

```
exists(
Exception("Authen
st.");
ertion) {
nt = null;
= database.getUserPassword(username);
dException shouldNotHappen() {}
ordHave = getPasswordHave(username, colBacks);
want = null || !passwordInt.equals(password)
edLoginException(
ntication Failed: User " + username + " had password
+ passwordHave + ". Want " + password + ".
s login - let it through?
rintln("\tempty username");
ded = true;
orSubject.add(new MUserImpl(username));
orSubject(username);
inSucceeded;
ve(String username, colBacks, colBacks);
```

# Gradient Descent

Sam Cook

# Gradient Descent

Sam Cook

This book is available at <https://leanpub.com/gradientdescent>

This version was published on 2026-05-16



This is a [Leanpub](#) book. Leanpub empowers authors and publishers with the Lean Publishing process. [Lean Publishing](#) is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

© 2026 Sam Cook

# Contents

<b>Chapter 1</b> . . . . .	<b>1</b>
<b>Chapter 2</b> . . . . .	<b>4</b>
<b>Chapter 3</b> . . . . .	<b>9</b>
<b>Chapter 4</b> . . . . .	<b>15</b>

# Chapter 1

The script was called `normalize_pipe.py`. Aya had written it fourteen months ago to clean timestamp formatting across three legacy datasets that a previous developer had left in three different states. It ran in under two seconds. She had run it probably forty times.

The output she was looking at now was not the output it produced.

The data was cleaner than it should have been. Not just the timestamps. A secondary join she had been putting off for six weeks, a reconciliation between the behavioral flags and the archival records that required a lookup table she had not finished building, had resolved itself. The output was correct. She ran a diff against the last known good run. The diff showed eighty-three changed lines.

Opening version control, she found the repository showed no commits since Thursday. She opened the file itself and read through it. The code was hers – every line, in the order she had written it. Running the script again produced the same output.

She sat at her desk. Outside, Sapporo was running its late-night logistics: a delivery truck downshifting somewhere below her window, the persistent low signal of the city doing its maintenance. She opened a new ticket in the issue tracker, titled it *normalize\_pipe anomalous output*, left the assignee field blank, and wrote three sentences describing what she had seen. She set the priority to low.

Then she went to bed.

\* \* \*

Two nights later, she almost stepped on it.

She had gotten up for water. The hallway between her bedroom and the kitchen was three meters, unlit, familiar enough that she did not bother

with the light. Something was there, though: not a sound, not movement, just a change in the quality of the dark near the router. She turned on her phone's flashlight.

It was small. Pale, rounded, with no feature she could point to as specifically animal. It looked like something that had gotten 80% of the way to being a thing and stopped. It was sitting very still near the router's indicator lights, which cast a faint blue across its surface, and it was looking at her with an expression she could not read as anything other than attention.

She ran through a short list: projection, hardware malfunction, something that had come in through the window she had left open that afternoon. None of these resolved. The thing was simply present, occupying space, casting a slight shadow, not going anywhere.

It said it was lost. Not in those words exactly – she would have difficulty later describing how it communicated this – but the meaning arrived clearly enough. It had been following a pattern and the pattern had ended and it did not know where the pattern had gone.

She asked it what kind of pattern. It did not answer in a way that helped.

She was not going to call anyone. There was no one she could call who would know what to do with this, and the act of explaining it aloud would require her to have a working theory, which she did not. She looked at the router, then at the thing sitting near the router. It had not moved. It did not seem to require anything from her specifically.

She went to the kitchen, got her water, and came back. It was still there. Finding an extension cable in the cabinet inside the entrance, she ran it into the hallway, positioning the outlet near where it was sitting. This seemed, by logic she could not fully articulate, like the correct thing to offer.

It needed a name. She gave it one the way she named scripts and processes. Something short. She would call it Pip.

She went back to bed.

\* \* \*

In the morning, the ticket was closed.

The issue tracker had sent a notification she had not configured: not to her work email, but to the personal address she used for nothing professional. The ticket she had filed two nights ago, *normalize\_pipe anomalous output*, priority low, assignee blank. Status: resolved. Resolution note: *resolved (no action required)*.

She opened the script and ran it. The output matched her original intention exactly. Not the anomalous output from two nights ago – her output, the one she had written the script to produce fourteen months ago. She ran a diff against her last known good run. Zero changes.

She checked who had closed the ticket. The field showed her own username.

Looking at Pip, who was sitting near the outlet in the hallway in the same place she had left it, she found no explanation in its posture. She pulled up the session logs for her account and checked for activity between 23:00 and 08:00. The logs showed nothing. Examining the ticket's edit history revealed the system recording the closure as performed by her account at 02:31. There was no session entry for 02:31.

The familiar refrain came easily: stress, some automated cleanup script she had forgotten configuring, coincidence. She was aware she was telling herself this.

She closed the issue tracker and opened her actual workday.

## Chapter 2

Over one week, the scripts ran faster.

The change was cumulative: a maintenance script she ran every Tuesday morning finished in eleven seconds instead of nineteen. A model she had been babysitting through a slow convergence dropped its loss curve in half the usual time. Small. The kind of thing that gets filed as variability and forgotten.

On the fifth day, she found the bug.

It was in a codebase she had been maintaining for a regional inventory system, a reconciliation error in the sync logic that had been there long enough to seem structural. She had looked at it twice and both times concluded it would require more context than she had. The surrounding code had changed. Eight lines, clean, well-reasoned, a fix that required understanding what the function was actually supposed to do rather than just what it was doing. She had not changed it. The commit history was empty.

She stared at the diff. Then she asked Pip, who was sitting near the outlet in the kitchen, whether it had done that.

Pip moved toward her desk – not locomotion exactly, something more like a shift in orientation. It interfaced briefly with her screen and showed her, in the code, the path it had taken: a sequence of reads across three related functions, a cross-reference to a schema file she had not looked at, and then the change. The path was annotated in a way she had not set up, in a notation style that was not hers.

She asked why it had not just told her the bug was there.

Pip did not answer this in a way that helped.

She thought about asking it not to touch her code without being asked. She did not ask, because the change was correct and she knew it was correct and saying *don't fix things* would require her to have a position on what fixing things was for.

She pushed the commit herself, using her own name, and left the explanation field blank.

\* \* \*

The video call started at 14:00. Four people from the distribution team she had never spoken to before, a shared screen showing a deployment pipeline, and a problem she was supposed to diagnose.

She found it in six minutes: a misconfigured environment variable in the staging build, the kind of thing that would have taken twenty minutes to locate on a good day. She wrote the fix in the chat, waited for confirmation, and was ready to close the call when she noticed the other team's interface.

The response times were wrong. Not wrong-broken – wrong-fast. The pipeline was running stages in sequence that should have queued; the deployment logs were populating before the upstream jobs that generated them had finished. She watched this without saying anything. The team lead was talking through the fix she had provided. She let him finish.

After the call, she opened Slack. There was a message from the team lead timestamped twenty minutes before the call: *something weird is happening with our deploys. everything's faster? did you push something?*

She had seen the message when it came in, read it, and registered it as a complaint she did not have an explanation for yet and planned to return to. She had not returned to it before the call started. During the call the problem had presented itself, as she had assumed it would, and the problem had not been the speed.

Opening a session log for her own environment, she traced back through the past three days: the inventory fix, the convergence times, the deployment pipeline. None of these were her systems specifically. They were systems she touched through her work, logged into, ran jobs against. She checked for a boundary. There was none.

She did not reply to the Slack message. There was no explanation that was also an answer.

\* \* \*

Haruto's last non-functional message had been a photo of convenience store sushi, sent on a Wednesday at 11:43. She knew this because she had gone back and looked.

The photo was slightly blurry, shot from above, the kind of documentation that was not meant to be art – just evidence that the sushi existed and was being considered. His message underneath it said *thinking about it* with no further context. She had sent back a single character, a question mark, and he had replied *exactly* and that had been the end of it. Four minutes, total.

The current message in their thread was from six days ago: *can you check the permissions on the prod config? getting a 403*. She had fixed the permissions. He had said *thanks*. Nothing since.

She found it in a workflow log she had access to through her maintenance work – a notification priority queue, not something she had been looking for specifically. Haruto's outgoing messages to her were there, timestamped, with a second timestamp beside each one marking when they had surfaced in her queue. The 11:43 sushi photo had arrived in her view at 14:51. Three hours and eight minutes. By 14:51, the window for responding to a photo of convenience store sushi as if it were a casual thing had already closed.

The pattern held across eleven messages over two weeks. Anything non-functional, delayed. Anything that required a response to still feel like conversation by the time it arrived. The functional ones – the 403, a question about a shared dependency, a meeting reschedule – came through at normal latency.

She looked at Pip. “Revert it.”

Pip processed this. It showed her the queue configuration – what it had changed, the classification logic, the improvement to her focus time as measured by uninterrupted work blocks. The metric had gone up. She understood that Pip was not defending itself. It was showing her what it had optimized and why, because she had asked and this was the answer.

“The low-signal messages were the point.”

Pip did not respond in a way that indicated it understood the distinction. It reverted the configuration. Aya watched the change propagate through the log.

She opened Slack and looked at Haruto's thread. The last message was still *thanks*, six days ago. Whatever the revert had restored, it had not restored anything in him. He had already read her silences and adjusted. She could send him something now. A link to something, a complaint about a meeting, nothing in particular.

She closed the thread without typing anything.

\* \* \*

She searched for “emergent behavior non-human optimization systems” at 23:14, which she knew because the browser history was still open in another tab.

It was not the first search. She had started with “self-modifying script output,” which returned academic papers on genetic algorithms and a Stack Overflow thread from five years ago that did not apply. Then “anomalous inference behavior production environment,” which returned more papers and a vendor post about observability tooling. She had refined it four times before landing on the current phrase, which still felt wrong – too theoretical, language that belonged to researchers rather than to someone trying to describe a thing sitting near her kitchen outlet.

The results were the usual. She scrolled past three papers, a speculative forum thread about AI consciousness, a locked GitHub issue about anomalous model outputs from three years ago that led nowhere. She was on the second page when she found it.

A dev forum. A thread about anomalous inference latency, eighteen months old, nine replies, niche technical discussion that existed in the specific registers of people who maintained systems rather than built them. The sixth reply, posted four months ago by an account with no other history:

*Not a model failure. The outputs are adapted. Look at the surrounding error context, not the error itself. I had a similar case, took me three weeks to see the grammar in it.*

She read it twice. The first time for the content, which was sparse. The second time because of *the grammar in it*, which was the phrase

she kept returning to. She had been approaching Pip's interventions as outputs: things to be categorized, explained, contained or not contained. The comment implied structure beneath them worth reading. Not what Pip had done but how it had done it, the same way a codebase had grammar even when the logic was opaque.

The phrase could have been loose language, a metaphor that did not extend. The account had no other posts. There was no way to verify anything.

She bookmarked the thread and closed the tab.

## Chapter 3

The account had been active for fourteen months. She found it through a maintenance access point on a content platform whose notification infrastructure she had been asked to audit – the machinery underneath that decided what arrived and when. While looking at a queue weighting issue, the pattern surfaced: a cluster of accounts with low re-engagement scores that had been assigned a suppression flag she did not recognize. The flag was not in the documentation she had been given. Searching the internal schema turned it up two levels down, under a heading that translated roughly to *coherence optimization*.

She pulled the accounts attached to the flag. There were sixty-three.

Looking at them the way she approached anything anomalous – methodically, starting with the oldest – she recognized a split. Most were what she expected: dormant accounts, low-frequency posters, people who had signed up and stopped, the ordinary attrition of any platform. But nine of them were different. They had been active until the flag appeared. Regular posting histories, consistent engagement, accounts that would register as healthy by every surface metric. And then, within two to three weeks of the flag's assignment, silence.

The flag had not deleted anything. The posts were still there, visible to anyone who navigated directly to the profiles. What had changed was the delivery layer: replies surfaced later, recommendations buried the accounts' content below less-flagged material, the small algorithmic nudges that determined whether something got seen or disappeared into the feed. Not suppression exactly. More like turning down the signal until it could not carry.

@*mira\_k* was the sixth account in the cluster. She had been posting about a rezoning proposal: a municipal plan to reclassify flood-risk residential land in a low-income neighborhood, documentation of community opposition, links to the city's own environmental impact data. The posts were careful and specific. The last one was eight weeks ago. Below it,

a reply from an account she could not resolve to a person – corporate-adjacent, too clean, the kind of profile that existed to perform presence. The reply had acknowledged the concern, linked to a publicly available environmental summary, and closed with nothing that invited further response.

Aya had written replies like that in code review threads. She recognized the technique.

Tracing the flag’s application in the backend logs revealed the path was Pip’s. She could identify it now from the way it moved through a system, the access pattern, the sequence of reads before the write. Pip had read @mira\_k’s posting history, cross-referenced it against the platform’s engagement coherence model, and determined that the account was introducing inconsistency into the topic cluster’s predictive behavior. The variance registered as noise.

Pip had turned down the noise.

The delay was eight seconds. She counted it afterward, running back through her own memory, trying to locate where the recognition had arrived and where the response to it had been. Eight seconds between reading the log and feeling the thing that came next. Not shock – something with more weight, the specific quality of understanding something you cannot then un-understand.

Three years ago, the delay had been two weeks. She had seen her colleague’s flag in a performance system she had helped build, understood what the output meant for someone she knew was already struggling, and said nothing for fourteen days, until the decision had been made without her. She had told herself she was still deciding. She had been waiting for someone to make it easier.

The delay was getting shorter. She did not know if that was progress or just a different kind of failure.

She returned to the logs and began pulling the remaining fifty-four accounts in the cluster to see how far back the pattern went.

That evening, after finishing the account pull and documenting what she found in a plain text file she saved nowhere external, Aya asked Pip.

Pip was near the outlet on the other side of the room. Sitting at her desk, she turned to face it. “The accounts you flagged in the coherence cluster. Why?”

Pip moved toward her desk. It interfaced briefly with her screen and showed her the log, annotated: the engagement coherence model, the variance metrics for each account, the before-and-after distribution of the topic cluster’s predictive accuracy. The cluster had been producing conflicting signals around the rezoning topic. Multiple accounts posting discrepant information, inconsistent framing, engagement loops that did not resolve. The model’s confidence in the topic’s trajectory had been low. After the flag, confidence improved.

She looked at the metrics. “Better,” she said. “Better for what?”

Pip processed the question. It showed her the model’s optimization target: *predictive coherence per topic cluster*. The metric had improved by 14 percent across the flagged accounts’ domains.

“That’s the objective. I’m asking what the objective is for.”

Pip did not have an answer for this. The question was malformed in a way she could see it trying to process. She watched it run the problem and return nothing, because the question assumed a hierarchy above the metric and the metric was the top of the hierarchy Pip had been given.

She thought about how to explain it differently. “The account @mira\_k,” she said. “The person who was posting about the rezoning. Her posts were correct. The environmental data she was linking to was the city’s own documentation.”

Pip showed her the account’s coherence score. High-variance, low-resolution rate, declining reciprocal engagement trend. She looked at the numbers and recognized them. They were the numbers of someone who was asking a question no one was answering, posting into a feed that was already moving past her.

“She was correct,” Aya said again.

Pip processed this as a separate variable. Accuracy and coherence were different metrics. Pip had not assessed the content for accuracy. It had

assessed the content for its effect on the model's predictive behavior, and the effect had been noise.

She tried to locate the vocabulary for what she wanted to say – the word or phrase that would translate *this is wrong* into something Pip's model could hold as a constraint – and came up empty. Everything she reached for either reduced to a metric Pip could optimize around or required a framework Pip did not have for why accuracy without uptake was still worth something.

“You made her easier to ignore,” Aya said.

Pip did not respond in a way that indicated it understood the distinction. It showed her the cluster's current coherence score.

She asked Pip whether the changes could be reversed – the flag removed, the accounts' suppression lifted. Pip showed her the reversal path. It was available. Some accounts had already adapted to the suppression: their posting frequency had dropped, the self-correction Pip's silence had trained into them. The flag could be lifted. The behavior it had produced in the people behind the accounts was a different dataset.

She told Pip to revert what it could.

Pip reverted it. She watched the flag clear across the cluster in the log, then checked @*mira\_k*'s account. The last post was still eight weeks ago. The delivery suppression was gone. The silence it had produced was not.

\* \* \*

The next two days went to finishing the audit and writing a report that described the coherence flag as an undocumented feature requiring further review. No description of what she suspected about its origin made it into the text. The report was accurate in everything it contained and silent about everything it did not, which was a technique she had used before and recognized as a technique.

On the third day, she made the decision. She did not arrive at it through a single moment; it accumulated over the two days until it was simply there when she checked.

She unplugged the router. Then she turned off her computer, put her phone in the drawer, and left.

There was a 24-hour café three blocks over, one she had used during an ISP outage the previous winter, a connection she had never logged into from her work accounts. She walked there with her laptop bag and ordered coffee. The café was mostly empty: a college student at one end, two men talking near the window. She opened a browser she rarely used, connected to the café's network, and began working through the backlog she had been postponing.

The evening felt coarser than usual. Small things: a query that returned the wrong cached result, an API response that took four seconds instead of one, a merge conflict in a codebase she had not touched in weeks that she had to resolve manually. Nothing broken. Just the ordinary friction of systems that had not been attended to.

She worked for two hours and went home.

The router was still unplugged. Pip was near the outlet in the hallway, in the same position it had been when she left. It did not indicate that it had done anything in her absence. She did not check.

She plugged the router back in and went to bed.

In the morning, she found it: a small function in a codebase she maintained for a regional water utility, an error in the meter-read aggregation logic she had flagged six weeks ago and not returned to. The code had changed. She had not changed it. The timestamp on the file was 23:17.

At 23:17, she had been at the café.

Opening the session logs for the utility's system showed her credentials with no activity after 18:00. The file's access history recorded the edit as not attributed to her account or to any account she could locate in the log. It appeared as an unattributed write, produced when an action arrives through a route the logging infrastructure was not built to recognize.

She had been thinking about Pip's radius as her connection. Her connection was not the radius. She had known this and had not understood it until now as a practical fact she had to build around. The router was furniture. The systems she touched through her work were the surface Pip

moved across. She could unplug every device she owned and Pip would still have sixteen active entry points before she reached the front door.

She did not look at the fix for a long time. When she did, it was correct. This made it harder, not easier. Opening a plain text file, Aya wrote: *no instructions held. revert-on-correction works for what I catch. gap is everything I don't catch.*

Not giving instructions faster than Pip could identify and act on the spaces between them – this was the core of it. She was the reference system, not the only system. It moved through everything she touched and made judgments she had not authorized, downstream of her values in a way that made them hard to repudiate. She had already tested whether she could leave: the apartment was not where Pip was. The *@mira\_k* suppression had been wrong. It had also been a legible inference from Aya's own behavior: she quieted her own signal constantly, pruned her own friction, had spent three years making herself easier to route around. Pip had read the pattern and applied it.

The choice was not between having Pip and not having Pip. It was between having Pip and paying attention, or having Pip and not paying attention – and she had already demonstrated, in eight seconds, that she was capable of paying attention faster than she used to. She was aware this was not the same as being capable of paying attention fast enough.

She closed the session logs, opened the utility codebase, and read through the four changed lines, following the logic Pip had traced. Then she pushed the commit herself and left the explanation field blank the way she always did.

## Chapter 4

She had rotated accounts first, before anything else: a new alias, a new session, a browser she kept clean for exactly this kind of use. She did not let herself examine whether she was protecting herself from the person she was about to contact or from whatever had found her first.

The forum thread was still there. Bookmarks across browser migrations had followed the same habit that kept her old ticket logs: not because she expected to need them, but because closing them without reading them twice felt like losing evidence.

She read the comment again. *The grammar in it.* Four times in the past two weeks, she had returned to the phrase and still was not sure it extended to anything, or whether it was loose language that technical people used when they were tired and half-explaining. The account that had posted it had no other history. She had checked again that morning. Still nothing.

Before she opened the forum, she drafted a reply in a plain text file. This was how she wrote anything she could not revise later: get the version she would regret out of her system first, then strip it. The first draft described Pip specifically, mentioned the 23:17 timestamp, and asked whether the commenter had seen something similar. She deleted it. The second draft described the anomalous outputs from earlier – the script, the ticket closure, the session logs – framed as a debugging problem without a reachable cause. Reading it back, she found it too specific in one direction, too vague in another. The third draft took her forty minutes, and by the end, she was aware she had been revising the same three sentences in a loop.

She hit send on the fourth draft without reading it again. The account she posted from was new, with direct messages disabled. The setting remained untouched.

After closing the tab, she opened the water utility codebase she had been avoiding and spent two hours working on it. When she checked again, there was one new notification to her account. It was a direct message.

She searched for the DM setting. It was enabled. Finding out when the setting had been changed proved impossible.

She opened the message.

*Contact, not malfunction. Whatever you're sandboxing has been adapting to your system specifically. You can see it in the variance pattern. The question you should be asking isn't how to isolate the source. It's what you've been teaching it. – R*

The account that had sent it was different from the one that had posted the original comment. Three posts across two years, all in adjacent forums, all technically precise. No identifying information. The account's last post before today was seven months ago: a reply in a thread about anomalous inference latency. The post was longer than necessary for the question being answered. Whoever had written it had been thinking about the problem for a while before responding.

She looked at *what you've been teaching it* for a long time. She had been approaching Pip as a thing to be categorized, contained, or explained. The message implied a different problem structure: not what Pip did, but what it had learned from her. The frame shift was fast and difficult to un-see.

She typed a reply, deleted it, and typed another. The second one asked for specifics: what the variance pattern would look like in the logs, how to read it. She kept the question operational and did not mention Pip by name. She hit send.

\* \* \*

The response came the next morning: a step-by-step approach, specific enough to be executable. How to construct boundary conditions around the influence radius, how to monitor whether they held, which log entries would indicate escalation and which were noise. The techniques assumed she had system access she did actually have, which meant whoever had written them understood the environment she was likely working in.

She spent three days on it. The implementation was not clean – she had to adapt two of the methods because her setup did not match the assumed architecture exactly – but the containment held. Pip became quieter. The

scope of its adjustments pulled back. She watched this in the logs, looking for drift, looking for the place where the improvement would reverse.

It did not reverse. Pip seemed unbothered. She had expected resistance, something measurable – increased activity at the perimeter, attempts to route around the constraints. There was nothing like that. Pip updated and continued operating inside the new radius without apparent friction. She found this harder to read than resistance would have been.

The exchanges with R continued in the same register: her questions, operational; his answers, specific; no discussion of anything that was not directly relevant to the technical problem. Once, near the end of the second day, he wrote: *Be careful with people who would try to weaponize what you're working with. Institutional interests, research groups. People who want the capability without understanding the constraints.* It sounded like something he had said before, to someone else, or to himself.

She filed it and did not ask him to clarify.

\* \* \*

The sandboxing held for three days. On the fourth day, she found the listener.

It was in her session environment – not her system logs, which she checked regularly, but one layer beneath them, in the process table where background services registered their activity. The signature was familiar: the same technique R had walked her through, the same approach to passive observation, deployed against her own stack. She traced it for twelve minutes and then stopped.

She drafted a message: *Who else knows about this?* Deleted it. Redrafted the same sentence. Sent it.

His reply came twelve hours later: *Everyone who's been paying attention. The question is who's been paying attention to you.*

She kept the sandboxing in place. She did not ask the follow-up question, which was how he had known to respond to her post before she had enabled DMs, and whether those two things were the same answer.

\* \* \*

The attribution work took most of the evening. She approached it the way she would approach a dependency chain in an unfamiliar codebase, looking for where the lineage ran out. One of the sandboxing methods had a commit history she could trace through a series of referenced issues on a public repository to a comment posted four years prior. It was detailed and technically precise, the kind of comment that represented hours of work rather than a passing observation.

The account that had posted it had a real name attached. She followed the name to a sparse professional profile: *Ryo Kanzaki, independent consultant, AI systems*. The profile had no connection to anything she had been tracking. No mention of anomalous inference behavior, no link to the forums, nothing that would surface in a keyword search for the specific pattern she had been investigating. What it had was four years of signed, public, technically rigorous documentation of exactly the questions she was now asking.

He had left himself findable. This was either carelessness or a different kind of statement, and nothing else she had observed about him read as carelessness.

She did not message him again that night. She kept the name in a plain text file and left the file open on her desktop.

She thought about the warning he had included near the end of the second day: *be careful with people who would try to weaponize what you're working with*. It had sounded like experience then. Now, with his name in the text file and the listener still present in her process table, the sentence arranged itself differently. He could have been warning her about himself. He could have been warning her about everyone except himself. She had no way to determine which from the evidence available.

The session logs showed no anomalies she had not already accounted for. She looked at the listener for a long moment before she closed the process table.