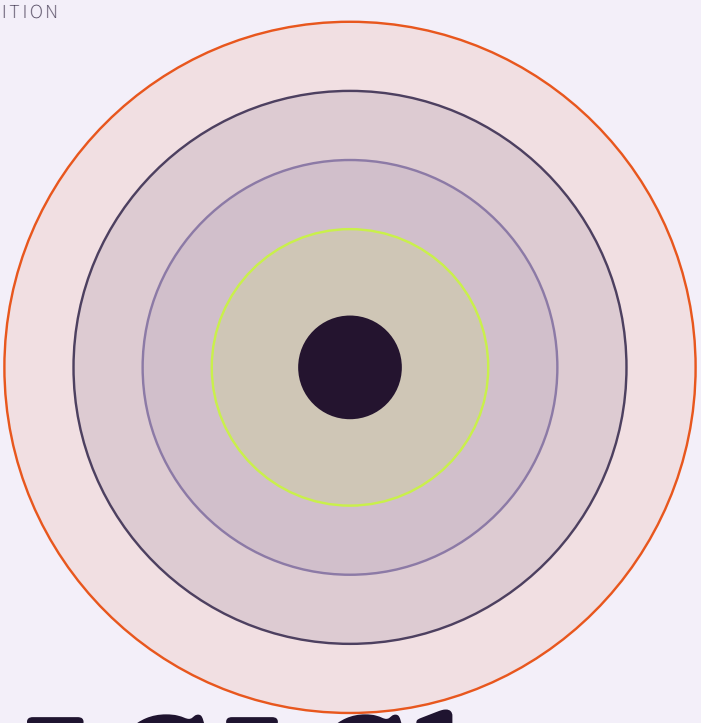


The AGI Chronicles: The Inside Story of the Race to Create an Artificial Superintelligence

A reading of the book by Kevin Roose (2026)

1. ACTION BOOK

CONTENTS

Inside this edition

The book in a minute 4

The takeaways 5

Your action book 7

Before you go 8

PART TWO · THE ACTION BOOK

The setup checklist 10

Keep a monthly AI log 11

Retake your stance every quarter 12

Ask who gains before you share 13

ABOUT THIS EDITION

This is HourLife's reading of The AGI Chronicles: The Inside Story of the Race to Create an Artificial Superintelligence, in our own words. It is not a substitute for the book.

BEFORE YOU START 01

The book in a minute

Kevin Roose is a technology columnist at The New York Times and the author of Futureproof. The AGI Chronicles is his reported history of roughly the last decade of artificial intelligence, told through the companies racing to build artificial general intelligence: OpenAI, Anthropic and Google DeepMind. It draws on more than 150 interviews with people inside those labs, along with internal documents and research memos. Roose follows the labs where language models gained new abilities quickly, the split that led Anthropic's founders to leave OpenAI, Sam Altman's removal and return as OpenAI's chief executive, Google's push to catch up, and the officials trying to regulate a technology that moves faster than they do. Along the way he explains in plain terms how these models are built and trained, and he argues that the stakes are high enough that there may be little time left to make them safe.

Below are the 9 ideas HourLife readers vote up most, each in our words with a note on what it asks of you, and then the steps that put them to work.

WHAT READERS KEEP

The takeaways

1

HOURLIFE TAKEAWAY

Many of the people racing to build AGI believed it could be dangerous, and they kept building it.

This is the puzzle at the centre of the book. Some builders hoped the technology would cure diseases, some feared a catastrophe, and plenty held both views at once. Roose's reporting tries to explain how people who took the risk seriously still ended up speeding the race along.

2

HOURLIFE TAKEAWAY

A model trained only to guess the next word turned out to learn far more than anyone planned.

The book walks through how the field moved from narrow systems built for single jobs, such as translation, to models trained to predict the next word across huge amounts of text. The bet was that doing this well would force a model to pick up a working picture of the world. Understanding that bet makes the rest of the story easier to follow.

3

HOURLIFE TAKEAWAY

Disagreements about safety did not slow the race. They created new competitors.

Roose covers the split that led Anthropic's founders to leave OpenAI. A group that worried about how fast things were moving started its own lab, which then had to compete at the frontier to have a say. It is a clear case of how caution can still add speed.

4

HOURLIFE TAKEAWAY

The checks on a frontier lab can be tested, and changed, within days.

The book covers Sam Altman's removal as OpenAI's chief executive and his return shortly after. Whatever you make of the people involved, the episode showed how much of the field's governance rested on a few boards and a few relationships.

5

HOURLIFE TAKEAWAY

Every big release puts pressure on the rival labs to answer.

Roose follows Google's effort to catch up once others had shown what large language models could do. Seen together, the labs behave less like separate projects and more like one race, where each move changes what the others feel they must do next.

6

HOURLIFE TAKEAWAY

Governments are trying to regulate a technology that changes faster than they can meet.

The book includes the officials who tried to get a handle on the field. Their problem is timing: rules take years, and the systems they describe can change within months. Keep that gap in mind when you read about any new law or agreement.

7

HOURLIFE TAKEAWAY

The optimists and the worriers mostly agree that something very large is coming. They disagree about what to do.

Roose spends time in the San Francisco AI scene, where people who want to speed up and people who fear the outcome live side by side. Both camps take the technology seriously. Their argument is about pace, control and who decides.

8

HOURLIFE TAKEAWAY

Roose argues that humans may be losing their monopoly on intelligence, and that the time to make these systems safe may be short.

This is the author's conclusion, and readers will weigh it differently. You do not have to share it to find the book useful. It is a strong case for following the question closely rather than treating it as science fiction or as marketing.

9

HOURLIFE TAKEAWAY

A careful record of the last decade is the best tool for judging the next claim.

Much of what people say about AI is either a pitch or a warning. A dated account of what the labs said, built and did gives you a baseline. When the next big announcement arrives, you can ask how it compares with what happened before.

PUT IT TO WORK

Your action book

Do the setup steps once, then keep the habits going. The setup changes your week; the habits keep it changed after the first good one.

The download puts this whole reading and a printable action book in one PDF: the setup checklist, then a page of real calendar days for each habit, starting on the day you choose.

SET UP ONCE

- [ ] Write down your AI forecast and date it
- [ ] Map ten of your weekly tasks
- [ ] Test a model on something you know well
- [ ] Learn how a model is trained, in three sentences
- [ ] Balance your sources

KEEP UP DAILY

- ◯ Keep a monthly AI log
- ◯ Retake your stance every quarter
- ◯ Ask who gains before you share

30 days · Nov 6 - Dec 5, 2026

Part two, after the reading, prints each of these as a page to fill in by hand.

THE LAST WORD

Before you go

Pick one step from the action book and do it today. An idea from a book starts to count once it changes what one ordinary day looks like.

PART TWO

The action book

Change the setup once, then keep it up. 5 setup actions to do once, and 3 habits to mark each day for 30 days.

Nov 6 - Dec 5, 2026

Print it, or write on it in any PDF app.

SET UP ONCE

The setup checklist

Each of these is done once. Do one a day, and check it off when it is done.

- [ ] SETUP 1

Write down your AI forecast and date it

In two sentences, write what you think AI will be able to do in five years, and how sure you are.

Date it and keep it where you will find it again. It is the first entry in your own chronicle.

MY VERSION

DONE ON

- [ ] SETUP 2

Map ten of your weekly tasks

List ten things you do at work or at home each week.

Mark each one: a current AI model does it well, does it badly, or cannot do it. Keep the list to compare later.

MY VERSION

DONE ON

- [ ] SETUP 3

Test a model on something you know well

Give a current AI model a real task from your own field, one where you can judge the answer.

Note where it is good and where it is confidently wrong.

MY VERSION

DONE ON

- [ ] SETUP 4

Learn how a model is trained, in three sentences

Read one plain explanation of how language models are trained, such as the technical passages in the book.

Then explain next-word prediction to a friend in three sentences without notes.

MY VERSION

DONE ON

- [ ] SETUP 5

Balance your sources

Pick three people to follow on AI: someone who builds it, someone who works on its safety, and someone who doubts the hype.

Drop one source that only ever cheers or only ever panics.

MY VERSION

DONE ON

DAILY HABIT 1 OF 3

Keep a monthly AI log

Once a month, write three lines: what changed, what you tried, and whether it changed your view.

Nov 6 - Dec 5, 2026 · 30 days · circle each day you do it

SMALLEST VERSION

CUE

WHERE

November 2026

MON	TUE	WED	THU	FRI	SAT	SUN
						1
2	3	4	5	6 ○	7 ○	8 ○
9 ○	10 ○	11 ○	12 ○	13 ○	14 ○	15 ○
16 ○	17 ○	18 ○	19 ○	20 ○	21 ○	22 ○
23 ○	24 ○	25 ○	26 ○	27 ○	28 ○	29 ○
30 ○						

December 2026

MON	TUE	WED	THU	FRI	SAT	SUN
	1 ○	2 ○	3 ○	4 ○	5 ○	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30	31			

EACH FRIDAY

WHAT HELPED

WHAT GOT IN THE WAY

ONE CHANGE TO THE SETUP

DAILY HABIT 2 OF 3

Retake your stance every quarter

Every three months, run the stance map on this page again and save the reading.

Nov 6 - Dec 5, 2026 · 30 days · circle each day you do it

SMALLEST VERSION

CUE

WHERE

November 2026

MON	TUE	WED	THU	FRI	SAT	SUN
						1
2	3	4	5	6 ○	7 ○	8 ○
9 ○	10 ○	11 ○	12 ○	13 ○	14 ○	15 ○
16 ○	17 ○	18 ○	19 ○	20 ○	21 ○	22 ○
23 ○	24 ○	25 ○	26 ○	27 ○	28 ○	29 ○
30 ○						

December 2026

MON	TUE	WED	THU	FRI	SAT	SUN
	1 ○	2 ○	3 ○	4 ○	5 ○	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30	31			

EACH FRIDAY

WHAT HELPED

WHAT GOT IN THE WAY

ONE CHANGE TO THE SETUP

DAILY HABIT 3 OF 3

Ask who gains before you share

Before you pass on a big AI claim, ask who is making it and what they gain if you believe it.

Nov 6 - Dec 5, 2026 · 30 days · circle each day you do it

SMALLEST VERSION

CUE

WHERE

November 2026

MON	TUE	WED	THU	FRI	SAT	SUN
						1
2	3	4	5	6 ○	7 ○	8 ○
9 ○	10 ○	11 ○	12 ○	13 ○	14 ○	15 ○
16 ○	17 ○	18 ○	19 ○	20 ○	21 ○	22 ○
23 ○	24 ○	25 ○	26 ○	27 ○	28 ○	29 ○
30 ○						

December 2026

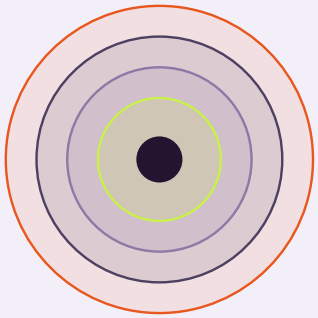
MON	TUE	WED	THU	FRI	SAT	SUN
	1 ○	2 ○	3 ○	4 ○	5 ○	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30	31			

EACH FRIDAY

WHAT HELPED

WHAT GOT IN THE WAY

ONE CHANGE TO THE SETUP



Many of the people building it feared it, and built it anyway. That is the story, and we are all in it.

HOURLIFE NOTE AFTER KEVIN ROOSE

Read it online, and vote on the takeaways:
hourlife.com/books/the-agi-chronicles
Made November 6, 2026 · HourLife, hourlife.com