




– Working Draft –

Note: If you are reading this as a PDF,
consider switching to the [Google Doc](#)
which may contain new updates.

Authored by [Justin Skycak](#)

Updated 9 October 2025

Authored by Justin Skycak

Copyright © 2025 Justin Skycak

First edition (working draft), updated 9 October 2025.

All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the author, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law. For permission requests, contact the author through the website below.

www.justinmath.com/contact

Contents

Preface.....	13
Chapter 1. Consistency.....	15
▶ You're Not Lazy, You Just Lack a Habit.....	15
▶ Don't Have a Passion? Go Create One.....	17
▶ Make the Habit Easily Repeatable.....	18
▶ The Hardest Part is Just Getting Started.....	19
▶ How to Overcome the Suck.....	20
▶ If You Struggle to Train Consistently, Do It Immediately After Waking Up.....	22
▶ Training Sessions Should be Short and Frequent as Opposed to Long and Sparse.....	23
▶ A Little Extra Consistency × A Little Extra Time = A Massive Increase in Volume & Progress.....	25
▶ Don't Overreact to Bad Days.....	25
▶ Seek Virtuous Cycles.....	27
▶ Protect The Habit.....	27
Chapter 2. Skills.....	29
▶ Hardcore Skills Are Always The Answer.....	29

▶ The “Alien-Level Skills” Hack.....	30
▶ The Importance of Having Your Prerequisites In Place.....	33
▶ Most People Don’t Realize They Can Systematically Climb Skill Trees.....	34
▶ The Least Efficient Learning Strategy.....	36
▶ The Skills You Need Should Sit Several Layers Deep.....	36
▶ Your Missing Foundations Will Wait For You.....	38
▶ Consuming is Only Helpful Insofar as it Enables You to Produce.....	39
▶ Actively Doing is the Key to Alpha.....	40
▶ Everything Matters.....	40
▶ You Want Exciting Opportunities? Learn Math and Coding.....	41
▶ If You're Not Both Technical and a Domain Expert, Then You're Underpowered...43	
▶ Domain Expertise, Math, Coding, Communication.....	44
▶ What Math To Learn for Skill Stacking.....	45
▶ Failure Modes in People Who Develop Math Skills but Don’t Capitalize On Them via Coding.....	47
▶ There is Always a Bad Excuse to	

Rationalize Not Skilling Up..... 49

Chapter 3. Discipline..... 51

- ▶ The Magic You're Looking For is in the Full-Assed Effort You're Avoiding..... 51
- ▶ Don't Put Off Training Because You Think Your Goals are Too Far Away..... 52
- ▶ At Some Point, Doing the Hard Thing Becomes Easier Than Making the Hard Thing Easier..... 54
- ▶ How to Cultivate Discipline.....55
- ▶ Keep Your Hands On The Boulder..... 57
- ▶ Just Do The F*cking Work..... 58

Chapter 4. The Grind..... 61

- ▶ Upskilling is Hard and That's a Good Thing..... 61
- ▶ The Most Superior Form of Training..... 61
- ▶ Outsized Success Requires Outsized Work..... 64
- ▶ Transformation Is Discomforting..... 65
- ▶ Enjoyment is a Second-Order Optimization..... 67
- ▶ Ability is Built, Not Unlocked..... 68
- ▶ What Max-Efficiency Training Feels Like.....69
- ▶ The Necessity of Grinding Through Concrete Examples Before Jumping Up a Level of Abstraction..... 71

▶ Be Willing to Do Tedious Work.....	73
▶ Don't Undervalue Turning Up the Dial on Your Grind, but Don't Overvalue the Last Turn.....	74
▶ When More Volume Equals More Progress.....	76
▶ Failure Is NOT the Key to Success.....	77
▶ Don't Drown Yourself in the Deep End.....	78
▶ Focus Less on Feelings and More on Measurable Progress.....	79
▶ The Problem with Overly Difficult Problems.....	81
▶ It's Not Just About Attempting Hard Things, It's Also About Successfully Achieving Them.....	82
▶ It Always Becomes a Battle of Willpower By the End.....	83
▶ The Cycle You Want To Be In.....	84
Chapter 5. The Journey.....	87
▶ Don't Get Hung Up on Youth Competitions.....	87
▶ The 3 Stages of Talent Development.....	89
▶ There Are No Shortcuts in Talent Development.....	91
▶ If You're Making Silly Mistakes Then You Need More Practice.....	94

▶ No Train, No Gain.....	95
▶ Why You Should Push Yourself.....	96
▶ Keep Your Foot On The Gas.....	96
▶ You Are a Car.....	97
▶ What to Do When You Hit a Ceiling.....	97
▶ Learn to Love the Games That You Were Built to Win.....	99
▶ Compound Hard Work and Luck.....	101
▶ Get Yourself In A Position Where You Can Eat Risk.....	102
▶ Tie Your Comfort to Real, Tangible Value.....	102
▶ If You Can't Find a Job That Really Excites You.....	104
▶ How to Allocate Your Bandwidth While Searching for Your Mission.....	104
▶ Repetition Can Lead to Expertise, Expertise Can Lead to Variety.....	106
Chapter 6. The Team.....	107
▶ If You're Asking Someone to Be Your Mentor then You're Doing it Wrong.....	107
▶ Put Pressure on Your Boss to Come Up with More Work For You.....	109
▶ Get On the Right Team.....	110
▶ Competition as a Means of Collaboration.....	111

- ▶ Your Goal is NOT to Prove You're Smart,
it's to Make Problems Go Away..... 113
- ▶ Make Your First Impression On a
Contribution, Not a Critique..... 114
- ▶ Never Come Up Empty-Handed..... 114
- ▶ You Need a Berserker At The Helm.....116

Chapter 7. The Mission..... 117

- ▶ Selecting a Good Problem to Work On..... 117
- ▶ You Have To Work Really Damn Hard to
Figure Out What Fulfills You..... 121
- ▶ The “Progress Equals Pressure” Formula..... 123
- ▶ Love What You Do.....123
- ▶ Be a Builder, Not Just a Fighter..... 125
- ▶ Build Where Building Creates More
Opportunities to Build..... 125

Chapter 8. Motivation.....127

- ▶ Disinterest is Often Just Overwhelm..... 127
- ▶ Why Extrinsic Motivation Matters..... 128
- ▶ How to Become a Super-Producer..... 129
- ▶ How The Highest Performers Sustain a
Massive Workload..... 130
- ▶ Overcoming the Paradox of Serious
Training..... 131
- ▶ The Joy of Upskilling is Like The Joy of

Exercise.....	133
▶ How Taxing Work Becomes Fun.....	135
▶ Why Try When Others Could Do Better?.....	137
Chapter 9. Learning.....	139
▶ The Greatest Educational Life Hack: Learning Ahead of Time.....	139
▶ What's the Point of Being Way Ahead? It Buys You Time.....	144
▶ When Does the Learning Happen?.....	147
▶ There is No Such Thing as Low-Effort Learning.....	148
▶ The Greatest Breakthrough in the Science of Learning Over the Last Century.....	150
▶ "Following Along" Versus Learning.....	153
▶ The Need to Practice Retrieving From Memory.....	158
▶ The Vicious Cycle of Forgetting.....	159
▶ Recall First, Reason Second.....	164
▶ The Vicious Cycle of Context Overload.....	165
▶ Prereq Yo' Self Before You Wreck Yo' Self.....	166
▶ Filling In Your Foundations is the Difference Between Omitted Steps Being Minor Potholes vs Uncrossable Chasms.....	167
▶ Plan Your Broad-Stroke Journey Top-Down, but Carry Out the Granular	

Steps Bottom-Up.....	169
▶ The Efficient Learning Loop.....	170
▶ Don't Bloat the Feedback Loop.....	172
▶ Some Pitfalls to Watch Out For When Learning From Projects.....	173
▶ Review Should Feel Challenging.....	174
▶ Learn Like You Lift.....	176
▶ Schooling Versus Talent Development.....	178
▶ Learning Doesn't Have to Be Synchronized for Camaraderie to Occur.....	182
▶ A Sanity Check for Effective Study Techniques.....	183
Chapter 10. Expertise.....	185
▶ The Driving Force Behind Expertise is Long-Term Memory.....	185
▶ Learning is Memory.....	187
▶ Learning is About Bridge-Building, Not Jumping.....	190
▶ It's All About Domain Knowledge.....	193
▶ Your Brain Should Surprise You.....	194
▶ Turn The Magical Into The Mechanical.....	195

Preface

From summer 2024-25 I wrote numerous scattered pieces on upskilling while answering questions and engaging in general discussion about *The Math Academy Way*. Eventually I found myself repeating common themes, analogies, and punch-lines, and I compiled them into this short book.

If you want to improve (or initiate) your upskilling regimen, then this book is a good place to start. It's short, easy reading, addresses the big ideas, and gets you in the right mindset. If you want a complete, comprehensive, academic understanding of everything down to the nuts and bolts, with hundreds of citations, then I would highly recommend to go on and read *The Math Academy Way* afterwards.

Chapter 1. Consistency

► You're Not Lazy, You Just Lack a Habit

If you're struggling to stick with demanding forms of practice, then temporarily forget about efficiency and just build a habit with some less effective but more enjoyable form of practice. Although the ultimate goal is to train efficiently and get the largest possible performance gain out of your limited training time, that's going to feel taxing, and you might not want to work that hard at first – not because you're inherently lazy, but because you haven't built a habit. You eventually want to get to the point where performance improvement is your primary focus and fun is a second-order optimization, but it's okay to optimize for fun at the beginning to help you build a habit.

Consider strength training, for example. If you're just starting out, but you're not looking forward to lifting heavy-ass weights, then that's okay! You don't have to lift them yet. Your #1 focus should be just getting your ass into the gym and doing some kind of activity that

loosely qualifies as exercise. After a week of, say, shooting hoops, you might be motivated to try some bodyweight exercises – and then the following week maybe some light weightlifting, and maybe the week after that you'll be ready to challenge yourself by putting some serious weight on the bar.

It's the same way with anything else – for instance, learning math. If you don't feel motivated to solve a high volume of problems that are challenging enough to make you mentally sweat, that's okay. You can start off watching math edutainment videos and exploring Wikipedia. The next week, maybe try to solve some “math meme” problems each day (and look at the comments to check if you got it right), and the following week, maybe work out some easy arithmetic or algebra problems each day (stuff that you still remember fairly well but haven't done in a while). By that point you've gotten yourself into the metaphorical weight room, you've done some light lifting, and you're ready to put some serious weight on the bar. And that's when you start working through a structured curriculum that engages you in taxing practice to pack the maximum possible learning into your practice time.

Once you get to that point, you've built a habit, and you need to do everything in your power to maintain it. If you want to take a day off, just do a quick 10 minutes –

something that feels negligible but keeps the habit going. The habit is a psychological force field that protects you from all sorts of negative feelings that try to dissuade you from training.

In summary: you're not lazy, you just lack a habit. So start simple, whatever gets the ball rolling. (But if you know this and you're still unwilling to build a habit... then yeah, you're lazy.)

► Don't Have a Passion? Go Create One.

Be disciplined, set up a habit, compound compound compound. Develop a relationship with it, put in extra time when you're bored, come to it when you need an emotional outlet. Trust that it will grow on you and seep into your identity as the days turn into weeks, months, years of practicing and developing serious expertise.

It's just like developing a close human relationship. You might not have a strong bond with the person initially, but you get along "enough" at the beginning, and over time you get to know each other so well, you go through so much shit together, that you are inseparable.

And don't worry about the long-term too much. The person who became your best friend, you probably didn't know it the day you met them. You probably got to know them better and better, week by week, month by month, until at some point you realized you couldn't imagine life without them. It's the same way with creating a passion. If you keep on making short-term progress then the long-term will sort itself out.

► Make the Habit Easily Repeatable

When you start a new routine, don't make it such a "big thing" that you do it one day and dread doing it the next day. You know what happens to people who start off their New Year's weight loss resolution with 3 hours at the gym every day? They come for one day and then don't come back! So don't do that. Instead, start out with a volume of work that's small enough that you don't dread doing it again the next day.

At the beginning, it doesn't matter if the volume of daily work is too small to achieve your long-term goals in the timeframe you want. Eventually, as you build up a habit and your mind and body adapt to whatever it is you're doing, it will feel easier to ramp up the volume of work until you're moving at a pace that puts you on track to accomplish your long-term goals.

So don't worry about the total volume of work at the beginning. Just focus on consistency. As the habit sets in and you adapt, you'll find it easier to increase your volume of work. And as the habit settles into your identity, you'll actually want to increase that volume of work.

► The Hardest Part is Just Getting Started

Most skills can be trained. But serious training usually isn't pleasant, so most people don't do it. That doesn't mean there's little benefit to training. It just means lots of people aren't willing to put in the work to capitalize on said benefit.

But the thing is, the hardest part is always just getting started. If you suck at writing, then just sit down and write for 15 minutes each day. It might be unpleasant fishing for cohesive thoughts in your brain stew, pulling them out, and translating them into text. But that doesn't mean it will always feel that way.

As you practice again and again, it will feel easier over time. And as it feels easier you'll free up more and more mental bandwidth to notice areas for improvement. And you'll get better. Will you become a

world-class writer? Who knows. Probably not. But will you open up opportunities that were previously closed to you? Probably.

And who's to say the training will actually be unpleasant? When you're not looking forward to training, you might think that negative feeling is going to intensify during your (physical or intellectual) workout, but often it just dissipates and you feel great 5 minutes in. Procrastinating builds up the dread but just getting started often makes it dissipate.

► How to Overcome the Suck

It gets easier to motivate yourself to work out once you start seeing results (e.g., clothes fit different, you can lift heavier weights or do calisthenics moves that you previously couldn't).

But what really, really makes a difference is when other people start noticing your results, complimenting you, sometimes even asking for your advice. You hear "damn you're strong/buff, what's your workout, any advice for me" enough times and it just seeps into your identity that you are a strong/buff person. And it's very easy to find motivation to do training activities that

help you maintain what you perceive to be a part of your identity.

By the way, this applies to all skill domains, not just physical exercise. You can do this with math, coding, music, whatever the heck you want.

But the hardest part is the beginning, before you've kicked off the virtuous cycle. Unless you have a ridiculous amount of intrinsic motivation, it's just a lot of suck at the beginning.

The way you overcome the suck is by setting a reasonable training schedule that you can actually stick to. At the beginning you need to focus on consistency, not training volume. You need to build a habit. As you build that habit, you'll not only find it easier to stick to your schedule, but you'll also find it easier to increase your volume, eventually up to a level that will put you on pace to reach your goals in a reasonable timeframe.

► If You Struggle to Train Consistently, Do It Immediately After Waking Up

If you struggle to train consistently, then the #1 remedy is to do it immediately after waking up. Here are several reasons why morning training works so well:

First, it shields you from those unexpected events that demolish whatever schedule you had planned for the day. Those unexpected events happen to everyone, even people who manage to train consistently, and they are not an excuse for inconsistency. If you're serious about training, you have to build chaos tolerance into your training schedule.

Second, there's no time to even consider skipping the training session. By the time you shake off the initial grogginess and become fully conscious of what you're doing, you've already started training. "I couldn't make myself do it" is not an excuse for inconsistency. If your conscious brain steers you away from training, then get the wheels rolling before your conscious brain even has a chance to take the wheel.

Third, there's a clear habit trigger and reward: waking up means it's time to train, and after you train you get your shower, meal, whatever other morning amenities

you enjoy, and you get to feel good about yourself all day. And if you skip the training, then there's the punishment of feeling bad about yourself all day. It's basically operant conditioning. You have every right, and every responsibility, to be blatantly and maximally manipulative with yourself.

► Training Sessions Should be Short and Frequent as Opposed to Long and Sparse

Suppose you're budgeting 3 hours per week to train. If circumstances allow, it would be better to train for 30 minutes six days per week, as opposed to 90 minutes twice a week. Below are several reasons why.

First, you want to form a habit. The more consistently you train, the more it will become a habit that you naturally do each day without thinking, just like (hopefully!) taking a shower and brushing your teeth. That habit is what will carry you through the long game once the initial adrenaline wears off.

Second, you want to operate at peak productivity during your session. It's easy to maintain high levels of focus and intensity throughout a short 30-minute session. However, during a long 90-minute session,

you'll become significantly less productive as fatigue sets in.

Third, you want to minimize the amount you forget between sessions. When you have multi-day gaps between sessions, you'll have to spend more time revisiting material you covered previously. (Just ask any teacher how much their students forget over weekends, and how much valuable class time they have to spend on Monday re-teaching the things that they covered on Thursday and Friday.)

However, there are also some caveats to consider.

Whenever you switch to a different cognitive activity, it may take a few minutes for your brain to spin up on the new context. Likewise, whenever you switch to a new physical activity, it may take a few minutes (or longer) to get set up. This is called “context switching cost,” and if you make your sessions too short, then the proportion of training time that is wasted on context switching will outweigh the other benefits of daily practice. While it's good to spread out your practice, each session should be long enough that the context-switching cost is proportionally negligible.

Additionally, if you have a hectic schedule and “six days per week” in theory ends up being just “three days

per week” in practice, then you’ll obviously need longer sessions just to achieve the same volume of practice.

► A Little Extra Consistency

× A Little Extra Time

= A Massive Increase in Volume & Progress

A little extra consistency and a little extra time per training session can compound into a massive increase in volume and progress.

You know what the difference is between 20 minutes 3 days per week and 30 minutes 4 days per week? Progressing twice as fast. That’s 60 versus 120 minutes every week. Up that to 40 minutes 5 days per week and you’re nearly doubling again at 200 minutes per week.

► Don’t Overreact to Bad Days

Even if you’re making the right decisions, you can still have bad days. Even if you play your cards in the statistically optimal way, you will not win every hand. So it’s important to stay consistent and not let a single bad outcome derail a good overall trend.

Yes, that can be difficult psychologically. We tend to be risk-averse and overreact to negative outcomes. But it can help to zoom out and look at your progress on a longer timescale. The long game is the only game that matters – short games only matter insofar as their contribution to the long game.

At the same time, though, you can't use that as an excuse to avoid measuring progress and thinking critically about it. Bad days don't matter if you're making statistically optimal decisions, but they sure as hell do if you're not. Every time there's a bad outcome, you have to ask whether there's anything you can learn from it to carry into the future.

Sometimes there's a flaw in your reasoning. Other times there's a flaw in your assumptions. Perhaps you didn't have all the key information to begin with and you should have done better due diligence. Or perhaps some information was revealed or changed after you made a decision and you were too slow to react.

It's not worth beating yourself up over mistakes, unless they're mistakes you're repeating over again. One-and-done mistakes won't keep you from making progress in the long-term, but repeated mistakes will.

► Seek Virtuous Cycles

It's a fact of life that things compound. You improve one aspect of your life, it will have carryover effects, which themselves will have further carryover effects, and so on.

It's important to take advantage of these feedback loops and orient them in a positive direction – because if they're not moving you in a positive direction, they're moving you in a negative direction. There is no stable equilibrium.

You're going to get pulled into self-perpetuating cycles whether you like it or not. So it's important to do all you can to get yourself pulled into virtuous cycles, not vicious ones.

► Protect The Habit

If I've learned anything about habit formation and maintenance, it's that you've absolutely got to stay consistent, even (and especially) during those times when you need to dial back the total volume. Obviously you'll move slower when you throttle the volume, but at

least you'll protect the habit. Yes, the wagon will be moving slowly, but you'll still be solidly on it.

When someone gets derailed from their journey to get better at writing, math, coding, an instrument, a sport, or whatever it may be, it's almost always the same story: at some point they fell off the wagon entirely and never managed to get back on.

When the time comes to get back into the swing of things, it's much easier to speed up a slow wagon that you're on, than to get back on a wagon that you've completely fallen off of. Every time you fall off the wagon, there's a chance you might not get back on. Don't compound that chance into a certainty. Once you've got a good habit going, do everything you can to protect it.

Chapter 2. Skills

► Hardcore Skills Are Always The Answer

Hardcore skills are the biggest bottleneck to improving your life and the lives of others. It doesn't matter where you fall on the spectrum from selfish to altruistic. Hardcore skills are always the answer.

So many people want to make a big impact on the world and in their own lives. But desire is not enough. You can't do anything big unless you have big skills. Sure, some people get lucky – born into the right family, in the right place, at the right time – and enjoy outsized influence with under-sized skills. But even for those people, the difference between a relatively large impact (relative to other people) versus an absolutely large impact (“put a dent in the universe”) still comes down to skill-building.

Hardcore skill development is also one of the greatest social mobility hacks. Even if your family is not well-connected, you can make up for it by developing

real skills. Sure, you have to develop more skills than well-connected people to reach the same level of opportunity, and you're going to have less guidance developing those skills and finding your way to the arena – but once you're in the arena, those extra skills pay big dividends.

► The “Alien-Level Skills” Hack

Having strong technical chops can be a gigantic power-up that sets you apart. You get to provide value that nobody else can, and you get recognized for it. That's what happens when you equip yourself with alien-level skills and solve problems for Earthlings.

But one of the things that keeps people from capitalizing on this hack is they don't invest in building broad technical foundations. For instance, a common argument against building broad mathematical foundations is “why not just wait to learn math until you have a job in a math-adjacent field, and then backfill all the useful math as you encounter specific problems on the job?”. On the surface, that might sound like a way to reduce the amount of work that it takes to develop alien-level math skills that set you apart and boost your career – however, in practice, what this “wait to backfill” approach *actually* reduces is

your chance of being able to capitalize on the alien-level skills hack. Here's why.

If you work in a math-adjacent field and don't have much math background, then:

1) You'll underestimate how often mathy tasks come up in your work. If a mathy task comes up, but you don't have plenty of math background, you probably won't even realize how mathy it is.

2) Even if you do come across a task you know is mathy, you likely won't have enough math background to identify what mathy approach you need to take to complete the task (i.e., what specific math you need to apply or spin up on).

3) Even if you do know what mathy approach you need to take, the task might be handed off to someone else who already has more math background (because it will take them much less time to spin up and solve the problem) – and as this happens repeatedly, their ability to capitalize on mathy tasks will expand, while yours will remain fixed. You'll be on the losing side of the Matthew effect of accumulated advantage.

4) Even if there is nobody else to steal the task, if there is time pressure, then you might not have enough time

to spin up and carry out that mathy approach. This can happen in a couple different ways:

- Defaulting to cumbersome methods: *"We don't have time for you to spin up on math for an ideal solution, we need to ship NOW. Just do the best you can in a week using what you currently know, even if it's not great, and we'll figure out a way to patch over whatever issues come up afterwards, even if the patches are complicated."*
- Passing up the problem: *"This would have been a great opportunity if we had someone who could solve this problem reasonably quickly, but we can't spend tons of time on it, so we're unfortunately just going to have to pass it up and focus on things that are closer to what we're able to do at the moment."*

I'll end with one caveat: depending on the field you're going into, you typically CAN do a reasonable amount of scoping down. For instance, if you want to work on ML/AI then you probably don't need to learn Abstract Algebra. But at the same time, there is still a mountain of math you'd benefit from knowing. Many topics in calculus, linear algebra, and probability & statistics tend to come up.

There's a balance to be had; there is an appropriate level of scoping. On one hand, you CAN skip out on math that is largely irrelevant for your math-adjacent field – BUT, whatever math *does* tend to show up in your field, if you don't have broad knowledge of it, you're going to struggle to pull off the alien-level skills hack. You CAN scope down and discard math that's absent from your field in general, BUT if you want to pull off the alien-level skills hack, then you should NOT scope down further and discard math just because it appears to be absent from a specific problem.

► The Importance of Having Your Prerequisites In Place

Having your prerequisites in place is the difference between something seeming confusing and inaccessible versus “wait... that's all it is?”. It's easy to think you lack learning ability when really you just lack prerequisite knowledge. Differences in learning ability do exist, but they're often conflated with presence or absence of prerequisite knowledge. (But beware: for the same reason, it's also easy to think you're a wicked fast learner when really you've just mastered more

prerequisites than your peers. Fill in your prerequisite knowledge, but don't get cocky.)

In general, the way to “unlock” things that feel inaccessible to you is to shore up your prerequisite abilities. This applies not just to learning tasks, but also to opportunities. Everybody knows that luck is where preparation meets opportunity, but fewer people understand that if you don't have the prerequisite abilities in place that prepare you to capitalize on an opportunity, then you probably won't see it in the first place. Imagine how many opportunities you're blind to because you don't have the prerequisite knowledge to even see them whiz by.

► Most People Don't Realize They Can Systematically Climb Skill Trees

They try a bunch of random stuff in random order, inconsistently, see little progress, and give up quickly.

They think they can't climb the tree when really it's just that they can't reach a tall branch directly from the ground.

They don't realize that every skill tree has a tree trunk and lower branches that you have to climb first before you can reach the higher branches.

They don't realize it's a tree. They just think it's some nebulous dispersed ether of scattered things that they either can or can't do.

They hit ceilings and think those ceilings are locked in place, impossible to break through, it either comes naturally to you or it doesn't.

Sometimes they realize this later in life, by chance, when picking up a new hobby or learning a new language or getting serious about work or having to teach someone.

And then it dawns on them that systematically climbing a skill tree is a generalizable approach, and they think about all the time they've wasted standing on the ground jumping at the highest branches of skill trees and doubting their ability to reach it, when all they had to do was climb.

► The Least Efficient Learning Strategy

The least efficient learning strategy is to ask "why am I so dumb" over and over. A more efficient strategy is to ask "what (prerequisite knowledge) am I missing" and go back and fill it in.

If you started a TV show midway through the season and couldn't make sense of what's going on, you wouldn't fault yourself for that, you wouldn't jump to the conclusion that *"this show is too smart for me and I'll never understand it."*

You would immediately realize *"well duh, of course this episode makes no sense, I haven't seen (or have forgotten) the previous episodes. I need to go spin up on them."*

It's the same situation in learning, even if lots of people don't realize it.

► The Skills You Need Should Sit Several Layers Deep

To have enough mental bandwidth to think deeply about a complex situation, you need total fluency in the fundamentals. They must lie well beneath your edge of

ability. You can operate at that edge – but not *comfortably* – and that makes all the difference. Your high-level train of thought will get continually derailed by the low-level details you have to manage. You'll have a hard time seeing the forest for the trees.

To hammer in your fundamental skills to the point of comfortable execution, it helps to not only get plenty of practice with those skills, but also layer plenty of more advanced skills on top.

For instance, consider figure skating. Yes, figure skaters get really good at skating in part because they skate a lot, but it's not just that. It's also that they continually layer more advanced jumps and spins. Skating around the rink will get you to a decent level of comfort in your basic skating skills, but being able to land jumps and spins will force a whole new level of robustness and fault-tolerance in those underlying skills. It's like those robot testing demonstrations where all the engineers stand in a circle shoving the robot around. It's not enough to just test that the robot can follow a predetermined path. You gotta bang it around a bit to make sure it's resilient.

It all comes down to forcing structural integrity of underlying skills. When you build advanced features on top of a system, they sometimes fail in ways that reveal

previously-unknown foundational weaknesses in the underlying structure. This forces you to fortify the underlying structure so that the system can accommodate new elements without compromising its integrity. And when you fortify the system to execute advanced tasks successfully, it becomes capable of executing simpler tasks *comfortably*. What's more, fortifying the underlying structure often requires improving its organization and elegance, which, in the context of knowledge, produces deep understanding and insight.

► Your Missing Foundations Will Wait For You

If you think you don't need to climb a skill tree and master your prerequisites, then don't. Just go ahead and do <insert hard thing here>. That's what you claim you're able to do, so do it.

Or try and fail enough times to gain some humility. Your missing foundations don't care. They'll wait for you patiently until you're ready to acknowledge them.

► Consuming is Only Helpful Insofar as it Enables You to Produce

If you want to make serious progress learning to skate, you have to obtain a pair of skates and a surface to skate on, and then do skating exercises that gradually increase in difficulty. You cannot learn to skate just by consuming skating videos and reading about it, even if you discuss with a friend afterwards. You might be able to become somewhat of a critic that way, but not an actual skater.

This phenomenon is completely obvious and generalizes to every single skill domain there is. If you want to gain the ability to produce something -- a movement of your body, a movement of your mind, whatever it may be -- you have to practice producing, not just consuming. Yet every so often I come across somebody who thinks they can have found a clever method to, e.g., learn math without actually solving problems.

If you feel personally attacked by this then please understand that the whole reason I'm saying this is because I want to see you achieve your goals. Passive consumption is not the way to do that. Consuming is only helpful insofar as it enables you to produce.

► Actively Doing is the Key to Alpha

Actively doing (as opposed to passively consuming) is the key to effective learning. But it's more than that. It's also the key to alpha, i.e., developing an edge.

Lots of people consume. Fewer people actively do. Even fewer people attempt challenging things. And even fewer people than that build up the foundational skills needed to *succeed* in doing those challenging things.

What's more, alpha compounds. When you succeed in doing a hard thing, the learning and resources you acquire will position you to succeed in doing even harder things and acquiring even more alpha. In other words, your edge gets sharper – not duller – with use.

► Everything Matters

Hard and soft skills, big-picture and detail-oriented thinking, technical and creative ability... they're ALL extremely valuable. We all have strengths and weaknesses, but the best results come from being solid

in every single one of these categories and world-class where we naturally excel.

Lean into your strengths AND shore up your weaknesses. You are maximizing a product, not a sum. Everything matters. Running your gifts and talents to the max means you have to also play defense against weaknesses that would otherwise try to interfere. The running back has to run but the blockers also have to block.

► You Want Exciting Opportunities? Learn Math and Coding

Learning advanced math and coding opens career doors you don't even know exist – sometimes even doors that the whole world isn't aware of yet.

Everyone knows that the future is here, it's just not very evenly distributed. You know who it's concentrated on? Those who are insanely skilled.

Likewise, everyone knows that the easiest way to become insanely skilled is to skill-stack. You know what skills pair really well with each other and basically everything else? Math and coding.

If you've got serious math chops, coding chops, and deep domain expertise in another discipline, then you're compounding three orders of magnitude. Be one-in-a-hundred on each and now you're one in a million overall. Be one-in-a-thousand in each and now you're one in a billion overall.

I don't think I've ever run into someone with serious math/coding chops and deep domain expertise, who wasn't working on something really exciting.

You take some area of interest, you go down the rabbit hole that's been dug by previous explorers, you run up against the rocky technical problems that prevented further digging, you smash those rocks to bits with your math/coding jackhammer, and you just keep going and smashing the crap out of any more problems that dare get in your way. What could possibly be more exciting than that?

Finally, keep in mind that in order to get yourself into the situation above, you have to actually be skilled. You can't just "appreciate" or "talk a good game" about math and coding.

When you run up against a rocky technical problem, nobody cares how amused you are by the rock, and

nobody cares how hype you get telling the rock how you plan to destroy it. The only thing that matters is that you can wield your math/coding tools masterfully enough to destroy the rocky technical problem.

Basically, you need to develop as strong a command over math and coding as a musician's command over their instrument, or a gymnast's command over their body. That takes a massive amount of consistent practice over a long period of time – which is hard, which is why most people don't do it, which is why you get such an outsized competitive advantage if you do.

TLDR: Pursue a domain you love, but simultaneously get so insanely technically skilled at math and coding that you can apply them to your domain of interest in an innovative way. Deep domain expertise plus alien-level technical skills equals lots of interesting and rewarding work to do.

► If You're Not Both Technical and a Domain Expert, Then You're Underpowered

If you're a domain expert but you can't build stuff yourself, then you're severely underpowered. Likewise, if you're technical but you rely on someone else for

domain expertise, then you're also severely underpowered.

If you check one box and find your complement in someone else, then that's great, but you're still somewhat underpowered. Yes, you can get things done together, but the cycles of communication and coordination will still turn much slower and cut much shallower than if they were happening entirely within your own brain.

It's best to be a technical domain expert and work with other technical domain experts when possible.

► Domain Expertise, Math, Coding, Communication

Domain expertise to identify an important problem and envision a solution, math and coding to build it, and communication to deliver it.

Without domain expertise you'll choose an unimportant or intractable problem or your solution won't really solve the problem (because you don't really understand the problem).

Without math and coding you'll be limited to whatever someone or something else (with comparatively little domain expertise) can build for you. You'll lack a technical edge against copycat competitors.

Without communication skills your solution won't be understood and adopted. You'll mistake lack of traction for lack of merit when it's really just a failure to articulate value.

► What Math To Learn for Skill Stacking

What math should you learn if you want to get the most ROI out of “skill stacking” it with another applied domain you’re interested in? Where’s the line between math being a true power-up vs a fascinating distraction?

In my experience, most of the value comes from the following:

1) *Being rock-solid on the applied math that shows up all the time across all fields of engineering:* linear algebra, calculus-based probability & statistics, algorithms, etc.

2) *Having broad knowledge of advanced approaches to mathematical modeling:* differential equations, machine

learning, etc. This includes not just conceptual understanding but also procedural fluency with foundational techniques.

3) *Knowing the basics of proofs.* It's less about knowing particular theorems and more about being able to follow lengthy logical manipulations and carry them out yourself in general. (However, the way you build up your general skills here is by practicing on particular instances – there's no escaping the particulars.)

4) *Going deep into mathematical modeling supporting the particular domain(s) you're interested in.* Here, it's not just about learning and developing modeling techniques, but also building domain expertise. Your models will only be as good as their data, and most data isn't available in machine-readable or even human-readable formats – it's siloed in the heads of domain experts who have extracted a massive amount of learning from a massive amount of hands-on experience.

► Failure Modes in People Who Develop Math Skills but Don't Capitalize On Them via Coding

I know some people who are skilled in math but never really capitalized on it. Two trends I noticed: 1) they were typically not very skilled at coding, and 2) they also lacked discipline to work on things that had to be done but weren't inherently intellectually enjoyable.

The idea of someone being good at math but not coding is sometimes surprising to readers who think that being good at math naturally transfers over to coding. While there is definitely a positive correlation, the correlation is nowhere near as high as you might expect. In particular, here are 4 failure modes that even mathematically inclined people can fall into, that prevent them from capitalizing on their quantitative skills in the form of coding:

1) Difficulty grappling with complexity when it grows so big that you can't fit everything in your head. Not organizing and naming things well, not understanding or maintaining scope, letting responsibilities bleed too much across scope. Accepting needless complexity, sometimes even viewing complexity as a friend rather than the enemy.

2) *Lack of understanding or willingness to accept practical constraints of the problem and incorporate them into the solution.* It's good to think about the Platonic ideal and build towards it by default, but you can't let it become a constraint in the sense of the great becoming the enemy of the good.

3) *Getting distracted by low-ROI features/details.* In math there's typically a clean line separating details that are absolutely critical versus completely irrelevant. But in reality there's more of a spectrum: lots of things matter at least a little bit, and some things matter way more than others, and it can be hard to tell them apart if you're an abstraction astronaut who lacks boots-on-the-ground experience. You need a good sense of what things matter a lot and how costly their implementation will be relative to their impact. You have to be able to sniff out the things that have a 100x or 1000x impact and ROI relative to other things.

4) *Being unwilling to do "tedious" work.* This plays into item #3 because in order to get that good sense of what really matters, you have to get your arms around the problem, which typically requires getting your hands dirty and doing enough manual grunt work to develop strong intuitions and gut feelings. Mathy people sometimes justify avoiding the grunt work because it's

tedious and they already have it all figured out in their head, but the issue is that the contour of the problem space in their head doesn't match up with reality. Their reasoning tends to be sound, but it's the assumptions that get them: there are parts of the real-life problem that they haven't loaded up in their head. Sometimes they think important things are negligible, other times they think negligible things are important.

► There is Always a Bad Excuse to Rationalize Not Skilling Up

Building strong foundational skills will always provide a massive edge. Most people just don't want to put in the work and will grasp at any excuse not to do so.

If you are choosing not to develop serious technical chops now because of AI or whatever, then don't kid yourself, you were never actually going to do it in the first place, you were always just going to grasp at the nearest excuse, and AI is the excuse you landed on.

Go listen to "Because I Got High" by Afroman and replace the word "high" with "AI". That's you.

Chapter 3. Discipline

► The Magic You're Looking For is in the Full-Assed Effort You're Avoiding

It's amazing how many people will rationalize avoiding skill development on the grounds of "nah bro, I'd rather have a life."

The thing is, you can make serious progress climbing pretty much any skill tree if you just put in 30 minutes of focused effort every day. But it has to be fully focused – a “full-assed” effort – and you have to be continually upping the level of challenge as your capabilities increase. You have to work intensely enough that you come out of each session seriously winded. Meaning that either your brain feels like mush or your body feels like jell-o (or both).

When someone fails to make decent progress towards their learning or fitness goals and cites lack of time as the issue, they're often wrong. It's often not lack of time but rather lack of willingness to put forth a

full-assed effort under a continually increasing level of challenge.

If you put in a half-assed effort then you get a quarter of the results at most. That's what causes the purported lack of time. To get the equivalent of 30 minutes full-assed, you have to put in at least 2 hours half-assed, which you quite reasonably might not have time for. Or you put in 30 minutes half-assed and get the equivalent of 7.5 minutes full-assed, which doesn't move the needle fast enough on your progress for you to reach your goal in a reasonable timeframe. The magic you're looking for is in the full-assed effort you're avoiding.

► Don't Put Off Training Because You Think Your Goals are Too Far Away

Once you actually get on the wagon and feel how fast you're moving, you might realize your goals are closer than you thought. Sometimes what you expect to take 5+ years turns out to only take a single year or less if you train seriously, consistently, efficiently.

This is especially true in academics where standard grade-level learning paces are based on unserious, inefficient training. A grade level's worth of learning

can be compressed much, much shorter than a year if you avoid wasting time.

Basically, if you are willing to train seriously to achieve a goal, don't let yourself get faked out and discouraged by how long it takes people who don't take their training seriously.

Don't get me wrong – there is no shortcut. Developing skills requires work and there is no way around it. You cannot skip out on the work and still come out with the skills. But what you CAN do is work efficiently and compress the necessary volume of work into a surprisingly small duration of time.

(All that may sound obvious, but the reality is that most people are unserious and inefficient with their training, and many people believe that they can find some shortcut by which to acquire skills despite training unseriously and inefficiently.)

► At Some Point, Doing the Hard Thing Becomes Easier Than Making the Hard Thing Easier

The condition for getting yourself to do something is simple: it's just $\text{internal willpower} \geq \text{external friction}$. If that condition is false then the way you make it true is by decreasing friction and/or increasing willpower.

It's helpful to think of this like balancing a budget: willpower is like your income and friction is like your spending. If your budget isn't balancing, then the first thing to do is cut out any dumb costs. Is there anything dumb about your environment that's causing needless friction? Cut it out. Your life is like a big codebase – if you're struggling to implement a new behavior in some area, then refactor that area to make it easier to build on.

But at the same time, you can only take cost-cutting so far. There are always going to be some basic expenses you have to cover. And there's a limit to how easy you can make it to add a new feature to the codebase. You can refactor all you want but there's always going to be some amount of complexity inherent to the new feature.

The trick is to be honest with yourself about when you start asymptoting off in your attempts to reduce environmental friction. At some point doing the hard thing becomes easier than making the hard thing easier. And that's when you have to muster up the willpower to overcome whatever friction is left over. That's when you have to say "enough refactoring, time to lock in and implement this sucker." That's when you have to change your focus from cost-cutting to producing extra income.

In any journey, you can chart an easier, more efficient course, but there's always going to be some serious trekking involved.

► How to Cultivate Discipline

When there's something that you know you should do, but you can't get yourself to, it means some habit is pulling you away from productive action. So what you need to do is tear down the unproductive habit and build up a counter-habit whose gravity eventually becomes strong enough to completely overtake the original habit. You disrupt momentum on your negative habit and create momentum towards a positive habit. And while you might not be able to do all that in one fell swoop, what you can do is iterate and

gradually ease into the transition one little step at a time.

Here's a concrete example which may or may not apply to you but will hopefully illustrate the main idea. Let's say you're having trouble cultivating the discipline to exercise every day. The first question is: what's keeping you from exercising? Maybe you plan to exercise after work but then things come up and you always find an excuse. Okay, so do it first thing in the morning. Why aren't you already doing it first thing in the morning? Let's say it's because you habitually wake up 15 minutes before you have to leave for work and there's not enough time for exercise.

That habit is pulling you away from your goal of daily exercise, so you need to gradually replace it with a more productive habit. Maybe instead of waking up 15 minutes before work, you wake up 20 minutes beforehand and spend 5 minutes doing jumping jacks as soon as you get out of bed. (Or if 5 minutes is too daunting then maybe you start with just 1 or 2 minutes and gradually build up to 5 minutes.) After enough days of waking up 5 minutes earlier for 5 minutes of jumping jacks, you'll have created a "new normal" morning routine, and you'll find it within yourself to wake up another 5 minutes earlier and replace your 5 minutes of jumping jacks with a 10-minute run.

You keep going this direction, gradually tearing down your habit of waking up just before you have to leave, and building up a habit of waking up earlier and earlier and doing more and more exercise with that extra time. Eventually you reach your desired fitness routine goal and then you just maintain that habit into the future.

► Keep Your Hands On The Boulder

People will do unbelievable mental gymnastics to convince themselves that doing an easy, enjoyable thing that is unrelated to their supposed goal somehow moves the needle more than doing a hard, unpleasant thing that is directly related to said goal.

If you want to move the needle on a goal, you have to concentrate your efforts directly on that goal. You can exhaust yourself doing other things, fulfilling other responsibilities and moving the needle on other goals – but at the end of the day, each goal has its own needle, and the general feeling of exhaustion doesn't imply you've successfully moved any needle in particular.

This can be a hard truth, especially for people who have taxing responsibilities that are separate from their aspirational goals. But the only way to achieve those

goals is to somehow find it in oneself to directly move the needle on them. There is no other way.

(When you're pushing the boulder, remember that the boulder is actually a gigantic reaction ball that's bouncing around everywhere trying to evade you, and you just have to keep at it and not give up no matter how many times it bounces sideways or even backwards. It will never move in exactly the direction you're trying to push it, but if you just keep pushing, then you WILL gradually wrestle it over to where it needs to go.)

► Just Do The F*cking Work

If you want to develop serious skills, you have to engage in intense, taxing workouts. Amateurs sometimes make up all sorts of excuses for why this rule doesn't apply to them, but real pros don't try to weasel their way out of the hard work.

You think you're too good for the grunt work? Too smart to listen to your coach's feedback? Then what are you waiting for – go on, succeed all by yourself in your current state. Either prove your inherent greatness, or fail and get your ass handed to you enough times to knock some humility into your head.

At the end of the day you can either waste time debating your coach on the training regimen, or you can use that time to just put your head down and do some f*cking work. One of those actions will turn you into a pro; the other will keep you tethered to amateur level for the rest of your life. It's your choice.

You want outsized results? Then you're going to have to put in an outsized amount of work. Achievement, expertise, greatness, whatever the hell you want to call it – it doesn't happen naturally. It's about transforming yourself from normal to abnormal in ways that confer a competitive advantage. There's nothing natural about it.

Chapter 4. The Grind

► Upskilling is Hard and That's a Good Thing

Upskilling takes work. That's a feature, not a bug. Upskilling increases agency and outsized agency is best held in the hands of people who know what it means to put in outsized work. We all know what happens when lazy people have power. The fact that upskilling takes work is a force against that.

► The Most Superior Form of Training

In the field of talent development, there is absolutely no debate about the most superior form of training. It's "deliberate practice": mindful repetition on performance tasks just beyond the edge of one's capabilities. Deliberate practice is about making performance-improving adjustments on every single repetition. Any individual adjustment is small and yields a small improvement in performance – but when

you compound these small changes over a massive number of action-feedback-adjustment cycles, you end up with massive changes and massive gains in performance.

Deliberate practice is superior to all other forms of training. That is a “solved problem” in the academic field of talent development. It might as well be a law of physics. There is a mountain of research supporting the conclusion that the volume of accumulated deliberate practice is the single biggest factor responsible for individual differences in performance among elite performers across a wide variety of talent domains. (The next biggest factor is genetics, and the relative contributions of deliberate practice versus genetics can vary significantly across talent domains.)

Why, then, does anyone seeking to attain a high level of skill engage in forms of practice other than deliberate practice? The answer might be the most hard-hitting 2 sentences in all of talent development research: “...[D]eliberate practice requires effort and is not inherently enjoyable. Individuals are motivated to practice because practice improves performance” (from Ericsson, Krampe, & Tesch-Romer, 1993, in *The Role of Deliberate Practice in the Acquisition of Expert Performance*). In other words, maximal performance does not happen naturally as a result of maximizing other things like enjoyment,

comfort, convenience, and ease of practice. In fact, maximal performance is at odds with some of these things. Sacrifices must be made.

Lots of people are unwilling to make sacrifices to engage in deliberate practice – and that’s fine. That’s a personal value judgment. But the problem is that many of these people still claim that they are doing their best to develop their talent. Typically, they will cut corners on one of the two requirements of deliberate practice – “mindful” and “repetition” – and then resist any form of objective, quantifiable measurement of their performance that would expose the ineffectiveness of their practice.

Deliberate practice is not mindless repetition. If you’re doing the same thing over and over again, then you’re doing deliberate practice wrong. Deliberate practice is about making performance-improving adjustments on every single repetition. Any individual adjustment is small and yields a small improvement in performance, but when you compound these small changes over a massive number of cycles, you end up with massive changes and massive gains in performance. None of this happens if you’re mindlessly doing the same thing over and over again without making adjustments.

Likewise, even if you're mindful during practice, you can't skip on repetition and still call it "deliberate practice." Deliberate practice necessitates a high volume of action-feedback-adjustment cycles in every single training session. Otherwise, the compounding doesn't happen. Any activity that throttles the number of these cycles cannot be described as deliberate practice.

Many heated debates in education stem from these misinterpretations of deliberate practice. Mindless repetition, doing the same thing over and over again without making performance-improving adjustments, is not deliberate practice. Likewise, any activity that throttles the volume of action-feedback-adjustment cycles (e.g., excessively challenging problems, or think-pair-share type of stuff) is not deliberate practice.

► Outsized Success Requires Outsized Work

You want to do something that sets you apart? You're going to have to work harder than most. There is no way around it. If you think you can achieve outsized success without putting in an outsized amount of work, then you will never achieve your goals because you will

never transform yourself into a person who is capable of achieving them.

And guess what? It's not enough to simply work hard. To achieve outsized success, it's critical to not only put in enough time and effort, but also to work productively. You have to work hard AND work smart. And furthermore, work in a direction where you have some competitive advantage (or, at least, you're not at a disadvantage). Part of this work involves engaging in activities that maximize the likelihood of you getting some lucky breaks. You have to work to maximize your luck surface area.

► Transformation Is Discomforting

If you don't push yourself to perform beyond your level of comfort, you don't improve your performance. Simple as that. Why? Because performance improvements come from your body adapting to additional strain. No strain, no gain.

Strain can be unpleasant. It's taxing and it leaves you fatigued. You may feel weak, untalented, even dumb if you're training an intellectual skill such as math. But it's completely necessary. To avoid the feeling of strain

is to avoid the process of adaptation, and thus, to avoid performance improvement.

What you want is a continual cycle of strain and adaptation. That's true in athletics, and it's just as true outside of athletics. You feel weak while exercising but you come back stronger. You feel dumb while studying but you come back smarter. The thing to remember when studying is that you are physically changing your brain to execute more complicated cognitive tasks. At a fundamental level it is just like lifting weights or practicing gymnastics. The keys to effective training are the same, and so is the feeling of effective training.

Transformation is discomfoting. But keep in mind that while discomfort is necessary for performance improvement, discomfort alone does not always indicate that you're engaging in performance-improving activities. You also need to be able to overcome the challenge that's inducing the discomfort. Think of it this way: dropping a 500lb-loaded barbell onto the shoulders of a novice lifter would be neither comforting nor productive. They'd get crushed, develop no strength, and the only thing they'd learn from the experience is that they hate strength training. The same is true for too-difficult math problems, too-difficult pieces of music... you get the idea.

► Enjoyment is a Second-Order Optimization

If you are seeking to maximize your “bang for buck” in terms of learning per unit practice time, then enjoyment is a second-order optimization that is often at odds with the first-order optimization, namely, deliberate practice. A key feature of deliberate practice is that it requires continually practicing beyond one's area of comfort, and this tends to be more effortful and less enjoyable (as one would expect of something that is by definition uncomfortable).

If you want to maximize your learning efficiency: 1) engage in deliberate practice, and 2) make the deliberate practice as enjoyable as possible (or, equivalently, as least unpleasant as possible). Ranked by efficiency, here's the whole spectrum: enjoyable deliberate practice > unpleasant deliberate practice » other enjoyable forms of training > other unpleasant forms of training.

Now, this is not to say that enjoyment is unimportant. If deliberate practice is not enjoyable for someone, then other enjoyable activities can sometimes be useful

for increasing motivation and softening the discomfort associated with deliberate practice. But it's important to realize that fun is a supplement, not a substitute, for deliberate practice.

Additionally, while deliberate practice is inherently uncomfortable, you can normalize yourself to it via repeated exposure – and once you begin to see your tiny improvements compounding into massive long-term gains, it can feel satisfying. As the saying goes, “nothing succeeds like success.”

► Ability is Built, Not Unlocked

One of the most harmful myths in education is that ability is something to be “unlocked” by curiosity and interest (which seems easy), not something “built” by deliberate practice (which seems hard). It's so funny when you imagine what this would sound like coming from an athletic trainer: “You want to get really good at basketball? Forget about practice drills – you were born to ball; all you need to do to unlock your inner baller is come in with the right attitude and play some pick-up ball at the park.”

Ability is built, not unlocked. You do not have a latent superpower that can be unlocked in a day. You have to

build up a skillset. You build it up high enough and then other people think it's a superpower you always had. It's crazy how many grown-ass adults think there's some kind of mutant motivational spider that's gonna bite them and turn them into intellectual Spiderman the next day, some kind of magical motivational unlock that will bring about a sudden phase transition in their ability. Don't bother. It doesn't exist. Iron Man is a much better mental model. Incrementally build up your skillset.

This is not to say that curiosity and interest don't matter. Just that these things do not themselves build ability. They don't move the needle directly. They motivate people to engage in deliberate practice, which is what directly builds ability. Curiosity and interest “grease the wheels,” so to speak, but they don't actually move the wheels.

► What Max-Efficiency Training Feels Like

There is sometimes a disconnect between what people think max-efficiency training should feel like, and what it actually feels like. It's common to think that max-efficiency learning should feel maximally scaffolded, perfectly smooth and easy the whole way through.

While this is more true than not, it misses an important nuance: max-efficiency training should feel just-enough scaffolded that the learning tasks are challenging yet still achievable in a reasonably quick timeframe. When you're developing skills at peak efficiency, you are maximizing the difficulty of your training tasks subject to the constraint that you end up successfully overcoming those difficulties in a timely manner.

A noteworthy corollary is that you are also minimizing your confidence in your ability to complete the training tasks (again subject to the constraint that you end up successfully completing them in a timely manner). In this way, confidence becomes more of a "hindsight" thing than an "in-the-moment" thing. If you feel confident while engaging in max-efficiency training, it's not because the task in front of you seems easy relative to your abilities, but because you've been in situations before where tasks felt challenging relative to your abilities but you've always managed to come out successful.

► The Necessity of Grinding Through Concrete Examples Before Jumping Up a Level of Abstraction

Many learners fail to understand that grinding through concrete examples imbues you with intuition that you will not get if you jump directly to studying the most abstract ideas.

If you go directly to the most abstract ideas then you're basically like a kid who reads a book of famous quotes about life and thinks they understand everything about life by way of those quotes. The way you come to understand life is not by just reading quotes. You have to actually accumulate lots of life experiences. And you might think you understand the quotes when you're young, but after you accumulate more life experience, you realize that you really had only the most naive, surface-level understanding of the quotes back then, and you really had no idea what the hell you were talking about.

It's the same way in any subject – even math, where information can be packaged into clean theorems that are provably correct. In general, the purpose and power of an abstract idea is that it compresses a zoo of concrete examples. But if you haven't built up that zoo

of concrete examples then you miss out on that power. If you study the theorems but shy away from grinding messy concrete problems, then you will never truly gain the deep intuition to know what the hell you're talking about.

Skipping the concrete examples is a one-way ticket to existential crisis. If you've lived and breathed concrete examples, they'll get compressed into tangible, meaningful abstractions that inject you with a dose of vitality every time you work with them – but if you haven't, then the abstractions will feel dull and lifeless, and you'll constantly wonder what's the point of pushing meaningless abstractions around in arbitrary patterns of allowed manipulations. For instance, a company's balance sheet can tell an incredibly interesting story if you have visceral experience with success and failure in business – but if you don't, then analyzing financials will make you feel like a robot checking whether numbers match semi-arbitrary conditions for being "good" or "bad".

Grinding the concrete examples is NOT about turning yourself into a robot and shielding yourself from intellectual awakening. It's the opposite. It's about equipping you with invigorating experiences that can live on through the abstractions, empowering you to actually know what the hell you're talking about.

► Be Willing to Do Tedious Work

In order to get a good sense of what really matters, you have to get your arms around the problem, which typically requires getting your hands dirty and doing enough manual grunt work to develop intuitions and strong gut feelings.

Many people justify avoiding the grunt work on the grounds that it's tedious and they already have it all figured out in their head, not realizing that the contour of the problem space in their head doesn't match up with reality.

Their reasoning tends to be sound, but it's the assumptions that get them. There's some parts of the real-life problem that they haven't loaded up in their head. Sometimes there are important things they think are negligible, sometimes negligible things they think are important.

► Don't Undervalue Turning Up the Dial on Your Grind, but Don't Overvalue the Last Turn

Regret minimization is often used to justify leaving a comfortable situation to grind towards an life-changing transition is uncertain and difficult in the short term. This might seem like flipping a switch towards 100% grind, constantly pushing the boulder, but it's important to keep in mind that regret minimization cuts both ways. Yes, grind grind grind, but also don't forget to take breaks to spend time with people you care about, especially around big events or when you're not going to see them again for a while.

If you're anything like me and have a nagging feeling that you should be making progress on "the thing" 100% of the time, what might help justify taking breaks is to think about relative speed increases: if you're pushing on the needle X% of the time, what kind of speed multiplier are you leaving on the table?

- If you're pushing 25% of the time, then there would be a 4x multiplier by pushing 100% of the time. 4x speedup is the difference between a decade of work vs a couple years.

- If you're pushing 50% of the time, then the multiplier drops to 2x. For me, at least, that's still leaving a lot on the table.
- If you're pushing 80% of the time, then the multiplier drops to 1.25x. You're getting fairly close to max capitalization.
- If you're pushing 90% of the time, then the multiplier is down to 1.1x. It's basically max capitalization with a slight rounding error.

Sure, if there is a make-or-break moment in your grind, then it might be worth temporarily turning the dial up to 100% to try to capitalize on it. But in the long run, outside of those situations, that last turn of the dial from 90% to 100% is not going to change the overall outcome – all it will do is create regret in other areas of your life. And that regret does not stay external. Even if you try to compartmentalize it, it will find a way to seep into your whole, detracting from your motivation & productivity, eventually conspiring to derail you early.

► When More Volume Equals More Progress

More volume equals more progress provided that you're working productively and not burning yourself out.

- If you're tired and your head is spinning and you're making tons of silly mistakes, then it's time to stop.
- If you're so fatigued that you can't help but zone out (or get distracted scrolling through memes) between questions, then it's time to stop.
- If you skip the next couple days because you're so blown out from the previous study session, then it's time to reduce the single-session duration and increase the consistency.

But until you hit those issues, doing more will have you truly learning more and making faster progress towards your long-term goals.

Basically, challenge yourself to put up some serious volume, but also be honest with yourself about whether you are working productively and showing up consistently, and don't lose the long game trying to win the short game.

► Failure Is NOT the Key to Success

Failure gets over-emphasized as being the key to success. LEARNING is the key to success. Failure only moves you towards success to the extent that you learn from it.

You should never fail twice for the same reason. Correcting a mistake will move you closer to success, but making the same mistake over and over again will lock you into losing.

Always analyze your mistakes. Whenever something doesn't go your way, deeply analyze the circumstances to identify what you should have done differently, or, even if you acted appropriately based on all the information available, how you could have gotten more information or taken intermediate actions to extract more signal to inform later bigger actions.

Maintain a gigantic locus of control and bias towards internalizing blame. Don't dwell or make excuses. Just learn and apply that learning to the future so you don't make the same mistake twice.

► Don't Drown Yourself in the Deep End

One of the most common – yet also most preventable – traps that people fall into, that leads them to give up on their aspirations, is that they go for the most advanced challenges and refuse to step back to a better calibrated level when they struggle. They keep attempting and failing the most advanced challenge with hopes that their practice will naturally fill in the swath of missing foundational skills that are preventing them from completing the advanced maneuver.

There are some people who can pull that off in some settings, but if you don't already have solid evidence indicating that you're one of them (or you don't know what I mean by that), then you're almost certainly not one of them.

For the vast, vast majority of people, jumping in headfirst without a baseline level of foundational ability is not a viable approach. It can be exciting at first, until you come to the realization that you're spending all your time flailing around and being confused. It's like going to the gym, getting super hyped up to lift the heaviest weight, and repeatedly failing to lift it. Not only are you not getting any stronger, but you're also wasting valuable time that you could be using to successfully develop your strength.

Now, I want to be clear: I'm not saying that you should avoid trying to jump into things that interest you. You've gotten really excited about something and you want to see how far you can pursue your interest jumping in headfirst? Go for it! All I'm saying is that you need to be honest with yourself about whether you're making progress – and if you're starting to flail (or, more subtly, doubt yourself and lose interest), then you need to put your ego aside and re-allocate your time into shoring up your foundations.

► Focus Less on Feelings and More on Measurable Progress

Anyone who knows about deliberate practice knows how important it is to spend your time practicing at the edge of your abilities.

But how do you really know when you're at the edge? Most people can tell when their practice is too easy – you're able to complete tasks effortlessly while thinking about other things. But what about when your tasks are too hard? That's often less obvious. Practice is supposed to challenge you, but how hard is too hard?

Here's my rule of thumb: Focus less on feelings, and more on measurable progress. When your practice is too difficult, you're going to be running in place and not making much measurable progress.

Think about what happens if you try to work out with a weight that's too heavy for you to lift. Yes, you might be able to tell that it feels excessively strenuous, but what really gives it away is that despite trying your hardest, you're not able to lift the weight. And not only are you unable to lift the weight, but you're not getting any closer to doing so.

The same thing happens if you work on a math problem or coding project that's too hard. Your brain goes into overdrive, and you work on it for a long time, but you just don't really get anywhere with it. You don't solve the problem, you can't point to any concrete skills you acquired during the process.

So, if you want to practice effectively, here are some things you absolutely must do: 1) have some concrete way of measuring your progress, 2) make sure that whatever you're doing is actually increasing that progress, and 3) make sure that the progress is increasing fast enough that you'll reach your goal in a reasonable (but realistic) amount of time.

By the way, if your goal is really lofty, then a reasonable amount of time might still be a long time – so long that it’s hard to tell whether you’re progressing fast enough. So if you have a lofty long-term goal, I would also recommend to decompose it into a series of shorter-term goals where it’s totally obvious whether you’re making fast enough progress to reach the next short-term goal in a timely manner.

► The Problem with Overly Difficult Problems

It’s tempting to think that to train up your skills, you should be focusing on the hardest training problems. But here’s the thing about “think really hard, struggle for a long time, solve it once in a while but usually look up the solution” problems: they can be fun (for a certain type of person), but they’re not an efficient way to learn.

Approaching challenging problems without having the subskills down pat is like jumping into a game of basketball without having developed dribbling and shooting skills. It might feel fun but you’re just going to be whiffing every shot and getting the ball stolen from you. You might make one layup the entire game

and feel good about it, but that's barely any training volume.

It's like going to the gym to lift weights but only eking out a single rep over the entire course of your workout. You need to be banging out more reps if you want to get stronger, and the only way you can bang out those reps is by working with a level of weight that's appropriate for you.

► It's Not Just About Attempting Hard Things, It's Also About Successfully Achieving Them

“Hard” or “easy” is conditional on your skill level and most skills can be improved through training. When you do hard things, those hard things become easier.

But it's not just about attempting hard things. It's also about actually doing them successfully.

- You can go to the gym, load 500lbs onto a barbell, try and fail to pick it up, and not gain any strength.
- You can open a quantum mechanics textbook, fail to understand any of the math that's going on,

faceplant on a practice problem, and not gain any learning.

The trick is to do things that are challenging but that you can achieve given your current skill level.

► It Always Becomes a Battle of Willpower By the End

If you want to upskill to an exceptionally high level, enough to build a career around it and achieve a high level of success in your field, then you're going to have to compete against people who are more advantaged than you.

There are many different types of upskilling advantages. Starting at an earlier age is an advantage, so is having access to superior training resources, so is having transferable skills from other skill trees that you've climbed. Biological factors (cognitive, physical, dispositional, etc.) often come into play, and some of these factors cannot be changed through extra training, or have soft limits to the range of improvement that can be expected.

Exceptionally advantaged upskillers tend to have some combination of all of these advantages. They are rare enough that you may not encounter them initially, but the higher you climb up the skill tree, the more frequently you'll run into them -- and even if you manage to climb high enough that you're no longer in direct competition with other people, then it means you're in direct competition with an even more challenging opponent, the universe itself.

So, while it's a good idea to upskill in directions where you have advantages, it's important to accept that you will eventually climb high enough that those advantages become table stakes. Lean into your relative advantages, but don't center your identity on them, because no matter how big of a leg up you have at the beginning of your upskilling journey, it always becomes a battle of willpower by the end.

► The Cycle You Want To Be In

You put in the work in a favorable practice environment, you build up your ability, what used to be hard becomes easy, you see the growth you've achieved, you imagine how much more growth is in your near future, you gain confidence in your present skills and

future trajectory, you gain motivation to keep on working hard...

This is the cycle to be in, but not enough people have a favorable practice environment and the willingness to put in the initial dose of hard work when confidence is low.

Chapter 5. The Journey

► Don't Get Hung Up on Youth Competitions

So what if you didn't win <insert competition for young people here>. It's not too late for you to make your mark.

For instance, young mathy people sometimes fall into this trap: they wanted to become a hardcore problem-solver, and they had fun participating in math competitions, which they view as the pinnacle of problem-solving – but they didn't win, they're now too old to compete, they're disappointed they can't try again, and they're a bit regretful in feeling that they could have practiced with fuller dedication. The same situation plays out with science fair, debate competitions, even elite college admissions.

If this is you, it's important to realize that these youth competitions are just practice arenas for early bloomers, and you can still win in the big leagues even

if you never won or even stepped foot in the youth practice arena. While youth competitions can help you build habits and connections and open some early doors, they don't gatekeep your future. You can still become insanely skilled and get recognized for it. This is just the beginning.

For instance, school-age math competitions are not the pinnacle of problem-solving. It just seems that way because in school, that's what lots of mathy people focus on and get recognized for, so it's always in your face. But think about it – of all the world's famous problem-solvers, how many of them gained their reputation from winning school-age math competitions? None of them. Even amongst the minority that did happen to win competitions in their youth, that's not what they're known for. They're known for their problem-solving success on more widely branching paths that they pursued after their initial schooling. Some published acclaimed research that pushed pure mathematics forward. Others solved hard practical problems in industry, applying math to push other fields forward.

All this to say: in the long run, the long game is the only one that matters, and missing the short game doesn't mean you've missed your shot.

► The 3 Stages of Talent Development

Across various talent domains, the journey to developing a high level of talent occurs through a similar general process and can be roughly divided into three stages. That's what Benjamin Bloom discovered in the 1980s while studying the training backgrounds of 120 world-class talented individuals across 6 talent domains (piano, sculpting, swimming, tennis, math, and neurology). Below are summaries of the stages:

- *Stage I: Fun and exciting playtime.* Students are just starting to develop awareness and interest in the talent domain. The teacher provides copious positive feedback and approval and encourages students to explore whatever aspects of the talent domain they find most exciting. Students are rewarded for effort rather than for achievement and criticism is rare.
- *Stage II: Intense and strenuous skill development.* Students are fully committed to increasing their performance. The teacher becomes or is replaced by a coach, who focuses on training exercises where the sole purpose is to improve performance. These exercises are demanding, and the coach provides constructive criticism to help the student perform the exercises properly. Positive feedback is

provided in response to achievement; effort is assumed.

- *Stage III: Developing one's individual style while pushing the boundaries of the field.* Students are proficient in all the foundational skills in the talent domain. They are so committed that they center their entire lives around the talent domain, no matter the sacrifice, and typically work with a world-class expert in the talent domain. The expert helps the student identify and lean into their individual strengths so that they can excel beyond perceived human capabilities.

However, there are several failure modes that one can run into when attempting to make the journey through these stages:

- *Failure Mode 1: The Permastudent.* The permastudent perpetually avoids the leap into creative production, opting instead to “expand sideways” and acquire skills that are not foundational for their talent domain.
- *Failure Mode 2: The Wannabe.* The wannabe jumps the gun on creative production before their foundational skills are in place. They build a portfolio of work that lacks substance and is made

trivial by foundational knowledge. Not only is it cringe, but it also has high opportunity cost because all this time could be put to better use actually acquiring said foundational knowledge.

- *Failure Mode 3: The Dilettante.* The dilettante cuts their journey even shorter than the permastudent – they never even make it past playtime, they never commit to serious foundational skill development in anything. The dilettante spends all their time in the land of diminishing returns, engaging in perpetual playtime across a large number of talent domains.

Basically, everybody wants to get to stage 3 in the talent development process, but nobody wants to grind through stage 2. Stage 3 can feel like a return to the playful nature of stage 1. But there is no stage 3 without stage 2. Without stage 2, stage 3 is just stage 1.

► There Are No Shortcuts in Talent Development

If Euler, Gauss, Maxwell, even Ramanujan, Galois, etc., needed to spin up on foundational knowledge, then so do you. Yes, these people all produced plenty of

original research at early ages, but they didn't skip the foundations.

Even Ramanujan self-studied. He didn't just sit there and think up math with no prerequisites in place. In his teens he reportedly worked through Carr's *A Synopsis of Elementary Results in Pure and Applied Mathematics*, consisting of 5,000 theorems.

And Galois? Same thing. Galois had a multi-year period of mathematical upskilling before making creative contributions. By his mid teens he had accumulated several years of formal schooling that included math. Then he caught the math bug, became obsessed with Legendre's geometry textbook, worked through it, and continued working through key research papers for a couple years afterwards. And THEN he started producing his own creative contributions.

(Keep in mind there were also far fewer prerequisites needed to make novel contributions to math back then, so a person could fill up their prerequisites to cutting-edge math in a shorter time. Additionally individuals were presumably highly cognitively advantaged with outsized learning rate and generalization ability, which would allow them to move faster through a curriculum, infer more knowledge

beyond it, and infer more knowledge from self-directed experimentation. Basically, across all math learners, these individuals would have been some of the *most* likely to have their prerequisites in place.)

In general, when it looks like someone progressed so fast they “must” have taken a shortcut, what really happened is they speed-ran the foundations. Either that or you’re overestimating their actual ability (likely because they’re exploiting signaling to trick you).

How can people speed-run the foundations? By way of a more efficient training environment, advantageous individual differences leading to more rapid skill acquisition, or by allocating way more of their time into training than is typical. Elite performers typically emerge from a combination of all three of those things.

Take Euler for example.

- *Advantageous training environment?* Check. Euler was taught by his father, who had studied under Jacob Bernoulli, and Euler enrolled in the University of Basel at 13 years old, where he studied under Johann Bernoulli and self-studied more advanced textbooks under Johann's guidance.

- *Allocating way more of their time into training than is typical?* Check. I mean, geez, what 13-year-old is taking university courses and self-studying even more math textbooks on top of that?
- *Advantageous individual differences leading to more rapid skill acquisition?* I don't think there's any formal record, but you'll see Euler on pretty much any "highest estimated IQ of all time" list. It seems pretty safe to conclude he was running on some exceptional cognitive machinery.

► If You're Making Silly Mistakes Then You Need More Practice

Climbing a skill hierarchy like math is not just about conceptual understanding, it's also about reliable execution. If you're making "silly mistakes," then you need more practice, simple as that.

If you don't clean up your silly mistakes on low-level skills, then you eventually hit a wall where no matter how hard you try, you're unable to reliably perform advanced skills due to the compounding probability of silly mistakes in the component skills.

Think about gymnastics: if you're "almost" able to land a backflip, then that's great – but at the same time, you're NOT ready to try any combo moves of which a backflip is a component. Even if it's a silly mistake keeping you from landing the backflip, you still have to rectify it.

And even that's the most optimistic scenario. Other times, silly mistakes indicate a deeper conceptual misunderstanding that you don't even know you have until you are held accountable for rectifying those mistakes.

► No Train, No Gain

If you're not measuring performance and taking actions to improve it then you're not seriously training, you're just playing around. Which is totally fine at the beginning to get a sense of what you like and dislike, what you're willing to commit yourself to training... but sometime you gotta grow up, ya know?

The world rewards those who train. No train, no gain. If you don't want to gain anything then sure, just play around forever, but don't get mad when nothing comes your way.

► Why You Should Push Yourself

Why push yourself? Because the road always stretches farther than what you can see in front of you, and you maximize your reward by traveling as far as you can.

Pushing yourself isn't about racing to a finish line. It's about avoiding stagnation. It's about efficiently growing your skills and continually leveraging them into new growth experiences. That's how you reach your potential. To do anything else is to fall short of your potential.

► Keep Your Foot On The Gas

Any time something initially comes to you easily, it's tempting to take your foot off the gas and try to coast the rest of the way. You coast for a while, just long enough for coasting to become the new normal, just long enough to forget that there's even a gas pedal. And eventually you grind to a halt just above the base of the mountain you're trying to climb.

You play the blame game – first you blame the mountain for existing, then you blame yourself for being incapable of climbing it. But at some point you

mature, come to terms with reality, and realize that all you have to do is put your foot back on the gas. And then you start making progress again.

► You Are a Car

You are a car. You go fast on paved roads and get stuck in mud.

Upgrade your engine to be as fast/powerful as possible, try to stay on the roads and avoid the muddy zones. (Train your strengths into superpowers and tailor your environment & goals to them.)

But at the same time, don't let a little mud (weakness) derail your journey. Look for the least muddy zones, try to cross there, and if you still can't, then upgrade your tires. (Shore up your weaknesses so they don't get in the way of your strengths.)

► What to Do When You Hit a Ceiling

In many talent domains, upskilling becomes hard and unintuitive for everyone at some point, and that point is different for everyone. In math, for instance – some

people start to experience major cognitive friction in algebra, for other people it's calculus, for others it's real analysis, for others it's algebraic topology, for others it's research-level math, and there's even this same gradation even within research-level math problems. The friction doesn't create a hard ceiling, a level at which one is suddenly incapable of further progress, but rather a soft ceiling, a point at which the amount of time and effort required to make further progress begins to skyrocket until upskilling in that direction is effectively no longer a productive use of one's time.

Sometimes people hit a ceiling early due to ineffective or inconsistent practice techniques. But even if you practice effectively and consistently, a ceiling still exists. It's just like sports: few people practice effectively and consistently enough to reach their athletic potential, but it's just a fact of life that most people could not become professional basketball players even with 100% effective and consistent practice. You can't detect ceilings with 100% certainty, but if you're practicing effectively and consistently, and you get stuck in a plateau, and one-on-one training with a coach or tutor doesn't break you out of that plateau, then it's pretty likely.

The natural question is: what do you do when you hit a ceiling? In general, when you feel yourself running up

against a ceiling in life, the solution is typically to pivot and into a direction where the ceiling is higher. For instance, the story of many a quantitative software engineer goes like this: 1) loved math growing up and wanted to be a mathematician, 2) realized during undergrad or grad school that they had lost their "edge" compared to other aspiring mathematicians, 3) also realized that they have a knack for coding and interest in some applied domain, and that the problems that need to be solved there boil down to interesting math that most people in software don't have the math chops for, and 4) pivoted in that direction where their ceiling is higher.

► Learn to Love the Games That You Were Built to Win

We all have limits, but they're often much higher than we think. Most people's ceilings are artificially lowered by inefficient practice: not mastering prerequisites, not working enough exercises, transcribing from a reference instead of pulling from memory, etc.

Push hard and train efficiently, but at the same time, if you're ambitious, then continually evaluate where your ceiling is highest, and learn to love the games that you

were built to win. Understand that games can change drastically as the level of competition rises. Your edge against 90th percentile competition might be table stakes against the 99th percentile, and the 99.9th percentile might have some kind of insurmountable edge over you.

At the same time, don't be too quick to call an opponent's edge insurmountable. Often, you can train up your abilities fast enough to stay competitive. Sometimes, what looks like inborn advantage in an opponent is actually the result of training. But not always. Sometimes you hit a soft ceiling, a point of severely diminishing marginal productivity, and there's nothing to do but pivot to a game where your ceiling is higher. Don't make your ceiling lower than it has to be, but when you run up against one, don't stay and spend the rest of your life hemorrhaging opportunity cost.

The good news is that there is a combinatorial explosion of games to play in life, and even if you pivot along some axis, you can and should continue leveraging much of the training and experiences you picked up along conserved axes. You can continue compounding an intersectional advantage. As reader Abhishek Singh once paraphrased: "The meta-skill is knowing which game deserves your training at this stage of life, and stacking your past experiences so

you're not starting at zero but at an intersection where few others can compete."

► Compound Hard Work and Luck

Both hard work and luck are necessary for greatness and neither is sufficient on its own. However, they are tangled up together, feeding into each other:

- If you work hard, you're more likely to get lucky – you create more lucky opportunities ("luck surface area") and are better prepared to capitalize on them.
- If you capitalize on a lucky opportunity and feel like you're succeeding and your hard work is making a difference, you get excited and motivated to lean into it further and continue working even harder. ("Nothing succeeds like success.")

Greatness emerges from a virtuous cycle of hard work and luck compounding on each other.

That said, while many forms of luck can be shaped by hard work, there are other forms of luck that cannot, such as some biological advantages. So, while it's necessary to work hard to capitalize on the

compounding nature of hard work and luck, it's also necessary to choose a direction that enables you to capitalize on any lucky innate edge you may have.

► Get Yourself In A Position Where You Can Eat Risk

An underrated component of finding career fit is building enough savings to pursue opportunities where reward is uncertain. Everybody knows that learning / upskilling is a huge component of career success, but so is the ability to eat risk. And it's also a huge competitive advantage. It doesn't matter how smart or skilled you are if you can't eat the risk.

► Tie Your Comfort to Real, Tangible Value

The whole point of skills, technology, superpowers, whatever, is to accomplish greater things. Not to stick to the status quo with less work.

Lots of people begin their upskilling journey with this mentality but lose it along the way. They get lazy and use their skills to minimize the amount of work needed for a baseline comfortable life. You follow that path for

a while, you end up stuck in a micro-optimization arena where more work has diminishing returns anyway. It's hard to escape from that quicksand.

The way to avoid the trap is to develop strong emotional ties to a lofty problem where incrementally solving it incrementally transforms the lives of some people you care about while incrementally making your life more comfortable. You need to tie your comfort to real tangible value that you can begin tasting relatively early.

The hardest part is the beginning, when there are other options that provide a more comfortable short term at the expense of 1) an asymptotic long term, or 2) de-coupling your comfort from real tangible value that you provide to the world. The emotional connection helps you stick with it until your current situation is even more comfortable than any alternative, which is the thing that really locks you in for the rest of the long term.

► If You Can't Find a Job That Really Excites You...

If you can't find a job that really excites you, then start (publicly) doing your own thing. You'll run into other people who are doing interesting things, and you can always re-evaluate whether you want to continue doing your own thing independently or team up / work for one of said people.

(Yes, depending on where you are in life you may have to take a job that pays the bills – but try to find a more “neutral” one that at least doesn't crush your soul and doesn't consume all of your time, so that you have the energy and bandwidth to simultaneously keep climbing towards a better fit.)

► How to Allocate Your Bandwidth While Searching for Your Mission

If you haven't found a single mission that you want to focus all your bandwidth on, and you're wondering how to distribute your bandwidth so as to pursue multiple interests while avoiding spreading yourself too thin, then here's the allocation I would recommend:

- one main focus (workload equivalent to a full-time job),
- one semi-focus (workload equivalent to a part-time job), and
- everything else a hobby with whatever time you have left over. (Your remaining bandwidth is about the equivalent of another part-time job, so depending on how many things there are in that "everything else," you might have a small number of serious hobbies or a large number of light hobbies.)

The rationale:

- You don't want to spread yourself too thin. You need to be moving at a competitive speed in at least one direction, i.e., your focus.
- The semi-focus is like a staging area for something that you want to eventually merge into your main focus. In order to successfully complete the merge you're going to have to develop a serious degree of expertise in it, so it has to be more than just a light hobby.

- Hobbies are mainly things that you just do for fun, but they can also serve as candidates to replace your semi-focus once you merge your existing semi-focus into your main focus.

► Repetition Can Lead to Expertise, Expertise Can Lead to Variety

Day-to-day variety can arise from focusing single-mindedly on one big mission. That's obvious if you have kids and they pull you into all sorts of things you never would have anticipated -- but this can happen professionally too.

Counterintuitively, in hierarchical domains, this transition to day-to-day variety is preceded by a lengthy upskilling ramp where you grind through repetitive (but challenging, not mindless) exercises. Eventually, you come out the other end with extreme capabilities that take you on extreme missions.

Chapter 6. The Team

► If You're Asking Someone to Be Your Mentor then You're Doing it Wrong

The #1 way to increase your productive output is to have it pulled out of you by an older experienced person who is unreasonably demanding, incredibly supportive, and has your deepest respect. Even if you think you're working hard already, a person like that can accelerate your output by multiple orders of magnitude by pointing you in the maximally productive direction and motivating you to sprint even faster and longer than you previously believed yourself capable.

But here's the catch: in order to find that person, be worth their time, and have that extra productive output pulled out of you, you typically have to be an incredibly hard and talented worker in the first place, already producing a solid level of productive output. You are not going to run into this person if you're just coasting. You have to turn the dial up from 1 to 10 yourself, and

hold it there for a while, before you meet the person who gets you turning the dial up to 100.

If you're asking someone to be your mentor then you're doing it wrong. It should look less like them helping you and more like you helping them. It starts with you bringing something to the table. You're a missing piece in a puzzle that they're trying to solve. Except, you're not a perfect fit initially. You kind of, sort of fill some gap initially, and you show signs of being able to grow to cover the remaining area. And they recognize that, so they invest their time into helping you grow into that perfect fit.

As you grow, you also expand into other gaps in the puzzle, maybe even some that you didn't anticipate, maybe even some that your mentor didn't anticipate. And you grow to fill those gaps in the puzzle as well. Eventually you cover enough of the puzzle that you yourself become a puzzle-master looking for puzzle pieces. And that's how the cycle continues.

A key takeaway is that when reaching out to someone you hope to work under, make sure you very clearly communicate the following: 1) show that you understand what puzzle they're solving, 2) state what missing piece of that puzzle you think you can fill, and

3) state what evidence there is that you can fill that piece.

► Put Pressure on Your Boss to Come Up with More Work For You

One of the best career hacks – especially for a junior – is to knock out your work so quickly and so well that you put pressure on your boss to come up with more work for you. It causes your projects to grow in scale, complexity, and responsibility.

When you're moving super quickly, your boss can't spend all their time communicating hyper-detailed specs to you, so they have to gradually pull back and offload more of the "scoping out" work to you. You get more responsibility to carry out the project with less scaffolding and supervision, and you build your boss's trust in your ability to execute.

And as you keep executing and forcing your boss to come up with more stuff for you to work on, your boss eventually gets to the point of thinking "I don't have time to scope out more work for them because I need to get X, Y, and Z done... huh, you know, things X and

Y are kind of advanced but I bet they could do thing Z for me with a little bit of coaching.”

Basically, you put so much pressure on your boss to come up with work for you to do, that your boss starts giving you work that they themselves need to do soon, which is really the exact kind of work that’s going to move your career forward.

(Note: There's an assumption here that your boss and organization are well suited for rapid career growth. If that assumption is false, then the very first step is to get yourself into a position where that assumption becomes true.)

► Get On the Right Team

When you’re on the right team, your working hard inspires your teammates to step up their game. When you’re on the wrong team, your working hard causes your teammates to lean back, do less, and let you make up the difference.

If you want to create serious change then you need to get on the right team. You can be the most committed and capable workhorse on the planet, but if you’re on the wrong team, the only thing you’ll change is your

team's allocation of work (i.e., now you do more work and other people do less).

You need to get on a team where increasing your effort produces an outsized gain in your team's collective effort. Why "outsized"? Because when you're on the right team, stepping up your game will inspire your teammates to step up their game as well.

► Competition as a Means of Collaboration

Competition and collaboration might sound like opposites, when there's actually a way in which competition can be reasonably viewed as a means of collaboration. It's when you're engaged in friendly competition with people that you're connected to and care about, where the point is to motivate each other and make each other better.

It's kind of like what you would expect on a serious sports team. During practice, teammates will be competing against each other, trying to create a high-intensity practice environment where they can make each other better. They might even do some light, joking trash-talk to get each other riled up and motivated to put their best foot forward -- not anything mean, of course, but just enough to get the other

person to react like "damn, let me show you what I got!"

But the thing is, it's not even about winning the competition. It's about growing and improving, and the competition is just a way to enter a psychological state where you're motivated to work hard and maximize your effort.

In this type of competition, it actually feels good to see the other person take the lead and raise the bar. The whole idea is that you want the other person to raise the bar on competition and pass you up, so that you're motivated to come right back and do the same to them. It's like you're creating a video game: each time one person passes another person up, a new level and challenge is created. Everyone has fun playing the game and wants to get to really high levels.

You could even call it teamwork: as a team, you try to maximize your total absolute level by having everyone compete on their individual relative levels.

► Your Goal is NOT to Prove You're Smart,
it's to Make Problems Go Away

When you're working on a team, your primary goal is not to prove you're smart. It's to make problems go away.

Yes, the more knowledgeable you are, the better you're equipped to solve problems, but if your primary focus is peacocking your intellect then you're going to create problems instead of make them go away.

You don't get points for creating an unnecessary problem on which to demonstrate your smarts. You don't get points for creating an overcomplicated solution to a simple problem. You lose points for these things.

What you get points for is taking a problem and making it go *poof,* completely solved and easy to maintain and nobody has to think about it anymore.

► Make Your First Impression On a Contribution, Not a Critique

When you join a new team, make your first impression on a contribution, not a critique. It's wiser to spin up on the surrounding context, demonstrate your ability and willingness to contribute, and THEN suggest the improvement. This shouldn't take long.

Almost always, your critique is either 1) completely off base, which you only realize after spinning up on the surrounding context, or 2) correct, and everyone on the team is already aware, but they haven't had the bandwidth to address it (possibly because it's not a top priority).

► Never Come Up Empty-Handed

When you are tasked with fixing a problem for someone else, and things are not going so well, it's best to keep a log of all the things you tried or looked into, like a lab notebook. Two reasons why:

1) Even if you can't solve the issue, if you can enumerate many instances of "here's what I tried and here's what happened, here's something weird I

noticed,” it might spur your supervisor to come up with a suggestion of what you might look into next.

2) It will also demonstrate how much honest work you put forth trying to solve the problem.

It’s crucial to recognize that working on a project for an extended period of time and not having anything to show for it is a hallmark of a lazy, dishonest, or incapable person – and such people are everywhere. They thrive in low-visibility environments like mold in a dark, damp basement. They gravitate towards projects that could reasonably not pan out, where they can maintain the perception of pulling their weight while continually failing to deliver real results.

You may not be a member of the mold, but if you find yourself in a situation resembling its natural habitat, then you need to actively demonstrate your work ethic, integrity, and capability with tangible receipts. If you’ve done honest work, you should be able to back it up. You can’t expect people to take your word at face value. You need to prove that your work survives when it’s exposed to the sunlight.

► You Need a Berserker At The Helm

Progress always comes down to a small number of people working heads down on a decades-long mission and seriously holding the line on what it means to join the team and contribute. It's a constant fight against mediocrity, and unless you have a berserker at the helm and an airtight chain of accountability, you will accomplish nothing.

Chapter 7. The Mission

► Selecting a Good Problem to Work On

What problem should you work on? At first, it doesn't matter. Just work on any toy problem that interests you so that you can build technical skills and gain domain knowledge. It doesn't matter if you solve it or not, whether it's been solved before, or how impactful it is. But once you start getting into problems that require many years of full-time work, selecting a good problem becomes very important. This is the land of startups and research labs, many of whose inhabitants regret that all the time and effort they invested did not yield a commensurate reward.

Personally, most of the toy problems I worked on were bad problems. That's okay because I learned a lot and gained a lot of skills – which is the whole point of a toy problem – but it made me painfully aware of two failure modes that can make a problem bad.

Failure Mode 1: You don't have an implementable vision of what the solution is. In particular, some of the resources you need (e.g. data, algorithms, compute power) do not actually exist yet and you don't have a good plan for obtaining them.

This means the problem is too hard and you probably won't be able to solve it. In my experience, many complex systems modeling problems lie here, e.g. creating useful predictive models of the human brain or macroeconomy. For instance, there was a time when I was interested in modeling the human brain. I framed it as a regression problem on a time-series data set containing the activities and connection weights of all the individual neurons in brain. It took me a while to realize that the data set I wanted did not exist, and creating it would require multiple lifetimes and revolutionary breakthroughs in wet-lab neuroscience (and I was not interested in wet-lab work).

It's worth noting that sometimes, Failure Mode 1 is an indication that you're not actually interested in the thing that you think you're interested in. In my case, I thought I was interested in neuroscience, but it turned out that I was interested in a lot of stuff that just happened to show up in neuroscience: multiscale modeling, connectionism, human learning/intelligence, etc. The most obvious thing that encapsulates all of

those interests is building a model of the biological brain, but it's not the only thing. What I'm doing now encapsulates all of those interests I listed and does not require any wet-lab work. I did still have to get my hands dirty with lots of teaching and content writing, but those were things that I enjoyed.

Failure Mode 2: People don't care about the problem. They are not willing to pay for a solution with whatever currency you're interested in (money, citations, their time/attention, etc).

This means that you're not going to experience any reward for solving the problem. In my experience, theoretical modeling problems can fall victim to this when the problem framing abstracts away details that make the problem intractable but are important for application to real life.

It's possible to argue that Failure Mode 2 doesn't apply to you if you're ahead of your time. However, there are two issues with that. First, you're probably not ahead of your time. Being ahead of one's time is rare, unverifiable, and tempting to believe. Talk about a recipe for flawed judgement! Second, even if it's true that you are ahead of your time, if you are too ahead of your time, then the reward will come too late in your

life to feel worth the sacrifice. You might not even live to experience it.

(That said, I have met some people who seemed entirely satisfied by exploring their intellectual curiosity without the prospect of receiving an external reward or making an external impact in their lifetime, if at all. These people might be legitimate exceptions to Failure Mode 2. But for the vast majority of people, exploring intellectual curiosity is not enough.)

How do you find a problem that avoids both failure modes? You need to find (or create) an intersection between your own interests/talents, the realm of what's feasible, and the desires of the external world. Unfortunately, it's rarely obvious where the intersection is. All the cards are stacked against its existence: you can't choose what you're interested in or what you're talented in, you can't choose what the rest of the world cares about, and if you're interested/talented in some area to the point that you want to solve problems in it, then your reasons for being interested in it are probably not shared by the rest of the world.

So how do you find (or create) the intersection? What's worked for me is to live two parallel lives – one in which you do solve problems that interest you, and another in which you solve problems that interest the

rest of the world. You continually try to push the parallel lives closer and closer together, and eventually, you figure out how to unify them.

► You Have To Work Really Damn Hard to Figure Out What Fulfills You

You know you have to work really damn hard to figure out what fulfills you, right? The solution to "I don't know what I want to do" is NOT "I guess I just won't do anything because I don't know what to do."

Even if you have literally zero self-knowledge, just try a bunch of random activities and reflect on what you liked and disliked. Boom! You now have some self-knowledge. Now consistently and seriously lean into a few of the things you liked the most (or disliked the least).

Just because something doesn't speak to you now doesn't mean it won't speak to you later. Developing baseline familiarity and competency can completely change the experience. Sometimes we think we dislike activities when we really just dislike doing unfamiliar things and sucking at them. Also, leaning into some activities now doesn't mean you're committing to lean

into them forever. Keep exploring on the side, and if you find some other activity you like better, switch it in!

You gotta come at the problem with a builder mentality, continually iterating. "What fulfills me?" is not some riddle that you can stare at until an epiphany comes to you. Epiphanies seldom happen, and when they do, they only come after you put a ton of work into building. The epiphany is just the final piece that snaps a bunch of infrastructure into place that you previously built. Sometimes you feel the snap when it happens, but other times you don't even realize the significance until later down the road when you look back and try to make sense of how you got to where you are.

TLDR: Self-knowledge is not part of your base install. You don't spawn with it. You gotta work your ass off to acquire it bit by bit, exercise by exercise, experience by experience, just like developing expertise in any other subject. It takes a massive volume of reps and there is no shortcut. All there is is compression. Whether it's self-knowledge, technical skill-building, whatever, you can compress the time needed to acquire it if you put in more volume and/or work more efficiently than is typical.

► The “Progress Equals Pressure” Formula

Most people know you have to build hardcore skills to do hardcore things. But what a lot of people with hardcore skills still don’t know is that you also have to shield yourself from nerd-sniping (xkcd.com/356).

When you build hardcore skills, you increase the surface area of things you can do – but many of these things are just fascinating distractions. And if you allow these side-quests to steal your attention, they dilute the power of your skills.

Making progress is all about putting pressure on a problem: applying the force of your skills to a specific problem area ($\text{pressure} = \text{force} / \text{area}$). Leveling up your skills will increase your force, but if you want that to carry over to an increase in pressure, you have to stay laser focused on the problem area that you’re trying to cut into.

► Love What You Do

If you want to succeed wildly and consider it a life well lived, you have to love what you do. There’s no way around it. If you want wild success, and not just on a

lottery ticket, then you have to put in such a high volume of work that it is life-consuming. And if your life is consumed by something you don't love, then it's a life thrown away.

This is not to say you must love an activity to get better at it. Effective practice will make you better at anything even if you don't love the thing. But if you don't love it, you'll never be able to keep up with the same volume of effective practice as someone who does have that love. You'll never outwork them.

Love is perpetual hardcore effort. Not as a descriptor, but as a definition. Love is being consistently hardcore. To say that a parent loves their kid is to say that the parent is consistently hardcore about raising their kid. That their kid is always on their mind and they are always putting max effort into raising their kid. You don't love something if you're not consistently hardcore about it, and you won't be consistently hardcore about it if you don't love it. It's a biconditional, a definition.

Consistently hardcore people achieve extraordinary outcomes through extraordinary actions; these actions go beyond the ordinary and are often seen as crazy. Framed as love, this is familiar: everyone knows that love makes people do crazy things.

► Be a Builder, Not Just a Fighter

Once you acquire hardcore skills and new career opportunities open up, try to avoid those that pit you against other hardcore-skilled people playing zero-sum games. What you want to do is create new value, not just fight over existing value. And the way to do that is to build infrastructure that solves people's yet-unsolved problems. Be a builder, not just a fighter.

► Build Where Building Creates More Opportunities to Build

Nobody who's building something with high future value is worried about AI taking their job. When you're in that setting, building stuff expands your opportunity surface area to build even higher-value stuff. The more you do, the more there is to do. The more of your work AI takes, the more work it creates. Every time AI improves and accelerates you towards your goals, it expands the scope of what's feasible for you to accomplish during your lifetime, and the goalposts move further out.

Ironically, this "high future value" stuff is the hardest to automate because you quickly reach a point where further iteration requires doing things that haven't been done or even thoroughly imagined yet (note that accomplishing these things requires you to be highly skilled yourself beyond what the AI is able to do). So you get in a situation where you have tons of work to do, that AI can't do yet, and you're sad about that.

In short: you're going to be sharing your pie with AI, and your feelings about this will depend on your ability to expand the pie.

Chapter 8. Motivation

► Disinterest is Often Just Overwhelm

It's amazing how many people turn out to enjoy <insert hard thing here> when the tasks they're asked to do are properly calibrated to their current level of skill. Making progress is really enjoyable.

Nobody really jumps off the upskilling train because they lose interest. They jump off because stagnated progress makes them doubt the train will ever reach the destination. They jump off because the level of overwhelm grows too large relative to their level of interest. They jump off because they hit a point of diminishing returns where further time investment is no longer worth the opportunity cost.

It's kind of like how sports and games in general are typically enjoyable by default – unless you're playing against competition that's way overpowered (or underpowered) compared to your skill level and as a result you're not improving, especially if you feel like

you're getting your ass handed to you every time you step up to the plate.

► Why Extrinsic Motivation Matters

It's common to think that "learning for the sake of learning" is better than "learning for the sake of achieving an extrinsic goal," but I'm not convinced that's true. People whose motivation is entirely intrinsic sometimes prioritize "fascinating distractions" over other things that would be more productive to their long-term happiness, in a sense "nerd sniping" themselves.

I think optimal motivation requires a balance of both intrinsic and extrinsic factors. Intrinsic motivation gets you working on interesting things with a unique perspective. Extrinsic motivation keeps you on the rails with your long-term goals and keeps you from falling victim to fascinating distractions.

Furthermore, intrinsic versus extrinsic is a false dichotomy. It's not like you have a limited amount of motivation to split between intrinsic and extrinsic factors. A percentage tradeoff is the wrong way to look at it. It's the other way around: there's no limit to how many motivational factors you can accumulate in each

category, and in turn, there's no limit to how motivated you can get.

So regarding intrinsic versus extrinsic motivation, here are two rules of thumb: 1) a balance across both sources is better than the same amount concentrated on just one source, and 2) more overall is better than less overall.

► How to Become a Super-Producer

The #1 trick to super-productivity is interleaving a wide variety of productive work that you enjoy. You get tired, bored, and unproductive if you're moving along one dimension for too long.

The solution is to get yourself in a situation where your production function has multidimensional input and all those inputs have large partial derivatives on the same order of magnitude. You follow the gradient in that space by cycling between these component activities. The activities are orthogonal, so whenever you start getting tired, bored, and unproductive from moving along one dimension for too long, you just switch to a different activity, moving along a different dimension.

That said, it's important to remember that "interleaving a variety of productive work" is different from running away from your problems. If you've hit a rut on some activity, and you're switching tasks just to avoid dealing with that problem, and it's been this way for a handful of cycles, then you're past the limit of "taking a break to freshen up" and what you really need to do is bunker down and bust through the plateau.

► How The Highest Performers Sustain a Massive Workload

In most domains, the most talented performers have some edge that allows them to level up faster than the average person, and despite having this edge, they still choose to put in much more time than the average person. They spend most of their waking hours pushing the boulder even though it's way beyond what most people consider the finish line. For this kind of person, there is no real concept of a finish line. The talent domain is a component of their identity, not just a skill equip.

The way they sustain such a high volume of work is by interleaving a wide variety of productive activities. They've gotten far enough in the skill domain that

they're well past the narrow tree trunk of fundamentals, and now they have many different branches they can be traveling outwards along. Of course, some branches are more productive than others, so it's necessary to focus one's efforts and avoid spreading oneself too thin, but even still, there are quite a few highly productive activities they can cycle between.

(Note that beginners may have a hard time imagining this because they're still climbing the tree trunk of fundamentals and haven't really experienced the "branch-out" effect where it feels like the more you do, the more there is left to do.)

► Overcoming the Paradox of Serious Training

One of the hallmarks of serious training is that past you would consider present you a lofty or even unattainable goal. That's the feeling of transformation over a long timescale. But the funny part is that over a short timescale, the feeling is entirely different. Day by day, you may not see it, you may doubt it's happening.

That's the paradox of serious training: the way you get stronger is by continually lifting weights heavy enough to make you feel weak. You never stop feeling humbled by your training, and the primary feeling you experience day to day is weakness, not strength.

That's why it's so important to periodically compare the capabilities of your present self to your past self. Periodically look back at stuff you originally found challenging a couple months ago, a year ago, a couple years ago, etc. That should make the growth obvious. There should be things you used to be really confused about (or maybe even confidently wrong about) that are way more clear now. Or things that used to take a lot of effort to accomplish, that would be way easier now.

Also – maybe this is less wholesome – but once you reach a high enough level of skill you can periodically compare yourself to other people who are clearly less skilled. Not saying things to make them feel bad, or even thinking poorly of them, just noticing evidence that your percentile has changed on the bell curve.

But, of course, you can't spend too long in this state. You dip in to get your confidence up again, and then you snap out and get back to lifting those metaphorical weights that are heavy enough to make you feel weak again.

► The Joy of Upskilling is Like The Joy of Exercise

The joy of upskilling is similar to the joy of exercise. You gotta earn it. Exercise gets way more fun once you notice a body transformation getting underway. It's the same way with upskilling.

But nobody feels the transformation on day one. Or day two, or even day three. You gotta stick with it for months before you really notice a difference. Cognitive/intellectual workouts are just like physical workouts. You can't expect your body or even your attitude towards exercise to transform in a single day.

And it's not just about time. You gotta spend that time doing well-calibrated exercises. If you spend months repeatedly trying and failing to lift the heaviest weight in the gym, you don't get better. If you spend months repeatedly lifting the lightest weight in the gym, you don't get better. You gotta 1) start at your level and 2) incrementally layer on more difficulty as your level increases.

Although you may not see a macro-scale transformation get underway until months have passed, you can really notice your level increasing over the scale of weeks. And that micro-scale progress can provide motivation and joy.

Making progress of any kind is inherently motivating and enjoyable. Lots of people think that motivation and enjoyment inspire progress, but it's really bidirectional and mostly the other way around: Progress inspires motivation and enjoyment. Don't underestimate how much motivation and enjoyment you'll feel when you can do exercises that you previously couldn't. Don't underestimate how much more motivation and enjoyment you'll feel when your clothes fit different.

Unless you've experienced it before while getting noticeably good in some other domain, you can't even imagine how much motivation and enjoyment you'll get when other people start noticing your results, complimenting you, sometimes even asking for your advice. You hear that enough times and it just seeps into your identity that you are a highly skilled person. And it's very easy to find motivation for and joy in training activities that help you maintain what you perceive to be a part of your identity.

Lots of people know this about physical body transformations but most people don't realize it applies to literally every skill domain. You can do this with math, coding, music, whatever the heck you want.

► How Taxing Work Becomes Fun

Initially you're missing a crap-ton of foundational knowledge. It's a rude awakening. Then you just focus on taking one bite at a time. Eventually it gets to be kind of fun. And at some point you look up and realize you've transformed. Not completely, but enough to know that it's really happening. "Wait, am I... cracked? No way. But I just did this thing that I've seen cracked people do and I wasn't able to do that before. Holy shit I'm actually getting cracked."

It's kind of like you show up to the gym weak and fat, not really looking forward to working out, but you just suck it up and do the workout and stick to the plan and eventually you get accustomed to it and it becomes kind of fun putting serious weight on the bar. Not fun in the sense of "lifting this heavy-ass weight feels so pleasurable" but in the sense of "it makes me feel legit and each time I put another plate on I feel really good about myself."

And then you notice your clothes fit differently, people make comments about how strong you look, and you're like "What? No, I'm fat and weak. But I guess I can lift some big weights now? Okay, fine, I'm no bodybuilder but yeah I guess I did put on some muscle. Holy shit, I'm actually getting ripped."

And then you realize that you're within striking distance of not just "getting" cracked, not just "getting" ripped, but you can actually close the loop in full if you stick with it and ramp up the intensity. And this is when it really feels like a video game. You're climbing this skill hierarchy, you're more advanced than most people, it's more than just a habit, it's starting to seep into your identity, you want to climb higher and see how far you can get.

And people start asking you for advice, they start looking up to you, you feel like your hard work is getting recognition, you're having a positive impact not just on your own growth but also on other people's growth, and it turns into this really positive feedback loop that continues compounding throughout the rest of life. You get caught in this virtuous cycle, it leads to more and more positive chance events you never would have anticipated, everything compounds and somehow everything is really fun despite being really taxing. Somehow the taxing becomes fun.

► Why Try When Others Could Do Better?

If you're not the smartest, highest-agency person you know, then what's the point of doing anything that they could do faster and/or better?

I'll tell you. Even the smartest, highest-agency people in the world are severely bandwidth constrained and don't get around to doing most of the things they have the potential to do. If you do something that someone else could have done faster but didn't, the reality is you've done it and they haven't.

Compound that over and over again in some niche and you get hundreds, thousands, millions of miles ahead of the fastest runners who aren't running down that niche.

Chapter 9. Learning

► The Greatest Educational Life Hack: Learning Ahead of Time

Why learn ahead of time? Because it guards you against numerous academic risks, opens all kinds of doors to career opportunities, and allows you to enter those doors earlier in life (which in turn allows you to accomplish more over the course of your career).

You know how, when you take a language class, there's often a couple kids who speak the language at home and think the class is super easy? You can do that with any other subject. When you pre-learn the material in a course before taking it at school or college, you're basically guaranteed an A in the class.

You guard yourself against all sorts of risks such as the course moving too quickly, brushing over concepts, explaining things poorly, assuming knowledge of important but frequently missing prerequisite material, not offering enough practice opportunities... There are

a hundred different ways to teach a class poorly, and most classes suffer in at least a handful of those aspects. This is especially helpful at university, when lectures are often unsuitable for a first introduction to a topic. But if you pre-learn the material, you're not depending on the teacher to teach it to you, which means you're immune to even the worst teaching.

Of course, the natural objection is “won't you be bored in class?” – but if you do super well in advanced classes, especially at university, then that opens all kinds of doors to recommendations for internships, research projects with professors, etc. It doesn't matter whether you're doing super well because you're learning in real-time or because you pre-learned the material.

When you blow a course out of the water while also interacting with the professor (answering questions in class, coming to every office hour with super insightful questions, etc.), that sets you up for a great recommendation letter – which is vital not just for high schoolers applying to college, but also for college students applying to summer research programs and graduate schools. Plus, it can open the door to working on a research project with the professor, or having them connect you to jobs, internships, and other opportunities with people in their network.

Basically, you can use pre-learning to kick off a virtuous cycle. Even if you aren't a genius, you appear to be one in everyone else's eyes, and consequently you get a ticket to those opportunities reserved for top students. Students who receive and capitalize on these opportunities can launch themselves into some of the most interesting, meaningful, and lucrative careers that are notoriously difficult to break into.

And why stop at pre-learning one year ahead? It's worth it to keep going, keep accelerating. The road always stretches farther than what you can see in front of you, and you maximize your reward by traveling as far as you can.

Let's consider math, for example. Many people think calculus is the "end of the road" for math, and that it doesn't matter if you get there many years ahead of schedule. But that's far from the truth! There are even more university-level math courses above calculus than there are high school courses below calculus.

After a single-variable calculus course (like AP Calculus BC), most serious students who study quantitative majors like math, physics, engineering, and economics have to take core "engineering math" courses including Linear Algebra, Multivariable

Calculus, Differential Equations, and Probability & Statistics (the advanced calculus-based version, not the simpler algebra-based version like AP Statistics). Beyond those core “engineering math” courses, different majors include plenty of specialized courses that branch off in various ways.

There are so many courses that a student could not fit them all into a standard 4-year undergraduate course load even if they overloaded their schedule every year – however, the more of these courses a student is able to take, the more academic opportunities and career doors are open to them in the future. (And while it’s true that students don’t need to know much beyond algebra to get a basic job a field like computer science, medicine, etc. – the people in such fields who do also know advanced math are extra valuable and in demand because they can work on projects that combine domain expertise and math.)

When a student learns a lot of advanced math ahead of time, they unlock the opportunity to dig into a wide variety of specialized fields that are usually reserved for graduates with strong mathematical foundations. This fast-tracks them towards discovering their passions, developing valuable skills in those domains, and making professional contributions early in their career,

which ultimately leads to higher levels of career accomplishment.

I'm not exaggerating here – this is actually backed up by research. On average, the faster you accelerate your learning, the sooner you get your career started, and the more you accomplish over the course of your career. For instance, in a 40-year longitudinal study of thousands of mathematically precocious students, researchers Park, Lubinski, & Benbow (2013) concluded the following:

"The relationship between age at career onset and adult productivity, particularly in science, technology, engineering, and mathematics (STEM) fields, has been the focus of several researchers throughout the last century (Dennis, 1956; Lehman, 1946, 1953; Simonton, 1988, 1997; Zuckerman, 1977), and a consistent finding is that earlier career onset is related to greater productivity and accomplishments over the course of a career. All other things being equal, an earlier career start from [academic] acceleration will allow an individual to devote more time in early adulthood to creative production, and this will result in an increased level of accomplishment over the course of one's career.

...

[In this study] Mathematically precocious students who grade skipped were more likely to pursue advanced degrees

and secure STEM accomplishments, reached these outcomes earlier, and accrued more citations and highly cited publications in STEM fields than their matched and retained intellectual peers."

► What's the Point of Being Way Ahead? It Buys You Time.

For anyone interested in a STEM career, acquiring advanced technical skills early unlocks the most valuable thing in existence. Something even money can't buy.

TIME.

If you acquire demonstrable, alien-level technical skills, you can get doors opened earlier. You don't have to wait until college to do research projects under professors, or even paid internships. If you're intentional about acquiring skills and putting yourself out there, you can kick-start a serious career before most people your age are even taking serious classes. And once you acquire early junior-level experience, that opens the door to early senior-level experience, and so on.

Compound this virtuous cycle over and over and you end up way ahead.

What's the point of being way ahead?

IT BUYS YOU TIME.

Unfortunately, lots of people misunderstand the point of compressing time. They think compressing time is about winning a rat race against your peers. But that's not really what it's about. There's a race, but the thing you're racing against is much scarier and much more powerful than any other human. You're racing against TIME ITSELF.

Time is the #1 killer of dreams and aspirations. When someone gives up on their dream, or gives up on figuring out what that dream is, it's typically a result of them losing the race against time. Pink Floyd put it best:

"You are young and life is long, and there is time to kill today.

And then one day you find ten years have got behind you.

No one told you when to run. You missed the starting gun.

And you run, and you run, to catch up with the sun, but it's sinking."

Whether you realize it yet, achieving your dreams is a race against time. Time forces convergence, and premature convergence is what kills dreams.

It's hard to understand this when you're young, before you have any sense of the wrath of time or the meaning of convergence. But no matter how many times you claim you'll never settle for something less than *ikigai*: it won't keep the sun from setting, it won't keep the time from passing, it won't keep you from increasingly desiring things that only a stable life can provide, and it won't keep you from gradually turning the dial from “explore” to “exploit.”

The further time gets ahead of you, the more likely you are to settle into a life that is “fine,” or even “good” – despite being unable to shake the feeling that you could have found something better if you had more time.

That is the point of compressing time.

That is the point of removing skill bottlenecks early.

It's about opening doors early and running down avenues that you might be interested in exploring, so that:

1) If you get the feeling the path you're going down has twisted and turned into something that's no longer a great fit for you, you can double back and explore other avenues before doors start locking behind you.

2) You can spend time trying to break down a wall instead of running through an existing door if that's something you want to do.

3) Once you find your path into a land that makes you as happy as you can imagine, you can maximize your time in that land.

That's why learning advanced math ahead of time is the greatest educational life hack for any student interested in a STEM career. Accelerating helps you find your place in the world before time closes in on you and forces you to settle for something else.

► When Does the Learning Happen?

Learning is the incremental gain in your ability to perform a tangible, reproducible skill. If you're not getting those gains, you're not learning.

Imagine signing up for tennis lessons with a personal coach. When does the learning happen? It's not when

you pay the coach the money. It's not when you watch the coach demonstrate a move. It's when you actually start doing things that you weren't able to do before. It's when you attempt a move, the coach corrects your form, and you attempt the move again with better results.

It's the same anywhere else. The keys to effective training in athletics, music, etc., are the same as the keys to effective training in any other skill-based domain (e.g., mathematics). Simply consuming information doesn't cut it. You have to actively practice the skills that you're hoping to acquire.

► There is No Such Thing as Low-Effort Learning

Talent development takes work – not just a little work, but a lot of work. There is absolutely no confusion about this in the talent development community. Can you imagine asking an athletic coach to help you become a star player using training methods that don't tire you out and make you sweat? No matter what skill is being trained, improving performance is always an effortful process.

A common theme in the science of learning is that effective learning feels like a workout with a personal trainer. It should center around deliberate practice, a type of active learning in which individualized training activities are specially chosen to improve specific aspects of performance through repetition and successive refinement. These practice activities are done entirely for the purpose of pushing one's limits and improving performance; consequently, they tend to be more effortful and less enjoyable.

Unfortunately, another common theme in the science of learning is that the perception of learning is often at odds with actual measurable learning. When using effective learning strategies, students perform better on assessments but may feel they've learned less. Why? Because effective strategies increase cognitive activation, enhancing learning despite students feeling it's harder. It's like weightlifting – the strongest people lift weights heavy enough to make them feel weak.

Many types of training methods are ineffective, but require little effort, and can therefore seem attractive to even the most well-intentioned, hardworking students because they create an illusion of comprehension. Examples include looking at notes, rereading course materials, and highlighting.

It is useful to familiarize oneself with instructional material before engaging in effortful practice, and it is also useful to revisit that material if one runs into issues while attempting to carry out the effortful practice – but it is not until effortful practice that true learning actually occurs. Familiarizing oneself with instructional material is similar to warming up before a workout: the warmup does not actually lead to improvements in strength or endurance, but it does help maximize performance and avoid injury during the workout.

► The Greatest Breakthrough in the Science of Learning Over the Last Century

The greatest breakthrough in the science of learning over the last century was characterizing the mechanics of learning in the brain. Learning is all about the interplay between working memory (WM) and long-term memory (LTM). If you understand that, then you can actually derive – from first principles – the methods of effective learning.

The goal of learning is to increase the quantity, depth, retrievability, and generalizability of concepts and skills in your long-term memory (LTM). At a physical

level, that amounts to creating strategic connections between neurons so that the brain can more easily, quickly, accurately, and reliably activate more intricate patterns of neurons. This process is known as consolidation.

Now, here's the catch: before information can be consolidated into LTM, it has to pass through working memory (WM), which has severely limited capacity. The brain's working memory capacity (WMC) represents the degree to which it can focus activation on relevant neural patterns and persistently maintain their simultaneous activation, a process known as rehearsal.

Most people can only hold about 4 chunks of coherently grouped items simultaneously in WM, and only for about 20 seconds. And that assumes they aren't needing to perform any mental manipulation of those items – if they do, then fewer items can be held due to competition for limited processing resources. (Note that this is an emergent behavior of a more complicated underlying mechanism: the actual WM limitation is not a fixed number of storage units, but rather, the ability to sustain relevant neural activity while suppressing interference from irrelevant activity.)

Limited capacity makes WMC a bottleneck in the transfer of information into LTM. When the cognitive load of a learning task exceeds your WMC, you experience cognitive overload and are not able to complete the task. Even if you do not experience full overload, a heavy load will decrease your performance and slow down your learning in a way that is NOT a desirable difficulty.

However, once you learn a task to a sufficient level of performance, the impact of WMC on task performance is diminished because the information processing that's required to perform the task has been transferred into long-term memory, where it can be recalled by WM without increasing the actual load placed on WM.

So, for each concept or skill you want to learn:

1. it needs to be introduced after the prerequisites have been learned (so that the prerequisite knowledge can be pulled from long-term memory without taxing WM),
2. it needs to be broken down into bite-sized pieces small enough that no piece overloads your WM, and

3. you need to get enough practice to achieve mastery on each piece (and that amount of practice may vary depending on the particular learning task).

But also, even if you do all the above perfectly, you still have to deal with forgetting. The representations in LTM gradually, over time, decay and become harder to retrieve if they are not used, resulting in forgetting.

The solution to forgetting is review – and not just passively re-ingesting information, but actively retrieving it, unassisted, from LTM. Each time you successfully actively retrieve fuzzy information from LTM, you physically refresh and deepen the corresponding neural representation in your brain. But that doesn't happen if you just passively re-ingest the information through your senses instead of actively retrieving it from LTM.

► “Following Along” Versus Learning

It's common to think that “following along” is the same as learning – like, if you can follow along with a video, book, lecture, whatever, without feeling confused, then you're learning. While this might “feel” like learning, it's not. The feeling is completely artificial.

Comfortable fluency in consuming information is not a proxy for actual learning. Learning is a positive change in long-term memory, and you haven't learned unless you're able to consistently reproduce the information you consumed and use it to solve problems. This doesn't happen when you just "follow along," even if you understand perfectly. That comfortable fluency you feel while following along is arising from the fact that the surrounding context is already on your mind – you're not made to pull it from long-term memory.

When you feel like you're absorbing information while passively following along, what you perceive is information sitting in your working memory, not your long-term memory. If you want to test whether information is in your long-term memory (i.e., whether you've actually retained it), you have to actively attempt to retrieve it when it's not already at the front of your mind. You have to put yourself in the position where it's not already in your working memory, and the only way to pull it out of your brain is from long-term memory.

If you don't practice retrieving information from memory, it dissipates quickly and almost entirely. Have you ever had the experience of being unable to remember something despite repeated exposures,

because you keep automatically looking it up from a reference instead of trying to retrieve it from memory? That's happened to me an embarrassing number of times with addresses, phone numbers, directions, etc. And any books you read, movies you watch – the only ones you remember in proper detail are the ones you periodically think about and replay in your head. If you just consume and don't reproduce then you forget almost entirely. I can't tell you how many times I've watched a movie and didn't even realize I'd seen it before until I got 20 minutes in and something felt familiar. And even then I could barely remember anything about the rest of the movie, just that it felt a bit familiar.

Retrieval is the act of pulling information from long-term memory into working memory. Practicing retrieval under challenging but achievable conditions (e.g., when your memory has gotten fuzzy or there is less priming) is what increases your ability to remember and use information. Each time you successfully recall a fuzzy memory, it stays intact longer before getting fuzzy again. Each time you successfully recall a memory with less priming, its recall becomes less dependent on priming in the future.

But if you don't practice retrieval, then this doesn't happen. The information quickly dissipates. It stays with you only briefly – just long enough to trick you into thinking it'll stick with you, when it's really on the way out the door. But, of course, you don't notice that it's gone if you're not actually testing whether it's there.

Consuming information without practicing reproducing it can produce an artificial feeling of fluency while the information is held and manipulated in working memory, but since retrieval practice is not occurring to extend the information's retention, the information dissipates quickly. The fact that it's in working memory can trick you into thinking it's going to stay there, but it doesn't.

Once it's gone, the only way you can bring it back without reloading it from an external reference is if you're able to retrieve the information from long-term memory. But if you don't practice retrieval, you won't be able to successfully retrieve. When all you do is consume information, you put yourself in a situation where the only way to load it back into your working memory is to re-consume it. This is why learning really amounts to increasing your ability to recall information from long-term memory unassisted, an action that can be trained by repeatedly performing said action in

gradually more challenging contexts, just like strength training.

Additionally, learners typically do not process all the key information as they consume, but they are unaware of this until they attempt to answer a question or solve a problem that requires them to retrieve some key information from memory. At that point, they realize that they never fully processed that piece of information to begin with, and they have to go back to find and properly process it. The same applies to generalization: learners typically do not fully generalize what they've consumed, but they are unaware until they attempt to answer questions or solve problems that require them to generalize their understanding.

The way to avoid this problem and maximize your learning is to switch over to active problem-solving immediately after consuming a minimum effective dose of information. I know that might feel a bit jarring, like it's slowing you down, but it isn't actually slowing down your learning – it's only exposing the fact that your perception of learning does not accurately reflect actual learning. Really, it's speeding up your actual learning, and the only thing it's slowing down is your perception of learning.

Now, you might say “but I had learned so much, and I had it down pat, and then I forgot it all when I focused my effort on solving a problem.” But the thing is, if you can’t retrieve that information from memory at the snap of a finger, after thinking about other stuff or zooming in to focus on a specific problem, it means you didn’t really have it down pat.

► The Need to Practice Retrieving From Memory

If you want to actually retain information you consume, you need to practice retrieving it from memory, not just re-consuming it. You have retained information if and only if you can regenerate it from memory. So naturally that is the exercise that must be practiced.

Yes, it is also helpful to make connections between isolated pieces of information.

Yes, those connections are themselves information.

Yes, if you want to retain the connections, then you have to exercise their retrieval, not just re-consume them from an external source.

Yes, this also applies to connections between connections.

Yes, this also applies to skills, concepts, procedures, applications, etc.

If you want to be able to generate it from your head, then you need to practice generating it from your head.

► The Vicious Cycle of Forgetting

Forgetting is frustrating. After putting forth the effort to learn something, who wants to waste time re-learning it later? To mitigate the effects of forgetting, it might feel helpful to solve problems alongside reference material.

But there's an issue: when you continually look back at a reference, the information doesn't stay in your brain. You hold the information in short-term memory, but only temporarily – it dissipates after your focus redirects elsewhere. The reference material becomes a crutch, and you're lost without it. You might think you need to spend more time reviewing the reference, but

really you just need to review properly, pulling information from memory.

Even people who are serious about their learning sometimes fall into this vicious cycle of forgetting. They might take great notes and then refer back to those notes all the time instead of trying to pull the information from memory.

The thing is, if you try to keep the information close by taking great notes that you can reference all the time, that just PREVENTS you from truly retaining it. That might seem counterintuitive, but it's actually pretty obvious.

What's the thing that transfers information to long-term memory? Retrieving from memory.

When you take notes, you know what you're NOT doing? Retrieving from memory.

When you take great notes and constantly refer back to them, you know what you're STILL not doing? Retrieving from memory.

Retrieval is not just any loading of information into your brain. Retrieval is the specific action of “pulling” information from one part of your brain (long-term

memory) to another part of your brain (working memory). It's like your brain is lifting a weight off the ground of long-term memory and raising it up into working memory. The fuzzier the memory, the heavier the weight – but just like weightlifting, as you practice lifting heavy weights, you get stronger, i.e., your brain becomes more easily able to activate the pattern of neurons that represent the information stored in long-term memory.

If you load information into working memory by looking at reference material instead of pulling from long-term memory, then you're not strengthening your retention. It's like you're going to the gym to lift weights, but you're just going through the motions and letting your spotter lift the weight for you. No strength is being developed. You end up throwing yourself into a vicious cycle of forgetting:

- You keep looking back at a reference because you can't remember things.
- You can't remember things because you're not transferring them to long-term memory.
- You're not transferring them to memory because you're not practicing retrieving them from memory.

- You're not retrieving them from memory because you're always looking back at the damn reference!

As you spiral into this vicious cycle of forgetting, your whole learning process completely falls apart. You learn slower, forget faster, and miss out on making connections that would deepen your understanding.

The only way to break this vicious cycle of forgetting is to engage in retrieval practice. Initially, that may seem like a paradox: “how can I engage in retrieval practice if I’m unable to retrieve?” But it’s not a paradox at all. Back to weightlifting – you just need to treat the reference material like a spotter. You try your hardest to lift the weight, and if you can’t, the spotter intervenes as a last resort, giving you just enough assistance to get you over the edge of lifting the weight. The spotter should be doing as little as possible while ensuring that you manage to eke out a successful rep.

In the same way, whenever you're about to look up information that you've seen before, that you would like to stick in your brain – always, always, always try your best to recall it from memory. DO NOT default to looking it up. If you cannot manage to retrieve it despite trying your best, then it’s okay to peek back at your reference material, but only as a last resort.

Peek once – just a little bit, just the tiniest bit of priming, just that specific piece of info that you were trying to remember, nothing else – and then close the reference, re-pull the information from memory, and try to recall the rest and proceed forward as far as possible without peeking back at the reference again. Never, ever transcribe from the reference. Your brain is lifting a weight and the reference material is your spotter – it's there as a last resort to help you get the weight up, only when you absolutely can't get it up yourself, and the amount of help should be kept to the bare minimum.

The goal is to wean yourself off of reference material, using it as sparsely as possible, until you don't need it at all. This may be very challenging if you've been relying on reference material as a crutch, but it's the only way out of the vicious cycle.

And you know what helps you wean yourself off of a crutch? Not having easy access to it. As long as you have a reasonable way to look up a piece of information if you forget it, then it's not worth optimizing for convenience. You WANT it to feel annoying to look stuff up, so that you're incentivized not to have to do that. And if you're engaging in proper retrieval

practice, you won't have to spend much time looking stuff up anyway.

► Recall First, Reason Second

If you want to build automaticity, then you need to practice automatic recall. That sounds obvious, but a common mistake even among serious learners is "I'll derive/reason the results from scratch to create a cheat sheet, and then refer to the cheat sheet during practice problems that involve applying those results."

It's good to practice deriving/reasoning results from scratch when doing so is within grasp of your skill level and you haven't done so in a while. However, that *does not count* as automatic recall practice on the result itself. It's just like how recalculating your times tables on a reference sheet that you constantly refer to will prevent you from developing automaticity with multiplication.

So here's what I recommend to do: instead of deriving/reasoning a result before applying it, apply it first and then derive/reason it after. Force yourself to recall the result from memory, and then justify the result afterwards. (By the way, this applies broadly, not just in math.)

► The Vicious Cycle of Context Overload

One of the least efficient ways to learn is to attack the most challenging “authentic” or real-world problem context right from the get-go. It creates a vicious cycle where you

1. struggle with the problem due to the additional complexity,
2. take so long to solve it that there’s time pressure to move on to new material,
3. struggle even more with the new material because you didn’t get enough reps in to master the previous skill,
4. and then the cycle repeats again starting from (1).

It’s a lot more efficient to strip skills down to the simplest possible context, get some reps in, and gradually increase the complexity of the context. When you get that scaffolding right, you can complete each rep reasonably quickly because the challenge is matched to your skill level, and you end up climbing

the skill tree even faster while building a solid foundation.

► Prereq Yo' Self Before You Wreck Yo' Self

Being out of your depth in skill training is a huge problem because your learning progress grinds to a halt. It's not like you're on a train that left late from the station. It's like you're on a train that's not even moving. The train might even start moving backwards: if you're so far out of your depth that you're just flailing around on new skills, then you're likely not absorbing much implicit review on the component skills you've previously learned, and as a result you're forgetting them.

The only way to get the train moving forwards again is to drop down and work at a level that's appropriate for you. You need to get yourself in a situation where you're successfully accomplishing new challenges. If you don't, the situation compounds into a vicious cycle, getting worse and worse. You continually come into new skills less and less prepared, getting more and more out of your depth.

When you skip prerequisites or otherwise don't master them, you don't have those skills available for

automatic execution, so when you're attempting to execute a new skill that depends on them, you exhaust all your focus and effort attempting to carry out the prerequisites. You might have the bandwidth to focus on a single prerequisite if you put all your focus and effort behind it, but not to execute multiple prerequisites in parallel, much less monitor and control the entire complex operation at a high level.

The only solution is to hammer in your prerequisite skills until they're rock-solid and easy to execute. That way, no individual prerequisite takes up much focus or effort, and you can execute multiple prerequisites in parallel while seeing the forest for the trees and strategizing at a high level.

► Filling In Your Foundations is the Difference Between Omitted Steps Being Minor Potholes vs Uncrossable Chasms

If you fill in your foundations, your missing prerequisite knowledge, then you'll be able to fill in the steps that are omitted in those cool technical books and projects you've wanted to work through. Filling in your foundations is the difference between those omitted steps being minor potholes you can get past

with a bit of a bumpy ride, versus uncrossable chasms that stop you dead in your tracks.

Too many people go straight for the coolest, most advanced books/projects, struggle to fill in the gaps in the advanced material, and then give up, thinking they're not smart enough. What they don't realize is that this is typically a solvable problem. Yes, fixed cognitive differences (IQ, working memory, generalization ability, etc.) are a thing, but think of it this way:

If you go into the gym, put a couple plates on the bar, and fail to lift it, then does it make sense to say "I guess I don't have enough physical strength encoded into my genes" and give up?" Of course not! You just need to train up your musculature, and you can do that by starting at your level and gradually adding weight to the bar until you reach your goal.

There is a point where this breaks down and genes become important, but it's a lot further than most people think. Most people vastly underestimate how far they can go with effective training. Can you become a world-record powerlifter if you're naturally skinny? Probably not. But can you turn yourself into one of the strongest people at the gym? Probably.

► Plan Your Broad-Strokes Journey Top-Down, but Carry Out the Granular Steps Bottom-Up

The top-down approach can be useful for planning a broad-strokes learning journey towards a goal. For instance, if you want to learn machine learning, then you can think top-down to figure out what fields of math you need to learn in order for machine learning to become accessible to you. You'll find that you absolutely need to learn calculus, linear algebra, and probability & statistics, and you can skip stuff like abstract algebra, number theory, etc.

However, the granular steps of the journey, the actual learning, needs to be carried out bottom-up. For instance, are you really going master computing neural net weight gradients via backpropagation by asking "what does that squiggly 'd' mean," "why do you have to chain-multiply the derivatives like that," "how do you calculate the derivative of any activation function," etc., all the way down to the depths of whatever is the last piece of math you've mastered?

No, all you're going to do with those questions is create a roadmap of what you need to learn. Which is essentially a calculus course. Except your roadmap will

be terrible because you don't actually know the subject yourself – it will have all sorts of gaps that you don't even realize are missing because, which is to be expected given that you don't actually know the subject.

You'll try to climb back up the skill tree implied by your incomplete roadmap and you'll repeatedly get stuck trying to climb up to the next branch that you can't reach because there are prerequisites that you don't realize you're missing.

Most people in this situation will eventually just give up due to all the friction. Only those who have extremely outsized perseverance and generalization ability have any chance of fighting through and making it to the other side. And even then, it will take longer (and they'll likely end up with more holes in their knowledge) than if they just sucked it up and worked through a well-sequenced calculus course.

► The Efficient Learning Loop

All the information you consume while learning, every problem you work out, it comes with the cost of using up more of your time. It has to be worth it. If you're

looking to maximize your learning efficiency, then what want is

- streamlined instruction ("no BS, just give it to me straight" explanations),
- most of the time focused on active problem-solving, and
- continually switching back and forth between instruction and problem-solving quickly enough that your attention span doesn't run out.

It's a continual cycle of minimum effective doses:

- minimum effective dose of streamlined "no BS, just give it to me straight" instruction,
- followed by minimum effective dose of problem-solving,
- then back to minimum effective dose of instruction to prepare you for slightly more challenging problems,
- followed by minimum effective dose on said problems,

and so on.

► Don't Bloat the Feedback Loop

Lots of people hear "repetitive" as a synonym for "mindless" when in fact repetition and mindlessness are completely orthogonal to each other. "Mindful repetition" is what you want – mindful because you're practicing just beyond your repertoire, getting feedback on every attempt so that you can improve on the next attempt. I.e., deliberate practice.

But even then, it's easy to go overboard and bloat the repetition with supplemental tasks that, while intended to support mindfulness, end up throttling the volume of practice, creating a severe bottleneck in the learning process. The feedback loop becomes too slow.

For instance, one of the worst offenders is "think-pair-share," which ensures that the amount of time wasted scales with the number of students in the class. If you do think-pair-share in a class full of 30 students, it's easy to burn most of the class time dragging out a single repetition, which is next to nothing in terms of training volume.

Kind of like if you go to the gym for a workout and only do 1 pushup the whole hour, it doesn't matter how perfect your form was on that pushup: you only did 1 pushup which is next to nothing. You're not going to get stronger without a serious volume of reps.

► Some Pitfalls to Watch Out For When Learning From Projects

Once you have acquired the foundational skills for a cool project that pulls those skills together, then by all means, go for it! However, here are some pitfalls to watch out for.

1) Don't use projects as a way to acquire fundamental skills. This is typically inefficient due to the higher-complexity setting. Instead, learn skills in a more scaffolded context and then use projects to pull those skills together.

2) Make sure the projects are guided. Don't put yourself in a position where you're spinning your wheels getting nowhere. That's just wasting time that would be better spent learning new stuff in a structured environment. Basically, don't fall into the "discovery learning" trap.

3) *Don't let the projects cut too much into your foundational skill-building.* Projects can be fun and productive for pulling existing skills together, but you don't want to let them become a distraction from further fundamental skill-building. For instance, you can imagine someone who learns algebra and then spends months doing algebra-based projects instead of learning the fundamentals of calculus, linear algebra, etc. And then once they gain these higher-level fundamental skills, many of these algebra-based projects are rendered trivial. While I wouldn't go as far to say those trivial projects were a complete waste of time, it's definitely inefficient compared to the alternative, which is to first focus on fundamental skill-building first, and then switch over to more focus on projects as you get closer to the edge of your field.

► Review Should Feel Challenging

Students often expect review to be easy. At least part of this expectation is due to conditioning: in school, when the teacher says it's a "review day," they might as well call it an "off day." But if you're actually trying to maximize learning efficiency, then reviews should feel tough. Why? Because recalling tricky information improves memory, while recalling easy information doesn't.

That's the whole idea behind spaced repetition: your memory has to get a bit fuzzy before the next repetition, otherwise the desired effect – slowing the rate of forgetting and remembering longer next time – doesn't happen (or at least not nearly as much). It's the act of successfully retrieving fuzzy memory, not clear memory, that extends the memory duration.

And if review problems are easy, not actually extending your memory duration, then what's the point? It's better to learn something new. A maximum-efficiency teacher will intentionally let your memory fade a bit before review so that the act of refreshing your memory actually deepens your long-term encoding, and they'll use the extra time to cover more new material.

In general, learning requires introducing “desirable difficulties” into the recall process, making it tough yet achievable. During an initial lesson, the desirable difficulty comes from manipulating new information. During review, the desirable difficulty comes from successfully recalling fuzzy memory – you've already learned how to manipulate the information, but now you're practicing in a trickier setting where enough time has passed for your memory to fade.

Consequently, reviews should feel as mentally taxing as initial learning. You're getting better, but the bar for success also is getting higher. Your brain has to hold the memory for a longer period of time – just like a muscle holding a weight.

The analogy to weightlifting runs deep. In the context of spaced repetition, the way you increase the weight is by waiting longer before retrieving the knowledge again. But you also don't want to wait too long before retrieving the knowledge, because then you won't be able to successfully retrieve it. This is just like how in weightlifting, you need to increase the weight to the point where you struggle to lift it, but you are able to overcome the struggle. That's how you build muscle, and that's also how you build long-term memory. Spaced repetition = "wait"lifting.

► Learn Like You Lift

Spaced repetition is so similar to weight training that you might as well call it wait training. You're lifting information off the floor of long-term memory and raising it up into working memory. The fuzzier that information, the harder it is to lift. The wait creates the weight. And just like successfully lifting a heavy weight strengthens muscles, successfully recalling fuzzy

information (lengthy wait) strengthens memory. (There are other ways to create weight too, such as interleaving between different contexts and layering more complex contexts on top of simpler ones.)

But you have to retrieve from memory. Spaced “re-reading” doesn’t count – that’s like letting your spotter lift the weight for you. The movement you’re trying to train is the lift from long-term memory into working memory. Re-reading brings information into working memory, but it doesn’t exercise the lift, and improving the lift is what improves retention.

The only time the spotter should help you lift the weight is when you can’t lift it despite trying your hardest. And even then, the spotter should only give you just enough assistance to get you over the edge of lifting the weight. The spotter should be doing as little as possible while ensuring that you manage to eke out a successful rep.

In the same way, the only time you should look at reference material during review is if you can’t recall something after trying your hardest. And while it’s okay to check reference material as a last resort, you should only peek once for a cue, and then try to recall the rest without looking again.

This weightlifting analogy generalizes beyond spaced repetition: in general, learning requires introducing “desirable difficulties” into the recall process, making it tough yet achievable. But remember: just like little strength is built by attempting and failing to lift a too-heavy weight, little knowledge is built by attempting and failing a too-difficult learning task. Even a desirable difficulty becomes undesirable if the learner is unable to overcome it. Additionally, not all difficulties are desirable. Plenty of difficulties are undesirable even if they can be overcome. For instance: sleep deprivation. Even if you overcome it, it’s not a productive challenge for building strength or knowledge.

► Schooling Versus Talent Development

The fundamental principles of effective training are similar across domains. But you only see this if you're actually optimizing for performance. That's what's done in the field of talent development: an individual's performance is to be maximized, so the methods used during practice are those that most efficiently convert effort into performance improvements. But elsewhere in education, the norm seems to be optimizing for fun and entertainment while, as a secondary concern,

meeting some low bar for shallowly learning some surface-level basic skills.

Schooling and talent development are completely different things. In schooling, students are grouped primarily by age, rather than ability, and each group progresses through the curriculum in lockstep. Each member of the group engages in the same tasks at the same time, and it is expected that different students will learn skills to different levels.

In talent development, students progress through skills at different rates, but learn skills to the same threshold of performance. Their progress is measured not by their level of learning in courses that they have taken, but rather by how advanced the skills are that they can execute to a sufficient threshold of performance. This is accomplished through completely individualized instruction. Learning tasks are chosen based on the specific needs of individual students, each student must learn each skill to a sufficient level of mastery before moving on to more advanced skills.

This contrast between schooling and talent development is not new. Researchers have known about it for many decades. For instance:

- *“Schools do not seem to have a great tolerance for students who are out of phase with other students in their learning process.” –Benjamin Bloom, 1985*
- *“In general, school learning emphasizes group learning and the subject or skills to be learned. Talent development typically emphasizes the individual and his or her progress in a particular activity.” –Bloom & Sosniak, 1981*

At the heart of it all, here’s the core difference: Outside talent development, lots of people in education disagree with the premise of maximizing learning. Whereas in talent development, an individual’s performance is to be maximized, so the methods used during practice are those that most efficiently convert effort into performance improvements.

Here’s a concrete example. On one hand, “testing” and “repetition” have become dirty words in education. However, practice testing and distributed practice (also known as spaced repetition) are widely understood by researchers to be two of the most effective practice techniques. Moreover, deliberate practice – which has been shown to be one of the most prominent underlying factors responsible for individual differences in performance, even among highly talented elite performers – is centered around using

repetitious training activities to refine whatever skills move the needle most on a student's overall performance.

So what gives? Why are there debates about scientifically proven learning techniques like testing and repetition? Because lots of people in education disagree with the premise of maximizing learning. The debates aren't about whether testing and repetition are effective learning techniques – the debates are about whether education should seek to maximize students' learning.

Outside of talent development, the typical approach to education involves maximizing other things like fun and entertainment while, as a secondary concern, meeting some low bar for shallowly learning some surface-level basic skills. I'll admit that de-prioritizing talent development ends up working out okay when students aren't expected to achieve a high level of success. For instance, if every student in gym class were expected to be able to do a backflip by the end of the year, things would have to change – but the expectations are so low that meeting them does not require talent development.

But serious skill development is different. Take math class, for example. Students are typically expected to

achieve a relatively high level of success in math: many years of courses increasing in difficulty, culminating in at least algebra, typically pre-calculus, often calculus, and sometimes even higher than that. As a result, in math, de-prioritizing talent development leads to major issues. When students do the mathematical equivalent of playing kickball during class, and then are expected to do the mathematical equivalent of a backflip at the end of the year, it's easy to see how struggle and general negative feelings can arise.

► Learning Doesn't Have to Be Synchronized for Camaraderie to Occur

You don't need to be moving in lockstep with your peers. There's plenty of shared experience to bond over when you're all working hard to climb the same skill tree along whatever individualized path and starting point is most efficient for you. You don't need to be taking each step and climbing each branch in sync.

► A Sanity Check for Effective Study Techniques

Every time you study, imagine the Grim Reaper is going to show up at the end of your session to quiz you on what you covered, and if there's any question you can't answer correctly, you die. Whatever study techniques you'd use in that situation, you better be using them already.

(I should emphasize that getting stuff wrong occasionally is totally okay and expected. The thought experiment here is more about the actions than the results: truly optimal study strategies would be conserved even in the theoretical highest-stakes scenario. When a high-accountability situation induces change in learning techniques, it exposes that the emperor was originally wearing no clothes.)

Chapter 10. Expertise

► The Driving Force Behind Expertise is Long-Term Memory

An important mechanism behind expert performance is "perceptual learning," the ability to extract key features from complex environments while filtering out irrelevant noise.

Whereas a beginner perceives individual isolated pieces of information, an expert perceives "chunks" of information organized into meaningful patterns and structures. These chunks are physically encoded as wiring in the expert's long-term memory, and they are the building blocks that make up the expert's representation of what they're looking at in working memory.

This is the critical point: it's not just that the expert actively thinks about things differently from the novice. It's that the expert literally *perceives* them differently to begin with. The same sensory signals are

processed into different working memory representations. What the expert holds in working memory is very different from what the novice holds in working memory.

What's more, these memory representations in the expert brain can also include predictive information that's not in the original stimulus. The stimulus activates a neural representation, and that neural representation may contain more information than is in the stimulus. It may include missing details or even future events associated with the stimulus. The expert will perceive all of this additional information while perceiving the stimulus, and they can use this information to make better decisions and sense/correct mistakes before they fully manifest.

For instance, consider an expert pianist. Without even pressing the keys of the piano, they can predict what sounds will be made purely based on hand position. If the hand position is wrong, they can quickly detect and correct the issue to avoid playing a wrong note. A beginner, on the other hand, will not realize that there is an issue until they hear the wrong note.

► Learning is Memory

Memory is the raw material of understanding, and we use it to build magnificent information structures in our brains. But when the towers grow high enough, intricate enough, we often look upon them as so magnificent, so divine, that we later claim they could not have emerged out of raw mechanical memory.

Our claims of "learning is not memorization" may sound poetic in the abstract, but when asked to break them down concretely, we find our explanations start to sound more and more like "tower-building is not nailing wood, not welding steel, not load-balancing, it is none of those mechanical things, it is the emergent grace of brushing the heavens." Which is completely devoid of instructional value, and entirely unhelpful to any aspiring tower-builder, aside from possibly inspiring a fleeting spark of motivation that spends itself searching for a place to spend itself.

At the end of the day, learning is memory. Understanding amounts to memory that is well-connected and deeply ingrained.

The difference between "just memorizing" and "deeply understanding" isn't the substrate of the representation, it's the depth of the representation.

Deep understanding consists of not only declarative facts, but also connections that link facts into related groups or "chunks" (think: concepts), connections that link smaller chunks into bigger chunks, and so on -- as well as procedures for operating on chunks (think: skills), connections that chunk sub-procedures into meta-procedures, and so on.

This is all raw mechanical memory. It's just storage and retrieval of information. The point of building superior representation is to build superior recall abilities, including broadening and fine-tuning the range of stimuli that activate the information. If someone is "just memorizing" as opposed to "deeply understanding," it really means they haven't stored enough information in memory.

"Learning is memory" might feel obvious, but many learners don't fully grasp the implications. If you don't realize that learning is memory, then you won't realize that the most effective way to learn is to use memory-supporting training techniques.

It's easy to get confused, thinking: "Truly understanding something is different from just memorizing it, so learning doesn't require memory-focused techniques like retrieval practice, spaced review, and interleaving (mixed practice). Those

are about memorization, not true understanding." And if that's what you think, then you'll likely shirk the hard work required to build memory, use fun/easy but ineffective training techniques instead, and end up not actually learning much.

I used to think resistance to "learning is memory" was genuine confusion, but now I think it's mostly laziness. If you accept that learning is memory, then you have to accept that maximizing learning requires memory-supporting training techniques. But those techniques are highly effortful and measurable, which make them unattractive to low-accountability / low-effort folks. The only way to reject the premise is to latch onto the idea that "understanding" is some supernatural thing that can't arise from raw mechanical memory. Which is problematic because there's decades of research showing how expertise arises from having lots of domain-specific information encoded into memory that is well-connected and deeply ingrained.

(A response to the most common genuine objection: Even learning to generate new ideas amounts to searching a space of possibilities, combining pieces of memory in ways that haven't been combined before. Now you might say "aha, the skill of searching/combining is something other than

memory," but let me ask you: when a someone trains the skill of coming up with novel ideas, such as a grad student learning to come up with research ideas that contribute to the cutting edge of knowledge in the field, where is that skill stored for future use? In memory.)

► Learning is About Bridge-Building, Not Jumping

In the cognitive science literature, there's a lack of research evidence that people can actually increase their "raw" working memory capacity and generalization ability through training general problem-solving skills. There's a mountain of evidence that you can increase the number of examples and problem-solving experiences in a student's knowledge base, but a lack of evidence that you can increase the student's ability to generalize from those examples. (For a brief overview, see Sweller, Clark, & Kirschner, 2010: *Teaching General Problem-Solving Skills Is Not a Substitute for, or a Viable Addition to, Teaching Mathematics*.)

However, there IS research evidence that you can effectively turn long-term memory into an extension of

working memory if you acquire domain-specific foundational skills and develop them to the point of automaticity. And as you layer more advanced skills on top, those foundational skills naturally get compressed into more generalizable neural representations that can be applied more flexibly across different contexts.

The phenomenon of turning long-term memory into an extension of working memory was observed as early as 1899 by Bryan and Harter: *"The learner must come to do with one stroke of attention what now requires half a dozen, and presently in one still more inclusive stroke, what now requires thirty-six. He must systematize the work to be done and must acquire a system of automatic habits corresponding to the system of tasks. When he has done this he is master of the situation in his [occupational or professional] field. ... Finally, his whole array of habits is swiftly obedient to serve in the solution of new problems. Automatism is not genius, but it is the hands and feet of genius."*

For about a century, the evidence was gradually bolstered, and the 1980s marked a tipping point that led to an explosion of research interest and behavioral studies. For instance, Chase and Ericsson (1982) found that *"rapid access to a sizeable set of knowledge structures that have been stored in directly retrievable locations in long-term memory ... produce[s] an effective increase in the*

working memory capacity for that knowledge base." And as explained by Unsworth & Engle (2005), "[I]ndividual differences in WM capacity occur in tasks requiring some form of control, with little difference appearing on tasks that required relatively automatic processing."

Today, in addition to behavioral studies, this phenomenon of turning long-term memory into an extension of working memory can be physically observed in neuroimaging. At a physical level in the brain, automaticity involves developing strategic neural connections that reduce the amount of effort that the brain has to expend to activate patterns of neurons. This has been observed, for instance, by Shamloo & Helie (2016), who studied functional magnetic resonance imaging (fMRI) brain scans of participants performing tasks with and without automaticity and found that only subjects who had developed automaticity were able to perform tasks without disruption to background thought processes.

In summary, it appears that skill development all comes down to building domain-specific chunks in long-term memory that allow you to bring more information into working memory without actually increasing the amount of cognitive effort you have to put forth to rehearse that neural activation. In other words, the way you increase your ability to make

mental leaps is not actually by jumping farther, but rather, by building bridges that reduce the distance you need to jump.

► It's All About Domain Knowledge

There is a lack of evidence that you can increase someone's ability to problem-solve by any other method than equipping them with more domain knowledge. If this seems counterintuitive, it's because 1) we tend to vastly underestimate how much of our problem-solving ability comes down to accumulated domain expertise, and 2) courses typically don't cover all the domain expertise that's needed.

For instance, an opponent might argue that debugging code is more than just domain knowledge: it's about the general ability to tinker around with things, and you don't learn that at school. However, what's really happening is they've amassed a vast zoo of tinkering techniques -- printing out specific relevant info, dumping all info into a logfile, setting a breakpoint, setting a conditional breakpoint, stepping through code, refactoring messy code in the hopes that reducing complexity will make the problem go away naturally or at least be easier to identify, reproducing the issue on a simpler case, etc. -- and they've gotten

tons of practice selecting and using these techniques in different problem contexts.

This is domain knowledge. It comes more quickly to some people than others, but there's nothing magical about it. If you learn these granular techniques, get really good at executing them in technically complex settings, and get really good at matching debugging problems to techniques, you'll be an expert debugger.

Unfortunately, this stuff is typically not covered in CS courses, and students typically get very little structured practice with it. It seems like most CS grads barely even know how to use a breakpoint, and that's just scratching the surface of debugging techniques. So what happens is people will sometimes learn an incomprehensive structured curriculum, learn a bunch of stuff outside of it, and then think that the stuff they learned is magical and cannot be taught in a structured manner -- when, in fact, it can, and the problem is just that nobody has put in the hard work to create that structured curriculum yet.

► Your Brain Should Surprise You

If you're not periodically surprised by the capabilities of your own brain, you're probably training it wrong.

Your brain is the most powerful machine learning model you'll ever use -- it's the deepest neural net in existence, with internal structure sophisticated beyond anyone's comprehension. If you're training it effectively, you will feel that power. Take advantage of it. You'll be surprised what performance feats you're able to pull off when you develop skills to the point of automaticity.

► Turn The Magical Into The Mechanical

When a process or phenomenon feels magical, that's typically an indication you don't really understand what's happening under the hood. You don't have the nuts and bolts in your head, so the outcome just feels like a magical result from some sorcerous incantation.

This magic can be exciting, even inspiring – but many learners make the mistake of leaning further into activities that increase this magical feeling, when really, the goal of learning is to turn the magical into the mechanical.

The shift from magical to mechanical doesn't diminish the beauty. It just changes the lens. You stop being a spectator of wonders and start becoming a builder of them.