

# Understanding the role of Interactive Machine Learning in Movement Interaction Design

MARCO GILLIES, Department of Computing, Goldsmiths, University of London, UK

Interaction based on human movement has the potential to become an important new paradigm of human computer interaction. However, high quality, mainstream movement interaction requires effective tools and techniques to support designers. A promising approach to movement interaction design is Interactive Machine Learning, in which designing is done by physically performing an action. This paper brings together many different perspectives on understand human movement knowledge and movement interaction. This understanding shows that the embodied knowledge involved in movement interaction is very different from the representational knowledge involved in a traditional interface, so a very different approach to design is needed. We apply this knowledge to understanding why interactive machine learning is an effective tool for motion interaction designers and to make a number of suggestions for future development of the technique.

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

## ACM Reference Format:

Marco Gillies. 2010. Understanding the role of Interactive Machine Learning in Movement Interaction Design. *ACM Trans. Comput.-Hum. Interact.* 9, 4, Article 39 (March 2010), 35 pages. <https://doi.org/0000001.0000001>

## 1 INTRODUCTION

Interaction paradigms that make use of human movement have been one of the most exciting frontiers of HCI research in the last decade. A range of tracking technologies have made it possible to use gestures and full body motion as a means of interaction. The increasingly low cost of the sensors means that this form of interaction could become very common and possibly a new major form of interaction after mouse and keyboard Graphical User Interfaces and touch screens.

However, the uptake of movement interfaces has been limited relative to volume of research prototypes. There are many potential reasons for this, but these include critiques by Norman [Norman 2010] that many of the movement interfaces that have developed have poor interaction design and fail to live up to their potential as “Natural Interaction” [Gillies and Kleinsmith 2014]. This points to a flaw in the way movement interactions are currently designed, which is itself potentially related to limitations in the tools currently used for that design.

The ubiquity of Graphical User Interfaces would not have been possible without the development of toolkits that make it easy to assemble new interfaces either by programming or using a graphical tool. Making movement interaction mainstream will mean making effective toolkits for designing it. While many sensor vendors such as leap motion (<https://www.leapmotion.com>) supply toolkits for designing interaction, these are often limited to a small set of standard gestures. A number of academic authors, on the other hand have developed methods that allow authoring of new

---

Author’s address: Marco Gillies, Department of Computing, Goldsmiths, University of London, New Cross, London, SE24 9PA, UK, [m.gillies@gold.ac.uk](mailto:m.gillies@gold.ac.uk).

---

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

© 2010 Association for Computing Machinery.

1073-0516/2010/3-ART39 \$15.00

<https://doi.org/0000001.0000001>

gestures and interactions by example [Hartmann et al. 2007; Wobbrock et al. 2007; Zamborlin et al. 2014] often using an approach called Interactive Machine Learning (IML) [Caramiaux et al. 2015a; Fiebrink et al. 2011] which is quite unlike the approach taken in a GUI toolkit. This paper seeks to understand how to make these toolkits. Can they work in a similar way to GUI toolkits or do they require a completely different approach? This paper explores the nature of movement interaction and uses this understanding to an analysis of why IML is suitable for movement interaction design and how it can be improved.

This paper makes the follow contributions:

- It brings together a number of different disciplinary perspectives to form an understanding of movement interaction, human movement knowledge and how people learn and teach movement knowledge.
- It uses this understanding to develop a theoretical model of how we should design and implement movement interaction.
- This model is applied to understanding why interactive machine learning is an effective way of designing movement interaction and highlights a number of features that support this design: designing by moving; a tight cycle of interaction, and supporting reflection on the experience of movement. It also uses this model to evaluate Interactive Machine Learning as an approach and suggest future research directions.

The aim of this work is to deepen our understanding of movement interaction and help guide the design of future movement interaction design toolkits.

## 2 WHAT IS MOVEMENT INTERACTION?

A simple definition of movement interaction is any technique for human computer interaction that makes use of body movement in a different, and fuller, way than mainstream interfaces such as mouse and keyboard or touchscreen. However, this encompasses a very wide range of possible applications, technologies, interaction styles and design techniques. While it is impossible to cover the full range of movement interaction in a paper of this length, the following section will attempt to provide and outline of the diversity of approaches.

### 2.1 Application areas

Movement interaction first entered the mass market in video games, and a lot of research still relates to games and entertainment, for example, Hämäläinen et al.'s work on movement and exertion games [Hämäläinen et al. 2015]. However, there are many other application areas. Much of the movement games research relates closely to sport, for example using projections to create a game out of a climbing wall activity [Kajastila et al. 2016] or allowing people to run together when geographically separated [Mueller et al. 2012]. Movement interaction has also been used in other ways in sport, particularly education, for example Velloso et al. [Velloso et al. 2013] developed a system that used motion sensing to give feedback to people learning weight training in order to improve their technique. Sport is not the only domain where physical movement and technique are important and so forms of movement interaction can help support many different forms of education, for example posture has a great effect on musical performance and movement feedback can be used effectively in initially training [Johnson et al. 2013]. The field of New Interfaces for Musical Expression (NIME) has seen many developments that use movement interaction as a means of musical performance with electronic instruments [Caramiaux and Tanaka 2013; Fiebrink et al. 2011; Françoise et al. 2013]. A closely related field is dance where movement interaction has been used to enhance the experience of professional dancers [Fdili Alaoui et al. 2015] and members of the public [Halpern et al. 2011]. Another important area for movement interaction is healthcare,

where movement games can support rehabilitation exercises for people with conditions such as parkinsons [McNaney et al. 2015] or stroke [Kirk et al. 2016]. Also health related is the development of custom interfaces for people with disabilities, for example the work of Katan *et al.* [Katan et al. 2015].

## 2.2 Sensing technologies

Movement interaction has seen mainstream interest in the last decade, primarily due to gaming devices such as the Nintendo Wii™ and Microsoft Kinect™. These are sensor technologies that are able to track human movement in detail, and their ready availability has triggered a large interest on the part of academic researchers in developing new forms of interaction that use human body movement. The sensors themselves take many forms. The Nintendo Wii™ primarily uses an accelerometer. Accelerometers are cheap, small and found in many devices including most mobile phones, making them an easily accessible way to sense aspects of human movement (though the data is limited, they cannot recognise full body pose). They have been used by many researchers working with movement interfaces [Caramiaux et al. 2015a; Zamborlin et al. 2014]. The Kinect™ on the other hand is based on computer vision, using an infrared structured light pattern to determine a depth image with is then used to determine a rough estimate of body pose. Vision based systems can give more detailed information and are therefore popular with researchers [Fdili Alaoui et al. 2015; Gillies et al. 2015a]. Other techniques include muscle sensing such as Electromyography (for example the commercial device Myo (<https://www.myo.com>), Mechanomyography (for example used by Caramiaux, Donarumma and Tanaka [Caramiaux et al. 2015b]) or two way interaction via electrical muscle stimulation [Lopes et al. 2015]. There has been also been considerable research into new techniques, such as sensing via the electrical noise that is ubiquitously present in modern homes [Cohn et al. 2011].

## 2.3 Interaction styles

These sensing techniques, and others, have been used to develop a wide range of interaction styles. For the purpose of this paper, we may distinguish three broad categories of movement interaction: object focused, direct mapping and movement focused.

**Object Focused Interaction** are forms of interface in which human movement is important, but the design is focused on an object of interaction, rather than the movement itself. Tangible user interfaces are a good example of this approach. The classic example of a tangible interface, URP by Underkoffler and Ishii [Underkoffler and Ishii 1999], is a tool for urban planning in which buildings are represented by physical models which are tracked by computer and overlaid with projected information such as daylight patterns. Users can interact with this in many ways by physically moving buildings around but also by moving their own viewpoint to the scene from different perspectives. These are interesting uses of body movement and are key to the effectiveness of the interaction. However, the focus of the design process is not the users' body movements, which are never represented or explicitly recognized by the system, but the objects themselves.

**Direct Mapping interaction** is a form of interface in which the movements of a user are directly mapped into some form of digital space. Examples, include a large proportion of Virtual Reality Interaction in which the important factor is seeing the user's body mapped into the VR space [Spanlang et al. 2014]. This type of interaction is often similar to object focused interaction, in that the key interaction design is focused on objects that users can interact with using whatever movements they choose (or at least those that are possible to track). The major difference is that in this case, the objects are virtual not physical. An

example of this form if interaction is a virtual button that is "clicked" by reaching out and touching it, the interaction is not determined by a specific movement by the user but simply the location of their hand that is mapped into virtual space. Unlike object focused interaction, there is a need to pay attention to human movement in design. However, the aims of this interaction design is quite clear and straightforward: to map movement accurately from the physical to digital domain (normally via tracking technology). This mapping is normally done in a standard way for a particular platform (e.g. for a particular Virtual Reality hardware system).

**Movement Focused Interaction**, on the other hand, is interaction design around specific body movements rather than objects. A typical example is the swipe gesture on a mobile phone. This does not rely simply on manipulating the phone directly, nor is it simply detecting the position of the user's finger. It is activated by a specific form of movement, and only that movement. The focus of interaction design in this case is not longer an object, real or virtual, but on the movement itself. For this reason movement focused interaction is the most radically divergent from traditional interaction design techniques which have focused on first 2D interfaces and then 3D objects.

These types are not mutually exclusive, there are many designs that include more than one type. For example, many wearable devices, particularly those used in dance [Karpashevich et al. 2018; Wilde 2010]. A wearable is a physical object that is amenable to object focused design, but since it is worn on the body it can affect movement. For example, in the work of Karpashevich et al. [Karpashevich et al. 2018] the design began with a physical object, a dress based on the Schlemmer's seminal costumes for the Bauhaus "Triadic Ballet". However, the costumes was designed in a way that would constrain and have implications for a dancers movements, so a part of the design was movement focused. This became particularly important as the focus shifted from design of the object, to its use in designing choreographies. The focus of this paper is on the design process, and a single artifact can go through multiple phases of design, in which it can be worked on as a form of object focused interactions and others where it is a form of movement focused interaction. For example, Márquez Segura et al. [Márquez Segura et al. 2013] use a pre-existing device, a motion sensing toy called BodyBug (now commercially available as Oriboo). This was presumably designed in an object focused way. However, they use this device to design new motion games, which could be classified as motion focused interaction design.

For the purpose of this paper, I will focus primarily on Movement Focused Interaction (or stages of the design processed that work on movement focused interaction). For both object focused and direct mapping interaction the locus of interaction design is on the object, whether real or virtual. The movement itself is either outside the digital domain (in the case of object focused interaction) or uses a highly standardized mapping to the digital domain (direct mapping interaction). This means that existing interaction design techniques can be well suited to these forms of interaction. While their 3D nature may make 2D paper sketches and prototypes unsuitable, 3D prototyping techniques work well. For movement focused interaction, the locus of design is on the human movement itself, which requires very different design approaches.

Movement focused interaction is itself very diverse, comprising a wide range of interaction styles. The most common being gestural interaction, in which the computer recognizes a number of gestures which are separate movements, typically performed with the hands, each having a characteristic shape. Much work on gesture has focused on 2D gestures on touch screens [Wobbrock et al. 2007]. However, there has also been considerable work on 3D free space gestures that make a fuller use of body movement [Caramiaux et al. 2015b; Thórisson 1998; Zamborlin et al. 2014]. Gestural interaction tends to be relatively constrained, both to a limited area of the body (typically

hands) and to a limited vocabulary of recognized gestures. However, other work has experimented with freer, full body movement, for example for dance [Fdili Alaoui et al. 2015; Halpern et al. 2011] or exercise [Kajastila and Hämäläinen 2015; Mueller and Isbister 2014]. Some of this work focuses on expressive and enjoyable aspects of movement allowing for relatively free form experience, while others are closer to quantified self applications that aim to track specific activities such as running [Mueller et al. 2012]. Most of the work just cited uses movement interactions that relate to our interaction with our physical environment or virtual proxies of it, but our body movements are also a key part of our *social* interactions. We use body language (or more formally, non-verbal communication) in all of our face-to-face communications and it can carry many implicit and explicit messages. For example, Thórisson [Thórisson 1998] used explicit gestures as a means of interaction with a virtual character, while Gillies, Brenton and Kleinsmith [Gillies et al. 2015a] used full body movements as an interaction in the absence of speech and Huang, Morency and Gratch [Huang et al. 2010] detected subtle non-verbal cues to allow a virtual character to give non-verbal feedback to a speaker. Social movement interaction has also been used together with natural language understanding systems, for example, using head movements to establish when a concept has been grounded [Vosoughi 2014].

Movement interaction is often associated with particular display technologies. For example, forms of movement interaction are considered by many [Gillies 2016; Slater 2009] to be as important to Virtual Reality. At the very minimum the ability to change your viewpoint on a scene by turning or moving your head is vital to creating presence. Other displays that are suited to body movement interaction are very large displays where users need to stand back in order to take in all of the information [Nancel et al. 2011].

The majority of examples above have involved explicit interaction in that users are consciously choosing to make movements (even if they are not conscious of the full details of that movement). This creates an explicit dialog for action and response between human and artifact. However, there are also forms of interaction that users are less consciously aware, and are less structured as a dialog. For example, a fitness tracker will collect data about a users activity without them ever consciously interacting or performing particular movements to be recognized. This is an example of monitoring a user, but this form of implicit interaction can also take the form of a computational system that has its own behavior which continues independently of users, but which can be modulated by user behavior, similar to what Ingold terms *Correspondence* [Ingold 2017]. An example of this might be a plant in a virtual reality simulation, which moves under a physics model independent of users, but may waft in a particular way as they pass, a form of responsiveness that is barely noticeable but can increase the sense of presence.

From the point of view of design these forms of interaction are not too different from more explicit interaction. While, users may not be consciously aware of them, designers must be in order to design them, and much of what follows applies equally. All three types of interaction described in this section can take this, more implicit form. The VR plant, might be a form of object focused interaction: it's behavior is determined by a physics engine, and any response to a user is a consequence of this and not explicitly designed. A fitness tracker that simply records intensity of movement is similar to direct mapping: the movement is simply turned into numbers without attempting to interpret it. On the other hand, if the fitness tracker attempts to recognize particular types of exercise, for example running or cycling, then the focus is on particular movement which would be a form of movement focused interaction.

## 2.4 Design Approaches

Movement focused interaction requires new approaches to interaction design. A design and prototype can no longer be embodied in a sketch or an object, but has to include the user's movement. It

is also not enough to view this movement from the outside. While a video prototype [Mackay and Fayard 1999], in which a person is shown performing the movement, may give the sense of what the movement is, but it does not give a sense of what it is like to do the movement. This requires a *first person* perspective on interaction design [Höök et al. 2018], in which designers focus on their own experience of movement. This implies that we should design by doing and moving [Hummels et al. 2007; Kleinsmith and Gillies 2013] rather than by representing (for example, with a sketch or video). The movement interaction community is gradually developing a range of design approaches focused on first person designing by doing. For example, *Embodied Sketching* [Márquez Segura et al. 2016] encourages a hands-on, activity centered approach to design in which designers or participants engage in physical activity at different phases of the design process including sensitization and brainstorming. Another important feature of embodied sketching is that it happens over time, not in a single instant. This means that designers can experience the changing experience of movement, but also that it supports many cycles of reflection and refinement, as designers reflect on those changing experiences of movement as they happen. Another related approach is Núñez-Pacheco and Loke's [Loke and Núñez-Pacheco 2018; Núñez-Pacheco and Loke 2018] use of Focusing in design workshops, which encourages designers to become aware of the first person experience of embodiment.

This paper will explore why these first person, embodied approaches work well for movement interaction.

### 3 INTERACTIVE MACHINE LEARNING

While embodied sketching is a good model for the ideation phases of movement interaction design, it cannot be entirely divorced from the process of implementation as noted by Françoise *et al.* [Françoise et al. 2017]. Designing by performing physical actions is pointless if it is not technically possible to recognize those action in a computer system. Also, performing an action without an implementation is fundamentally different from performing that action with feedback from technology and in the presence of tracking errors and limitations that are inevitable with movement based technology. Implemented prototypes become more important for movement interaction because there is no equivalent of a "paper prototype" that can capture the experience of movement interaction. Implementation techniques have to be fast if they are to enable iterative prototyping, however, this often limits interactions to simple recognition techniques that do not enable the full potential of embodied, movement based interaction.

A popular emerging approach to prototyping movement interaction is Interactive Machine Learning (IML) [Amershi et al. 2014; Fails and Olsen 2003]. Machine learning is a class of technology in which recognition systems are learned from a number of examples rather than being explicitly programmed. Machine learning is traditionally a batch process in which a large dataset is used to learn a model automatically, with little or no human interaction. Interactive Machine Learning, on the other hand, is an approach to machine learning which places much more importance of the role of the human in machine learning. Rather than the batch process of standard learning, the learning process is viewed as highly interactive, with the human user guiding the learning process by carefully selecting data and constantly testing the results. This creates a tight loop of interaction in which human and machine collaborate to create a model:

Rapid, focused and incremental learning cycles result in a tight coupling between the user and the system, where the two influence one another. As a result it is difficult to decouple their influence on the resulting model and study such systems in isolation. [Amershi et al. 2014]

IML is part of a broader movement of Human-Centered Machine Learning [Gillies et al. 2016], which foregrounds the role of humans in machine learning and also Machine Teaching [Simard et al. 2017] which the role of the human "machine teacher" as key to the success of learning.

Interactive Machine Learning is normally framed in terms of supervised learning. Supervised Learning is a sub-class of machine learning problems in which the computer is presented with a number of examples of data each of which has a "label" representing the correct output of the system. The computer must learn a mapping from input data to output labels. The labels could be discrete categories in which case the task is called classification (for example, classifying whether an image contains a cat or a dog) or a continuous signal, in which case the task is called regression (for example, a musical interface that learns a continuous mapping between body movement and sound). In IML, a designer can "teach" a system by providing labeled training examples, and further refine the system by adding new example in cases where the current output is wrong or undesirable.

IML may be seen as similar to Active Learning [Settles 2010], which also features a tight loop of interaction between human and machine learner. However, the fundamental difference is where the initiative lies. In active learning the initiative lies with the learning algorithm. The human acts as an "oracle" for the learner, labeling any data that the algorithm requests. The human is in the loop of the machine learner, but the loop still belongs to the machine. In IML, on the other hand, the human has the initiative, selecting data and choosing what to label. They correct behavior of the learner with more data, rather than relying on the machines judgment of what has been learned well or not. The machine is now in the loop of human activity, making the learning process much closer to a design process.

Interactive Machine Learning makes it possible to design a system by providing example of the correct behaviour. Designing by example goes back to early work on programming by demonstration [Hartmann et al. 2007; Lieberman 2001] and is in fact a popular method for designing gestural and movement interfaces, for example defined gestures by drawn templates [Wobbrock et al. 2007]. This style of design by example has been implemented using IML in many gesture recognition and movement interaction systems [Caramiaux et al. 2015a; Fiebrink et al. 2011].

IML has been applied in a wide range of domains from email classification [Kulesza et al. 2015] to network alarm triage [Amershi et al. 2011]. It has been particularly popular in movement interaction. For example, Fiebrink *et al.* [Fiebrink et al. 2011] and Zamborlin *et al.* [Zamborlin et al. 2014] used it as a means of electronic musicians to design new movement based instruments. Françoise *et al.* [Françoise et al. 2017] also develop musical interfaces, but in this case use IML for the design of an interactive, public installation. Caramiaux *et al.* [Caramiaux et al. 2015a] use interactive machine learning for rapid prototyping of movement in a participatory design context and Gillies *et al.* [Gillies et al. 2015a,b] have used it for interactive 3D characters. We will return to some of these examples later in order to better understand them, but they point to IML as a successful method for designing movement interaction. The rest of this paper will attempt to explain the success of IML through a better understanding of the nature of movement interaction and the challenges involved in designing it.

#### 4 WHY MOVEMENT INTERACTION?

Even this short and far from exhaustive overview shows that movement interaction is a large and complex domain that uses many techniques and spans many areas of application. It is also very different from traditional desktop or touch screen interaction. This raises the question of why people have chosen to implement movement interaction, which is far more technically difficult than a traditional interface. Is it simply a gimmick or are there important benefits? Only by understanding this can we think about how to design movement interaction.

To many the answer to the question of whether and why movement interaction is a good idea is a straightforward yes. Recent years have seen movement interaction commonly (particularly in industry) described as “Natural User Interfaces”. The implication being that there is something innately natural for us about using body movement and that using a wider range of movement will automatically make human computer interaction easier. Norman [Norman 2010] has criticized this view, pointing out that many gestural interfaces are in fact very hard to use because they lack a number of features of good interfaces that Norman proposed in his seminal work [Norman 2013]. For example they are not discoverable in the sense that it is hard to know what gesture to do and we often cannot look it up in the way we would look at a menu to see what options are on it. This also makes them less memorable: there is no support in remembering the correct gesture. Finally, many gestural interfaces do not provide good feedback on whether they are being performed correctly. There is nothing so natural about body movement that means that the rules of good interaction design are suspended, some movements may be natural to learn by many are awkward, tiring and hard to remember. So if movement interfaces are not innately natural and can be harder to use than a screen based interface, are there reason, at least in certain situations, why they are valuable and worth doing? <sup>1</sup>

A simple answer to this question is that full body exertion can be healthier than sitting virtually motionless in front of a screen. This certainly seems to be the reason for some of the work described above, for example the exertion games [Hämäläinen et al. 2015], but for most examples there seems to be more to it than simply exercise.

#### 4.1 Reality-based interaction

Another common argument is that movement interaction mimics our interactions with the physical (and possibly social) world. Slater [Slater 2009] proposes that movement is key to virtual reality because the relationship between our movements (turning our head) and what we see on the changing display (our view of the scene turning) reproduces the *sensori-motor contingencies* that we experience in the real world. Sensori-motor contingencies are a term from O’Regan and Noë’s theory of perception [O’Regan and Noë 2001]. They propose that our perception of the world is based on learned relationships between our movements and our senses. According to this theory our experience of living in a 3D world is at least in part determined by the way our retinal image changes as we move our head. Traditional screen displays do not change in this way as we move our head and so they cannot recreate important aspects of our experience of being in the world. However, a virtual reality head-set can do so as it includes both head mounted displays in front of our eyes and a gyroscope to detect head rotation. In a properly set up virtual reality environment our head movements and view direction are linking in exactly the same way as they are in the real world. Slater proposes that this is a feature that creates a sense of presence he calls “*Place Illusion*”, the illusion of physically being in a virtual space that can be created in VR but not on screen.

A more general statement of this view can be found in Jacob *et al.* [Jacob et al. 2008]. They identify a number of interaction techniques, including movement interaction, tangible interaction and virtual reality, that they class as “*Reality-Based Interaction*”. These are forms of human computer interaction that reproduce at least some element of our interactions with the real world, whether that relates to the physical environment, our bodies or our social interactions with other people. A reality-based interface does not have to reproduce every single aspect of the real world, in fact to do so would be pointless as the technology needs to add something over and above what the real world can do, but reproducing key aspects of our interactions with the world can create usable forms of

---

<sup>1</sup>Gillies has examined this question in detail [Gillies 2016; Gillies and Kleinsmith 2014] and this section and the next will describe an expanded version of his argument.



interaction. The reason for this, according to Jacob *et al.* is that reality-based interaction allows us to make use of our existing skills for interacting with the world including our understanding of physics, our own bodies, our awareness of our environment and our social skills. These are all skills that we have learned from infancy and are strongly grounded in us, so we should be able to easily learn any interaction that makes use of them (as long as, following Norman, that interaction is discoverable and provides sufficient feedback).

This reality-based view helps explain why movement interaction can be valuable and also gives some guidance on what makes a good or bad movement based interaction. Movement based interaction is useful if it reproduces key aspects of our interaction with the real world and allows us to use our existing skills of interacting with the world.

## 4.2 Embodied Cognition

Another argument for movement interaction is that our body movements are closely related to our cognition and mental state. The theory of embodied cognition [Kirsh 2013] states that our physical movements are central to our thinking and cognition: we think, at least in part, by moving. Kirsh [Kirsh 2013] gives the example of assembling flat-pack furniture. We (or at least most people) cannot understand how to assemble the furniture by sitting in a chair thinking about it abstractly with our eyes closed, only moving to act when we have completely solved the problem in our head. We understand the problem as much with our movements as with our thoughts, or in fact we understand it with the two inseparable combined. We think about how to fit two sections together by moving to look at them from different angles; picking them up and manipulating them to see how they might join. The concept of embodied cognition can help us understand better the real world skills that Jacob *et al.* see as central to reality based interaction. In the most part they involve exactly the combination of movement and cognition that embodied cognition proposes. This also suggests that movement can also be valuable if they allow us to think about problems in an embodied way. In fact this is a key benefit of a lot of tangible interaction including the URP system [Underkoffler and Ishii 1999] described earlier. This might be less applicable to the types of movement based interaction we are discussing in this paper, but a focus on movement could be important if the cognition relates to the body more than the physical world.

The embodied nature of cognition is closely related to the embodied nature of emotion. A number of studies have shown that the relationship between emotional experience and body movement is complex and two way, with the body having a powerful influence on emotion. For example, Bianchi-Berthouze and colleagues [Bianchi-Berthouze et al. 2007; Lindley et al. 2008] have shown in number of studies, that using full body movements when playing video games results in greater emotional engagement. Similarly, Wells and Petty [Wells and Petty 1980] found that participants who moved their heads in a nodding movement while listening to an audio of speech agreed more with the message than those who moved their head in a shaking motion. These results link to more theoretical ideas of the role of the body in emotion such as Damasio's Somatic Marker Hypothesis [Damasio 1994] and even going back as far as William James [James 1890]. All this suggest that movement interfaces can have powerful emotional effects that more traditional interfaces cannot, a factor that is used extensively by many User Experience researchers, for example Höök and her team [Hook 2009].

So to summarize, it is naïve to assume that movement interaction will automatically be more natural than a traditional user interface, but it does bring a number of advantages, if well designed. It can tap our embodied cognition via our embodied skills, knowledge and emotions. However, the movement interaction must be designed in such a way that it can make use of our embodied knowledge and support our reflection on it. In order to do that we need a better understanding

of embodied knowledge and how it differs from the type of knowledge we typically bring to a graphical user interface.

## 5 MOVEMENT KNOWLEDGE

The idea that there is a distinctive type of knowledge associated with the body goes back at least to Merleau-Ponty [Merleau-Ponty 2002]. Tanaka [Tanaka 2013] defines this type of knowledge as “*embodied knowledge*” drawing on Merleau-Ponty’s philosophy. He describes embodied knowledge in this way:

Embodied knowledge is a type of knowledge in which the body knows how to act. A simple and general example is riding a bicycle. Most of us know how to ride a bicycle, and we can do so without any deliberation. There is no need to verbalize or represent in the mind all the procedures required. The knowledge of how to ride a bicycle seems to be imprinted in one’s body and just lived through it, without being consciously represented. Thus, the knowing subject here is not the mind but the body. [Tanaka 2013]

Tanaka shows that this type of knowledge draws together a number of strands of thought about knowledge. It is closely related to embodied cognition, in which, as we have seen, our cognitive processes are tightly entwined with our bodily actions and interactions with our environment. Tanaka also quotes Ryle [Ryle 1949], who distinguishes “*knowing that*” (for example, know that a bicycle has two wheels) from “*knowing how*” (for example, knowing how to ride a bicycle). Embodied knowledge is of the “*knowing how*” type which has traditionally been ignored by more declarative theories of knowledge. Embodied knowledge is also *tacit*, defined by Polanyi [Polanyi 1966] as the way we know “*more than we can tell*”: knowledge we have but cannot express in language. In general Tanaka sees embodied knowledge as non-representational, i.e it does not involve representation in language or other symbolic forms. This seems to make embodied knowledge *implicit* which Kirsh [Kirsh 2009] defines as “*knowledge that cannot be elicited, cannot be made directly conscious, and cannot be articulated*” (though, as we shall see later, Kirsh is actually quite critical of this idea).

Tacit embodied knowledge is a key aspect to understanding movement interaction and its design. However, this view of movement knowledge is, in fact, a simplification in a number of important ways. In particular we must understand how explicit knowledge interacts with tacit knowledge. The following sections will help understand that interplay by talking about how we learn motor skills (individually), how we reflect on our body and movement and how we teach and communicate movement knowledge.

### 5.1 Motor Learning

A mature motor skill, such as riding a bicycle, is likely to be entirely tacit and subconscious, however, this is not the case while we are learning that skill. Motor skill learning is an interplay between conscious and subconscious. Fitts and Posner’s classic model [Fitts and Posner 1967] of motor learning consists of three stages, in which explicit and implicit knowledge are used differently.

**The Cognitive Stage** involves explicit setting of goals and planning of the actions. The knowledge here is largely explicit and the work is primarily thinking rather than doing. In terms of riding a bicycle, this is planning how to mount the bicycle and checking the location of pedals and breaks.

**The Associative Stage** is where the learner begins the physical action, and involves an interplay of explicit and implicit. They must practice movements that can only be fully understood implicitly, but they must also plan the detailed sequence of movements, and reflect on cases where they are not successful, both of which are explicit. In terms of riding a bicycle, this

is the stage where learners must concentrate consciously on actions such as pedaling or breaking.

**The Autonomous Stage** is where the action is becoming fully (or almost fully) automatic. Learners must practice the skill in order to improve it but most of this practice involves embodied action rather than explicit cognition. This is the stage at which a learner can comfortably ride a bicycle and can concentrate on where they are going rather than on the pedals, but still need practice in order to improve.

So in the Fitts-Posner model, explicit, cognitive thought is as important as implicit embodied practice, though one gradually morphs into the other. Authors such as Taylor and Ivry [Taylor and Ivry 2012] stress the importance of cognitive strategies in motor skill learning. So movement knowledge is not purely tacit, even if there are many aspects that we can never put into words. The interplay of explicit and implicit knowledge is vital to learning in terms of both planning and reflection. But, if we cannot fully explain our movement in words, how can we think about them and how can we teach them? The following two sections will explore these questions.

## 5.2 Somatic Reflection

Deeper understanding of the role of the body in HCI has been developed by Schiphorst [Schiphorst 2009, 2011]; Núñez-Pacheco and Loke [Loke and Núñez-Pacheco 2018; Núñez-Pacheco and Loke 2018], and Höök and colleagues [Höök et al. 2016, 2015] by introducing Richard Shusterman's philosophy of Somaesthetics [Shusterman 2008]. Somaesthetics is the combination of two words: *Soma* meaning body and *aesthetics* meaning sensory experience. Like embodied cognition, Shusterman's philosophy places the body at the center of our understanding and cognition, but he also draws on movement or bodily practices such as yoga, meditation and Feldenkrais technique (of which Shusterman is a trained teacher). These techniques not only make use of the body as sport or dance might, they encourage reflection on the body. Feldenkrais involves slow movements that draw attention to the feelings of the body and can make people aware of uncomfortable positions or problems such as poor use of posture. Shusterman proposes that this form of reflection on our bodily experience can help improve our use of our bodies and break us out of bad habits of movement by making us more aware of them, and then changing the way we move and act. Both Schiphorst [Schiphorst 2009, 2011] and Höök [Höök et al. 2016, 2015] suggest that technologies can help to encourage this reflection if they are designed appropriately, and they show a number of designs, often movement based, that exemplify this idea. Movement interaction might, therefore, be able to not only use our embodied skills and knowledge but also improve them.

So while somaesthetics is closely aligned with embodied cognition it brings us to a very different picture from the previous section, in which knowledge is entirely tacit; we can only act without knowing how or why we act, and we cannot teach that knowledge only learn it from experience. Shusterman shows us that while it is not possible to put certain body movement experiences into words, we can be conscious of them and reflect on them. This is supported by Kirsh's critique of the concept of implicit knowledge [Kirsh 2009]. He proposes that knowledge can be explicit but non-verbal. Polanyi's original definition of tacit knowledge is that we know "*more than we can tell*" [Polanyi 1966]. This does not, in fact, imply that the knowledge is implicit in the sense that we cannot think about that knowledge explicitly, simply that we cannot put it into language. Shusterman's somatic reflection seems to be exactly that: knowledge that we can be conscious of but cannot put into words. In fact, even this is too strong, as some aspects can be put into words, all of the somatic practitioners we have discussed in this section, Shusterman, Feldenkrais, Schiphorst, Höök, have written extensively about their experiences. None of these writings can capture the

experience of movement in its entirety but they all capture something useful that contributes to that experience.

Another notable feature of bodily practices such as Feldenkrais is that they are done in the presence of, and learned from a teacher or facilitator, who is an expert, or in Schiphorst's phrase, a *connoisseur* of the practice. The role of the facilitator is vital, they guide learners in correct movements and help them be aware of the correct feeling of those movements. Without facilitators, students could easily fall into incorrect habits and fail to realize the subtle but essential details of the movements they should be doing.

Not only can we be conscious of embodied knowledge, but we can also communicate it and teach it. Feldenkrais teachers do teach embodied knowledge, even though they cannot fully explain it in words. How they do this is a complex combination of many elements. The next section will explore how this is done, as it can serve as a good model for movement interaction design (which will be explained in the follow section).

### 5.3 How do we communicate embodied knowledge?

The fact that we can reflect on our bodily knowledge and that we be taught it by being guided by an expert therefore challenges the idea that we can only learn by experience. While we cannot learn by being told, we can learn through a complex interaction with a teacher. What happens in this interaction? How is an embodied skill communicated if not (entirely) through language?

Höök has published a valuable autoethnography of learning an embodied skill: horseback riding in the English style [Höök 2010] (she was already proficient in riding the Icelandic style). This example illustrates the process well.

Her description certainly supports the idea that personal experience is vital to this kind of learning, and we cannot learn by being told:

This insight did not come from being told a visual metaphor ... or from Christian [her teacher] showing me through riding himself what it should look like. In my case, I had to experience it myself in order for it to be meaningful to me. [Höök 2010]

and

Suddenly, without consciously trying to, I shifted my weight further back in the saddle and suddenly it worked. ... Thinking is precisely not what is required, I was simply entirely absorbed by the situation. My muscles were doing what they should be doing, my eyes were directed towards where we should be going, my hearing followed the rhythm of the horse's hooves on the ground, and I felt as one with the horse. [Höök 2010]

However, Höök was not (and probably could not) learning to ride on her own, she had a teacher, Christian. What was his role and how did he support her in learning? Some of the interaction is verbal:

Christian kept saying that it was important that I put weight into my heels... [Höök 2010]

However, purely verbal instruction often fails:

... I could not understand what he meant. I tried making my heel the lowest point of my body rather than my toes, but I was still not getting it. [Höök 2010]

Simply telling can fail because understanding what the instructor means relies on correctly interpreting what is said in terms of an embodied experience which cannot directly be put into words. At this stage a different type of intervention is needed:

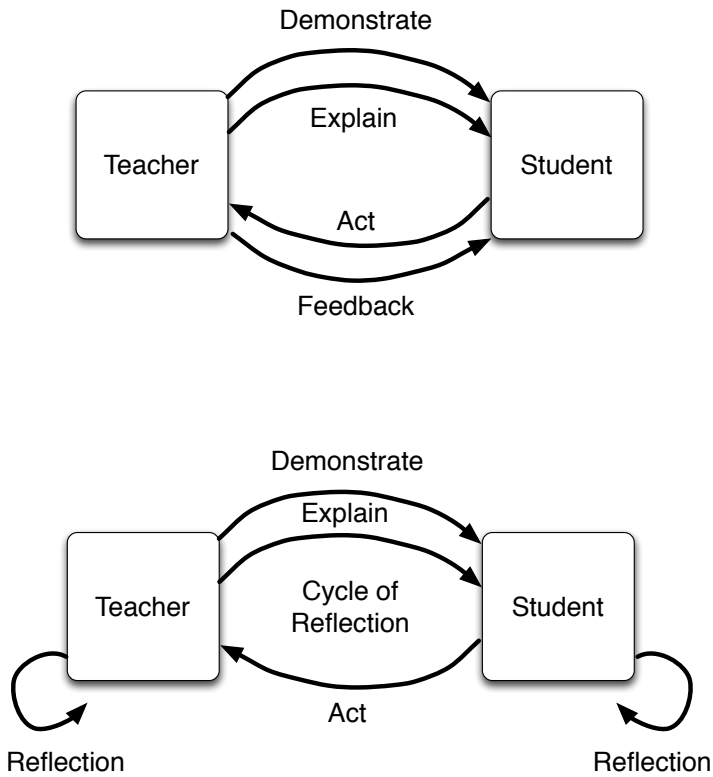


Fig. 1. Schön’s model of a reflective practicum. At first sight it appears to be a simple model of demonstration, action and feedback (top) but that tacit nature of the knowledge complicates this, because any stage of the interaction can lead to failures in understanding which require reflection and repair (bottom).

Christian decided to help me understand where my weight should be through a special exercise. He asked me to stand up in the saddle ... Suddenly I realized what Christian had meant by putting weight into my heels. My whole weight had moved down to my heels and it was from there that I found my balance. [Höök 2010]

This example shows that that teaching and learning an embodied skill is a complex interaction. The student acts (rides a horse) and the teacher diagnoses a problem with the action. The teacher intervenes verbally with a suggestion but the communication fails because it relies on an embodied experience that the student does not have. The teacher and student must therefore now diagnose a problem of communication and attempt to rectify it using an intervention based on action, not language. The interaction is a complex cycle of action, reflection and feedback in which teacher and student must try to come to a common understanding of concepts that neither of them can fully articulate.

Donald Schön has studied this type of teacher student interaction in his work on professional (though generally not movement) practice [Schön 1983, 1987]. He noticed that while there are theoretical/scientific elements of professional practice that can be taught explicitly (biology to doctors, physics to engineers) there are important professional skills of action that are tacit and cannot be taught in this way (diagnosis for doctors, design for engineers). This has a strong analogy to the embodied skills we are interested in.

Schön proposes that this tacit knowledge can be taught via what he calls a “Reflective Practicum”, a concept he developed based on his observation of architecture studio teaching (and which he then demonstrates in several other fields). A reflective practicum is one to one coaching where a teacher and student work together on a piece of the students independent work, giving feedback and suggesting improvements. Again, this is strongly analogous to the type of coaching we have seen in the somatic practices that Shusterman talks about [Shusterman 2008] or Höök’s example of horse riding [Höök 2010].

The basic process of a reflective practicum is shown in figure 1 (top). The teacher demonstrates a particular behaviour (Christian shows Höök how to sit in the saddle) or possibly explains it verbally (as we have seen this is likely to be difficult for many forms of embodied behaviour). The student then attempts to act either by directly copying the teacher or producing similar work of her/his own. The teacher observes this action and gives feedback either verbally or by demonstration (“*Christian kept saying that it was important that I put weight into my heels*” [Höök 2010]).

However, as we have seen in Höök’s example this process can fail because of the tacit nature of the knowledge involved. Without the tacit knowledge needed to understand the teacher’s demonstration or feedback, it can be impossible to decipher (“*I could not understand what he meant.*” [Höök 2010]). In the case of verbal explanation or feedback this happens because the teacher’s words refer to an embodied experience which cannot be fully explained in language. This means that students can only understand the comments fully via their own experience, but since they are learning they have not yet had that experience (or at least not fully understood it) and so they cannot understand what the teacher is saying. Demonstration may seem to get around this problem, but without personal experience the student does not know what aspects of the complex physical action to attend to and so to reproduce [Schön 1983].

Schön proposes that this is solved via a complex process of reflection, which as we saw is also key to Shusterman’s philosophy [Shusterman 2008]. Both teacher and student must reflect on their own actions and those of the other and attempt to understand what those actions mean to themselves and the other person. This reflection leads to a cycle of attempts to repair the failed understanding through further interaction: the student tries to act again or asks questions, the teacher adds further demonstrations and explanations. These attempts take many forms both verbal and non-verbal and a difficult problem might require the entire interaction to be reframed in order to achieve shared understanding, for example Christian, when his verbal instructions fail, attempts a new exercise that will lead to an experiential understanding (sitting in the saddle). This cycle of reflection and interaction, in difficult but successful cases, becomes what Schön calls the “*ladder of reflection*” [Schön 1987]. The teacher and student begin by reflecting directly on the students work, but in order to repair failed understanding they must reflect on their and the other’s understanding of that action, then on their communication of that understanding so that they can have a meta-conversation about the interaction, and then reflect in turn on this meta-conversation. These progressively higher levels of interaction happen in the context of a close cycle of interaction between teacher and students where each act (whether verbal or non-verbal) is an attempt to understand and communicate.

This process is illustrated in figure 1 (bottom). A teacher must communicate a way of behaving that cannot be explained directly in words. This leads to a situation where teacher and students do

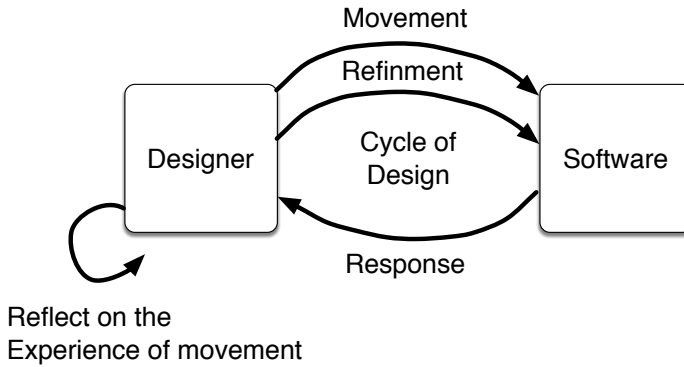


Fig. 2. The Schön’s reflective loop model adapted to movement design

not share an understanding and so cannot understand what the other is saying or doing. This must be repaired via a cycle of verbal and non-verbal interactions and personal reflection. The teacher demonstrates and/or explains and the student acts. Each must reflect on what happened before demonstrating, explaining or acting again, leading to an interactive cycle of mutual reflection. The main requirements of this process are a tight interaction cycle where both student and teacher attempt to communicate aspects of their understanding and the ability for each to reflect on that interaction to adjust their own understanding.

## 6 WHY IS INTERACTIVE MACHINE LEARNING SUITED TO MOVEMENT INTERACTION?

The previous sections have developed an understanding of movement knowledge as largely implicit and embodied but with explicit elements that can be thought about and talked about. It has also developed a model of how we communicate movement knowledge through a mix of demonstration, explanation and reflection, as exemplified by Schön’s reflective loop. What does this imply for the *design* of movement interaction?

Design and implementation of movement interaction involves movement knowledge in the sense of understanding a particular movement that is designed, and also communication of that knowledge, both to other people (designers and the public) and also, importantly to computers and related technological artifacts as the interaction is implemented. This requires a reflective loop similar to that prosed by Schön [Schön 1983, 1987], but it is no longer a loop between teacher and student but between designer and artifact (figure 2). While a designer might begin with an explicit idea for an interaction they can only fully understand it through physical practice and reflection on that practice. This becomes a loop because the designer can refine the interaction based on this reflection. However, because this explicit reflection is necessarily partial and incomplete, results of a change to the artifact can only truly be understood through another cycle of practice.

First person, embodied design processes such as Embodied Sketching [Márquez Segura et al. 2016] are valuable in this context as they allow designers to design in an embodied way through movement and reflection on that movement, and they also support the communication of movement knowledge with other designers. However, on their own they do not support implementation of

the movement interaction: communication of movement knowledge to a machine. For that we need appropriate implementation tools.

## 6.1 Implementation and Toolkits

Embodied Sketching is therefore an effective means of ideation that leverages movement knowledge. However, it is useless if the resulting designs cannot be implemented in practice. For traditional graphical interfaces this would rely on Graphical User Interface (GUI) toolkits that allow for quick prototyping and implementations of 2D interfaces.

However, if embodied knowledge is not representational then it is very different from the kind of representation we normally associate with GUI toolkits. Computers are generally based on explicit representations [Gillies and Kleinsmith 2014], and this is equally true of GUI toolkits with their symbolic buttons, sliders and menus. Designing interfaces generally involves creating representations, whether they are in the form of program code or graphical user interface layouts. If we are unable to make explicit representations of our embodied knowledge then this means that traditional styles of interface design toolkit will not work as they rely on representations. They will only be able to make use of a small set of bodily actions that we can describe explicitly (for example, raise your arms above your head), not more complex ones (e.g. the movements we make when we pedal a bicycle). So a designer who has developed an interaction using embodied sketching will find themselves severely limited when it comes to implementing that interaction if they have to rely on a symbolic or code based toolkit. All the benefits of embodied knowledge would become instantly lost.

It may be possible to implement a more sophisticated form of interaction with complex code. However, this is likely to be a slow process that does not support fast design iteration. The above discussion should make clear that fast iteration is even more important for movement interaction than it is for graphical interaction. Understanding the nature of the interaction that is being designed relies on the tight reflective loop of action and reflection based on that proposed by Schön. Without that iterative loop, designers would find themselves unable to fully understand the possibilities of interaction, let alone improve it.

## 6.2 Interactive Machine Learning

If we cannot create toolkits based on representations, how can we support designers of movement interaction? If embodied knowledge is implicit it means that we can *do* movements but cannot say *how* we do them. This implies that we should not only design by doing and moving [Hummels et al. 2007; Kleinsmith and Gillies 2013] but we should implement those designs in the same way. Implementing by doing means implementing by giving examples of actions. Designing by example goes back to early work on programming by demonstration [Hartmann et al. 2007; Lieberman 2001] and is in fact a popular method for designing gestural and movement interfaces, for example defined gestures by drawn templates [Wobbrock et al. 2007]. This style of design tool is mostly associated with machine learning which is used in many gesture recognition and movement interaction systems [Caramiaux et al. 2015a; Fiebrink et al. 2011]. Machine learning is a class of technology in which recognition systems are learned from a number of examples rather than being explicitly programmed. The focus on examples means that machine learning based design can support design based on embodied knowledge, rather than symbolic or linguistic knowledge, and the popularity of machine learning based toolkits for movement interaction supports this idea.

There is also a pleasing analogy to how people learn embodied skills. The idea is that people can only learn a skill like riding a bicycle through experience (c.f. learning from examples) rather than by being told how to do it (c.f. being explicitly programmed). This analogy captures the strong



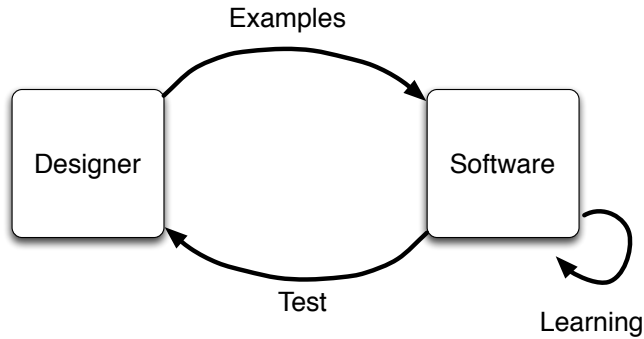


Fig. 3. The standard process of design via machine learning.

relationship between the idea of implicit, embodied knowledge and the process of design through machine learning.<sup>2</sup>

Machine learning makes it possible to design by giving examples of movement: designing by moving rather than having to put our tacit, embodied knowledge into explicitly linguistic or symbolic (i.e. program code) form. However, standard machine learning does not include the cycles of reflection and repair discussed in the previous section.

Let us compare Schön's model with the standard process of using machine learning to create a gesture recognizer, illustrated in figure 3. The designer's role is to provide a (normally large) number of examples of the gestures to be recognized. Once this data set has been collected they are given to the learning algorithm which learns a model from the data, which can be tested by the designer. If we are to translate this into human learning it corresponds to a teacher designing a large number of readings and exercise, giving them to the student and then leaving them to work through them on their own, learning independently. They may possibly return at the end to administer an exam. Not only does this not model the teaching process well, most teachers will realize that it is unlikely to succeed.

Schön's model calls for a much tighter loop of interaction, which is precisely what is found in Interactive Machine Learning. In IML, rather than a designer creating a large data set before hand, they would iteratively interact with the system to gradually guide learning with a number of example, in a workflow that is analogous to Schön's cycle of reflection. Figure 4 shows a proposed design process based on IML and Schön's model. The monolithic collect data - train - test sequence is replaced with a tighter cycle of design in which data is added interactively with the ability to test at any time and repair any problems through further interactions. These interactions can potentially take many forms including demonstration and examples but also communicating explicit aspects of the knowledge (corresponding to verbal explanations in human teaching).

The interactive machine learning process should promote reflection on the interaction between human and software. This reflection should not simply be, as is the case in most software, about learning what the system can recognize and adapting movements to the software. The designer is

<sup>2</sup>In fact, it is not a perfect analogy. We do not learn to ride a bicycle by viewing example, but by trial and error, a process much closer to reinforcement learning [Sutton and Barto 1998] than the types of learning discussed in this paper. A better analogy might be language learning or learning to paint.

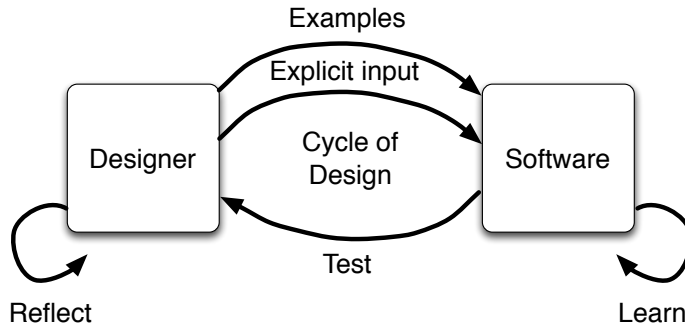


Fig. 4. A model of interaction with a movement interaction design tool based on Schön's reflective practicum.

changing the system through the design process to recognize more movements or recognize more accurately. This is certainly about reflecting on the system but it is also about reflecting on the movements made to the system because the system is defined by the movements it recognizes (in machine learning this is literally the case: the system is designed by giving movement examples). In some cases this reflection will be largely about the system: the sensors cannot distinguish two movements so the designer must design new movements. Some times the reflection will be purely about movement: during testing a design might decide that a movement quickly becomes tiring or uncomfortable (this case is very close to Shusterman's model of reflection). However, most of the time the reflection will be about the two together: a movement is not recognized because the designer is in fact performing it in a way that seems very similar to their original examples but is different in important ways. This requires reflecting on the details of the two movements and understanding how they are different, but also reflecting on the system and how it requires diverse examples (this model is very close to Schön's model of reflection).

## 7 EXAMPLES

This section will analyze three examples of using interactive machine learning to design movement interaction to see if they can enable the form of reflective design proposed in the previous section. Each examples presents results from user studies in which users were asked to perform interaction design tasks. In the majority of cases the participants in those studies had some experience of interaction design, even if their main expertise was in another domain. The studies all, therefore, give some idea of the role of a designer within the system.

### 7.1 The Wekinator

The Wekinator (<http://www.wekinator.org>) is a software platform developed by Fiebrink [Fiebrink 2011; Fiebrink et al. 2011] for designing gestural and movement based interfaces, primarily for performing electronic music (though it has been used in a range of other ways). The software uses and interactive machine learning process. Users provide examples of gestures or movements and can quickly test the results and provide more examples in order to correct or refine the system.

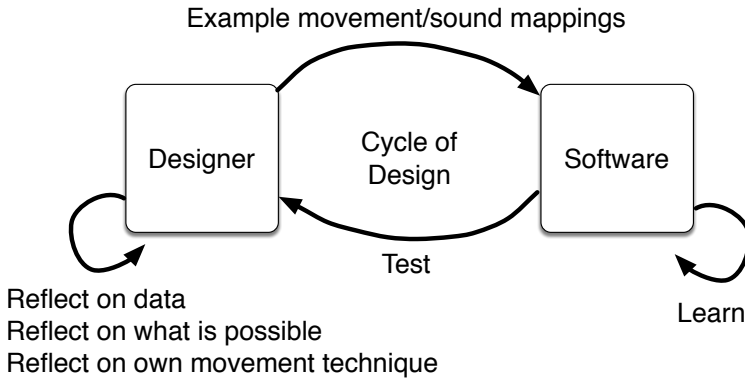


Fig. 5. The reflective design loop model applied to the wekinator

The discussion below is based on a number of user studies with a range of participants, including both students and professional musicians and designers. The majority had considerable experience of electronic or acoustic music and some experience of interaction design for musical applications.

The ability to design by example was popular with Fiebrink’s participants, who felt that it allowed them to design in more embodied ways than with explicit programming systems such as Max/MSP (which many were familiar with):

“With [theWekinator], it’s possible to create physical sound spaces where the connections between body and sound are the driving force behind the instrument design, and they \*feel\* right. ... it’s very difficult to create instruments that feel embodied with explicit mapping strategies, while the whole approach of [the Wekinator] ... is precisely to create instruments that feel embodied.” participant quoted in [Fiebrink 2011]

Fiebrink *et al.*’s participants interacted extensively with the learning process and preferred interactive testing to more traditional machine learning metrics like cross validation [Fiebrink *et al.* 2011]. This interactive testing served to improve the learned models but it also encouraged users to adapt their own behaviour and learn themselves. Fiebrink *et al.* [Fiebrink *et al.* 2011] list a number of ways in which their participants learned through interaction with the Wekinator.

**Teaching Users to Provide Better Data.** “ten students indicated that they had learned during their interaction with the software to provide training data that more clearly expressed their intentions” [Fiebrink *et al.* 2011]. This type of learning corresponds closely to Schön’s reflective cycle [Schön 1983, 1987]: users are learning to communicate better with the system.

**Teaching Users What is Possible** “Users often adapted their goals for the system based on what they discovered through direct evaluation ... When discovering that their efforts were failing to produce a model that worked how they wanted ... users adjusted their goals. Another reason for adaptation was that, through hands-on experimentation with the system, users discovered that models performed in unexpected ways that they actually liked better than their initial goals” [Fiebrink *et al.* 2011]. This form of learning is different from that described by Schön and shows that as well as reflecting on the interaction, interactive machine learning can also support reflection on the design.

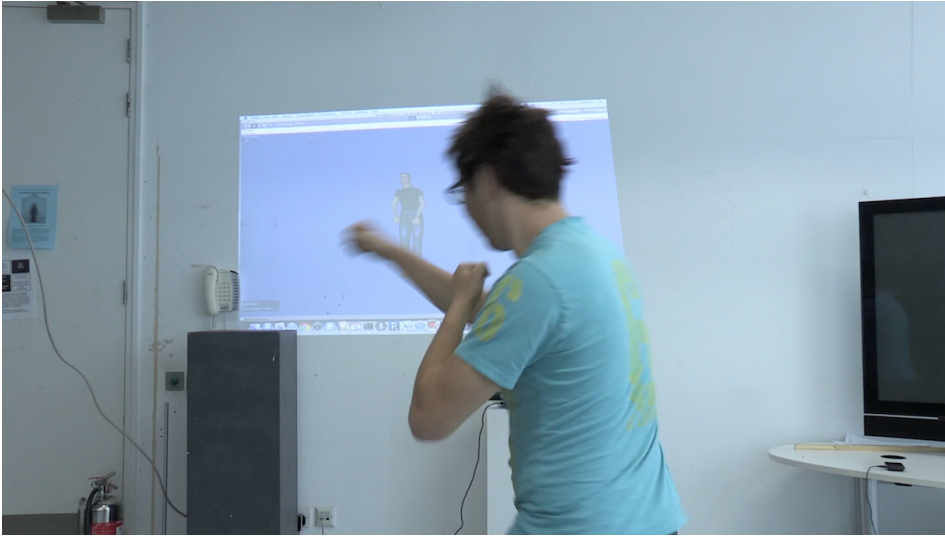


Fig. 6. Body language interaction with a virtual character

**Providing Feedback on Users’ Gestural Techniques** *“the cellist ... noticing that the bowing articulation model was not discriminating well between riccocet and spiccato strokes, she reexamined her own technique for those strokes and discovered that her spiccato technique actually needed to be improved”* [Fiebrink et al. 2011]. Again this shows that interactive machine learning can support the kind of reflection on movement technique and embodied knowledge that Shusterman [Shusterman 2008] advocates.

Figure 5 shows the reflective loop model applied to the Wekinator. The design loop is relatively simple, consisting of the designer providing examples and then testing the result. However, this interaction supports reflection on a number of different factors.

The example of the Wekinator therefore shows that interactive machine learning can support design based on embodied movement knowledge and also reflection on that design process. This reflection is not only on the communication (as proposed by Schön [Schön 1983, 1987]) but also on the design itself and the users own embodied knowledge and technique (as proposed by Shusterman [Shusterman 2008]). The reflective design process is therefore multi-faceted and allows for reflection on multiple different levels.

However, the problem of communication still remains. The Wekinator must attempt to learn what the designer intends and the designer must attempt to understand what the Wekinator has learned in order to correct it, without either being able to explicitly explain to the other. While we have established that entirely explicit communication is not possible, most of the teaching of movement practice does involve some form of explicit explanation as well as demonstration. The next two examples explore whether some form of explicit representation can help support the design process.

## 7.2 Body language based gameplay

The second example is work by Gillies *et al.* [Gillies et al. 2015a,b] on a platform for designing body language interaction with a virtual character in a game-like environment. A player (figure 6) can see a virtual character on screen. They can perform actions with their body that are tracked




Fig. 7. Visual feedback to support debugging in the body language based gameplay example. Each stick figure represents an item in the original data set. The colors of the figures represent the classes they were labeled with. The size represents how much they contributed to the classification. Designers are therefore able to see that, in this instance, the current pose is being classified as the blue class based primarily on their similarity to the two figures on the far right.

using a Microsoft Kinect™ and recognized by a machine learning based classification algorithm. The character will respond to the player's actions with movements of its own.

The purpose of the design tools was to enable people to design the set of actions from the player that should be recognized, the responses by the character and train the machine learning classifier to recognize the actions. The system was specifically designed to be used by actors since they are experts in creating characters, and part of the design process was a series of participatory design sessions with a professional actor, who had limited experience of interaction design but considerable experience of “designing” characters in a theatrical context.

These sessions showed that the actor could easily understand the machine learning process of providing example data and classifying it. However, problems emerged when the system did not classify an action correctly. The actor could interactively test the system and see there was a problem, but did not know how to fix it:

[the actor was] not able to clearly identify specific causes of problems (“about the clips I’m still not clear what ...the problem [is]”) and therefore any proposed solutions were vague and often inappropriate. [Gillies et al. 2015b]

This situation seems clearly analogous to the failures in communication described by Schön [Schön 1983, 1987]: the actor sees there is a problem but has no means to understand the actions of the system. What was needed was a way to enable him to reflect on the interaction and formulate a suitable correction. However, in this case simply iterating was not sufficient as the problem was too obscure to understand and any attempts at repair would have been essentially random.

The solution was to provide visual feedback to help users understand how the machine learning algorithm was making its decisions (figure 7). However this was not possible without changing the learning algorithm used. The first prototype used a decision tree algorithm whose classifications were based on a series of thresholding operations on joint rotation values. However, these rotation values, for example the y-orientation of the left shoulder, are not what we naturally think about when we move our bodies. Our sense of our bodies is much more holistic and we are more likely to think of the shoulder, clavical and upper arm (and maybe more of the body) as a single unit. For that reason, following a heuristic analysis using Blandford's CASSM framework [Blandford et al. 2008], Gillies et al. redesigned the system to use a machine learning algorithm that worked on holistic poses (a soft nearest neighbour based on Gaussian Mixture Models). The result is shown in figure 7: users are shown poses from the original training data which are scaled to show how much they contribute to the classification, so they can conceptualize the action of the system in

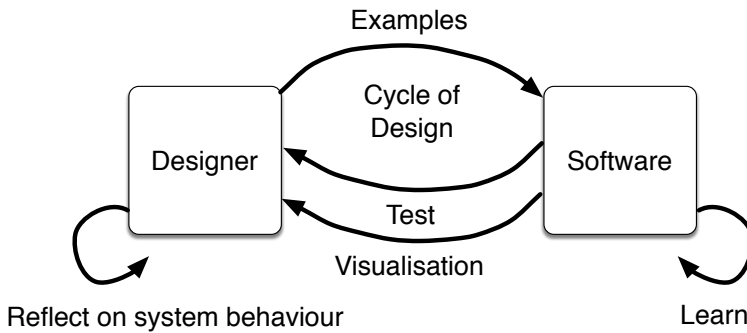


Fig. 8. The reflective design loop model applied to the body language interaction

terms of whole body poses not the fine detail of individual joints. This visualization aimed to allow users to more easily reflect on the behavior of the system and so debug problems, and their results seemed to suggest this was effective [Gillies et al. 2015a]. Figure 8 shows how the reflective loop model applies. The visualization provides an extra loop from the software designer that supports reflection on the internal behavior of the system.

An important point about the visualization used is that it is graphical not verbal, it allows users to tap their non-verbal knowledge of body movement. It is also real time: users can see how the system responds to their movements as they are doing those movements. This has the potential to set up a feedback loop between users bodily experience of moving and state of the system. As I move my arm I immediately see how the system interprets it and so I can learn to relate my action to the system's interpretation. This loop has the potential to support reflection on their bodily actions (though this was not investigated by Gillies *et al.*).

This example shows that interaction is not necessarily enough to enable reflection and that reflection is vital when things go wrong and need debugging. Effective reflection on the behavior of software system requires multiple forms of feedback, both implicit in terms of actions by the system and explicit in terms of (in this case visual) representations of the state of the system. However, these representations can still be ineffective if the way the system represents movement is too different from how people think about and experience movement. This would require the designer to bridge a large conceptual gap between two very different means of representation. Machine learning is most effective if it is not a black box but it represents movement in ways that are in some way analogous to how we as humans experience it.

### 7.3 Gesture Interaction Designer

The previous two examples show the importance of interaction and the combination of implicit and explicit feedback from the toolkit when designing movement based interaction. However, in both cases all of the feedback from users to the system is implicit: in terms of labeled examples. Can users be allowed to also give explicit feedback to the system?

In human-to-human teaching explicit feedback is generally verbal and in language chosen by the teacher, leading to the risk that the student will not understand. In human-computer interaction this form of verbal feedback is far too challenging: the designer must give feedback in terms defined



Fig. 9. Visual feedback in the GIDE system (top). A video of the original recorded image (1) is shown on the left hand side of the view for each gesture, together with a visualization of the waveform (2). At the top right of the screen there are controllers to adjust the parameters (3). (bottom right) The current most likely gesture is highlighted in green (4) while the others are darkened (5). The current positions is shown by a horizontal red bar on the waveform view (6). (bottom left) The effect of tolerance parameter is shown by the width of the green line (7) which demonstrates how likely the system is to detect divergence from the expected signal.

by the system, for example tuning parameters. However, this leads to the risk that the designer will not understand how to give effective feedback.

Zamborlin *et al.*'s Gesture Interaction DEsigner (GIDE) [Zamborlin *et al.* 2014] is a platform for designing gestural interfaces, primarily for musical and audio-visual performance. Users are able to record a single example for gesture to be recognized and then test by performing more gestures. The system uses Bevilacqua *et al.*'s Gesture Follower [Bevilacqua *et al.* 2010] recognition algorithm, which enables it to recognize gestures continuously while they are being performed and not simply at the end of the gesture. It is also able to estimate how far through the gesture the user is at any given time. This allows GIDE to give detailed visual feedback about the gesture that the user is currently performing as shown in figure 9. As in the previous example the feedback is real time so users can relate their bodily experience of moving to the state of the system.

As with the Wekinator studies, participants were all experienced electronic musicians with many having considerable interaction design experience. They were all professional musicians or graduate students in a subject related to electronic music or interaction design. Participants in Zamborlin *et al.*'s trials [Zamborlin *et al.* 2014] found this visual feedback useful in testing and debugging their gesture systems.

If the recognizer was not performing correctly users could record another example, but they could also tune a number of parameters: *tolerance* affected how different new gestures were allowed to be from the original; *latency* was a window size for recognition, high values resulted in stable classifications, high values allowed rapid changes based on new input, and *contrast* a relative

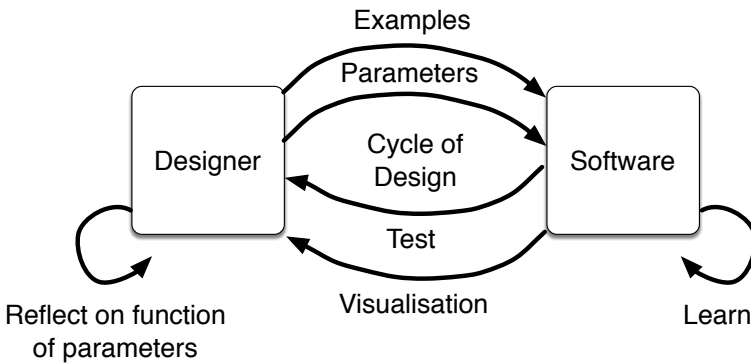


Fig. 10. The reflective design loop model applied to Gesture Interaction Designer

weighting on the probabilities of gestures that emphasized the best performing gestures. Tolerance and contrast had corresponding visualization (see figure 9 for the visualization of tolerance). Each of these parameters is relatively technical and related to detailed functioning of the algorithm, nonetheless the study participants were able to use tolerance and contrast effectively, though not latency. The qualitative results of the study seem to show that the visual feedback was vital in helping participants understand how to use the parameters. For example:

By increasing [contrast] they could see the likeliest gesture more clearly referring to it as the green gesture, pointing out that the association between the likelihood and this specific color was pretty clear. On the other hand, when colors started flickering too much, they knew quickly that it meant that it was a good idea to decrease the value of the contrast.

On the other hand latency did not have a clear visual analogue in the interface and the authors concluded that this was the reason that it was less used. As in the previous example in which visual feedback helped users make sense of the behavior of the system, visual feedback is also able to support users in effectively given explicit feedback to the system. The results is a complex two way interaction between designer and interaction toolkit that combines both implicit and explicit interaction to build a shared understanding of the movement interaction to be designed, exactly as we proposed in our model, as shown in figure 10.

## 8 DISCUSSION

The previous section describes three examples of the use of interactive machine learning in movement interaction design. They show the importance of being able to design by moving and how interactive machine learning enables this. For example, the participants using the Wekinator were enthusiastic about being able to design musical instruments in an embodied way. However, these examples also showed the importance of being able to bring some aspects of the interaction design into an explicit form. These explicit representations can help designers reflect on their work and improve their designs.



Beyond these three examples how applicable is interactive machine learning? What are its current limitations? What future research would enable interactive machine learning to be applied to a wider range of movement interaction design tasks? This section will address these questions.

### 8.1 Applicability of Interactive Machine Learning

Section 2 discussed a number of application techniques and interaction styles used in movement interaction, and these will help us understand the applicability of interactive machine learning. We defined three major styles of interaction: object focused, direct mapping and movement focused. As discussed in section 2, interactive machine learning is primarily applicable to movement focused interaction in which the system is required to recognize particular specific movements.

If we turn to the various application areas we've discussed, each sits somewhat differently in terms of these styles of interaction. For example sports and exertion games often rely primarily on direct mapping. The players must simply move as appropriate and scores often relies simply on detecting whether the player enters a particular area or not. However when there are cases in which it is important to judge exact movements movement then interactive machine learning can play a part and this may be particularly true of training in particular sports movements, such as the work of Velloso *et al.* [Velloso *et al.* 2013].

In the domain of New Interfaces for Musical Expression (NIME) Interactive Machine Learning is very applicable because instruments tend to require very specific movements and very detailed recognition of these movements and in fact many of the applications cited in this paper do fit in domain of NIME [Caramiaux and Tanaka 2013; Fiebrink *et al.* 2011; Françoise *et al.* 2013]. A related domain is dance and theater, where physical performer's movements can interact with projections or other aspects of scenography. There have been a number of uses of IML in this context [Fdili Alaoui *et al.* 2015; Katan *et al.* 2015].

Medical and health interaction often relies on the need to recognize specific movements for example whether it is the specific rehabilitation exercise as described by Kirk *et al.* [Kirk *et al.* 2016]. Often these must be personalized to the needs of a particular patient who may have more or less movement capabilities. These are good examples of where interactive machine learning can be very valuable. The need to recognize specific movements which may be hard to put in an explicit form and in particular the need to personalize requires a light weight method of changing the recognition method.

However, the medical domain also points to another type of movement interaction which has less need for IML. Many exercises where there is a very well defined measure of the quality of a movement, for example the amount of weight lifted or the angle of the shoulder. In these cases, an explicitly programmed solution is likely to be more reliable than machine learning.

Section 2 also discussed types of movement interaction that are more implicit, in that users are not consciously interacting with a system in a action-response dialog, but the system responds to the ongoing activity of the user. IML has been little used for this type of interaction, so it is hard to say how applicable it is. Much of the argument of this paper does apply to this case, for example, it relies on embodied human movements, and therefore on embodied knowledge of those activities that can be difficult to communicate explicitly. So IML could be very applicable as part of the design process. However, as this type of interaction will involve spontaneous, rather than explicitly designed movements, it may be that traditional batch machine learning approaches will work better as they can capture a more representative sample of natural movement.

Different applications of movement interaction are likely to use different types of sensor. The examples described in this paper cover a wide range from depth cameras [Gillies *et al.* 2015a] to accelerometers [Zamborlin *et al.* 2014] and custom interaction devices [Fiebrink *et al.* 2011]. Other papers cited have used many other sensors. One of the key benefits of IML and Machine Learning

in general is that it is relatively agnostic to sensing technologies. As long as the sensed data can be put in the form of a vector of numerical values then most machine learning algorithms can be applied to it, so it is unlikely that this approach will be limited by sensing technology. The Wekinator is a case in point, the software is designed to be independent of sensing technology and new sensors can easily be used with it by sending data via the OSC protocol. Many of the examples that Fiebrink discusses [Fiebrink 2011; Fiebrink et al. 2011] involve users linking the wekinator software to new sensing devices as part of the interaction design process. However, this is not to exclude the possibility that data from a particular sensor would be particularly hard for a machine learning algorithm to interpret, for example, if it is very noisy.

This section, and this paper more generally, has discussed a number of applications of movement interaction that have used Interactive Machine Learning, and that broadly follow the model presented in this paper. Much of this work consists of research prototypes and is relatively preliminary. These examples support the use of IML in these domain, and the ideas presented in this paper reinforces their usefulness, but it is too early to draw definitive conclusions about the long term success of IML (or the ideas presented here), but it can be expected that as these approaches are applied in practice, new difficulties and problems will emerge.

The application areas presented are only a subset of the possible uses of movement interaction, and movement interaction is likely to expand in the coming years, particular with the introduction of new approaches to computing such as Virtual and Augmented Reality, physical computing or the quantified self. For example, the Body Language Based gameplay example in this paper shows the potential of using IML for designing interactive characters in Virtual Reality. While it is hard to predict how these new movement interactions will be designed, the ideas presented in this paper suggest that a combination of embodied design approaches and Interaction Machine Learning will be valuable under these three conditions:

- The interaction is movement focused, i.e. the focus is designing a type of movement. Other types of interaction, like direct mapping are likely to require less sophisticated techniques
- The movements rely on embodied movement knowledge that cannot fully be communicated explicitly. There would be no need for IML in contexts where the movements are, for example, based on an established scientific theory, such as the example given above of medical exercises that can be assessed with well established quantitative measures.
- Designers are able to design movement interaction by performing, and reflecting on examples of movements. This will apply in most cases, but may be problematic in cases where the movement capabilities of designers is different from those of users (for example users with disabilities, or users that are highly trained athletes) or where designers experience of movement is different from users (for example if designers are more physically fit, or younger than the average user). In many cases these problems could be resolved by using the approaches described in this paper in a participatory design setting.

## 8.2 End-to-end IML

One drawback of a lot of interactive machine learning systems is that they often focus only on the recognition of movement. Most machine learning algorithms are classifiers which will recognize particular movements or possibly regression models which map a movement to a continuous number. For example, their output might be a label such a "running" or "jumping" or a continuous value indicating a feature such as how "angry" or "light" a movement is. These outputs are typically need to be turned into outputs.

Recognition of movement is a vital part of movement interaction design, however it is not the whole story. Movement interaction requires a movement but it also requires a response from the

system. The outputs typically need to be turned into new forms, whether that is graphics in a VR environment, sound in a musical performance interface or even physical changes in a shape changing interface [Rasmussen et al. 2012]. This output interaction design of the movement-response loop is often not handled in machine learning applications. The mapping of machine learning outputs to responses is typically handled by standard programming approaches. However, this does seriously undermine the embodied design process, which should design the entire interaction, not simply the recognition aspect. What is needed is an end-to-end embodied interaction design process which goes from user sensing to system response. This requires end-to-end interactive machine learning tools, in which movement is not simply mapped to a label or number but it is mapped directly to system responses such as graphics, sound or shape changes. This is an under-researched area, but examples such as the work of of Françoise *et al.* [Jules Françoise 2012], is likely to be increasingly valuable.

### 8.3 Implicit and Explicit

An important theme of this paper is that movement knowledge is, in large part, implicit and non-verbal. We know how to move in an embodied way, through moving, not in an explicit symbolic way. Machine learning offers great promise because it allows us to design interactions by giving examples of movement. That means we can design by moving not by putting movement into verbal or symbolic form.

However another theme of this paper is that, while movement interaction is largely non-verbal, words and explicit representations do play an important part in how we learn and communicate movements skills. Höök's [Höök 2010] example of learning horseback riding with a teacher shows their struggle to communicate verbally things that are understood only in an embodied way, but also it shows the importance of this struggle. Learning could not have happened if Höök and her teacher were not able to communicate their experience linguistically. This very much connects with Schön's model [Schön 1983] in which student and teacher must struggle to acquire a common language in which to talk about their non-verbal experiences.

Treating embodied knowledge as something that cannot be accessed directly and only through examples of action (treating the learning algorithm as a "black box") is therefore missing a lot. It is likely that some information can be put into a symbolic form that can be read by a computer, for example Silang Maranan *et al.* [Silang Maranan et al. 2014], in their machine learning system, use examples of dance which are created using a dancers embodied, but both the dancers movements and the computer system are structured by the explicit, symbolic Laban Movement Notation. More importantly, much of the knowledge that cannot be put into words or symbolic form can still be made conscious and reflected on. This implies that tools for designing movement interaction should make support and make use of this form of non-verbal reflection

### 8.4 Opening the black box

So explicit representations must play a part in movement interaction design. These could be verbal, linguistic tags representing certain aspects of the interaction, but they may also be other forms of representation, for example visualizations can play an important part in how we do interactive machine learning as demonstrated by Gillies [Gillies et al. 2015a]. These forms of visualization make it possible to open the "black box" of machine learning and allow users to understand how it functions. However, this poses considerable challenges, since the workings of machine learning algorithms are often very difficult to interpret. A typical deep neural network could have thousands of parameters, so visualization would be very difficult to interpret. This is an important drawback in the use of machine learning in design. Ensuring that models can be interpreted puts an important

constraint on the algorithms that can be used. If the workings of a learning algorithm need to be explicitly represented, the algorithm needs to be of a form that can be interpreted.

There has been considerable research in recent years about interpretable machine learning. One approach is to use learning algorithms that are readily interpretable to humans [Caruana et al. 2015; Letham et al. 2015; Wang and Rudin 2015]. For example, Wang and Rudin [Wang and Rudin 2015] have developed a machine learning system for medical diagnosis, and use an algorithm that generates check lists of a form similar to those used in hospitals. An example from movement interaction is Gillies *et al.* [Gillies et al. 2015b], described above, who propose using an interpretable nearest neighbor algorithm over potentially more powerful, but less interpretable algorithms. This approach has the benefit that the workings of the learned model are easy to understand and normally to visualize as is the case of Gillies *et al.* However, it does constrain the type of model learned, which may not be as effective as, for example, a complex, but uninterpretable, neural network. The assumption, particularly in an interactive setting, is that having designers being more able to effectively guide the learning will make up for the learning algorithm itself being less powerful.

An alternative approach is to use a standard, complex learning algorithm, and, instead of attempting to explain the workings of the algorithm as a whole, to explain individual decisions [Baehrens et al. 2009; Kulesza et al. 2015; Patel et al. 2010; Ribeiro et al. 2016]. For example, authors, such as Baehrens *et al.* [Baehrens et al. 2009] and Ribeiro *et al.* [Ribeiro et al. 2016] explain individual decisions by fitting a local intelligible model. This means that they take a particular decision and present users with a model, which is much simpler than the original learned model, but approximates it well, if only for data similar to the one the decision is being made on. This makes it possible for users to understand the behavior of far more complex learning algorithms, but it does so at the cost of only understanding individual decisions, not of understanding the working of the system as a whole.

## 8.5 Emergent representations

One important point to note is that the representations we use in movement learning are often emergent<sup>3</sup>. In dialog, these representations are typically words or phrases we use to describe movement sensations for which we have no pre-defined vocabulary (or we do not know that vocabulary), but in design they can also include visual representations and diagrams. Rather than a predefined and shared vocabulary, our ways of talking about movement experience emerge through interaction. As both Schön and Höök have discussed teacher and student must negotiate to develop a common language in which to discuss their experiences. This is particularly true in the case of interaction design when we are often designing new embodied experiences. While teacher and student may be able to fall back on established vocabulary of their domain, when we are designing something new we must be able to create new words and new representations with which to think about those interactions. That means that the representations used in interactive machine learning should not be a fixed set of words or a single visualization. They should be interactive and supports users in adding new aspects to any representations and creating their own variants of all representations. So, if we follow the argument of this paper, a key future area of research in interactive machine learning is not simply visual representations or verbal representations for the behavior of interactive machine learning system but ways in which users can manipulate, customize and appropriate [Dix 2007] representations for new emergent needs. This itself raises important questions. For example, how easily can representations move from one domain or other? Are we able to kick start the representation building with a set of existing representations, or

---

<sup>3</sup>Thank you to an anonymous reviewer for raising this important point

should the design process always start nonverbally and move onto a representational stage as our understanding of the interaction matures?

## 8.6 Beyond Supervised Learning

The relationship between implicit and explicit representations echoes another important development in interactive machine learning research, which is thinking of interaction with machine learning as more than simply providing example data.

Most interactive machine learning approaches are based on supervised learning that is learning from examples which include labels of the correct output of the system. However some researchers are looking at other aspects of the problem. For example before supervised learning can be performed the data must be transformed into a number of features that are suitable for learning. A feature is a numerical representation of the data and certain features are more informative than others, for example edges in an image may be more informative than raw pixel values. This feature design is often left out of interactive machine learning or done explicitly such as the work of Fails and Olson [Fails and Olsen 2003]. However including humans in the process of designing features could be an important way of leveraging human movement knowledge.

Other researchers have used different approaches to interactive machine learning for example Thomaz and Breazeal [Thomaz and Breazeal 2008] use reinforcement learning, where the human feedback is judgments of whether the machine's behavior is correct or incorrect rather than examples of correct behavior. And finally other approaches may give more explicit information about how a system should learn. For example Talbot *et al.* [Talbot *et al.* 2009] allows users to combine simple classifiers into different ensembles and Zamborlin *et al.* [Zamborlin *et al.* 2014] system supports adjusting parameters of the machine learning algorithms.

All of these factors have the potential to allow a richer interaction with machine learning algorithms that allow the kind of interplay of explicit and implicit knowledge that characterizes movement. However many of the examples are primarily focused on the needs of machine learning, for example, adjusting algorithm parameters and we need to develop methods that are focused on the knowledge of movement interaction designers and how interaction with machine learning fits within their design process.

## 8.7 Reflection

Another important theme of this paper is that the ability to reflect on one's movements is a key part of movement learning and movement knowledge and therefore must be a key part of movement interaction design. Without the ability to reflect consciously it is unlikely to be possible to refine and improve interaction designs and therefore this form of reflection is also a key part of the design process.

The use of explicit representations are extremely important in supporting reflection. One of the key challenges of much machine learning is that the models generated can be highly uninterpretable. They are complex combinations of parameters and the reasons for the decisions of a machine learner can be highly obscure to a human user even a machine learning expert. This lack of interpretability is an obstacle to reflection since it prevents a designer from understanding what is happening in their interaction design. Visualization and explicit representations can help make machine learning more interpretable and so are also important as a key way of enabling reflection.

Interacting with machine learning is itself also key to reflection, as shown in Fiebrink's work on the Wekinator [Fiebrink *et al.* 2011] discussed above. The tight loop of giving examples and testing the results enables designers to understand the impact of their actions in the kind of reflective loop that we see in Shön's model. The combination of representation of machine learning and a tight

loop of interaction is therefore a way of fostering the forms of reflection that are key to movement knowledge.

## 8.8 Evaluation

Reflection is also closely related to the issue of evaluating machine learning: part of reflection is in evaluation of how well the system is working. Evaluation of interactive machine learning is often a challenge because there are very different evaluation cultures in the domains of machine learning and of human computer interaction that form the twin intellectual poles of IML. Machine Learning is generally evaluated in a very quantitative way using clearly defined fixed measures. However these often do not capture the the nuance of human behavior. For example, Fiebrink [Fiebrink et al. 2011] has found that the users of her work IML system use very different approaches to evaluation than would normally be used by a machine learning engineer. Having said, that the metrics used in machine learning can be a valuable source of information about how well a system is recognizing a particular behavior and therefore it is very important to find methods of evaluation that can bridge the gap between machine learning and movement knowledge [Boukhelifa et al. 2018].

## 8.9 Toolkits

This paper has compared interactive machine learning to the graphical user interface toolkits used in traditional interaction design. GUI toolkits are a vital part of how we design interaction because they allow us to prototype graphical interfaces quickly and to use standardized components that may be familiar to users. However the fact that they rely entirely on symbolic representations means that they do not carry over well to the embodied nature of movement interaction. If movement interaction is to become successful and mainstream, we do need the equivalent of toolkits which will allow designers to rapidly prototype and implement movement interaction designs.

IML promises to be an answer to this. It is a generic approach to movement interaction that can be implemented into relatively standardized software to be used by designers to prototype and implement a wide range of behaviors. The rapid interaction and design process supports the needs of prototyping well and the ability of machine learning algorithms to work equally well with diverse sensor inputs and diverse outputs, has the potential to make toolkits entirely generic.

Having said that there is a lot of work still to be done. Toolkits are still relatively underdeveloped they're often very domain specific and generic ones are often aimed at engineers rather than then designers [Patel et al. 2010]. In many ways this is an engineering challenge that needs to be solved as interactive machine learning moves from the research lab into into professional use. However, there is also a lot of interaction design work to be done to ensure that interactive machine learning toolkits are as usable as they need to be for designers. It is also important to consider how IML might integrate with other forms of toolkit, for example, hardware toolkits such as Arduino (<https://www.arduino.cc>). IML is primarily a software framework but it is vital that it integrates effectively with the hardware that supports sensors and actuators.

## 8.10 Skills and Education

Even with usable toolkits, machine learning in interaction design is still a very unfamiliar way of working and as Dove and colleagues [Dove et al. 2017] have shown, though interaction designers are increasingly using machine learning, they often lack the education and expertise to use it effectively. What we have discussed here is a very different style of designing than we might use in a traditional interface and both movement interaction itself and the use of machine learning require the require new skills and new forms of education. It therefore seems that in the next few years it will be a vitally important to improve education in the use of machine learning for those

outside of the circles of machine learning engineering experts and new ways of teaching how to use machine learning that are applicable to designers are likely to emerge.

### 8.11 End user design and appropriation

This goes further than simply educating designers. Though it has not been the focus of this paper, Interactive Machine Learning can also open up the possibilities of end users continuing to adapt interactions designs. While many designers will chose to simply ship the results of an interactive machine learning process, leaving the IML code in a product opens up the possibility of end users redesigning it. This has many benefits, including the ability to adapt technology to a users particular physical and other needs, a very important enabler of accessible technology. It is also a good example of design for appropriation [Dix 2007; Dourish 2003], in which systems are specifically designed so that end users can “appropriate” them, which means adapting for users that were not envisaged by the original designer (a prosaic example being using a shoe to hammer a nail). Another important future direction for IML is, therefore, ensuring that is is also usable by end users.

## 9 CONCLUSION

Interactive Machine Learning is a popular approach to building toolkits for movement interaction, that has the potential to make the design of movement interaction as ubiquitous as the design of 2D Graphical User Interfaces. This paper has studied the nature of human movement knowledge and how we teach movement knowledge in order to understand what makes a good tool for designing movement interaction. This understanding highlights some of the important features that a movement interaction design system needs:

**Designing by moving.** The system should learn from examples of movement to avoid making users describe movement entirely explicitly

**A tight cycle of interaction** to correct errors and ensure that system’s behavior corresponds to the designer’s understanding

**Support reflection on the experience of movement.** This reflection is both on the actions of the system but also on the designers understanding of movement and can be supported by the tight cycle of interaction and by explicit verbal or graphical representation of the state of the system.

These features can help us understand why Interactive Machine Learning has often been used for movement interaction design, which situations it has been successful and what future research is needed for it to work well in other situations.

We hope that this analysis can help guide researchers in developing future interactive machine learning toolkits that can help people design better, more embodied, movement interfaces.

## REFERENCES

- Saleema Amershi, Maya Cakmak, William Bradley Knox, and Todd Kulesza. 2014. Power to the People: The Role of Humans in Interactive Machine Learning. (12 2014), 105–120 pages. <https://doi.org/10.1609/aimag.v35i4.2513>
- Saleema Amershi, Bongshin Lee, Ashish Kapoor, Ratul Mahajan, and Blaine Christian. 2011. Human-guided machine learning for fast and accurate network alarm triage. *IJCAI International Joint Conference on Artificial Intelligence* (2011), 2564–2569. <https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-427>
- David Baehrens, Timon Schroeter, Stefan Harmeling, Motoaki Kawanabe, Katja Hansen, and Klaus-Robert Mueller. 2009. How to Explain Individual Classification Decisions. (2009).
- Frédéric Bevilacqua, Bruno Zamborlin, Anthony Sypniewski, Norbert Schnell, F Guédy, and Nicolas Rasamimanana. 2010. Continuous realtime gesture following and recognition. *Gesture in Embodied Communication and Human-Computer Interaction* (2010), 73–84. <http://www.springerlink.com/index/V8PN885111256625.pdf>
- Nadia Bianchi-Berthouze, Whan Woong Kim, and Darshak Patel. 2007. Does Body Movement Engage You More in Digital Game Play? and Why?. In *Affective Computing and Intelligent Interaction, Second International Conference, ACII 2007*,

- Lisbon, Portugal, September 12-14, 2007, *Proceedings (Lecture Notes in Computer Science)*, Ana Paiva, Rui Prada, and Rosalind W Picard (Eds.), Vol. 4738. Springer, 102–113.
- Ann Blandford, Thomas R G Green, Dominic Furniss, and Stephann Makri. 2008. Evaluating system utility and conceptual fit using CASSM. *Int. J. Hum.-Comput. Stud.* 66, 6 (6 2008), 393–409. <https://doi.org/10.1016/j.ijhcs.2007.11.005>
- Nadia Boukhelifa, Anastasia Bezerianos, and Evelyne Lutton. 2018. Evaluation of Interactive Machine Learning Systems. *ArXiv* (2018), 1–20. <http://arxiv.org/abs/1801.07964>
- Baptiste Caramiaux, Alessandro Altavilla, Scott G. Pobiner, and Atau Tanaka. 2015a. Form Follows Sound. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*. ACM Press, New York, New York, USA, 3943–3952. <https://doi.org/10.1145/2702123.2702515>
- Baptiste Caramiaux, Marco Donnarumma, and Atau Tanaka. 2015b. Understanding Gesture Expressivity through Muscle Sensing. *ACM Transactions on Computer-Human Interaction* 21, 6 (1 2015), 1–26. <https://doi.org/10.1145/2687922>
- Baptiste Caramiaux and Atau Tanaka. 2013. Machine Learning of Musical Gestures. *Proceedings of the International Conference on New Interfaces for Musical Expression* (2013), 513–518. <http://nime2013.kaist.ac.kr/>
- Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. 2015. Intelligible Models for HealthCare. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '15*. 1721–1730. <https://doi.org/10.1145/2783258.2788613>
- Gabe Cohn, Dan Morris, Shwetak N Patel, and Desney S Tan. 2011. Your Noise is My Command : Sensing Gestures Using the Body as an Antenna. (2011), 791–800.
- Antonio Damasio. 1994. *Descartes'Error: Emotion, Reason, and the Human Brain*. 312 pages. [https://books.google.co.uk/books/about/Descartes\\_Error.html?id=SLdYHllhgKMC&redir\\_esc=y](https://books.google.co.uk/books/about/Descartes_Error.html?id=SLdYHllhgKMC&redir_esc=y)
- Alan Dix. 2007. Designing for appropriation. In *Proceedings of the 21st British HCI Group Annual Conference on People and Computers: HCI...but not as we know it - Volume 2 (BCS-HCI '07)*. British Computer Society, Swinton, UK, UK, 27–30. <http://dl.acm.org/citation.cfm?id=1531407.1531415>
- Paul Dourish. 2003. The Appropriation of Interactive Technologies: Some Lessons from Placeless Documents. *Comput. Supported Coop. Work* 12, 4 (9 2003), 465–490. <https://doi.org/10.1023/A:1026149119426>
- Graham Dove, Kim Halskov, Jodi Forlizzi, and John Zimmerman. 2017. UX Design Innovation: Challenges for Working with Machine Learning as a Design Material. In *CHI '17 Proceedings of the 2017 annual conference on Human factors in computing systems*. 278–288. <https://doi.org/10.1145/3025453.3025739>
- Jerry Alan Fails and Dan R. Olsen. 2003. Interactive machine learning. In *Proceedings of the 8th international conference on Intelligent user interfaces - IUI '03 (IUI '03)*. ACM Press, New York, New York, USA, 39. <https://doi.org/10.1145/604045.604056>
- Sarah Fdili Alaoui, Frédéric Bevilacqua, and Christian Jacquemin. 2015. Interactive Visuals as Metaphors for Dance Movement Qualities. *ACM Transactions on Interactive Intelligent Systems* 5, 3 (2015).
- Rebecca Fiebrink. 2011. *Real-time Human Interaction with Supervised Learning Algorithms for Music Composition and Performance*. Ph.D. Dissertation. Princeton University, Princeton, NJ, USA.
- Rebecca Fiebrink, Perry R. Cook, and Dan Trueman. 2011. *Human model evaluation in interactive supervised learning*. ACM Press, New York, New York, USA. 147 pages. <https://doi.org/10.1145/1978942.1978965>
- P M Fitts and M I Posner. 1967. *Human Performance*. Brooks/Cole Publishing Company. <https://books.google.co.uk/books?id=XtFOAAAAMAAJ>
- 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. <https://doi.org/10.1145/3025453.3025714>
- Jules Françoise, Norbert Schnell, and Frédéric Bevilacqua. 2013. A multimodal probabilistic model for gesture-based control of sound synthesis. *Proceedings of the 21st ACM international conference on Multimedia - MM '13* (2013), 705–708. <https://doi.org/10.1145/2502081.2502184>
- Marco Gillies. 2016. What is Movement Interaction in Virtual Reality for?. In *Proceedings of the 3rd International Symposium on Movement and Computing - MOCO '16*. ACM Press, New York, New York, USA, 1–4. <https://doi.org/10.1145/2948910.2948951>
- Marco Gillies, Harry Brenton, and Andrea Kleinsmith. 2015a. Embodied design of full bodied interaction with virtual humans. In *Proceedings of the 2nd International Workshop on Movement and Computing - MOCO '15*. ACM Press, New York, New York, USA, 1–8. <https://doi.org/10.1145/2790994.2790996>
- Marco Gillies and Andrea Kleinsmith. 2014. Non-representational Interaction Design. In *Contemporary Sensorimotor Theory*, John Mark Bishop and Andrew Martin (Eds.). Springer-Verlag, 201–208.
- Marco Gillies, Andrea Kleinsmith, and Harry Brenton. 2015b. Applying the CASSM Framework to Improving End User Debugging of Interactive Machine Learning. In *Intelligent User Interfaces*.
- Marco Gillies, Bongshin Lee, Nicolas D'Alessandro, Joëlle Tilmann, Todd Kulesza, Baptiste Caramiaux, Rebecca Fiebrink, Atau Tanaka, Jérémie Garcia, Frédéric Bevilacqua, Alexis Heloir, Fabrizio Nunnari, Wendy Mackay, and Saleema Amershi.



2016. Human-Centred Machine Learning. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems - CHI EA '16*, Vol. 07-12-May-. 3558–3565. <https://doi.org/10.1145/2851581.2856492>
- Megan Halpern, Jakob Tholander, Max Evjen, Stuart Davis, Andrew Ehrlich, Kyle Schustak, Eric P S Baumer, Geri Gay, and College Ave. 2011. MoBoogie : Creative Expression Through Whole Body Musical Interaction. (2011), 557–560.
- Perttu Hämäläinen, Joe Marshall, Raine Kajastila, Richard Byrne, and Florian "Floyd" Mueller. 2015. Utilizing Gravity in Movement-Based Games and Play. *Proceedings of the 2015 Annual Symposium on Computer-Human Interaction in Play - CHI PLAY '15* (2015), 67–77. <https://doi.org/10.1145/2793107.2793110>
- Björn Hartmann, Leith Abdulla, Manas Mittal, and Scott R Klemmer. 2007. Authoring sensor-based interactions by demonstration with direct manipulation and pattern recognition. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI '07)*. 145–154. <https://doi.org/10.1145/1240624.1240646>
- K. Hook. 2009. Affective loop experiences: designing for interactional embodiment. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364, 1535 (2009), 3585–3595. <https://doi.org/10.1098/rstb.2009.0202>
- Kristina Höök. 2010. Transferring qualities from horseback riding to design. *Proceedings of the 6th Nordic Conference on Human-Computer Interaction Extending Boundaries - NordiCHI '10* (2010), 226. <https://doi.org/10.1145/1868914.1868943>
- Kristina Höök, Baptiste Caramiaux, Cumhur Erkut, Jodi Forlizzi, Nassrin Hajinejad, Michael Haller, Caroline Hummels, Katherine Isbister, Martin Jonsson, George Khut, Lian Loke, Danielle Lottridge, Patrizia Marti, Edward Melcer, Florian Müller, Marianne Petersen, Thecla Schiphorst, Elena Segura, Anna Ståhl, Dag Svanæs, Jakob Tholander, and Helena Tobiasson. 2018. Embracing First-Person Perspectives in Soma-Based Design. *Informatics* 5, 1 (2018), 8. <https://doi.org/10.3390/informatics5010008>
- Kristina Höök, Martin P Jonsson, Anna Ståhl, and Johanna Mercurio. 2016. Somaesthetic Appreciation Design. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 3131–3142. <https://doi.org/10.1145/2858036.2858583>
- Kristina Höök, Anna Ståhl, Martin Jonsson, Johanna Mercurio, Anna Karlsson, and Eva-Carin Banka Johnson. 2015. Somaesthetic Design. *Interactions* 22, 4 (2015), 26–33. <https://doi.org/10.1145/2770888>
- Lixing Huang, Louis-Philippe Morency, and Jonathan Gratch. 2010. Learning backchannel prediction model from parasocial consensus sampling: a subjective evaluation. In *Proceedings of the 10th international conference on Intelligent virtual agents (IVA'10)*. Springer-Verlag, Berlin, Heidelberg, 159–172. <http://dl.acm.org/citation.cfm?id=1889075.1889095>
- Caroline Hummels, Kees C Overbeeke, and Sietske Klooster. 2007. Move to get moved: a search for methods, tools and knowledge to design for expressive and rich movement-based interaction. *Personal Ubiquitous Comput.* 11, 8 (12 2007), 677–690. <https://doi.org/10.1007/s00779-006-0135-y>
- Tim Ingold. 2017. On human correspondence. *Journal of the Royal Anthropological Institute* 23, 1 (2017), 9–27. <https://doi.org/10.1111/1467-9655.12541>
- Robert J K Jacob, Audrey Girouard, Leanne M Hirshfield, Michael S Horn, Orit Shaer, Erin Treacy Solovey, and Jamie Zigelbaum. 2008. Reality-based interaction: a framework for post-WIMP interfaces. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*. ACM, New York, NY, USA, 201–210. <https://doi.org/10.1145/1357054.1357089>
- William James. 1890. *The principles of psychology, Vol I*. <https://doi.org/10.1037/10538-000>
- Rose Johnson, Nadia Bianchi-Berthouze, Yvonne Rogers, and Janet van der Linden. 2013. Embracing calibration in body sensing: using self-tweaking to enhance ownership and performance. *UbiComp '13: Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing* (2013), 811. <https://doi.org/10.1145/2493432.2493457>
- Frédéric Bevilacqua Jules Françoise, Baptiste Caramiaux. 2012. A hierarchical approach for the design of gesture-to-sound mappings. *Smc'12* (2012), 1 – 8.
- Raine Kajastila and Perttu Hämäläinen. 2015. Motion games in real sports environments. *interactions* 22, 2 (2015), 44–47. <https://doi.org/10.1145/2731182>
- Raine Kajastila, Leo Holsti, and Perttu Hämäläinen. 2016. The Augmented Climbing Wall: High-Exertion Proximity Interaction on a Wall-Sized Interactive Surface. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 758–769. <https://doi.org/10.1145/2858036.2858450>
- Pavel Karpashevich, Eva Hornecker, Michaela Honauer, and Pedro Sanches. 2018. Reinterpreting Schlemmer ' s Triadic Ballet : Interactive Costume for Unthinkable Movements. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18* (2018), 1–13. <https://doi.org/10.1145/3173574.3173635>
- Simon Katan, Mick Grierson, and Rebecca Fiebrink. 2015. Using Interactive Machine Learning to Support Interface Development Through Workshops with Disabled People. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15* (2015), 251–254. <https://doi.org/10.1145/2702123.2702474>
- Pedro Kirk, Mick Grierson, Rebeka Bodak, Nick Ward, Fran Brander, Kate Kelly, Nicholas Newman, and Lauren Stewart. 2016. Motivating Stroke Rehabilitation Through Music. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*. ACM Press, New York, New York, USA, 1781–1785. <https://doi.org/10.1145/2858036.2858376>
- David Kirsh. 2009. knowledge, explicit and implicit. 397–402.
- David Kirsh. 2013. Embodied Cognition and the Magical Future of Interaction Design. (2013).

- Andrea Kleinsmith and Marco Gillies. 2013. Customizing by doing for responsive video game characters. *International Journal of Human-Computer Studies* 71, 7 (2013), 775–784. <https://doi.org/10.1016/j.ijhcs.2013.03.005>
- Todd Kulesza, Margaret Burnett, Weng-Keen Wong, and Simone Stumpf. 2015. Principles of Explanatory Debugging to Personalize Interactive Machine Learning. In *Proceedings of the 20th International Conference on Intelligent User Interfaces - IUI '15*. ACM Press, New York, New York, USA, 126–137. <https://doi.org/10.1145/2678025.2701399>
- Benjamin Letham, Cynthia Rudin, Tyler H. McCormick, and David Madigan. 2015. Interpretable classifiers using rules and bayesian analysis: Building a better stroke prediction model. *Annals of Applied Statistics* 9, 3 (2015), 1350–1371. <https://doi.org/10.1214/15-AOAS848>
- Henry Lieberman (Ed.). 2001. *Your wish is my command: programming by example*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Siân E Lindley, James Le Couteur, and Nadia L Berthouze. 2008. Stirring up experience through movement in game play: effects on engagement and social behaviour. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 511–514. <https://doi.org/10.1145/1357054.1357136>
- Lian Loke and Claudia Núñez-Pacheco. 2018. Developing somatic sensibilities for practices of discernment in interaction design. *The Senses and Society* 13, 2 (2018), 219–231. <https://doi.org/10.1080/17458927.2018.1468690>
- Pedro Lopes, Alexandra Ion, Willi Mueller, Daniel Hoffmann, Patrik Jonell, and Patrick Baudisch. 2015. Proprioceptive Interaction. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15* (2015), 939–948. <https://doi.org/10.1145/2702123.2702461>
- Wendy E Mackay and Anne Laure Fayard. 1999. Video Brainstorming and Prototyping: Techniques for Participatory Design. In *CHI '99 Extended Abstracts on Human Factors in Computing Systems (CHI EA '99)*. ACM, New York, NY, USA, 118–119. <https://doi.org/10.1145/632716.632790>
- Elena Márquez Segura, Laia Turmo Vidal, Asreen Rostami, and Annika Waern. 2016. Embodied Sketching. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 6014–6027. <https://doi.org/10.1145/2858036.2858486>
- Elena Márquez Segura, Annika Waern, Jin Moen, and Carolina Johansson. 2013. The Design Space of Body Games: Technological, Physical, and Social Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 3365–3374. <https://doi.org/10.1145/2470654.2466461>
- Roisin McNaney, Patrick Olivier, Madeline Balaam, Amey Holden, Guy Schofield, Daniel Jackson, Mary Webster, Brook Galna, Gillian Barry, and Lynn Rochester. 2015. Designing for and with People with Parkinson's : A Focus on Exergaming. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15* (2015), 501–510. <https://doi.org/10.1145/2702123.2702310>
- Maurice Merleau-Ponty. 2002. *Phenomenology of perception*. Routledge. 544 pages.
- Florian Mueller and Katherine Isbister. 2014. Movement-based game guidelines. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14* (2014), 2191–2200. <https://doi.org/10.1145/2556288.2557163>
- Florian Mueller, Frank Vetere, Martin Gibbs, Darren Edge, Stefan Agamanolis, Jennifer Sheridan, and Jeffrey Heer. 2012. Balancing exertion experiences. *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12* (2012), 1853–1862. <https://doi.org/10.1145/2207676.2208322>
- Mathieu Nancel, Julie Wagner, Emmanuel Pietriga, Olivier Chapuis, Wendy Mackay, F Orsay, and F Orsay. 2011. Mid-air Pan-and-Zoom on Wall-sized Displays. (2011), 177–186.
- Donald A Norman. 2010. Natural user interfaces are not natural. *interactions* 17, 3 (5 2010), 6–10. <https://doi.org/10.1145/1744161.1744163>
- Donald A Norman. 2013. *The Design of Everyday Things: Revised and Expanded Edition*. Basic Books. <http://books.google.co.uk/books?id=nVQPAAAAQBAJ>
- Claudia Núñez-Pacheco and Lian Loke. 2018. Towards a technique for articulating aesthetic experiences in design using Focusing and the Felt Sense. *The Design Journal* 6925, May (2018), 1–21. <https://doi.org/10.1080/14606925.2018.1467680>
- J K O'Regan and A Noë. 2001. A sensorimotor account of vision and visual consciousness. *The Behavioral and brain sciences* 24, 5 (2001), 939–973.
- Kayur Patel, Naomi Bancroft, Steven M Drucker, James Fogarty, Andrew J Ko, and James Landay. 2010. Gestalt: integrated support for implementation and analysis in machine learning. *Proceedings of the 23rd annual ACM symposium on User interface software and technology* (2010), 37–46. <https://doi.org/10.1145/1866029.1866038>
- Michael Polanyi. 1966. The Tacit Dimension. *Knowledge in Organizations* (1966), 135–146. <https://doi.org/10.1016/B978-0-7506-9718-7.50010-X>
- Majken K. Rasmussen, Esben W. Pedersen, Marianne G. Petersen, and Kasper Hornbæk. 2012. Shape-Changing Interfaces: A Review of the Design Space and Open Research Questions. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)* (2012), 735–744. <https://doi.org/10.1145/2207676.2207781>
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *ACM SIGCHI Workshop on Human-Centered Machine Learning*, Marco Gillies and Rebecca Fiebrink (Eds.). ACM, San Jose. <http://hcml2016.goldsmithsdigital.com/program/>

- G Ryle. 1949. *The Concept of Mind*. Hutchinson's University. 286–292 pages. <https://doi.org/10.7208/chicago/9780226922652.001.0001>
- Thecla Schiphorst. 2009. soft(n): Toward a Somaesthetics of Touch. *Extended Abstracts on Human Factors in Computing Systems - CHI EA'09* (2009), 2427–2438. <https://doi.org/10.1145/1520340.1520345>
- Thecla Schiphorst. 2011. Self-evidence. *Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems - CHI EA '11* (2011), 145. <https://doi.org/10.1145/1979742.1979640>
- D A Schön. 1983. *The Reflective Practitioner: How Professionals Think in Action*. Basic Books. <http://books.google.co.uk/books?id=ceJIW4y-jgC>
- Donald A Schön. 1987. *Educating the reflective practitioner: Toward a new design for teaching and learning in the professions*. Jossey-Bass. 355 pages. <https://doi.org/10.1182/blood-2010-02-266338>
- Burr Settles. 2010. Active Learning Literature Survey. *Machine Learning* 15, 2 (2010), 201–221. <https://doi.org/10.1.1.167.4245>
- Richard. Shusterman. 2008. *Body consciousness : a philosophy of mindfulness and somaesthetics*. Cambridge University Press. 239 pages.
- Diego Silang Maranan, Sarah Fdili Alaoui, Thecla Schiphorst, Philippe Pasquier, Pattarawut Subyen, and Lyn Bartram. 2014. Designing for movement. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*. ACM Press, New York, New York, USA, 991–1000. <https://doi.org/10.1145/2556288.2557251>
- Patrice Y. Simard, Saleema Amershi, David M. Chickering, Alicia Edelman Pelton, Soroush Ghorashi, Christopher Meek, Gonzalo Ramos, Jina Suh, Johan Verwey, Mo Wang, and John Wernsing. 2017. Machine Teaching: A New Paradigm for Building Machine Learning Systems. (2017). <http://arxiv.org/abs/1707.06742>
- Mel Slater. 2009. Place illusion and plausibility can lead to realistic behaviour in immersive virtual environments. *Philos Trans R Soc Lond B Biol Sci* 364, 1535 (12 2009), 3549–3557. <https://doi.org/10.1098/rstb.2009.0138>
- Bernhard Spanlang, Jean-Marie Normand, David Borland, Konstantina Kilteni, Elias Giannopoulos, AusiÀ s PomÀ©s, Mar González-Franco, Daniel Perez-Marcos, Jorge Arroyo-Palacios, Xavi Navarro Muncunill, and Mel Slater. 2014. How to Build an Embodiment Lab: Achieving Body Representation Illusions in Virtual Reality. *Frontiers in Robotics and AI* 1 (11 2014), 9. <https://doi.org/10.3389/frobt.2014.00009>
- R S Sutton and A G Barto. 1998. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA.
- Justin Talbot, Bongshin Lee, Ashish Kapoor, and Desney S Tan. 2009. EnsembleMatrix: interactive visualization to support machine learning with multiple classifiers. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*. ACM, New York, NY, USA, 1283–1292. <https://doi.org/10.1145/1518701.1518895>
- Shogo Tanaka. 2013. The notion of embodied knowledge and its range. *Encyclopaedia* (2013). [https://doi.org/10.4442/ency\\_{37}\\_{13}\\_{03}](https://doi.org/10.4442/ency_{37}_{13}_{03})
- Jordan A. Taylor and Richard B. Ivry. 2012. The role of strategies in motor learning. *Annals of the New York Academy of Sciences* 1251, 1 (2012), 1–12. <https://doi.org/10.1111/j.1749-6632.2011.06430.x>
- Andrea L. Thomaz and Cynthia Breazeal. 2008. Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence* 172, 6-7 (2008), 716–737.
- K Thórisson. 1998. Real-time Decision Making in Multimodal Face-To-Face Communication. In *second ACM international conference on autonomous agents*. 16–23.
- John Underkoffler and Hiroshi Ishii. 1999. Urp. In *Proceedings of the SIGCHI conference on Human factors in computing systems the CHI is the limit - CHI '99*. ACM Press, New York, New York, USA, 386–393. <https://doi.org/10.1145/302979.303114>
- Eduardo Velloso, Andreas Bulling, and Hans Gellersen. 2013. MotionMA: Motion Modelling and Analysis by Demonstration. (2013), 1309–1318. <https://doi.org/10.1145/2466110.2466171>
- Soroush Vosoughi. 2014. Improving automatic speech recognition through head pose driven visual grounding. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14* (2014), 3235–3238. <https://doi.org/10.1145/2556288.2556957>
- Fulton Wang and Cynthia Rudin. 2015. Causal Falling Rule Lists. (2015). <http://arxiv.org/abs/1510.05189>
- Gary L. Wells and Richard E. Petty. 1980. The Effects of Over Head Movements on Persuasion: Compatibility and Incompatibility of Responses. *Basic and Applied Social Psychology* 1, 3 (6 1980), 219–230. [https://doi.org/10.1207/s15324834bas0103\\_{3}](https://doi.org/10.1207/s15324834bas0103_{3})
- Danielle Wilde. 2010. Moving To Move :. *Second Nature* 1 (2010), 164–197.
- Jacob O Wobbrock, Andrew D Wilson, and Yang Li. 2007. Gestures Without Libraries, Toolkits or Training: A \$1 Recognizer for User Interface Prototypes. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology (UIST '07)*. ACM, New York, NY, USA, 159–168. <https://doi.org/10.1145/1294211.1294238>
- Bruno Zamborlin, Frédéric Bevilacqua, Marco Gillies, and Mark D'Inverno. 2014. Fluid Gesture Interaction Design: Applications of Continuous Recognition for the Design of Modern Gestural Interfaces. *ACM Trans. Interact. Intell. Syst.* 3, 4 (1 2014), 20:1–20:36. <https://doi.org/10.1145/2395123.2395125>