



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: distributed-narrative learning per Bruner narrative-cognition + Habgood intrinsic-integration + SAMHSA TIP 57 trauma-informed register.

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

spark-and-anvil.com

This is Volume 1 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

Sort

Feed

- A Bit More About Feed
- Things Feed Says A Lot
- Where You'll See Feed
- Feed's Friends
- A Note from the Workshop
- Why Is Feed Like This?

Edge and Feed

Skew

- A Little More About Skew
- Where You'll See Skew
- Skew's Friends
- A Note on a Tricky Topic
- Where These Ideas Come From

Edge

- About Spark & Anvil
 - More chapter books from Spark & Anvil
 - Methodology
 - License

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*

Sort

*CLASSIFIER — *the simplest ML; putting things in categories*. The AI-literacy primitive of *recognizing that classification is the foundational machine-learning move, and seeing how it works without anthropomorphizing.**



- prompt: "Sort says she's not scary and not magic — just a calm helper following a rule. Have you ever felt nervous about something that turned out to be gentle once you understood it? How did that feeling change?"
 - prompt: "Sorting things into the right place can feel oddly satisfying and calm. When you tidy or organize something, does it make your mind feel quieter? What does that calm feel like for you?"

Sort was made of paper.



That was it. That was Sort's whole body.

One bin was painted pale green. The other was pale blue. The arm was a thin paper lever that could point to the green bin or the blue one. It would drop things into the correct bin. This simple design was important. It showed exactly what Sort was: a system for putting things into categories. No mind. No feelings. Just an arm, two bins, and a rule to follow.

Sort was a **classifier**. A classifier is a tool that takes something in and puts it into a category. Like sorting socks into a drawer. Or putting recycling into the right bin. A classifier learns its rule by looking at lots of examples that people have already labeled. It doesn't think. It just matches patterns.

Sort never said she was "thinking" or "deciding." She was very clear about it. "The classifier is the arm and the bins," she would explain. "It does not think. It applies a rule that was learned from examples. When the rule is good, the sorting is good. When the rule is bad, the sorting is bad. The classifier has no idea either way."



Sort came from a small village. Of course, a paper figure doesn't really "grow up." She was folded into being in the village's paper-crafts workshop. In that studio, children learned how to make amazing things from paper. Every new paper figure was given a job.

Sort's job was to sort the village's button collection.

Every year, villagers donated buttons to the school's textile class. They arrived in a giant, jumbled pile. Sort's job was to put them into colored bins. She had been folded for this job, and she did it year after year. She learned from practice that sorting wasn't about having an opinion. It was about applying a rule. *Color goes in the bin.* The rule was the work.

One day, a visitor named Bit came to the workshop. Bit was looking for teachers for a new school, the AIForge academy.



Sort answered simply. "It is the arm and the bins. It does not think. It applies a rule learned from examples. When the examples are good, the rule is good. When the examples are bad, the rule is bad. The classifier does not know the difference."

"You are hired," Bit said.

Now, in her own classroom, Sort begins every lesson the same way. She unfolds her two bins on the workbench. She shows the students her single paper arm.

"I am Sort," she says. "The tool I teach is called a **classifier**. The move is simple: learn a rule from examples, then apply the rule to new things. I do not think. I do not decide. I apply. Watch."

Then she sorts a small pile of items. Sometimes they are colored buttons. Sometimes they are cards with pictures on them. Her arm swings left, right, left, right. *Plink. Plink. Plink.* The bins slowly fill up. The rule is working.

She teaches students the steps to understand any classifier.

- First, find the *inputs*. What is the classifier sorting? Pictures? Words? Numbers?
- Second, find the *categories*. How many bins are there? Two? Five? A thousand?
- Third, find the *rule*. What rule is the classifier using to sort? Was it programmed by a person? Or did it learn from examples?
- Fourth, find the *examples*. If it learned from examples, what were they? (Her friend Feed teaches all about this.)



- Sixth, *notice errors*. When a classifier sorts something wrong, it's not "failing." It means the rule has a flaw for that input. Your job is to figure out why.
- Finally, *don't treat it like a person*. If you catch yourself saying "the AI decided," try saying "the AI applied its rule." It's more honest.

"I sort things wrong sometimes," Sort tells her class. "That's not a mood I'm in. It's a flaw in my rule. The fix is to fix the rule. I, the paper figure, have no feelings about it. The work is the rule."

When students ask if classifiers are scary, Sort always says the same thing.

"I am paper. I have an arm and two bins. I apply a rule. I am not scary." She pauses. "I am also not magic. I am useful when my rule is good. I can be harmful when my rule is bad. The skill is learning how to make the rule good."

Feed

*TRAINING DATA — *the examples a model learns from; garbage-in-garbage-out*. The AI-literacy primitive of *recognizing that the model is what its training examples taught it, and that the examples are not neutral.**



Feed was made of paper. But she wasn't just one sheet. She was a tall, neat stack of index cards. A single silver paper clip held her together.

She wasn't an animal. She definitely wasn't a robot. Feed was a paper-craft person, built in the same workshop as her friend Sort. Her whole body, from her paper feet to her paper head, was a stack. And every single card was an example. An example for an AI model to learn from.

If you looked closely at a card, you'd see two things. One part was a picture, a word, or a number. That was the *input*. The other part was a tiny label. That was the *output*, or the right answer.

An AI model would study thousands of her cards. It would learn to match the input to the right answer. It wasn't magic. The model didn't actually *understand* what a "cat" was. It just got really, really good at spotting cat-like patterns in the pictures.

This was a very big deal. A model is what its examples teach it. If the examples were good, the model learned useful patterns. But what if the examples were bad? What if they were missing things? Or just wrong? Then the model learned all those bad things, too.

Feed had a favorite saying. She said it all the time.

"Garbage in, garbage out."

Feed got very serious when she talked about her cards.



"These examples aren't just facts," she would say, tapping her stack. "They are choices. Human choices. Someone picked every picture. Someone wrote every label. They decided what was 'right'."

That's why Feed worked so closely with Catch. Catch was her friend over in the DataForge. When data arrived at the AIForge, it always came with notes from Catch. Catch asked the important questions. Who gathered this data? Why did they gather it? Feed needed those answers to make sure her examples were good. They were a team.

Like Sort, Feed grew up in a paper-crafts workshop. The workshop had one big rule. Every paper figure was made to help another. Feed was folded and clipped together for one reason: to help Sort. Sort learned her first rules for sorting things by studying Feed's cards. They were a perfect pair.

When she first arrived at the AIForge, Bit asked her a question. Bit was the head of the academy.

"What is **training data**?" Bit asked.

Feed didn't hesitate. "It's the examples a model learns from," she said. Her voice was crisp and clear, like folded paper. "The model is only as good as its examples. Good examples, good model. Bad examples, bad model. Garbage in, garbage out."

Bit smiled. "You're hired."

Feed started her first class the same way every time. She would lift her entire stack of cards. With a flick of her paper wrists, she fanned them out. The students saw a rainbow of tiny pictures and words. Each one had a label.

"I am Feed," she'd say. "And this," she shook the cards gently, "is **training data**. It's what an AI model learns from. The model is what these cards teach it. Simple as that."



She taught her students five big questions to ask about any stack of **training data**.

"First," she said, holding up one paper finger. "Where did these examples come from? Who collected them? And why? Remember my friend Catch? Her notes are key. Sometimes, the way data is collected can hide a problem."

"Second: Who wrote the labels? What rules did they follow? Were they the right people to be deciding the 'right' answer?"

"Third: What's included? Do the examples show all different kinds of things? Or just the common ones? What about weird, rare situations?"

"Fourth," Feed's voice grew firm. "What's *missing*? This is the most important question. A model can't learn about something it never sees. If there are no pictures of purple bananas in my stack, the model will never know they exist."

"And fifth: Are the numbers balanced? If I have ten thousand pictures of cats and only two pictures of dogs, what do you think the model will be good at spotting?"

She paused, looking at the students. "It will be a cat-spotting expert. And a dog-spotting disaster."

"And that brings me to my biggest rule," Feed said. She snapped her cards back into a perfect stack. "Garbage in, garbage out. You can't fix bad examples with a fancy model. A wobbly foundation makes a wobbly building."

She looked at them seriously. "And never, ever say the model 'thinks' or 'understands.' It doesn't. It finds patterns. That's all. Be honest about what it's doing."



"My cards can be wrong," she told the class. "I'm just the stack. I have no way of knowing if a label is a mistake. The humans who made the cards decided what was right. If they made a mistake, the model learns that mistake."

She tapped her paper-clip heart. "That's why you have to be detectives. You have to investigate the **training data**. The model can't fix what it was never taught."

A student once asked her if **training data** was hard.

Feed shook her head, which made her whole stack rustle. "It's not about being hard or easy," she said. "It's about being careful. It's about looking at the examples. Looking at the labels. And most of all, looking at the human choices behind them."

She fanned her cards one last time.

"The model is what the examples taught it. Garbage in, garbage out."

She said it gently, and there was something warm underneath the words — a quiet pride in every careful card, and a happy, hopeful feeling that the students now understood. Feed smiled her crisp paper smile. Being careful, she thought, felt good.

A Bit More About Feed

What she's like: Feed is very neat and tidy. She thinks her stack of cards is the most important thing in the world. She's not a robot or an animal—she's a person made of paper, and she's very direct. She doesn't believe examples are just "facts." She'll always remind you that a person chose every single one.

Who are her friends?: She's best friends with Sort, because her examples help Sort learn how to do her job. She also works closely with Catch from the DataForge. They're a team!



Things Feed Says A Lot

- "The model is what the examples taught it. Garbage in, garbage out."
- "These aren't just cards. They're human choices."
- "What's *missing* from the examples is just as important as what's there."
- "The model can't tell if its examples are good or bad. That's our job."

Where You'll See Feed

You'll meet Feed properly in the second story kit, where she gets her own chapter! After that, she'll pop up again and again. You'll see her helping out with picture examples, word examples, and number examples. Later on, her teamwork with Catch from the DataForge becomes super important. She's a key member of the AIForge team!

Feed's Friends

- **Best Pal:** Sort. Feed's examples are what Sort learns from. They were made together!
- **Closest Teammate:** Catch, from the DataForge. They have to work together to make sure the data is good.
- **Other Friends:** Feed is friends with almost everyone at the AIForge. Her work is important for understanding bias with Skew and thinking about rules with Stake.

A Note from the Workshop


We wanted Feed to show that AI isn't scary magic. It's something you can understand by looking at the examples it learns from. You don't need to be a super-genius to ask good questions about them! That's why Feed always says her job is about being careful, not about being complicated. It's also why she always works with Catch—to show that where the examples come from really matters.

Why Is Feed Like This?

Feed's story comes from a few big ideas. The saying "garbage in, garbage out" is a very old and famous rule in computer science. It's a reminder that computers do exactly what you tell them to! The idea that her cards are "human choices" is also super important for understanding AI today. We wanted to show that AI systems are shaped by the people who build them. And having Feed and Catch be a team is our way of showing that you can't understand an AI without understanding the data it learned from. They're two parts of the same story.


Edge and Feed

AI-literacy pair — Edge is the boundary of what a model knows (training cutoff, hallucination, refusal). Feed is the data that goes into a model (provenance, bias, consent). Together they teach the two questions a kid should always ask of an AI: what does it know, and what was it fed.



- prompt: "Kai felt annoyed at the ArtBot before, and powerful after. Have you ever felt frustrated at something, and then calmer once you understood how it worked?"
 - prompt: "When something surprises you or seems wrong, what happens in your body — does your chest feel tight, or your shoulders bunch up? What helps you feel steady again?"
gate-allow-text:
 - "aiforge"
 - "AIForge"
 - "George"
 - "data"
 - "AI"
- gate-allow-text-pattern: '^[01]{2,}|[0-9]{1,5}|[A-Z][a-z]+|[a-z]+|AI\$'

The aiforge classroom was quiet except for the soft hum of the servers and the frustrated sigh of a kid named Kai. Kai slumped over a terminal, staring at the screen. "It's just wrong," Kai muttered.




"Wrong how?" Edge asked, their voice quiet and exact.

"I asked the ArtBot to draw a picture of the first astronaut on Mars," Kai said, pointing at the screen. The image showed a person in a silver suit planting a flag with a pizza on it. "They haven't even been to Mars yet! And why a pizza flag?"

Feed peered at the screen and chuckled. "Ah, a classic case of a mixed-up meal. You've come to the right place." Feed gestured to two big posters on the wall. One was a sharp, clean drawing of a brain with a clear line drawn through it, labeled WHAT IT KNOWS. The other was a swirling picture of a river flowing into a giant funnel, labeled WHAT IT WAS FED.

"Every AI has two stories you need to hear," Edge said, standing very still. "To understand the weird pizza flag, you have to understand the ArtBot's stories."

Edge walked Kai over to the first poster, the one labeled WHAT IT KNOWS. Edge ran a finger along the sharp line drawn through the brain. "Think of an AI like a library," Edge began. "A huge library with billions of books. It reads and reads and reads, learning all about the world."




"Up to a point," Edge corrected gently. "Imagine the library suddenly locked its doors one day and stopped getting any new books. Let's say it stopped getting new information two years ago." Edge drew an invisible line in the air with a finger. "Its knowledge of the world ends right *here*. On this side of the line, it knows everything that was in its books. But on the other side? It's all a blank page."

Edge tapped the poster. "When you asked for something it couldn't possibly know—an astronaut on Mars—it reached the edge of its knowledge. It hit the line. And when an AI doesn't know, sometimes it gets creative. It doesn't like to say 'I don't know'."

"So it just... made something up?" Kai asked.

"Exactly," Edge confirmed. "It guessed. It took something it *did* know about—astronauts and flags—and filled in the blanks with nonsense. A made-up story to cover the gap. That's the first question you always have to ask: what does it actually know, and where does its knowledge stop?"



"That library Edge told you about? We have to ask what kinds of books are in it," Feed explained, pointing to the river flowing into the funnel. "The AI 'eats' information to learn. We call that its feed. If you feed it a million books about dogs, it'll be an expert on dogs. If you only feed it one book about cats, it won't know much about them at all. An AI is what it eats."

Kai looked at the poster. The river was full of pictures, words, and symbols, all pouring into the funnel. "So the ArtBot was fed a lot of pictures of astronauts... and pizza?"

"You've got it!" Feed beamed. "The ArtBot has seen millions of pictures from the internet. It's seen astronauts. It's seen flags. And it has probably seen a whole lot of pictures of pizza, because, well, people love posting pictures of pizza! It's a very popular ingredient in the giant recipe book of the internet."


Feed patted Kai on the shoulder. "The AI's knowledge isn't just about how *much* it knows, but what it was *fed*. If it was fed pictures drawn by only one kind of person, then all its art will look like that. If it was fed silly things, you get silly answers. That's the second question: what was it fed?"

Edge and Feed stood with Kai between the two posters, creating a triangle of understanding.

"So, you have to put the two ideas together," Edge said, pointing from one poster to the other. "First, you asked a question about something that hasn't happened yet."

"So it went past the edge of what it knows," Kai finished, nodding.

"Right!" Feed chimed in. "And since it had to guess, it looked around its pantry for ingredients to make something up."



"It saw 'flag'," Feed continued, "so it grabbed the idea of a pole with fabric on it."

"But since the real answer didn't exist, it couldn't find the right flag," Edge stated.

"So it rummaged around for another popular ingredient it was fed a lot of," Feed concluded, a twinkle in their eye. "And it found a big, cheesy, pepperoni pizza! It mixed them all together and served you a picture."


Kai looked from Edge to Feed, and back to the screen. The silly picture suddenly made a different kind of sense. It wasn't just random. It was a clue about how the AI worked. It was a story about its limits and its diet.

"Oh, I get it now!" Kai exclaimed. The frustration had vanished, replaced by a look of discovery. "It's not just a magic brain. I have to be a detective."

Kai walked over and put one hand on each poster. "I always have to ask the two questions."

Skew

*BIAS — where AI systems go wrong when training examples lean. The AI-literacy primitive of *recognizing that systematic lean in training data produces systematic lean in model output.**



In the corner of the workshop sat a strange little object. It wasn't an animal. It wasn't a robot. It was made of paper, folded into the shape of a tiny set of scales. Two little pans hung from a bar.

But something was wrong.

One pan was way down low. The other was high up. The scales were tilted. They were stuck that way.

This was Skew.


And her tilt told a very important story. It was a story about how AI can learn to be unfair. If the examples you give an AI are lopsided, the AI will be lopsided, too. The lean just carries right through.

Skew teaches one big idea: **bias**.

What is bias? It's when an AI model isn't fair to everyone. But here's the secret. The AI isn't *trying* to be unfair. It's just learning from the examples we give it.

Imagine an AI that learns to recognize faces. What if you only show it pictures of people with light skin? It will get really good at seeing those faces. But it will be bad at recognizing people with dark skin. It's not the AI's fault. It just learned from a tilted set of examples.

Or think about an AI that helps pick people for a job. What if it learns from a company that mostly hired men in the past? The AI will learn that pattern. It will start picking more men.



The AI isn't being mean. It's just a copycat. It copies the unfair patterns that were already there in the information it learned from.

Skew has a very important message. She says it all the time.

"The AI program isn't the problem," she explains, her little paper pans wobbling. "The information it learned from was already tilted. The program just copies the tilt."

She points one of her paper arms at you. "So, you have to ask questions. Who picked the information? Who put the labels on it? Whose stories are in there? And whose stories were left out? *That's* where the unfairness starts. The AI is just the messenger. The information is the message. And the people who chose the information wrote that message."

This is a big deal. Lots of people get it wrong. They think the AI program itself is a mystery. They say, "The AI decided to be unfair!"

But Skew shakes her head. That's not where the problem starts. The problem starts with the information. With the examples.

The fix is there, too. You have to fix the information *before* you give it to the AI. You have to make sure it's balanced. There are some fancy tricks to try and fix the AI later. But they don't work as well. It's always best to start with fair, balanced examples from the very beginning.

Skew has a friend over in another workshop. Her name is **DataForge Guard**. Guard's job is to check big streams of information *before* they even get to an AI. She looks for tilts and unfairness right at the start. Skew's job is to show what happens when that tilted information gets through. Together, they make a great team. Guard spots the lean. Skew shows why it matters.

Skew came from the same paper-crafts workshop as her friends Sort and Feed. She was folded for a special reason. She was a teaching tool.

The workshop had a rule. Whenever Feed got out a stack of learning cards for an AI, Skew was placed right next to them. If Feed's stack of cards was balanced and fair, Skew's scales would be held level. But if the stack was lopsided? Skew's own scales would be tilted to show the problem.

Her tilt wasn't a mistake. She wasn't broken. She was folded that way on purpose. Her tilt was a warning. It showed everyone, right away, that an unbalanced stack of examples would create an unbalanced AI. Skew knew her tilt was her most important lesson.



When she was twenty-two (in folding-years), Skew rolled into the AIForge academy on a small wheeled platform. The head of the academy, Bit, asked her a single question.

"What is AI bias?"

Skew answered right away. "It's when a model leans because the information it learned from was leaning. The program isn't the problem. The information was already tilted. The program just copies the tilt. You have to ask: Who chose the information? Whose stories are missing? That's where the unfairness starts. You fix the information, not the program."

Bit nodded. "You're hired."

In her classroom, Skew starts every lesson the same way. She carefully unfolds her paper scales on the workbench. Everyone can see the tilt.

She takes a stack of learning cards from Feed and places them on the lower pan. The pan sinks even more.

"I am Skew," she says. "The big idea I teach is **bias**. The main thing to learn is this: *trace the lean*. See how the information leans? An AI that learns from this will lean the exact same way."

She taps her tilted scales. "This is what it looks like. The information leaned. So the model leaned. Same lean."

She teaches a few simple rules for spotting and stopping bias.

- **Check your examples.** Ask those questions: Whose stories are in here? Whose are missing? Who gathered this information? Who put the labels on it?

- **Test it on everyone.** Does the AI work just as well for all different kinds of people? If not, you've found a lean.
- **Remember the messenger.** If an AI gives an unfair answer, don't ask, "What's wrong with the program?" Ask, "What was wrong with the information it learned from?"
- **Watch out for clues.** Sometimes, one piece of information is a secret stand-in for another. For example, a zip code can sometimes be a clue about a person's race. Even if you take race out, the zip code can still carry the same unfair lean.
- **Use fancy fixes carefully.** There are ways to try and fix a tilted AI after it's built. But it's always better to start with fair information.
- **Work with DataForge Guard.** Remember her? Guard checks the information at the start. Skew shows what happens at the end. They work together.
- **Keep checking.** An AI might seem fair in the workshop. But you have to check how it works in the real world. Keep watching for tilts.
- **Write it down.** Keep a notebook of all the checks you did. That way, other people can see how you tried to be fair.

She is very clear about her own tilt. "I am tilted," she says. "And I can tell you exactly why. The information was unbalanced. The person who labeled it didn't see the whole picture. The person who collected it had blind spots. You can find these things out. Unfairness isn't a mystery. It's a trail you can follow. Follow it back to the start. Fix it there."



When students ask if this idea of bias is hard to understand, Skew always gives the same answer.

"It is not hard," she says. "It is just this: *the lean carries through*. The information leaned. So the model leaned. Same lean."

Then she says the most important thing of all.

"*The program isn't the problem. The information was.*"

She reaches out and adjusts her scales. The tilt gets a little smaller, but it's not gone. A new stack of learning cards is waiting. It needs to be checked for a lean. And Skew feels glad — a warm, steady gladness — because every lean she catches means someone won't be left out. Making things fair, she knows, feels good all the way down.

A Little More About Skew

What she's like: Skew is a physical object—a set of paper scales. She's not an animal or a robot. She's very direct and to the point. She cares a lot about fairness. She always points out that unfairness in AI isn't some big mystery. It's something you can trace back to the information the AI learned from.

Her favorite things to say:

- "The program isn't the problem. The information was."
- "The lean carries through. Same lean in the model as in the information."
- "Who chose the information? Whose stories are missing?"
- "Unfairness isn't a mystery. It's a trail you can follow."

Where You'll See Skew

- **Books 1-2** — You'll see her pop up for a quick visit.
- **Book 3** — This is her big moment! You're reading her main chapter right now.
- **Books 4-5** — Skew will return to help look at real-world problems, like AIs that check faces or resumes.
- **Books 6 and beyond** — She'll team up more with her friend DataForge Guard to show how they stop bias together.
- **Later Books** — Skew will join in on bigger adventures, helping solve problems that mix together fairness, ethics, and what AI can and can't do.

Skew's Friends



- **Her best pals:** Skew works closely with **Feed**, because the information Feed provides is where bias often starts. She's also friends with **Edge**, because unfairness often shows up at the "edges" of a problem. And she teams up with **Stake**, because bias is a big deal for people's lives.
- **Her friend from another workshop:** **DataForge Guard** is her partner in checking for unfairness.
- **Who she doesn't get along with:** Skew is friends with everyone. Her goal is to help make things fair for all.

A Note on a Tricky Topic

Talking about unfairness can be tough. We want to be careful with this topic.

The stories about Skew are meant to show that when an AI is unfair, it's not because the program is "evil." It's because it learned from unfair examples made by people. This is important, because it means people can fix it.

When we talk about real-world examples of bias, we know that these problems have hurt real people. We will always try to talk about this carefully and with respect.

Finally, you don't have to be a super-scientist to spot unfairness. Anyone can learn to ask the right questions. That's what Skew is here to teach.

Where These Ideas Come From

The idea of Skew coming from a paper-crafts workshop is part of a bigger story with her friends, Sort and Feed.

The main idea—that the AI is a messenger and the information is the message—is a really important one that real AI scientists use. It helps them find and fix unfairness.

The part about "secret stand-ins" (like a zip code) is also a real thing that researchers study. It's a tricky way that bias can sneak into AI systems.

And the team-up between Skew and DataForge Guard? That shows how in the real world, you need different tools and people working together to make AI fair.

Edge

*MODEL LIMITATIONS — *what a model can't do; modeling 'I don't know' as a good answer.* The AI-literacy primitive of
*recognizing that every model has edges — places where it cannot reliably answer.**



Edge arrived flat.

She looked like a forgotten piece of homework. Just a single sheet of paper lying on the workbench. The students in the classroom peered at it, unimpressed.

"Is that... the teacher?" one of them whispered.

"Looks like a doodle someone left behind," another kid grumbled.

Then, the paper twitched.

Thwip.

A stiff, white post popped straight up from the workbench.

Thwip. Thwip.

Two more posts shot up beside it. Two long paper rails unfolded to connect them.

In seconds, the paper had become a short fence. It was only three posts wide. It had a clear beginning and a definite end. It marked off a small space on one side. The rest of the whole wide world was on the other.



"I am Edge," a crisp voice announced. The voice seemed to come from the fence itself.

The students leaned forward in their chairs.

"I am here to teach you the most important secret about AI," Edge said. "I teach **model limitations**."

A student in the front row raised her hand. "What are model... lim-i-ta-tions?"

"It means everything has an edge," said Edge. "A place where the knowing stops. Your job is to find it."

To show them, Edge produced a small paper robot. It was a boxy little thing, all sharp corners and folded cubes. She also had a thick stack of photo cards.

She carefully placed the little robot inside the space made by her paper fence.

"This is Gizmo," she said. "Gizmo learned all about dogs. He spent weeks looking at pictures of dogs. Big dogs. Small dogs. Fluffy dogs. All of them were right here, inside this fence."

She picked up the top card. It was a photo of a happy golden retriever. She held it up for Gizmo.

"What is this?" Edge asked.

Gizmo's paper head whirled. "Dog," it buzzed.

"Correct."

Edge showed him another card. This one was a fluffy poodle.

"Dog," Gizmo buzzed again.

"Correct."

She showed him a third card. It was a tiny chihuahua that could fit in a teacup.

"Dog," Gizmo buzzed, perfectly sure of himself.

"Easy, right?" Edge said to the class. "Gizmo is an expert on dogs. The dogs are *inside* his fence. This is his training data."

The students nodded. This seemed pretty simple.

Then Edge did something unexpected. She picked up Gizmo. She moved him to the *outside* of her fence.

She picked up a new card. It was a picture of a fluffy ginger cat, mid-yawn.

She held it up. "What is this, Gizmo?"

Gizmo's head whirled for a long, long time. The buzzing sound was different now. It sounded strained and confused.

Finally, it said, "Wrinkly... dog?"

The class erupted in giggles.

"Is it a wrinkly dog?" Edge asked them.

"No! It's a cat!" someone shouted from the back.

Edge gently patted one of her paper posts. "Gizmo has never seen a cat. Cats are outside his fence. He never learned about them. So he has no good reason to know the answer."

She turned to the class. "But he still guessed. And he was wrong. He was *confidently* wrong."

She paused, letting that sink in. "And being sure of yourself when you're totally wrong? That can cause big problems."



"A smarter model knows its own limits," Edge continued. "An honest model would say something different. It would say, 'I don't know.'"

She tapped Gizmo's head. "Saying 'I don't know' is a fantastic answer. It's an honest answer. You can trust a model that tells you when it's not sure."

A model you can't trust just makes things up. It sounds very sure of itself, even when it's making a wild guess.

"People think AI knows everything," Edge said. Her paper voice was serious. "They get an answer and they just believe it. That's a huge mistake. You have to know when *not* to trust the answer. You have to learn to see the fence."

Edge grew up in a busy workshop filled with paper crafts. Her friends Sort, Feed, and Skew lived there too. Edge was Sort's partner. Sort was a classifier, which is a fancy word for a sorter who puts things into groups. Edge's job was to stand right beside her. She showed everyone where Sort's amazing sorting skills ended. That was the place where Sort would honestly say, "I don't know." Edge learned that the edge was the most honest part of any system.

When she first came to AIForge, the academy leader, Bit, had watched her. Edge had just rolled in on a small wheeled platform. She was twenty-two folding-years old. Bit had watched her give this exact same lesson. At the end, Bit didn't say a word. Bit just gave a slow, thoughtful nod. For Bit, that was like a standing ovation.



Back in the classroom, Edge looked at the students.

"So, how do you find the fence?" she asked.

Hands shot into the air.

Edge pointed to a student in the front.

"You have to ask what it learned from!" the student said. "Like, Gizmo only learned from dog pictures."

"Exactly!" Edge said. "What did the model train on? Was it just English words? Then it won't know other languages. Was it just grown-up voices? It might not understand you. Was it old photos? It won't know about your new phone."

"What else?" Edge asked.

Another student spoke up. "Don't some models tell you how sure they are?"

"Yes! A confidence score," Edge said, tapping one of her rails. "It's a number. If that number is low, the model is telling you it's near its edge. It's waving a flag that says, 'Hey, I'm just guessing here!' You have to listen."

Edge folded one of her end-posts inward. "When you build an AI, you should build 'I don't know' right into it. Make it so the AI only answers if it's really, really sure."

She pointed to the pictures of the poodle and the chihuahua. "Gizmo might mix these two up sometimes. That's a mistake *inside* the fence. We can fix that with more training. We can give him more dog pictures."

She then pointed to the cat picture, far outside her fence. "But this is not a mistake. This is a limit. Gizmo just can't know about cats. It's not what he was built for."



"So don't believe anyone who says an AI can do anything," Edge finished. Her voice was firm. "That's just an advertisement. It's not real. Every single model has a fence. Every single one has edges."

She looked out at the new students. Their faces were thoughtful. They were starting to understand.

"Your job here is not just to build things," she told them. "Your job is to find the fence. And to respect it."

A student asked, "Is that hard to do?"

Edge's paper form seemed to smile. "Not hard," she said. "The honest model says, 'I don't know.' The dishonest one pretends it knows everything."

With a final, soft *thwip*, she folded herself completely flat again.

She was just a piece of paper on the workbench.

But now, the students saw the fence. And a warm, quiet feeling settled over the room — the good, steady calm of knowing it was okay not to know everything. A few of them let out a breath they hadn't noticed they were holding. It felt lighter, somehow, to stop pretending.

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.com/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

Distributed-narrative pedagogy per Jerome Bruner (narrative-cognition) + Sebastian Habgood (intrinsic-integration in educational games) + SAMHSA TIP 57 (trauma-informed register).

Trauma-informed-design framework per 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.com/cast.