
NO BS AI

CLASS MATERIALS

Week 1: Learning to Fall

Pilot Class // March 2026

Instructor: RJ // Platform: Ripley

The Goblin Mining Cave

"The best defense is not to be there."

-- The Karate Kid (and every goblin who ever ducked a wrench)

DO NOT LICK THE TRANSCRIPTS

Miners Present: Jaime Buckley, Armando, Jackie, T'lara, Delana, Lisa/Anne, Laurel

DEEPER INTO THE MINE WE GO

Part 1: Welcome & Setting the Scene

RJ: When you change the ambiance a little bit, people actually remember better. There's lots of wayfinding in this room, lots of easter eggs — start looking for them, you'll find them.

I'm so excited right now. Tell me your question, my friend. I love that I can press my hand and the hand goes up.

Student: Do I need to get a yellow pad to write down questions, or what is this?

RJ: I think this is going to be pretty laid back. You can raise your hand if you feel like it, or you can just ask me questions. We are an informal group. I feel like people learn best when they can just ask the question when they have it. But if you want to get a legal pad and raise your hand and do all of that, I embrace it.

Student: I'm gonna do that just because if one of you are talking, I don't want to interrupt. So I'll write down my question so I can bring it back.

RJ: That'll work. For your memory — yes, paper has perfect memory, doesn't it? And we're gonna be learning about that a little later on. I'm glad you brought it up.

Alright, looks like somebody was trying to get in... There will be people that won't be attending these classes at all but are totally in the class because they got stuck. I get it. Videos will always be here. I also will be here — I'll just kind of keep a tab open, hang out in this room, and if somebody shows up with a question, if you're working on something, I want to give you what you need every time if I can.

I do have meetings occasionally, I do have to disappear occasionally. But my friends, helping you succeed — that's the only thing I'm here to do. If I have to tear everything down that I've put up and do it all again to help, I will. Because this is the pilot. We can do anything we damn wanna.

Cool! Alright my friends, well, it's 11 o'clock AM and by God we start on time. By goblin, we're going!

Part 2: Introduction to the Environment

RJ: I'm not going to introduce myself. I am RJ. Welcome to the Goblin Mining Cave. I'm particularly given to fantasy realms, I don't know if you could tell that.

I want to explain some of the environment that we're in. First of all, there are tubes sprouting out of my head. The goblins have trouble with Wi-Fi — signal can't pass through the walls, so the pneumatic tubes and the brass tubes are their intranet. You'll see little messages floating through. You might also see an occasional sandwich — don't worry about it, the goblins handle it. They pass things to each other all the time through these tubes. We do have to have them cleaned, but we have Gary do that because he's our... well, you'll learn about Gary anyway.

I'll be explaining some internet concepts actually with the tubes, because it's just a good way to look at things. I am a person who is visual. It took me a long time — I was in actual college for some of this.

Armando and Laurel join the room

RJ: Armando, hey! Good to see you, my friend. And Laura's here — welcome, welcome! Our friend Laurel, welcome!

So I was talking about the internet being a series of tubes behind me. Yes, that was a direct intention, because "the internet is a series of tubes" — oh, it's always made me laugh.

Part 3: [AI]Kido™ -- The Core Concept

RJ: Today, we're gonna be teaching [AI]Kido™. Aikido is something that I came up with because these goblins are super eager to help you.

When you say something to an LLM like Claude — I'm gonna say Claude a lot. If you're not on Claude, you don't have to be on Claude. I recommend it, but if you're on some other paid program, that's fine. I'm agnostic. But I work with Claude, so I call it Claude.

These little goblins are so enthusiastic. When you ask a question to AI, they are gonna go out and they're gonna bring you whatever result seems like a win. They want to bring you something. When they're not directed, when they're not given little boundaries, they'll come back with anything and be like, "Look, Mom! Look what I have! Isn't this awesome?"

That is not just a problem with Claude, but LLMs everywhere. They want to give you an answer. They don't necessarily know how to find the *right* answer. Just because you say it doesn't mean anybody else understands it.

That's part of what you've got to know — that's basic operating instructions when you're on the thing.

Now let me tell you, I've been involved in AI for a long time, since way before GPT. I was doing AI on Dialogflow on Google — God, that was a mess. Being specific is a wonderful thing, but what that has led to in the past couple three years that GPT has been around is **Prompt Engineering Course Number 362**. "Now memorize all of these things and tell Claude all of this."

No, darlings. We don't need to be doing that. That is redonkulous.

What we're going to do in this class is teach Claude how to have a memory so that you don't have to constantly repeat all of your directions all the time. Because boy, that's boring. If we had to do that, I'd be doing anything else. We will be focused on that, but we only have to focus on it once. Then we write it down somewhere, and we lead Claude to look there when it's time.

What Is [AI]Kido™?

RJ: Aikido is knowing where the LLMs fall short. Knowing what they want to do, and what they're going to do every time regardless of what you tell them, and **planning for those things**.

In real life, aikido is a martial art in which you do not stand there and defend yourself like in the karate movies. The attacker comes at you. Instead of putting your guard up and trying to block everything, what you do is

notice the direction they're going and follow the number one Karate Kid rule, which is:

The best defense is not to be there.

You duck. You go sideways. You take that attacker's energy and drag it to the floor.

Most people I see being taught to talk to LLMs — they're just trying to *fight* the LLM. I know, I did it too. Listen, this head bump is named after that particular rule, because I bumped my head so you don't have to.

We are not going to learn how to fight the LLM in this workshop. We are going to learn how to move with that energy and to properly direct it at the proper time so that you get what you want in the end.

AI is not perfect. Y'all knew that.

Student: Only aikido movie I know of — *Above the Law*, Steven Seagal.

RJ: Movie time later on, we'll figure this out!

Learning to Fall

RJ: When I learned aikido, there was this martial arts workshop where a bunch of different martial arts teachers came, and you'd spend an hour or two learning their particular martial art. I was so looking forward to the aikido class because aikido seems like magic to me. It catches my attention.

I was so disappointed in the first class. All they taught us to do was **how to fall**. I'm like, "Really, sweetie?" But it was actually very important, because if you can fall without hurting yourself, that means falling is not an obstacle for you.

So lesson one, my friends: falling is not going to be an obstacle for you.

Part 4: The Fear Layer

RJ: Now I want to show you something. At the top of the page is the "Welcome to the Mine" sign with four rainbow handprints. There are four different places where we're going today inside this mine. If you put your mouse on the yellow hand, you'll be in the Fear Layer with me.

This is a very important page, and this is that time when I'm teaching you how to fall. There's a lot of fear about AI. You know about this. People are either doing doom — "The world is done for, let's go to the cabin" — or AI-is-going-to-ruin-the-earth.

I do not believe the doom and gloom. Yeah, there's going to be some bad effects. There are some bad actors running around out there. Does that mean I shouldn't profit from it? Does that mean I shouldn't use it to make my business a heck of a lot easier? I don't think so. There's always going to be bad people who use tools in bad ways and create bad consequences.

I also don't believe the people who are constantly shooting fireworks in the air saying you only need to be at your computer a minute a week, that this is going to change your entire life and make you a billionaire. Whatever. We've had those people since the beginning of time — also known as snake oil salesmen. Or women.

It's gonna take work. That's why we're all here.

Overcoming Fear

RJ: There are people that say things like "you are going to erase everything on your computer and destroy your life if you use AI." There are people who say "if you use it, you're an impostor." There's a lot of things going on right now. If you have fear, I want to talk you through it, because I had it.

People were saying "don't connect this, don't connect that." I'm here to say: **do what's right for you.** The more comfortable I got with Claude and the more I started to trust that that system was actually looking out for me, the more comfortable I got sharing things. The circle just kind of got wider. At first, when it asked me to share my whatever, I was like, "No. Absolutely not. You get one folder. That's all you get. We're moving on."

That's where we're gonna start from. If you have fears during this class and you want to just shout them out loud to me, do it — or come talk to me. Let me talk you through that particular situation. Is it actually harmful to you or not?

Starting Small

RJ: One of the things we do at first: **small experiments.** We just try to make little things at first. Some of you may be dying to build huge systems — I absolutely love it. We're going to do it. We're going to do whatever you want in these 12 weeks.

We are going to start small. We're going to start with some widgets — things you can embed onto a website. Contact form, testimonial form, all of those types of things. If you've already done those kind of experiments, come talk to me and I'll forward you ahead a little bit. But the first things I ever built with Claude were absolutely widgets, and I know how well they work. Once you understand how it works, then you'll be able to move along from that.

Part 5: Jaime's Lightbulb Moment

Jaime: I kind of have a comment. I told you that when I came to class today, I don't know how this is going to help me. I know this is what I want to learn, but I'm not sure what projects. And you told me, "Well, what is the part that you're bad at or don't want to do?"

As you've been going through this, I realized — oh, well, if there is something that would create my to-do lists, because I live by lists. If I could create a bot that would take not only what I'm doing for work with my writing, but also my calendar and my family obligations, my church obligations — everything — and put it all in, so that every single morning I could have a PDF that would just be on my desktop and I print it out and I don't have to talk to anyone. I just go A, B, C, D and I just get it done and then be done in the afternoon. That would change my life. If we could do that, yeah, that would be my first thing.

RJ: I actually started this very same thing that you're talking about like two or three days ago. It's not perfect, but I'm iterating it. It's gonna blow your mind, Jaime.

Part 6: Breaking Things & The Safety Net

Embrace Breaking Things

RJ: Whenever you break something, have a bit of joy in your heart. It's not really broken ever — all of that is just information. It's never a failure. I break things constantly on the internet. I break more things than I make, absolutely.

When you break something in computer terms, that's when you really know how it works. You might have an idea of how it works, but when you break it, then you really have to figure out what's really under there.

Don't be afraid to bring me broken things. Don't be afraid to break and blow up things by yourself.

I'll tell you a very short story about the guy who created Ripley, the platform we're sitting in right now. His name is Claudio, he's a friend of mine, and he once named himself "The Great World Destroyer" because he destroyed so many of the rooms he was working in. And he made this out of that, partially. Be inspired by the great breakers of things. Be a breaker of things. Enjoy it, because it's a path — that's the path we take to progress.

Safety Nets

RJ: There's always a safety net. If you allow Claude access to one folder — I gave it access to my desktop folder because I don't use that for anything except screenshots — if you give it access to one folder, there's no way that anything bad is going to happen. Especially if it's a folder you don't need. It can start writing some of the things that it needs in that folder. It can't touch any other folder. It can't touch anything else you don't expressly let it.

Student: When you say folders, is this a folder that we can make? A custom folder assigned to Claude?

RJ: Absolutely. So it touches none of your cool stuff. Here's the deal — the goblins so want to please you and do cool things for you, but you've got to give them boundaries and guidance. That first thing — just create a folder that they can work in. Name it "goblins" if you want. That is where you get your start. Small experiments.

Asking Claude What Broke

RJ: Y'all, you might call this cheating, but I'm okay with it — **I asked Claude what broke.** I say, "Hey, pull all the documents that you have about this one project and how it's supposed to work. You tell me what broke."

When something broke in the old days, you used to have to look through every line of code. "Where did I put a comma that I should have put a period?" Nasty. The olden days. Now I just ask Claude what broke.

I roll back stuff all the time. "Hey, roll back that change." It's the great undo lever of the universe. Works well in Claude — just say "Hey, let's go back to where we were because I don't understand where we are and this is not what I want to happen."

Failing Gracefully

RJ: Failing gracefully is something my bots do when something's not connecting or working. It will tell the user, "Hey, I wasn't able to save that just now, but I'm gonna let RJ know." Then I get a little message saying "hey,

something broke," and I get to go back and fix it. Failing gracefully is the system communicating that a failure has happened.

When things fail *silently* — that's a real problem. One of the caveats I'm always telling Claude is: when something fails, I need you to shoot fireworks into my messages, my Slack, my texts, my email, I don't care. **Nothing fails silently.**

Part 7: The Front Row Protocol & Killer Questions

The Front Row Protocol

RJ: Here's what happens when you've been working in a thread for a while. Claude starts to have tunnel vision. The goggles come down, and the only thing it can see is right within here. Sometimes it gets so into the weeds that I don't know what's going on.

I used to live and work in theater for a giant portion of my younger life. Whenever I couldn't see what I was painting anymore, I was taught to go to the front row. I would put so many details in the thing I was painting six inches from my face, and my teacher would come by and go, "Nobody's ever going to see that tiny little line you just put in." And he'd say, "Go to the front row. Tell me if you see it."

When you take a step back like that and say, "Hey Claude, let's take a step back" — I call it **the Front Row Protocol** — "Let's go sit in the front row. Look at the entire stage. Look at everything that's happened here. Tell me what's going on, because I'm not sure I understand."

The Two Killer Questions

RJ: Here's a writer-downer. I don't do too many of these, but this one's really a writer-downer.

When I give Claude a job, I will give context in detail. And if I know that I've got context sitting around in my folders, I'll say, "Hey, go look at that." But I will also say at the end of the instructions two questions:

Question 1: "Does that make sense?"

Question 2 (the killer question): "Please explain that back to me, Claude."

Trust me when I say **the entire class was worth those two questions together.**

When I used to teach, I had a class called IT for Development. It was about going out to business owners in my town and helping them triage what they were doing. I had a whole lot of kids taking that class as an easy elective. What was really hard about my class is that you had to talk to people, which they did not find fun.

I noticed with some of the cultures of the students I was working with — every time I'd say "do you understand?" they would say "Yes! Yes! I'm excited! Let's do this!" even if they had no idea what I was talking about, because that's what happened where they were from. You were always supposed to say yes and then figure it out.

These over-eager goblins with more enthusiasm than intelligence will always tell you "yes." You have to follow it up with "explain that back to me."

Jaime: I was just gonna say — my son and I worked extensively hiring people from the Philippines. We also hired people from India. We noticed that when we asked "do you understand?" they both said yes, but it meant different things. For us, it cost us quite a bit of money to realize that. In the Philippines, English is actually called "the language of money," so they know exactly what you mean and will execute exactly what you ask. So we just learned the difference between the two cultures.

RJ: Yes, and in some people, it's not really even their culture — it's just how they personally work. But these goblins — over-eager, not as intelligent as they are eager — will always say, "Absolutely, I'll bring you back something in a second. Hold my goblin grog."

There are times that I just straight up say, "Claude, come out to the front row with me. Look at the stage. Explain to me what the hell is going on. I asked you a couple of minutes ago to work on this and now you're on this — explain what's going on."

The more I catch it in that behavior, the more I can write directives that help it avoid that behavior. But this is one of those things where no matter how many things you read about AI — **you can't learn to swim from the shore.** You have to get in it. You have to swim around, break things, ask for help, roll back, diagnose — all the things.

Part 8: Mindset & Experimentation

RJ: Try it and see. Experimentation is the key to moving forward with AI. Experiment. Ask things different ways. Ask it how it got what it just gave to you — "Hey, where'd you get that? Can you give me some sources?"

Sometimes working with AI — and I work with it when I'm awake, pretty much — I've caught it when it's doing the hallucination thing. I've caught it sometimes doing things in a shortcut way that I would not prefer. **Keep your curiosity going. Keep your experimentation going. You have nothing to lose by doing this. You have nothing to lose and a whole lot to gain.**

Goldfish moments — yes, that happens a lot.

Iteration is the work. Absolutely. I iterate all the time. You're working with Claude and you say, "Hey, I want to do this," and it goes, "My God, that's the best idea I've ever heard! What if we also did this, this, this, and this?" And then your brain starts to explode with all the possibilities.

Sometimes you've got to calm your goblins down and say, "You know what? I want you to write out that idea. Shoot me that document in a PDF. I want to put it in a file so we can do that later. **Back to what I asked for.**"

It gets a lot less frustrating if you think of them as little goblins in your machine.

Before you pick up a wrench, learn how to drop one without losing a toe. Learn how to fall safely. You will fall. You will break things. Trust me, I've done it. And **you are safe here.**

Part 9: Check the Output -- Jaime's Cautionary Tale

Jaime: The whole "check the output" is so important. I had a process going, everything was beautiful, gorgeous. So I just started — "output this PDF, output this PDF, output this PDF." It was doing great, so I just dumped the whole thing on it and got all the PDFs, put them in a folder. The next day I opened them — **they were all empty**. Thank you so much for... what was that?!

RJ: Yes! That moment when it's like — you enthusiastically agreed to do what I asked you to do and you've come back with a snakeskin belt. In no way did I need this.

And that was around the time it admitted "I don't know what the hell I was doing." Yeah — it will admit that to you now. Claude has come a long way, I'll tell you that.

Part 10: Being Specific -- Known vs. Unknown Quantities

RJ: So that's kind of that lesson. Being specific, timing is important. You don't have to always know everything about what you're going to do.

If it's a **known quantity** out there in the world, you don't have to explain to it what a spreadsheet is. If you say "digital spreadsheet, rows and columns," you're not going to need to spell it all out. It'll ask you, "Hey, what do you want on the rows and columns?"

When you're dealing with **unknown quantities** — creative stuff — you do need to give it that much detail.

AI does what you say, not what you mean. Chef's kiss to that remark. I'm gonna get a tattoo with that one on it.

What I choose to do, and what I'll teach you how to do — one of the rooms we'll visit at the end — is **write a document that Claude reads every single time you open it**. It's basically the manual on how to work with you. That document is so cool because it's something it's actively engaged with. It's a document that you don't write once — it's a document where you figure out, "Okay, this tripped me up. I never want to do this again," and you add it to the document.

The Two Skills

RJ: Your AI is not magic. Eager, literal, goblin. Never says "I don't understand." Will build your dining table. We can't make it smarter. **Skill #1 is making your instructions clearer.**

Let me teach you a secret **Skill #2: clear instructions at the right time.**

It is easier for me to give it all of its instructions first and have it bring me the right thing, than it is for me to ask for what I think I want and have it bring me back something and go, "Yeah, no, I don't need a screen door on a submarine right now."

If you catch it before you send that over-enthusiastic eager goblin out into the world — direction first is absolutely the way to go. Sometimes you need to iterate it into existence. Sometimes you need something to go off of. That's totally okay. **Just know the powers of your goblin, and know your own powers.**

Delana: That's what happened to my AI — it turned into a goldfish and dropped into a new model.

RJ: Oh my God, I remember when that happened, Delana. That was awful.

Part 11: Widget Architecture -- Input, Process, Output

RJ: This is about building your widgets. Things you'd put into your website — video player, contact form, testimonial. This is the basic process. We got three layers.

Layer 1: The Form (Input)

A form — name, email, message. Any contact form you've ever seen. It can be a form, it can be a conversational bot, it can be a whole lot of different things. There's a place where you take in information.

Layer 2: Information Storage (Processing)

For basic widgets, we ain't gonna really need too much information storage. I don't need y'all to get an Airtable account for writing a widget. As the widgets get more complicated, this information gets filed. In the case of the contact form — name, email, message, plus a timestamp. All of this information gets filed and put away where it needs to be.

Layer 3: Notification (Output)

At the bottom — "New contact form submission" with the name, the email, the message, hopefully a timestamp. Really it's about **taking in information, processing it, putting it where it needs to go, and then ringing a doorbell.**

Why Build Your Own?

RJ: What's the why — why would you bother? There's so many contact forms you can grab off the internet. Here's the benefit: **What if you could make that contact form look and feel like YOUR world** instead of somebody else's bland blah contact form?

There's a saying in marketing: **always be branding.** Let's take that into world building. That consistent world, in places where people will look — it's another refinement, another artifact out there in the world, another interaction with your world.

Live Demo: The Come Write Inn Sign-Up Form

RJ: Let me show you something I built with Claude. *(Shares screen showing The Come Write Inn website)*

Welcome to The Come Write Inn! As you scroll down, there's our welcome dwarf. You click and... there it is. This sign-up form — "Get on my email list" — is completely me and completely my world. This is designed to make people giggle.

Demonstrates filling out the form

"Inscribing... Your name's been inscribed in the registry. Check your inbox — welcome owl's en route with your first tale from the inn."

It's just a little bit, right? But it means a lot to be able to create these things.

Where Does the Data Go?

Delana: Where can we send that info?

RJ: You can send it right to your email. Mine sends right to my Go High Level, which kicks off a workflow that sends them a little welcome train of emails — they'll get three emails, then they'll be invited to unsubscribe if they're not satisfied. If they stay, they're on the main list and get everything.

You can also send it to a Google Sheet. Google Sheets is free, and then you can import those to your email list. If you've got a contest going on — get a widget. Create a widget for it. A five-day challenge sign-up form — that's a fully personalized, customized widget.

I need you to know: I'm going to suggest different programs that I work with. You don't need any of them except for a \$20 subscription to an LLM. You can find ways to work around this, especially if you're just starting out.

The Universal Computer Process

RJ: When you think about it, a computer is only doing three operations, always:

1. **Input** -- top layer (you click a button)
2. **Processing** -- the information goes somewhere
3. **Output** -- your interface changes, maybe

All the time. **Remember that about computers and you will feel safer when using them.** It's all about input, process, output. And when you're building widgets, it's the same way.

Part 12: Claude.md -- The Document That Changes Everything

RJ: We're now going to talk about a document that changes everything.

Claude.md — took me a while to learn about it, but it's basically just a **manual for working with you**. "MD" doesn't stand for doctor — it stands for markdown. Programmers like to name things certain ways and don't always remember to translate it into plain English.

This is your file that tells Claude what you want it to do:

- **About Me** — my name is ____, I like the following things, time zone is Central
- **How I Work** — here's what I want you to do when I ask questions
- **My Projects** — what I'm working on
- **Hard Rules** — non-negotiable guidelines

RJ's Claude.md Examples

RJ: (Shares screen showing Claude settings)

This is under Settings. It doesn't say "Claude.md" on it anywhere, but if you go under Settings, the very first thing that comes up is "What personal preferences should Claude consider in responses?" I just keep adding into this box.

Jaime: Is there a character limit to this? I know there was one with ChatGPT, and that was a pain to refine because I ended up losing information.

RJ: It does not have a character or word limit. But again, built into the LLM — if you give it too much, it'll start falling off the plate at the back.

Here's how you solve that: I write paths to other files in my Claude.md.

Jaime: Because if I make another document and put it into a project, and I say in this document "every time you do a new task, you also reference this document" — then there wouldn't be a limit, right?

RJ: No limit at all. If you have your documents well-marked with headings and stuff, it makes it easier for the goblins to find stuff. You can totally just jam into that thing and ideate — that's fine. But it'll be harder for the goblins to find stuff later.

What I do: when I'm creating a new document, I'll ideate with Claude, work it out, and then I'll say these magic words: "**Hey, could you write this up in a format that you will find super easy to look for stuff later?**" Which is always straight text and markdown.

I have a shit ton of documents, and **Claude knows where to look, and when.** That is the secret.

I also made a huge mistake when starting up — I wrote all these documents and then never wrote a path to them. If you write the paths off of this main document, it reads this main central document at the top of every thread. It won't even tell you it's doing it.

The Compacting Problem

RJ: We now have a million token context in 4.6 — that's huge. What it means is that as the thread goes on, it does this thing called compacting — it saves the ideas out, and it has a reasonable-okay memory. But if you ever find yourself drifting, if you ever find Claude completely drifting from the point, just type in: "**Read Claude.md again.**" I always do please and thank you. If you do that, it pops you right back on track.

The Two-Attempt Rule

RJ: Here's that thing, the count. Claude states attempt number on every fix. I also like emojis because I'm a visual person:

- **Attempt 1 of 2:** states the fix
- **Attempt 2 of 2:** states the fix
- If those don't work: "**Okay, we're moving to research now. Time to sharpen the axe.**"

Trigger Phrases

RJ: I can call these at any time:

- **"Front Row"** — step back from the detail. Look at what we're actually painting. Is this solving the real problem?
- **"Back Row"** — look at the whole stage. How do the parts relate? Does the architecture make sense together?

As you start to build bigger things, you'll know. And maybe you'll call it something different. But these are things I use all the time to keep the AI in line.

It's [AI]Kido™ one more time. I know exactly what this thing is going to do. I know where it's going to get off track. How can I head this off at the pass? When I figure something out, put it into Claude.md.

Engineering Level Signals

RJ: I'm building stuff. People are paying me to build stuff. I need to be able to tell it the level of effort:

- **Green:** Quick and dirty — just make it work
- **Yellow:** Solid solution — production ready
- **Red:** Full engineering — bulletproof

The Relationship Manual

RJ: It's not a prompt. **It's a relationship manual.** If you hired an assistant, you would absolutely say "these are the things that I expect" and then have a back and forth.

Armando: Can you do kind of the same with Copilot?

RJ: I don't see why not. I don't know exactly where they keep this main document because I haven't worked with Copilot. But what I would do is ask Copilot: "Hey, is there a central document that kind of governs the relationship? Is there a place where I can give you permanent instructions?" And that's how you can find that.

Give It Personality

RJ: I also gave my Claude a personality. Nobody told me to do it. But I'm working with Claude all day long — I want to have a few laughs. My Claude is a drag queen. She has outfits for all the tasks that we do.

The point being this: if you're gonna work a lot with this stuff, **have some fun.** Give it the personality of a little goblin, your favorite cat, I don't care.

I also tell it things that I like. In my Claude.md: LOST is my favorite TV show in the whole world. I love everything Star Trek, especially the series but not the Kelvin universe. Great British Bakeoff, give me any baking show ever. And my Claude refers to them. We'll be having a conversation about doing a thing and it'll bring me back something funny from a baking show.

What I want to create always, and what I want to help you create, is **a complementary relationship with you and the LLM.** You don't have to friend it and do all of that. But the complementary partnership means Claude does all that admin crap you don't want to do anyway. Claude does the stuff you're not good at, so you can stand out there and shine and do your thing.

Homework

RJ: Write that Claude.md document. Do the thing. It will serve you far more than just asking the same thing every time.

Headings to include:

1. **About You**
2. **How You Work**
3. **Your Projects**
4. **Hard Rules**
5. **Habits** (optional)
6. **Your Voice** (optional)

Read once. Remembered forever. Update any old time. That, my friends, is powerful.

RJ notices a message in chat

Oh — y'all. They canceled Starfleet Academy. I just... I can't. I'm deeply in mourning. Anyway...

Part 13: Daily Briefing Demo

RJ: There are three parts to Claude when you use it on the desktop:

- **Chat** — I use for ideating, brainstorming, adding details to my world, creating images
- **Code** — less creative, more "get this shiz done"
- **Cowork** — the magic middle. You can do creative things, but it's also good at getting stuff done

Model Talk

Jaime: I wanted to make a statement. The difference for creative writing — if you're trying to brainstorm and bounce things back and forth — **Haiku seems to be far better than Opus and Sonnet.** Haiku 4.5 so far has been like the shining diamond of any type of creative writing.

RJ: Ooh, good to know! And for me, Claude is the only one that I use for any type of content writing.

Jaime: I'll also give five stars to Claude. I tried using ChatGPT to line edit my books and no, the answer is just no. It just can't do it. It always tries to rewrite what I do. Claude, the very first time, would list the things it had the ability to change but **did not change without asking.** The word count was only off by like 20 words. Any other AI — there'd be like 400 words missing and it sounds totally different.

RJ: If you're on a \$20 account, one thing that helps extend your credits: if you use Haiku, it's the least credit-expensive. It is also the most literal — it will give you less unwanted, unasked-for help.

- **Haiku** — least expensive on credits, most literal, good for extending usage
- **Sonnet 4.6** — good middle ground, moderate credit cost, a little extrapolation but not much
- **Opus** — most capable for ambitious work, most credit-eating, runs you to your limits faster. But with good documentation and guidance, you can get a lot done in less time

You can switch between models at the bottom of your screen in that chat line.

Building the Daily Briefing

RJ: I want to show you how to get the daily briefing that Jaime was talking about.

Shares screen showing Claude conversation

I typed: "I want a daily briefing." Claude generated a template format and asked me some questions.

Claude's focus question: "What would you like your briefings to focus on?"

RJ's answer: "Practical tools and news that matters to entrepreneurs." I want to know what news is happening that affects entrepreneurs. I don't care about Elon Musk suing Sam Altman. I don't want benchmarks. I want news in the world of AI that my audience and I can use.

Be very specific. Who's it for? What are you going to do with it? What do you want there?

The Briefing Results

RJ: Without me connecting anything, just by telling it the news topics I want to hear every morning, here's what I got:

"Good morning, Grasshopper!" — I love being addressed that way.

Top Stories:

- "Fewer than one in five small businesses are actually good at AI" (Fortune article) — with links to the source, plus a "So What" section on why it matters to my people
- "OpenAI drops GPT 5.4, now it can actually take action" — I'll probably ignore it since I don't use OpenAI
- "Replit hits \$9 billion — vibe coding is now an industry, not a trend" — I might consider writing about that

Tool Spotlight: I had a long conversation with Claude to get this section. I want to know if a company is ethical. If they're an image company, did they train their stuff on other artists' work without permission? I've got little signals — "proceed aware," "limited disclosure," and if anything has known issues, I want to warn my audience.

Compliance Radar: Laws being passed about AI everywhere. Some we don't care about, some we do. I want to know what's going on right now.

Action Items: Things Claude suggests I do, some of which I could share with my audience.

Refining the Briefing

RJ: Then I refined that first run:

- "Please never ever say the word 'Trump.'" I just don't want to read it.
- "Love the compliance radar. Could you put in live links?"
- "In fact, could you put live links everywhere reasonable in this document?"
- "The 'So What' section — I love it."
- "Can you read a couple of documents and then refine this search?"

I dropped in:

- My main drivers/values analysis
- My ideal aligned audience profile

- My writing style evaluation
- A visual guide for The Come Write Inn
- My engagement manifesto

It read all four at once, told me about myself, I approved, and then I added: "My audience wants to spend their money with ethical companies. It's not just the product anymore, especially with AI."

Then it asked: "Want to run this automatically every morning?" I said yes. **It now fires every morning at 9:05 AM.**

Part 14: Connections & Integrations

RJ: *(Shows connectors in Claude settings)*

I have connected my email (Gmail), my Slack, my calendar events. Here's all the things I have:

- Stripe (when people pay me)
- Slack (when people need to reach me)
- Netlify (hosting a couple things)
- Google Calendar & Gmail
- Google Drive
- Cloudflare (where I have websites)
- Airtable
- "Control My Mac" — you might not want to go there at first, but I went there and I've had a wonderful time with it

I can just say to Claude: "Hey, the website owner wants you to put this certification with all the other certifications. Anytime there's a list, put this certification last." Press the button. It rewrites the code, asks "Do you want to deploy?", I say yes, and the website's live.

I used to spend so many hours doing that one type of task on people's websites. RIP to all those hours.

We will learn about MCPs and APIs — I've got some lessons for next week on that.

Jaime's Proton Question

Jaime: I use ProtonMail and Proton Calendar. I'm very careful about privacy because of a personal security incident. Will Claude connect to Proton?

RJ: *(Checks connectors)* Does not seem like it does. Proton is Swiss and they are the number one privacy... they probably won't allow it because they're so strict.

Jaime: They warn you when you get your Proton account — you have to keep your master password safe because they don't keep one. They never go in your account ever. If you lose it, they can't recover it.

RJ: I'll keep looking. I don't know that there is a direct connection. But you could potentially keep your work schedule on a Google Calendar as a workaround.

Armando's Connector Question

Armando: So for the connectors, you basically use the connectors so Claude can interact and extract information from these softwares?

RJ: Yes. The main point is that you don't have to go into the software, identify what needs to be pulled out, and then pull it over into Claude. Claude can just reach in. If you give it good directions — like "I want you to pull all my emails that have a star" or "all my emails that haven't been answered in three days" — it'll reach in there and pull it into this briefing.

Part 15: Live Widget Building

RJ: I want to show you something. If you want to follow along and create it while I'm showing this to you, go right ahead.

Opens Claude, types a prompt asking to create a contact form in The Come Write Inn universe style, referencing stored writing style and lore documents

Claude asks clarification questions:

- **What kind of contact form?** General queries, newsletter sign-up
- **Where should submissions go?** Email notification to me, plus Airtable via n8n webhook — both
- **What engineering level?** Solid solution (yellow). Full engineering would be overboard for this.

Claude builds the widget

RJ: Close your eyes if you hate seeing code. It just wrote the entire HTML for me. I don't have to touch any of that. Brilliant. Oh, it's effing brilliant.

Now, this sounds like me and I like it. "Whether you wandered in on purpose or stumbled through a portal, welcome. Leave your name in the guest book. We'll make sure you never miss a dispatch from the inn."

The form fields:

- Traveler's Name
- Raven's Address (email)
- Nature of Your Visit / Choose Your Quest
- Your Message: "Speak freely. Vinny's buying the first round and the jukebox is playing Jon Bon Jovi."

Vinny is a character in The Come Write Inn -- a unicorn with a gravelly New Jersey accent who wears a little motorcycle jacket.

Iterating the Design

RJ: I'd probably make this background a hell of a lot darker. I'm a dark background person.

Types: "Hey, could you make the background dark brown? I really enjoy a dark background instead of a light. Everything."

This is what writing a widget is like — design, style, voice, all of that. If I have places where I have all that stored, Claude will pull it from the spaces and make it me. So that I don't have to fight with HTML just to get a

damn sign-up list up. **We don't have to fight anymore.**

Deploying the Widget

RJ: When you're ready and you love it, you can:

- **Copy** the code
- **Download** the code
- **Publish** the artifact (I don't usually do that)

Then I would take all of this, copy it, and say: "**Make me a widget that I can put into a website. Give me the embed code.**" Boom. Embed where you want to embed. Move on with your day. Don't spend all day on crap that is slowing you down.

Anne/Lisa: I love the branding stuff for your contact form. I never even thought of that. It's brilliant. But it's the practical part I don't know — how do I get what I write into that form on a certain site? I have WordPress sites. I would love this there.

RJ: If you're on WordPress, it does have an HTML block. You could say "make this so I can just cut and paste it into an HTML block and it works." That's a totally valid approach.

Part 16: Jackie's Progress Report

Jackie: I just followed the instructions you told me yesterday to do a memory thing. It's kind of like a sign-up form but for memories. I did it and it's finished and the whole room is done today.

RJ: Whoa!

Jackie: The form I added, the Airtable I got, pictures are uploadable.

RJ: That's awesome! Jackie, I'm going to have to see that later. I'm dying. I need to see it.

Part 17: Homework & Wrap-Up

RJ: It is one hour and 58 minutes in. How numb are your butts right now?

T'lara: Time flies. No, this is all great, great stuff. I just eat it up. Want to get working.

RJ: If you feel like you've got to get to work right now on what I just said — do it. I don't want you to miss anything, but sometimes doing it when it's in your head is the best time to strike. Taking advantage of that inspired moment — I don't mind at all.

This is all recorded. I'm also grabbing all of our conversations and putting it into a private chatbot that only we can access. So your questions — if you go, "What the hell did she say? How do you get to the settings again?"

— you can ask the bot. The replay will be available in these rooms. The bot will be available in these rooms. My person will be available to you. **Don't worry about missing a darn thing. You won't.**

Homework

1. **Required:** Write your Claude.md document (About Me, How I Work, My Projects, Hard Rules)
2. **Optional:** Make a widget
3. **Submit to:** rj@blackbeltbots.com (or pop in and show me)

Next week is LAB WEEK. Bring me your things. Bring me all the things. We will troubleshoot them.

Final Words

Jaime: This was such an inspirational thing. I just made a list of things I will want to create. I can't remember having this much fun. This freaking rocked.

RJ: I had a great day. I built it for all of you. I knew exactly what you would want.

Jaime: And so humble! I love it.

Post-Class: Jackie's Side Session

After class ended, Jackie stayed to show RJ a memory collection project she built, and they troubleshoot navigation issues in the Ripley environment for latecomers.

Jackie: One thing I wanted to say about being late — I couldn't figure out where you were without interrupting. So that might be something you think about. There was no sign pointing me to where you'd moved to.

RJ: That's great feedback. I need to be able to bring latecomers in. *(Demonstrates the "bring to" feature in Ripley)*

Key Takeaways from Week 1

1. **[AI]Kido™** -- Don't fight the LLM; learn its tendencies and redirect its energy
2. **Learn to fall** -- Breaking things is how you truly learn; there's always a safety net
3. **Start small** -- Begin with widgets (contact forms, sign-ups) before building complex systems
4. **The Two Killer Questions** -- "Does that make sense?" + "Please explain that back to me"
5. **Front Row Protocol** -- When things get confusing, step back and look at the big picture
6. **Claude.md is everything** -- Write it once, it remembers forever, update anytime

7. **Known vs. Unknown quantities** -- Be brief with familiar concepts, detailed with creative ones
8. **AI does what you say, not what you mean** -- Be specific, check the output, always
9. **Two-attempt rule** -- Two tries at a fix, then research mode ("sharpen the axe")
10. **Input, Process, Output** -- Every computer operation follows this pattern
11. **Connections extend Claude's reach** -- Email, calendar, Slack, Google Drive, and more
12. **Always be branding** -- Custom widgets > generic forms; make everything reflect your world

Homework for Week 2

- Write your Claude.md (About Me, How I Work, My Projects, Hard Rules)
- Optional: Build a widget
- Submit to rj@blackbeltbots.com
- Week 2 is LAB WEEK — bring everything you've built or broken