Shifting Ambiguity, Collapsing Indeterminacy: Designing with Data as Baradian Apparatus

COURTNEY N. REED*, Institute for Digital Technologies, Loughborough University London, United Kingdom
ADAN L. BENITO*, Centre for Digital Music, Queen Mary University of London, United Kingdom
FRANCO CASPE*, Centre for Digital Music, Queen Mary University of London, United Kingdom
ANDREW P. MCPHERSON*, Dyson School of Design Engineering, Imperial College London, United Kingdom

This paper examines how digital systems designers distil the messiness and ambiguity of the world into concrete data that can be processed by computing systems. Using Karen Barad's agential realism as a guide, we explore how data is fundamentally entangled with the tools and theories of its measurement. We examine data-enabled artefacts acting as Baradian apparatuses: they do not exist independently of the phenomenon they seek to measure, but rather collect and co-produce observations from within their entangled state: the phenomenon and the apparatus co-constitute one another. Connecting Barad's quantum view of indeterminacy to the prevailing HCI discourse on the opportunities and challenges of ambiguity, we suggest that the very act of trying to stabilise a conceptual interpretation of data within an artefact has the paradoxical effect of amplifying and shifting ambiguity in interaction. We illustrate these ideas through three case studies from our own practices of designing digital musical instruments (DMIs). DMIs necessarily encode symbolic and music-theoretical knowledge as part of their internal operation, even as conceptual knowledge is not their intended outcome. In each case, we explore the nature of the apparatus, what phenomena it co-produces, and where the ambiguity lies to suggest approaches for design using these abstract theoretical frameworks.

CCS Concepts: • **Human-centered computing → Interaction design theory, concepts and paradigms**; *HCI theory, concepts and models*.

Additional Key Words and Phrases: ambiguity, research through design, agential realism, entanglement, mapping, digital musical instruments

1 INTRODUCTION

To live in the world is to engage in a constant process of meaning-making amidst ambiguity [51]. Sensory stimuli do not come pre-labelled with objective, universally-agreed meanings. Instead, we navigate and make meaning of the world through embodied lived experience, both conceptual and pre-reflective, individual and social [41, 78, 84, 103, 116]. HCI research has long been interested in how the messiness of the world can make its

---

[*] All authors have contributed equally to this research.

Authors' addresses: Courtney N. Reed, c.n.reed@lboro.ac.uk, Institute for Digital Technologies, Loughborough University London, Queen Elizabeth Olympic Park, The Broadcast Centre Here East, Lesney Ave, London, United Kingdom, E20 3BS; Adan L. Benito, a.benitotemprano@qmul.ac.uk, Centre for Digital Music, Queen Mary University of London, Mile End Road, London, United Kingdom, E1 4NS; Franco Caspe, f.s.caspe@qmul.ac.uk, Centre for Digital Music, Queen Mary University of London, Mile End Road, London, United Kingdom, E1 4NS; Andrew P. McPherson, andrew.mcpherson@imperial.ac.uk, Dyson School of Design Engineering, Imperial College London, South Kensington Campus, London, United Kingdom, SW7 2AZ.

ACM 1557-7325/2024/8-ART
https://doi.org/10.1145/3689043

way into design processes of seemingly rigid digital systems [20, 43, 44, 105]. Research through Design (RtD) in particular seeks new forms of knowledge through design practices themselves [11, 46], often following winding paths [48] rather than the neat goal-directed explorations that science and engineering at least purport to offer. However, worldviews of positivism and representationalism received from the sciences always stand ready to influence design thinking, particularly when working with data and computation.

This paper examines how data comes to be created and interpreted in the design of digital artefacts. Building on the *diffractive* account by Sanches et al. [105] which takes inspiration from Barad's philosophy of agential realism [8], data can only exist in relation to a far-reaching entanglement of matter, context, history, and frameworks of conceptual knowledge; a 'raw' form cannot exist [47]. Designing digital artefacts, including interactive systems, which encode or process data (hereafter referred to as *data-enabled artefacts*) thus always involves an active choice of model or frame for data collection and processing. This framing process is never neutral, but where and how it happens is often overlooked.

To distil the messiness of the world into a form that can be processed by symbolic digital systems, designers often organise their thinking into conceptual spaces, either deliberately or implicitly [40]. A data-enabled artefact might encode one conceptual space describing the expected input from the user, perhaps in the form of sensor data [15], and another space concerning the possible outputs or behaviours of the artefact. This pre-conceptualisation then permits a more deterministic engineering-related process of *mapping* [62] by which numerical associations are established between the two conceptual spaces. Mapping decisions are highly important to the identity of an artefact, but so too are the contexts and concepts that underlie them, which are crucial to co-constructing the meaning of the data in the first place.

We are interested in how these conceptualisations come to exist and also in how they affect the patterns of meaning and ambiguity that emerge when a data-enabled artefact is placed in a context of use. Building on Gaver et al.'s influential call for making ambiguity a resource for design [44] and drawing an analogy from Giaccardi's three shifts in data-enabled RtD practice [46], we query how data-enabled artefacts enact *ambiguity shifts*, in which patterns of ambiguity are variably altered, collapsed, or created anew.

In Section 3 we propose two related mechanisms for this ambiguity shift. The first mechanism is a human-centred epistemological argument about how design makes some concepts explicit while leaving others implicit, and that explicit concepts tend to draw the attention of the designer (and user). The second mechanism ventures more deeply into agential realism, examining the data-enabled artefact as a Baradian *apparatus* [7, 8]. The apparatus does not exist independently of the phenomenon it seeks to measure, but rather collects and co-produces observations from within an entangled state; the phenomenon and the apparatus co-constitute one another [6]. The act of observing collapses a fundamental indeterminacy in the world into particular fixed measurements already laden with meaning through a process Barad calls the *agential cut*. In this second telling, the explicit concepts encoded in a data-enabled artefact are not pre-existing abstract quantities that have been called to human attention; their very existence is inseparable from the apparatus of measurement. Our arguments about data in design resonate with Sanches et al. [105], but our focus is less on diffraction as a knowledge production process and more on querying the identity of the apparatus, the agential cuts it enacts, and the implications for ambiguity as a design resource.

Throughout the paper we draw examples from the area of digital musical instrument (DMI) design, including three case studies from our own design practices. DMIs are data-enabled artefacts: they necessarily encode symbolic and music-theoretical knowledge as part of their internal operation [77, 81, 117], but unlike neighbouring research in music information retrieval (MIR), conceptual knowledge is not the intended outcome [104]. This sets up a delicate dance between scientific and artistic modes of thinking [49, 69], where the ambiguity and indeterminacy of musical practice is transiently distilled into symbolic conceptual spaces and mappings in order to be projected back out again in a way that retains fidelity to the aesthetic, cultural and social goals of the designer and their constellation of collaborators [128].

In the tradition of foregrounding the particulars in design research [24, 114], our case studies are not intended as instantiations of a broader methodological framework. Rather, in Section 4 we seek to narrate how we grapple with the unavoidable tensions between (un)certainty, ambiguity and indeterminacy in working with data, while highlighting in Section 5 how abstract theoretical frameworks can ultimately suggest real and actionable activities for design.

2 BACKGROUND

To frame our design investigation, we first discuss ambiguity in design and the history of its study as a factor in the design process. Then, we introduce literature around Barad's agential realism [8] and describe the process of making agential cuts to explore how materiality, context and the designer's positionality are entangled with their mappings through the apparatus. Finally, we present the design of musical instruments and related work to centre DMI research as a noteworthy subfield for the exploration of ambiguity in design.

2.1 Ambiguity in Design

Gaver et al. [44] famously outline how ambiguity is a resource for design and establish a taxonomy of ambiguity. There are three general categories of ambiguity, depending on where in the design process it arises [44]. Ambiguity can come from the artefact or apparatus itself and the way it represents information; this *ambiguity of information* arises when the results and interpretations of conceptual mappings are not clearly or at all prescribed. *Ambiguity of context* arises when the particular mapped contexts or conceptual spaces provide unexpected or novel meaning to apparatuses, and *ambiguity of relationship* when there is conflict between the understanding of the designer and the end user due to the unique meaning-making each agent will make.

Ambiguity is in no way inherently "bad" for design. Gaver et al. recognise that making use of ambiguity inherent in all systems is a valuable strategy for establishing a connection with users in a non-prescriptive manner. This approach embraces the apparatus's own limitations to distil the ambiguity present in the world. It then enables users to construct their own interpretations, find personal meaning or invite them to reflection when engaging with artefacts [44]. Ambiguity in communication between human and non-human agents can be intentionally manipulated as a resource for negotiating lived experiences and reaching common ground [103]. Ambiguity can also be useful when working with a lack of data or data which cannot be well-expressed through the quantitative distillation needed for symbolic manipulation or mapping. Relationality and the way information is disambiguated, either by the designer in a system's representation and mapping of data or by users in interpretation of the designer's mappings, determine the effectiveness of this communication in a system. Boehner [20], for example, presents approaches for using ambiguity as a design resource to generate *space for interpretations* for Personal Communication Systems users, explaining how these cannot disambiguate a situation, but instead can constrain and shape the set of possible interpretations users may have about it. The *ambiguity of information* provides a solution to the impossibility to create a one-size-fits-all design for personal communication. The designer assists the interpretation and use of the system through its constraints, but it is ultimately left up to each end-user.

Although facts and situations do not have unique interpretations, designers can create opportunities for meaning in ambiguity to be negotiated with others. Howell analyses this in situated conversations between two friends, both wearing a thermochromic t-shirt [56, 59]. This garment features patterns that change color when its wearer's skin conductance increases, which acts as marker for sudden arousal. Sense-making in this case strongly relies on the context of the conversation and the background and relationship of the participants themselves when interacting with and interpreting the artefact; in this case, the participants learned to exploit the ambiguous change of color to convey a range of emotions to the other party, using the biosignal display as a "social cue".

This *ambiguity of context* depends on the connection between the participants and their relationship as friends, with the artefact's role being defined as a factor within a pre-existing space.

Designers can open the same negotiation between user and artefacts, and this meaning can be ultimately projected to be jointly interpreted by their audiences. Devendorf [33] studies how clothing-based ambiguous displays can function for personal clothing as meaningful communicators. Participants in the study envisioned usages for the displays to foster social encounters or support irony or reflexive experiences; in this way, the designer has left an open space for interpretation of the artefact, wherein users take the role of mediators between their artefact's ambiguity and the public.

Finally, ambiguity can be used in design to negotiate relationships between users and artefacts [124]. Wakkary designed a set of counterfactual bowls that communicate between themselves. This *ambiguity of relationship* [44] can add dimensionality and complexity to the notion of artefact's existence showing how relationships do not need to be clear in order for us to maintain attachment to things. In the same way, Wakkary also explores questions about presence of artefacts through defamiliarisation by bringing new ambiguous functions to everyday objects [125].

2.2 Data-Enabled Artefacts as Apparatuses

Simply put, we are unable to design independently from our biases. Our conceptualisation of the world arises from our understanding, either conscious or not, of societal roles, identity, equity, socio-cultural background, and physical environments around us that we experience through our individual bodies [10, 41]. This entanglement prevents us from being completely unambiguous about ambiguity as designers. In design, the individual perspectives and assumptions of the designer are codified and influenced by the larger world context. Ambiguity, like data, falls victim to the "commodity fiction" [47]; Ambiguity is not a universal, determinate quantity that can be manipulated directly in a design. Rather, the meaning and form of ambiguity are emergent from the myriad of entangled factors present in the design process of an artefact and its context of use.

2.2.1 Karen Barad's Agential Realism and Diffraction. In order to examine these entangled factors, we adopt a Baradian framing [8] to identify the roles of different agents in the designer's decisions and assumptions through the design process. Barad draws on Niels Bohr's quantum philosophy-physics to move beyond entrenched debates about nature versus culture – on the one hand, scientific positivism, in which we pursue underlying objective truths about the universe, and on the other hand, constructivism, in which all knowledge is the product of social processes. In both cases, Barad critiques a "representationalist belief in the power of words to mirror preexisting phenomena" [7]. Instead, "for Barad, data is not a measurement of an external world, or a representation of a thing that exists independently of the act of measuring; instead, it is produced by instruments - the instruments are part of an apparatus - entangled with the world. It is in the meeting point, or the boundary, between the apparatus and the rest of the world that the world and the apparatus are made, in a process of mutual becoming" [105].

The emergence of practices in design emerge from the intra-actions within phenomena, such as data, sensors, algorithms, and human perceptual and social factors. This entanglement of factors in even a single, specific design case could extend outward ad infinitum as we examine more and more phenomena. This *agential realism* as a relational theory ascribes phenomena as interdependent and arising concurrently, rather than within a cause-effect or representational relationship. Barad introduces the idea of "agential cuts" to describe the process of making an active choice between the phenomena being examined and the outside world [9]. This is not to say that elements of phenomena in this relationship are ignored; rather, the cut is a conscious and intentional decision to separate cause and effect and allow us to focus on particular interactions of interest [9].

This Baradian framing also offers an alternative to the view that data and mappings must represent the world, instead providing space for ambiguity in design [105]. Barad uses the idea of diffraction as a method to gather insights about meaning and understand insights and agents through one another. With agential cuts, the world need not be distilled into something neat and explicit; rather, particular intra-actions within phenomena can be isolated and investigated within the larger mess of the world. Through a diffractive approach, we are able to attend to the relationships different agents have with one another and how they come to matter in interaction [7].

2.2.2 Diffraction in Design. Hill's diffraction-in-action [52] extends the idea of reflection-in-action to capture how designers can make space for these entanglements between different agents and examine phenomena not only by their form but also by the feedback loops between them. In Research through Design (RtD) practices, diffraction-in-action captures a perspective wherein lived experience can be used to make meaning of phenomena and different agents, rather than attempting to use a particular agential cut that removes technology from the world in which it exists or assuming particular reflections of the world outside of our living bodies [3, 91].

For instance, Sanches et al. [105] use diffraction-in-action to examine "data as a material," with its own agency in interaction, and how designers designate meaning and knowledge about human bodies. Through a series of design case studies, the familiarisation of and "living-with" biodata changes the view of data from a representational one to one that acknowledges that meaning originates in the entanglement between that data and the world [105]. These different phenomena have unstable, fluid relationships with each other and can only come to have meaning through one another [7]; for instance, data only comes to have meaning based on defined concepts about what is being measured. The sensor itself also exists because of assumptions made in its design about the world [91], human bodies, and what we interpret its readings to mean or represent. By giving respect not only to data, but also to the relationship between maker and material itself, as co-creator in the design process, RtD examines the reciprocal negotiation between the experience of data and the data of experience [42].

Some examples of interest are in Howell et al.'s work with biodata, particularly in expressions of physiological arousal through measurements of heart rate [61] and skin conductance [59, 60]. Howell's perspective of data-as-material focuses on the intimate yet anonymous qualities of biodata and utilises unclear mappings so that humans can interpret their own connections to others' lives and emotions [61]. Skin conductance is only one particular distillation of emotion, and Howell notes the material transformations of this data source and intra-actions in the ambiguous biosensing along agential cuts (cf. [57, p. 13]). Examination along these cuts helps to examine how data-*becomes*-material, and how this mapping functions from the viewpoint of the designer. The designer's specific assumptions of the world are however not always accessible to an end user; Howell also notes that the assumption of skin conductance to represent emotion as a factor of physiological arousal neglects that calmer feelings can be overlooked [60].

2.2.3 Distillation through the Apparatus. Our focus, rather than on the process of diffraction, is on the idea of the *apparatus*. Barad uses the concept of apparatuses to describe the means of distilling and reconceptualising specific aspects of the world with specific agential cuts and making explicit [6]. Apparatuses are themselves phenomena, and they dynamically shift within their entanglement with other apparatuses. This means they, like the measurements they take and the ambiguity they negotiate, are neither objective nor neutral but rather tied to the concepts and boundaries they address. As well, they can themselves contain ambiguity; as Barad describes, apparatuses are not deterministic and do not function as specific instruments, rather being open-ended practices of conceptualisation. From the quantum physics origin of Barad's theory, a suitable analogy is the inability to measure position and momentum simultaneously; they are not just uncertain (unknowable), but are in fact ontologically indeterminate. Classical physics allows the measurement of one quantity or the other, but the apparatus needed to give meaning to one concept necessarily excludes the other.

For instance, this non-objectivity presents itself in the designation of mappings, where the apparatus is the medium through which meaning is made. Homewood examines changing data sources and their meanings in a non-passive, changing body [53] as dynamic properties of the apparatus. RtD has explored how, by attending to materials, including data as a material, as non-human agents [4, 34, 90], materiality can lead disambiguation in the design process. By allowing this perspective of the material into the narrative and engaging in reflection-in-action [106], the material provides reference to phenomena outside of itself. The apparatus then distills ambiguity in a way that reflects not only the designer's view, but also the influence of the material itself. Views of materiality are then a way of making this disambiguation by letting the material lead the meaning-making. Examination through the living body encourages the emergence of the designer's ideas over time through stages of intermediate knowledge [122] and shifts in infrastructural and temporal characters of the resulting artifacts [46]. This shift in design allows for other agents to become collaborators in meaning-making [70, 89] and for cues from lived experience to guide how designers handle ambiguity [95].

2.2.4 The Artefact as a Locus of Meaning. Giaccardi [46] offers an account of how artefacts can become partners in design (the *agential shift*) over an extended time (the *temporal shift*). This is dependent on the context in which they operate and instability of values that form within them (the *infrastructural shift*). These three shifts occur through interconnections with other things made possible by the use of data within the artefact. The data-enabled artefact is at once a computing system and an RtD product. Working with data, and particularly data connected to human users or other artefacts, blurs the traditional boundaries of where the artefact begins and ends. This argument suggests analogies to the Baradian apparatus, which we will unpack in Section 3.

A central issue in working with data is the representational nature of computing systems. According to Agre, designers of computational systems "often start by constructing representations of the activities that are found in the sites where they will be used.... A computer, then, does not simply have an instrumental use in a given site of practice; the computer is frequently *about* that site in its very design" [1]. The data-enabled artefact thus finds itself in philosophical limbo between its representational activities as a computational system and its anti-representationalist status as a Baradian apparatus, intra-acting with other objects and apparatuses as part of a phenomenon.

Our argument is not to avoid designing computational systems (a pointless exhortation in any case). Rather, in Section 3.2 we seek to call attention to the frictions between the neatness of meaning that are apparently represented within a data-enabled artefact and the messiness and contingency of the phenomena it co-produces. Since human experience is often a primary concern of designers, we examine what effect these frictions might have in a context of use and argue that they are best understood as shifting patterns of ambiguity.

2.3 DMIs as Data-Enabled Artefacts

The relationship between musician and instrument is sometimes identified as one of bodily extension, intimacy, and transparency, where performers perceive and act in a musical environment through a merger of the instrument and the self [88]. At other times, particularly with digital technology, the instrument might assume a symbolic or *hermeneutic* character [75], or a quality of otherness (*alterity* [63]) akin to a creative collaborator [117]. Waters describes the nature of instruments as a process, not an object [128]. Because the instrument is inherently entangled with the performer, whose embodied cognition and mediation construct the relationship with the instrument, [64], the nature of the instrument is always changing with that lived experience. The instrument and performer therefore co-construct one another [119]. Rather than being a stagnant entity, the instrument's agency in interaction shifts over time. Ihde's position is that tools acquire identity and meaning in interaction through use rather than through preordained status [63]. Moreover, the instrument is embedded in cultural and individual identity [12], not just in a societal or contextual relationship but also to the individual and even the perception of the self and one's own body [93, 99].

2.3.1 Representations in Broader Music Research. The motivations for DMI research to make sense of a performative context depart from other research contexts where formal concepts and theories are an end unto themselves, rather than a means to an end. Music Performance Analysis (MPA), for example, focuses on analysing performance recordings to extract parametric descriptions and characteristics (such as tempo, timing, dynamics, pitch accuracy, etc.) for understanding and assessing the actions and intentions of a performer or the effects of certain interpretation on the audience [74]. This relates directly to the field of Music Information Retrieval (MIR), where researchers work on mathematical representations and computational models which have the objective of describing audio signals in terms of sonological or musical features which are extracted from them. Although these features and descriptors can be used for performing specific tasks, the underlying models are generally not ascribed to these from their conceptualisation.

However, the interpretation of musical interaction from these symbolic representations is not without difficulty: Particularly within MIR and its application to further Music Emotion Research (MER), there is a tension in the dialogue between the objectivity of the ascribed models and the subjectivity and diversity of human perception [5, 22]. Aucouturier notes that a major issue in MIR is the lack of agreement between what MIR researchers do and how information is distilled into representations, as opposed to how cognitive science researchers discuss the subjectivity of music perception [68]. Extending the application of MIR and MER techniques in DMI design, mappings and models that include information beyond what humans perceive subjectively can lead to systematic misrepresentation of these relationships in-context [32].

Inscribing a DMI with some internal representation of music-theoretical meaning is only useful insofar as it helps produce a musically compelling result. From this perspective, Jack et al. suggest a view where DMIs are interpreted as research products, used to inquire about the entire ecosystem of an instrument [65]. This approach is predicated on the idea that designer and performer have independent views of the artefact itself: musicians conceptualise the instrument through performative engagement with the instrument as it is rather than becoming preoccupied what it *might become* (as in a technology prototype) nor being primed by the designer's concerns about what the instrument *should do* [76]. No understanding of the instrument is preordained; it is developed over time and in specific contexts through processes of co-production and co-control [63, 119].

2.3.2 DMIs as Representational Artefacts. From a Baradian viewpoint, the instrument (*apparatus*) is entangled with the player, the musical context, materials, audience, etc. [104]. Engineering-driven approaches to DMI design, where representational models are generated and embedded within the device without further consideration of such entanglements, do not account for the designer, performer and artefact's role in a temporally sustained design process. A more productive approach would be to consider DMIs as data-enabled artefacts which are partners in the design process [46] and contribute to our understanding of what we do as designers and musicians.

DMIs, like any computational system, are representational artefacts; the precise form of representation of music theory and human interaction matter enormously [85]. The keyboard is a ubiquitous example: Dolan [36] explains how the keyboard helped regulate the invention of many new musical instruments through the efficient instrumentalisation of sound, from its idealisation as a metaphor for total control in the eighteenth and nineteenth centuries to its inclusion as a control mechanism on many early electric and electronic instruments to its role in the standardisation of the MIDI protocol. Representations forged around the keyboard will make some concepts explicit (e.g., discretisation of pitch and onset) while leaving other concepts implicit or entirely unspecified.

DMI design always entails a negotiation, whether or not musician 'users' are present in the process. Lepri and McPherson [73] suggest that the design of a DMI entails an in-situ negotiation between designer and their technical tools (and by extension the designers of those tools). Designers and performers alike react to suggestions offered by the tools based on their previous experience and personal knowledge [100]. By making specific routines and processes more accessible or explicit, musical tools might bring attention to these, implicitly blocking different methods and techniques which do not fully adjust to such suggestions. Visual computer music languages (e.g. PureData (Pd) or Max) which are commonly used in the development of DMIs, for example, suggest specific ways of working with them due to the underlying motivations behind their creation, adoption and evolution (e.g. a focus on "rapid experimentation") while making other workflows harder to access. In Pd, the temporal dimension of music is left implicit in its design, while the patching mechanism for data flowing promotes a strong use of linear gesture-to-sound relationships [73]. This critique is taken further by Born and Snape [23], who analyse how the explicit objects and components that Max offers to artists and the prevailing infrastructures around the software have translated into the development of certain techniques for composition and interaction which lead to specific aesthetic expressions that are idiomatic of the software. As a result, pieces and instruments which use Max might exhibit characteristics that make them "sound like Max".

In summary, designers negotiate with the materials and tools we choose for the creation of DMIs and the systems of knowledge and theory embedded in them. Design attention might be guided by the explicit concepts encoded in digital tools and analogously, the eventual musical use of an instrument will be strongly influenced by the concepts and theories encoded within it, even as frictions emerge between the contexts of design and performance.

3 CHARTING THE AMBIGUITY SHIFT

Building on Giaccardi's three shifts in data-enabled RtD practice [46] (*agential shift*, *temporal shift*, *infrastructural shift*), we propose a fourth type, the *ambiguity shift*. Following Section 2.1, we view ambiguity in its various forms as both inescapable and useful. In this section we argue that designing with data leads to a reconfiguration of patterns of ambiguity, attributable to conceptual knowledge encoded in the digital artefact and the frictions between that knowledge and other possible systems of knowledge (including those of human users). We propose that data-enabled artefacts neither eliminate ambiguity nor transparently preserve it. Rather, the practices of designing with data produce phenomena internal to the artefact in which ambiguity is temporarily collapsed then recreated in a different form when the artefact is used in context.

After a definition of terms (Section 3.1), we propose in Section 3.2 that the way meaning is precipitated within a data-enabled artefact gives rise to frictions that lead to an ambiguity shift. We then propose two related mechanisms by which the shift takes place, which are mutually compatible, but as we will see in Section 5, present different implications for designers.

3.1 Terminology and Theoretical Framing

Our argument will touch on several terms with many disparate definitions as we link together different theoretical frameworks, so here we provide brief working definitions of some key concepts:

Phenomenon: Karen Barad's theory of agential realism [7], building on the philosophy-science of Niels Bohr, identifies the *phenomenon* as the "primary epistemological unit", the most primitive form of existence rather than things with fixed boundaries. Phenomena are "ontologically primitive relations — relations without preexisting relata", brought into existence through a process of *intra-action* between mutually inseparable components.

Apparatus: Barad [6], drawing on Niels Bohr and Michel Foucault, construes the apparatus as a measuring instrument, but not a passive one nor something that can be treated as an abstraction. Rather, the apparatus forms a part of the phenomenon, inseparable from the object it seeks to measure and the surrounding context and discourse: "apparatuses are not preexisting or fixed entities; they are themselves constituted through particular practices that are perpetually open to rearrangements, rearticulations, and other reworkings. This is part of the creativity and difficulty of doing science" [6, p. 102].

Data-enabled artefact: We take a *data-enabled artefact* to be a physical-digital system that collects, encodes or processes data. Data-enabled artefacts constitute *apparatuses* in Barad's sense of the term: "With this sensibility, the very act of making a distinction between the world and the apparatus, the subject and the object, the inside and outside of the phenomena we are interested in measuring, is an active choice that Barad calls an '*agential cut*'" [105]. Thus, rather than being passive observers, a data-enabled artefact collects observations from within phenomena, co-creating and interpreting information as part of its entanglement with a phenomenon, and producing output in the world according to that interpretation.

Indeterminacy is an ontological characteristic of phenomena. According to Barad [7], different apparatuses enact different agential cuts to phenomena into distinguishable subjects and objects, producing a "local resolution within the phenomenon of the inherent ontological indeterminacy" and local patterns of cause and effect. The lack of any single underlying model of a phenomenon, independent of the apparatus of measurement, constitutes its indeterminacy.

Uncertainty: We take a prevalent view in HCI for uncertainty, previously denoted as the *disciplining mode* of uncertainty [111] and describe it as an epistemic characteristic of phenomena. Once a phenomenon is observed, we model the uncertainty of its observations as a quantity that represents an estimation of the imperfection or incompleteness of the data collected.

Ambiguity: Human beings live experiences, learn and make meaning through embodied cognition. Although ambiguity and uncertainty are terms that are used interchangeably in literature [112], we understand ambiguity as an inherent property of human experience in the world, as each person can distil their own interpretation from similar situations. Designers of data-enabled artefacts choose conceptual systems and numerical representations in an attempt to organise the messiness of the world into compact and quantifiable measurements by which actions can be taken. In this case, the ambiguity comes from the possibility for each observer to have a differing interpretation of a phenomenon, or for a single observer to sustain multiple possible interpretations of the same phenomenon.

3.2 The Precipitation of Meaning within in an Artefact

Traditional research through design, according to Giaccardi, produces "an object around which behaviours and values are meant to precipitate and converge, if then to diverge again at use time" [46, p. 141]. Giaccardi's word *precipitate* is well-chosen, implying a diffuse or immanent set of behaviours or values which are made manifest through the catalysing action of the design object. This precipitating effect can be found within methods such as cultural probes [45], co-speculation [123], slow technology [92], accountable artefacts [14] and within DMI practice, hackable instruments [130] and instruments-as-probes [117], amongst others.

In the same paper, Giaccardi argues that data-enabled artefacts give rise to temporally sustained and reciprocal infrastructures of humans and things [46] which drive the proposed shifts in data-enabled RtD practice. Why would the use of data produce such an effect? In our interpretation, it is because the data-enabled artefact acts as an alternate site of precipitation, necessitated by the active perception and (typically) symbolic reasoning encoded in and by the artefact. As discussed in Section 1, data always comes with a frame, and designing data-enabled artefacts inherently requires a local making-determinate of particular concepts or representations within the artefact itself. The result is two competing loci of perception and action – one situated in a context of use as in traditional RtD, the other encoded in the artefact itself – each with their own distinct but partially overlapping webs of relationships.

The internal state of a data-enabled artefact is not transparently observable in use, and even if it were, the way in which the artefact derives symbolic meaning from the messiness of the world reflects the context and assumptions found in the design process which might differ from the context of use. Hence, competing interpretations and frictions emerge between the two loci of precipitation. These frictions make Giaccardi's three shifts possible, but they also reconfigure the patterns of ambiguity around the artefact, thus producing an *ambiguity shift*. Per Section 3.1, we treat ambiguity as a property of human experience, and thus even the most strictly deterministic digital system need not produce unambiguous outcomes. The term ambiguity *shift* makes no attempt to claim that the overall amount of ambiguity (however such a concept might be metricised) would go up or down; rather, it signifies that the location and form of the experienced ambiguity will be different than an alternate context that does not include the artefact. Working within this theoretical context, we can examine the mechanics of ambiguity shifts through explicit and implicit concepts and measurement as a determinant of meaning:

3.3 Explicit and Implicit Concepts

We build on the suggestion by Gitelman et al. [47] and Sanches et al. [105] that 'raw' data is a fiction: data always exists in relation to particular contexts of collection and particular interpretive frames. Interpretive frames are alluring. They offer the promise of representing and making sense of the world. However, frames carry a risk of uncritical acceptance: Agre's critique of AI researchers' "tendency to conflate representations with the things that they represent" [1] could apply equally to many forms of technology-enabled design practice.

To choose a frame – any frame – constitutes an agential cut, producing locally separable entities within a phenomenon. The biggest risk for designers is that they do not realise where and how they chose that frame, and that they consequently come to believe that the frame represents reality itself. To draw on a distinction by Sanches et al. [105], a keen sense of what data is *intended to* represent should not occlude a view of what else that data *could* represent.

Any frame will have the effect of making certain quantities or qualities *explicit*, brought to the foreground with apparent specificity: these are typically assigned words or symbols. Other quantities or qualities may remain *implicit* in the data; they are things that the data could represent but without their own obvious labels. Implicit quantities are often found only through carefully inspecting relationships within the data.

For the designer, explicit concepts easily come to lead the design process [85, p. 7]. In specification-led engineering processes, the specifications are by necessity explicit. Even in less goal-directed scenarios such as *bricolage* processes of interaction design [120], explicit concepts invite attention, exploration and manipulation. Metaphorically speaking, *knobs invite turning*, and the choice of frame in a data-enabled artefact essentially determines the knobs that will be placed in plain view. Explicit concepts also risk becoming reified as essential and constitutive components of the underlying phenomenon [121], rather than the local result of one of many possible agential cuts.

Several authors have noted the ways that explicit concepts in programming languages come to shape the ideas programmers express with them [81, 87, 110]. Magnusson [p. 123][77] offers similar ideas concerning music notation and transcription: "Notational systems define which parameters are 'valued', and these get abstracted out and assigned a symbol.... The transcribers will have to 'fit' what they hear into the symbolic language available." What is explicit or implicit in a data-enabled artefact is also entangled with its *inscription*, scenarios envisaged by the designer "out of which the future history of the object will develop" (though not necessarily in a straightforward way) [2, p. 214].

3.3.1 *MIDI as an Example.*

The MIDI (Musical Instrument Digital Interface) specification, first formalised by an industrial consortium in the early 1980's and still ubiquitous today, encodes a keyboard-centric view of music theory as a digital representation [35]. The foundational unit of MIDI is the *message*, of which the *note message* is the most common. A *note* in MIDI is given two explicit quantities: its *note number* (which key on a standard piano keyboard it corresponds to), and its *velocity* (how hard that key is pressed, typically related to the dynamics of the resulting sound). Notes in MIDI are inherently discrete, with a temporally instantaneous onset given by one message, and an instantaneous release given by another, later message.

Other concepts are representable in MIDI, but exist implicitly in the relationship amongst discrete messages: timing (in live MIDI performance, time is not a quantified part of a message but depends on when the message is received by the synthesiser); rhythm (which is implicit in the temporal relationships between successive messages); harmony (implicit in pitch relationships between messages); and phrasing (an ambiguous quality dependent on pitch, dynamic and timing relationships of longer arcs of messages). Other concepts, such as articulation and timbre, may not have any straightforward representation in MIDI at all, other than as indexes to an externally defined bank of sounds (e.g. MIDI program change messages).

The programmer or instrument designer working with MIDI messages is naturally drawn to the explicit qualities of note numbers and velocities. Common interactive digital artefacts include rudimentary keyboards [36] or theremin-like interactions where XY parameter spaces are mapped to pitch and amplitude [80]. Data-flow programming languages like Max or Pd similarly encourage manipulation of these basic elements over the subtler implicit relationships [80, 110].

The allure of MIDI's interpretive frame goes beyond data streams themselves. A common design pattern is the analysis and segmentation of audio streams into a series of discrete events defined by onset, release and pitch value. In music informatics research, the motivation is often *transcription* of a numeric score from audio; in instrument design, a common task is to turn an instrument into a *controller* of other sounds. Indeed, MIDI guitars have existed for decades based on analysing audio signals through a frame consisting of pitch, loudness and onset information. The design of such an artefact becomes its own process of transcription, through which the designer decides what is and is not important [77].

3.3.2 Mechanism of Ambiguity Shift. Playing a MIDI-based instrument, following Section 3.2, entails two loci where musical meaning precipitates. The first is the DMI's internal transcription according to the symbolic language of MIDI (which is non-neutral and shaped by political and economic as well as aesthetic forces [35]); the second is the contextual interpretation of the performance negotiated between musicians operating in a shared social and aesthetic environment.

Providing an explicit symbolic representation of a concept entails a process of disambiguation, in whole or in part. Out of the messy space of possible interpretations of data, a frame is chosen which yields a concrete observation stripped of its original indeterminacy. It is not sufficient to say that our frame might have error or bias with respect to some underlying reality; rather, the underlying phenomenon itself may be irresolvably ambiguous in the general case.

For example, what constitutes a 'note onset' in musical performance may not have any agreed meaning across all players, instruments and contexts. Similarly, not every quasi-harmonic sound lends itself to an obvious and indisputable measurement of 'pitch' (consider for example a sustained feedback tone on an electric guitar, which may appear to shift between harmonics depending on the listener and the context). For a digital instrument to internally encode performance as pitches and onsets requires a local disambiguation into one particular representation (or at least a parameterised probabilistic model). But meanwhile, the reductive nature of that framing means that multiple disparate phenomena which were previously distinguishable in their original context get modelled the same way by the digital system. Since the internal numerical model of the digital instrument determines its behaviour, this means that previously distinct phenomena become indistinct at its output.

Thus, a new *ambiguity of context* [44] is produced: the symbolic MIDI representation within the instrument entails a different context and a different set of assumptions than its situation of use, and therefore the instrument might act on its internal representation in surprising ways compared to how a human (or a different digital instrument designed in a different context) would handle the same indeterminate source material. Arguably, this ambiguity shift is produced because the instrument knows less than it "should": the digital system has no means of accounting for uncertainty or contextual factors of the particular environment it is played in when performing its internal transcription.

3.4 The Measurement Co-Determines the Phenomenon

In the preceding example, the process of transcription and representation entailed making certain concepts explicit and determinate, but this process also produced an ambiguity of context when the artefact's internal representation was unable to account for all the contextual and environmental factors in how it was used. The designer-engineer's response might be to attempt to explicitly model an ever-greater number of factors to better represent an underlying reality. In other words, perhaps a more "intelligent" instrument might be aware of embodied, stylistic and even social factors of its use, in order to generate more situationally appropriate responses.

The trouble with this approach lies in the assumption that there exists a stable, observation-independent reality to represent in the first place. As we argued in Section 2.2, a data-enabled artefact acts as a Baradian apparatus, entangled with its object of measurement as part of a phenomenon that co-constitutes both object and apparatus. As the quantum experiments of the early 20th century showed, reality is indeterminate, becoming determinate in different ways through different measurement practices. This is to say a data-enabled artefact cannot neutrally and objectively represent all of the contextual factors of its use because the artefact forms an inseparable part of the context and ecology of use.

It is also important to avoid a constructivist view which makes the apparatus the sole causative agent of its own observations. Barad cautions [6, p. 96]:

> "Since there is no inherent distinction between object and apparatus, the property in question cannot be meaningfully attributed to either an abstracted object or an abstracted measuring instrument. That is, the measured quantities in a given experiment are not values of properties that belong to an observation-independent object, nor are they purely artifactual values created by the act of measurement (which would belie any sensible meaning of the word 'measurement'). My reading is that the measured properties refer to phenomena, remembering that phenomena are physical-conceptual 'intra-actions' whose unambiguous account requires 'a description of *all relevant features of the experimental arrangement*' " (emphasis added).

What might constitute "all relevant features" of a data-enabled artefact in use? Surely it must encompass both loci of meaning (Section 3.2): internal symbolic representations and interpretations in a context of use. But to suppose that the artefact itself can represent all relevant features *as determinate, symbolic data* is to impose a circular dependency. There must be some limit on what can be made determinate within the artefact or with how much certainty some measurements can be made. Additionally, apparent certainty within the internal symbolic representations of the artefact need not translate to clear and unambiguous human experience; to suppose otherwise would once again assume that the artefact can internally represent all the external conditions of its use.

Thus, we argue that to design a data-enabled artefact is to make an agential cut, choosing a way of distilling the indeterminacy of the world into particular observable phenomena. The data-enabled artefact collapses a space of potentialities into determinate measurements, which does not imply that greater overall certainty has been gained nor that ambiguity has disappeared. Like the famous dual-slit experiments demonstrating the particle-wave duality of light [7, p. 815] wherein a different and mutually incompatible apparatus is needed to elicit each phenomenon, increased determinacy in one aspect of measurement necessarily implies greater indeterminacy in another.

3.4.1 FFT and Pitch Detection as Examples.

Agential realism provides a useful perspective on certain familiar conundrums in engineering. In signal processing, Fourier theory stipulates that signals can be mathematically represented either as time series or as a weighted sum of sinusoids of different frequencies, amplitudes and phases, and that lossless transformations between these two domains exist (the Fourier transform and its inverse) [109]. The special case of the Discrete Fourier Transform (and its efficient computational implementation in the Fast Fourier Transform or FFT) has become a bedrock of signal processing, including in audio. Here, successive short time windows of a signal are expressed as the sum of a finite number of frequency components. From this *frequency domain* representation, various analytical features can be efficiently extracted, or the signal can be modified before a transformation back into the *time domain*. Signal processing engineers are familiar with the inherent trade-off between time and frequency resolution: a longer time window gives greater precision in frequency at the cost of a reduced precision at localising events in time, and also a greater latency when such a system is deployed in real time. In audio, high frequency resolution can be useful for applications such as pitch tracking, while temporal resolution can be useful for detection of musical events (e.g. note onsets).

In the Baradian view, the time-frequency tradeoff is not simply a practical limitation of the algorithm. Rather, frequencies and temporal events are fundamentally different phenomena that only come into existence through different, mutually incompatible apparatuses. Frequency has no temporally instantaneous meaning; it is only meaningful with respect to a sufficient window in time. This irreducibility has far-reaching implications for the design of interactive music systems, where detecting 'notes' (i.e. onsets with a definite pitch) is a common task. The apparatus that yields an effective onset detector is unlikely to be the same one that produces the best pitch detection, and the friction between them and the different phenomena they enact yields a number of problems familiar to music technologists: apparent inaccuracy, where the system disagrees with the expectations of human musicians; temporal sloppiness, including latency and jitter, which reduce the perceived quality of performance [66]; and surprising discontinuous behaviours resulting from artificial boundaries (for example, it is always possible to find edge cases of signals without an obvious frequency or onset).

The entanglement grows to encompass sociotechnical factors when we consider that pitch (or even fundamental frequency) doesn't simply fall out of an FFT as a neutral, obvious property. With or without a Fourier transform, pitch detection algorithms are carefully tuned to operate in familiar settings through feature engineering [30, 79] or training on corpora of labelled examples [72]. In an echo of Sterne's format theory [113], normative psychoacoustic experiments and aesthetic judgments can be reified into constitutive components of "pitch" as an idealised property of musical signals, overlooking the diversity of sonic practices and ways of hearing [37]. When such a detector is deployed in real-time transcription, as in the audio-to-MIDI converter, it will inevitably fail in certain situations. This is not because the detector is insufficiently refined, but because *pitch doesn't exist* independently of a specific material-discursive framework that encodes assumptions about humans, technologies and cultural systems, while also precluding the making-determinate of other aspects of the same phenomenon the system seeks to quantify.

3.4.2 Mechanism of Ambiguity Shift. Where the example in Section 3.3.2 focused on *ambiguity of context* in Gaver et al.'s taxonomy [44], the pitch detector can be seen to (also) produce an *ambiguity of information*. While the ambiguity of context arose from the data-enabled artefact knowing less than it should, here the ambiguity of information arises from the artefact knowing *more* than it should: it offers apparent certitude about quantities with no universal meaning. Different pitch detection algorithms and settings will produce different (but equally confident) values from the same signal, leading (as in Gaver's examples) to quantities that are simultaneously over-determined and untrustworthy, leaving the user to decide what to make of it all. (Bowers and Green use this very ambiguity to artistic effect in a series of performance art pieces based on disagreeing pitch trackers [25]).

To summarise the Baradian argument, the phenomenon of musical pitch is mutually constituted by the audio (the object of measurement) and its measuring apparatus, which are inseparable from one another. The apparatus is a complex amalgam of technical and discursive elements, including means of capturing and sampling audio signals, mathematical formalisms such as the Fourier transform, algorithms for identifying features of interest, implementations on particular hardware, theories of human hearing, and musical ideology around what constitutes "pitch" and how it might manifest with common signals in familiar musical scenarios. To the extent the term "pitch" can be said to have any abstract meaning, it will be indeterminate until the apparatus collapses that indeterminacy by enacting a particular agential cut. At that point, "pitch" is best understood not as a musical universal but as a situated measurement whose meaning is dependent on a complete accounting of the system that produced it.

However, using this situated measurement to generate new musical material obscures the full accounting of its origins, and "pitch" takes on a universal status it doesn't deserve. Since there never existed an abstract, determinate, unambiguous concept of pitch to begin with, the result is that musicians perceive the output of the system as ambiguous, unpredictable or even wrong in certain cases. The locations of those ambiguities may be different than where they would have appeared without the intervention of the artefact (i.e. how a particular human, in a particular context, would have perceived the pitch of the input), and determinacy of a concept within the artefact is no guarantee of unambiguity in use. Ambiguity has not been eliminated, but it has shifted. Better engineering can move the ambiguity to more desirable places, but the problem cannot be fully solved because it is an ontological one which cuts to the core of what we take symbolic concepts to mean.

Thus, to use a data-enabled artefact such as a digital musical instrument is to grapple with the repercussions of the conceptual frameworks that underpin its operation, whether or not conceptual knowledge is the intended output of the artefact. In the following section, we present three case studies of our own artefacts and how we grappled with the conceptual frictions of working with data.

4 CASE STUDIES

Design can be an intuitive process. Designers often make decisions without specific analytical attention (whether reflective or diffractive), operating on a 'felt sense' of meaning-making rather than rational processes [43, 94]. We describe here a set of three case studies from our own practices in digital musical instrument (DMI) design, each of which acts as a different apparatus producing its own forms of ambiguity shift. These designs precede this paper, and each was created for different musical and research reasons. In this paper, we reflect on each artefact and the design decisions that underpin it, seeking to elucidate the agential cuts through its use of data, mappings and conceptual spaces. To direct this exploration we focus on five questions:

1. What concepts or ideas are being made explicit?
2. What concepts or ideas are left implicit?
3. What things were indeterminate prior to observation or conceptualisation?
4. What ambiguities are generated or reconfigured by the particular apparatus?
5. What assumptions were made in the design?

4.1 VoxEMG and the VoxBox

The VoxEMG (Figure 1), designed by Reed and McPherson, is a platform for control and interaction with muscular movements through surface electromyography (sEMG) [101]. Neural electrical impulses, which cause muscles to contract, are measured through electrodes placed onto the skin. This measurement can then be used for real-time feedback. The VoxEMG was designed specifically for interfacing with muscles active while singing, such as the laryngeal muscles [98]. The voice is a unique case for DMIs, as there is little externally available feedback about the vocalist's movements or behaviour; the use of sEMG allows for interaction with covert movement while singing [98].

4.1.1 Mapping Vocal Muscular Movement to Sound. This case study designed a basic sonification mapping for interaction with the laryngeal sEMG while singing; rather than manipulating the vocal signal itself, used the sEMG signal to control a soundscape in which a singer might explore their action and movement without changing something about the vocal signal. The goal of the sonification itself was to create a non-musical backdrop, to avoid assuming particular vocal styles would be used. The sound should represent the body in some way, and simulating a sort of wind or breath with a slightly scratchy quality to indicate tension within the muscles during contraction. The synthesis, performed on Bela, an open-source embedded computing platform for creating real-time audio- and sensor-based interactive systems [82], is based on a physical model of a rubber duck, which calculates the differential of an input signal and simulates the sound of the airflow through a rubber duck — the greater the change in squeezing pressure (based on the input signal), the greater the filtering (in this case, the squeak).[1] Here,

---

[1] Rubber Duckie by Christian Heinrichs (DC:A) https://learn.bela.io/tutorials/pure-data/synthesis/rubber-duckie/

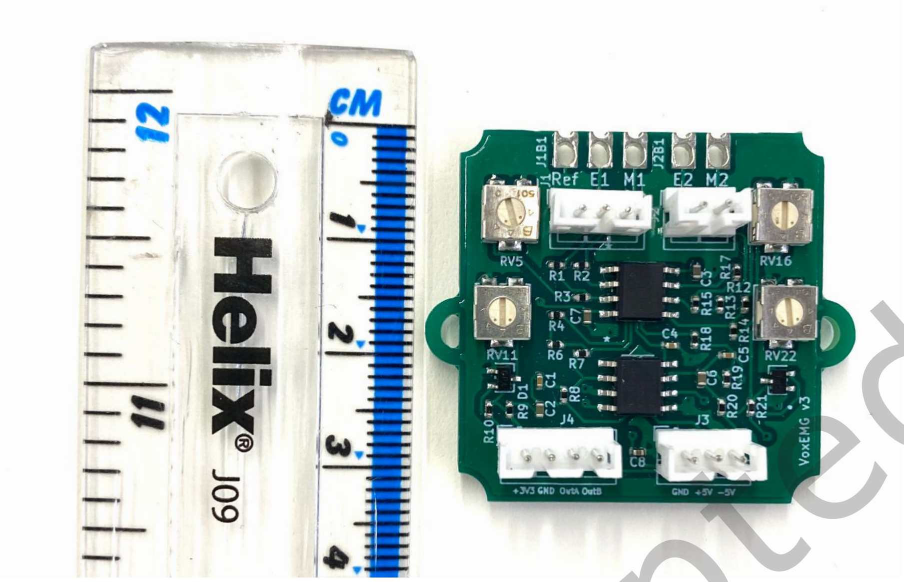

Fig. 1. The VoxEMG board, a small dedicated PCB for differential amplification of sEMG input, specifically designed for capturing vocal muscular movement. Schematics and images are provided open-source at https://github.com/courtcourtaney/voxEMG/.

a single muscle is measured with the VoxEMG for the input signal. The model is adjusted so the differential of the sEMG signal is calculated and mapped to the cutoff of a high-pass filter applied to a white noise generator. This causes a sort of *whooshing* when the muscle contracts and there is a large change in the signal. Another noisy drone is added as a base layer; together, in the intended design, the sonification sounded a bit like a crackly wind with changing intensity.

At the time of prototyping, Reed required a general output feedback that would allow for exploration of what kinds of vocal techniques could be measured effectively. The basic requirement was that the output should indicate that a movement had occurred and that sEMG could be gathered from particular muscle sites; this needed to be done before attempting to map explicit vocal movement to distinct outputs. Only one muscle was measured at a time during the design to limit the interaction. It would be possible to determine gestures during which the muscle was active, but the implementation at this point did not endeavor to determine what kind of activity occurred. However, the activation can occur in any number of typical and extended vocal techniques, meaning that many different gestures produced similar sounds. This became interesting and enjoyable to use during the design process.[2] The mapping was retained to explore the ambiguous output data and the many-to-few style representation of the vocal movement that lead to the resulting sonification.

As a part of the design process, Reed conducted an autoethnographic study to examine the use of the system in practice. She incorporated the VoxEMG into her own vocal practice while iterating on the design and documented her experience with it. The ambiguity of the output and not immediately knowing which aspects of her actions had caused a response in the sonification allowed her to explore her technique and form her own understanding of the mapping, based on what she knew about her own body and vocal practice [99]. There were moments in the interaction where Reed did not immediately understand the connection between her action and the resulting

---

[2] A demo of this sonification, presented at the ACM Augmented Humans (AHs) International Conference 2022: https://youtu.be/grvRBR5DjRs

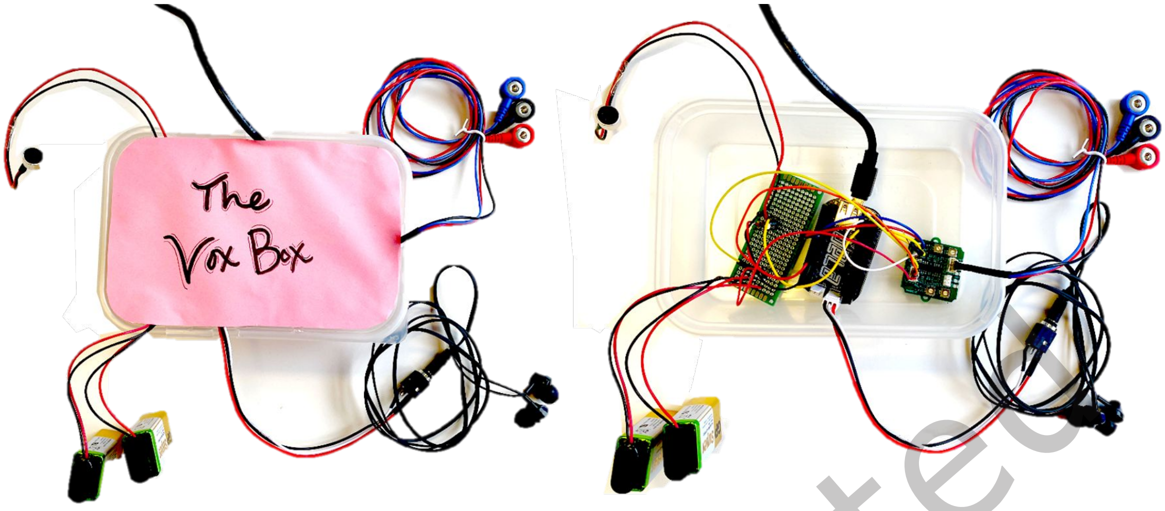

Fig. 2. The VoxBox kit given to other vocalists, containing necessary components for at-home vocal sEMG measurement and sonification with the VoxEMG and Bela.

sonification; for instance, she described an instance of being concentrated on her sound and preparatory actions, like posture and breath, which typically originated in a more unconscious manner. When Reed received feedback as a result of these non-vocalised movements, it took time and understanding of the mapping to deduce the source of the change in feedback. Through ongoing exploration and design, the mapping was successful in allowing her to make connections between the sonification, her movements, and her vocal practice [99]. The additional feedback allowed Reed to explore actions she had been doing subconsciously, like engaging in supported breathing. This allowed her to consider vocal mechanics that had become more automatic through her experience as a singer [99].

4.1.2 Adapting the Mapping for Other Users. In a second implementation of the VoxEMG, Reed adapted the system into the sEMG-at-home VoxBox kit (Figure 2) [100]. Two other vocalists received the VoxBox as part of a study to explore how they would interact with sonified vocal sEMG and what those interactions would reveal, both from a research perspective and to the vocalists themselves, about vocal embodiment.

The sonification and measurement of a single muscle was left as described above (Figure 3), with the intention to allow other vocalists to form their own meaning about what they were doing in their vocal practice to cause changes in the feedback. Using the VoxBox as a probe into this embodied vocal practice, the study aimed to observe connections other vocalists made between their own actions and the sonified feedback [100].

The outcome of this mapping was that the vocalists were able to make connections between the sonification and some of their movements. This included some of the same non-voiced actions Reed had observed, like supported breathing, which was evidenced through interviews with the vocalists about their experience [100]. However, it was also found that there was mismatch between the vocalists' expectations of the system and the feedback they experienced using the VoxBox. For both vocalists, as for Reed, there were moments of confusion about the mapping; they expected the sonification to change at moments where it did not, and vice versa. The uncertainty within the ambiguous response and divergence from anticipated behaviour from the system lead to confusion and self-blame, rather than meaning-making [100]. In the end, this study revealed some larger truths about the role of technology in interaction. Namely, the vocalists' responses indicate that technology can be perceived to

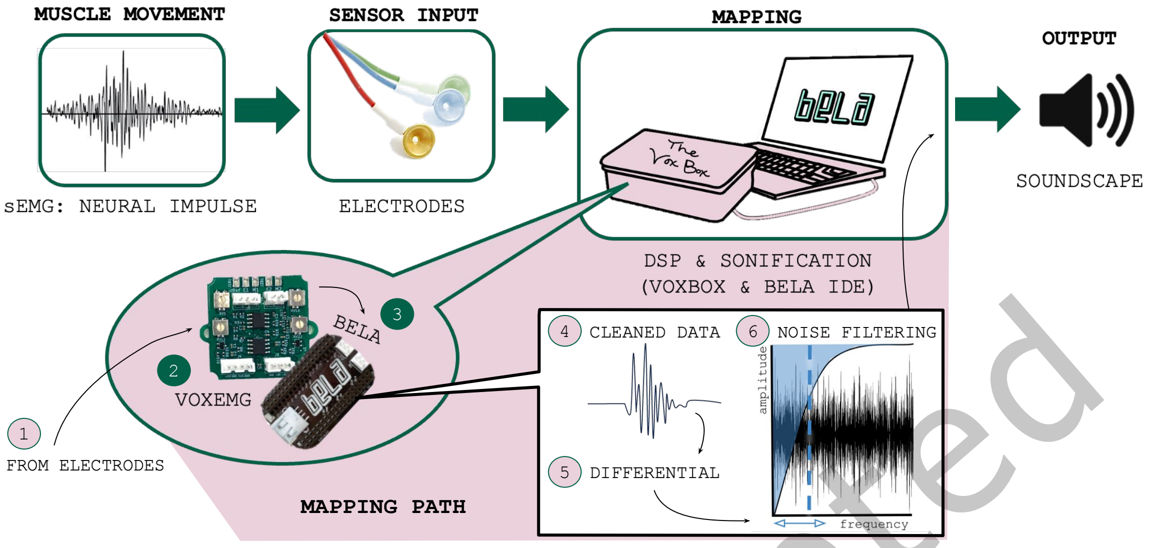

Fig. 3. The process of disambiguating laryngeal muscular activations and mapping them to sonifications with the VoxBox: Muscle movements are measured with sEMG through electrodes. The signals are processed through VoxEMG and Bela. The differential of the cleaned signal used to control the cutoff frequency of a white-noise filter. This filtered noise is then presented with another noisy drone to form a soundscape for vocal interaction.

be "correct" and it tells us some truth about the information we as designers supply through our mappings; in this case, information about our bodies and action.

4.1.3 Reflecting on the VoxBox's Design. The VoxBox deliberately used a simple signal mapping and intentionally ambiguous response, which did not attempt to distil clear conceptual meaning from the EMG signal (in contrast to the subsequent case studies, which aimed to disambiguate and provide explicit meaning to the data and are discussed below). The goal of the data sonification was rather to create an environment where the end user was responsible for the disambiguation.

Biophysically, EMG signals are notoriously complex and stochastic; while the density of events generally corresponds to the amount of muscular activation, the local details, whether measured in the time or frequency domain, are unpredictable. sEMG is known to be noisy; other bodily activity can be easily captured and we found this to be especially the case when working with the vocal musculature. The muscles are small and there are several large arteries in the measured area, along with other muscles that are responsible for moving the head and neck in addition to the larynx. Although it is possible to capture vocal movement, there are other body aspects represented in the signal. As a result, a requirement to produce clear conceptual meaning from the signal would require long temporal averaging whose resulting latency and low bandwidth would work against the goal of building a system for live performance.

Instead, the simple adjustable high-pass filter on a familiar white noise signal encourages a vocalist to construct their own meaning and experience from the phenomenon, while the avoidance of overtly 'musical' qualities of pitch and rhythm attempt to avoid privileging particular notions of music theory. Nonetheless, the result is never aesthetically neutral, with the VoxBox tending toward certain free improvisation scenarios — perhaps because a free improvisation aesthetic actively allows noise-like sounds to be part of the musical argument in a way that traditional Western practice does not. This worked for the experience of the VoxEMG system during its design, as Reed is more familiar with this type of free improvisation in her own vocal practice; while aiming to produce a neutral environment in which vocalists could experiment and explore their vocal behaviour, Reed reflected that she had assumed this particular context would be accessible to others without realising it. This is key to understanding what happened with the other vocalists using the VoxBox.

A critical moment in the design process occurred when shifting perspectives: During the initial design, she moved from a third-person to a first-person perspective, which allowed a letting-go of certain representationalist ideals of the sEMG signal: what it conceptually 'meant' was subsidiary to what effect it produced in rehearsal or performance. Rather than attempting to classify explicitly vocal behaviours using the sEMG signal, which likely would have been inaccurate due to the noise and overlapping behaviours of the muscles in different actions, Reed intentionally focused on providing an ambiguous output through the VoxBox's mapping. Instead, she left it to the user to determine what aspects of movement created changes in the feedback.

Having shifted views to first-person, it then became possible to return to a third-person perspective through querying the extended experience of the other two vocalists in the VoxBox study. However, as seen in the experience of the vocalists, these observations acknowledge Reed's role as an entangled participant-observer, rather than an ostensibly neutral 'modest witness' [50]. The mapping, and therefore the interaction, are entangled with her own understanding of sEMG signals in general, her experiences while prototyping, and her background as a vocalist working in more ambiguous and free-improvisation settings. Returning to our four aspects of ambiguity in the VoxBox, as the other vocalists experienced it, we see the ambiguity is shifted almost completely onto the user:

What concepts or ideas are being made explicit? The VoxBox's design was intended to make the presence of movement explicit. Where vocal muscular movement is typically internal and observed only through interoceptive understanding of the body, the sEMG mapping allowed for covert muscular movement to be measured and externalised through sonification.

What concepts or ideas are left implicit? With this many-to-few mapping of non-explicit vocal gesture to sound, the design intention was to leave meaning-making up to the end user. There is no expected action from the technology; rather, the sonification was designed to highlight existing understanding of the body during vocal practice so that users could make their own judgements of the output.

What things were indeterminate prior to observation or conceptualisation? The decision to sample a single point on the body was made to limit the scope of interaction. The activation of this single muscle is generally ambiguous in vocal interaction; the laryngeal muscles perform specific tasks but are used in combination to execute different vocal gestures and the same muscle will activate in very different behaviour. By examining a single muscle, the VoxBox collapses this ambiguous and indeterminate action into a single point for mapping. The presumption that this spatial location represents certain underlying actions in vocal technique, but in practice different individuals' physiology might differ and their technique and background might lead to different activity. Therefore, this presumption might not be universally or straightforwardly true. This ambiguity originates in our understanding of our actions — we cannot explicitly determine how these small muscles move in our own perception.

What ambiguities are generated or reconfigured by the particular apparatus? It is important to again acknowledge that the VoxBox also contains the context of its deployment, both of Reed's own design and vocal background in the development of the VoxEMG and also of the specific probe used in study with other vocalists with the VoxBox. In a "typical" DMI case, which will be seen in the subsequent examples in this section, the designer must make a designation of what parameters will be controlled and how they will be mapped, with an endeavor to create some explicit meaning for the end user. In this case, the mapping intentionally introduces ambiguity at the output by having the user determine what aspects of their action is responsible for the feedback they hear.

The presentation of a complex movement (which is largely ambiguous to the performer anyway, as they are understood through higher-level awareness of movement) is as a time series signal, of a particular bandwidth, with a particular noise floor that determines the meaningful dynamic range. This aspect of the apparatus is also determined by the electrodes themselves, the analog preamplifier components making up the VoxEMG, and by digital conversion on Bela, meaning the particular physical aspects of the measurement also introduce their own influence and ambiguities on the resulting feedback. A more extended discussion of electromyography as Baradian apparatus can be found in [102].

What assumptions were made in the design? The equipment used in the VoxBox cannot be disentangled with its development, especially before reaching the designer's hands. For instance, the cup electrodes used were mainly cultured in a medical context and expect particular environments of use that do not introduce as much ambiguous input and noise as vocal performance might. The design assumes the electrodes will function within its intended environment and culture. The sonification as a non-musical soundscape centred the interaction in a free-improvisation context and assumed this would avoid bias in favour of one particular musical theory. It was thought this would provide space any vocal practice to be used as input. Additionally, we assume that the measurement of a single-muscle input can represent the movement of larger groups of muscles needed for particular techniques. This extends to assuming that the activation of the muscle, even in different individuals' physiologies, will be at an amplitude greater than surrounding physiological noise and therefore be captured when calculating the differential of the signal.

4.2 Tone Transfer with Neural Networks

This project (technical details of which can be found in [28]) was motivated by the desire to control a synthesizer using a familiar musical instrument rather than a digital controller. Franco Caspe, the main designer, plays electric guitar, an instrument whose history is inextricably intertwined with the development of sound effects and musical experimentation. However, notwithstanding a number of commercial MIDI guitar controllers (e.g. Roland GR and GK series), the variety of integration between electric guitars and synthesizers has been limited [83].

The primary vehicle for Caspe's exploration was the development of ML algorithms within audio software for assembling a system, composed of a guitar, audio interface, computer, software, and amplifier, that can be considered as a single instrument. The treatment of audio signals and symbolic synthesizer controls constitutes the primary data within this data-enabled artifact, and the design of the ML algorithm that processes this data is the main focus of his research, which he initially approached as an engineering task to be solved.

Recently, the role of materiality of ML algorithms has been addressed in design research and HCI. This view escapes the technical reflexive approach of ML development as an optimisation task within a lab, and it rather explores complex modes of intra-action [7] of humans and non-humans agents within scope. Scurto et al. [107] explore diffraction as a framework for analysis and design of art-practice with ML, identifying a set of socio-technical conditions where art-based ML prototypes interfere with normative ML conceptual approaches, and outline *intra-active machine learning* as a process to conceptually reconfigure the boundaries of designers and engineers, encouraging hands-on experimentation and iterative designs. Under this approach, designers can take an embodied perspective over which every element (data, model, learning algorithm, context of use) can be leveraged as design material. On the other hand, Benjamin et al. [18] employ a post-phenomenological [63] view to investigate design research projects and understand ML uncertainty, proposing three ways on which it can be a generative factor of phenomena in the world, and therefore exploited as design material. Understanding the intra-actions and the inherent materiality within the components of a ML algorithm is key for analysing the design process that Caspe conducted and the experiences he had while playing the instrument.

Most digital synthesizers are controlled with explicit symbolic instructions, typically using the industry-standard MIDI protocol. As discussed in Section 3.3.1, MIDI foregrounds certain explicit parameters such as note beginnings and endings, sound intensity (MIDI velocity), and fundamental frequency (pitch, via MIDI note number). Using audio signals to control synthesizers presents challenges in generating appropriate conceptual meaning from an unlabelled time series [96], in the process reducing the information complexity but also potentially losing critical nuance. One longstanding approach is to directly extract explicit control parameters from audio, deriving a compact representation of salient characteristics that match the expected controls of the synth. Guitar-to-MIDI converters typically do precisely this, using music informatics techniques to identify onsets and frequencies and converting these to MIDI control messages.

The approach here presented begins similarly with audio analysis algorithms for estimating pitch and loudness. However, Caspe did not implement a method for explicit onset detection, instead producing pitch and loudness estimates as a time series, and relying on a Machine Learning algorithm for generating musical instrument sounds based on such series, while also providing the capabilities of control with other instruments, such as a guitar.

At the time of this project's inception, Caspe and his co-designers noted an interesting application where neural networks are used as a control processor and synthesizer, making it an interesting candidate for an audio-to-audio mapping strategy. They based their initial design on one of the chief examples of such an application, called Differentiable Digital Signal Processing (DDSP) Decoder [39], which learns a scheme for the control and synthesis of musical instrument sounds from an audio corpus. After training, the neural net can be controlled with pitch and loudness extracted from a variety of audio signals, while the sound it produces retains the timbral characteristics of the original training set. This approach is known as *Tone Transfer*, as it allows the transformation of the timbre of an input audio source while preserving its basic music form (fundamental frequency and dynamics).

An interesting characteristic of DDSP is how information is conceptually modeled across the neural network architecture, through distinct agential cuts that assign explicit meanings and functions for the control inputs, the internal sound representations, and the outputs, which we briefly review as follows:

- *Pitch and loudness extractors as input layers:* The neural network model features two input layers that implement signal processing algorithms, that *explicitly* extract pitch and loudness sequences from an audio input. During the training process, the network *implicitly* learns to map these two variables to a set of predicted synthesizer parameters.
- *Explicit modelling of audio components:* The predicted parameters are fed to another signal processing block, a set of audio oscillators and noise generators, that are implemented as internal layers in the neural network. This creates a strong inductive bias in the learning process toward audio generation in terms of the specific parameters required by the audio generator layers. In other words, the architecture learns how to control the DSP blocks using a standard Deep Neural Network (DNN) training procedure.
- *Spectral similarity metric:* With the input layers and audio generators in place, the neural network is trained with a reconstruction objective called *resynthesis*: given a target audio signal and a sequence of pitch and loudness estimates extracted from it, find a suitable set of synthesis parameters that in turn, generate new audio that is as close as possible to the target. This process requires a similarity metric, typically called a loss function, that guides the training as desired. The architecture is trained using a spectrogram loss as an objective function [127], which is minimised during training, and computes the difference between the magnitude spectra of the target and synthesized audio.

As quantities such as pitch, loudness, and frequency bands are modeled explicitly with algorithms acting as apparatuses, other important characteristics of the system are left to be learned implicitly by matching metrics within the loss function, such as the control of the synthesizer.

Fig. 4. A typical setup for FM Tone Transfer using an electric guitar and the neural network plugin.

The project builds on the approaches developed in DDSP, particularly on the idea of using audio components as layers within a neural network. Caspe decided to use Frequency Modulation (FM) synthesis, as it is a well-known synthesis architecture whose relatively few sound design primitives can be easily manipulated by players [28].

Since the technical goals were set beforehand and Caspe counted on a reference design, fixing the model architecture and employing data from different musical instruments was the main material for exploration with the algorithm. He trained different networks with corpora of data from different single instruments, and then implemented the result as a real-time audio plugin. Figure 4 shows a setup using the plugin for controlling the FM synthesizer using a common, unmodified electric guitar. The instrument is connected to a USB audio interface, and the captured audio is then processed by the plugin, which maps the pitch and loudness of the audio signal to synthesizer parameters using the neural network. The computer screen (Figure 5) shows the GUI and a diagnostic screen that visualises the pitch and loudness inputs and synthesizer controls predicted by the neural network.

Once deployed into a usable plugin, Caspe was able to spend a long time playing with the system. Initially, he struggled to get the instrument to react as he wanted, particularly in terms of the consistency between the actions he executed on the guitar, and the synthetic sound result, and especially when executing fast passages, such as sweep picking and pedal notes, or low-intensity hammer-on and pull-off. Adjusting the individual parameters associated with the pitch and loudness detectors could only take him so far, and with further use, Caspe noticed how his playing technique changed to accommodate to the interaction possibilities of the system. For instance, he became physically aware of the limitations of the monophonic pitch tracker, improving his legato technique, making it “cleaner”, and trying to ensure that only one string would sound at a time.

Furthermore, during live use, he started to play phrases with longer notes, not only as a way to showcase the timbral differences between a regular guitar and the system but also to have more control over the way they are

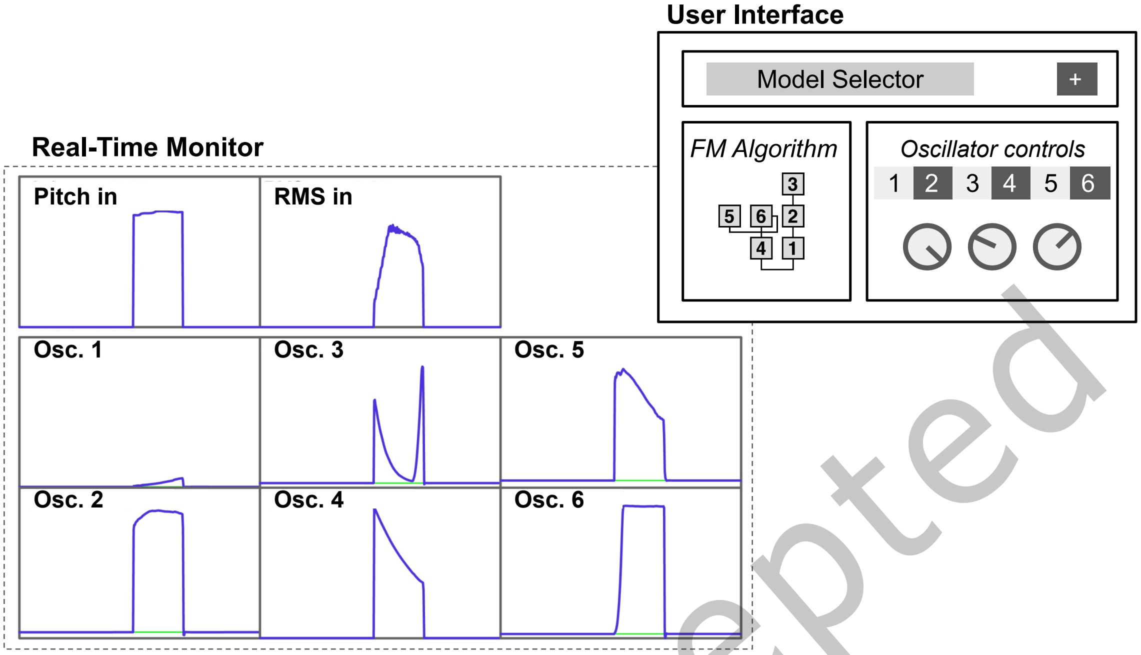

Fig. 5. A depiction of the plugin's user interface, and real-time monitor, showing the quantities made explicit within the system, (pitch and loudness/amplitude, and synthesis parameters for six oscillators), whereas timbre and note beginnings and endings are implicitly encoded within the temporal progression of such sequences, shown on the monitor.

executed, in order to maximise the chances of the system to interpret each action executed on the guitar as he intended as a designer.

Caspe understands this phenomenon as a *pattern leakage* [18] , where his playing experience is affected by the system being highly uncertain at times when his inputs are ambiguous to the algorithm that tries to explain them. Benjamin et al. identify this concept as a phenomenon that can potentially be leveraged for design, and indeed, Caspe observed that these modes are generative of unexpected behaviour that can be sonically interesting; some models can be induced to an low frequency oscillator (LFO)-like self-oscillation when the loudness trajectories are noisy or highly dissimilar to the ones shown during training. Nevertheless, the way this system reconfigures his relationship with the guitar does not lead him to a sympathetic use of the instrument as a controller.

As a guitarist, Caspe found that the system presents itself as a set of overly strict rules to comply with, and breaking at any point any of these could break the musical phrase that may be taking place. On the other hand, as an engineer, Caspe values the conceptual simplicity of the pitch and loudness over-time control scheme; this three-dimensional control space has the potential to afford nuanced results facilitated by lightweight systems that can run in real-time. In the end, he found that he could not entirely blame his playing nor blame the system approach for the system breaking his expectations as a player.

Only when the system is deployed in the wild does it shows its full dimension of interdisciplinary values that often collide. This condition has been identified previously by Scurto et al. as a situation the authors call "situational whole" [107], which they identify as an entry point for appropriation. And indeed, this tension between Caspe's points of view as a designer and musician, forced him to think about changing the data as design material on the data-enabled artefact, which finally lead him to a new tone transfer system design, using a completely different data representation [29].

What concepts or ideas are being made explicit? During inference (i.e., when played in real-time), the system requires two explicitly defined quantities alongside the audio signal, namely pitch and loudness, which are mapped into synthesizer control parameters. The synth control parameters are also explicit quantities, as their function is defined beforehand by placing the audio generation components within the network. It is also worth noting that all these quantities are analysed *over time*; the network can account for time dependence on the values, which are crucial for modeling the implicit characteristics of the system (see next). Finally, because the network training process seeks to minimise a spectrogram loss function, spectral content also becomes an explicit part of the internal conceptual representation of the artefact, even though it is only analysed during training and not during inference.

What concepts or ideas are left implicit? An important concept that is left implicit is timbre, which is itself ambiguous and lacking in an agreed definition other than by exclusion (everything that is not pitch and loudness) or by effect (the quality that allows one to distinguish two different instruments playing similar musical content). An important characteristic of this work is that we make no attempt to explicitly define timbre, instead concentrating on one quantifiable contributor to it, the audio spectrum. This has a notable implication: the spectrogram-matching objective, in prioritising frequency over temporal events, cannot account for fine temporal information. Transient characteristics of the instruments' audio, present as part of the audio signal, can be partially rendered by the network due to its capacity to model time dependence on pitch and loudness signals, even though they are never explicitly modelled.

The system does not learn the dynamic profile of the notes, as this is explicitly provided by the explicit input layers. As timbre is left implicit, we found that the timbral qualities of the learned instruments can be replicated only when the system is provided with pitch and loudness contours of the same instrument. During inference, the system is driven by an instrument that features a different loudness profile than that of the dataset, forcing the player to adjust their playing to better resemble the original training conditions if a convincing timbre is wanted.

What things were indeterminate prior to observation or conceptualisation? Pitch, loudness, and spectrum are modeled as independent concepts, but in actuality, they are not separable. As discussed in Section 3.4.1, pitch and loudness are not objective independently-existing entities but are brought into existence in relation to a particular apparatus of measurement. On the input side of the network we have explicit algorithms for constructing estimates of pitch and loudness from temporal windows of audio, and during training the spectral characteristics are also rendered observable in relation to a particular spectral measurement algorithm. However, this simple three-dimensional decomposition (pitch, and loudness, over time) of music performance is too reductive to accommodate a wide variety of interactions with the guitar and is only useful as long as the temporal trajectories are idiomatic to the instrument and player.

What ambiguities are generated or reconfigured by the particular apparatus? A musician's performance unfurls through embodied cognitive processes that involve a complex language of sensorimotor imagination and internal narrative. This is not expressed in a compatible vocabulary with that of the system's pitch and loudness trajectories, which results in only a subset of actions being processed and turned into sound by the system.

Throughout a performance, musicians may impart meaning (or not) to their sound-producing actions, but this is not necessarily understood by the audience in the same way, which could be explained in terms of Gaver et al. as an *ambiguity of information* [44] . The introduction of the Tone Transfer system is motivated by an extension of the timbral possibilities of the guitar but comes with an additional (and in this case unwanted) interpretation stage, that intends to derive meaning from every single player action, while rejecting seemingly "meaningless" phenomena that do not fit within its explicit (pitch and loudness) or implicit (timbre and transients) models of the world, and thus breaking Caspe's expectations as a performer. This takes place as a shift from an ambiguity of information to one of relationship, where Caspe found himself in friction with the system while trying to accommodate his playing so that he could establish a cause-effect relationship between his actions and sonic outcome that he could understand and harness for music performance.

The process of adapting his playing to this software has altered Caspe's connection with the instrument and reconfigured his role within the system: he reflects that "I am not only trying to play the guitar but also trying to comply with a third-party agent." The agential cuts established by the pitch, loudness, and spectral loss, have all played a role in reshaping his performance techniques. The software designed to modify the tone also modifies the player.

Encountering the system as a musician yielded a conceptualisation much different from the one Caspe had as a designer: what he thought was a simple and tidy parameterisation of performance became unwieldy through an ambiguity shift. This intra-active understanding of the system prompted him to reconsider the data he employed for the tone transfer problem. This stand has similarities with the practice of *intra-active machine learning* [107]. In this case, the ambiguity shift pointed him to iterate over the data design within the system. The technical details [29] are out of scope for this work, but this allowed him to re-negotiate the elements that remain explicit and implicit, and the ambiguity shift that they generate, yielding a more versatile system.

What assumptions are being made in the design? The entire project is naturally situated in certain styles of music characterised by guitars, synths, and monophonic playing techniques. Indeed the first two of these point in somewhat contradictory directions: guitar is often played polyphonically even in solo playing, while monophony is more common in classic synthesizers for technical and economic reasons. To assume that architecture of this sort would generalise to any musical context or any instrument would be a mistake. A further assumption appears in the modular architecture, where the software is deployed as an audio plugin. In publishing the technical work [28], Caspe pushed to the background any effects or contingencies to do with what software, hardware, or instrumental environment the plugin is deployed in, adopting a scientific logic of isolated but interacting components rather than entangled *intra-action* [7].

4.3 The (Un)ambiguous Guitar

The provocatively named *(Un)ambiguous Guitar* [16], designed by Adan Benito, is an augmented musical instrument designed primarily by Benito on the material basis of an electric guitar, but modified to alter the behaviour of gestures that would otherwise be sonically equivalent so that the sound they produce becomes distinct. In particular Benito attempted to "disambiguate" string bending gestures by programming different behaviours depending on the direction and quantity of the physical bend. String bending is a *modification gesture* [26], a common idiomatic resource electric guitar players employ on demand to exert precise, continuous control over pitch. The physics behind string stretching makes bending gestures inherently ambiguous: the pitch of the played note will always be raised regardless of the bending direction (upward/downward).

The (Un)ambiguous Guitar has been modified with custom magnetic bend sensors that measure angular alterations on a magnetic field produced by horizontal string displacement transverse to the longitude for the string [17], a hexaphonic pickup[3] and a dedicated processing unit to extract audio and string movement information for each string (see Figure 6). Benito designed the instrument to present itself to the player like any other electric guitar; no additional controls or means of interaction are made obvious, and the guitar can be connected to any standard processing chain one could normally use (e.g. pedals or amplifiers). The sensing apparatus and processing unit are conceptualised as a 'black box' to alter the sound produced by the interaction between the player and the strings based on conceptual features of the digital system, while the locus of interaction remains on the guitar itself.

---

[3] A pickup with six discrete outputs, one per string, that allows independent processing of signals from each individual string.

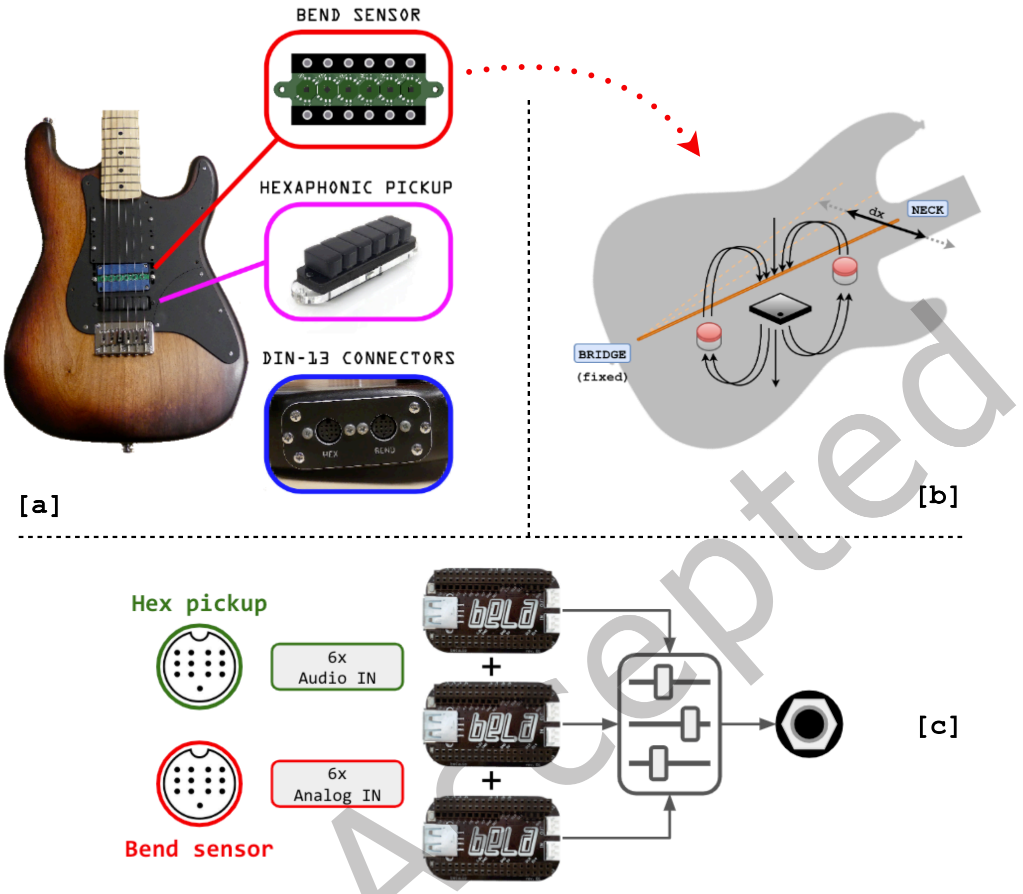

Fig. 6. [a] Components of the *(Un)ambiguous Guitar*. [b] Depiction of sensing device for capturing string displacement. [c] Internal diagram of the processing unit including three Bela processing units which run the signal processing algorithms that alter interaction based on sensed inputs.

Even though the sensor system provides a measurement that is taken to represent string displacement, the correlation between such measure and the different gestures that produce it is not obvious. Many interactions between instrument and performer, including simple string plucking, result in a lateral displacement of the string. Therefore, in an attempt to isolate and analyse bending gestures and to elicit a clear separation into upward and downward bends Benito, as designer, ended up building conceptual models of 'gestures' as explicit events with clear boundaries. This concept of 'gesture' thus becomes inscribed in the instrument, locked in by the parameters of the analysis and the interactions programmed as a result of that analysis; within the artefact, 'gesture' is no longer ambiguous but a specific measurable quantity. However, the guitarist playing the instrument might have a different conceptual framing of gesture, or might not concern themselves with such a conceptual approach at all. Hence, the (Un)ambiguous Guitar is neither entirely unambiguous, nor ambiguous in the original sense of the instrument; rather it produces an ambiguity shift for the player.

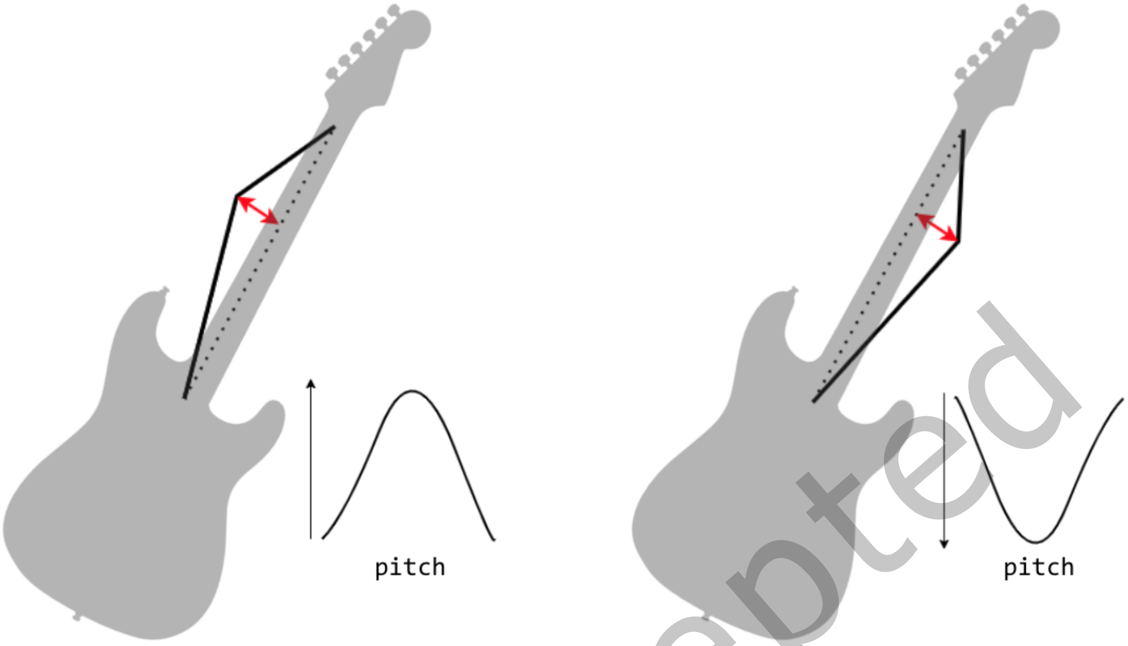

Fig. 7. Depiction of the behaviour for upward and downward bends on the unambiguous guitar (with respect to pitch).

The model of string bending inscribed in the instrument combines magnetic sensor signals with features extracted from the audio of each string. Benito's signal processing algorithm combines sensor data with real-time features extracted from the audio signal of the guitar to draw an explicit boundary of when bending gestures occur based on a model for transients which determines the beginning of a note and a pitch detection algorithm employed to estimate stable and smooth pitch trajectories of the vibrating string. These boundaries are learned empirically: several separate upward and downward bends are captured to fit a regression model that maps string displacement to pitch shift (produced by bending the string) from the reference note of each fret and string in the guitar [17]. This mapping is then used to alter the behaviour of downward bends in real time following a model of interaction which divides each note into 5 possible discrete states: *transient*, *note on*, *note off*, *bend on* and *bend off*. These states are navigated, and downward bend behaviour determined, by combining the transient model described above from the audio signal, a differentiated version of the displacement signal representing the velocity of the string bend, the estimated frequency and the confidence estimated of the pitch detector (more details are presented in [16]). The decision to encode a concept of 'transient' required heuristic tuning to decide what did and did not qualify based on these parameters. Similarly, experimentation was needed to decide the boundaries of *bend on* and *bend off* gestures, navigating a trade-off between response latency and transient leakage into the bending state.

For upward bends the behaviour remains identical to a standard guitar (raising the pitch), while for downward bends new behaviours are programmed, with bend amount controlling either a pitch shift algorithm which modifies the frequency of the note in real-time or an artificial feedback simulator. The pitch shifter allows the (Un)ambiguous Guitar to re-scale and re-map the pitch range of the bend (Figure 7), while the feedback simulator produces an idiosyncratic 'howling' sound reminiscent of the coupling between guitar pickup and very high-gain amplifiers, as popularised by Jimi Hendrix and other guitarists.

Benito put the (Un)ambiguous Guitar in the hands of expert players during a series of exploratory musical experiments [16] to investigate how performers adapt to *disambiguations* that either follow or confound their mental imagery of bending. The core interest was in whether specific patterns of interaction were elicited by different modalities across participants. Benito identified different levels of disruption for different participants and modalities ranging from the inability to adapt to the augmentation to being able to perform as if the instrument was unmodified and then recall the specific interaction on demand. Notably, almost all participants were able to recognise the subtlety of the interaction with the new gestural paradigms offered by the instrument and to identify the mechanisms of the augmentation during the unguided exploratory exercise, including developing nuanced understandings of the limitations and idiosyncrasies of the gesture detection mechanism (such as discontinuities in the interaction and differences in response along the fretboard). Participants rapidly came to terms with the modified behaviour of the instrument even when the setting attempted to disrupt their familiar musical imagery.

What concepts or ideas are being made explicit? Beyond the explicit features noted above, there are two main qualities made explicit during the design of the (Un)ambiguous Guitar. First, following the motivation to disambiguate upward and downward bending gestures, the apparatus explicitly conceptualises these as two separate entities with separate meanings that should yield different sonic results through the model described above. Second, the instrument brings the correspondence between bend displacement quantity and control of sonic parameters (pitch-shift or feedback) to the forefront of interaction, which is something that we could confirm during subsequent studies [16].

What concepts or ideas are left implicit? As a result of trying to isolate a specific set of bending gestures and quantify these in order to make explicit mappings, we leave the separation between excitation and modification gestures implicit in the (Un)ambiguous Guitar. For downward bending gestures to behave as we presume in our model (Figure 7), we expect that a transient event – normally an onset produced by plucking the string – is detected prior to the modification gesture so that a specific fret can be estimated and the bend quantity mapped to a control parameter. By employing this model of interaction to explicitly modify the behaviour of specific bending gestures, we implicitly create boundaries around other aspects of the gestural vocabulary of the performer and presume that our model will fit within that context. The system disrupts any pre-existing gestural inter-relationships possibly related to string displacement that do not fit our explicit model of bending gestures.

What things were indeterminate prior to observation or conceptualisation? Bending quantity and direction, from the apparatus' perspective, are indeterminate qualities of guitar performance that the (Un)ambiguous Guitar interprets as parametric quantities of gesture. The conceptual model embedded in the apparatus localises these through the combination of sonic features with observations of string movement in order to be able to draw explicit mappings. At a lower level, the model of string bending that gives an interpretation of data gathered through our sensing device is used to project performer actions through measurements of angle shifts in a magnetic field at a specific point along the guitar strings to obtain a measurable observation of string displacement that can be collapsed into a single variable. Similarly, to create mappings from gesture to pitch and to separate the different states of our gestural model, features are extracted from the sampled audio signal captured via the hexaphonic pickup. Bend note onsets are derived from a transient detector modelled through amplitude envelopes of this signal and analysed as discrete observations symbolically separable from other audio properties. Moreover, a pitch detector is employed to estimate fret position and to establish a correspondence between displacement observations and pitch-shifts produced by these. Neither onsets nor pitch are pre-existing qualities of musical expression and are therefore in-determinant from a phenomenological perspective, but here we conceptualise them through mathematical models and use the observations elicited by these as higher-level parameters of a gestural model.

What ambiguities are generated or reconfigured by the particular apparatus? Our approach to categorisation and subsequent 'disambiguation' of gesture entails a reconfiguration of the performer-instrument entanglement that gives the guitar its identity in use. By focusing on characterising and splitting a specific performative gesture, we enact certain boundaries that become sources of ambiguity on their own. The first of these boundaries appears through the piece-wise characterisation of bending gestures by which a different mapping is created for each string and fret in the guitar. The changes in response across the different dimensions of the fretboard produced by this approach were quickly appreciated by several performers who performed on the instrument without explicit knowledge of how it was configured [16] and created a sense of ambiguity regarding how the apparatus responds to gesture. Participants of the study described in [16], noticed changes in the response of the bending gestures for different parts of the fretboard and different strings with respect to the behaviour of an unaltered guitar and also how the speed at which a string bend would be produced would its behaviour and found that this would conflict with their mental representation of the bending the gesture. The second, and more disruptive, boundary comes from the implicit separation of excitation and modification gestures necessary to observe bending gestures as single conceptual entities. This separation is enacted by a heuristic mechanism whereby the speed (velocity) of displacement (first-order difference of the displacement measurements) is combined with our transient model to derive when a bending gesture in a particular direction starts and ends. As a result, performers were presented with a change of interaction where the speed of bending and the intensity of plucking had a direct effect on the response of the apparatus with respect to the expected behaviour. Moreover, a lack of continuity could be appreciated over the decision boundary which shifted the way in which performers approached certain techniques [16]. Although certain participants felt that they could adapt or navigate around these changes in response, some felt that their relationship with the instrument changed up to the point where they could not perform fluidly. By trying to somehow resolve the ambiguity of certain gestures, our design ended up generating new sources of ambiguity which reshaped the expectations of the performer and inadvertently created new tensions in the design of the apparatus.

What assumptions were made in the design? The audio and sensor processing assumes that excitation gestures (i.e., plucking) and modification gestures (i.e., bends) can be cleanly separated from one another and treated as discrete events. The system also expects a definable boundary between upward and downward bends as being modification gestures starting from the resting position of the string, which is conceptually linked to the first assumption. The system also assumes that pitch contours and transients are separate events and that the whole interaction between performer and string for this particular gesture (bending) can be modelled through a finite-state-machine model with clearly defined transitions.

The sensor approach presumes that alterations in magnetic field are representative of string displacement and that gestural information can be derived from the displacement of the string at one fixed point, even though the location of finger-string contact varies along the fretboard. Moreover, this approach assumes a correlation between such displacement and audio features (such as pitch and transients). Furthermore, Benito assumed that the pitch detection algorithm produced a frequency that is representative of the pitch of the string as intended by the performer, and that this pitch can be captured during the transient state of the model. It was also also assumed that transients are discrete events representative of pluck onsets that can be extracted from a peak or mean-square envelope of an audio signal from a string. It is taken as given that the output of each individual pickup and bend sensor contains mostly just information about the manipulation of that specific string without crosstalk.

5 DISCUSSION

The three preceding case studies describe some of our approaches, from our perspectives as designers, end users, and researchers, to crafting data representations and handling ambiguity. We reflect on the designs of these DMIs and the trade-offs between specificity and ambiguity, support for specific contexts of music-making without creating overly reductive, theoretically-laden artefacts "where there is not enough space left to act and improvise" [97] and places where artistically-productive ambiguity has been shifted into confusing or frustrating behaviour. Our contribution to RtD is more reflective (or indeed diffractive) than methodological: the three musical artefacts differ in their design process, use of data and musical context; other artefacts could also be examined along similar lines. It should also be noted that, as each of the three case studies and researchers' perspectives arise from different musical backgrounds and research focuses, there is a natural disparity between them. Although we have applied the same reflective questions to each of the case studies, the authors' own narration and reflection on their work is imbued with this plurality of experience. This disparity is significant to the diffractive approach, wherein the examination and reflexivity of the designs also generate more of their own agential cuts. These reflections are then inseparably entangled with the designers and authors themselves. This being outlined, the discussion presented in this section will attempt to unify our perspectives in order to propose reflection useful for others in HCI, who will bring yet more disparate perspectives in work of their own, to acknowledge this entanglement.

We endorse the suggestions by Sanches et al. that "the design researcher should resist the impulse for actionable insight from day one" and that designers should "hold space for messy, ambiguous data that requires active interpretation" [105]. More than some data-driven systems, DMIs provide fertile ground for exploring such approaches, since the embodied act of performance need not entail a conceptual or symbolic relationship to the technology of the instrument [63]. If, in the moment of performance, the instrument and performer dissolve from coherent entities into a "bundle of affects" [64, p. 111], what then is the need for the "clean and tidy data" that Sanches et al. critique?

However, resisting tidiness or actionable insight is harder than it seems. If "raw data" is a fiction [47], then encountering data anywhere in the pipeline necessitates an interpretive frame, even if that frame is solely internal to the artefact. Even without subscribing wholesale to an ideology of mapping conceptual spaces (Section 1), *any* conceptualisation of data within the design process is an agential cut, opening certain possibilities while foreclosing others. We now return to our two arguments from Section 3 on the mechanisms of ambiguity shift in data-enabled design, where we will see that the philosophical viewpoint leads to different practical suggestions for designers.

5.1 Designing for Explicit and Implicit Concepts

In Section 3.3 we argued that designing with data always makes certain concepts explicit and others implicit, and that the process of making explicit entails a local collapsing of ambiguity. Taking the example of Tone Transfer (Section 4.2), pitch and loudness were made explicit through particular numerical models, as was a particular conceptualisation of timbre (inherent in the spectral loss function used to train the network). Making these concepts explicit supported an engineering process intended to preserve pitch and loudness while varying timbre. On the other hand, techniques that were previously ambiguous in pitch or loudness (e.g., muted notes, certain types of polyphony) are forced into artificial categories while multiple playing techniques that produce identical measurements of pitch and loudness will become indistinguishable even if their original effect was notably different (e.g., plucking with fingers vs. plectrum). Reducing the space of playing techniques that produce sonically distinct results can in turn lead to a progressive reduction in the variety of interaction [67].

If the fundamental problem is that a representation privileges certain concepts over others, a possible response would be to increase its dimensionality or resolution: in other words, *make everything explicit.* This does not necessarily imply more complex interfaces with elaborate multi-parametric mappings (see Tanaka [118] for a critique of this longstanding technological impulse), but rather to build empirical models of progressively finer details of phenomena. Within such a representationalist view, we could draw on the cognitive or social sciences to provide numerical models of human factors, an outlook which can be found in music informatics [5, 71] and in scientifically controlled user studies [38, 126].

Making implicit qualities explicit is valuable: it promotes transparency and accountability of the researcher's assumptions, and it can lend crucial nuance to otherwise crude concepts. However, in light of the drawbacks of fixing an explicit interpretation of data, is it possible to make the *existence* of a concept explicit without providing any model or measurement whatsoever? The VoxEMG (Section 4.1) would seem to do this, treating the EMG signal as a measurement of muscular activation without attempting to interpret the contents of that signal. On the other hand, the process by which the time-series data is smoothed and used to control an audio filter – and indeed the process of acquiring the signal to begin with – make assumptions about what is meaningful, which may not be universal (e.g., perhaps in singing, particular types of motion of the suprahyoid muscle are more perceptually salient than the average neuromuscular activation). Conceptual frames are likewise unavoidable in designing with biodata in the artefacts described in Sanches et al. [105], no matter how that data might be presented back to the people whose bodies produced it. The same is true of any data-enabled design process. Messiness is not necessarily the default entropic state of a data-enabled artefact; it is something to work at, and perhaps requires identification of the places we inevitably fail to be messy.

5.2 Towards Design for Indeterminacy

In Section 2.2.3, we proposed that the data-enabled artefact acts as a Baradian apparatus — that is, the means of distilling and reconceptualising particular aspects of the messy world — brings into existence the very phenomena it purports to measure. The apparatus, through its observation of the phenomenon it is entangled with, collapses the indeterminacy of the phenomenon but gives rise to local patterns of cause and effect and also local patterns of ambiguity. In attempting to create representations and mappings of complex, entangled relationships, we create new, entangled relationships. This view suggests an alternate path for design researchers which might sidestep some of the debates around clean or messy data: *don't observe what you don't need.*

Unlike music informatics, the conceptual knowledge contained within a digital musical instrument is often a means to an end, an engineering convenience to support conceptual design processes such as mapping. Whether the performer, listener, or any other entity understands the world on the same terms as the instrument is of secondary importance at best. Indeed, creating and exploiting tensions between human and machine understanding is a common form of artistic practice or design provocation [25, 27, 31, 115]. In DMI performance, it is perhaps more the rule than the exception that the performer pushes back against the apparently limited affordances of the instrument to define a personal artistic space of possibilities [76, 129].

Nonetheless, drawing a Baradian analogy from quantum physics, what has been observed can no longer be returned to its indeterminate state. The act of observing requires a specific agential cut to be made, where measuring certain aspects means we are unable to measure others. For example, in the (Un)ambiguous Guitar (Section 4.3), providing different sonic behaviours for upward and downward bends was the desired outcome. However, along the way, the implementation depended on observations and models of musical gesture, 'identifying' (and thereby enacting) boundaries as to where bending and plucking actions begin and end. Having brought those boundaries into being, the question of precisely where to place them can become preoccupying. Every possible answer provides a different ambiguity shift, where some apparently similar actions will fall on opposite sides of an artificial boundary. Players of the (Un)ambiguous Guitar often noted quirks in behaviour related to these boundaries. Each of the three case studies show how performers' existing skill and expectations recontextualise the designer's approach to interpretation of phenomena, generating an ambiguity shift. The designers do not have the possibility of dictating how the data will be interpreted, even in the cases where the designers are also the performers. Existing musical practice and performance knowledge act as an additional agent, pulling the end-user's interpretation in individual ways.

What would it mean to leave some things deliberately indeterminate in a data-enabled artefact? One temptation might be to work with apparently simpler data, for example designing with time-domain signals without post-processing or feature extraction. However, even the collection of the signal itself constitutes a form of observation and hence an agential cut. In musical programming practice, providing ostensibly low-level languages and starting from a blank page do not necessarily result in culturally-neutral or aesthetically-varied music making, but rather in aesthetics that draw heavily on readily available primitive building blocks such as sine oscillators and white noise [80, 110]. Reductionism is not the same as indeterminacy.

Many critical design practices have been proposed in HCI which could provide alternate sources of inspiration, including anti-solutionism [19], removal of technologies [54], embracing failure [58], or simply not designing [13]. Leaving aside the null solution of avoiding data entirely, a productive mindset could be to simply *not care* about any intrinsic conceptual meaning of the data or its derived features. This outlook resonates with Giaccardi's infrastructural shifts as creating "unstable forms of value" [46] and with Sanches et. al's rejection of "actionable insight" [105]. Value would be ascribed based on the individual interpretation, which is the designer may work to constrain or guide but ultimately cannot dictate.

In DMI design, note onset detection stands out as a prime candidate for replacement by such *non-insightful data*. As discussed in Section 3.3, a symbolic or linguistic understanding of 'onset' is neither necessary nor particularly useful for musicians, since meaning is formed in the context of performing and listening. The meaninglessness of 'onset' holds even for notated music from the Western classical canon, where the sonic rendering of boundaries between notated symbols is often ambiguous, especially on wind or string instruments, and where what is written as several discrete note events might be perceived as a single event (e.g., tremolo on violin or glissando on harp). In these cases, it is not helpful for the instrument to enact any conceptual boundary between an onset or non-onset. We might go so far as to say that increasing the quantity or specificity of conceptual understanding encoded in the digital instrument proportionally *decreases* the specificity of meaning the musician can construct for themselves because of the way ambiguity is shifted: smarter instruments do not produce smarter players, nor smarter music.

Nonetheless, homogeneous featureless spaces tend to be boring, and creative practitioners are often drawn toward edges and discontinuities [86]. Thus, we speculate that data models and features without clear externally-referential meaning could help support the creation of interesting, idiosyncratic and sometimes ambiguous interactions [73]. AI models like the DDSP framework [39] found in the Tone Transfer project are a partial step in that direction, since models are trained implicitly from a corpus of examples and because they often contain internal latent spaces that defy conceptual explanation. However, AI systems are still at the mercy of how they metricise distance (or loss) or the purpose of training, which bring explicit concepts into play through the back door. If the inclusion of additional frames in depicting a phenomenon poses the risk of reducing space for personal meaning to emerge, and the omission of explanations may lead to an uninteresting outcome, the solution lies not in adding more or less but in reevaluating the agential cuts. Rather than focusing on the number of cuts, our attention should be directed towards understanding the ambiguity shift they create. We can then choose which agential cuts can be softened with increased ambiguity, allowing for a system that can still be engineered but that permits a wider interpretive stance.

5.3 Incentives and Modes of Research

If we can indeed produce rich entanglements by designing digital artefacts which deliberately avoid distilling well-defined conceptual knowledge as part of their operation, we face a challenge. The pull of *generalisable knowledge* is deeply ingrained across research domains, including in HCI [41, 55]. In engineering, this often unfolds through a logic of iteration and comparison, where the value of new research is demonstrated by improvements over a previous state of the art on accepted standard evaluations.

If an artefact does not distil or encode any concepts which have an external point of reference, and if conceptual knowledge is not even the intended output (as in the case of musical instruments), how do we demonstrate the value of the research contribution? How can we even make a convincing case that the artefact is working correctly, or that others should learn anything from our design decisions? The positivist approach, dominant in early HCI but still prevalent in the community, shifts the problem toward empirical user studies: the value of a system is demonstrated not in the design process, but in how it is used.

Research through design, in arguing that knowledge is created through the design process itself, has led to an *infrastructural shift* [46] toward other forms of contribution, including pictorials, annotated portfolios [24], accounts of 'ultimate particulars' [114] and circuitous or unplanned research journeys [48]. Practice-led research presents similar aims [108], making arts practice "not only the result of the research, but also its methodological vehicle, when the research unfolds in and through the acts of creating and performing" [21].

Incorporating agential realist viewpoints into such frameworks, we suggest there may exist value in what we *don't* conceptualise as well as what we do, and that publications can highlight areas where meaning has been deliberately left indeterminate – not just to the user, but in the treatment of the data within a digital system. Far from being an omission, to leave certain qualities indeterminate can be a crucial part of apparatus design which generates and preserves rich and productive forms of ambiguity.

5.4 Examining Approaches to Ambiguity

Our three case studies varied in their level of encoded symbolic representations, from a relatively open-ended interpretation of neuromuscular data in VoxEMG to familiar (though by no means universal) metrics of pitch and loudness in the Tone Transfer project to discretised and labelled gestures in the (Un)ambiguous Guitar. The relationship between the designer and players also differed, with varying degrees of autoethnographic study (Tone Transfer), open-ended use in the wild (VoxEMG) and structured laboratory studies ((Un)ambiguous Guitar). In each case, we see that the artefacts are far from passive observers of pre-existing qualities: the very actions they attempt to measure (vocal muscular movements, note onsets, plucking gestures) end up changing significantly under the influence of the artefact.

We speculate that this is a common result of any artefact that uses real-time data to change its behaviour in the world: the more important a given quality is to the measurement apparatus, the more likely that quality is to change as a result of performing the measurement. This effect stems from the balance between the two loci of meaning in Section 3.2. Perhaps the more tightly the designer attempts to forge determinate symbolic representations within the artefact's data processing system, the more difficult it becomes to allow space for meaning-making within a more holistic context of use. In a kind of conservation of ambiguity, unequivocal insistence on the meaning of data within an artefact might lead inevitably to a user perception of glitchiness or over-determination. Taken this way, the ambiguity shift can be problematic, leading to uncontrolled or counterproductive forms of ambiguity and the perception of poorly-operating technologies. As Gaver et al. caution [44, p. 240], "Many ambiguous systems are merely confusing, frustrating, or meaningless."

On the other hand, the same authors argue that ambiguity can be empowering and encourage critical questioning and reflection, suggesting perspectives without imposing solutions. In that case, an ambiguity shift could be useful, and theorising it through agential realism might suggest new design approaches to explore different forms of ambiguity. Without attempting to proclaim general rules, we would argue for a letting go of the need to precisely and deterministically understand the world within the symbolic system of an artefact, and to make room for competing loci of meaning within and outside the artefact without attempting to prescribe a single correct viewpoint. If it is not important for humans to understand the world on identical terms as symbolic representational computing systems, then spaces of possibility remain open for interpretation, curiosity and exploration.

6 CONCLUSION

Ambiguity can be both a resource and a challenge for designers and users of digital artefacts. Using Barad's agential realism as a guide, and drawing analogies to Giaccardi's three shifts in data-enabled RtD practice, this paper has explored how ambiguity is (re)configured as part of a process of designing with data. We chart the subtle relationship between ambiguity as a facet of human experience and deeper issues of ontological indeterminacy, arguing that data-enabled artefacts act as *apparatuses* in the Baradian sense, inextricably entangled with the phenomena they purport to measure. To design a data-enabled artefact involves making certain things determinate while precluding other things from becoming determinate, in what Barad calls an *agential cut*. It also involves making certain concepts or qualities explicit while leaving others implicit. As we show, even the most deterministic technical system is not necessarily free of ambiguity. Rather, designing data-enabled artefacts inevitably entails an *ambiguity shift* with potentially far-reaching consequences for interaction.

We have exemplified our ideas through three case studies, all pertaining to the design of digital musical instruments (DMIs) but otherwise varied in their contexts and methods. The tacit and aesthetic facets of musical experience provide a ready canvas for exploring ambiguity, but in Section 5 we address a wider community of HCI researchers working with data; our suggestions include paying close attention to what is made explicit versus implicit, allowing human and machine understandings to diverge, and avoiding over-measurement and over-conceptualisation to leave space for productive indeterminacy.

ACKNOWLEDGMENTS

AB and FC are supported by the UKRI Centre for Doctoral Training in Artificial Intelligence and Music (EP/S022694/1). AM is supported by a UKRI Frontier Research grant (EP/X023478/1, "RUDIMENTS: Reflective Understanding of Digital Instruments as Musical Entanglements").

REFERENCES

[1] Philip E Agre. 1997. Toward a critical technical practice: Lessons learned in trying to reform AI. *Social science, technical systems, and cooperative work: Beyond the Great Divide* 131 (1997).

[2] M Akrich. 1992. The De-Scription of Technical Objects. *Shaping Technology Building Society: Studjes in Sociotechnical Change* (1992), 205–224.

[3] Miquel Alfaras, Vasiliki Tsaknaki, Pedro Sanches, Charles Windlin, Muhammad Umair, Corina Sas, and Kristina Höök. 2020. From Biodata to Somadata. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3313831.3376684

[4] Kristina Andersen, Ron Wakkary, Laura Devendorf, and Alex McLean. 2019. Digital Crafts-Machine-Ship: Creative Collaborations with Machines. *Interactions* 27, 1 (Dec 2019), 30–35. https://doi.org/10.1145/3373644

[5] Jean-Julien Aucouturier and Emmanuel Bigand. 2013. Seven Problems That Keep MIR from Attracting the Interest of Cognition and Neuroscience. *Journal of Intelligent Information Systems* 41, 3 (Dec. 2013), 483–497. https://doi.org/10.1007/s10844-013-0251-x

[6] Karen Barad. 1998. Getting Real: Technoscientific Practices and the Materialization of Reality. *Differences: a Journal of Feminist Cultural Studies* 10, 2 (1998), 87–91.

[7] Karen Barad. 2003. Posthumanist Performativity: Toward an Understanding of How Matter Comes to Matter. *Signs* 28, 3 (2003), 801–831. https://doi.org/10.1086/345321

[8] Karen Barad. 2007. *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning*. Duke University Press.

[9] Karen Barad. 2014. Diffracting Diffraction: Cutting Together-Apart. *Parallax* 20, 3 (July 2014), 168–187. https://doi.org/10.1080/13534645.2014.927623

[10] Shaowen Bardzell. 2010. Feminist HCI: Taking Stock and Outlining an Agenda for Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) *(CHI '10)*. Association for Computing Machinery, New York, NY, USA, 1301–1310. https://doi.org/10.1145/1753326.1753521

[11] Ditte Amund Basballe and Kim Halskov. 2012. Dynamics of research through design. In *Proceedings of the Designing Interactive Systems Conference* (Newcastle Upon Tyne, United Kingdom) *(DIS '12)*. Association for Computing Machinery, New York, NY, USA, 58–67. https://doi.org/10.1145/2317956.2317967

[12] Eliot Bates. 2012. The social life of musical instruments. *Ethnomusicology* 56, 3 (2012), 363–395.

[13] Eric P.S. Baumer and M. Six Silberman. 2011. When the implication is not to design (technology). In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) *(CHI '11)*. Association for Computing Machinery, New York, NY, USA, 2271–2274. https://doi.org/10.1145/1978942.1979275

[14] Steve Benford, Adrian Hazzard, Alan Chamberlain, Kevin Glover, Chris Greenhalgh, Liming Xu, Michaela Hoare, and Dimitrios Darzentas. 2016. Accountable Artefacts: The Case of the Carolan Guitar. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 1163–1175. https://doi.org/10.1145/2858036.2858306

[15] Steve Benford, Holger Schnädelbach, Boriana Koleva, Rob Anastasi, Chris Greenhalgh, Tom Rodden, Jonathan Green, Ahmed Ghali, Tony Pridmore, Bill Gaver, Andy Boucher, Brendan Walker, Sarah Pennington, Albrecht Schmidt, Hans Gellersen, and Anthony Steed. 2005. Expected, sensed, and desired: A framework for designing sensing-based interaction. *ACM Trans. Comput.-Hum. Interact.* 12, 1 (mar 2005), 3–30. https://doi.org/10.1145/1057237.1057239

[16] Adan L. Benito Temprano, Teodoro Dannemann, and Andrew McPherson. 2023. Exploring the (un)ambiguous guitar: A Qualitative Study on the use of Gesture Disambiguation in Augmented Instrument Design. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Miguel Ortiz and Adnan Marquez-Borbon (Eds.). Mexico City, Mexico, Article 61, 10 pages. http://nime.org/proceedings/2023/nime2023_61.pdf

[17] Adan L. Benito Temprano and Andrew McPherson. 2021. A TMR Angle Sensor for Gesture Acquisition and Disambiguation on the Electric Guitar. In *Proceedings of the 16th International Audio Mostly Conference* (virtual/Trento, Italy) *(AM '21)*. Association for Computing Machinery, New York, NY, USA, 256–263. https://doi.org/10.1145/3478384.3478427

[18] Jesse Josua Benjamin, Arne Berger, Nick Merrill, and James Pierce. 2021. Machine Learning Uncertainty as a Design Material: A Post-Phenomenological Inquiry. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA) *(CHI '21)*. Association for Computing Machinery, 1–14. https://doi.org/10.1145/3411764.3445481

[19] Mark Blythe, Kristina Andersen, Rachel Clarke, and Peter Wright. 2016. Anti-Solutionist Strategies: Seriously Silly Design Fiction. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 4968–4978. https://doi.org/10.1145/2858036.2858482

[20] Kirsten Boehner and Jeffrey T. Hancock. 2006. Advancing ambiguity. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Montréal, Québec, Canada) *(CHI '06)*. Association for Computing Machinery, New York, NY, USA, 103–106. https://doi.org/10.1145/1124772.1124789

[21] Henk Borgdorff. 2010. The production of knowledge in artistic research. In *The Routledge companion to research in the arts*. Routledge, 44–63.

[22] Georgina Born. 2020. Diversifying MIR: Knowledge and real-world challenges, and new interdisciplinary futures. *Transactions of the International Society for Music Information Retrieval* 3, 1 (2020).

[23] Georgina Born (Ed.). 2022. *Music and Digital Media: A Planetary Anthology*. UCL Press. https://doi.org/10.14324/111.9781800082434

[24] John Bowers. 2012. The logic of annotated portfolios: communicating the value of 'research through design'. In *Proceedings of the Designing Interactive Systems Conference* (Newcastle Upon Tyne, United Kingdom) *(DIS '12)*. Association for Computing Machinery, New York, NY, USA, 68–77. https://doi.org/10.1145/2317956.2317968

[25] John Bowers and Owen Green. 2018. All the Noises: Hijacking Listening Machines for Performative Research. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Thomas Martin Luke Dahl, Douglas Bowman (Ed.). Virginia Tech, Blacksburg, Virginia, USA, 114–119. https://doi.org/10.5281/zenodo.1302699

[26] Claude Cadoz. 1988. Instrumental Gesture and Musical Composition. In *ICMC 1988 - International Computer Music Conference*. Cologne, Germany, 1–12. https://hal.science/hal-00491738/file/Cad88_Conf_ICMC.pdf

[27] Kim Cascone. 2000. The Aesthetics of Failure: "Post-Digital" Tendencies in Contemporary Computer Music. *Computer Music Journal* 24, 4 (2000), 12–18.

[28] Franco Caspe, Andrew McPherson, and Mark Sandler. 2022. DDX7: DIFFERENTIABLE FM SYNTHESIS OF MUSICAL INSTRUMENT SOUNDS. In *Proceedings of the 23rd ISMIR Conference* (Bengaluru, India). International Society of Music Information Retrieval (ISMIR).

[29] Franco Caspe, Andrew McPherson, and Mark Sandler. 2023. FM Tone Transfer with Envelope Learning. In *Proceedings of the 18th International Audio Mostly Conference* (Edinburgh, United Kingdom) *(AM '23)*. Association for Computing Machinery, New York, NY, USA, 116–123. https://doi.org/10.1145/3616195.3616196

[30] Francis Charpentier. 1986. Pitch detection using the short-term phase spectrum. In *ICASSP'86. IEEE International Conference on Acoustics, Speech, and Signal Processing*, Vol. 11. IEEE, 113–116.

[31] Nicolas Collins. 2020. *Handmade electronic music: the art of hardware hacking*. Taylor & Francis.

[32] Roddy Cowie, Gary McKeown, and Ellen Douglas-Cowie. 2012. Tracing Emotion: An Overview. *Int. J. Synth. Emot.* 3, 1 (jan 2012), 1–17. https://doi.org/10.4018/jse.2012010101

[33] Laura Devendorf, Joanne Lo, Noura Howell, Jung Lin Lee, Nan-Wei Gong, M. Emre Karagozler, Shiho Fukuhara, Ivan Poupyrev, Eric Paulos, and Kimiko Ryokai. 2016. "I Don't Want to Wear a Screen": Probing Perceptions of and Possibilities for Dynamic Displays

on Clothing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 6028–6039. https://doi.org/10.1145/2858036.2858192

[34] Kristin N. Dew and Daniela K. Rosner. 2018. Lessons from the Woodshop: Cultivating Design with Living Materials. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal, QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3173574.3174159

[35] Ryan Diduck. 2018. *Mad skills: MIDI and music technology in the twentieth century*. Watkins Media Limited.

[36] Emily I. Dolan. 2012. Toward a Musicology of Interfaces. *Keyboard Perspectives* 5 (2012), 1–13.

[37] John L Drever and Andrew Hugill. 2022. *Aural Diversity*. Taylor & Francis.

[38] Alexander Eiselmayer, Chat Wacharamanotham, Michel Beaudouin-Lafon, and Wendy E. Mackay. 2019. Touchstone2: An Interactive Environment for Exploring Trade-offs in HCI Experiment Design. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–11. https://doi.org/10.1145/3290605.3300447

[39] Jesse Engel, Lamtharn Hantrakul, Chenjie Gu, and Adam Roberts. 2020. DDSP: Differentiable Digital Signal Processing. In *8th International Conference on Learning Representations* (Addis Ababa, Ethiopia). ICLR, 19. http://arxiv.org/abs/2001.04643

[40] Melanie Feinberg. 2017. A Design Perspective on Data. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 2952–2963. https://doi.org/10.1145/3025453.3025837

[41] Christopher Frauenberger. 2019. Entanglement HCI The Next Wave? *ACM Trans. Comput.-Hum. Interact.* 27, 1, Article 2 (nov 2019), 27 pages. https://doi.org/10.1145/3364998

[42] Mikhaila Friske, Jordan Wirfs-Brock, and Laura Devendorf. 2020. Entangling the Roles of Maker and Interpreter in Interpersonal Data Narratives: Explorations in Yarn and Sound. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference* (Eindhoven, Netherlands) *(DIS '20)*. Association for Computing Machinery, New York, NY, USA, 297–310. https://doi.org/10.1145/3357236.3395442

[43] William Gaver, Peter Gall Krogh, Andy Boucher, and David Chatting. 2022. Emergence as a Feature of Practice-Based Design Research. In *Designing Interactive Systems Conference* (Virtual Event, Australia) *(DIS '22)*. Association for Computing Machinery, New York, NY, USA, 517–526. https://doi.org/10.1145/3532106.3533524

[44] William W. Gaver, Jacob Beaver, and Steve Benford. 2003. Ambiguity as a Resource for Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, Florida, USA) *(CHI '03)*. Association for Computing Machinery, New York, NY, USA, 233–240. https://doi.org/10.1145/642611.642653

[45] William W. Gaver, Andrew Boucher, Sarah Pennington, and Brendan Walker. 2004. Cultural probes and the value of uncertainty. *Interactions* 11, 5 (sep 2004), 53–56. https://doi.org/10.1145/1015530.1015555

[46] Elisa Giaccardi. 2019. Histories and Futures of Research through Design: From Prototypes to Connected Things. *International Journal of Design* 13, 3 (2019), 139–155.

[47] Lisa Gitelman, Virginia Jackson, Daniel Rosenberg, Travis D. Williams, Kevin R. Brine, Mary Poovey, Matthew Stanley, Ellen Gruber Garvey, Markus Krajewski, Rita Raley, David Ribes, Steven J. Jackson, and Geoffrey C. Bowker. 2013. *Data Bite Man: The Work of Sustaining a Long-Term Study*. 147–166.

[48] Bruna Goveia da Rocha and Kristina Andersen. 2020. Becoming Travelers: Enabling the Material Drift. In *Companion Publication of the 2020 ACM Designing Interactive Systems Conference* (Eindhoven, Netherlands) *(DIS' 20 Companion)*. Association for Computing Machinery, New York, NY, USA, 215–219. https://doi.org/10.1145/3393914.3395881

[49] Michael Gurevich. 2016. Diversity in NIME research practices. *Leonardo* 49, 1 (2016), 80–81.

[50] Donna J Haraway and Thyrza Nichols Goodeve. 2018. *FemaleMan© _Meets_OncoMouse™: Mice Into Wormholes: A Technoscience Fugue in Two Parts*. Routledge. 46–118 pages.

[51] Steve Harrison, Deborah Tatar, and Phoebe Sengers. 2007. The Three Paradigms of HCI. In *Alt. Chi. Session at the SIGCHI Conference on Human Factors in Computing Systems, San Jose, California, USA*. 1–18.

[52] Cher M Hill. 2017. More-than-reflective practice: Becoming a diffractive practitioner. *Teacher Learning and Professional Development* 2, 1 (2017).

[53] Sarah Homewood. 2018. Designing for the Changing Body: A Feminist Exploration of Self-Tracking Technologies. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI EA '18)*. Association for Computing Machinery, New York, NY, USA, 1–4. https://doi.org/10.1145/3170427.3173031

[54] Sarah Homewood, Amanda Karlsson, and Anna Vallgårda. 2020. Removal as a Method: A Fourth Wave HCI Approach to Understanding the Experience of Self-Tracking. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference* (Eindhoven, Netherlands) *(DIS '20)*. Association for Computing Machinery, New York, NY, USA, 1779–1791. https://doi.org/10.1145/3357236.3395425

[55] Kristina Höök and Jonas Löwgren. 2012. Strong concepts: Intermediate-level knowledge in interaction design research. *ACM Trans. Comput.-Hum. Interact.* 19, 3, Article 23 (oct 2012), 18 pages. https://doi.org/10.1145/2362364.2362371

[56] Noura Howell. 2016. Representation & Interpretation of Biosensing. In *Proceedings of the 2016 ACM Conference Companion Publication on Designing Interactive Systems* (Brisbane, QLD, Australia) *(DIS '16 Companion)*. Association for Computing Machinery, New York, NY,

USA, 31–33. https://doi.org/10.1145/2908805.2909422

[57] Noura Howell, John Chuang, Abigail De Kosnik, Greg Niemeyer, and Kimiko Ryokai. 2018. Emotional Biosensing: Exploring Critical Alternatives. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 69 (Nov 2018), 25 pages. https://doi.org/10.1145/3274338

[58] Noura Howell, Audrey Desjardins, and Sarah Fox. 2021. Cracks in the Success Narrative: Rethinking Failure in Design Research through a Retrospective Trioethnography. *ACM Trans. Comput.-Hum. Interact.* 28, 6, Article 42 (nov 2021), 31 pages. https://doi.org/10.1145/3462447

[59] Noura Howell, Laura Devendorf, Rundong (Kevin) Tian, Tomás Vega Galvez, Nan-Wei Gong, Ivan Poupyrev, Eric Paulos, and Kimiko Ryokai. 2016. Biosignals as Social Cues: Ambiguity and Emotional Interpretation in Social Displays of Skin Conductance. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems* (Brisbane, QLD, Australia) *(DIS '16)*. Association for Computing Machinery, New York, NY, USA, 865–870. https://doi.org/10.1145/2901790.2901850

[60] Noura Howell, Laura Devendorf, Tomás Alfonso Vega Gálvez, Rundong Tian, and Kimiko Ryokai. 2018. Tensions of Data-Driven Reflection: A Case Study of Real-Time Emotional Biosensing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3173574.3174005

[61] Noura Howell, Greg Niemeyer, and Kimiko Ryokai. 2019. Life-Affirming Biosensing in Public: Sounding Heartbeats on a Red Bench. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland, UK) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–16. https://doi.org/10.1145/3290605.3300910

[62] Andy Hunt, Marcelo M Wanderley, and Matthew Paradis. 2003. The importance of parameter mapping in electronic instrument design. *Journal of New Music Research* 32, 4 (2003), 429–440.

[63] Don Ihde. 1990. *Technology and the lifeworld: From garden to earth.* Indiana University Press.

[64] Tim Ingold. 2017. *Correspondences.* University of Aberdeen.

[65] Robert Jack, Jacob Harrison, and Andrew McPherson. 2020. Digital Musical Instruments as Research Products. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Romain Michon and Franziska Schroeder (Eds.). Birmingham City University, Birmingham, UK, 446–451. https://doi.org/10.5281/zenodo.4813465

[66] Robert H Jack, Adib Mehrabi, Tony Stockman, and Andrew McPherson. 2018. Action-sound latency and the perceived quality of digital musical instruments: Comparing professional percussionists and amateur musicians. *Music Perception: An Interdisciplinary Journal* 36, 1 (2018), 109–128.

[67] Robert H Jack, Tony Stockman, and Andrew McPherson. 2017. Rich gesture, reduced control: the influence of constrained mappings on performance technique. In *Proceedings of the 4th International Conference on Movement Computing*. 1–8.

[68] Jean-Julien Aucouturier and Emmanuel Bigand. 2012. Mel Cepstrum & Ann Ova: The Difficult Dialog Between Mir And Music Cognition. (2012). https://doi.org/10.5281/ZENODO.1417179

[69] Sergi Jordà. 2004. Instruments and players: Some thoughts on digital lutherie. *Journal of New Music Research* 33, 3 (2004), 321–341.

[70] Elvin Karana, Elisa Giaccardi, Niels Stamhuis, and Jasper Goossensen. 2016. The Tuning of Materials: A Designer's Journey. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems* (Brisbane, QLD, Australia) *(DIS '16)*. Association for Computing Machinery, New York, NY, USA, 619–631. https://doi.org/10.1145/2901790.2901909

[71] Kevin Kilgour, Mauricio Zuluaga, Dominik Roblek, and Matthew Sharifi. 2019. Fréchet audio distance: A reference-free metric for evaluating music enhancement algorithms.. In *INTERSPEECH*. 2350–2354.

[72] Jong Wook Kim, Justin Salamon, Peter Li, and Juan Pablo Bello. 2018. Crepe: A convolutional representation for pitch estimation. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 161–165.

[73] Giacomo Lepri and Andrew McPherson. 2022. Embrace the Weirdness: Negotiating Values Inscribed into Music Technology. *Computer Music Journal* (2022), 1–19.

[74] Alexander Lerch, Claire Arthur, Ashis Pati, and Siddharth Gururani. 2020. An Interdisciplinary Review of Music Performance Analysis. 3, 1 (2020), 221–245. Issue 1. https://doi.org/10.5334/tismir.53

[75] Thor Magnusson. 2009. Of epistemic tools: Musical instruments as cognitive extensions. *Organised Sound* 14, 2 (2009), 168–176.

[76] Thor Magnusson. 2010. Designing constraints: Composing and performing with digital musical systems. *Computer Music Journal* 34, 4 (2010), 62–73.

[77] Thor Magnusson. 2019. *Sonic writing: technologies of material, symbolic, and signal inscriptions.* Bloomsbury Publishing USA.

[78] Paul Marshall and Eva Hornecker. 2013. Theories of Embodiment in HCI. *The SAGE handbook of digital technology research* 1 (2013), 144–158.

[79] Matthias Mauch and Simon Dixon. 2014. pYIN: A fundamental frequency estimator using probabilistic threshold distributions. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 659–663.

[80] Andrew McPherson and Giacomo Lepri. 2020. Beholden to our tools: negotiating with technology while sketching digital instruments. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Romain Michon and Franziska Schroeder (Eds.). Birmingham City University, Birmingham, UK, 434–439. https://doi.org/10.5281/zenodo.4813461

[81] Andrew McPherson and Koray Tahıroğlu. 2020. Idiomatic patterns and aesthetic influence in computer music languages. *Organised Sound* 25, 1 (2020), 53–63.
[82] Andrew McPherson and Victor Zappi. 2015. An Environment for Submillisecond-Latency Audio and Sensor Processing on BeagleBone Black. In *Audio Engineering Society Convention 138*. Audio Engineering Society.
[83] Eduardo Meneses, Sergio Freire, and Marcelo Wanderley. 2018. GuitarAMI and GuiaRT: two independent yet complementary projects on augmented nylon guitars. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Thomas Martin Luke Dahl, Douglas Bowman (Ed.). Virginia Tech, Blacksburg, Virginia, USA, 222–227. https://doi.org/10.5281/zenodo.1302559
[84] Maurice Merleau-Ponty. 2014. *Phenomenology of Perception.* Routledge.
[85] Landon Morrison and Andrew McPherson. 2024. Entangling Entanglement: A Diffractive Dialogue on HCI and Musical Interactions. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 493, 17 pages. https://doi.org/10.1145/3613904.3642171
[86] Tom Mudd, Simon Holland, and Paul Mulholland. 2019. Nonlinear dynamical processes in musical interactions: Investigating the role of nonlinear dynamics in supporting surprise and exploration in interactions with digital musical instruments. *International Journal of Human-Computer Studies* 128 (2019), 27–40.
[87] Chris Nash. 2015. The cognitive dimensions of music notations. In *Proceedings of the International Conference on Technologies for Notation and Representation (TENOR).*
[88] Luc Nijs, Micheline Lesaffre, and Marc Leman. [n. d.]. The Musical Instrument as a Natural Extension of the Musician. In *The 5th Conference of Interdisciplinary Musicology* (2009). LAM-Institut jean Le Rond d'Alembert, 132–133.
[89] Charlotte Nordmoen, Jack Armitage, Fabio Morreale, Rebecca Stewart, and Andrew McPherson. 2019. Making Sense of Sensors: Discovery Through Craft Practice With an Open-Ended Sensor Material. In *Proceedings of the 2019 on Designing Interactive Systems Conference* (San Diego, CA, USA) *(DIS '19)*. Association for Computing Machinery, New York, NY, USA, 135–146. https://doi.org/10.1145/3322276.3322368
[90] Charlotte Nordmoen and Andrew P. McPherson. 2022. Making Space for Material Entanglements: A Diffractive Analysis of Woodwork and the Practice of Making an Interactive System. In *Designing Interactive Systems Conference* (Virtual Event, Australia) *(DIS '22)*. Association for Computing Machinery, New York, NY, USA, 415–423. https://doi.org/10.1145/3532106.3533572
[91] Charlotte Nordmoen and Andrew P. McPherson. 2023. What the Sensor Knows: More-Than-Human Knowledge Co-Production in Wood Carving. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (Pittsburgh, PA, USA) *(DIS '23)*. Association for Computing Machinery, New York, NY, USA, 779–789. https://doi.org/10.1145/3563657.3596075
[92] William Odom, Richard Banks, Abigail Durrant, David Kirk, and James Pierce. 2012. Slow technology: critical reflection and future directions. In *Proceedings of the Designing Interactive Systems Conference* (Newcastle Upon Tyne, United Kingdom) *(DIS '12)*. Association for Computing Machinery, New York, NY, USA, 816–817. https://doi.org/10.1145/2317956.2318088
[93] Jessica O'Bryan. 2015. "We ARE our instrument!": Forming a singer identity. *Research Studies in Music Education* 37, 1 (June 2015), 123–137. https://doi.org/10.1177/1321103x15592831
[94] Claire Petitmengin. 2007. Towards the source of thoughts: The gestural and transmodal dimension of lived experience. *Journal of consciousness Studies* 14, 3 (2007), 54–82.
[95] Jonathan Pigrem, Andrew McPherson, Nick Bryan-Kinns, and Robert Jack. 2022. Sound -> Object -> Gesture: Physical Affordances of Virtual Materials. In *Proceedings of the 17th International Audio Mostly Conference* (St. Pölten, Austria) *(AM '22)*. Association for Computing Machinery, New York, NY, USA, 59–66. https://doi.org/10.1145/3561212.3561230
[96] Cornelius Poepel and Roger B Dannenberg. 2005. Audio signal driven sound synthesis. In *ICMC*.
[97] Johan Redström. 2006. Towards user design? On the shift from object to user as the subject of design. *Design studies* 27, 2 (2006), 123–139.
[98] Courtney N. Reed and Andrew P. McPherson. 2020. Surface Electromyography for Direct Vocal Control. In *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME '20)*, Romain Michon and Franziska Schroeder (Eds.). Birmingham City University, Birmingham, UK, 458–463. https://doi.org/10.5281/zenodo.4813475
[99] Courtney N. Reed and Andrew P. McPherson. 2021. Surface Electromyography for Sensing Performance Intention and Musical Imagery in Vocalists. In *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction* (Salzburg, Austria) *(TEI '21)*. Association for Computing Machinery, New York, NY, USA, Article 22, 11 pages. https://doi.org/10.1145/3430524.3440641
[100] Courtney N. Reed and Andrew P. McPherson. 2023. The Body as Sound: Unpacking Vocal Embodiment through Auditory Biofeedback. In *Proceedings of the Seventeenth International Conference on Tangible, Embedded, and Embodied Interaction* (Warsaw, Poland) *(TEI '23)*. Association for Computing Machinery, New York, NY, USA, Article 7, 15 pages. https://doi.org/10.1145/3569009.3572738
[101] Courtney N. Reed and Andrew P. McPherson. 2024. *VoxEMG.* https://github.com/courtcourtaney/voxEMG
[102] Courtney N. Reed, Landon Morrison, Andrew P. McPherson, David Fierro, and Atau Tanaka. 2024. Sonic Entanglements with Electromyography: Between Bodies, Signals, and Representations. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (IT University of Copenhagen, Denmark) *(DIS '24)*. Association for Computing Machinery, New York, NY, USA, 2691–2707. https://doi.org/10.1145/3643834.3661572

[103] Courtney N. Reed, Paul Strohmeier, and Andrew P. McPherson. 2023. Negotiating Experience and Communicating Information Through Abstract Metaphor. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 185, 16 pages. https://doi.org/10.1145/3544548.3580700

[104] Matthew Rodger, Paul Stapleton, Maarten van Walstijn, Miguel Ortiz, and Laurel S Pardue. 2020. What Makes a Good Musical Instrument? A Matter of Processes, Ecologies and Specificities . In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Romain Michon and Franziska Schroeder (Eds.). Birmingham City University, Birmingham, UK, 405–410. https://doi.org/10.5281/zenodo.4813438

[105] Pedro Sanches, Noura Howell, Vasiliki Tsaknaki, Tom Jenkins, and Karey Helms. 2022. Diffraction-in-action: Designerly Explorations of Agential Realism Through Lived Data. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 540, 18 pages. https://doi.org/10.1145/3491102.3502029

[106] Donald A Schön. 1968. The Reflective Practitioner. *New York* (1968).

[107] Hugo Scurto, Baptiste Caramiaux, and Frederic Bevilacqua. 2021. Prototyping Machine Learning Through Diffractive Art Practice. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference* (Virtual Event, USA) *(DIS '21)*. Association for Computing Machinery, New York, NY, USA, 2013–2025. https://doi.org/10.1145/3461778.3462163

[108] Hazel Smith. 2009. *Practice-led research, research-led practice in the creative arts.* Edinburgh University Press.

[109] Julius Orion Smith. 2008. *Mathematics of the discrete Fourier transform (DFT): with audio applications.* Julius Smith.

[110] Joe Snape and Georgina Born. 2022. Max, Music Software, and the Mutual Mediation of Aesthetics and Digital Technologies. In *Music and Digital Media: A planetary anthropology*, Georgina Born (Ed.). UCL Press.

[111] Robert Soden, Laura Devendorf, Richmond Wong, Yoko Akama, and Ann Light. 2022. Modes of Uncertainty in HCI. *Foundations and Trends in Human-Computer Interaction* 15, 4 (2022), 317–426. https://doi.org/10.1561/1100000085

[112] Robert Soden, Laura Devendorf, Richmond Y. Wong, Lydia B. Chilton, Ann Light, and Yoko Akama. 2020. Embracing Uncertainty in HCI. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '20)*. Association for Computing Machinery, New York, NY, USA, 1–8. https://doi.org/10.1145/3334480.3375177

[113] Jonathan Sterne. 2012. *MP3: The meaning of a format.* Duke University Press.

[114] Erik Stolterman. 2008. The nature of design practice and implications for interaction design research. *International Journal of Design* 2, 1 (2008).

[115] Bob LT Sturm and Oded Ben-Tal. 2021. Folk the algorithms:(Mis) Applying artificial intelligence to folk music. *Handbook of Artificial Intelligence for Music: Foundations, Advanced Approaches, and Developments for Creativity* (2021), 423–454.

[116] Dag Svanæs. 2013. Interaction design for and with the lived body: Some implications of merleau-ponty's phenomenology. *ACM Trans. Comput.-Hum. Interact.* 20, 1, Article 8 (apr 2013), 30 pages. https://doi.org/10.1145/2442106.2442114

[117] Koray Tahıroğlu, Thor Magnusson, Adam Parkinson, Iris Garrelfs, and Atau Tanaka. 2020. Digital Musical Instruments as Probes: How Computation Changes the Mode-of-Being of Musical Instruments. *Organised Sound* 25, 1 (2020), 64–74. https://doi.org/10.1017/S1355771819000475

[118] Atau Tanaka. 2000. Musical performance practice on sensor-based instruments. *Trends in gestural control of music* 13, 389-405 (2000), 284.

[119] Kai Tuuri, Jaana Parvainen, and Antti Pirhonen. 2017. Who Controls Who? Embodied Control Within Human–Technology Choreographies. *Interacting with Computers* (2017), 1–18. https://doi.org/10.1093/iwc/iww040

[120] Anna Vallgårda and Ylva Fernaeus. 2015. Interaction Design as a Bricolage Practice. In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction* (Stanford, California, USA) *(TEI '15)*. Association for Computing Machinery, New York, NY, USA, 173–180. https://doi.org/10.1145/2677199.2680594

[121] Dylan van der Schyff. 2015. Music as a manifestation of life: exploring enactivism and the â€˜eastern perspectiveâ€™ for music education. *Frontiers in Psychology* 06 (March 2015). https://doi.org/10.3389/fpsyg.2015.00345

[122] 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

[123] Ron Wakkary, Doenja Oogjes, and Armi Behzad. 2022. Two Years or More of Co-speculation: Polylogues of Philosophers, Designers, and a Tilting Bowl. *ACM Trans. Comput.-Hum. Interact.* 29, 5, Article 47 (oct 2022), 44 pages. https://doi.org/10.1145/3514235

[124] 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* (Edinburgh, United Kingdom) *(DIS '17)*. Association for Computing Machinery, New York, NY, USA, 503–514. https://doi.org/10.1145/3064663.3064734

[125] Ron Wakkary, Doenja Oogjes, Henry W. J. Lin, and Sabrina Hauser. 2018. Philosophers Living with the Tilting Bowl. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3173574.3173668

[126] Marcelo Mortensen Wanderley and Nicola Orio. 2002. Evaluation of input devices for musical expression: Borrowing tools from HCI. *Computer Music Journal* 26, 3 (2002), 62–76.

[127] Xin Wang, Shinji Takaki, and Junichi Yamagishi. 2019. Neural Source-Filter-Based Waveform Model for Statistical Parametric Speech Synthesis. ICASSP. http://arxiv.org/abs/1810.11946

[128] Simon Waters. 2021. The entanglements which make instruments musical: Rediscovering sociality. *Journal of New Music Research* 50, 2 (2021), 133–146.

[129] Victor Zappi and Andrew McPherson. 2014. Dimensionality and Appropriation in Digital Musical Instrument Design. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Goldsmiths, University of London, London, United Kingdom, 455–460. https://doi.org/10.5281/zenodo.1178993

[130] Victor Zappi and Andrew McPherson. 2018. Hackable Instruments: Supporting Appropriation and Modification in Digital Musical Interaction. *Frontiers in ICT* 5 (Oct. 2018). https://doi.org/10.3389/fict.2018.00026

Received 6 July 2023; revised 5 May 2024; accepted 9 July 2024