

Challenges Program

Thinking Outside the Basin

October 2019 to July 2026

33 challenges, 886 solvers, and $809,000 in prize money

Contents

In brief	2
The program at a glance	3
Seeking answers	4
Finding the form	4
Coming back bigger	5
The busiest years	6
How a challenge was made	8
The community	10
What the challenges delivered	11
What changed	11
Acknowledgements	12
Notes	12

In brief

- From 2019 to 2026, ThinkOnward ran 33 open data science challenges on problems from the energy industry, most of them in the subsurface. Teams joined 6,181 times, 886 people posted a scored entry, and ThinkOnward paid $809,000 in prizes.
- The program's bet was that people from outside the industry could solve problems the industry was stuck on. They did: winners included signal-processing and computer-vision specialists with no geoscience background, and a community model for oil slicks raised detection precision from 11% to 89%.
- By 2026, agents could run much of the optimization loop that challenges had asked people to run. ThinkOnward closed the Challenges program and focused on its Projects offering.

The program at a glance

Challenges

33

From October 2019 to July 2026

Solvers

886

People with a scored entry on a public leaderboard

Prize money

$809,000

Awarded to community members across all challenges

Teams joined

6,181

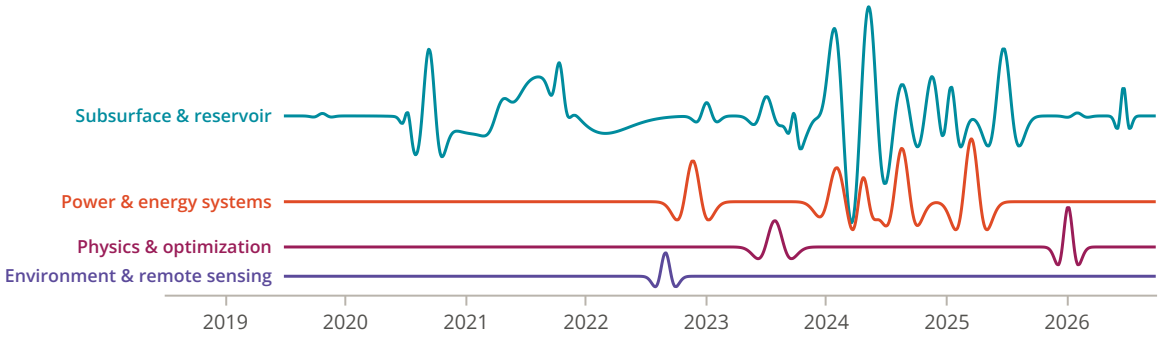
Team sign-ups across all challenges

Twenty-five of 33 challenges were subsurface problems



Each tile is a challenge's cover, colored by its domain.

The program ran in waves, and most of its challenges were subsurface problems



Each wavelet is one challenge, as tall as the number of teams that joined it.

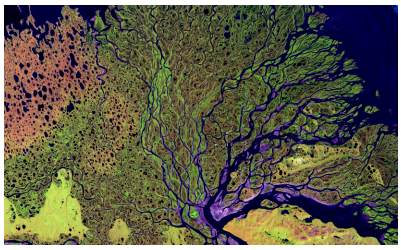
Seeking answers

Around 2019, ThinkOnward was tasked with solving two kinds of problems. Some needed one specific expert, and that work became Projects. The rest already had a data science solution that fell short in practical use, and those became Challenges. Challenges also needed data that could be released openly or replaced with synthetic data. Most produced a machine learning model.

Geoscientists and engineers across the industry had picked up coding and were building useful tools. They knew the subject, but their algorithms often stalled because the code ran too slowly, cost too much, or couldn't scale to a full dataset. Kaggle had shown what a crowd of data scientists could do with a well-posed problem. ThinkOnward paired the industry's subject knowledge with a large crowd. People from other fields brought divergent approaches, and sometimes one of them was the breakthrough an existing solution needed. As Jesse Pisel, ThinkOnward's Director of Data Science Challenges, put it: "The answers lived outside oil and gas, so we opened the doors."

The program began under the name Xeek, from "seek", as in seeking an answer, on a platform wired together by hand.

The first challenges were simple. ThinkOnward wanted to learn whether people would take part, and whether they would write interesting code. **Find Ancient River Channels**, in October 2019, asked teams to find depositional patterns from gamma ray logs. The early challenges also tested how to engage the community, how to value code contributions, and how far synthetic data could go and still produce effective answers.

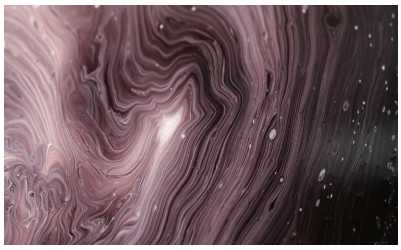


Find Ancient River Channels
The first challenge, October to November 2019

Of the early challenges, **Find Ancient River Channels** proved the most instructive. Participants reverse-engineered the dataset to recover the parameters that generated it, and some submitted code that took days to run on a fairly simple problem. ThinkOnward introduced runtime limits in response, and the issue never came back.

Finding the form

Three challenges later, the program found its form. **Go with the Flow**, in mid-2020, asked teams to model how hydrocarbons move through the subsurface with simple agents instead of full simulations. The problem fit in a sentence, took a different approach from traditional tools, and came with a dataset no one else could offer. Its results were easy to share with other teams, and they helped those teams think differently about their own problems.



Go with the Flow
June to July 2020. 73 teams joined.

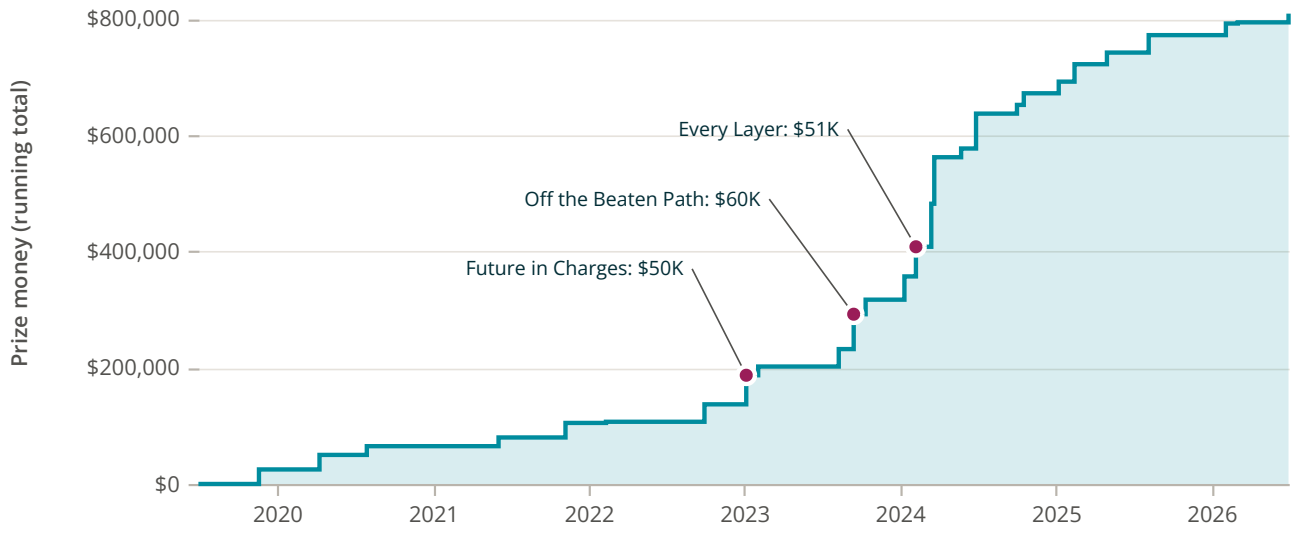
Some entries were fast, while others were more faithful to the geology or closer to reservoir engineering. Experts who ran traditional basin models, which can take a month, were not pleased to see answers that ran in about 20 minutes. ThinkOnward took that reaction as a sign the challenge had hit something real. Two members of the winning team, Bastian Steffens and Quentin Corlay, went on to publish an agent-based tool for screening hydrocarbon migration with colleagues, including Nate Suurmeyer, ThinkOnward's Head of Innovation.[1]

Later in 2020, FORCE, a Norwegian consortium of oil and gas operators, brought two datasets of its own and asked ThinkOnward to run them on its platform: one on predicting lithology from well logs and one on mapping faults in seismic data. The lithology challenge drew 329 teams, and an undergraduate in Nigeria beat the field of industry data scientists. A great reminder that a good idea can still move a hard problem, whoever has it.

Coming back bigger

In 2021, ThinkOnward rebuilt the platform it had first wired together by hand, and no new challenge launched between September 2021 and August 2022. By the time challenges resumed, the team had grown, bringing in data scientists with geoscience backgrounds who could work a problem through with a client. Prize pools rose to match the major platforms, and mlcontests.com, which lists data science competitions across the web, sent a stream of new community members to compete for them.

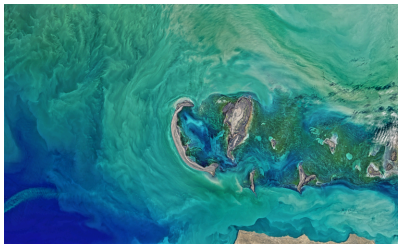
$809,000 awarded to community members since 2019



FORCE 2020 set and paid the prizes for its two challenges, which aren't included.

Slick in a Haystack launched in August 2022. A client team had already built an algorithm to find oil slicks in radar satellite images, done the market research, and proven the need. The team struggled to make the algorithm efficient, so ThinkOnward launched a challenge to help. Participants raised

the algorithm's precision from 11% to 89% and found a way to run the analysis on lower-resolution images. That cut the processing cost from about $60 a scene to under $1. The new method was in production eight weeks after the challenge closed.



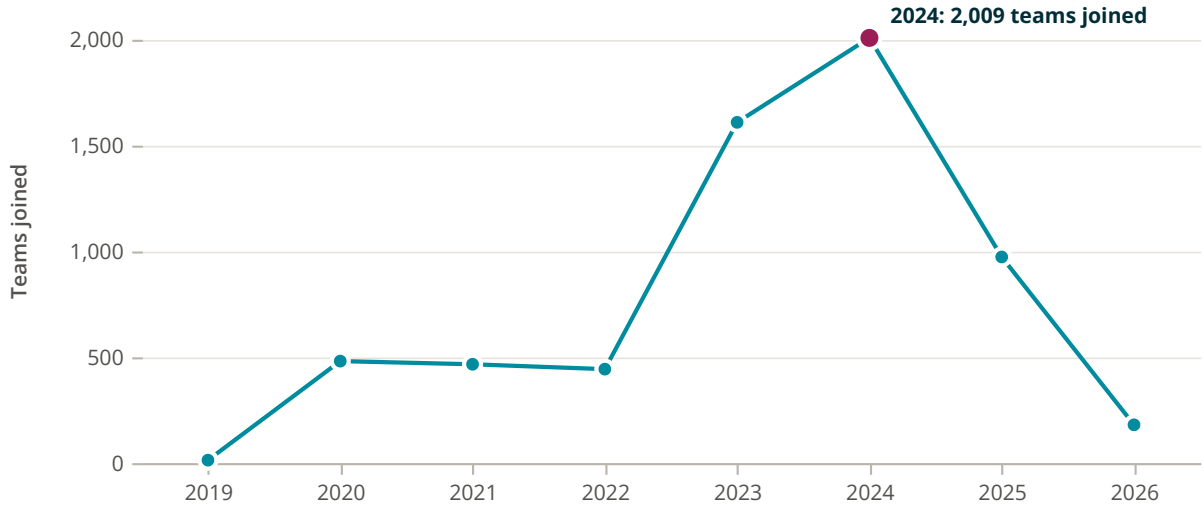
Slick in a Haystack
August to September 2022. 134 teams joined.

Future in Charges, later in 2022, was harder. It asked teams for fast ways to compute the electrical potential in the thin layer of charge that forms where a fluid meets a charged surface, a building block of battery models. The challenge ran in three stages that ended in three dimensions, and only a fraction of the first-stage participants reached the last stage. The smaller final group still produced a wide range of solutions from people with different backgrounds.

The busiest years

ThinkOnward ran 10 challenges in 2023 and seven in 2024. Teams joined 1,610 times in 2023 and 2,009 times in 2024. By late 2023, the program had moved from niche challenges, where most community members knew each other, to larger ones that newcomers could pick up.

Participation grew more than fourfold from 2020 to 2024



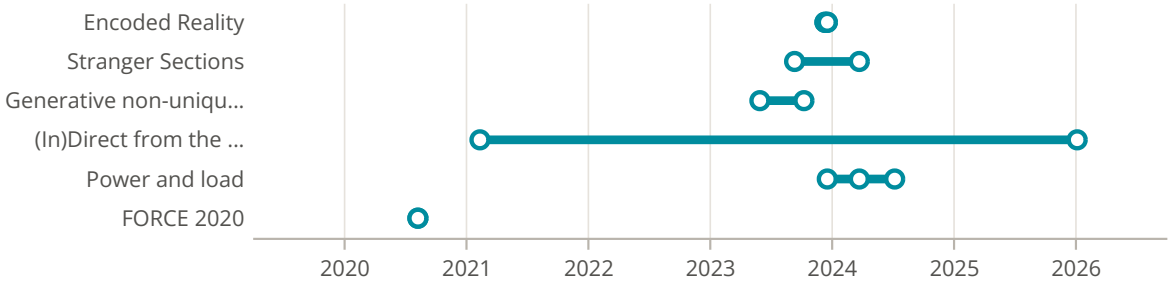
Teams joined per year, counted by the year each challenge started.

Many of the larger challenges ran in series. **Stranger Sections** and the generative pair **From One to Many** and **Two Birds, One Neural Network** built one challenge on the last. **Encoded Reality** posed four seismic problems at once, and one sponsor funded all four to get the industry thinking about foundation models and vision transformers for seismic data.



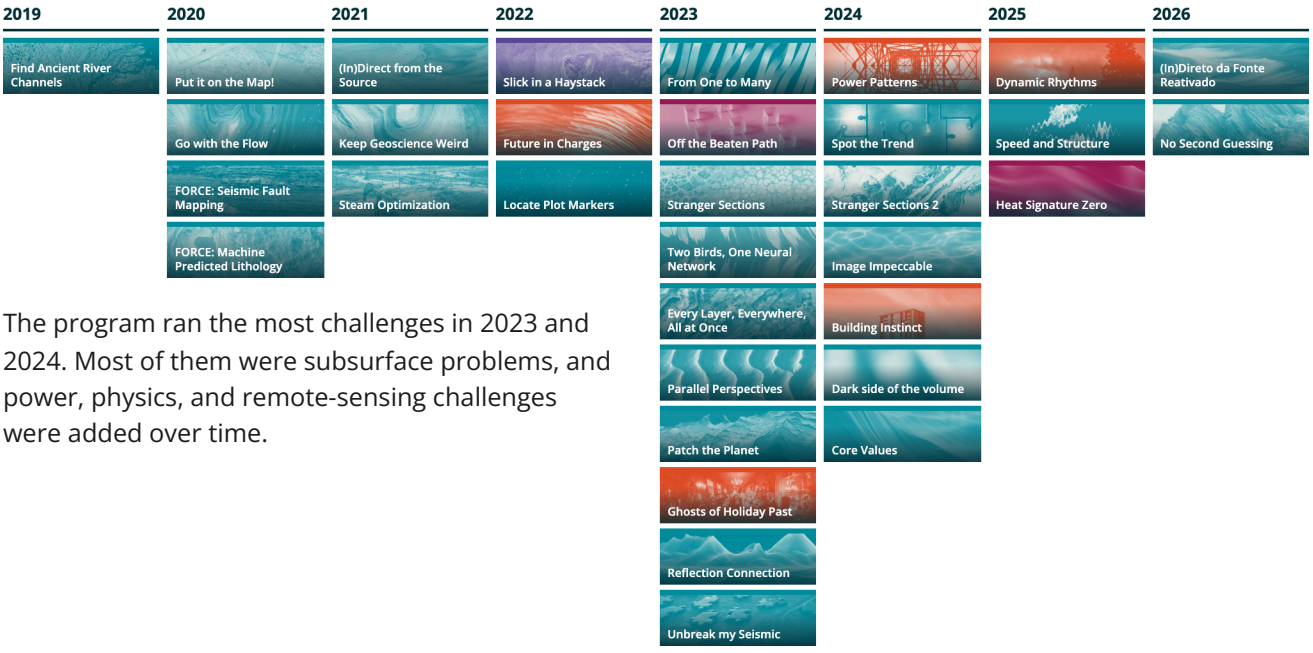
Unbreak my Seismic
One of the four Encoded Reality challenges, December 2023 to March 2024. 139 teams joined.

Fifteen challenges ran as part of a series



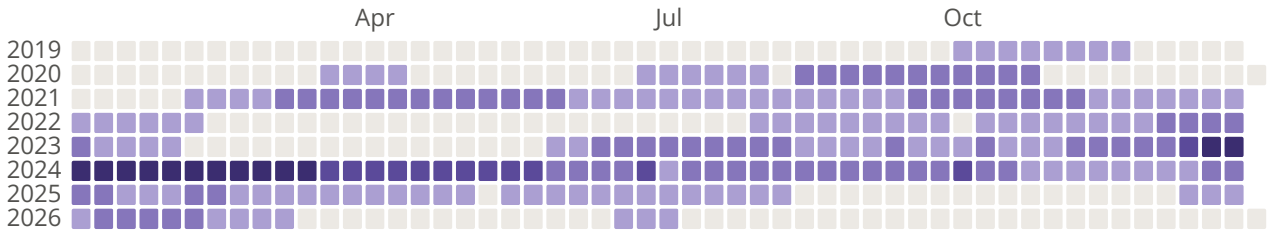
Each stop is a challenge.

ThinkOnward also opened the program to research teams outside the subsurface. **The Energy-Efficient Ghosts of Holiday Past**, in winter 2023, grew out of Ashkan Hosseinloo's focus on power problems, and power, physics, and remote-sensing challenges followed. With so many running at once, the team learned to stagger launches, because several in one week made both the start and the close hectic. In 2024 the team also added launch-event webinars and office hours to meet the community.



The program ran the most challenges in 2023 and 2024. Most of them were subsurface problems, and power, physics, and remote-sensing challenges were added over time.

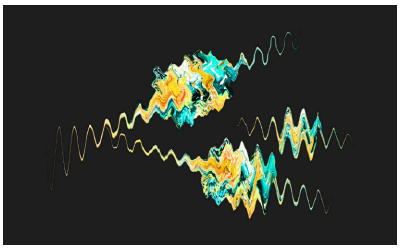
Up to seven challenges ran at once, from December 2023 to March 2024



Each cell is a week; darker weeks had more challenges running.

How a challenge was made

Each challenge started with data. The team used anonymized real data when it had it, synthetic data when the real thing was too sensitive to release, and analogous data when it was too sensitive even to describe. Building a synthetic world as complex as the real one was a project of its own. The **Speed and Structure** dataset needed a seismic forward-modeling pipeline, software that simulates the seismic data a given subsurface would produce, and its scoring used full waveform inversion, which fits a model of the subsurface by matching simulated seismic waves to recorded ones, alongside vision transformers.



Speed and Structure
May to August 2025. 386 teams joined.

The team ran two formats. Structured challenges gave labels and a scoring metric up front and took more work to build. Unstructured challenges started from an idea and launched faster, but the work moved to judging and to understanding what participants had done. Both formats took about the same total hours. In **Dynamic Rhythms**, for example, a week of building was followed by three weeks of judging.

Six challenges with six different profiles



The outer ring marks the top among the 25 challenges with a public score and a prize pool. Repeat solvers are scorers who also scored in another challenge.

The team tested scoring algorithms, hunted for weaknesses before the participants did, drew judges from across the organization, and spent weeks reading final submissions. It also built its own tools. Mike McIntire, with help from Chad Crawford, built a seismic "printer" on Synthoseis[2] that created synthetic seismic volumes in batches, in parallel in the cloud. Later the team built a neurosymbolic scoring system that judged final submissions automatically and wrote a results report for sponsors, and methods to generate synthetic data for new challenges.

Every challenge had a pun for a name. Suurmeyer said the team kept that rule because humor and playfulness matter for creativity.

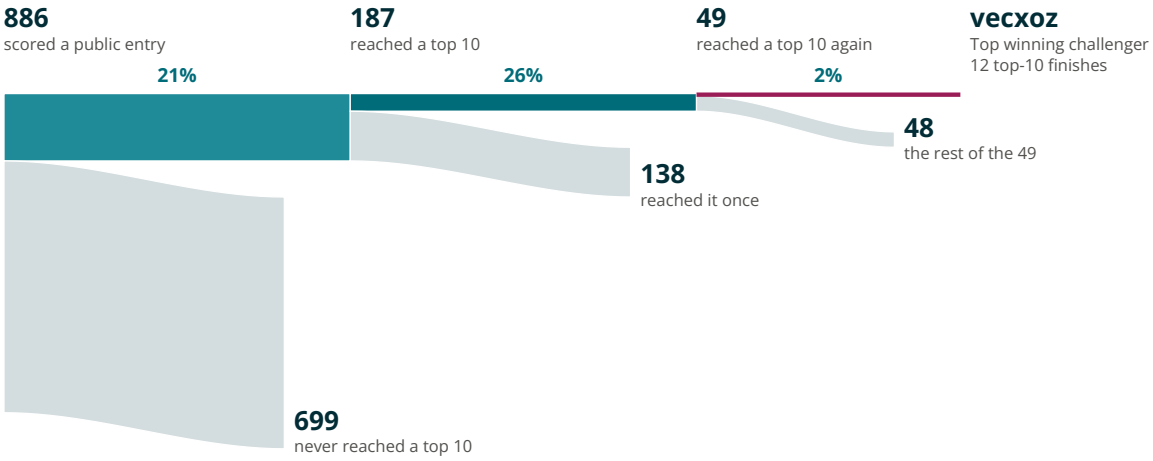


> Launch day was the most fun I've had at ThinkOnward: watching a challenge roll out, knowing the team had built everything from the data science and engineering to the marketing, the glossy photo, and the punny name. The second-best day came about two weeks after it closed, when we started figuring out what people had sent us. I've lost weeks poring over the different ways they went about things. It's amazing how many different ways there are to see the same data. I always ask how we can bring more brains to a problem, and this was the highest throughput we could get: massively parallel intelligence.
>
> **Nate Suurmeyer**, Head of Innovation, ThinkOnward

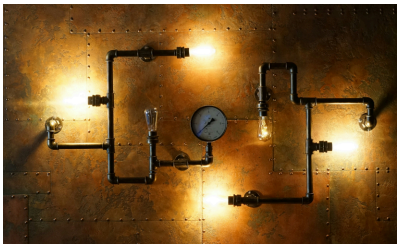
The community

One of a small group of community members who kept coming back, vecxoz entered 20 challenges, finished in the top 10 in 12 of them, and won cash in eight.

Of 886 people who posted a score, 49 reached a top 10 more than once



Counts cover public leaderboards.



Spot the Trend
March to June 2024. 387 teams joined, the most of any challenge.

Several winners came from outside geoscience. Some came from banking and radar engineering, had never studied geoscience, and saw the data as a signal-processing problem. **Every Layer, Everywhere, All at Once** was won by a medical student and a hospital machine learning engineer, who mapped seismic layers the way medicine segments MRI and CT scans. **Spot the Trend**'s winner, a machine learning engineer at a video-streaming platform, saw a well's pressure as a video player's buffer. **Future in Charges** went to a researcher who had modeled space debris. That winner treated the spread of electric charge as a set of orbital paths.

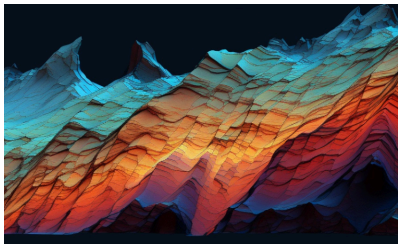
One participant, on sabbatical from a startup, used his own computing power to sharpen his skills and tore through the **Encoded Reality** challenges.



> The community was the best part of the job. I got to watch people grow: students, founders on sabbatical from startups, people who had never touched a well log, all improving with every submission. One winner walked in as a learner and ended up a world expert in data science, and he was not our only Challenger to dramatically accelerate. That's what Challenges got right: it grew the people who found some amazing solutions.
>
> **Jesse Pisel**, Director of Data Science Challenges, ThinkOnward

What the challenges delivered

Slick in a Haystack's method went into production. **Patch the Planet**'s 450 synthetic seismic volumes became pre-training data for ThinkOnward's open Geophysical Foundation Model[3], released with open weights, which later challenges such as **Image Impeccable** and **Dark Side of the Volume** put to work. Some community members crossed to ThinkOnward's other path: past winners joined Projects as experts in residence, and two winners with no geoscience background joined a project on tomography inversion and found efficiencies that experts in the field had not considered for years. Starter notebooks and winning solutions for many of the challenges are public in ThinkOnward's challenges repository on GitHub.[4]



Patch the Planet
December 2023 to March 2024. 226 teams joined.

What changed

By 2026, agentic systems had begun to run the optimization loop that challenges had asked people to run. Andrej Karpathy's autoresearch, released in March 2026, is one example: an agent edits a model's training code, runs a short experiment, keeps the change if the score improves, and repeats overnight.[5] Some of the program's earliest challenges could now be solved in minutes by a coding agent.

The program's last challenge, **No Second Guessing**, in June 2026, accepted agent submissions. The company then turned to Projects, the other of the two paths it had started with. ThinkOnward no longer hosts stand-alone challenges, but it will always look to both AI and the community to help shape the best solutions for its customers.

Acknowledgements

First, thanks to the 886 solvers and thousands of members who took these problems on, and to the many more who came to learn. They showed that people with a fresh angle could solve problems that experts had been stuck on.

Thanks to the customers and sponsors who trusted that community with their hardest problems and their data.

Thanks to the team that built and ran the program (alphabetical): Rodrigo Assar, Mike Campbell, Chad Crawford, Ashkan Hosseinloo, Karol Kawalerczyk, Wenwen Kong, Nataliia Manakova, Mike McIntire, Jakub Mizera, Bao Nguyen, Vitalii Panchyk, Jesse Pisel, Kenton Prindle, Jeff Roth, Matthew Schultz, Yaroslav Sihetii, Nate Suurmeyer, Ognjen Tanovic, and Ashley Yanner.

Thanks also to the broader team at ThinkOnward and SoftServe for their support of the program's operations.

Notes

1. B. Steffens, Q. Corlay, N. Suurmeyer, J. Noglows, D. Arnold, and V. Demyanov, "Can Agents Model Hydrocarbon Migration for Petroleum System Analysis? A Fast Screening Tool to De-Risk Hydrocarbon Prospects," Energies 15, no. 3 (2022): 902. Abstract retrieved October 21, 2026. https://doi.org/10.3390/en15030902 back
2. ThinkOnward, "synthoseis," GitHub repository. Retrieved October 22, 2026. https://github.com/thinkonward/synthoseis back
3. ThinkOnward, "geophysical-foundation-model," GitHub repository. Retrieved October 22, 2026. https://github.com/thinkonward/geophysical-foundation-model back
4. ThinkOnward, "challenges," GitHub repository of starter notebooks and winning solutions. Retrieved October 22, 2026. https://github.com/thinkonward/challenges back
5. A. Karpathy, "autoresearch," GitHub repository, March 2026. Retrieved October 22, 2026. https://github.com/karpathy/autoresearch back