COACHING *for* LEADERS

Ultralearning: Master Hard Skills, Outsmart the Competition, and Accelerate Your Career

by Scott Young

- Featured on episode 451
- Purchasing this book? Support the show by using the Amazon link inside our book library.

Dave's Reading Highlights

Then I got a message from a friend: "You're on the front page of Reddit, you know." The internet had found my project, and it was generating quite a discussion. Some liked the idea but doubted its usefulness: "It's sad that employers won't really treat this in the same way as a degree, even if he has the same amount (or more) knowledge than a graduate does." One user claiming to be the head of R&D for a software company disagreed: "This is the type of person I want. I really do not care if you have a degree or not." The debate raged. Had I actually done it or not? Would I be able to get a job as a programmer after this? Why try to do this in a year? Was I crazy?

What if, instead of forming a group of English-speaking friends and struggling to pop out of that bubble once my French was good enough, I emulated Benny Lewis and dived straight into immersion from the very first day? How much better could I be, if as in my MIT Challenge, I held nothing back and optimized everything around learning a new language as intensely and effectively as possible?

Every ultralearner I encountered was unique. Some, like Tamu, preferred punishing, full-time schedules to meet harsh, self-imposed deadlines. Others, like Jaunzeikare, managed their projects on the

side while maintaining full-time jobs and work obligations. Some aimed at the recognizable benchmarks of standardized exams, formal curricula, and winning competitions. Others designed projects that defied comparison. Some specialized, focusing exclusively on languages or programming. Others desired to be true polymaths, picking up a highly varied set of skills.

Despite their idiosyncrasies, the ultralearners had a lot of shared traits. They usually worked alone, often toiling for months and years without much more than a blog entry to announce their efforts. Their interests tended toward obsession. They were aggressive about optimizing their strategies, fiercely debating the merits of esoteric concepts such as interleaving practice, leech thresholds, or keyword mnemonics. Above all, they cared about learning. Their motivation to learn pushed them to tackle intense projects, even if it often came at the sacrifice of credentials or conformity.

ULTRALEARNING: A strategy for acquiring skills and knowledge that is both self-directed and intense. First, ultralearning is a strategy. A strategy is not the only solution to a given problem, but it may be a good one. Strategies also tend to be well suited for certain situations and not others, so using them is a choice, not a commandment. Second, ultralearning is self-directed. It's about how you make decisions about what to learn and why. It's possible to be a completely self-directed learner and still decide that attending a particular school is the best way to learn something. Similarly, you could "teach yourself" something on your own by mindlessly following the steps outlined in a textbook. Self-direction is about who is in the driver's seat for the project, not about where it takes place. Finally, ultralearning is intense. All of the ultralearners I met took unusual steps to maximize their effectiveness in learning. Fearlessly attempting to speak a new language you've just started to practice, systematically drilling tens of thousands of trivia questions, and iterating through art again and again until it is perfect is hard mental work. It can feel as though your mind is at its limit. The opposite of this is learning optimized for fun or convenience: choosing a language-learning app because it's entertaining, passively watching trivia show reruns on television so you don't feel stupid, or dabbling

instead of serious practice. An intense method might also produce a pleasurable state of flow, in which the experience of challenge absorbs your focus and you lose track of time. However, with ultralearning, deeply and effectively learning things is always the main priority.

Professional success, however, was rarely the thing that motivated the ultralearners I met—including those who ended up making the most money from their new skills. Instead it was a compelling vision of what they wanted to do, a deep curiosity, or even the challenge itself that drove them forward.

There's an added benefit to ultralearning that transcends even the skills one learns with it. Doing hard things, particularly things that involve learning something new, stretches your selfconception. It gives you confidence that you might be able to do things that you couldn't do before. My feeling after my MIT Challenge wasn't just a deepened interest in math and computer science but an expansion in possibility: If I could do this, what else could I do that I was hesitant to try before? Learning, at its core, is a broadening of horizons, of seeing things that were previously invisible and of recognizing capabilities within yourself that you didn't know existed. I see no higher justification for pursuing the intense and devoted efforts of the ultralearners I've described than this expansion of what is possible. What could you learn if you took the right approach to make it successful? Who could you become?

De Montebello's story illustrates that it's possible to decide to become an ultralearner. But ultralearning isn't a cookie-cutter method. Every project is unique, and so are the methods needed to master it. The uniqueness of ultralearning projects is one of the elements that ties them all together. If ultralearning could be bottled or standardized, it would simply be an intense form of structured education. What makes ultralearning interesting is also what makes it hard to boil down into step-by-step formulas.

There are nine universal principles that underlie the ultralearning projects described so far. Each embodies a particular aspect of

successful learning, and I describe how ultralearners maximize the effectiveness of the principle through the choices they make in their projects. They are: 1. Metalearning: First Draw a Map. Start by learning how to learn the subject or skill you want to tackle. Discover how to do good research and how to draw on your past competencies to learn new skills more easily. 2. Focus: Sharpen Your Knife. Cultivate the ability to concentrate. Carve out chunks of time when you can focus on learning, and make it easy to just do it. 3. Directness: Go Straight Ahead. Learn by doing the thing you want to become good at. Don't trade it off for other tasks, just because those are more convenient or comfortable. 4. Drill: Attack Your Weakest Point. Be ruthless in improving your weakest points. Break down complex skills into small parts; then master those parts and build them back together again. 5. Retrieval: Test to Learn. Testing isn't simply a way of assessing knowledge but a way of creating it. Test yourself before you feel confident, and push yourself to actively recall information rather than passively review it. 6. Feedback: Don't Dodge the Punches. Feedback is harsh and uncomfortable. Know how to use it without letting your ego get in the way. Extract the signal from the noise, so you know what to pay attention to and what to ignore. 7. Retention: Don't Fill a Leaky Bucket. Understand what you forget and why. Learn to remember things not just for now but forever. 8. Intuition: Dig Deep Before Building Up. Develop your intuition through play and exploration of concepts and skills. Understand how understanding works, and don't recourse to cheap tricks of memorization to avoid deeply knowing things. 9. Experimentation: Explore Outside Your Comfort Zone. All of these principles are only starting points. True mastery comes not just from following the path trodden by others but from exploring possibilities they haven't yet imagined.

I find it useful to break down metalearning research that you do for a specific project into three questions: "Why?," "What?," and "How? " "Why? " refers to understanding your motivation to learn. If you know exactly why you want to learn a skill or subject, you can save a lot of time by focusing your project on exactly what matters most to you. "What? " refers to the knowledge and abilities you'll need to acquire in order to be successful. Breaking things down into concepts, facts,

and procedures can enable you to map out what obstacles you'll face and how best to overcome them. "How? " refers to the resources, environment, and methods you'll use when learning. Making careful choices here can make a big difference in your overall effectiveness. With these three questions in mind, let's take a look at each of them and how you can draw your map.

Reaching out and setting up a meeting with an expert isn't hard, either, but it's a step many people shy away from. Many people, particularly the introverts among us, recoil at the idea of reaching out to a stranger to ask for advice. They worry that they'll be rejected, ignored, or even yelled at for presuming to take up a person's time. The truth is, however, that this rarely happens. Most experts are more than willing to offer advice and are flattered by the thought that someone wants to learn from their experience. The key is to write a simple, to-the-point email, explaining why you're reaching out to them and asking if they could spare fifteen minutes to answer some simple questions. Make the email concise and nonthreatening. Don't ask for more than fifteen minutes or for ongoing mentorship. Though some experts will be happy to help you in those ways, it's not good form to ask for too much in the first email.

If I'm trying to learn a nonacademic subject or a professional skill, I'll probably instead do online searches for people who have previously learned that skill or use the Expert Interview Method to focus on resources available for mastering that subject. An hour spent searching online for almost any skill should turn up courses, articles, and recommendations for how to learn it. Investing the time here can have incredible benefits because the quality of the materials you use can create orders-of-magnitude differences in your effectiveness. Even if you're eager to start learning right away, investing a few hours now can save you dozens or hundreds later on.

The Emphasize/Exclude Method involves first finding areas of study that align with the goals you identified in the first part of your research. If you're learning French with the idea of going to Paris for two weeks and speaking in shops and restaurants, I would focus a lot more on pronunciation than being able to spell correctly. If you're

learning programming solely to make your own app, I'd focus on the inner workings of app development more than theories of computation.

The second part of the Emphasize/Exclude Method is to omit or delay elements of your benchmarked curriculum that don't align with your goals. For example, one common recommendation for learning Mandarin Chinese, advocated by people such as the renowned linguist and Sinologist Victor Mair, is to focus on learning to speak before you try to read characters. This isn't the only route available, but if your main goal is to speak, then this path to fluency might be more effective.

A good rule of thumb is that you should invest approximately 10 percent of your total expected learning time into research prior to starting.

In the realm of great intellectual accomplishments an ability to focus quickly and deeply is nearly ubiquitous. Albert Einstein focused so intensely during his formulation of the general theory of relativity that he developed stomach problems. The mathematician Paul Erd ős was a heavy user of amphetamines to increase his capacity for focus. When a friend bet him that he could not give them up, even for a short time, he did manage to do so. Later, however, he complained that the only result had been that mathematics as a whole was set back a month in his unfocused absence

However remarkable this is, I'm more interested in the kind of focus that Somerville seemed to possess. How can one in an environment such as hers, with constant distractions, little social support, and continuous obligations, manage to focus long enough not only to learn an impressive breadth of subjects, but to such depths that the French mathematician Siméon Poisson once remarked that "there were not twenty men in France who could read [her] book"? How did Somerville become so good at focusing? What can we glean from her strategies in getting difficult mental work done in less-than-ideal conditions? The struggles with focus that people have generally come in three broad varieties: starting, sustaining, and optimizing the

quality of one's focus. Ultralearners are relentless in coming up with solutions to handle these three problems, which form the basis of an ability to focus well and learn deeply.

Eventually, you may reach Mary Somerville's level of focus, one that she could activate on a moment-to-moment basis, making a decision as to whether she had time to spare. Despite her formidable capacity for focus, it seems that even Somerville would deliberately block out time for the study of particular subjects. Therefore it was a conscious habit, not merely spontaneous studying, that enabled her many successes. For myself, I find that some learning activities are so intrinsically interesting that I can focus on them for a long time without pressure. I generally had no problem watching lectures during the MIT Challenge, for instance. Other tasks, however, required the five-minute rule for me to get past my desire to procrastinate. If I had to scan and upload my files, they'd often build up in a pile before I would finally tackle them. Don't ever feel bad if you have to back up a stage, either; you cannot control your aversions or tendency to distraction, but with practice you can lessen their impact.

My own thought is that a flow state is not impossible during ultralearning. Many cognitive activities associated with learning are in the range of difficulty that makes flow possible or even likely. However, I also agree with Ericsson that learning often involves entering into situations in which the difficulty makes flow impossible. Additionally, the self-consciousness that is absent in flow may need to be present in both ultralearning and deliberate practice, as you need to consciously adjust your approach.

Directness is the idea of learning being tied closely to the situation or context you want to use it in. In Jaiswal's case, when he wanted to get enough architectural skill that firms would hire him, he opted to build a portfolio using the software those firms used and design in the style those firms practiced. There are many routes to self-education, but most of them aren't very direct. In contrast to Jaiswal, another architect I spoke with aimed to improve his employability by deepening his knowledge of design theories. Though that might have

been interesting and fun, it was disconnected from the actual skills he would be using in entry-level work.

We want to speak a language but try to learn mostly by playing on fun apps, rather than conversing with actual people. We want to work on collaborative, professional programs but mostly code scripts in isolation. We want to become great speakers, so we buy a book on communication, rather than practice presenting. In all these cases the problem is the same: directly learning the thing we want feels too uncomfortable, boring, or frustrating, so we settle for some book, lecture, or app, hoping it will eventually make us better at the real thing.

The opposite of this is the approach so often favored in more traditional classroom-style learning: studying facts, concepts, and skills in a way that is removed from how those things will eventually be applied: mastering formulas before you understand the problem they're trying to solve; memorizing the vocabulary of a language because it's written on a list, not because you want to use it; solving highly idealized problems that you'll never see again after graduation.

During the MIT Challenge, I recognized that the most important resource for being able to eventually pass the classes wasn't having access to recorded lectures, it was having access to problem sets. Yet, in the years since this project, when I am asked for help by students, they often decry the absence of lecture videos from some classes, only rarely complaining about incomplete or insufficient problem sets. This makes me think that most students view sitting and listening to a lecture as the main way that they learn the material, with doing problems that look substantially similar to those on the final exam as being a superficial check on their knowledge. Though first covering the material is often essential to begin doing practice, the principle of directness asserts that it's actually while doing the thing you want to get good at when much of learning takes place. The exceptions to this rule are rarer than they may first appear, and therefore directness has been a thorny problem in the side of education for over a century.

The twin challenge of directness is that sometimes the exact situation in which you want to use the skill isn't available for easy practice. Even if you can go straight into learning by doing, this approach is often more intense and uncomfortable than passively watching lecture videos or playing around with a fun app. If you don't pay attention to directness, therefore, it's very easy to slip into lousy learning strategies.

Transfer has been called the "Holy Grail of education." It happens when you learn something in one context, say in a classroom, and are able to use it in another context, say in real life. Although this may sound technical, transfer really embodies something we expect of almost all learning efforts—that we'll be able to use something we study in one situation and apply it to a new situation. Anything less than this is hard to describe as learning at all. Unfortunately, transfer is also something that, despite more than a century of intense work and research, has largely failed to occur in formal education. The psychologist Robert Haskell has said in his excellent coverage of the vast literature on transfer in learning, "Despite the importance of transfer of learning, research findings over the past nine decades clearly show that as individuals, and as educational institutions, we have failed to achieve transfer of learning on any significant level." He later added, "Without exaggeration, it's an education scandal."

Providing multiple examples seems to aid transfer a bit, yet the cognitive science researcher Michelene Chi noted that "in almost all the empirical work to date, on the role of example solutions, a student who has studied examples often cannot solve problems that deviate slightly from the example solution." In his book The Unschooled Mind: How Children Think and How Schools Should Teach, the developmental psychologist Howard Gardner pointed to the body of evidence showing that even "students who receive honors grades in college-level physics courses are frequently unable to solve basic problems and questions encountered in a form slightly different from that on which they have been formally instructed and tested." Nor has this failure of transfer been limited to schools. Corporate training also suffers, with the former Times Mirror Training Group chairman John H.

Zenger writing “Researchers who rigorously evaluate training have said that demonstrable changes following training are hard to find.”

So what explains the disconnect? Why have educational institutions struggled to demonstrate significant transfer, if transfer is something we all need to function in the world? Haskell suggests that a major reason is that transfer tends to be harder when our knowledge is more limited. As we develop more knowledge and skill in an area, they become more flexible and easier to apply outside the narrow contexts in which they were learned. However, I’d like to add my own hypothesis as an explanation for the transfer problem: most formal learning is woefully indirect.

Given the problem of transfer and the importance of learning directly, let’s look at some of the ways that this is managed in different ultralearning projects. The simplest way to be direct is to learn by doing. Whenever possible, if you can spend a good portion of your learning time just doing the thing you want to get better at, the problem of directness will likely go away. If this isn’t possible, you may need to create an artificial project or environment to test your skills. What matters most here is that the cognitive features of the skill you’re trying to master and the way you practice it be substantially similar.

Many ultralearners who have specialized in a smaller subset of fields are masters at transfer; no doubt this is largely due to their depth of knowledge, which makes transfer easier to accomplish.

Tactic 1: Project-Based Learning Many ultralearners opt for projects rather than classes to learn the skills they need. The rationale is simple: if you organize your learning around producing something, you’re guaranteed to at least learn how to produce that thing. If you take classes, you may spend a lot of time taking notes and reading but not achieve your goal.

Tactic 2: Immersive Learning Immersion is the process of surrounding yourself with the target environment in which the skill is practiced. This has the advantage of requiring much larger amounts of practice

than would be typical, as well as exposing you to a fuller range of situations in which the skill applies. Learning a language is the canonical example of where immersion works. By immersing yourself in an environment where a language is spoken, not only do you guarantee that you'll end up practicing the language a lot more than you would otherwise (since you have no choice), but you also face a broader diversity of situations that require learning new words and phrases. However, language learning is not the only place where you can apply immersion to learn more. Joining communities of people who are actively engaged in learning can have a similar impact, since it encourages constant exposure to new ideas and challenges. For example, novice programmers might join open-source projects to expose themselves to new coding challenges.

Tactic 3: The Flight Simulator Method Immersion and projects are great, but for many skills there's no way to actually practice the skill directly. For skills such as piloting a plane or performing surgery, it's not even legal to practice them in a real situation until you've already invested considerable time into training. How can you overcome this?

Tactic 4: The Overkill Approach The last method I've found for enhancing directness is to increase the challenge, so that the skill level required is wholly contained within the goal that is set. Tristan de Montebello, when preparing to compete in the World Championship of Public Speaking, pushed to speak at middle schools, giving early versions of his talk. His feeling was that the feedback he received at Toastmasters clubs might be too soft or congratulatory to cut deep at what worked and didn't work in his speech. Middle school students, in contrast, would be merciless. If a joke he said wasn't funny or his delivery was boring or cheesy, he would be able to tell immediately from their faces what needed to be reworked. The overkill approach is to put yourself into an environment where the demands are going to be extremely high, so you're unlikely to miss any important lessons or feedback.

Another friend of mine decided to exhibit her photography as a means of pushing her skills and talent. Deciding in advance that your work will be viewable publicly alters your approach to learning and

will gear you toward performance in the desired domain, rather than just checking off boxes of facts learned.