

Catalist for Opportunities: Bulgarian AI Factory BRAIN++

Prof. Simeon Stoyanov

Managing Director, BRAIN++ AI Factory | Business Development Director, Discoverer

Catalist for Opportunities: Bulgarian AI Factory **BRAIN++**

Why, What & How: applied AI on sovereign resources



Prof. Simeon Stoyanov

Managing Director, BRAIN++ AI Factory | Business Development Director, Discoverer

5 October 2026
Hotel Chernomorec, Varna



Co-funded by
the European Union



EuroHPC
Joint Undertaking

Prove of Navier–Stokes existence & smoothness issues September 2026

Two mathematicians spent a year on this problem, feeding their drafts into an AI coding tool. Twelve hours after they announced their result, the AI company announced the solution.

88

hours

an AI system worked on a 90-year-old maths problem

10,000+

agents

working in parallel on the proof

12

hours

between the mathematicians' announcement and the AI company's

OpenAI, 8 Sept 2026: *“We cannot rule out that de-identified data derived from their usage of our products helped improve our models.”*

Do you know if someone or something is learning from your work?



Hands up

Who used an AI tool for work today?



Keep it up

...if you know whether that conversation is used to train the model.

Every prompt, every agent, every automated workflow teaches an AI how your business works and how it solves problems.

From our data to our know-how



Yesterday: personal data

Large providers often protect it better than a small company's own server.



Today: our workflows and ways of thinking

Every prompt, agent and correction teaches AI how we work and solve problems, in more and more industries. That know-how stays with the provider.

OpenAI pays \$150 an hour for it. We give it away for free.

Project Mercury (Bloomberg, Oct 2025): 100+ former investment bankers paid to teach an AI how financial models are built



Imagine your grid-balancing or forecasting workflow becoming a feature of someone else's model.

Sovereign AI in practice: five layers of control



Data

Where is it processed, and under which jurisdiction?



Models

Can you run, inspect and keep the model you depend on?



Compute

Who decides who gets the GPUs, and on what terms?



Operations

Can you keep running if a provider changes or leaves?



NEW

Know-how

What does the AI learn from your workflows, and who owns it?

Sovereignty is not autarky (economic independence and complete self-sufficiency) -> Europe still runs on non-European chips. Sovereignty means the ability to choose, to control and to exit.

And the classic risks still apply



Jurisdiction

“No, I cannot guarantee it.”

Microsoft France on US access to EU data, French Senate, June 2025



Switch-off

E-mail cut after sanctions

ICC Chief Prosecutor, after US sanctions on the court, 2025



Compute

Bulgaria in Tier 2

US AI Diffusion Rule, Jan 2025 (rescinded May 2025)



Continuity

Gone overnight

Models retired, one cloud outage stops thousands of services

For operational results, predictability matters as much as performance.

Europe's answer: AI Factories

19

AI Factories

13


AI Factory Antennas

up to 7

AI Gigafactories

- ✓ **What they are**
Public AI supercomputing plus data, skills and services for startups, SMEs, industry, science and the public sector
- ✓ **Free for innovation**
Access for startups and SMEs through open calls and fast lanes
- ✓ **AI Gigafactories**
Call launched 30 July 2026, deadline 12 Nov 2026; expected to unlock more than €20 bn of private investment

Sources: EuroHPC JU; AI Gigafactories call (30 July 2026)



BRAIN++: the Bulgarian AI Factory


Selected by EuroHPC JU

in March 2025, part of Sofia Tech Park, co-funded by the EU and Bulgaria

Infrastructure

Discoverer and Discoverer+ in operation

Discoverer++ in deployment:
600+ GB200/GB300 superchips,
FPGA partition, two Cerebras systems



Five questions before you deploy AI

- 1 Where is your data processed, and under which jurisdiction?
- 2 Does the provider learn from your prompts, agents and workflows?
- 3 Can you run the model yourself, for example with open weights?
- 4 Can you switch provider tomorrow (AI router, exit plan)?
- 5 Can the solution be audited under the EU AI Act?

Your workflow is the new training data. Frontier closed and open weight models & where they add value; core know-how on sovereign infrastructure.

Three stories of applied AI



ALSO Bulgaria

Nadezhda Dinisheva

My experience with artificial intelligence

My question:

Which tasks did you keep away from public AI tools, and why?



TiD-Engineering

Iliya Zhelyazkov

What if energy could think?

My question:

Where do your energy models run, and who learns from the data?



Energopro

Intelligent energy management: from data to decisions

My question:

If your AI provider switched off tomorrow, what would stop?

Backup: Navier–Stokes timeline, September 2026

2025–2026

- T. Buckmaster and L. Alpöge build on the Córdoba–Martínez-Zoroa programme; for a year they enter drafts and research into OpenAI's Codex

1 Sept

- OpenAI starts an 88-hour run with 10,000+ agents, after hearing rumours of other solutions

7 Sept

- Buckmaster and Alpöge announce their results (Euler, Boussinesq; verified in Lean)

8 Sept

- About 12 hours later, OpenAI announces a solution of the Navier–Stokes problem (166-page proof, Lean)

10–13 Sept

- OpenAI: recent user prompts could not have influenced the system

11 Sept

- Clay Institute: problem “apparently settled”; verification “deliberately unhurried”

Sources: Axios (8 Sept 2026), CNBC (9 Sept 2026), Clay Mathematics Institute statement (11 Sept 2026)

Backup: BRAIN++ in a nutshell

