Chapter 5: The Learner-Centered Technology Design Framework

Richard Davis, Engin Bumbacher, Veronica Lin, Paulo Blikstein, Chris Proctor, Leah Rosenbaum

5.1. Why we need a new design methodology for learning technology

Designers' *theories of learning* shape the design of learning technologies and, therefore, the learner's educational experiences. This happens because the process of design requires many consequential *choices*, in which theories play a key role. For example, theories inform which design parameters need to be considered, what characteristics of the learners and their environments to pay attention to, or how to measure the learning outcomes. As we discussed in Chapter 3, this effect occurs whether or not the theories are *explicitly evoked* in the design process. For example, a designer who believes that learning occurs through a stream of reward signals (i.e., behaviorism) will focus on creating learning environments in which a learner's actions are interpreted as correct or incorrect. By contrast, a designer who believes that learning occurs through engaging in complex cognitive processes (i.e., diverse flavors of constructivism) will design learning environments that are likely to activate these cognitive processes, such as exploring natural phenomena in an interactive simulation to develop and test hypotheses. Without opportunities to surface and reflect on the theories we hold, our understanding of the range of possibilities in the design process remains limited, and our ability to design effective or new learning technologies remains stifled. At worst, we might design technologies that ultimately harm users' learning potential and contribute to "design hegemony."

How can designers create opportunities to explicitly engage with theories of learning in the design process and use them to inform their decisions? This is where *design methodologies* come into play, by structuring and guiding this process. In general, these methodologies help designers approach a design problem, choose what factors and parameters to consider, and decide how to develop, test and refine prototypes. There are well-known general-purpose design methodologies, such as Human/User-Centered Design, Design Thinking, or Activity-Centered Design, that have proven useful in a wide range of design contexts, from creating new models of smartphones to improving the workflows in a hospital.

But not all design methods are general-purpose. Many fields have developed domain-specific, formalized methodologies. For example, Evidence-Based Design provides a structured process for integrating empirical research into healthcare architecture. Similarly, Service Design offers a methodological approach tailored to the creation and refinement of complex service experiences in service-

oriented sectors such as public administration or hospitality. In education, several domain-specific design methodologies have gained prominence, including Learner-Centered Design, Universal Design for Learning, and Understanding by Design. Learner-Centered Design (LCD), adapted from user-centered design and grounded in constructivist theory, positions learners as active participants in shaping their educational experiences. It emphasizes agency, contextual relevance, and iterative co-design, supporting motivation and deeper engagement. Universal Design for Learning (UDL), by contrast, is rooted in cognitive neuroscience and promotes inclusive instructional design by offering multiple means of engagement, representation, and expression. Originally developed to enhance accessibility, it now serves as a general framework for addressing learner variability in general.

In light of this proliferation of design methods and their successes, you might be tempted to ask why then we had to write a book proposing yet another design method. Through our work as scholars and educators of *learning technologies*, we have come to understand that the design of technology-infused learning environments is particularly challenging, and that none of the existing methodologies adequately assist designers in developing effective learning technologies. First, the more general design methods such as user-centered design do not sufficiently account for the intricacies of education and learning. Learning mathematics is not quite the same as learning how to use a new vacuum cleaner. Second, the existing education-specific methods were created for more generic learning environments and do not sufficiently account for the intricacies of learning with technology. Technology has so many new learning affordances that we need special design methods to help us navigate amongst the myriad of new possibilities brought about by them.

We need methodologies that help designers make explicit the theories of learning and technology they draw on and that help translate the theories into design decisions that guide the design process and outcome. We developed our new design methodology with the goal to address these challenges, and we call it **Learning-Centered Technology Design** (LCTD).

We describe LCTD in two parts. In this chapter, we will discuss the first part: the different types of goals that must be considered when designing a learning technology and the relationships between these goals. In subsequent chapters, we talk about the factors that should be considered when setting these goals and how a designer can meet them.

5.2. LCTD goals and how they relate to each other

There are three kinds of goals in the Learning-Centered Technology Design framework:

- **Learning goals** - what you want people to learn
- **Learning Experience (LX) goals** -the learning experiences you want to make possible with your technology that people will learn from
- **Learning mechanisms (LMs) -** Though not goals themselves, **learning mechanisms (LMs)** - the activities, interactions, and experiences that result in learning, according to learning theory - derive from learning goals and directly inform learning experience goals.
- **User Experience (UX) goals -** the ways people will interact with your technology

5.2.1. What is a goal in the design process, and why do we care?

Design is the process that achieves a desired goal embodied in a product or process. G**oals** are central to this process. Whenever you design, you have a goal in mind—whether it's to build a bridge helping people cross a river or to develop an app that improves mental health. Your intention as a designer is to accomplish this goal. However, there is no exact, fixed way to define a design goal because design is a process that creates something new that does not yet exist. At best, we can identify a goal area or territory described by criteria and metrics. How you formulate and specify these criteria and metrics affects both the design process you follow and how you assess your designs.

In the case of bridge design, the criteria and metrics are mostly quantifiable. For example, a bridge should be able to support a minimum load, withstand certain wind speeds, or use the least amount of materials. The design process is more like a multi-parameter optimization problem; different criteria or metrics influence how a bridge is evaluated but are generally straightforward to interpret and assign.

Things get difficult when design goals defy obvious quantitative metrics. A good example is the domain of mental health. Improving mental health as a design goal for an app does not translate into a straightforward and measurable optimization function. What does it mean for a person to be mentally healthy? Does it mean being more aware of one's mental state so that one knows when to seek help, or does it mean adopting healthy behaviors, such as regular physical workouts? And how can an app help a person become mentally healthy? Is the goal of the app to help users become mentally healthy without it, or to use it continuously to track data for personalized recommendations? Each of these possible goal descriptions has dramatically different consequences for the design and its evaluation. For example, a focus on healthy behaviors might lead to an app that supports individual users in changing their habits. In contrast, a focus on self-awareness and help-seeking might lead to an app that facilitates connections with the local community.

Thus, well-defined design goals are central to the design process, influencing every aspect from development to evaluation. And clearly defining design goals

for complex domains can be challenging. In this sense, designing for education is more akin to designing for mental health than designing a bridge.

5.2.2. The Multiple Goals of Learning Technologies

When it comes to learning technologies, speaking of a design goal as a single (albeit multi-dimensional) construct does not do justice to the complexities of learning. Learning through the use of technology is different from using technology to achieve a non-learning outcome. Learning is both a process and an outcome, an experience and a goal.

In LCTD, there are two types of goals that are inextricably linked in designing learning technologies: **learning goals** and **design goals**. **Learning goals** capture **what** you (the designer) want people to learn as a result of using said technology. **Design goals** describe **how** you (the designer) intend your technology to be used in service of these goals. There are also two kinds of "hows" when it comes to using the learning technology: "how" in the classical sense of usability (UX design goals) and "how" in the sense of a learning experience (LX design goals). There is also a fourth element, learning mechanisms, that specifies at a higher level "how" learning experiences result in learning goals, according to learning theory. We elaborate on each of these goals and elements below.

5.2.3. Learning Goals

A learning goal answers the question "What do you want people to learn?"

Put less tautologically, **a learning goal captures how you want people to change from using your technology.** We interpret "change" in this book in an expansive sense that includes one's perspectives, beliefs, attitudes, emotional or physical capacity, in addition to a change in knowledge or skills. In other words, a learning goal **does not** have to only take the form of a learning outcome like "*Solve a simple system consisting of a linear equation and a quadratic equation in two variables algebraically and graphically.*" This is an example of a content-based learning goal that is often found in educational standards and explicitly describes a change in procedural or conceptual knowledge. We also consider other types of changes (e.g., "become more empathetic") as equally valid learning goals, and a number of learning technologies - such as immersive VR or role-playing card games - have been designed to meet this goal. Even though becoming more empathetic does not necessarily involve a change in knowledge (as we usually define it in schools,) it is still a type of change that we consider as learning.

Learning involves a starting point (the knowledge, skills, attitudes, perspectives, etc. people already have) and an endpoint (the knowledge, skills, attitudes, perspectives, etc. you want them to develop). The goal of a learning technology is to help bridge these two points - on its own or through the facilitation by teachers and tutors. Learning technologies succeed by inducing this change through particular learning mechanisms.

5.2.4. Learning Mechanisms

A learning mechanism answers the question, "Through what interactions or experiences are learners most likely to achieve the learning goals, according to theory?"

While not goals per se, learning mechanisms are sufficiently tied into the LCTD goals that they deserve brief treatment here. Learning mechanisms are the interactions and experiences that, according to learning theory, foster particular learning outcomes. As we discussed in the Introduction, learning is not deterministic, so learning outcomes are never guaranteed. That said, learning mechanisms are the paths by which perspectives, beliefs, physical capacity, knowledge, or skills can change, according to accumulated empirical evidence and theory building.

As you might guess, what learning mechanisms to consider is directly informed by learning goals. Consider the learning goal of "knowing multiplication tables." If "knowing" in this context means instantaneous recall (i.e., saying that 56 is the result of 8x7), an instructivist-informed learning mechanism would be passive exposure; students are repeatedly told (and perhaps repeat themselves) the answers to various multiplication problems and practice recalling those results on their own. A cognitivist-informed learning mechanism for this same "knowing" goal would repeat the information retrieval practice at timed intervals (e.g., every morning for 15 minutes). If "knowing" instead means being able to computecomputing a result (i.e., following a protocol to obtain an answer), a behaviorist-informed learning mechanism would demonstrate the procedure and give students repeated opportunities to practice that procedure, giving positive reinforcement for correct execution and negative reinforcement for incorrect execution. If "knowing multiplication" actually means understanding the operation as adding equal groups ("2 x 3" means "3 groups of 2" or "2 added to itself 3 times"), a learning mechanism informed by embodied cognition would be for students to manipulate blocks, beads, or other tangible objects, arranging and rearranging them in 3 groups of 2 and counting them by 2s to yield an answer. As this admittedly protracted example illustrates, learning mechanisms rely on precise learning goals, and even the same learning goal can be achieved through different learning mechanisms depending on your preference of theory. All the more reason for specificity as you define and redefine your learning goals (more in this in Chapter N). Learning mechanisms are usually formulated broadly and abstractly; they describe the steps that learners go through regardless of their context. The learning experience, on the other hand, is the *concrete implementation* of the learning mechanisms where designers include considerations such as culture, age, socio-economic background, resources, to design an experience that is tailored to a specific audience and context. We discuss below what this process might look like in the LCTD framework.

5.2.5. Design Goals: Learning Experience (LX) Goals and User Experience (UX) Goals

A design goal answers the question "What kind of experience do you want people to have with your technology?"

In other words, what (and how) do you want them to do, to think about, to notice, to talk about, to grapple with, to feel, or to perceive? Critically, learning technologies "work" by helping to create experiences that people can learn through. Of course, the main objective of any learning technology is to help people meet learning goals. However, a learning technology can't change a person *directly* in the way that a drill can change a piece of wood by making a hole in it. While technologies with this capability may exist in movies (yes, the Matrix, we know), there are no existing technologies that can change the connections between neurons in the brain so that someone suddenly knows Kung Fu[1]. *Learning technologies can provide experiences; through engaging in those experiences, people learn.* Different designs provide different types of experiences. The choices that one makes when designing a learning technology shape these experiences. We further separate these kinds of experiences into two kinds of goals: learning-experience (LX) goals and user-experience (UX) goals.

2.4.1. UX goals

UX goals concern the usability of a technology. Usability is arguably the primary concern of user-centered design methods. A usable design is one that is easy to understand and use while performing its function. *In other words, a usable design is one that minimizes the user's effort while maximizing output.* As with every technology, usability is an important factor in the design of learning technology. However, a singular focus on usability risks minimizing effort to the point of precluding learning. Learning is an inherently effortful process because it is about inducing change. You need to exert some effort - what we call *meaningful effort* - to learn. However, a badly designed learning technology might require a learner to remember too much information at once or to click 10 times to get to the desired function. This kind of *extraneous effort* is detrimental to learning because it overcharges the cognitive load and limits the capacity available for engaging in *meaningful effort*. UX goals help designers remove unnecessary distractors, minimize extraneous effort, and maximize meaningful effort.

2.4.2. LX Goals

M*eaningful effort* is captured by **LX goals**. The LX goals consist of the types of interactions with the technology and the cognitive and discursive processes involved in these interactions. It is only through providing these specific kinds

[1]Technologies such as Neuralink are in development to the other way; they aim to use neural impulses to produce change in objects outside the human brain, allowing people with paralysis to communicate and control devices.

of experiences that a learning technology is able to support people in meeting the learning goals. You can have a system with great usability (UX design) and no meaningful learning (i.e., all interactions are so easy that you never engage with the core tasks that would make you learn something). You can also have a system with lots of learning experiences that is so hard to use that you need tens of hours just to understand how the system works even before you start engaging with the learning activities themselves. That is why LX goals and UX goals are different dimensions that both are relevant to the design: having one does not imply having the other and, of course, we want to have both!

2.5. Alignment Between Learning Goals and Design Goals

For a learning technology to be effective, learning goals and design goals must be aligned. It isn't enough for a design to have stated learning goals, learning mechanisms, LX goals, and UX goals. If these goals are misaligned, the technology is less likely to have a meaningful impact. Misalignment occurs whenever one set of goals makes it harder to meet another set of goals. There are two kinds of misalignment that need to be considered. The first is a misalignment between the Learning Goals and the LX design goals, and the second is a misalignment between LX goals and UX goals.

Let's consider a made-up example[2] to illustrate the first kind of misalignment between LX goals and learning goals. This example is about a technology that is meant to help people learn to play the guitar. Imagine the technology's designer noticed that children often have trouble learning to play the guitar because they have trouble remembering the chords. For example, when asked to play a C chord, many children guess and play a G chord. So, the designer set out to create a technology with the following learning goal: **help children remember the mapping between chord names and the way they are played.**

To achieve this, she designed a device that clamps onto the neck of the guitar and makes it possible to play chords on the guitar by pressing a single button (Figure 5.1.). Each button is labeled with the chord it plays. There is a button for the G chord, for the C chord, for the A minor chord, and so on. When asked to play a G chord, a child no longer would need to guess, but could simply press the right button. The implicit LX goal in this design is that **students would no longer have to guess when asked to play a chord**, which, the designer stipulates, would increase their confidence and enjoyment in playing the guitar.

[2]The technology we describe is real (see Figure 1 below), but the scenario we are describing is made up.

Figure 5.1. A technology for playing guitar chords by pressing a single, labeled button.

Can you see how the learning goal (learn to map the chord name to the way it is played) and this LX goal might be misaligned? Yes, people using this tool will likely be able to play the correct chord when asked to play a C, G, or A minor chord, and it is likely that this would increase their confidence and enjoyment while playing the guitar. But, would that help them learn how to actually *map* the chord names to how they're played? If the tool was taken away, would you be confident that the person would still be able to play a G chord correctly when asked? Likely not.

Figure 5.2. The example misalignment of the LX of pressing a labeled button to the learning goal of mapping chords to how they're played.

What are some ways that the LX goals could become better aligned with the original learning goals? One way would be to change the learning goal. Instead of trying to meet the learning goal of helping children remember the mapping between chord names and the ways that they are played, the designer could reformulate her learning goal to be about increasing confidence and enjoyment of playing the guitar for beginners (see Figure 5.2.). This is a simple solution, but in most cases, a designer of learning technologies is invested in meeting the original learning goal and is not willing to change it so easily. If she were unwilling to change the learning goal, then the tool itself needs to be changed.

Let's continue with this example to demonstrate the second type of misalignment between UX goals and LX goals. During testing, the designer noticed that some of the children were trying to see which of the strings was being pressed when they pressed one of the buttons but that the black plastic material (Figure 5.1)

made this impossible. She realized that if the plastic were clear, then children would be able to see which strings were being pressed when they were asked to play a specific chord. She remade a prototype with clear plastic and reformulated a new LX goal: **when asked to play a specific chord, beginners would press the button and inspect how and where the strings were being pressed**. Thus, she decided to replace the black plastic casing with a transparent one that allows to see the link between the buttons and the chords. By giving them simple songs to play (like "Let It Be" as shown in Figure 5.3.) students could experience the satisfaction of playing a real song while also being able to inspect how the chords in the song were being played.

Figure 5.3. The clear plastic allows users to inspect where each string is pressed.

Now, the LX goal and learning goals seemed more aligned. However, when the designer tested her new design with children, she found that they still weren't learning the chords. They were more interested in playing songs by pressing the buttons than inspecting how the chords were being played. And even those students who did inspect how the chords were being played could not remember how to play them once the tool was removed from the guitar.

Figure 5.4. Example misalignments of UX to LG and of LX to LG, here for the

clear prototype of pressing a button to play a chord.

One of the reasons for this failure are misalignments between UX goals and LX goals (Figure 5.4.), just like in the first prototype. The tool simply made it too easy to play chords, which removed all of the effort from the learning experience. Meeting the learning goal, in this case, requires trying, failing, and trying again, but in ways that encourage the student to continue and to experience progress. However, with the new version of the tool there was little room for failing to hit the right chord. For a G chord, just press the blue button.

A second misalignment between UX goals and LX goals is that there was nothing in the tool design that required students to investigate the chord being pressed; one could play the chords without ever looking more closely beyond the buttons themselves. But also, learning to play chords on the guitar is not simply a matter of looking at which strings are being pressed when a particular chord is being played. There is a great deal of muscle memory involved as well. It is important that a beginner learns what it feels like to play a chord, and associates that feeling with the name of the chord. Thus, there is another fundamental misalignment between LX goals and the learning goal, in addition to the misalignments between UX goals and LX goals.

Figure 5.5. This technology provides lights as indicators of where to press to produce each chord, requiring users to physically produce the chord themselves.

To address these misalignments, the designer created a radically new prototype (Figure 5.5.). Instead of designing a tool that could be placed on the neck of a guitar, she designed a new kind of guitar neck with LED lights built into it. Buttons with the different chord names were built into the headstock of the guitar such that when a button was pressed, the LEDs on the fretboard lit up, indicating where to press each of the strings. Now, pressing a button would not play the chord, which addressed the misalignment between UX and LX goals. And in order to play the chord, the beginner actually needed to press the strings themselves, addressing the misalignment between the LX goals and the learning goals. Notice that this is the first prototype that *considered a learning mechanism* for how people learn to play guitar chords (Figure 5.6.): Practicing finger placement to develop muscle memory. As a result, when testing with beginners, the designer found that after a short time practicing with this tool, beginners were able to play many chords even when the lights were turned off.

Figure 5.6. Mapping the goals of the illuminated fingerboard design.

We present a simple diagram of a well-aligned set of goals in Figure 5.7. This alignment serves as the backbone of learning-centered technology design, and we will continue to build on this backbone throughout the rest of the book.

Figure 5.7. Elements and relationships of the LCTD framework.