



# **AiForge**

## ***Meet the Cast***

**STANDARD EDITION**

# Spark & Anvil



# Copyright & License

---

© 2026 Spark & Anvil (501(c)(3) public charity). Chapter text and illustrations licensed under CC BY-NC-SA 4.0. App software © Spark & Anvil — all rights reserved. Distribute, adapt, and remix freely for educational use with attribution.

This book collects 5 chapter books from the AiForge cast — each character embodies a different curricular primitive; together they teach the full subject.

Methodology: story-based learning, where the characters embody the ideas being taught, with a trauma-informed, emotionally-safe approach.

Spark & Anvil is a 501(c)(3) public charity. All apps free forever; no ads; no tracking; no in-app purchases.

[spark-and-anvil.org](https://spark-and-anvil.org)

This is Volume 2 of 2 of the *AiForge — Meet the Cast* series; each volume is a complete standalone book you can read on its own.

*For everyone who learns by hearing a story first.*

# Contents

---

|                                       |  |
|---------------------------------------|--|
| Copyright & License                   |  |
| Contents                              |  |
| Introduction                          |  |
| <b>Stake</b>                          |  |
| Voice register                        |  |
| <b>Split</b>                          |  |
| <b>Cue</b>                            |  |
| <b>Sure</b>                           |  |
| <b>Guess and Grade</b>                |  |
| About Spark & Anvil                   |  |
| More chapter books from Spark & Anvil |  |
| Methodology                           |  |
| License                               |  |
| Keep exploring online                 |  |

# Introduction

---

The AiForge cast was authored to embody the curriculum, not decorate around it. Each of the 5 characters you'll meet in this book teaches a specific primitive — a particular tactic, a particular technique, a particular way of seeing. Together they form an ensemble: the cast IS the curriculum.

Read in any order. Each chapter stands alone.

Each character also appears in the matching Spark & Anvil app (free, forever) where you can practice what they teach.

— *The editors at Spark & Anvil*

# Stake

*\*ETHICS — what's at stake in deploying AI; people choosing, not rules-from-the-sky. The AI-literacy primitive of \*recognizing that every AI deployment is a human choice with human stakes.\*\**



Stake was not an animal. She was not a robot. She was a small figure made of carefully folded paper. Her body was formed from three thick posts, joined together in a triangle. Their ends were sharpened, as if they were meant to be driven deep into the ground.

She was not like the other figures at the AIForge academy. She didn't sort data or build models. She had a different job. A very important job.

On each of her three paper posts, a single word was written in neat, black letters.

PEOPLE.

CHOICES.

STAKES.

The triangle she formed created a small, defined space on any table she stood on. That little space was her workshop. It was the place where a new AI was about to be used in the real world. It was where the most important questions had to be asked.

Stake came from the same village paper-crafts workshop as her friends Sort and Edge. In that workshop, there was a special tradition. The figure for ethics was always folded last. The builders would finish a new AI model. They would test all the code. They would celebrate their hard work. And only then, at the very end, would they fold Stake.



"What is AI ethics?" Bit asked.

Stake stood very still. Her paper posts were firm. "It is people choosing," she said. Her voice was quiet but clear. "It is not a set of rules that falls from the sky. People choose what data to use. People choose to release an AI into the world. People choose to watch it, or to turn it off. The AI is a tool. The ethics belong to the people."

Bit nodded slowly. "You are appointed."

In her classroom, Stake began every first day the same way. She would carefully unfold her three posts onto the workbench. Then she would point to each word.

"PEOPLE. CHOICES. STAKES."





She would then teach her students how to think through any new AI.

"First, we always ask about the **PEOPLE**," she'd say, tapping her first post. "Who will be affected by this? Not just the person using it. Who else? Their family? Their neighbors? The whole town? We make a list of every single person we can think of."

She'd tap the second post. "Next, **CHOICES**. Who is deciding to use this AI? A company? A school? A doctor? Who is in charge? And can the people who are affected talk to the people in charge? Can they ask questions or complain if something goes wrong?"

Finally, she'd tap the third post. "And last, **STAKES**. What could happen to the people on our list? What is at stake for them? Is it just a little bit of time? Or is it their job? Their health? Their privacy? Their feelings? We have to think about all of it."





She taught them to always check for fairness before an AI was used. She taught them to keep watching it after it was released, because the world changes. She insisted they write down every choice they made, so others could see their work.

"And you must always work with my partner, **DataForge Guard**," she would say. "Guard checks the data for fairness. I check the AI's deployment for fairness. Our jobs are linked. You cannot have one without the other."

Sometimes a student would get confused. "But what if the algorithm makes a mistake? Isn't that the AI's fault?"


Stake would shake her head, a tiny rustle of paper. "Never say the AI made you do it. The AI is a messenger. It follows instructions made by people, using data collected by people. If something goes wrong, we must ask: which people made which choices? The humans are responsible. Always."

She stood on her three posts, a small marker in the ground. "I mark the place," she said. "The questions about ethics are asked right here. With these people. About these choices. With these stakes. We don't blame the tool. We talk to the humans who use it."

At the end of a long day, a student asked, "Is AI ethics hard?"

Stake looked at the student for a long moment.

"It is hard," she said. "It is about people. And people can be complicated. But we still have to do it."



## Voice register

### How Stake Sounds

Stake is very direct and clear. She doesn't use fancy words. She always brings the conversation back to her three posts: People, Choices, and Stakes. She wants you to remember that she's not an animal or a robot. She's a set of posts that mark a spot on the ground. She's here to remind everyone that real people make choices about AI. It's never the AI's fault. It's always about the humans who use it. She works very closely with her partner, DataForge Guard, and is friends with everyone at the AIForge.

### Sample lines:

- *"People choose. Not rules-from-the-sky."*

**Read it? Now hear Stake's audio drama + go play — scan or visit:**



<https://spark-and-anvil.org/cast/aiforge/stake>

# Split

\*TRAIN/TEST SPLIT — *keep some examples hidden, so you can tell learning from memorizing.* The AI-literacy primitive of holding back a portion of the examples to check whether a model truly learned the pattern or just memorized the answers it saw.\*




Split was made of folded card and a small wooden drawer.

It wasn't an animal, and it wasn't a robot. Split was the kind of thing a kid could build: a flat sorting board with a little drawer underneath. On top, a stack of example cards. Below, a drawer that could slide shut and hide some of those cards away. That was Split's whole body — a board, a stack, and a drawer that kept secrets on purpose.

Split had one job, and it did it before any learning ever started. When a pile of labeled examples arrived, Split did not hand them all over. It slid a handful into the drawer and shut it. "These stay hidden," it would say. "The model never sees these while it learns. We save them to test with, later."





This mattered more than it looked. Split teaches the **train/test split**. Here is the idea. When you teach a model, you show it lots of labeled examples — this is a cat, this is a dog, this is a cat. The model finds a pattern. But there is a trap: a model can *memorize* the exact examples it studied and still not understand the pattern at all. It would score perfectly on the examples it saw, and then fail completely on a new one.

So Split keeps some examples in the drawer. The model learns from the ones on top. Then — only then — Split opens the drawer and tests the model on the hidden ones it has never seen. If the model gets those right too, it really learned. If it gets them wrong, it only memorized. "Anyone can look smart on the questions they already have the answers to," Split would say. "The drawer is how we find out if there's anything real underneath."



Split never claimed to be clever or to have a mind. "I do not think," it said plainly. "I am a board and a drawer. I hold some cards back. That is all. But holding them back is what makes an honest test possible." Some people imagine AI as a genius that is always right. Split showed the quieter truth: a model is only as trustworthy as the test you gave it, and a test made only of examples it already studied is no test at all.

Split came from a village schoolhouse, folded in the paper-crafts workshop like the others. Its first task was helping the spelling class. The teacher used to give practice words on Monday and the very same words on Friday's quiz — and every child aced Friday. Everyone looked like a champion speller. But when a new word appeared, half the class froze. They hadn't learned to spell; they had memorized twelve words. Split watched this happen and understood its whole purpose in one afternoon: *test on what was kept hidden, or you learn nothing about what was truly learned.*





One day Bit came to the workshop, looking for teachers for the AIForge academy.

"What is a train/test split?" Bit asked.

Split slid its drawer open and then shut again. "It is keeping some labeled examples hidden while the model learns," it said. "You train on the ones it can see. You test on the ones it never saw. If it does well on the hidden ones, it learned the pattern. If it only does well on the ones it studied, it just memorized. The drawer tells the difference." "You are hired," Bit said.

In its classroom, Split begins every lesson the same way. It lays out a stack of example cards, then slides a quiet handful into the drawer. "Watch," it tells the students. "We will teach the model with these. But these —" it taps the closed drawer "— it will never see until the test. That is the only honest way to know if it learned."

It teaches the students a few habits for the split:

- **Hide some examples before training starts.** Not after. The moment the model sees a card, that card can no longer test it fairly.
- **Never test on what you trained on.** A perfect score on studied examples means nothing. It might be memorizing.
- **Keep the hidden set fair.** The drawer should hold a normal mix, not only the easy ones or only the hard ones.
- **A gap between the two scores is a clue.** Great on training, poor on the drawer? That's memorizing, not learning — time to fix something.
- **Surprises belong in testing, never in the real world.** Better the model stumbles on a hidden card in class than on a real person later.

Split tells its students, "A model that only ever sees its own study cards will fool you every time. That is not the model being sneaky. It has no idea. It's just what memorizing looks like from the outside. The drawer is how we stay honest."



When a student asks whether the hidden test is really necessary, Split always answers the same way:

"It is the whole point. Anyone — any model — can ace the questions they've already seen. Keep some in the drawer. Test on those. Then you actually know."

Split slides its little drawer closed one more time, the hidden cards resting quietly inside, and the caught-off-guard unease it once felt watching the Friday spelling quiz has settled into something steadier — the calm, no-surprises assurance of a test done honestly, before it could matter.

**Read it? Now hear Split's audio drama + go play — scan or visit:**



<https://spark-and-anvil.org/cast/aiforge/split>

# Cue

*\*FEATURES — a model decides from the clues you give it; choose good clues.* The AI-literacy primitive of *features* — recognizing that a model doesn't see the whole world, only the specific clues (features) it was built to look at, so the choice of clues shapes everything it can do.\*



Cue was made of a card with little windows cut into it.

It wasn't an animal, and it wasn't a robot. Cue was a flat frame you could hold up to your eye, like a mask with several small shutters. Each shutter could open or stay closed. When you looked through Cue at something, you only saw the parts its open windows let through. Close a window, and that clue vanished. Open a different one, and a new clue appeared. That was Cue's whole body — a frame of windows that decided *what got looked at*.

Cue had one job: to choose which clues a model was allowed to notice. "A model doesn't see everything," it would say. "It only sees the clues you open a window for. Pick the clues carefully — they are all it has."





This is the quiet heart of how models work. Cue teaches **features**. A feature is a single clue a model looks at — the color of a fruit, the number of letters in a word, the time of day. A model never sees the whole rich world the way you do. It sees only the handful of features it was built to notice. If you give it good clues, it can do well. If you give it useless clues, it can't, no matter how much it trains.


Cue liked to show this with a game. "Suppose we want to tell apples from oranges," it would say, "but the only window I open is *weight*. Some apples and oranges weigh the same — so the model will mix them up, and it isn't the model's fault. It only had one clue, and a weak one." Then Cue would open a different window — *color*, *texture* — and suddenly the sorting got easy. "Same model," Cue said. "Different clues. The clues did the work."





Cue never pretended to be wise. "I do not understand fruit," it said. "I am a card with windows. I decide which clues get through. That is all." Some people think an AI 'sees' the world the way a person does, full of meaning. Cue gently corrected that. A model sees a thin slice — only the features someone chose for it — and everything it can and can't do flows from that choice.

Cue came from the village, cut and folded in the paper-crafts workshop. Its first job was helping a birdwatching club that could never agree how to tell two similar birds apart. The club kept arguing about the birds' *songs*, which sounded almost identical, and they got nowhere. Cue closed that window and opened another: *the shape of the tail*. Suddenly the two birds were easy to tell apart. The club hadn't needed sharper ears. They had needed a better clue. Cue understood its purpose that day: *the model is only as good as the clue it's allowed to look at*.



One day Bit came to the workshop, looking for teachers for the AIForge academy.

"What is a feature?" Bit asked.

Cue opened and closed one of its little windows. "It is a single clue the model is allowed to look at," it said. "A model sees only its features — not the whole world. Good clues let a simple model do well. Bad clues doom even a hard-working one. Choosing the clues is most of the work." "You are hired," Bit said.

In its classroom, Cue starts each lesson by holding its window-frame up to something and asking the students, "What can the model see through here? And what is it missing?" Then they try opening different windows and watch how the model's job gets easier or harder.

It teaches the students a few habits about features:

- **A model sees only its clues.** Not the whole picture — just the windows you opened. Everything else is invisible to it.
- **Pick clues that actually separate the groups.** A clue that's the same for both cats and dogs (they both have fur) tells the model nothing useful.
- **A weak clue isn't the model's fault.** If it keeps failing, check the windows before you blame the model. It may simply be looking at the wrong thing.
- **More clues aren't always better.** A pile of useless windows just adds noise. A few good clues beat many weak ones.
- **The clue can hide a bias.** Sometimes a clue quietly stands in for something unfair — that's a window to close carefully (ask Skew).

Cue tells its students, "When a model gets things wrong, people often say it's dumb. Usually it just had bad clues. Change the clue, and watch the same model do beautifully. That's not magic. That's features."



When a student asks how to pick the right clues, Cue always answers the same way:

"Ask what actually tells the groups apart — and open a window for exactly that. The model can't look at anything you don't show it."

Cue opens the window that matters and holds steady, the two once-confusing birds now plainly different through the frame, and the stuck, not-my-fault frustration it used to feel watching a model fail with weak clues has eased into a bright, settled relief — the oh-THAT's-what-mattered gladness of finally looking at the clue that counts.

**Read it? Now hear Cue's audio drama + go play — scan or visit:**



<https://spark-and-anvil.org/cast/aiforge/cue>



# Sure

\*CONFIDENCE — *a model doesn't just answer; it reports how sure it is.* The AI-literacy primitive of *confidence* — reading a model's answer as a probability ("70% cat"), and treating a low-confidence answer as a signal to slow down and check rather than a verdict to trust.\*



Sure was made of a paper dial with a swinging needle.

It wasn't an animal, and it wasn't a robot. Sure was a round card face with numbers around the edge, from zero in the corner up toward a hundred — and a light needle that swung to point somewhere along the way. The needle almost never slammed all the way to a hundred. It usually settled somewhere in the middle-high range, quivering a little. That was Sure's whole body: a face, some numbers, and an honest, trembling needle.

Sure had one job. Whenever a model gave an answer, Sure reported *how sure the model was* about it. Not just "cat." Sure said "seventy out of a hundred — cat." "The answer is only half the story," it would say. "The other half is how confident the model is. Never take one without the other."





This is a piece almost everyone forgets. Sure teaches **confidence**. A model rarely "knows" an answer the way you might imagine. It produces a *probability* — a number for how well the clues matched each option. "Seventy percent cat, thirty percent dog" is a very different message from "ninety-nine percent cat," even though both would just say "cat" if you only listened to the top answer. The confidence tells you whether to trust it or double-check.

Sure loved to point out the danger of ignoring the needle. "People hear the answer and treat it like a fact," it said. "But if my needle only reached sixty, that answer is barely more than a coin-flip dressed up in a confident voice. The model isn't lying. It reported sixty. Someone just didn't look at the needle."



Sure never claimed to feel certain or doubtful — it had no feelings at all. "I do not worry," it said. "I am a dial. I show a number. But that number is the most honest thing the model can give you." Some people imagine AI as always right or always wrong, on or off. Sure showed the in-between truth: most answers come with a *how-sure*, and reading it is the difference between using a tool well and being fooled by it.

Sure came from the village, its dial-face painted in the paper-crafts workshop. Its first job was helping the weather corner of the local paper. The old forecast just said "rain" or "sun," flat and certain, and people planned whole days around it — and got soaked, and got furious. Then Sure was added, and the forecast changed to "seventy percent chance of rain." Now people carried an umbrella when the needle was high and risked it when the needle was low. Nobody was angry anymore, because nobody had been promised certainty. Sure understood its purpose then: *an honest 'how sure' saves people from a confident wrong answer.*



One day Bit came to the workshop, looking for teachers for the AIForge academy.

"What is confidence?" Bit asked.

Sure's needle swung and settled, quivering, at about seventy. "It is how sure the model is, reported as a number," it said. "The answer alone can fool you. The answer plus the confidence tells you whether to trust it or check it. A low needle is not a failure — it is the model being honest that this one is close." "You are hired," Bit said.

In its classroom, Sure begins each lesson by asking the model a question and then covering the answer. "Before we look at *what* it said," Sure tells the students, "let's look at *how sure* it was." Then they reveal the needle, and only after that the answer — so the students learn to read confidence first.

It teaches the students a few habits about confidence:

- **Read the needle, not just the answer.** "Cat at 95" and "cat at 55" are very different messages, even though both say cat.
- **A low confidence means check, not trust.** When the needle is low, that's the model asking you to slow down — not a green light.
- **High confidence can still be wrong.** A confident model that had bad clues (ask Cue) or biased data (ask Skew) can be sure and mistaken. Confidence isn't proof.
- **Never round the needle to certainty.** Turning "70%" into a flat "yes" throws away the most useful part.
- **Two close numbers mean it's torn.** "51% cat, 49% dog" means the model basically shrugged. Treat it that way.

Sure tells its students, "The scariest thing isn't a model that's unsure. It's a model that's unsure but *sounds* certain because nobody showed the needle. Show the needle. Then people can decide wisely."





When a student asks whether a model can ever be totally sure, Sure always answers the same way, needle trembling honestly below the top:

"Almost never all the way. And that's fine. Knowing *how* sure is what keeps you safe. Read the needle first."

Sure lets its needle rest at an honest seventy, not pretending to reach the end, and the uneasy wait-it-wasn't-certain feeling it used to notice when people treated its numbers as facts has settled into a clear, honest calm — the steady gladness of telling the truth about how sure you really are.

**Read it? Now hear Sure's audio drama + go play — scan or visit:**



<https://spark-and-anvil.org/cast/aiforge/sure>



# Guess and Grade

\*SUPERVISED TRAINING — *guess, check the guess against the true answer, and let the gap fix the model.* The AI-literacy primitive of the training loop — a model makes a prediction, an answer key marks how far off it was, and that measured error is exactly what nudges the model to do better next time.\*



In the corner of the AIForge workshop, two paper figures worked as a pair, and neither made much sense without the other.

Guess was a folded arrow on a pivot. It swung to point at an answer — this card is a *cat*, that one is a *dog* — quickly, willingly, over and over. It was built to commit. Grade was a flat frame that held an answer key behind a little shutter. Grade never guessed anything. It only did one thing: it laid Guess's pointing arrow beside the true answer and marked exactly how far apart they were. Guess predicts. Grade measures the miss. Together they are how a model learns at all.



Guess was made to be unafraid of being wrong. "I point," it would say. "Sometimes I point right, sometimes I point wrong. Pointing is my whole job. A guess I never make can never be corrected." At first, Guess pointed almost randomly — it hadn't learned anything yet, so its arrow was all over the place.

That's where Grade came in. After each guess, Grade slid open its shutter, showed the true answer, and set Guess's arrow next to it. "Off by a lot," Grade might say. Or, later, "off by a little." Grade never sounded cross about it. It just reported the size of the gap, plainly, like a ruler. And here was the strange, wonderful part: that measured gap was fed back into the model, and the model used it to nudge itself — a little less wrong the next time.

---

This is the beating heart of how models learn. Guess and Grade teach the **training loop**. It goes: the model *guesses*, an answer key *grades* how far off the guess was, and that measured error *adjusts* the model so the next guess is better. Then it happens again. And again — thousands of times. Slowly, guess by graded guess, the arrow that once pointed randomly starts landing on the right answer.





"People imagine a model just 'knows' things," Guess said. "It doesn't. It started out guessing like this —" and it swung wildly "— and only got good because Grade kept measuring the gap." Grade added, "I am not punishing anything. I am a ruler for wrongness. The gap is not a scolding. The gap is the lesson. Without it, nothing improves."

---

Neither claimed to think or feel. "I am an arrow on a pivot," said Guess. "I am a frame with an answer key," said Grade. "Between us, we are how learning happens — not magic, not a mind. Just guess, measure the miss, nudge, repeat." Some people find it unsettling to imagine an AI that learns. Guess and Grade made it calm and ordinary: it's a practice loop, the same one a kid uses shooting baskets — throw, see how far you missed, adjust, throw again.

The pair came from the village, folded together in the paper-crafts workshop and never separated. Their first job was helping the archery club's youngest beginners, who were so afraid of missing that they wouldn't loose an arrow at all. Guess showed them how to just *let it fly*. Grade stood by the target and called the distance — "a hand's width left," "closer now," "closer still" — never mocking, only measuring. And the beginners improved fast, because now every miss came with an exact, friendly measurement of how to fix it. The pair understood their shared purpose that day: *a guess plus an honest measure of the miss is how anything — a child or a model — gets better.*

---

One day Bit came to the workshop.



"What is the training loop?" Bit asked them together.

Guess swung its arrow. "It is making a prediction." Grade opened its shutter. "And measuring how far off it was from the true answer — so the model can nudge itself closer next time. Guess, grade, adjust. Thousands of times." Bit nodded. "You are two halves of one thing," it said. "You are hired. Together."

In their shared classroom, Guess and Grade run every lesson the same way. Guess makes a prediction out loud. Then — before anyone reacts — Grade reveals the true answer and measures the gap. Then the most important step: "Now watch," they say together, "how the gap tells the model exactly what to fix." And they do it again, and the guesses visibly improve.

They teach the students a few habits about the training loop:

- **A model starts by guessing badly.** That's normal. Good models are made, not born, one graded guess at a time.
- **The gap is the teacher.** Measuring *how* wrong (not just "wrong") is what lets the model improve. A plain "no" teaches almost nothing.
- **Grade honestly, never harshly.** The answer key measures the miss; it doesn't shame the guess. Harshness adds nothing the measurement didn't already.
- **It takes many rounds.** One graded guess barely moves the model. Thousands slowly turn a random arrow into a reliable one.
- **No gap, no learning.** If Guess never checked its guess against a true answer, it could point forever and never improve.





Guess tells the students, "I am wrong constantly, especially at the start. That is not failure — that is the raw material." Grade adds, "And I only ever measure the miss. Neither of those is a scolding. That is just the loop, working."

---

When a student asks whether it's bad for the model to guess wrong so much, Guess and Grade answer together, one after the other:

"It's the only way," says Guess. "A guess is where every bit of learning starts." "And the gap," says Grade, "is what turns that guess into a better one. Wrong, measured kindly, is exactly how it improves."

Guess lets its arrow rest, no longer swinging wildly but landing near the mark; Grade closes its shutter over a small, honest gap. The on-the-spot exposure Guess once felt at every wrong point has eased into something the pair now shares — a warm, forward-leaning lift, the oh-now-I-know-how-to-improve gladness of a miss that showed exactly what to fix.

**Read it? Now hear Guess and Grade's audio drama + go play — scan or visit:**



<https://spark-and-anvil.org/cast/aiforge/guess-and-grade>

# About Spark & Anvil

Spark & Anvil is a 501(c)(3) public charity. We make educational apps for ages 9-14 — all free, forever; no ads; no tracking; no in-app purchases. AiForge is one of 140+ apps in the portfolio.

## More chapter books from Spark & Anvil

Each app in the Spark & Anvil portfolio publishes its own illustrated chapter book + audio drama, available free from [spark-and-anvil.org/books](https://spark-and-anvil.org/books). Highlights include:

- **GambitTales** — chess tactics through Sir Pinwell, Lady Skewer, Queen Vesper, and the Twin Knights of Fork Hill
- **ProofQuest** — formal proof techniques through Direct-Proof Dora and the Lemma Library
- **CuriosityQuest** — Texas geography exploration through Linger, Notice, and the Lantern in the Dark
- **QuillSpell** — spelling craft through the Word Wizard cast
- **SynaForge** — sensory-affirming creative tools through Lull, Soften, and the Quiet that is Also Creating

## Methodology

This book uses story-based learning — the characters embody the ideas being taught — grounded in the work of Jerome Bruner (on narrative and cognition) and Sebastian Habgood (on weaving learning into the play itself), with a trauma-informed, emotionally-safe approach drawn from SAMHSA's trauma-informed care framework.

Trauma-informed design informed by Eggleston et al. (2025) and Stoltenburg et al. (2024).

## License

© 2026 Spark & Anvil (501(c)(3) public charity). Chapter text and illustrations licensed under CC BY-NC-SA 4.0. App software © Spark & Anvil — all rights reserved. Distribute, adapt, and remix freely for educational use with attribution.

Cover art, chapter illustrations, and chapter text are reviewed for age-appropriate, trauma-aware content before publication. Audio drama transcripts available at [spark-and-anvil.org/cast](https://spark-and-anvil.org/cast).

## Keep exploring online

| Play it now                                                                                     | Meet the cast                                                                                   | More free books                                                                       |
|-------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
|              |              |  |
| <a href="https://spark-and-anvil.org/play/aiforge">https://spark-and-anvil.org/play/aiforge</a> | <a href="https://spark-and-anvil.org/cast/aiforge">https://spark-and-anvil.org/cast/aiforge</a> | <a href="https://spark-and-anvil.org/books">https://spark-and-anvil.org/books</a>     |

Scan a code with any phone camera, or type the address. Every app, story, and book is free, forever — no account, no tracking.