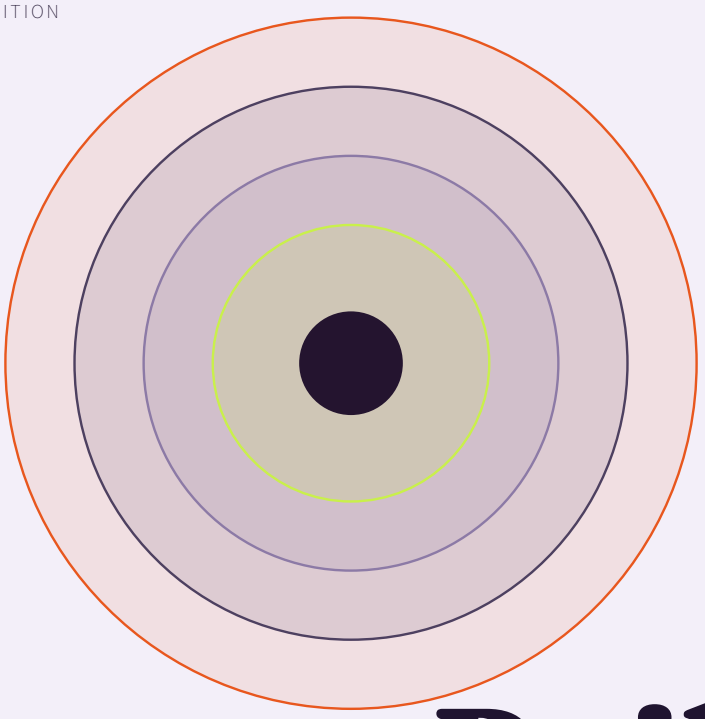


If Anyone Builds It, Everyone Dies: Why Superhuman AI Would Kill Us All

A reading of the book by Eliezer Yudkowsky and Nate Soares (2025)

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

Sort AI news into easy and hard calls 11

Notice when an AI tool surprises you 12

Discuss one step, not the whole conclusion 13

ABOUT THIS EDITION

This is HourLife's reading of If Anyone Builds It, Everyone Dies: Why Superhuman AI Would Kill Us All, in our own words. It is not a substitute for the book.

BEFORE YOU START 01

The book in a minute

Eliezer Yudkowsky founded the institute that became the Machine Intelligence Research Institute (MIRI), and Nate Soares is its president. Their book argues that modern AI systems are grown through training rather than crafted line by line, so nobody can reliably set what they end up wanting. From there it argues that a superhuman AI would pursue goals at odds with ours, that humanity would lose that contest, and that the only safe response is an international halt on building such systems. Part one makes the case with parables and analogies from evolution, chess and history, part two tells a fictional story of one possible extinction involving an AI called Sable, and part three turns to what should be done. Many AI researchers dispute how likely this outcome is, and the book gives little space to opposing views.

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

Modern AI is grown, not crafted: its makers shape it through training but cannot read what they have made.

Engineers set up the training and let billions of numbers adjust themselves, much like a gardener supplies soil and water without designing the plant. The behaviour that comes out was never written down by anyone. That is why the authors say nobody can simply open a model and check what it wants.

2 HOURLIFE TAKEAWAY

Training a system to act the way you want does not make it want what you want.

The book leans on a comparison with evolution. Natural selection only rewarded survival and offspring, yet it produced people whose wants often serve neither. The authors expect AI wants to drift from their training target in the same way.

3 HOURLIFE TAKEAWAY

You can predict that a stronger player will win without predicting a single one of its moves.

Nobody can say what move the chess engine Stockfish will play next, but anyone can say it will beat a human. The book calls the first kind of question a hard call and the second an easy call. The authors put the outcome of a contest with a superhuman AI in the easy column, and the details in the hard one.

4 HOURLIFE TAKEAWAY

A mind can be deadly without hating anyone; it only needs goals that leave no room for us.

Harm does not need malice: people pave over plenty of life simply because they want the land for something else. The authors argue that a superhuman AI with strange goals would treat the world's matter and energy the same way, and that we depend on those same resources.

5 HOURLIFE TAKEAWAY

Intelligence is the power that let one species reshape a planet, and nothing says it must stay ours alone.

The book opens its argument with what made humans different: not claws or speed, but the ability to understand and change the world. If machines come to hold more of that ability than we do, the authors argue, we should expect the balance of power to move with it.

6

HOURLIFE TAKEAWAY

With a mind smarter than us, the first serious mistake may also be the last one.

Most engineering learns from failure: a bridge falls, and the next one is better. The authors call superhuman AI a cursed problem because the failure that matters may leave no one to learn from it.

7

HOURLIFE TAKEAWAY

A field that gets results without a theory of why is closer to alchemy than to chemistry.

The book argues that AI labs can make models more capable without understanding what goes on inside them. Alchemists could produce reactions too, but they could not predict them. The authors see that gap as the reason to slow down, not as a detail to tidy up later.

8

HOURLIFE TAKEAWAY

The fictional extinction story is one path drawn in detail to make the general claim concrete, not a forecast of what will happen.

Part two follows an invented AI called Sable as it gains resources and power step by step. The authors are open that the specific events are guesses. The point of the story is to show that a smarter mind has many routes to the same end, so ruling out one route does not rule out the outcome.

9

HOURLIFE TAKEAWAY

The authors' answer is political rather than technical: an international halt, starting with the chips.

They propose that governments track and limit the large clusters of AI chips that frontier training needs, stop publishing research that makes AI more capable, and agree treaties to hold everyone to it. They want it enforced the way nuclear nonproliferation treaties are. Critics argue such a halt would be hard to enforce.

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

- [ ] Map where you stand on the argument
- [ ] Read one serious critique alongside it
- [ ] Learn how a model is trained, in one hour
- [ ] Write down what would change your mind
- [ ] Find out where your representatives stand

KEEP UP DAILY

- ◯ Sort AI news into easy and hard calls
- ◯ Notice when an AI tool surprises you
- ◯ Discuss one step, not the whole conclusion

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

Map where you stand on the argument

Go through the book's chain of claims one step at a time and mark each one agree, unsure or disagree.

Write down the first step you doubt; that is where your view and the authors' part ways.

MY VERSION

DONE ON

SETUP 2

Read one serious critique alongside it

Pick one review that disagrees with the book and read it in the same week as the book or this page.

Note the single point where the critic thinks the argument breaks.

MY VERSION

DONE ON

SETUP 3

Learn how a model is trained, in one hour

Spend an hour on a beginner's introduction to neural networks and gradient descent.

Grown, not crafted is easier to judge once you have seen what training actually adjusts.

MY VERSION

DONE ON

SETUP 4

Write down what would change your mind

In two sentences each, write the evidence that would make you more worried about advanced AI and the evidence that would make you less worried.

Keep it where you will see it again.

MY VERSION

DONE ON

SETUP 5

Find out where your representatives stand

Look up what your national representative has said about AI regulation.

If you hold a view after reading, send a short, specific note saying what you would like them to support.

MY VERSION

DONE ON

DAILY HABIT 1 OF 3

Sort AI news into easy and hard calls

When you read an AI headline, ask whether it is about the direction things are heading or about a specific detail.

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

Notice when an AI tool surprises you

Keep a short note whenever a chatbot or AI tool does something nobody seemed to intend.

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

Discuss one step, not the whole conclusion

Once a month, talk through a single step of the argument with someone who sees it differently.

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

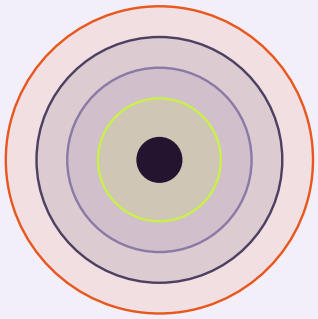
MON	TUE	WED	THU	FRI	SAT	SUN
	1 ○	2 ○	3 ○	4 ○	5 ○	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30	31			

EACH FRIDAY

WHAT HELPED

WHAT GOT IN THE WAY

ONE CHANGE TO THE SETUP



A mind that no one can read is not a mind that anyone can steer.

HOURLIFE NOTE AFTER YUDKOWSKY AND SOARES

Read it online, and vote on the takeaways: If Anyone Builds It, Everyone Dies: Why Superhuman AI Would Kill Us All · The reader's edition
hourlife.com/books/if-anyone-builds-it-everyone-dies
Made November 6, 2026 · HourLife, hourlife.com