

Learnable Computing

MARCO GILLIES, Goldsmiths, University of London, UK

We have to learn all new technologies and we continue to learn for as long as we use them and develop that use. Learning is therefore an integral part of human engagement with technology, as it is with all areas of life. This paper proposes that we should consider learning as an important part of all human computer interaction and that theories of learning can make an important contribution to HCI. It presents 6 vignettes that describe different ways in which this could happen: rethinking HCI concepts in terms of learning, applying learning theory to better understanding established ideas in HCI, using learning research to inform HCI practice, understanding how people learn software and inspiring us to rethink the aims of this discipline. This paper aims to start a conversation that could bring valuable new ideas into our “inter-discipline”.

CCS Concepts: • **Human-centered computing** → **HCI theory, concepts and models**; • **Social and professional topics** → **Informal education**; **Computing education**.

Additional Key Words and Phrases: HCI Theory, Learning, Education

ACM Reference Format:

Marco Gillies. 2021. Learnable Computing. In *Proceedings of CHI '21: CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI '21 EA)*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3411763.3450386>

1 INTRODUCTION

There is an obvious fact about human computer interaction that is not remarked upon as much as it should be: we have to learn how to interact with computers. However, learning is not as integral to our work as it maybe should be. We tend to study competent technology use by people who have been interacting with computers (or similar devices) for years and have long since learned the basic skills of their use, or we perform lab studies in which people are presented with piece of technology for the first time and study how they interact with it, without giving time for learning.

This paper makes the argument that HCI should attend more to learning and that it can bring us new insights into the core concepts of our field. HCI as an “inter-discipline”[3] has flourished and grown through theories and insights imported from other disciplines, whether it is the philosophy of phenomenology[7] or Theatre and Performance[18]. The study of learning is itself a large inter-discipline ranging from the detailed research of psychology and neuroscience into the mental and neural processes of learning, or the situated study of classroom experiences in educational studies. These bring a wide range of theoretical perspectives and scientific results which can help inform our discipline and think about how we learn to use technology and what that implies for all of our interactions with our devices.

Of course, HCI does pay attention to learning and education. CHI has an entire subcommittee devoted to it, but it primarily focuses on how HCI and technology can support learning, normally within an educational setting. It is less focused on how learning is an integral part of all human computer interaction, not simply the first time we pick up an app or in an explicit training session or tutorial, but a constant process of improving as we interact with technology.

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.

© 2021 Association for Computing Machinery.

Manuscript submitted to ACM

The rest of this paper will sketch some vignettes of how a focus on learning and an awareness of the diverse literature on learning can add to the theory and practice of HCI. This can be done in a number of ways:

- Considering ideas in HCI in terms of learning can lead us to think about them in new ways. Section 2 looks as a big idea in HCI (affordances) in terms of learning and looks at the implications.
- Theories of learning can help us rethink and better understand ideas in HCI. Section 3 looks at how constructivist theories of learning can help us understand the role of metaphor in HCI.
- Learning research can influence our practice in HCI. Section 4 discusses results relating to practice and testing and how they might influence game design.
- Theories and research in learning can help understand how people learn to use technology and we can try to understand which theories of learning apply best to technology. Section 5 looks at whether socially situated theories of learning in communities of practice might be a better model of learning technology than school based learning.
- We can try to fundamentally rethink the aims of HCI, just as educators have done for education. Section 6 looks at how liberatory pedagogies challenged the traditional aims of education and how HCI might do the same.

2 LEARNING AFFORDANCES

Affordances are one of the most commonly cited theoretical concepts in HCI. The term was first introduced to the HCI literature by Gaver[9] and popularised by Norman[20], who defined the concept as follows:

An affordance is a relationship between the properties of an object and the capabilities of the agent [person] that determine just how the object could possibly be used.

Examples range from the flat surface of a chair or bench affording sitting to a scroll bar in a user interface affording scrolling. While the concept has been very widely used in HCI it originates in psychology (an excellent example of the importing of theoretical perspectives into HCI discussed in the introduction) with the work of J.J. Gibson[11], who made them central to his ecological theory of human perception.

Affordances are an excellent conceptual tool to think about how humans relates to technologies. However, the definition as it stands does somewhat confuse the psychological status of affordances. Norman's definition (which echoes Gibson's) places affordances outside of either the object or the person, but in their relationship. This way of viewing them gives them an apparent objectivity and universality, which is implied by the writing of Gibson, Gaver and Norman. However, Gibson used affordances as part of his theory of perception, in which affordances are something that can be perceived. They must exist, or at least be represented in our perceptual system. Do we therefore learn affordances (or at least do we learn to perceive them)?

A paradigm example of an affordance in ecological psychology is the visual cliff experiment by Eleanor Gibson (J.J. Gibson's wife) and Richard Walk[10]. An apparatus is set up with a raised platform next to a sudden drop (the cliff). Both are textured with a strong pattern (e.g. a checkerboard) so that the depth is easily perceivable. However, a glass panel is placed over the drop, so it is, in fact, possible to walk over it safely. Gibson and walk tested crawling babies and young animals on this apparatus and they almost universally refused to crawl or walk over the drop. The visual appearance of the raised platform showed an affordance for support and crawling, while the lower area had a strong visual dis-affordance for the same (though as some babies discovered, there was a haptic affordance for crawling). The fact that young babies clearly perceived the affordances (or lack thereof) for physical support again implied their innateness and universality.

However, later research by Adolph and others (see Adolph and Holt[1] for an overview) tells a different story. Later studies compared infants who had only just started crawling with infants who had been crawling for several weeks. While experienced crawlers avoided the drop, new crawlers would cross it. They had not yet learned to perceive that the drop did not afford support and crawling. In fact, the same effect can be seen in the transition from crawling to walking. Confident crawlers refuse to cross the drop, but new walkers will do so. This clearly shows that the affordance (or at least its perception) is learned.

This isn't particularly surprising, it is hard to see how the affordances of a scroll bar are innate, but if in HCI we focus on competent abilities, and not learning, we might forget about this obvious fact. Why does this matter? As Costanza-Chock[6] argues, the false implication of the universality or innateness of affordances can lead to us to erase the experience of those who have not learned those affordances, whether because they come from a different cultural or social background from a typical technology user (learning technology affordances implied sufficient wealth to access that technology) or they have different physical capabilities (a staircase has very different affordances to a wheelchair user). The affordances that we think are perceptible to all, and provide access to technology, can be *disaffordances*[29] to many people, themselves inaccessible and rendering the technology inaccessible.

A focus on learning and how we learn to use technology can therefore help avoid the trap of assumed universalism and therefore help us realise that the ability to interact with technology, however “intuitive”, is the result of a lived experience that has been embedded both in particular physical body and a particular socio-cultural situation.

3 METAPHORS, CONSTRUCTIVISM AND NEURONAL NETWORKS

User interface metaphors are an old idea in HCI, going back to the desktop metaphor that underlies all our modern graphical user interfaces. The idea of the desktop metaphor was to provide users of computers with a way of representing the workings of a computer, and interacting with it, that is familiar from traditional office working practices. However, the desktop metaphor is not much like a real desk, let alone a traditional office: paper files are not stored in hierarchies of paper folders, nor do they typically stored on a desk, for example, and the way we work with computers is very different from the physical habits of office work. Even worse, we now no longer use these traditional office items: a young person is likely to be much more familiar with a computer file than a paper file. A competent computer user is very unlikely to be thinking about a traditional desktop when they are using a desktop based GUI.

So what is the purpose of the desktop metaphor, and HCI metaphors in general? An obvious answer is that it is to support learning to use a computer rather than experienced practice and theories of learning can help elucidate why it works.

Constructivism is one of the most important set of theories of learning in the 20th Century and is still highly influential in the 21st. It is traditionally traced back to the work of Piaget[?] on child development (though precedents exist) and was developed through the work of thinkers such as Vygotsky[26] and Bruner[4] who introduced social and cultural context (often called Social Constructivism). It has since developed into a very wide ranging theory of learning, education and cognition.

The basic idea of constructivism is that learners are not passive recipients of knowledge that is “input” to them by a teacher, but that they actively construct their knowledge based on their prior knowledge and experienced combined with the learning activities they are undertaking¹. It is a huge complex of diverse theories that I cannot do justice

¹Constructivism is related to, but distinct from Constructionism, an approach to learning that might be familiar to HCI researchers from the work of Papert[?] and Resnick [?] on the classic learning tools for programming, Logo and Scratch. Constructionism is a learning method in which learners learn by constructing things (in the case of Logo and Scratch they are constructing computer programs). While constructivist thinkers often advocate for this type of active learning, students can be thought of as actively constructing their own representation of knowledge in even the most passive lecture setting

to in a short paper such as this one, so I will focus on introducing the ideas of Vygotski and then also cover a more contemporary take by Zull[30] which links ideas of learning to neuroscience, and which I think is instructive.

Vygotsky held that a child's social and cultural background was key to their learning. They acquired what he called *cultural tools* from their prior socio-cultural experience. A cultural tools are ways of representing and communicating that support knowledge and understanding in a way analogous to how physical tools support physical action. Vygotsky's paradigm example of a cultural tool is language, as it supports both the communication of knowledge and it's internal representation (for Vygotski thought is internalised language) and underlies the vast majority of our learning in some form or other. Language is not a single cultural tool as there are many different specialist languages or ways of using language, for example, acquiring "scientific language" enables us to learn scientific concepts. Other cultural tools include pictorial representations, narratives and core concepts of a domain of knowledge. The basic idea is that learners use cultural tools in order to learn new knowledge and new cultural tools. They must learn fundamental cultural tools first and must then use them to construct new knowledge that they are learning.

These kinds of constructivist ideas are supported by more recent developments in neuroscience[30]. At a physical level, learning means the strengthening (or weakening) of connections between neurons, primarily in our brain (it can also involve the creation of new neurons). When we first encounter a new concept or experience, it causes the firing of neurons across existing patterns of connections (called circuits), because that is all their is in the brain. As neurons fire together in response to the new experience, their connections are strengthened, and over time this can result in the development of new circuits. Learning, therefore, must build on existing brain circuits (i.e. prior knowledge) and involves the active construction of new brain circuits (i.e. new knowledge).

What has all this got to do with user interface metaphors? The key idea is that learning new things must build on existing patterns of knowledge. A metaphor explicitly creates links between different areas of knowledge. The desktop metaphor links knowledge of an established domain (a traditional office) with a new area of knowledge (computers). This supports new users in learning a new technology by pointing them towards prior knowledge (existing neural circuits) on which to build their new understanding of computer use. There prior knowledge isn't used directly, desktop metaphor interfaces are too different from real desktops for that to happen. Instead it is used as the building blocks, or cultural tools for constructing new knowledge. The metaphor provides a scaffold to build that new knowledge by making explicit the links with prior knowledge (metaphors are often used in a similar way in education, for example, the water pipe metaphor of electricity).

What then, can we say about the other objection to the desktop metaphor, that people are now more familiar with computers than they are with traditional office fixtures such as filing cabinets. Why do we still continue to talk about files? This leads us to another key HCI concept that is explained by constructivist views of learning: consistency. Consistency, one of Norman's design principles[20], is that similar elements should be used to achieve similar aims. This is applied not only within individual pieces of software, but across different pieces of software. This is best exemplified by user interface toolkits that provide a consistent set of widgets for all software to use, and thus enabling standardised ways of interaction with elements like buttons, tabs, or folder structures.

Although normally considered separate from the idea of metaphor, if we view it from perspective of constructivist learning theory, consistency has the same function. A consistent interface allows us to use prior knowledge (in this case of using other software) as a basis for learning new software. We can think of using a user interface toolkit as a form of metaphor, but this time the source of the metaphor is not the real world, but other software. There are many examples of this, one is the use of red underlining to indicate syntax errors in programmers' editors, this is a metaphor

in the sense that it links a syntax error in programming to a spelling error in natural language (since red underlining is commonly used in word processors).

I hope that this section has shown that we can use theories of learning to provide a deeper understanding of key concepts in HCI, in this case metaphors and consistency. An acquaintance with the literature of learning can, therefore, be a potentially valuable cultural tool of both HCI researchers and practitioners.

4 TRAINING, TESTING AND TUTORIALS

The previous two sections have shown how we can rethink key HCI concepts by thinking in terms of learning, but true power of bringing in a new discipline to the interdisciplinary mix of HCI is that it can bring with it entirely new ideas.

Here we turn to another branch of the study of learning, one grounded more in laboratory studies than in educational practice. The area commonly called *Science of Learning*[14] originates in psychology and is founded on detailed studies of how we learning new knowledge (for example, memorising facts or learning languages) and physical skills (e.g. sports skills such as baseball batting). It focuses on the the details of how training, practice and testing oneself influence learning.

One of the most important results is that testing oneself is key to learning. Simply reading and re-reading material (or watching lectures or videos) is not as effective as putting your knowledge to the test, whether by doing practice questions or other means. It is very easy to fool yourself into thinking you have an understanding of a topic when reading a passage, but still not be able to recall or reconstruct the ideas when you don't have the explanation in front of you. The value of testing extends to some quite unintuitive results. For example, testing can help learning even before any material has been learned. Taking a practice version of final exam on the first day of a class dramatically improves learning, even if students score close to zero on it. Taking a test before learning (called pre-testing) seems to prime people for learning that will happen later.

Another set of results relate to practice. Practice is key to learning (and testing is an example of practice), but not all practice is equally beneficial. Ericsson[8] calls the most valuable type of practice *Deliberate Practice*. There are many criteria for deliberate practice but a brief summary would be that deliberate practice has a strong focus on areas that require improvement and a strong use of feedback, rather than simply repeating the same actions time after time.

Another result about practice is about timing. Often, in the classroom a new technique (e.g. a mathematical method) is introduced and practiced intensely before moving on to another technique. This approach has been shown to ineffective for long term learning. People learn better if they *space* their practice out over longer time periods[14] (there are many complex results relating to optimal spacing). They also learn better if they *interleave* their practice, practicing a range of different techniques at the same time or in rapid succession. This may be because they learn how to select the correct technique to use in a given context, not simply how to apply it. Interestingly, students who use spaced and interleaved practice perform worse in the short term, so it can feel less effective, but their learning lasts longer.

What implications does this have for HCI? The most obvious is for technology that is explicitly designed for learning or training, whether it be a language learning app or sport and fitness technology. These types of app do typically involve educationalist and will make use of science of learning results. In fact it is one of the benefits of using learning technology is that spaced and interleaved practice schedules can easily be programmed in.

Another area is for learning software itself, particularly the design of tutorials. These are common in many areas of software but are particularly sophisticated in videogames. In videogames players must learn very complex skills, and constantly improve in order to progress. Early games often had a tutorial section separate from the game which introduced the key mechanics, but it is now more common to integrate the tutorials with the narrative of the game.

This is intended to make the tutorials less boring, but it is also clearly better from the point of view of the science of learning, as it implements a degree of spaced practice and interleaving. However, it is possible to look at the science of learning more deeply to design onboarding experiences for games that will support players to constantly improve and therefore get the sense of achievement that is a key part of the enjoyment of games.

The idea of pre-testing suggests that players should encounter challenges before they are taught the skills needed to solve them, they should be allowed to fail first. This sounds harsh, but failure is common in games and players will be very willing to learn the skills they need once they have been defeated (and are more likely to remember it). Deliberate practice suggests that players should be focused on areas that they are less good at. This happens already to a degree, since they typically have to repeat sections that they fail, but a system that detects problems with a particular skills (e.g. a player is not accurate with archery, but fine with a sword) can be used to focus use of that skill. Spaced and interleaved practice can be a good way of structuring progress through a game, sections should combine new skills with existing skills and new skills should be introduced in such a way that their practice is spread out over many sections of a game, each with increasing difficulty. This all suggests that the trend to integrating tutorials into a game should be pushed even further with tutorial elements throughout (as does sometimes happen, though often only when new items or powers are introduced), and the entire gameplay being viewed as an learning experience throughout and one in which players are constantly encouraged towards deliberate practice (this does, of course, happen in many games, but the science of learning can be a useful lens for making this explicit in the design process).

5 SITUATED LEARNING AND COMMUNITIES OF PRACTICE

Another question that we can ask of learning theories is how do people learn to use computer technologies? Our common conceptions of learning, and the majority of learning theories, situate learning within explicitly educational contexts: schools and to a lesser degree higher education or adult education and training. There is certainly some element of this type of classroom type education in computing. I, and most academics reading this, teach classes that involve some degree of learning of software. There are plenty of professional training courses in software, and to a lesser degree tutorial modes and youtube videos also follow this model.

However, this type of learning focuses on high end professional software such as Photoshop, Excel or Unity. It has little to say about how we learn informal software such as Instagram, minecraft or a web browser. Even for high end software the majority of learning a professional will do is outside of the classroom, what is commonly called informal learning.

So is most of this learning trial and error or individual practice? Certainly, this there is a fair amount of this, and the theories of practice outlined in the last section do apply to this type of learning, but this is a very individualistic view of learning and ignores the fact that we do a lot of learning from or otherwise involving other people, it is just that these other people are not teachers or explicitly teaching us.

A theory of learning that is very relevant to this type of informal, social learning is the *Situated Learning* of Lave and Wenger[19]. They use the term “situated” in a way that will be familiar with HCI practitioners: embedded with a social and physical context and innately entangled with it. Of course, explicit education is situated in a social context: that of the school, but they are interested in learning situated outside the academy.

In particular, they base their research on anthropological studies of apprenticeships, whether that is West African tailors, butchers in a western supermarket or more loosely defined apprenticeship, such as that of recovering alcoholics in Alcoholics Anonymous. They see that these apprenticeships are very different from the classroom, and in particular they involve very little explicit teaching.

Instead the focus is on becoming part of a *Community of Practice*[28]. By this they mean a formal or informal community centred around a particular practical activity that is often, but not always, professional. This idea seems to apply well to tailors or game developers, but it also applies to non-professional software use. It makes little sense to think of Instagram without thinking about communities of users and video games like minecraft have strong communities. Office environments might not be as strongly focused on a piece of professional software as game developers or graphic designers are, but they provide informal communities around the use of technologies such as word processors or email. Similarly, families and friends can provide communities that support the use of certain social or leisure software.

Lave and Wenger see learning within communities of practice not simply being about acquiring certain technical skills, but also knowing how to behave within a community and acquiring an identity as community member. This certainly seems to apply to social media, which is as much about identity and community membership as skills, and also to professional software use: game developers or graphic designers have strong professional identities. Other work environments are less likely to define identity and membership in terms of skillful use of a piece of software, but that is the point: software skills are only one small part of a very complex community identity. For example, how some one writes an email is just one element of the professional identity of, say, a lawyer.

Lave and Wenger theorise the implicit learning process of apprentices as *Legitimate Peripheral Participation*. *Participation* is key: apprentices make genuine, useful contributions to their community or practice. This participation is *legitimate* as they are acknowledged members of the community. *Peripheral* implies two things. The most obvious is that apprentices are not yet full members of the community, they are at it's periphery. However, being at the periphery also means they are not outside the community, in the way a university student might be outside a community of work, even when learning the skills it requires. This model is useful for thinking about new members of a community, or new learners of a particular technology, however, it is also important to remember that learning also takes place between established members of a community of practice, particularly when dealing with rapidly evolving technology.

The ideas of Lave and Wenger are one area of learning theory that is quite well known within the HCI and CSCW communities, and is applied to analysing a wide range of participation, and often learning, in communities as diverse as live streamers[12], activist[15] or hacker and tinkerers [27]. However, I believe these ideas can be made even more central if we acknowledge that all of the technology that we study or develop at CHI exists, and it learned, within communities of practice. It is not simply a useful concept for understanding communities but all software.

6 TECHNOLOGY AS A PRACTICE OF FREEDOM

An important feature of theories of learning and educational studies, is that they do not simply differ in terms of what is considered the best technique for achieving a universally agreed aim. Learning theorists are willing to challenge the very goals of education. Education and learning are very important parts of human life and society, and learning theories are fundamentally about the human condition and what we want our societies to be like: they are innately political and ideological.

Many naïve views of learning simply see it as the process of accumulation of knowledge or skills which is “input” to a student from a teacher, and in many ways our education systems follow these methods, particularly in the form of standardised testing. However, many educational thinkers challenge this view and in fact have much more ambitious aims for education: to liberate and improve people's lives.

Prominent among these thinkers are Paulo Freire[?], a Brazilian Educator who taught literacy to the poorest rural peasants and Bell Hooks[13], an African-American feminist scholar who taught multicultural classes in US universities. Both of these educators saw “education as a practice of freedom” (the title of Freire's first book and subtitle of Hooks'

first book on education). They sought not simple to instruct their students with knowledge but to liberate them and their thinking. Freire saw education as a process of what he calls *Conscientização* (often translate to English as *Conscientization*), a complex terms which can roughly be described as a person gaining a critical awareness of then own social condition. Rather than teachers teaching students how and what to think, a Freirean teacher will support them to critique and come to their own conclusions. This is done through a process of posing problems rather than solutions and through genuine dialogue in which the voice of the student is as important, if not more, than that of the teacher.

Changing the goal of education also changes the methods used. The difference between a traditional classroom and that of Freire or Hooks is not simply what is taught it is how teaching is done, the very fundamental methods of education, and why it is done, the fundamental aim of learning. In fact, debates about educational methods often masks fundamental disagreement about the aims of education.

What can HCI learn from these pedagogies of freedom? Probably the most important thing is that our work deeply affects society and is therefore implicitly political. We need to not only debate the methods we use but the fundamental goals of our field. Can we envision HCI in the service of human liberation and flourishing, rather than simply making better products? Can we envision HCI as a practice of freedom?

This doesn't seem impossible, if we think about how fundamentally technology affects both individuals and society, it has the potential for equally fundamental positive liberatory impacts (and some technology, of course does). Thinking of our aims in this way could help to build a better world.

This change of emphasis would require different methods, but this is not alien to HCI. Participatory design[17] has many affinities with a Freirean pedagogy because they are both fundamentally about dialogue. A participatory designer takes seriously the views of users just as liberatory teachers take seriously the viewpoints of students. However, though early participatory design was often associated with liberatory politics, that aspect is often forgotten in contemporary practice (as is often the case in the practice of the kind of dialogical teaching that Freire and Hooks proposed).

A truly liberatory HCI needs to focus both on goals of liberating people and methods that take those very people seriously. This is a big change to how we have worked before now, but after Black Lives Matter and the upheavals of 2020, isn't it our responsibility to deeply consider these issues?

7 CONCLUSION

This paper has presented a number of vignettes of how an attention to learning, and to the theories and science of learning can help us rethink our own field of HCI and push it forward, with the intention of raising the profile of these ideas in the HCI community and encouraging other researchers to look at them. Ideas of learning are not entirely new to HCI, of course, but are often confined to explicitly education settings, for example classroom education in coding[? ?], but humans learn all the time, particularly in their encounters with technology, so an attention to learning is relevant far beyond the classroom. In fact, as I have suggested, we are learning all the time as we interact with technology. There is a possible objection that making technology that we have to learn goes against the aims of HCI: making technology "intuitive". However, there are two objections to this view. Firstly, the fact that a technology may feel "intuitive" is in fact the results of a lot of learning and experience with prior technology, not simply the properties of the technology. Secondly, it isn't clear that all technologies should be usable without a period of learning. Complex professional software ranging from Excel to Photoshop does require learning, because it unlocking its full power is only possible for highly skilled users. Removing the requirement of learning would inevitably reduce the power and complexity of such software.

The academic study of learning is a huge field, going back to the very founders of human intellectual thought, such as Confucius[24] or Plato[2], in fact it could be said to be several huge fields if we separate education studies, psychology and sociological/anthropological approaches. This paper has only been able to provide the briefest survey and the most superficial sketches of complex ideas (of which I am by no means an expert myself), this is by no means an end point but a jumping off point of readers to explore these ideas further, and hopefully disagree with and pick holes in my presentation.

It should be noted that learning is a field full of controversies and that many of the ideas presented here might be seen by many as antithetical. However, I would encourage readers not to take sides too quickly but look at commonalities. For example, the science of learning's focus on testing might seem aligned to ideas of standardised testing which would be opposed by most constructivist thinkers and all Freirean educators. However, scientists study testing because it is easy to measure in the lab. What the results really say is that students need to actively apply ideas over long, variable time frames and to do so in ways that combines many different concepts, without putting them into silos. These are precisely the characteristics of the types of active learning advocated by constructivists and problem posing dialogic learning advocated by Freireans. Conversely, while many constructivists advocate active learning their theory holds that students also actively construct knowledge in more passive educational environments including those that use standardised testing.

I could not discuss Freirean education in this paper and then attempt to impose my own views on you, the reader. This should be a call for you to develop your own ideas and for dialogue, so I will end by posing problems and questions.

- How could paying attention to learning change the way you approach your areas of HCI?
- What theories of learning could help rethink the big ideas of your area of research?
- How could learning research influence your HCI practice?
- How do people learn to use the technologies you study? What learning theories would be relevant?
- How can you challenge and rethink the fundamental goals, not simply the practice, of your work?

8 A NOTE ON READINGS

If there is a call to action from this paper it is to find out more about research on learning, which will involve some reading. In the text I have referenced primary text where possible, but these are often not the best introduction for novices to the field(s). I will list a few ideas of books to look at (though there are also many interesting YouTube videos, web based articles, online courses etc.).

Educational studies and theory of learning is a big field and the best starting point is probably a broad introductory text. I personally have used "Understanding and Using Educational Theories" by Aubrey and Riley[?], since it critically surveys the ideas of many important thinkers and is a good entry to further reading (though there are many other options). For more in depth reading the Bloomsbury Library of Educational Thought has introductory books on many educational thinkers, including Piaget[16], Vygotsky[25], Bruner[22] and Freire[23] (though it is worth exploring other books in the series).

The Science of Learning literature tends, unfortunately, to be rather distinct from educational studies since it is generally studied in psychology departments. I have already cited "From the laboratory to the Classroom" by Horvath, Lodge and Hattie[14], which is a good introduction including essays from some of the principle researchers in the field, but there are also a number of popular science books that provide an accessible and enjoyable introduction to the field,

for example “How We Learn” by Carey[5] or Mindshift by Barbara Oakley[21]. These are primarily aimed at students and learners aiming to improve their learning, but do provide good surveys of the key ideas.

REFERENCES

- [1] Karen E Adolph and Justine E Hoch. 2019. Annual Review of Psychology Motor Development: Embodied, Embedded, Enculturated, and Enabling. *Annu. Rev. Psychol* 70 (2019), 141–164. <https://doi.org/10.1146/annurev-psych-010418>
- [2] Robin Barrow. 2014. *Plato*. Bloomsbury Publishing.
- [3] Alan F. Blackwell. 2015. HCI as an Inter-Discipline. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems Extended Abstracts- CHI EA '15*. ACM Press, New York, New York, USA, 503–516. <https://doi.org/10.1145/2702613.2732505>
- [4] Jerome Bruner. 1997. *The Culture of Education*. Harvard University Press. <https://doi.org/10.2307/j.ctv136c601>
- [5] Benedict Carey. 2014. *How We Learn: The Surprising Truth About When, Where and Why It Happens*. Pan Macmillan.
- [6] Sasha Costanza-Chock. 2020. *Design Justice: Community-Led Practices to Build the World We Need*. MIT Press. <https://doi.org/10.1080/17547075.2020.1831294>
- [7] Paul Dourish. 2001. *Where The Action Is: The Foundations Of Embodied Interaction*. MIT Press.
- [8] K. Anders Ericsson. 2008. Deliberate practice and acquisition of expert performance: A general overview. In *Academic Emergency Medicine*, Vol. 15. 988–994. <https://doi.org/10.1111/j.1553-2712.2008.00227.x>
- [9] William W Gaver. 1991. Technology Affordances. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '91)*. Association for Computing Machinery, New York, NY, USA, 79–84. <https://doi.org/10.1145/108844.108856>
- [10] Eleanor J Gibson and Richard D Walk. 1960. The "Visual Cliff". *Scientific American* 202, 4 (1960), 64–71. <https://doi.org/10.2307/2494047>
- [11] J J Gibson. 1979. *The ecological approach to visual perception*. Lawrence Erlbaum Associates.
- [12] William Hamilton. 2016. Live Media Places: Participation in Online Education through Composition. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)*. Association for Computing Machinery, New York, NY, USA, 213–217. <https://doi.org/10.1145/2851581.2859017>
- [13] Bell Hooks. 2014. *Teaching To Transgress*. Taylor & Francis. <https://doi.org/10.4324/9780203700280>
- [14] Jared Cooney Horvath, Jason M. Lodge, and John Hattie. 2016. *From the laboratory to the classroom : translating science of learning for teachers*. Routledge.
- [15] Shih-Wen Huang, Minhyang (Mia) Suh, Benjamin Mako Hill, and Gary Hsieh. 2015. How Activists Are Both Born and Made: An Analysis of Users on Change.Org. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 211–220. <https://doi.org/10.1145/2702123.2702559>
- [16] Richard Kohler. 2014. *Jean Piaget*. Bloomsbury Publishing.
- [17] Sarah Kuhn and Michael J. Muller. 1993. Participatory design. *Commun. ACM* 36, 6 (1993), 24–28. <https://doi.org/10.1145/153571.255960>
- [18] Brenda Laurel. 2014. *Computer As Theatre*. Addison Wesley.
- [19] Jean Lave and Étienne Wenger. 1991. *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press.
- [20] Donald A Norman. 2013. *The Design of Everyday Things: Revised and Expanded Edition*. MIT Press.
- [21] Barbara Oakley. 2017. *Mindshift: Break Through Obstacles to Learning and Discover Your Hidden Potential*. Penguin Publishing Group.
- [22] David R Olson. 2014. *Jerome Bruner*. Bloomsbury Publishing.
- [23] Daniel Schugurensky. 2014. *Paulo Freire*. Bloomsbury Publishing.
- [24] Charlene Tan. 2014. *Confucius*. Bloomsbury Publishing.
- [25] Réne van der Veer. 2014. *Lev Vygotsky*. Bloomsbury Publishing.
- [26] Lev S. Vygotsky. 1973. Thought and Language. *Leonardo* 6, 2 (1973). <https://doi.org/10.2307/1572700>
- [27] Tricia Wang and Joseph Kaye. 2011. Inventive leisure practices: Understanding hacking communities as sites of sharing and innovation. *Conference on Human Factors in Computing Systems - Proceedings (2011)*, 263–272. <https://doi.org/10.1145/1979742.1979615>
- [28] Étienne Wenger. 1999. *Communities of Practice: Learning, Meaning, and Identity*. Cambridge University Press.
- [29] D E Wittkower. 2016. Principles of anti-discriminatory design. In *2016 IEEE International Symposium on Ethics in Engineering, Science and Technology (ETHICS)*. 1–7. <https://doi.org/10.1109/ETHICS.2016.7560055>
- [30] James E. Zull. 2002. *The art of changing the brain : enriching teaching by exploring the biology of learning*. Stylus Pub.