



HAL
open science

Prototyping Machine Learning Through Diffractive Art Practice

Hugo Scurto, Baptiste Caramiaux, Frédéric Bevilacqua

► **To cite this version:**

Hugo Scurto, Baptiste Caramiaux, Frédéric Bevilacqua. Prototyping Machine Learning Through Diffractive Art Practice. Proceedings of the 2021 ACM Conference on Designing Interactive Systems (DIS 2021), Jun 2021, Virtual Event, Unknown Region. 10.1145/3461778.3462163 . hal-03267282

HAL Id: hal-03267282

<https://hal.science/hal-03267282v1>

Submitted on 22 Jun 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Prototyping Machine Learning Through Diffractive Art Practice

Hugo Scurto
Reflective Interaction group,
EnsadLab, EnsAD-PSL Univ.,
UMRS1158, Inserm-Sorbonne Univ.,
ISIR, CNRS-Sorbonne Univ.
Paris, France
hugo.scurto@ensad.fr

Baptiste Caramiaux
ISIR, CNRS-Sorbonne Université
Paris, France
baptiste.caramiaux@sorbonne-
universite.fr

Frédéric Bevilacqua
STMS, IRCAM-CNRS-Sorbonne
Université
Paris, France
Frederic.Bevilacqua@ircam.fr

ABSTRACT

In this paper, we outline a diffractive practice of machine learning (ML) in the frame of material-centered interaction design. To this aim, we review related work in ML, HCI, design, new interfaces for musical expression, and computational art, and introduce two practice-based studies of music performance and robotic art based on interactive machine learning tools, with the hope of revealing the computational materiality of ML, and the potential of embodiment to craft prototypes of ML that reconfigure conceptual or technical approaches to ML. We derive five interference conditions for such art-based ML prototypes—*situational whole*, *small data*, *shallow model*, *learnable algorithm*, and *somaesthetic behaviour*—and describe their widening of design and engineering practices of ML prototyping. Finally, we sketch how a process of *intra-active machine learning* could complement that of interactive machine learning to take materiality as an entry point for ML design within HCI.

CCS CONCEPTS

• **Applied computing** → **Media arts; Sound and music computing**; • **Human-centered computing** → **Interface design prototyping**.

KEYWORDS

Machine Learning; Art Practice; Design; Diffractive Methods.

ACM Reference Format:

Hugo Scurto, Baptiste Caramiaux, and Frédéric Bevilacqua. 2021. Prototyping Machine Learning Through Diffractive Art Practice. In *Designing Interactive Systems Conference 2021 (DIS '21)*, June 28–July 2, 2021, Virtual Event, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3461778.3462163>

1 INTRODUCTION

In recent years, design research started approaching machine learning (ML) as a design material [89]. Grounded in studies of how user experience designers effectively collaborate with ML engineers [27, 91], it seeks to develop conceptual frameworks to envision divergent and socially-situated designs for ML [9, 26]. Yet, design research still lacks specific tools to materialise such ML concepts into design artifacts [27]. Researchers stated that “ML challenges the

general idea of prototyping” [89], which is certainly true when considering big data and deep learning approaches pushed by modern ML engineering, but seems to overlook alternative implementations of ML that could emerge from creative practice with ML. For example, performers in the field of NIME—acronym for *New Interfaces for Musical Expression* [67]—developed an interactive account of ML, which can, for instance, be used to tacitly train ML models using small, user-provided data, to craft custom gestural controllers of sound [11, 30, 53]. More generally, embodied or somaesthetic design approaches could help design research go beyond language and logic to describe and explore ML as design material [48].

We believe that art practices that use code and ML as creative material can be one such embodied approach to implement concrete prototypes of ML. In code-based art practice [55], artists approach ML not only as a technique to engineer logical concepts or functionalities—such as automatic classification or reward maximisation by learning agents—; they also approached ML as a *computational material*, whose raw properties, such as adaptive learning [20], model extrapolation [32], algorithmic exploration [76], or probabilistic uncertainty [9], can be crafted and experienced within installations, performances, and other hybrid embodiments [1, 5, 25, 39, 64]. By focusing on ML materiality, code-based art practice may thus help design research explore ML for what it is—*i.e.*, a set of computational material possessing specific properties—rather than on what it is currently used for—*e.g.*, a set of computational techniques contributing to socio-cultural discourses on artificial intelligence.

In this paper, we rely on our own art practice and that of other artists to further understandings of ML materiality for design. Specifically, we propose to explore *diffractive methods* [6] to frame analysis and production of art-based ML prototypes within design research, by specially attending and responding to social discourses and material configurations in which ML is embedded. According to Karen Barad’s material feminist theory [6], diffraction intends to displace reflection, which assumes pre-existing subjects and objects interacting with each others, as a dominant model of inquiry. It does so by assuming that humans and non-humans are bound together within complex *socio-material* practices, which are fluid and ever evolving; and refers to their entanglement as *intra-action*, as opposed to interaction [6]. As such, we believe that diffractive methods may help design research to consider performative and transformative socio-material phenomena of code-based art practice, such as artists becoming engineers, algorithms living things, or humans embodying computational material [68, 69], as part of ML prototyping practices. By doing so, we do not want to impose a

DIS '21, June 28–July 2, 2021, Virtual Event, USA

© 2021 Association for Computing Machinery.

This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Designing Interactive Systems Conference 2021 (DIS '21)*, June 28–July 2, 2021, Virtual Event, USA, <https://doi.org/10.1145/3461778.3462163>.

static understanding of the theoretical interest of diffractive practices for design research. Rather, our wish is to start a dialogue with designers and researchers about the generative potentials of diffractive practices, and how they may expand or restrain more widespread methods of participatory design or material-centered interaction design within HCI. In this, we continue prior work from the HCI community on exploring art-based methods to reconfigure core computing notions and enrich interaction design [8, 29, 50].

To do so, we provide a review of ML practices within art, design, engineering and HCI, and two stabilised accounts of instances of art-based ML prototyping of our own, which we refer to as *diffractive art practice*, a term recently introduced by artists-researchers Jane Prophet and Helen Pritchard to describe code-based art practices that attend and respond to entanglements of social and technical—thus material—aspects of computation [68, 69]. These two interrelated empirical studies explore diffractive methods in two complementary ways for ML-based art practice. In the first study, diffraction is used as a method to analyse the role that ML materiality played along the prototyping of a NIME instrument. In the second study, diffraction is used as a method to practice with ML and other practitioners and material agencies along the prototyping of a robotic art installation.

Our work led us to frame art-based ML prototyping as (a) a mode of craft that takes embodiment of ML computational properties as entry point for design and engineering, and (b) a socio-material practice that reveals and reconfigure the fluid boundaries between humans and ML technology. Specifically, our two studies suggest five conditions for art-based ML prototyping to interfere with conceptual and technical approaches to ML and widen design or engineering practices of ML prototyping: namely, *situational whole*, *small data*, *shallow model*, *learnable algorithm*, and *somaesthetic behaviour*. Finally, we rely on Barad’s notion of *intra-action* [6] to sketch how a process of *intra-active machine learning* may help to analyse practices of ML prototyping from a more-than-human perspective—by analogy with human-centred perspectives on ML brought by interactive machine learning [43, 70]. To sum up our contributions, we: (1) present two artistic works made by crafting and experimenting with ML materiality; (2) explore diffraction as conceptual and methodological framework for both art practice with ML and analysis of art practice with ML in the frame of HCI; and (3) identify five socio-technical conditions that contribute to an understanding of ML prototypes as valid research sites in wider material-centered interaction design.

2 RELATED WORK

In this section, we start by introducing elements of ML that constitutes its materiality. Then, we review previous work investigating ML as design material, with its promises and caveats. We then present how ML can be considered as creative material in art practice. We finally describe the methodological concepts that we bring forward to frame code-based art practice within design research.

2.1 Material Elements of ML

The next paragraphs introduce components of ML as defined by engineering approaches. These definitions are useful to understand how art practice has enabled to switch from a technical perspective

to a *material* perspective on ML components, and how this switching may imply adopting novel methods for design. We identified four elements constitutive of ML materiality (schematised in Figure 1): *techniques*, *data*, *models* and *algorithms*. Each element possesses different structures on their own, as well as different functionalities related to others.




Figure 1: Diagram for ML materiality. In bold: elements; in italics: structures; in grey arrows: functionalities.

2.1.1 Techniques. ML defines different *techniques* that enable computers to perform specific tasks in relation to a goal by directly learning from data. In *supervised learning*, the goal is to learn a function that maps inputs to outputs based on example pairs of input-output data. The learned function enables to execute tasks such as automatic classification and recognition of new input data. *Unsupervised learning* relies on input data only: the goal is to learn patterns within the data to enable complex tasks such as clustering or generation. In *reinforcement learning*, the goal is to learn an action policy within an environment so as to maximize some notion of reward [77]. Other types of learning techniques exist, but in the scope of this paper—and the projects presented below—, we focus on these three.

2.1.2 Data. ML relies on *data* to perform the tasks defined above. The data-driven nature of ML is what makes it different from programming, which requires explicitly specifying all rules of a program to make it work. It also makes ML attractive in a design context, as data can be used to tacitly convey desired forms or behaviours to a computer. In modern ML, *big data* may be harvested or synthetically created to support learning [82]. In some cases, *small data* may suffice, which may involve designing appropriated data features—*e.g.*, gestural features, such as shape or frequency, to use ML for gesture recognition.

2.1.3 Models. *Models* designate mathematical functions that embody a set of statistical assumptions on data. Taken alone, models cannot afford learning: they rather act as probabilistic representations of data, and formal structures for the learning process. Some models can be trained from small amounts of data—*e.g.*, *shallow models*—, but in turn are highly sensitive to the choice of data features. Others require large amount of data to be trained—*e.g.*, *deep learning models* [44]—but may automatically compute data features in return.

2.1.4 Algorithms. *Algorithms* define a set of rules used for optimizing a model's parameters in relation to data. Importantly, algorithms may be *interpretable* in the sense that their internal functioning can be intuitively understood by humans [71]—*e.g.*, decision trees. Others (the large majority of them) remain *black-boxes*, and are subject to research to make their decisions explainable from both human and algorithmic perspectives [70, 87].

In most modern applications of ML, the incremental and progressive learning process is performed offline and separated from usage. Typically, a deep learning model is trained by ML engineers over large data sets; once trained, the model can be employed to perform its task without further training. In the emerging field of interactive machine learning, adaptation is a key property of ML, as learning is performed online: typically, people can add, delete, or modify small amounts of data to tacitly train an ML model [28], or experiment with various parameters of an algorithm to fine-tune a model's learning [51].

2.2 ML as Design Material

Modern engineering of ML techniques has arguably provoked and prevented many conversations on how ML could be designed, and on what makes an ML system optimal, harmful, or ethical. However, the discourse has often focused on technical considerations related to ML engineering. This notably involves the big amount of data required to optimise the millions of model parameters, and the resources needed to collect and annotate them. This also involves the computational resources needed to fulfill these storage and learning activities, which can become quite high compared to standard computers' performance [46]. As such, many ML-related research in HCI has so far been driven by advances in ML engineering rather than design innovation [90]. This resulted in the creation of ML systems whose design was partly specified by computational ML models. For instance, endeavours have focused on developing improved user models [56, 63], affective-enabled systems [19], or intelligent agents [58, 62].

As a consequence, we can observe that most of ML prototyping activities is often hidden in ML research publications, as highlighted by research in the emerging field of human-centred machine learning [43]. ML prototyping is here understood as the steps and decisions taken before the communication and dissemination of the final results. This leaves all the human work at stake in ML prototyping activities, such as carefully labelling training data sets, or choosing the best hyperparameters for a given algorithm, opaque to the public or fellows in the field. Model comparison remains the main aspect of ML prototyping that is published. Yet, it is often led based on quantitative concepts, such as algorithmic performance on standardised data sets, and as a result, only led within engineering labs [86].

In recent years, design research has been concerned by developing divergent ways of analysing and practicing with ML. Researchers sought to develop alternative conceptual frameworks for both ML and artificial intelligence. Such approaches are essential to challenge the networks of socio-material practices in which ML is embedded. Socio-material practices of ML are here understood as human ways of engineering ML models—as described above—, as well as the technical elements constituting the materiality of

ML. However, design research still lacks specific tools to go beyond Wizard of Oz ML prototypes [17]—where expected functionalities are simulated and therefore biased by the researcher's inquiry—, and materialise its alternative ML concepts into fully-functioning design artifacts [27, 89]. As a result, most research on ML as design material remain restrained in conceptual approaches, leaving hands-on exploration of ML materiality to software engineers who have expertise in ML.

In this context, we propose to look at the emergent field of human-centred ML, which applies user-centred and participatory design methods to design interactive machine learning tools [31, 76, 81], highlighting the qualitative concepts used by non-ML experts users to evaluate ML compared to quantitative concepts used by ML engineers [2]. Interestingly, Yang et al. observed how small data approaches of interactive machine learning helped designers who are not ML experts to craft and experiment with ML to build functional systems [92]. We are interested in deepening understanding of such hands-on, material-centered approaches to ML design.

Material-centered interaction design sought to bring material perspectives into HCI inquiry [24, 88]. Within this practice-oriented paradigm, physical and computational material are often placed on an equal footing for design [85]. Scholars have sought to produce conceptual frameworks for materials to be integrated earlier in the design process [41]. Interestingly, researchers explored ways to provide materials with more agency along research through design processes, approaching them as co-ethnographers [42]. We argue that this approach is the one implicitly brought by computational artists and NIME practitioners, as presented in the following section.

2.3 ML as Creative Material

Practices in NIME and computational art have explored ML materiality beyond conceptual or technical lenses, by adopting a crafting approach to code and ML as creative material [55]. In these practices, roles of artists, designers, or engineers may often be fluid along creative processes, with artists practicing with code, engineers exploring unconventional implementations, or even sometimes, artists, designers, and engineers being one same person. Such hands-on approaches to ML materiality have brought about novel configurations for ML that extend or even challenge design or engineering approaches to ML—*e.g.*, human-machine co-creativity [22, 23].

Computational artists have crafted ML materials within hybrid physical artworks and mediums to reveal and inquire our intimate entanglement with ML technology [4]. Poems by Allison Parrish reveal the intuitive coherence of deep learning models trained on large text data sets [64]. Video work by Memo Akten explores how images generated by deep learning models trained on large data sets push us to reflect on our collective representations of the world [1]. Sound sculptures by Stephanie Dinkins crafted through small data harvested by diverse communities call for better consideration of race, gender, and aging in the building of ML technology [25]. Robotic artworks by Sofian Audry explore the affective potentials of adaptive behaviors produced by shallow models of reinforcement learning when embodied by objects [3, 5]. Installations by Petra Gemeinboeck and Rob Saunders explore how ML-driven algorithmic exploration may provide robots with a sense of agency

and aliveness [39, 40]. Collective practice of *CoMo*, a smartphone-based web application based on interactive supervised learning, enabled a wide range of movement practitioners to experience ML computational properties through performance [59].

In addition, the NIME community has explored embodied ways of approaching ML, focusing on music performance as one specific case of interaction with sound and technology [11, 30, 53]. Earlier work from the MIT Media Lab explored ML to craft hyper-instruments [57], computer theaters [66], as well as novel gestural interaction within virtual musical environments [38]. More recently, deformable objects were crafted to explore pre-trained ML models of sound by mapping performers' haptic data to ML [79, 80]. Live coding environments were developed to explore data and algorithms as raw material for audiovisual improvisation [52, 94]. Interestingly, interactive machine learning was also explored to craft NIMEs, which essentially concerns the creation of a mapping between input movement data and sound output. Bevilacqua et al.'s *Gesture Follower* pioneered the use of interactive supervised learning to craft custom gestural controllers by iteratively training a real-time movement recogniser model using movement demonstrations [10]—an approach followed by Françoise et al.'s *XMM* [34], and many other tools [12, 21, 75]. Fiebrink et al.'s *Wekinator* [31, 33] was designed to allow performers and musicians to craft sound interactions through demonstrations and allowed for incrementally act upon ML components as materials constituting the interaction design. Recently, Scurto et al.'s *Co-Explorer* used interactive reinforcement learning to let people perform sound space exploration by communicating positive or negative reward data to an algorithmic agent [76]. Reinforcement learning was here used to craft sounds as well as exploration behaviours produced by the algorithmic agent.

In this paper, we propose to make use of interactive machine learning tools developed within the NIME community—specifically, *XMM* for interactive supervised learning and *Co-Explorer* for interactive reinforcement learning—in order to facilitate hands-on exploration and craft of ML materiality. Our aim now is to identify a methodological framework in which to situate such code-based art practices as potential approach for design research.

2.4 On Design Research Methods

Artistic methods and processes were recently introduced into HCI research to enrich both conceptual and technical advances on computation [8, 50]. Specifically, *reflexive methods* were employed to study participatory design of interactive machine learning with artists and HCI researchers in contexts of art practice. Caramiaux and Donnarumma used subjective inquiry to reflect on their five-year collaboration on ML for music performance [20]. Fiebrink and Sonami used interviews to relate their eight-year collaboration on ML for instrument design [32]. While highlighting the shifting roles and perspectives of both artists and researchers along the ML design process, reflexive methods put the emphasis on humans as cultural entities able to produce reflection, leaving little room for agency of ML as a material entity.

This paper proposes to explore *diffractive practices* as method to attend and respond to such specific entanglements. Diffractive practices combine two practical advantages of participatory design

and material-centered interaction design: namely, recognition of differences in human experiences as generative phenomena, and recognition of non-humans as producers of material agency [6]. It does so by relying on the practitioner's *embodied engagement* with diverse human and non-human entities, including data [78] and other practitioners [47]. Crucially, through her material feminist theory, Barad supports the idea that diffraction creates something ontologically new, breaking out of the cyclical, inductive realm of reflection [6].

The following sections introduce two interrelated studies exploring the potential of diffractive methods for prototyping ML through art practice. In the first study, diffraction is used as a method to analyse ML materiality in a NIME instrument, created following a linear process of conceptualisation, ML prototyping, and experimentation with performers. In the second study, diffraction is used from the starting of the project, as a method to iteratively and recursively practice with ML with other human and non-human actors, in the context of a robotic art project. As mentioned above, we propose to use the term *diffractive art practice*, introduced by artists-researchers Jane Prophet and Helen Pritchard [68, 69], to refer to our such art practice with ML following a diffractive method. Rather than deepening theoretical understandings of diffractive methods, our wish is to explore how they could reconfigure practices of design and engineering usually employed with ML.

As will be detailed in each section's introductions, each study was conducted by the first author in a specific academic institution—both located in France—, the first dedicated to music research, and the second to art and design research. Their common aim was to design novel structures and functionalities for ML by approaching it as a creative material following a diffractive method. In Sections 3 and 4, we will see that these novel designs for ML were tightly entangled with the design of novel human-machine interactions, and that both fluidly emerged through interference of human and non-human agencies along design processes. Thus, as will be discussed in Section 5, we believe that these two studies open up new spaces for material-centered interaction design, as well as for design research on, with, and through ML.

3 STUDY 1: SOMASTICKS

This section relates the crafting of an ML-driven NIME, *somasticks*, conducted between 2017 and 2019 by the first and third author at IRCAM, Paris, France, and its experimentation by performers during the “movA lab days”, an international workshop held in Nantes, France. In this study, the first author acted as a NIME performer and designer working with ML as creative material. The third author participated mainly in the form of conceptual and technical advice on ML implementation. The study adopted a linear approach to ML prototyping: first, (1) crafting an ML prototype for gestural interaction with sound based on the first author's embodied perspectives on ML materiality; then, (2) using a diffractive method to analyse differences in performers' embodied perspectives on ML as valid and generative results for prototyping.

3.1 Description

somasticks are augmented drumsticks that seek to emphasize the somatic side of drumming practice (see Figure 2). Contrary to standard drumsticks, *somasticks* do not need to hit any physical object to produce sound, but rather may be continuously waved in the air to trigger recorded drum sounds. Such a workflow may encourage performers to explore various movements in reaction to the internal sensations that sounds may produce within their body, thus helping cultivate their somatic knowledge over time.

Real drumsticks were used as tangible objects creating gestural affordances and suggesting movements related to drumming. Wireless sensors¹ were embedded to the sticks using a 3D-printed support specifically designed for drumsticks to harvest movement acceleration data from a performer. As will be detailed below, recorded drum samples were selected and curated to make the NIME sound like an actual drum, and ML materials were crafted to adapt sound generation to the performer's movement.




Figure 2: Rendering for *somasticks*.

3.2 ML Prototyping

Our wish was to adopt an embodied approach to ML to craft gestural interaction with sound within *somasticks*. We thus built on the first author's practice as NIME performer to elicit embodied concepts related to drumming, and used an interactive machine learning tool to craft these concepts within a concrete prototype.

3.2.1 Eliciting Embodied Concepts. The first embodied concept that we identified consisted in designing a form of "kinaesthetic zoom" [35], where the sticks increase in sensibility when movement activity decreases. This embodied concept can be found in drumming, where smaller gestures often require the performer to focus deeper on the perception of their body movement in space, in relation to sound.

A second embodied concept consisted in designing interactions with sound that account for the periodicity typical of drumming movements. Instead of raw accelerometer data, we decided to design *data features* related to movement frequency. We identified online wavelet analysis² as one method to compute a spectrogram of performers' movement in real-time. The resulting workflow for performers thus consists in producing stable movement qualities

to attempt to control the *somasticks*' sound, and slightly varying movement frequency to explore ML materiality in relation to sound.

3.2.2 Crafting ML to Materialise Embodied Concepts. We identified *unsupervised learning* as one candidate technique to craft such a kinaesthetic zoom. Specifically, we thought about using unsupervised learning in an online setting; rather than separating training and performance steps—as in standard interactive supervised learning workflows [31, 36]—, we decided to merge both to have ML continuously learn and dynamically adapt to input movement data. Our intention was to let performers physically interact with a fluid and adaptive model that constantly generates sound, depending on both previous and current movement. As such, performers could experiment with *learning adaptation* through the metaphor of movement consistency: the more their movement would be consistent, the more the sticks would focus in small variations in these movements, helping performers zoom in specific qualities of their movement. Thus, such a crafting differs from modern engineering of unsupervised learning, which leaves users passive during learning adaptation: here, performers actively produce new unlabeled data to shape learning adaptation, in what could be called an interactive unsupervised learning workflow—by analogy with interactive supervised learning workflows, in which users produce new labeled data to interact with ML.

We used a *shallow model*, called *Gaussian Mixture Model*, to perform the unsupervised learning task. Its simple structure, based on Gaussian components, let us craft the kinaesthetic zoom in easier ways than complex structures of deep neural networks. After some initial tests, we opted for a *small data* approach, using wavelet features extracted over a 10-second sliding window as training data accounting for movement periodicity. Lastly, we crafted an *online learning algorithm* by retraining the ML model every 100 milliseconds over the full data set. While non-optimal from an engineering perspective, this implementation enabled to rapidly experiment with classification enabled by unsupervised learning.

We used the *XMM* library³ for interactive machine learning [36] and the Max/MSP visual programming environment—emerging from NIME practices—to craft ML for the *somasticks*. Crucially, using this interactive machine learning tool enabled us to fully focus on experimenting and performing with ML, by rapidly implementing the envisioned ML technique, model, and data features, while skipping the engineering of online learning algorithm—deemed to be a joint mathematical and technical challenge for ML [61]. Full details of the implementation can be found in [74].

We used concatenative synthesis to generate new sound patterns from recorded samples [73]. We created six sound corpuses related to six elements commonly found in standard drum kits: bass drum, snare drum, rack toms, hi-hat, crash, and ride. Each corpus was carefully designed to contain a wide variety of performing modes related to drumming—*e.g.*, from soft to hard hitting. These modes are easily captured by descriptors of concatenative synthesis [73].

The first author then experimented and performed with the *somasticks* to iterate and converge on a fixed mapping between sound and ML. This mapping makes an extensive use of the uncertainties produced by the *Gaussian Mixture Model* as creative material: probabilities will dynamically adapt as performs continuously interact

¹<http://ismm.ircam.fr/riot/>

²<http://ismm.ircam.fr/mubu/>

³<http://ircam-md.github.io/xmm/>

with the online learning algorithm, producing variations in sound synthesis. The instantaneous class defined the sound corpus from which samples are played. Gaussian probabilities set the respective gains at which samples are played. The model probability set the rhythmic period at which new samples are played: the higher the probability is, the faster samples are played. The probability weights set a random temporal variation for playing a sample: the more weighted a cluster is, the more rhythmic the playing of its samples is. Statistical means and covariances set the choice of sound samples: we scaled means over spectral and loudness sound descriptors, and used covariances as radius to search sound space.

3.3 Diffractive Analysis of Performers' Data

We experimented *somasticks* in a workshop involving six expert performers—two dancers, two somaticists, two musicians (see Figure 3). All performers agreed to participate without compensation to the workshop. Performers were asked to spend between 5 and 20 minutes experimenting *somasticks* relying on the following advice: listen to the produced drum sounds, focus on their bodily sensations, and move freely with the sticks. Importantly, performers were not explained technical details on ML to fully focus on embodied interaction and performance with ML, following a NIME approach (similar to Section 2.3). Informal discussion was then introduced by the first author to disrupt hierarchies between them and performers and open up analysis from a variety of perspectives. The first author used diffractive analysis [6] to read through transcripts of performers' audio-visual data. Rather than seeking to normative framings, diffractive analysis sought to illuminate differences in performers' embodied perspectives as valid and generative results for prototyping.



Figure 3: Pictures of performers experimenting with *somasticks* in the workshop.

A first observation lies in the use of bodies and perception by both the first author and performers to experiment with the sticks. For example, one dancer tried not to move as she was first “*in quest of silence*”, then seemed to understand that movement periodicity was crucial to the model’s functioning: “*There is a kind of obstacle*

course that appeared. [...] It was when I made small movements that I managed to refocus on what I was doing, and to take back control over the system”, she said. The other dancer seemed to grasp kinaesthetic zooms enabled by ML: “*I understood that there was a delay. This meant that if I am doing the same thing during three seconds, the totality [sic] will come after*”, she said. The metaphors they used suggest that the shallow model possesses rich material properties: despite having elementary structure from a computational perspective, it managed to produce a diversity of experiences in performers, who engaged their imaginary and bodily sensibilities to explore ML.

The drum sticks and sounds had all performers adapt their movement in different ways, depending on their specific perception. For example, two participants reported an asymmetry in their motion: “*The sticks, really, define their very own geometry, so I focused a lot on the trajectories defined by the lines [of the sticks]. [...] I rapidly realized that I was very lateralized, in the sense that my right hand dominates, and I am not trained in drumming*”, one somatic practitioner analysed. Another dancer modified its movement exploration, this time in relation to the act of drumming: “*As soon as I saw it, I thought that I can play [the drums]. Thus, I begin to sit down*”, she commented. Such specific patterns of movement data could not have been anticipated if we opted for a traditional ML engineering approach.

Eventually, the uncertainty of unsupervised learning classifications generated strong differences in experiences between performers, with some of them praising it—“*For me, there is something very attractive in the fact that there are moments of synchrony and moments of autonomy. And in its autonomy, there were lots of variation still, so it was always interesting to listen to, and move along with*”, one somaticist reported—, and others not hesitating to complain about it—“*It’s kind of weird. Yeah, it’s really high level, and then it’s like, you know... You feel like you’re kind of blind, and then I have no idea what’s going on and why*”, one musician confessed. These observations may push one to improve engineering of the learning algorithm—which we intentionally implemented non-optimally—as next design step for *somasticks*, with the goal of finding some normative formulation to improve its interpretability. Or, one may take this materiality of ML as it is, and embrace differences in experiences that it produces in performers as characteristic of *somasticks*. Going in this sense, one musician suggested to test *somasticks* on an actual drumset: “*I think it could be great, even to refine the system in a kind of design process, to play it with something, you know, just, like an augmented instrument*”, he said.

Through this study, we demonstrated how an embodied approach led the first author to prototype ML in ways that valued computational properties—real-time adaptation of online learning, classifications of unsupervised learning, uncertainty of a *Gaussian Mixture Model*, and qualities of movement data—over modern engineering. We also showed how diffractive analysis enabled to value different embodiments produced by different performers as situated ways of evaluating an early ML prototype, beyond normative framings.

4 STUDY 2: THE APPRENTICES

Witnessing their relevance for attending to differences in human experiences of ML, the first author decided to leverage diffractive methods throughout the prototyping of a collaborative robotic art project (*The Apprentices*), started in 2020 at EnsadLab, Paris, France, in collaboration with the Art Direction Nods Team (Xdlab) of Orange, Châtillon, France. In this project, the first author acted as computational and sound artist working with ML as creative material. The interdisciplinary team, composed of members of the Reflective Interaction research group of EnsadLab, and of the Art Direction Nods Team (Xdlab) of Orange, included one media artist, one UX designer, one product designer, one anthropologist, one cognitive psychologist, and two engineers. The creative process that drove the following ML prototyping relied on the first author's diffractive art practice, who moved within and beyond team members' and other material agencies' perspectives to sustain conversations on ML and produce novel material configurations of ML. The next sections describe the iterative and recursive journey of diffractive ML reading, material-centered design space stabilisation, and diffractive ML prototyping in which the interdisciplinary team engaged⁴.

The project started from a common wish to engage in the crafting of a collective of *behavioral objects*—*i.e.*, robotic objects that avoid symbolic or anthropomorphic representations to diverge from engineering-driven robotics [7, 13–15, 54]. Such an artistic approach to robotics would enable all team members to explore novel ways of designing objects, going beyond their physical and interactive aspects to embrace temporal aspects of movements, sounds, and collective behaviours of objects. A long-term goal of the project lies in evaluating audience experience of social behaviours emerging from spatio-temporal configurations of objects. Specifically, a shared interest resided in the exploration of ML materiality to generate movement-sound behaviors within our object collective, as will be detailed in the following sections.

4.1 Diffractive ML Reviewing

To start discussion and collaboration on exploring ML for behavioural objects, an interrelated understanding of ML was required. Initial meetings thus aimed at producing a *diffractive review* of ML. Six group meetings—involving individual presentations followed by group discussions—as well as six informal brainstorming sessions let the first author introduce data-driven techniques inherent to ML (similar to Section 2.1), along with design concepts counterbalancing modern ML engineering (similar to Section 2.2), and artistic works exploring material configurations of ML (similar to Section 2.3). This interdisciplinary space, designed to match other team members' interests, was essential to avoid technical barriers and provoke conversations on how ML could be designed.

The diverse expertise inherent to our team in turn interfered with these reviews. One UX designer situated ML within the lens of social and cognitive theories of learning, spanning pedagogical methods, affective theories, and animal behaviours, to imagine new opportunities for ML. One media artist presented a range of artistic works in the areas of cybernetics and behavioural aesthetics [65]

⁴Due to the global pandemic situation, collaboration progressively switched from in-person meetings and workshops to virtual discussions and presentations.

to situate ML in a historical context, and call for material-centred approaches to ML to explore novel configurations between humans and machines. One anthropologist introduced practices of animism led by diverse human communities to have team members reflect on their perception of ML and robotic objects as socio-material artifacts. All in all, these diffractive readings and conversations helped us embrace a multiplicity of different perspectives on ML, which remained entangled throughout practice.

4.2 Stabilising a Material-Centered Design Space

Progressively, conversations alone seem to prevent team members from going further in ML design in relation to our object collective. We thus decided to define a material-centered design space for *The Apprentices*. While entangled in ML prototyping, object design remains out of the scope of this paper. We will only sum up the embodied concepts for ML that emerged from material practice within the team, leaving details of object, sound, and movement design as well as fabrication for a next paper.

We opted for a collective of three objects (see Figure 4). This number was chosen to reduce complexity in the design of collective object behaviours, while also remaining sufficiently high to study audience perception of collective object behaviours in future work. We decided that all three objects would share same shape (with a diameter of 15 centimeters), so that perception of specific object behaviours would only emerge through temporal aspects of movements, sounds, and learning. Crucially, we opted for a hybrid *installation-performance* format, presenting our object collective on a dark and circular stage with a diameter of 2 meters. We opted for a structured performance format based on three short parts—approximately five minutes each, resulting in a total duration of 15 minutes—, as an audience would spend relatively short amounts of time facing the objects, reducing the probability for complex human-objects relationships to emerge. This design constraint raised specific design challenge for ML, as will be detailed below.



Figure 4: Rendering for *The Apprentices*.

We opted for *soft materials* to realise the objects' organic, uncluttered, and curved shape (see Figure 4), providing them with a sense of aliveness throughout their materiality [41]. We chose to use the *sonic artifacts* produced by objects' motorised movements

as material for sound design, conveying a sense of physical, internal effort for the robotic objects [93]. Lastly, we thought about using *rhythm* as temporal form of expression for the objects, as well as of interaction modality for the audience [60], who will be able to softly hit the edge of the stage to communicate with objects using rhythm during the third part of the performance, as we will see. Details on object fabrication and sound and movement design will be provided in a next paper.

To guide the design of rhythmic structures for each of the three parts, we elicited three concepts freely inspired by social perception of animacy and agency [83, 84]: *mutual aid*, *competition*, and *attachment*. These social concepts would also help to formulate hypotheses on how collective object behaviours would be experienced by an audience in future experimentations. For *mutual aid*, objects would simultaneously produce metronome-like oscillatory movements (with a one-second period): we hypothesized that the audience would perceive a growing solidarity between objects as their oscillatory movements would phase. For *competition*, objects would successively move to produce rapid rhythmic patterns (with a tempo of 6 beats per second): we hypothesized that the audience would ascribe a growing strength to a given object as its rhythmic pattern would densify. For *attachment*, objects would first wait that the audience softly hit the edge of the stage to perform a rhythmic pattern, then successively move to reproduce it: we hypothesized that the audience would be affected by how objects would succeed or struggle to accurately imitate their rhythm.

4.3 Diffractive ML Prototyping

Motivated by our first study, the first author proposed to use prototypes of ML to have all team members experience material configurations of ML and share their different experiences of it, without necessarily delving into technique. Yet, crafting ML for each of the three social concepts would have required the first author to spend an important amount of time on implementation (as in *so-masticks* linear prototyping process, described in Section 3.2), and consequently, could have disentangled team members from the ML prototyping activity. We thus opted for Wizard of Oz ML prototypes: rather than individually implementing ML, we sought to collectively embody ML computational properties by rapidly creating mock-up collective behaviours based on abstract graphical and sonic representations (see Figure 5). Crucially, rather than simulating functionalities of modern ML engineering [17], we specifically designed these Wizard of Oz ML prototypes so that they reproduce structural properties of ML material elements, as will be described below.

4.3.1 Wizard of Oz ML Prototyping. The first author started by identifying three ML techniques to materialise the three social concepts. For *mutual aid*, objects' task would be to learn to synchronize their oscillatory movements by sharing their knowledge of others—a problem that can be addressed by *multi-agent reinforcement learning* [18]. For *competition*, objects' task would be to learn to attain maximum rhythmic density before others—a process that can be realised by *reinforcement learning* [77]. For *attachment*, objects' task would be to imitate rhythms produced by the audience by learning a model of vibration data harvested through the stage—a workflow enabled by *supervised learning* [16].

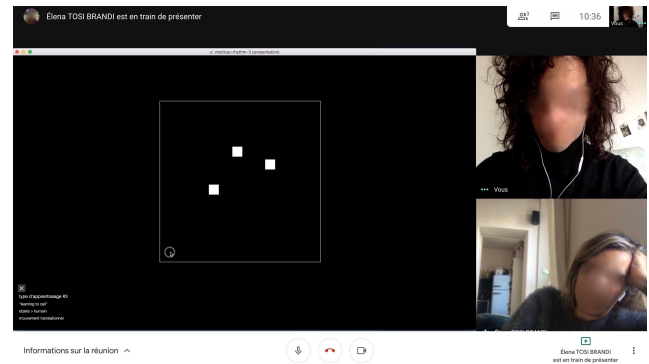


Figure 5: Wizard of Oz ML Prototyping for *The Apprentices*, led through virtual conferencing.

The first author then used Max/MSP to implement these Wizard of Oz ML prototypes. Data workflows specific of each ML technique were reproduced for each part to recreate ML structural properties; yet, learning algorithms were replaced by random generators, thus acting as Wizards in the prototypes. Abstract white squares and synthesised sounds were employed to create graphic and sonic representations of robotic objects. To start a dialogue, the first author crafted initial object configurations for each of the three parts. Members of the team reacted positively to these prototypes, iteratively generating ideas for objects' movements and spatial arrangement, while recursively adjusting the three concepts of mutual aid, competition, and attachment in relation to structural constraints raised by ML techniques and data. For example, one cognitive psychologist suggested an ML implementation for “competition”, proposing that different exploration-exploitation trade-offs of reinforcement learning could embody different adaptation abilities for robots, thus echoing human and animal learning behaviours [77].

4.3.2 Material ML Prototyping. Along Wizard of Oz prototyping, the need to experience ML materiality emerged within the team. We decided to directly craft ML within robotic objects in order to experience ML materiality through embodiments that are specific to our project. We used the *MisB Kit*⁵, an open-source, modular robotics kit enabling rapid prototyping of robotic objects using servo-motors and velcro bars [14]. Before crafting ML, we introduced the *MisB Kit* to members of the team in an experimental workshop (see Figure 6). Material prototyping was useful to deconstruct symbolic or anthropomorphic representations of robotics among the team. One product designer and one engineer, who were accustomed with familiar designs of robots and ML [91], generated many ideas for life-like behaviours witnessing objects moving and making sound in chaotic ways.

The first author eventually involved in material ML prototyping by exploring and experimenting with the *MisB Kit* to converge to a final specification for ML. We used the *Co-Explorer* library⁶ for interactive reinforcement learning [76], the *XMM* library⁷ for interactive supervised learning [36], and the Max/MSP visual programming

⁵<https://misbkit.ensadlab.fr/>

⁶<https://github.com/Ircam-RnD/coexplorer>

⁷<http://ircam-md.github.io/xmm/>




Figure 6: Material prototyping for *The Apprentices*, led through an in-person workshop.

environment to craft ML for *The Apprentices*. Using these interactive machine learning tools enabled to rapidly experiment with ML techniques, despite the conceptual complexity described above.

We crafted a *shallow model*, called *Sarsa* [77], to perform multi-agent reinforcement learning and embody the concept of mutual aid. In *Sarsa*, an agent acts within an environment and updates its learning based on positive or negative reinforcement received from the environment. Here, the task for object i is to learn to attain the phase of object $i + 1$. For object i , reinforcement is obtained by taking opposite euclidean distance to object $i + 1$'s phase. While non-optimal from an engineering perspective, this circular definition entangles the objects' cooperative learning. As we wanted our objects to learn in less than five minutes—due to temporal constraints raised by the hybrid installation-performance format—, we opted for a *small data* approach, taking as discrete state the object's phase of actuation, and the corresponding actions as increasing or decreasing the phase by one discrete step. We used different discount factors for different objects as *learning algorithms* for multi-agent reinforcement learning: a factor approaching 0 will make the agent short-sighted by only considering current rewards, while a factor approaching 1 will make it strive for a long-term high reward. As such, objects with best memorising abilities will cooperate with more forgetful objects by guiding their learning through positive or negative reinforcement.

We also crafted a *shallow Sarsa* model to perform reinforcement learning and embody the second concept of competition. The objects' task is to individually learn to attain maximum rhythmic density as quickly as possible. While usually hidden in engineering approaches, such real-time learning would materialise the antagonisms of these objects. Due to the five-minute duration constraint, we opted for a *small data* approach, taking as discrete state a four-beat rhythmic pattern for object actuation, and the corresponding actions as activating or deactivating one beat in the pattern. Reinforcement is positive when an action densifies the rhythmic pattern,

and negative otherwise. We used an ϵ -greedy policy as *interpretable algorithm* for real-time learning: an ϵ value approaching 1 will make the agent only explore its environment (equivalent to a random behaviour), while a value approaching 0 will make it only take actions that yield the most reward (but will risk leading it to remain stuck in local rhythmic optima). As such, objects that balance between exploration and exploitation will be firsts to find maximum rhythmic density, while other will struggle or fail to learn to attain it.

We crafted a *shallow model*, called *Hierarchical Hidden Markov Model* [36], to perform supervised learning and embody the concept of attachment. Our idea here is to use supervised learning to have objects imitate rhythms produced by an audience, thus engaging the audience in a communicative relationship with objects, iteratively exploring the objects' different learning abilities, as well as experiencing the clumsy moves and sounds produced by the objects' materialities. The objects' task is to learn to reproduce the audience's vibration pattern through some sequence complexity. While usually employed to fine-tune a model's learning in engineering, different sequence complexities provide different expressive abilities to objects linked with timing precision. We opted for a *small data* approach, taking time as input data, and audience vibration energy over time as output data. We used an interactive supervised learning workflow as *interpretable algorithm*, combining a demonstration step—the audience performing a rhythmic pattern—and a performance step—objects imitating the rhythmic pattern.

Members of the team engaged their bodily sensibilities to interpret learning behaviours embodied by objects' rhythmic movements and sounds. One engineer reported that the progressive alignment of objects' sound and movement in the first part somehow surprised him, as it embodied some kind of team effort across objects. The first author noticed the empathetic relations emerging between objects in the second part, observing how the objects succeeded, struggled, or failed to produce dense rhythms. One media artist, one UX designer, and one anthropologist were surprised to witness some sweetness emerging from objects' tentative imitations in the third part, comparing them to child objects that strive to communicate with their vibrating bodies, with different expressivities depending on their sequency complexity. Current work includes preparation of an experimental setup to investigate how members of an audience experience these ML-driven rhythmic behaviours and compare it to our design hypotheses and own experiences with the robotic object collective.

The study reviewed above demonstrated how displacing and diffracting the practitioners helped produce novel configurations for ML, situating their prototyping within socio-material practices by opening space for experimentation and communication on computation. It also showed how embodied perspectives led members of the team to prototype ML in ways that valued computational properties—probabilistic decisions of reinforcement learning, exploration and learning abilities of algorithms, imitation complexities of *Hidden Markov Models*, and performativity of vibration data—over modern engineering.

5 DISCUSSION

In this section, we discuss our diffractive art practice with ML to (1) identify socio-technical conditions for art-based ML prototypes to produce constructive interferences between conceptual and technical approaches—thus becoming valid research sites on their own for material-centered interaction design—, and (2) sketch contours of an *intra-active* understanding of ML to start a dialogue on using diffractive practice in wider design research. We finish by (3) stating the limitations of the present inquiry, calling for contributions from other designers and researchers to deepen theoretical understandings of how diffractive practices may extend or restrain other design methods.

5.1 Interference Conditions for Art-Based ML Prototypes

The interactive machine learning tools used in our two studies enabled to craft ML in ways that favoured embodiment of ML materiality through mediums such as sound, visuals, or movement. As such, the produced ML prototypes slightly differ from modern ML systems found in engineering and design: while the latter often privilege the realisation of logical concepts or functionalities through technical engineering, our prototypes emphasised the experience of raw computational properties through incremental crafting. The next sections derive five socio-technical conditions for such art-based ML prototypes to interfere with conceptual and technical approaches, based on qualitative analysis of our two practice-based studies.

5.1.1 *Situational whole.* Through our studies, we were able to observe that art-based ML prototypes take their whole dimension as interdisciplinary research sites when they are deployed “in the wild”, that is, when they are practised out of engineering labs where ML materials are usually developed—as already demonstrated by Wagstaff [86]. Such a *situational whole*, defined by Giaccardi et al. as the network of social and cultural values and individual experiences in which materials are embedded [41], should be understood as the entry point for art-based ML prototypes to be appropriated by researchers and practitioners from diverse fields, in tight relation with other materials—*e.g.*, objects, sounds, configurations—, as our studies have shown. Situational wholes thus open pathways to the design and development of ML materials that are entangled with social and cultural factors [6].

5.1.2 *Small data.* In addition, we also witnessed that art-based ML prototypes should go widespread big data approaches and build on *small data* approaches, which allow for intuitive apprehension of ML pipelines by researchers and practitioners who are not ML experts [31]. The performers in our first study reported that they took control over their movement data as our ML prototype accounted its periodicity over a short temporal scale. The artists and designers in our second study were able to experience the data-driven structures and functioning of both Wizard of Oz and material ML prototypes of our robotic object collective. Small data thus opens possibilities for designing ML systems that are scaled to their human users, and that can be collectively crafted by communities of people under-represented within ML research [25]. Additionally,

small data may act as intermediary step for non-ML experts to engage in modern big data approaches, providing them with further computational properties for material design.

5.1.3 *Shallow model.* The findings of our studies suggest that art-based ML prototypes’ ability to adapt to humans and non-humans was realised by *shallow models* able to rapidly learn online [36]. Contrary to deep learning models, which afford rich computational properties at the cost of large amounts of data and ML expertise [1, 64], shallow models’ elementary computational tasks, such as classification or decision-making, provide probabilistic properties that remain to be engaged as material by designers [9], building on advantages of small data. At the same time, shallow learning models require standard computational power to be trained, which may reduce joint economical and environmental cost of building ML systems [46]. In a nutshell, shallow models offer researchers and designers possibilities to prototype ML systems without relying on large human and computer infrastructures. Additionally, shallow models may enable crafting of deep learning models based on small data, by being used in conjunction with pre-trained deep learning models—an approach known as transfer learning [72].

5.1.4 *Learnable algorithm.* Our two studies also shows the importance of art-based ML prototypes to rely on *learnable algorithms* to keep their learning abilities appreciable to non-ML experts as much as to ML experts [87]. Our observations suggest that non-ML experts may seize a learning algorithm when the actions that it produces are related toward a rational task. This task can either be goal-oriented, as in most engineering approaches—*e.g.*, classifying human data in our first study—or open-ended, as in some design approaches [76]—*e.g.*, exploring an environment in our second study. The notion of learnable algorithms thus open ways to move beyond notions of interpretability or black-box to approach ML as a material, whose properties should be explored by artists, designers, or other communities of people, as much as by engineers.

5.1.5 *Somaesthetic behaviour.* Beyond engineering aspects, our work illuminates the importance of ML prototypes’ *somaesthetic behaviour*—understood as the temporal aspects of computational properties of ML experienced through some medium—, which helped artists project themselves into ML data, models, algorithms, techniques, and computations using their bodies and perception in addition to their rational thinking [48]. The interactive learning of our ML prototypes, including the errors and strivings of non-optimal learning, created behavioural morphologies that were intuitively embodied by art practitioners [3]. As such, somaesthetic behaviour should be understood as one intrinsic feature of ML prototypes, enabling to deconstruct symbolic or anthropomorphic representations of intelligent behaviours in ML by including the weaknesses and vulnerabilities that make us humans learn. Somaesthetic appreciation of ML prototypes thus opens holistic approaches to the design of ML systems, suggesting that the behaviours produced by ML computational properties may be as valid as its engineering to be approached as design material within HCI inquiry [76].

5.2 Intra-Active Machine Learning

The understanding of diffractive art practice with ML discussed here assembles complex disciplinary concepts, techniques, and practices,

as well as diverse roles of artists, designers, and engineers, into one fluid and ever evolving socio-material phenomenon [6]. The next sections sketch contours of *intra-active machine learning*, a diffractive practice that may attend and respond to assemblages of design and engineering beyond the realm of art to produce novel ways of becoming ML practitioners.

5.2.1 Embracing Differences between Design and Engineering. In most ML design processes, including interactive machine learning, designers often work with ML engineers through predefined collaboration processes, where conceptual and practical roles are static and well delineated. Yet, our art practice with ML reveals how ML can also be designed iteratively and recursively, through hands-on experimentation with ML materiality. The shallow models resulting from such a practice may evolve into more developed structures through iterative design and engineering with communities of users and stakeholders over longer time periods [76]. Yet, the resulting crafts may be perceived as elementary, or even erroneous, by modern practices of engineering or design mainly oriented toward established approaches of computational modelling.

An intra-active understanding of ML calls for resisting such normative framings of ML by producing novel material configurations that embrace differences between design and engineering. Within such a process, art-based ML prototypes may thus be considered as valid research sites in their own, supporting the reconfiguring of core computing notions by living at the boundaries of more established practices. By making transparent their underlying assumptions using small data and learnable algorithms, practitioners may entangle qualitative and quantitative concepts of established design or engineering frameworks, and create space for emergence of novel concepts for ML, specifically produced by differences, interferences, and intra-actions between these disciplinary practices, as our studies have shown.

5.2.2 Becoming A Diffractive ML Practitioner. In the two interrelated accounts of art-based ML prototyping developed here, the first author sought to report the process of moving within and beyond reflexive practice to become a diffractive ML practitioner, who materially constitutes themselves through intra-action among practices, materials, techniques, as well as artists, performers, designers and engineers. It is not to say that this displacing and diffracting the self who practices ML is without costs. Fragile configurations of computational materials can push established ML researchers and engineers to question their relevance or accuracy. As such, tensions can occur between the diffractive practitioner and other practitioners not necessarily aware of the diffractive method being applied. Yet, such costs of being simultaneously nowhere and everywhere within ML must be considered with respect to the costs required to independently design or engineer ML: innovating ML as design material implies large-scale research and teaching effort [89], and engineering modern ML techniques requires large human and computational infrastructure or institutions [43].

By paying attention to small details of intra-actions and focusing on becoming-with others [47], diffractive methods can contribute to reconfigure complex socio-material practices of ML. As sketched above through the concept of situational whole, appropriation of ML computational properties by artists and designers emerged through tight engagement and collaboration with ML engineers,

where individualities fluidly adjust to each other, going up to displacing and diffracting creative roles along the process [20, 32]. Embracing somaesthetic behaviours can be one way to facilitate communication between individuals, including computational and robotic entities, relying on pre-reflexive and non-verbal cues as basis for embodiment of material properties [48], and performance for fluid alterations of judgement and configurations between people, materials, and practice [41]. By extending reflexive methods of interactive machine learning, diffractive methods for intra-active machine learning may have the potential of assembling socio-material practices of ML in a process of more-than-human design in HCI [37].

5.3 Limitations

In this paper, we took the decision to put processes and materials involved in ML-based diffractive art practice forward. This approach enabled us to identify features and notions that may be seen as intermediary knowledge, or strong concepts [49], capturing generative qualities of art-based ML prototyping and making them explicit for the field to build upon. Yet, we did not attempt to deepen our explorations of diffractive methods at a theoretical level. Rather, our wish was to start a dialogue with other designers and researchers from the field on how diffractive methods may expand or restrain other design methods. As a specific group of artists, designers, researchers, and engineers, our intuition is that diffractive methods better recognise the fluidity between human and non-human actors than participatory design—emphasising humans—or material-centered interaction design—emphasising non-humans. In this sense, it may share similarities with framings of research through design as post-phenomenological practice [45], while however putting more emphasis on performativity inherent to human-things relations, thus opening ways to better trace responsibility in ML design processes [37]. Future research may investigate such ethical and theoretical aspects of diffractive methods, especially on the topic of ML, whose embedding in social and political dimensions is critical.

6 CONCLUSION

This paper explores the potentials of diffractive art practice as method to prototype ML in design research. Through a review of technical, artistic, and conceptual approaches to ML, and two accounts of practice-based studies with ML of our own, we argue that novel ways to prototype ML can emerge through embodiment of ML materiality and craft of ML techniques, data, models, and algorithms. Thus, ML prototypes may be realised based on small data, shallow models, and learnable algorithms, to produce novel somaesthetic behaviours that can be experienced and discussed by a diversity of practitioners. The socio-technical conditions derived for producing art-based ML prototypes pave the way for an intra-active understanding of ML that embraces differences between practices of design and engineering through diffractive practice. While produced by a specific group of artists, designers, researchers, and engineers, we hope that the present work will encourage other individuals from a diversity of communities, both academic and non-academic, human and non-human, to produce novel understandings of our intimate entanglements with ML technology.

ACKNOWLEDGMENTS

The authors would like to thank Jules Françoise, Riccardo Borghesi, Djellal Chalabi and Emmanuel Flety for their support in prototyping *somasticks*, and all members of the movA network for their precious time and feedback experimenting with them. The authors would also like to thank Elena Tosi-Brandi, Samuel Bianchini, Joffrey Becker, Florent Levillain, Olivain Porry, Didier Bouchon, Victor Audouze, Corentin Loubet, Joséphine Mas, Franck Weens, Catherine Ramus, and Camille Dauhut for their fundamental contributions and fruitful engagement within *The Apprentices*. The latter project was supported by Orange's Art Direction Nods Team (Xdlab) and EnsadLab's Reflective Interaction research group, the research laboratory of the École nationale supérieure des Arts Décoratifs.

REFERENCES

- [1] Memo Akten, Rebecca Fiebrink, and Mick Grierson. 2019. Learning to see: you are what you see. In *ACM SIGGRAPH 2019 Art Gallery*. ACM, 1–6.
- [2] Saleema Amershi, Maya Cakmak, William Bradley Knox, and Todd Kulesza. 2014. Power to the people: The role of humans in interactive machine learning. *AI Magazine* 35, 4 (2014), 105–120.
- [3] Sofian Audry. 2019. Behavior Morphologies of Machine Learning Agents in Media Artworks. *Leonardo* (2019), 1–10.
- [4] Sofian Audry. 2020. La matérialité révélatrice de l'apprentissage automatique. *Espace: art actuel* "IA - Art sans artistes", 124 (2020), 44–49.
- [5] Sofian Audry, Rosalie Dumont-Gagné, and Hugo Scurto. 2020. Behaviour Aesthetics of Reinforcement Learning in a Robotic Art Installation. In *4th Workshop on Machine Learning for Creativity and Design at NeurIPS 2020*.
- [6] Karen Barad. 2007. *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning*. duke university Press.
- [7] Joffrey Becker, Virginie André, and Alain Dutech. 2019. Qualcomm: une expérience sur la qualification des comportements d'une lampe robotique. *Techniques & Culture. Revue semestrielle d'anthropologie des techniques* (2019).
- [8] Steve Benford, Chris Greenhalgh, Andy Crabtree, Martin Flintham, Brendan Walker, Joe Marshall, Boriana Koleva, Stefan Rennick Egglestone, Gabriella Giannachi, Matt Adams, et al. 2013. Performance-led research in the wild. *ACM Transactions on Computer-Human Interaction (TOCHI)* 20, 3 (2013), 1–22.
- [9] Jesse Josua Benjamin, Arne Berger, Nick Merrill, and James Pierce. 2021. Machine Learning Uncertainty as a Design Material: A Post-Phenomenological Inquiry. *arXiv preprint arXiv:2101.04035* (2021).
- [10] Frederic Bevilacqua, Fabrice Guédy, Norbert Schnell, Emmanuel Fléty, and Nicolas Leroy. 2007. Wireless sensor interface and gesture-follower for music pedagogy. In *Proceedings of the 7th International Conference on New Interfaces for Musical Expression (NIME)*. ACM, 124–129.
- [11] Frédéric Bevilacqua, Rémy Müller, and Norbert Schnell. 2005. MnM: a Max/MSP mapping toolbox. In *Proceedings of the 2005 conference on New interfaces for musical expression*. National University of Singapore, 85–88.
- [12] Frédéric Bevilacqua, Bruno Zamborlin, Anthony Sypniewski, Norbert Schnell, Fabrice Guédy, and Nicolas Rasamimanana. 2009. Continuous realtime gesture following and recognition. In *International gesture workshop*. Springer, 73–84.
- [13] Samuel Bianchini, Rémy Bourganel, Emanuele Quinz, Florent Levillain, and Elisabetta Zibetti. 2015. (Mis) behavioral Objects. In *Empowering Users through Design*. Springer, 129–152.
- [14] Samuel Bianchini, Hiroshi Ishii, Bourganel Rémy, Mahé Emmanuel, Labrune Jean-Baptiste, and Quinz Emanuele. 2014. The misbehavior of animated object. In *Proceedings of the 8th International Conference on Tangible, Embedded and Embodied Interaction*. 381–384.
- [15] Samuel Bianchini and Emanuele Quinz. 2016. *Behavioural Objects I: A Case Study*. Vol. 1. Sternberg Press.
- [16] Christopher M Bishop. 2006. *Pattern recognition and machine learning*. springer.
- [17] Jacob T Browne. 2019. Wizard of Oz Prototyping for Machine Learning Experiences. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–6.
- [18] Lucian Buşoni, Robert Babuška, and Bart De Schutter. 2010. Multi-agent reinforcement learning: An overview. *Innovations in multi-agent systems and applications-1* (2010), 183–221.
- [19] Erik Cambria. 2016. Affective computing and sentiment analysis. *IEEE intelligent systems* 31, 2 (2016), 102–107.
- [20] Baptiste Caramiaux and Marco Donnarumma. 2020. Artificial Intelligence in Music and Performance: A Subjective Art-Research Inquiry. *arXiv preprint arXiv:2007.15843* (2020).
- [21] Baptiste Caramiaux, Nicola Montecchio, Atau Tanaka, and Frédéric Bevilacqua. 2015. Adaptive gesture recognition with variation estimation for interactive systems. *ACM Transactions on Interactive Intelligent Systems (TiIS)* 4, 4 (2015), 18.
- [22] Nicholas Davis, Chih-Pin Hsiao, Kunwar Yashraj Singh, Lisa Li, and Brian Magerko. 2016. Empirically studying participatory sense-making in abstract drawing with a co-creative cognitive agent. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*. ACM, 196–207.
- [23] Nicholas M Davis, Yanna Popova, Ivan Sysoev, Chih-Pin Hsiao, Dingtian Zhang, and Brian Magerko. 2014. Building Artistic Computer Colleagues with an Enactive Model of Creativity. In *ICCC*. 38–45.
- [24] 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*. 1–12.
- [25] Stephanie Dinkins. 2020. Community, Art and the Vernacular in Technological Ecosystems. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*. 221–221.
- [26] Graham Dove and Anne-Laure Fayard. 2020. Monsters, metaphors, and machine learning. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [27] Graham Dove, Kim Halskov, Jodi Forlizzi, and John Zimmerman. 2017. UX Design Innovation: Challenges for Working with Machine Learning as a Design Material. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 278–288.
- [28] Jerry Alan Fails and Dan R Olsen Jr. 2003. Interactive machine learning. In *Proceedings of the 8th international conference on Intelligent user interfaces*. ACM, 39–45.
- [29] Sarah Fdili Alaoui. 2019. Making an interactive dance piece: Tensions in integrating technology in art. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. 1195–1208.
- [30] Sidney S Fels and Geoffrey E Hinton. 1993. Glove-talk: A neural network interface between a data-glove and a speech synthesizer. *IEEE transactions on Neural Networks* 4, 1 (1993), 2–8.
- [31] Rebecca Fiebrink, Perry R. Cook, and Dan Trueman. 2011. Human Model Evaluation in Interactive Supervised Learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Vancouver, BC, Canada) (CHI '11)*. ACM, New York, NY, USA, 147–156. <https://doi.org/10.1145/1978942.1978965>
- [32] Rebecca Fiebrink and Laetitia Sonami. 2020. Reflections on Eight Years of Instrument Creation with Machine Learning. In *Proceedings of the 20th International Conference on New Interfaces for Musical Expression (NIME'20)*.
- [33] Rebecca Fiebrink, Dan Trueman, and Perry R Cook. 2009. A Meta-Instrument for Interactive, On-the-Fly Machine Learning. In *NIME*. 280–285.
- [34] Jules Françoise and Frederic Bevilacqua. 2018. Motion-sound mapping through interaction: An approach to user-centered design of auditory feedback using machine learning. *ACM Transactions on Interactive Intelligent Systems (TiIS)* 8, 2 (2018), 1–30.
- [35] Jules Françoise, Yves Candau, Sarah Fdili Alaoui, and Thecla Schiphorst. 2017. Designing for kinesthetic awareness: Revealing user experiences through second-person inquiry. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 5171–5183.
- [36] Jules Françoise, Norbert Schnell, Riccardo Borghesi, and Frédéric Bevilacqua. 2014. Probabilistic models for designing motion and sound relationships. In *Proceedings of the 2014 international conference on new interfaces for musical expression*. 287–292.
- [37] Christopher Frauenberger. 2019. Entanglement HCI The Next Wave? *ACM Transactions on Computer-Human Interaction (TOCHI)* 27, 1 (2019), 1–27.
- [38] Martin Friedmann, Thad Starner, and Alex Pentland. 1992. Device synchronization using an optimal linear filter. In *Proceedings of the 1992 Symposium on Interactive 3D Graphics*. 57–62.
- [39] Petra Gemeinboeck. 2020. The aesthetics of encounter: a relational-performative design approach to human-robot interaction. *Frontiers in Robotics and AI* 7 (2020), 217.
- [40] Petra Gemeinboeck and Rob Saunders. 2013. Creative Machine Performance: Computational Creativity and Robotic Art. In *ICCC*. 215–219.
- [41] Elisa Giaccardi and Elvin Karana. 2015. Foundations of materials experience: An approach for HCI. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 2447–2456.
- [42] Elisa Giaccardi, Chris Speed, Nazli Cila, and M Caldwell. 2016. Things as co-ethnographers: Implications of a thing perspective for design and anthropology. *Design anthropological futures* 235 (2016).
- [43] Marco Gillies, Rebecca Fiebrink, Atau Tanaka, Jérémie Garcia, Frederic Bevilacqua, Alexis Heloir, Fabrizio Nunnari, Wendy Mackay, Saleema Amershi, Bongshin Lee, et al. 2016. Human-centred machine learning. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 3558–3565.
- [44] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.
- [45] Sabrina Hauser, Doenja Oogjes, Ron Wakkary, and Peter-Paul Verbeek. 2018. An annotated portfolio on doing postphenomenology through research products. In *Proceedings of the 2018 designing interactive systems conference*. 459–471.

- [46] Peter Henderson, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, and Joelle Pineau. 2020. Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning. *arXiv preprint arXiv:2002.05651* (2020).
- [47] Cher M Hill. 2017. More-than-reflective practice: Becoming a diffractive practitioner. *Teacher Learning and Professional Development* 2, 1 (2017).
- [48] Kristina Höök. 2018. *Designing with the body: somaesthetic interaction design*. MIT Press.
- [49] Kristina Höök and Jonas Löwgren. 2012. Strong concepts: Intermediate-level knowledge in interaction design research. *ACM Transactions on Computer-Human Interaction (TOCHI)* 19, 3 (2012), 23.
- [50] Laewoo Kang, Steven J Jackson, and Phoebe Sengers. 2018. Intermodulation: Improvisation and Collaborative Art Practice for HCI. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [51] Ashish Kapoor, Bongshin Lee, Desney Tan, and Eric Horvitz. 2010. Interactive optimization for steering machine classification. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1343–1352.
- [52] Chris Kiefer and Thor Magnusson. 2019. Live coding machine learning and machine listening: a survey on the design of languages and environments for live coding. In *Proceedings of the International Conference on Live Coding*. ICLC.
- [53] Matthew Lee and David Wessel. 1992. Connectionist models for real-time control of synthesis and compositional algorithms. In *Proceedings of the International Computer Music Conference*. 277–277.
- [54] Florent Levillain and Elisabetta Zibetti. 2017. Behavioral objects: The rise of the evocative machines. *Journal of Human-Robot Interaction* 6, 1 (2017), 4–24.
- [55] Golan Levin and Tega Brain. 2020. *Code as Creative Medium: A Handbook for Computational Art and Design*. MIT Press.
- [56] Wanyu Liu, Rafael Lucas d'Oliveira, Michel Beaudouin-Lafon, and Olivier Rioul. 2017. Bignav: Bayesian information gain for guiding multiscale navigation. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 5869–5880.
- [57] Tod Machover. 1989. Hyperinstrument: Musically Intelligent and Interactive Performance and Creativity Systems. *Proceedings of 1989 ICMC* (1989), 186–190.
- [58] Nirav Malsattar, Tomo Kihara, and Elisa Giaccardi. 2019. Designing and Prototyping from the Perspective of AI in the Wild. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. 1083–1088.
- [59] Benjamin Matuszewski, Joseph Larralde, and Frédéric Bevilacqua. 2018. Designing Movement Driven Audio Applications Using a Web-Based Interactive Machine Learning Toolkit. In *Web Audio Conference (WAC)*.
- [60] Marek P Michalowski, Selma Sabanovic, and Hideki Kozima. 2007. A dancing robot for rhythmic social interaction. In *Proceedings of the ACM/IEEE international conference on Human-robot interaction*. 89–96.
- [61] Kevin P Murphy. 2012. *Machine learning: a probabilistic perspective*. MIT press.
- [62] Iohanna Nicenboim, Elisa Giaccardi, Marie Louise Juul Søndergaard, Anuradha Venugopal Reddy, Yolande Strengers, James Pierce, and Johan Redström. 2020. More-Than-Human Design and AI: In Conversation with Agents. In *Companion Publication of the 2020 ACM Designing Interactive Systems Conference*. 397–400.
- [63] Antti Oulasvirta. 2018. *Computational interaction*. Oxford University Press.
- [64] Allison Parrish. 2018. *Articulations*. Counterpath Press.
- [65] Simon Penny. 2017. *Making sense: Cognition, computing, art, and embodiment*. MIT Press.
- [66] Claudio S Pinhanez. 1996. Computer theater. In *Proc. of the Eighth International Symposium on Electronic Arts (SEA'97)*. Citeseer.
- [67] Ivan Poupyrev, Michael J Lyons, Sidney Fels, and Tina Blaine. 2001. New interfaces for musical expression. In *CHI'01 Extended Abstracts on Human Factors in Computing Systems*. 491–492.
- [68] Jane Prophet and Helen Pritchard. 2015. Diffractive art practices: Computation and the messy entanglements between mainstream contemporary art, and new media art. *artnodes* 15 (2015).
- [69] Jane Prophet and Helen Pritchard. 2015. Performative apparatus and diffractive practices: An account of artificial life art. *Artificial Life* 21, 3 (2015), 332–343.
- [70] Gonzalo Ramos, Jina Suh, Soroush Ghorashi, Christopher Meek, Richard Banks, Saleema Amershi, Rebecca Fiebrink, Alison Smith-Renner, and Gagan Bansal. 2019. Emerging Perspectives in Human-Centered Machine Learning. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, W11.
- [71] Cynthia Rudin. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* 1, 5 (2019), 206–215.
- [72] Téó Sanchez, Baptiste Caramiaux, Jules Françoise, Frédéric Bevilacqua, and Wendy Mackay. 2021. How do People Train a Machine? Strategies and (Mis)Understandings. In *Proceedings of the 24th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*. ACM.
- [73] Diemo Schwarz, Grégory Beller, Bruno Verbrugge, and Sam Britton. 2006. Real-time corpus-based concatenative synthesis with catart. In *9th International Conference on Digital Audio Effects (DAFx)*. 279–282.
- [74] Hugo Scurto, Frédéric Bevilacqua, and Jules Françoise. 2017. Shaping and Exploring Interactive Motion-Sound Mappings Using Online Clustering Techniques. In *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME'17)*.
- [75] Hugo Scurto, Rebecca Fiebrink, et al. 2016. Grab-and-play mapping: Creative machine learning approaches for musical inclusion and exploration. In *Proceedings of the 2016 International Computer Music Conference*.
- [76] Hugo Scurto, Bavo Van Kerrebroeck, Baptiste Caramiaux, and Frédéric Bevilacqua. 2021. Designing Deep Reinforcement Learning for Human Parameter Exploration. *ACM Trans. Comput.-Hum. Interact.* 28, 1, Article 1 (Jan. 2021), 35 pages. <https://doi.org/10.1145/3414472>
- [77] Richard S Sutton and Andrew G Barto. 2011. *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- [78] Hillevi Lenz Taguchi. 2012. A diffractive and Deleuzian approach to analysing interview data. *Feminist theory* 13, 3 (2012), 265–281.
- [79] Koray Tahiroglu, Miranda Kastemaa, and Oskar Koli. 2020. AI-terity: Non-Rigid Musical Instrument with Artificial Intelligence Applied to Real-Time Audio Synthesis. In *Proceedings of the 2020 International Conference on New Interfaces for Musical Expression*.
- [80] Koray Tahiroglu, Thomas Svedström, and Valter Wikström. 2015. Noisa: A novel intelligent system facilitating smart interaction. In *Proceedings of the 33rd annual acm conference extended abstracts on human factors in computing systems*. 279–282.
- [81] Samuel Thompson, Hugo Scurto, and Rebecca Fiebrink. 2019. Sound Control: Supporting Custom Musical Interface Design for Children with Disabilities. In *Proceedings of the 19th International Conference on New Interfaces for Musical Expression (NIME 2019)*.
- [82] Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 5026–5033.
- [83] Elena Tosi Brandi. 2018. *Design de l'expérience sensible: une philosophie des sens pour le design et la création*. Ph.D. Dissertation. Paris Sciences et Lettres.
- [84] Nikolaus Troje, Francesca Simion, Lara Bardi, Elena Mascalon, Lucia Regolin, Emily Grossman, Maggie Shiffar, James P Thomas, Phil McAleer, Scott Love, et al. 2013. *Social perception: Detection and interpretation of animacy, agency, and intention*. MIT Press.
- [85] Anna Vallgård. 2014. Giving form to computational things: developing a practice of interaction design. *Personal and ubiquitous computing* 18, 3 (2014), 577–592.
- [86] Kiri L Wagstaff. 2012. Machine Learning that Matters. In *In Proc of ICML*. Citeseer.
- [87] Danding Wang, Qian Yang, Ashraf Abdul, and Brian Y Lim. 2019. Designing theory-driven user-centric explainable AI. In *Proceedings of the 2019 CHI conference on human factors in computing systems*. 1–15.
- [88] Mikael Wiberg. 2014. Methodology for materiality: interaction design research through a material lens. *Personal and ubiquitous computing* 18, 3 (2014), 625–636.
- [89] Qian Yang. 2018. Machine learning as a UX design material: How can we imagine beyond automation, recommenders, and reminders?. In *AAAI Spring Symposia*.
- [90] Qian Yang, Nikola Banovic, and John Zimmerman. 2018. Mapping Machine Learning Advances from HCI Research to Reveal Starting Places for Design Innovation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 130.
- [91] Qian Yang, Alex Scuito, John Zimmerman, Jodi Forlizzi, and Aaron Steinfeld. 2018. Investigating How Experienced UX Designers Effectively Work with Machine Learning. In *Proceedings of the 2018 Designing Interactive Systems Conference*. ACM, 585–596.
- [92] Qian Yang, Jina Suh, Nan-Chen Chen, and Gonzalo Ramos. 2018. Grounding interactive machine learning tool design in how non-experts actually build models. In *Proceedings of the 2018 Designing Interactive Systems Conference*. 573–584.
- [93] Mo H Zareei, Dale A Carnegie, Ajay Kapur, and Dugal McKinnon. 2014. Mutor: Drone chorus of metrically muted motors. In *International Computer Music Conference*.
- [94] Michael Zbyszynski, Mick Grierson, Matthew Yee-King, et al. 2017. Rapid prototyping of new instruments with codecircle. In *Proceedings of the 17th International Conference on New Interfaces for Musical Expression*.