43] Q. Vera Liao, Daniel Gruen, and Sarah Miller. 2020. Questioning the AI: Informing Design Practices for Explainable AI User Experiences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, Honolulu, HI, USA, 1–15. <https://doi.org/10.1145/3313831.3376590>
- [44] Youn-kyung Lim, Sang-Su Lee, and Da-jung Kim. 2011. Interactivity attributes for expression-oriented interaction design. *International Journal of Design* 5, 3 (2011).
- [45] Zachary C. Lipton. 2016. The Mythos of Model Interpretability. *arXiv:1606.03490 [cs, stat]* (June 2016). arXiv: 1606.03490.
- [46] Susan Loh, Marcus Foth, Glenda Amayo Caldwell, Veronica Garcia-Hansen, and Mark Thomson. 2020. A more-than-human perspective on understanding the performance of the built environment. *Architectural Science Review* 63, 3-4 (July 2020), 372–383. <https://doi.org/10.1080/00038628.2019.1708258> Publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/00038628.2019.1708258>
- [47] Davor Löffler. 2018. Distributing Potentiality. Post-capitalist Economies and the Generative Time Regime. *Identities: Journal for Politics, Gender and Culture* 15, 1-2 (2018), 8–44.
- [48] Betti Marenko and Philip van Allen. 2016. Animistic design: how to reimagine digital interaction between the human and the nonhuman. *Digital Creativity* 27, 1 (Jan. 2016), 52–70. <https://doi.org/10.1080/14626268.2016.1145127>
- [49] Warren S. McCulloch and Walter Pitts. 1943. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics* 5, 4 (Dec. 1943), 115–133. <https://doi.org/10.1007/BF02478259>
- [50] Dan McQuillan. 2018. Data Science as Machinic Neoplatonism. *Philosophy & Technology* 31, 2 (June 2018), 253–272. <https://doi.org/10.1007/s13347-017-0273-3>
- [51] Nick Merrill and John Chuang. 2018. From Scanning Brains to Reading Minds: Talking to Engineers about Brain-Computer Interface. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*. ACM Press, Montreal QC, Canada, 1–11. <https://doi.org/10.1145/3173574.3173897>
- [52] Nick Merrill, John Chuang, and Coye Cheshire. 2019. Sensing is Believing: What People Think Biosensors Can Reveal About Thoughts and Feelings. In *Proceedings of the 2019 on Designing Interactive Systems Conference - DIS '19*. ACM Press, San Diego, CA, USA, 413–420. <https://doi.org/10.1145/3322276.3322286>
- [53] Tim Miller. 2017. Explanation in Artificial Intelligence: Insights from the Social Sciences. *arXiv:1706.07269 [cs]* (June 2017). arXiv: 1706.07269.
- [54] Masaki Nakada, Han Wang, and Demetri Terzopoulos. 2017. AcFR: Active Face Recognition Using Convolutional Neural Networks. In *2017 IEEE Conference on*

- Computer Vision and Pattern Recognition Workshops (CVPRW)*. 35–40. <https://doi.org/10.1109/CVPRW.2017.11> ISSN: 2160-7516.
- [55] Iohanna Nicenboim, Elisa Giaccardi, Marie Louise Juul Søndergaard, Anuradha Venugopal Reddy, Yolande Strengers, James Pierce, and Johan Redström. 2020. More-Than-Human Design and AI: In Conversation with Agents. In *Companion Publication of the 2020 ACM Designing Interactive Systems Conference (DIS '20 Companion)*. Association for Computing Machinery, New York, NY, USA, 397–400. <https://doi.org/10.1145/3393914.3395912>
- [56] Doenja Oogjes, William Odom, and Pete Fung. 2018. Designing for an other Home: Expanding and Speculating on Different Forms of Domestic Life. In *Proceedings of the 2018 Designing Interactive Systems Conference (DIS '18)*. Association for Computing Machinery, Hong Kong, China, 313–326. <https://doi.org/10.1145/3196709.3196810>
- [57] Doenja Oogjes, Ron Wakkary, Henry Lin, and Omid Alemi. 2020. Fragile! Handle with Care: The Morse Things. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference (DIS '20)*. Association for Computing Machinery, Eindhoven, Netherlands, 2149–2162. <https://doi.org/10.1145/3357236.3395584>
- [58] James Pierce. 2019. Smart Home Security Cameras and Shifting Lines of Creepiness: A Design-Led Inquiry. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, 45:1–45:14. <https://doi.org/10.1145/3290605.3300275> event-place: Glasgow, Scotland UK.
- [59] A. Rakotomamonjy, V. Guigue, G. Mallet, and V. Alvarado. 2005. Ensemble of SVMs for Improving Brain Computer Interface P300 Speller Performances. In *Artificial Neural Networks: Biological Inspirations – ICANN 2005 (Lecture Notes in Computer Science)*, Włodzisław Duch, Janusz Kacprzyk, Erkki Oja, and Sławomir Zadrozny (Eds.). Springer, Berlin, Heidelberg, 45–50. https://doi.org/10.1007/11550822_8
- [60] Johan Redström. 2005. On Technology as Material in Design. *Design Philosophy Papers* 3, 2 (June 2005), 39–54. <https://doi.org/10.2752/144871305X13966254124275>
- [61] Johan Redström and Heather Wiltse. 2015. Press Play: Acts of defining (in) fluid assemblages. *Nordes* 1, 6 (May 2015). <http://www.nordes.org/opj/index.php/n13/article/view/432>
- [62] Robert Rosenberger and Peter-Paul Verbeek (Eds.). 2015. *Postphenomenological Investigations: Essays on Human-Technology Relations*. Lexington Books, Lanham, MD.
- [63] Frank Rosenblatt. 1958. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review* 65, 6 (1958), 386–408. Publisher: American Psychological Association.
- [64] Tamar Sharon. 2013. *Human nature in an age of biotechnology: the case for mediated posthumanism*. Springer Berlin Heidelberg, New York, NY.
- [65] William V. Sweet, Greg Dusek, Jayantha Obeysekera, and John J. Marra. 2018. *Patterns and Projections of High Tide Flooding Along the U.S. Coastline Using a Common Impact Threshold*. Technical Report 086. National Oceanic and Atmospheric Administration, Maryland, USA.
- [66] Peter-Paul Verbeek. 2006. Materializing Morality: Design Ethics and Technological Mediation. *Science, Technology, & Human Values* 31, 3 (May 2006), 361–380. <https://doi.org/10.1177/0162243905285847>
- [67] Peter-Paul Verbeek. 2006. *What things do : philosophical reflections on technology, agency, and design*. Penn State Press, University Park.
- [68] Peter-Paul Verbeek. 2008. Cyborg intentionality: Rethinking the phenomenology of human–technology relations. *Phenomenology and the Cognitive Sciences* 7, 3 (Sept. 2008), 387–395. <https://doi.org/10.1007/s11097-008-9099-x>
- [69] Peter-Paul Verbeek. 2015. Beyond Interaction: A Short Introduction to Mediation Theory. *Interactions* 22, 3 (April 2015), 26–31. <https://doi.org/10.1145/2751314>
- [70] Peter-Paul Verbeek. 2015. Designing the Public Sphere: Information Technologies and the Politics of Mediation. In *The Onlife Manifesto: Being Human in a Hyper-connected Era*, Luciano Floridi (Ed.). Springer International Publishing, Cham, 217–227. https://doi.org/10.1007/978-3-319-04093-6_21
- [71] Ron Wakkary, William Odom, Sabrina Hauser, Garnet Hertz, and Henry Lin. 2015. Material Speculation: Actual Artifacts for Critical Inquiry. *Aarhus Series on Human Centered Computing* 1, 1 (Oct. 2015), 12. <https://doi.org/10.7146/aahcc.v1i1.21299>
- [72] Ron Wakkary, Doenja Oogjes, Sabrina Hauser, Henry Lin, Cheng Cao, Leo Ma, and Tijs Duel. 2017. Morse Things: A Design Inquiry into the Gap Between Things and Us. In *Proceedings of the 2017 Conference on Designing Interactive Systems (DIS '17)*. ACM, New York, NY, USA, 503–514. <https://doi.org/10.1145/3064663.3064734>
- [73] Danding Wang, Qian Yang, Ashraf Abdul, and Brian Y. Lim. 2019. Designing Theory-Driven User-Centric Explainable AI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, 601:1–601:15. <https://doi.org/10.1145/3290605.3300831> event-place: Glasgow, Scotland UK.
- [74] Norbert Wiener. 1961. *Cybernetics or control and communication in the animal and the machine* (2. ed ed.). MIT Press, Cambridge, Mass. OCLC: 990311734.
- [75] Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* 8 (1992), 229–256.
- [76] Christine T. Wolf. 2019. Explainability scenarios: towards scenario-based XAI design. In *Proceedings of the 24th International Conference on Intelligent User Interfaces (IUI '19)*. Association for Computing Machinery, New York, NY, USA, 252–257. <https://doi.org/10.1145/3301275.3302317>
- [77] Jason Shun Wong. 2018. Design and fiction: imagining civic AI. *Interactions* 25, 6 (Oct. 2018), 42–45. <https://doi.org/10.1145/3274568>
- [78] Richmond Y. Wong, Vera Khovanskaya, Sarah E. Fox, Nick Merrill, and Phoebe Sengers. 2020. Infrastructural Speculations: Tactics for Designing and Interrogating Lifeworlds. (April 2020). <https://doi.org/10.1145/3313831.3376515>
- [79] Richmond Y. Wong, Nick Merrill, and John Chuang. 2018. When BCIs Have APIs: Design Fictions of Everyday Brain-Computer Interface Adoption. In *Proceedings of the 2018 Designing Interactive Systems Conference (DIS '18)*. ACM, New York, NY, USA, 1359–1371. <https://doi.org/10.1145/3196709.3196746>
- [80] Qian Yang, Nikola Banovic, and John Zimmerman. 2018. Mapping Machine Learning Advances from HCI Research to Reveal Starting Places for Design Innovation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, 130:1–130:11. <https://doi.org/10.1145/3173574.3173704> event-place: Montreal QC, Canada.
- [81] Qian Yang, Alex Scuito, John Zimmerman, Jodi Forlizzi, and Aaron Steinfeld. 2018. Investigating How Experienced UX Designers Effectively Work with Machine Learning. In *Proceedings of the 2018 Designing Interactive Systems Conference (DIS '18)*. Association for Computing Machinery, New York, NY, USA, 585–596. <https://doi.org/10.1145/3196709.3196730>
- [82] Qian Yang, Jina Suh, Nan-Chen Chen, and Gonzalo Ramos. 2018. Grounding Interactive Machine Learning Tool Design in How Non-Experts Actually Build Models. In *Proceedings of the 2018 Designing Interactive Systems Conference (DIS '18)*. ACM, New York, NY, USA, 573–584. <https://doi.org/10.1145/3196709.3196729> event-place: Hong Kong, China.