

Why things get better

EFFICIENCY + OPTIMIZATION • 1881 TO 2026

For 140 years, every company has been playing the same game: find the slowest part, fix it — then discover the slowness has quietly moved somewhere new. This is that story, in plain words, and where the slow part sits right now.

What I did: read 67 original sources — company engineering write-ups, research papers, government energy data — and put 140 years of it in one place, with every number traced back to whoever actually measured it.

12.5 hrs to 93 min

Time to build a Ford car, before and after 1913

9x - 900x

Yearly fall in the price of the same AI ability, by task

1.09 vs 1.54

Power used per unit of computing: Google vs the industry

5+ years

The wait to connect a new power plant to the US grid

1 in 20

New medicines that make it. So failing early is the win.

THE STORY, IN FOUR STEPS

1 — It starts with a car

In 1913 a Ford took twelve and a half hours to build. Then Ford moved the car past the workers, instead of the workers to the car. One finished car every 93 minutes. Nobody worked harder. The waiting was removed.

2 — Then the problem moved

A factory that fast needed parts that fast. Now the suppliers were the slow part. So Ford bought mines, forests and a railway — not to build an empire, but to chase the new slow part.

3 — And it kept moving

Toyota found it was piles of stock nobody used. Others found it was guessing instead of measuring. Then computers, then software, then memory. Every generation thought they had finished. Every one was wrong.

4 — Today it is electricity

A better chip takes about two years to design. A new power plant takes more than five years to connect to the grid. The cleverest computer ever built cannot do a thing about that.

THE ONE IDEA

Fix the slowest part of anything, and the slowness does not disappear. It moves.

So you go and find it again. That is the whole job — and people are paid very well to do it.

IS THIS FOR YOU? (ESPECIALLY AT 40)

You do not need to be a programmer. The valuable person is the one who can look at a whole system and say: the problem is not here, it is over there. That takes years of watching real work go wrong. Anyone can learn to code. Far fewer can see a system. At 40 that is your advantage, not your obstacle.

US median pay — Industrial Engineer \$102,440, Operations Research \$88,940. Both growing 12%. The AI end of the same skill pays well above it.

80 SECONDS: WHAT JUST HAPPENED

- OpenAI let its own AI rewrite the code that runs it. Result: 20% cheaper to run. Their own figure, July 2026.
- NVIDIA stopped selling chips and started selling whole racks — claiming 10x more work per unit of electricity. That is their claim, not an independent test.
- Everyone changed the question. Not "how fast?" any more. "How much work per watt?"
- AI is running out of human writing to learn from, somewhere between now and 2032. So companies now make their own practice material — carefully. Done carelessly, the AI quietly gets worse.
- And the real limit left the building. It is now the queue for electricity.

IF YOU WANT TO DO THIS FOR A LIVING — THE SHORTEST ROUTE, ABOUT SIX MONTHS

WEEKS 1-2

Learn one idea

Read The Goal. It is a novel, not a textbook. One idea: the slowest step decides everything.

MONTH 1

Learn to measure

Before and after. If you cannot measure it, you cannot claim it — and nobody will believe you.

MONTH 2

Fix one real thing

At the job you already have. Something small. Write the number down.

MONTH 3

Learn the new words

Tokens, GPU, cache, latency, watt. That is a week of reading, not a degree.

MONTHS 4-5

Do it again, in public

One page: what was slow, what you changed, what it saved. Put your name on it.

MONTH 6

Aim at the overlap

Where your old industry meets the new one. Hospitals, factories, logistics, energy all need this now.

Every generation thought their slow part was the last one. None were right — and that is the good news: there is always another, and somebody has to find it.

What to say

90 SECONDS • 225 WORDS • READ IT, DO NOT MEMORISE IT

PAGE ONE IS WHAT THEY SEE. THIS PAGE IS FOR YOU.

0:00

HOOK — slowly, and look up at the end

In 1913, building a Ford car took twelve and a half hours. Then Ford changed one thing. He moved the car past the workers, instead of the workers to the car. Ninety-three minutes. Nobody worked harder.

0:15

THE TWIST — this is the whole talk

But here is what almost nobody notices. The slowness did not go away. It moved. Now the suppliers were too slow — so Ford bought mines, forests and a railway. Toyota found it hiding in the stockpiles. Then it was computers. Then software. Then memory.

0:40

TODAY — point at the screen

Today it is electricity. A better chip takes about two years to design. A new power plant takes more than five years just to get connected to the grid. The cleverest computer ever built cannot do a thing about that.

0:55

TURN TO THEM — slow right down here

So — for anyone sitting there thinking they are too late. Every single one of those fixes needed a person who could look at a whole system and say: the problem is not here, it is over there. You cannot learn that in a twelve-week course. You learn it by watching real work go wrong for twenty years.

1:15

CLOSE — then stop talking and let it sit

If you are forty, you are not behind. You are holding the half that is rare. Writing the code is the easy half, and it is the half being automated fastest. Finding the slow part is the whole job. And somebody has to do it.

THE THREE QUESTIONS YOU WILL GET — SHORT ANSWERS

“Won’t AI just do this job too?”

It already writes the code. It does not decide which slow part is worth money. OpenAI’s own 20% saving came from people pointing a model at the right problem.

“Do I need to be able to code?”

Enough to read it and to measure it. Not enough to be hired for it. The scarce half is the judgement, and that half is not getting cheaper.

“Am I too old to start?”

The skill is recognising patterns across systems you have personally watched fail. That is the one thing nobody can buy at twenty-two.

IF THEY ASK WHAT TO ACTUALLY LEARN

“Cost per good answer”

Every company now has AI in its product. Almost none of them can tell you two things: what one genuinely useful answer costs them, and whether it was even right.

Learn to measure both. Then learn to make the cost smaller without making the answers worse.

That is measurement and judgement — not maths you missed at school. It is the least crowded, highest-demand skill in AI right now, and twenty years of watching real systems fail is the entry ticket, not a handicap.

WHO IS PAID TO DO EXACTLY THIS

AI Product Manager

Measured on cost per completed task.

Technical Program Manager

AI infrastructure, where nothing reports to you.

Inference / ML performance

Tokens per second per GPU, inside a latency budget.

Capacity & fleet planning

The oldest operations maths, on the newest asset.

Biotech runs the identical play. About 1 in 20 medicines make it, so the prize is not cheaper trials — it is knowing a molecule will fail before Phase 3 rather than during it. Same skill, different units.

MONDAY MORNING

The six steps are on page one. Step one is one book and one number — that is the entire barrier to entry, and it is why most people never start.