



Preparing Europe for Transformative AI

A tech tree strategy for agency, resilience and flourishing

Arq Foundation

Contents

1 Introducing Arq Foundation	1
--------------------------------	---

2 What's at stake	3
Agency	4
Resilience	6
Flourishing	8

3 Our approach	11
Europe's tech tree	12

4 Topics Arq will work on	15
AI Infrastructure	16
Middle-Power Coordination	17
AI Resilience R&D	18
Breakthrough Innovations	19
Statecraft	20

5 How Arq operates	23
Full-Stack Model	23
Forward Deployed Impact	23
Philanthropically Funded	23
AINative	23

6 Conclusion	24
----------------	----

Introducing Arq Foundation

We are excited to launch Arq Foundation, a new think tank based in Brussels and active throughout the continent. Our mission is to help Europe maintain and improve its agency, build resilience, and flourish in an age of transformative AI.

“Our mission is to help Europe maintain and improve its agency, build resilience, and flourish in an age of transformative AI.”

- Arq Foundation

We believe AI is not just another normal technology. Today's systems are already transforming cybersecurity, solving 80-year old mathematical problems with no human input, and are starting to disrupt entry-level job markets. We, like Commission President von der Leyen, expect future systems could approach, or even match and surpass human performance across cognitive work, a milestone some refer to as Artificial General Intelligence (AGI).

We don't know precisely when this threshold will be crossed, or whether it makes sense to speak of a sharp border at all. AGI could arrive within years, or it could take a decade to arrive and diffuse. We are uncertain about the exact path, but across most scenarios, we believe Europe is underprepared to mitigate the risks and ill-placed to capture the benefits.

We like to see ourselves as pragmatic optimists: European Union values and way of life are worth protecting, but in a hardened geopolitical environment, protecting them means facing difficult trade-offs and offering a positive vision to win public support. That, oddly, is where our optimism comes from. Preparing for transformative AI is not separate from the competitiveness push already underway across the institutions, nor the movement towards European sovereignty and resilience across energy, security and technology. It is part of the same agenda, viewed from a different angle.

Our vision is a Europe that grows the industries AI will reshape rather than watching them grow elsewhere, that builds the safeguards to weather what's coming, and that puts its science and its state to work at the pace the moment demands. The ideas mostly already exist, scattered across reports, labs, and movements throughout the continent. What's been missing is a way to see them as one agenda, which is what this essay tries to offer.

What's at stake

Europe faces three interconnected challenges.

1 | Agency

Europe must maintain agency over its AI future by building international leverage on AI;

2 | Resilience

Europe must build resilience to AI- and AI-amplified risks;

3 | Flourishing

Europe must create a flourishing society from the benefits of AI and frontier technologies.

Agency

Europe risks losing access to frontier AI systems and in turn losing a foundation of its agency.

Frontier AI capabilities are becoming increasingly important for national security and economic competitiveness. The two countries leading their development, the US and China, may therefore decide to limit or revoke access and keep the technology largely to themselves. This becomes more likely if the supply of compute (the AI datacentres used to train and serve models) continues to fall well short of demand. Because compute supply is highly inelastic — new fabs and power take years to come online — the pie is effectively fixed in the near term, making allocation zero-sum: every chip serving a competitor or ally is one not serving your own economy and military.

Europe has already seen access restrictions. Before launch, Anthropic's most capable AI model, Claude Myths, was used to identify thousands of software vulnerabilities across every major operating system and web browser. Select US companies and cybersecurity agencies received advance access to help safeguard their own software; the EU did not. A subsequent request from Anthropic to extend access to a small number of European organisations was vetoed by the White House on national security grounds. Even though the EU ultimately received access to Myths, access restrictions will remain a threat for every future model release.



At this point, we think it is unlikely that European AI companies will catch up to frontier capabilities. Even Chinese competitors are already falling further and further behind the American frontier, and on independent benchmarks, Mistral's strongest model scores only about 2/3 that of the leading US frontier models. Mistral's emphasis on open source may yield a profitable, successful business as a sovereign second mover — but it will not close the frontier capability gap. With this economic and military vulnerability, even short of fully revoking access, the US and China may still threaten to do so to extract geopolitical concessions, as they have with tariffs and rare earth minerals.

RAND research suggests that frontier AI will be transformative for economic and military power. Industrial and R&D bases could have their competitiveness hinge on access to frontier capabilities, and AI-augmented militaries could be a step-change similar to drone technologies. If this is the case, losing frontier access could amount to losing something as fundamental as electricity. By default, this dependence will leave Europe vulnerable, without a seat at the table, and seeing both its sovereignty and agency slowly erode.

“Losing frontier access could amount to losing something as fundamental as electricity.”

- Arq Foundation

But Europe isn't powerless.

Two responses stand out:

- dramatically growing Europe's share of the AI stack to build leverage and
- pooling said leverage with the handful of other middle powers in similar positions.

We explore both in: Topics Arq Will Work on – AI Infrastructure and Topics Arq Will Work on – Middle Power Coordination.

Resilience

Europe must become resilient to AI- and AI-amplified risks.

The Mythos episode has shown that today's AI systems are capable enough to autonomously breach well-protected software; on the horizon are AI agents with even more destructive potential, for instance, ones that could help design novel and dangerous biological agents.

Misuse is not the only concern. AI systems remain very poorly understood: developers know how they are trained but not why they behave the way they do once training is complete. Despite fast capability development, none of the leading AI companies has a detailed plan for how to robustly align their models' values with human interests;

all rely on "automated AI alignment": training AI systems to handle the problem themselves.

“Misuse is not the only concern. AI systems remain very poorly understood.”

- Arq Foundation

Even the most responsible developers operate under competitive pressure where tripling revenue can hinge on shipping months ahead of peers. This most often makes slowing down and investing in more safeguards infeasible.



As capabilities advance, the risk of dangerous misuse may drive further securitisation — governments treating frontier AI as national-security assets to be classified, export-controlled, and shielded from foreign scrutiny. Meanwhile, at the major AI companies, the majority of code is now written by AI (more than 70% at Anthropic), and these systems generate the training data their successors are trained on and monitor deployed AI systems in real time. All of this leaves Europe with less visibility and oversight over leading AI companies.

“Europe can still maintain resilience — by deliberately building it.”

- Arq Foundation

However, Europe can still maintain resilience — by deliberately building it. By accelerating the technologies that keep defenders ahead of attackers, Europe can directly shape technological progress in a better direction. We explore promising approaches to this and what Europe’s role could be in: Topics Arq will work on – AI Resilience R&D.

Flourishing

Europe must learn to integrate AI to underpin a flourishing society.

Europe today offers a way of life that many enjoy: generous welfare, short working weeks, long holidays, beautiful cities, and comparably safe societies. But a key foundation of that flourishing — economic growth — is increasingly weak. Several economies like Poland (and Ireland and Malta) have converged to the highest income Member States impressively. But converging toward the frontier is not the same as the frontier advancing. Worryingly, IMF projections show France falling from 86% of US per capita GDP in 2000 to 71% by 2030, Italy from 93% to 68%, and Spain from 72% to 61%. As Draghi and a growing body of work make clear, this is overwhelmingly a productivity gap, largely driven by Europe's lag in the technologies that power frontier growth. Some argue the gap is exaggerated by how we measure prices across countries, but the most careful attempts to check confirm it is real — and AI, on its present trajectory, deepens rather than closes it.

With public spending being stretched by rearmament and demographic change, economic growth and competitiveness are not just a nicety; they are essential to maintaining and improving Europe's present quality of life.




“Economic growth and competitiveness are not just a nicety; they are essential to maintaining and improving Europe’s present quality of life.”

- Arq Foundation

We worry that the bottlenecks causing stagnation will be all the more painful with AI: dismissal procedures that make labour reallocation prohibitive, punitive bankruptcy laws that disincentivise risk-taking, and capital markets too fragmented to fund European firms at scale. The cost of these bottlenecks is missing out on education that adapts to each student, healthcare that catches disease earlier, more effective public services and faster progress on climate. These are the things AI-led growth most directly delivers, and the things Europeans most clearly want.

Public backlash compounds this problem. AI is already unpopular on the continent, and that wariness is hardening into organised opposition to new datacentres. To tap into the benefits of AI, European leaders need a positive story that brings voters along — one about shaping the technology’s impacts on citizens’ lives, building new industries, and supporting those displaced. This story needs to be built through public input: where datacentres are built, how local communities are compensated, and which AI systems enter public life. A story the public helps write survives contact with the communities it affects; a story the public is only told does not.

The good news is that the pieces for this story exist already. The Draghi, Letta, and Heitor reports lay clear many of the reforms that, if made, would create a competitive, prosperous and equitable Europe: they just need to be connected to urgency AI presents. (See: Topics Arq will work on – Breakthrough Innovations.)




Our approach

Tackling these three challenges will require many parallel approaches: Classical policymaking at not-so-classical speeds, democratic strengthening, grassroots movements, public-engagement experiments, and forms of action we can't yet anticipate. As a small organisation, we can't do all of these at once, and we don't believe any single lever will be enough.

However, across all three problems, we've identified a category of solutions that we are particularly excited about: Europe **using technology** to build agency, resilience and flourishing.

Alternative solutions exist — proposals built around regulatory leverage, moral suasion, or the economic mass of the single market alone — but none of them alone leaves us as confident about Europe's fate. Solving all three of these problems appears only doable if Europe can harness the immense power of technology. We hold this view with caution and humility, out of necessity, not blindness; if others can make other approaches work, we'd welcome the company.



Europe's tech tree

Even if technology is central to Europe's managing transformative AI, not all technology is built equal. In particular, a naive view may be that to properly prepare for transformative AI, one must build it. While this would certainly help, it is neither absolute nor likely feasible.


It's important for Europe to compete where it has a chance. On frontier AI development, this will prove challenging. At today's technological maturity, state-led efforts likely face a choice between being too slow, prohibitively expensive, or unlikely to succeed. Mistral's sovereign second mover strategy is unlikely to offer Europe frontier capabilities. There are a few calibrated bets to leapfrog the frontier that Europe can and should take. However, Europe's future requires robust back-up plans if these long-shot approaches fail.

So where, if not in frontier AI development, should Europe build? It helps to imagine technology development as a tree: new discoveries unlocking further discoveries that branch outwards from them. If advanced AI systems are the trunk of the 21st-century technology tree, we think Europe should strategically pursue three distinct parts of it: the **Roots**, **Branches**, and **Flowers**.

The roots of the AI tech tree are the foundational technologies that make up today's AI stack. Whether it's the chips powering AI training or inference, the machines making those chips, or the parts that go into those machines, the AI supply chain is complex and highly concentrated. Europe already has some excellence here, with Netherlands-based ASML (who make the machines that make the chips) and Germany's Zeiss (who make mirrors that go in those machines). However, when you zoom out, Europe is a small fraction of the supply chain. Strategically growing its share of this upstream market is, we think, the most direct way to build short-term leverage and ensure continued access to AI-powered R&D and manufacturing.

The branches of the AI tech tree are the possible parallel technologies developed alongside AI. The tech tree is often thought of as linear, but there are often hundreds of technologies that could be differentially pursued in parallel. mRNA vaccines were developed entirely independently of existing vaccine paradigms, the entirety of the open-source Linux movement started as a hobby, and cheap solar panels only came into existence through a combination of German political will and Chinese manufacturing prowess. This matters especially for resilience: the market's default isn't always the socially optimal one. It's often up to governments and philanthropy to steer, cultivate and foster technologies in these differential directions. In the medium-term, strategically accelerating parallel technologies can give Europe resilience in a world increasingly transformed by AI.

The flowers of the AI tech tree are the downstream technologies that advanced AI will accelerate or enable. Some flowers will wilt; many will bloom beautifully. A few will yield so much fruit that they will go on to seed entirely new trees on their own. Identifying and building these downstream technologies in Europe is, we think, vital for its long-term relevance, and provides a path to not just follow but lead. These technologies are still early enough that government intervention can do a lot: R&D funding and procurement incentives can derisk and meaningfully accelerate technological development. While picking winners is hard, picking winning fields is much more doable. Europe could be the home of fusion, photonics, or who-knows what else.



Topics Arq will work on

With this strategic frame in mind, most of Arq's focus will be on five policy areas.

1. **AI Infrastructure** builds Europe's presence in the roots of the tech tree.
2. **Middle-Power Coordination** is the geopolitical strategy that turns those root assets into negotiable leverage.
3. **AI Resilience R&D** focuses on the branches that keep Europeans safe as capabilities grow.
4. **Breakthrough Innovations** reforms the R&D engine that produces the flowers.
5. **Statecraft** creates the institutional foundation enabling all of the above.



AI Infrastructure

Building AI inputs to build AI leverage

The upstream AI stack — chips, datacentres, energy — is where much of the leverage currently sits. AI companies desperately need more supply, and whoever provides it will be able to exert significant influence over these companies and the development of their technology. Unfortunately, Europe isn't presently on track to doing this. The US controls almost 80% of global AI compute, China 14%, and the EU just 5%, and the gap is widening.

Europe's flagship response, the AI Gigafactories programme, has already had its call for proposals delayed and construction pushed to 2029, all for only some 0.6 GW of datacentre capacity. By the time they are operational, Arq projects the US will have installed more than 100 GW.

The private sector is ready to invest in Europe, but runs into bottlenecks at every turn. Permitting can take 48 months. Grids are overstressed, and even European champions like ASML have signalled potential moves abroad over housing, immigration, and grid constraints. While China adds as much electricity generation as Germany's entire grid every year, European production has stagnated. If datacentre operators cannot build in Europe, they will build in the US, often with more environmental harm in the process.

“AI companies desperately need more supply [of compute], and whoever provides it will be able to exert significant influence over these companies and the development of their technology.”

- Arq Foundation

These are solvable problems: compute demand is tripling every year, which Europe could harness and turn its 5% share into, say, 15–20% within a few years. But, the continent will have to make uncomfortable trade-offs. European providers cannot build sufficient compute on their own. Member states will have to give the private sector more room to build, and that includes American hyperscalers building compute for American labs. If Europe makes it easy for foreign companies to build on European soil, it can demand access and sovereignty guarantees in return, good for both sides of the Atlantic.

Middle-Power Coordination

Pooling leverage to keep a seat at the AI table

European countries don't individually hold enough leverage over AI to negotiate on equal terms with the US and China, as shown by the EU's experience with Mythos. But they hold important pieces of the AI supply and value chain, and so do a handful of other middle powers in similar positions, like the UK, Canada, Japan and South Korea. Coordinated, those pieces add up to something substantial: enough, plausibly, to deter coercion from either Washington or Beijing.

“Coordinated, those pieces add up to something substantial: enough, plausibly, to deter coercion from either Washington or Beijing.”

- Arq Foundation

A natural starting point is a Middle-Power Coalition that begins small and focuses narrowly on AI. Its members would secure as much of the upstream supply chain on their soil as possible, partnering with both non-American and American companies, and helping domestic players in the semiconductor supply chain expand, and bundle those assets into a joint bargaining position. Over time, a broader and more formal coalition could branch out: mediating between the US and China, or playing an active role in shaping binding international AI frameworks. Arq is exploring these questions and aims to deliver concrete proposals later this year.

The instinct is already in the air: at Davos, Mark Carney made the case for middle powers banding together to defend their shared interests, and the framing has continued to circulate since. The task now is to turn that instinct into something operational.

AI Resilience R&D

Steering technology to build societal resilience

As AI capabilities accelerate, the threats they enable will outpace existing defences. The continent's resilience will depend in part on whether it can deliberately accelerate technologies that keep defenders ahead of attackers, sometimes called differential technology development or d/acc. Several candidates are already ripe for attention: interpretability methods, which help researchers understand why AI systems behave the way they do; formal verification techniques could harden the software running critical infrastructure against AI-enabled attack; epistemic tools could help societies make sense of an AI-polluted information environment; and personalised agent paradigms — where every user has their own AI working in their interest — could prevent the concentration of power that comes from a handful of providers mediating everyone's access. Built out deliberately — each citizen with an agent that answers to them, and communities able to see and contest how AI shapes their information and their choices — these could become not only defences but public infrastructure in their own right.

Counterintuitively, the bottleneck isn't funding — it's ideas, and the people willing to carry them forward. Philanthropy, public funding, and markets for technologies that empower societies will together mobilise over hundreds of billions of Euros. The harder problem is determining what to build, finding the people to build it, and helping them start. We are particularly inspired by the Institute for Progress's Launch Sequence, and see room for similar work in Europe: not only to identify the technologies society needs today rather than tomorrow, but to surface and back the talent willing to make them real.

“Governments have plenty of levers”

- Arq Foundation

Governments have plenty of levers: standing up new R&D agencies dedicated to resilience-enhancing technologies, earmarking part of the 5% NATO spending commitment for resilience R&D, partnering with philanthropists to channel pledged capital toward the technologies Europe most needs, or using procurement reforms like advance market commitments to derisk these technologies before they're market-ready.

Breakthrough Innovations

Reforming R&D to lead in tomorrow's technologies

Every technological breakthrough starts with research and development, but not all R&D leads to breakthroughs. Through a mix of risk-aversion, excessive bureaucracy, and limited room for reflection and creativity, science, Europe's engine of progress, has slowed down. We think that's fixable.

Europe spends over €125 billion annually on public R&D, a foundation worth defending, but much of it is undermined by the system that allocates it. A single grant application takes 36–45 person-days to write. The process is so complex that many applicants hire specialist consultants, paying up to 12% of the total grant in success fees just to navigate it. The EU's research funding system discourages risk-taking, and has never systematically tested whether its own processes work.

AI raises the stakes here: advanced systems could significantly accelerate science itself, compressing discovery cycles and automating much of experimentation — and forcing funders to adapt, as AI-assisted proposals overwhelm existing review systems. The countries able to absorb all this are likely to gain a competitive advantage. This will require clearing practical bottlenecks — datasets, hardware, intellectual property — ahead of time, before they bite.

The necessary reforms aren't hypothetical. In the US, Focused Research Organisations (FROs) and ARPA-like agencies already successfully fund science differently — fast, mission-driven, willing to fail — and the UK has set up a Metascience unit on the premise that "the scientific method... should be systematically and routinely applied to how we practise, fund and support science itself." Europe could do the same: freeing researchers from excessive paperwork, backing young scientists to pursue ambitious ideas, experimenting with new funding models, and testing whether they actually work.

“A rigid research system doesn't only cost Europe breakthroughs; it slows the resilience technologies that Europe most needs to stay safe.”

– Arq Foundation

A rigid research system doesn't only cost Europe breakthroughs; it slows the resilience technologies that Europe most needs to stay safe. If Europe acts now, it can help shape the science AI is poised to transform, rather than having to adopt versions others build.

Statecraft

Reforming the state to act at the speed of AI


Nobel laureate and Google DeepMind CEO Demis Hassabis describes AGI as "ten times the industrial revolution at ten times the speed — unfolding over a decade instead of a century." If true, this calls for a century's worth of reform and intervention from the state in the same window, at a time when measured public service productivity has declined across France, Germany, and Italy. None of the other topics in this essay matter without states well-equipped to act on them. That's what Statecraft is about.


“None of the other topics in this essay matter without states well-equipped to act on them.”

- Arq Foundation

As AI capabilities improve, and governments face increasingly difficult challenges under increasing uncertainty and time pressure, it will be essential to have first-rate expertise inside government; external advice will be hard to procure at pace and will lack the deep context necessary for agility.

That requires designing institutions that can hire elite talent and empower them. The UK's AI Security Institute (AISI) has poached world-class talent directly from AI labs by paying well and doing away with bureaucratic friction; when Anthropic released Claude Mythos, the UK was the only European country given advance access, because it was the only one with a body able to evaluate it. France's Beta.gouv.fr makes a complementary bet on its operating model: since 2013, it has run roughly 150 "startups d'État" — small teams embedded inside ministries that ship public digital services product-style, with autonomy to iterate and explicit permission to fail. Both show that world-class technical talent will join the state if given the conditions they need: pay, autonomy, and a clear public mission.





How Arq operates

These are large and challenging problems. We're setting up Arq in a way that gives us the best shot at addressing them.

- 1. Full-stack model.** A lot of consequential policy work happens at a highly technical level where decision-makers want to engage directly with the person who spearheaded the analysis, not an intermediary. Therefore, we hire individual subject-matter experts who own our work on specific policy areas end to end: from ideation to research to outreach. We call them Full-Stack Policy Entrepreneurs. The model also keeps overhead low, so the vast majority of marginal funding goes toward expanding into new policy areas.
- 2. Forward-deployed impact.** Other consequential decisions are made by senior, politically engaged decision-makers who work across many issues at once and require the broad political picture. Engaging here takes more than subject-matter depth, so we also hire Forward-Deployed Policy Entrepreneurs: experts in science communication and political dynamics, based in select capitals. They act as advisors and force-multipliers, strengthening the feasibility of our experts' recommendations and grounding them in the broader political picture. They also tailor Arq's analysis to local contexts and lead research and outreach that spans multiple experts' fields, like public consultations, so that Arq's work stays connected rather than siloed.
- 3. Philanthropically-funded.** Arq aims to be funded by unrestricted, mission-aligned philanthropy. Two failure modes follow when think tanks aren't: prioritising restricted funding can distort focus, and funding from the institutions they advise may produce audience capture and watered-down recommendations. Being selective about funding comes at a cost, but we think it's the right trade-off. Indeed, several of the workstreams in this essay remain prospective, pending additional funding. If our theory of change excites you and you could help us scale, [please get in touch](#).
- 4. AI-native.** We believe organisations designed to be AI-native from day one will have a meaningful edge. We are intentionally building our infrastructure, safeguards and processes to scale with the technology we are trying to grapple with. We're currently hiring for a builder-in-residence to help us accelerate this work.

Conclusion

We worry that Europe is currently on a trajectory to lose out on its agency, resilience and flourishing as transformative AI arrives. Arq has singled out five areas we believe are most important to address those challenges: AI Infrastructure, Middle-Power Coordination, AI Resilience R&D, Breakthrough Innovations, and Statecraft. A changing world is being forced upon the continent, and it's up to Europeans to decide whether to prepare or merely react.

We are not alone in pushing this agenda. Inside the European institutions, a growing consensus is forming that the status quo is untenable: [Mario Draghi](#) has made the case on competitiveness, [Enrico Letta](#) on the single market, and [Manuel Heitor](#) on research and innovation. EVP [Henna Virkkunen](#) has built her Commission portfolio around tech sovereignty, and EPP MEP [Christian Ehler](#) is pushing for ARPA-style reform of how the EU funds research.

Others are pushing in the same direction. On AI safety, Turing laureate [Yoshua Bengio](#) is advancing safe-by-design AI through LawZero. On democratic renewal, Audrey Tang has spent years showing how civic AI — deliberation at scale, and tools that give citizens a hand in steering decisions — can strengthen democracies rather than corrode them.

In the US, the Institute for Progress's [Launch Sequence](#) lays out concrete proposals for AI-for-science and AI-for-resilience. Closer to home, think tanks like the [KIRA Center](#) in Berlin and Carnegie's [Anton Leicht](#), and publications like the Garicanos' [Silicon Continent](#) and [Works in Progress](#), are building the analytical base — while emerging coalitions like [Accelerate Europe](#) are turning ideas into political momentum.

“Our hope is for Europe to regain this role for transformative technologies like AI and beyond.”

– Arq Foundation

The agenda is ambitious and it's reasonable to ask whether Europe is realistically capable of delivering on it. A reason for optimism is that Europe has done this all before. Europe was instrumental in building the modern world: the steam engine in Glasgow and penicillin in London, vaccines in Paris and GLP-1s in Copenhagen, EUV lithography in Eindhoven and the web at CERN. Each of these seeding a technology tree that remade its century, and for three hundred years Europe fueled the frontier of human technology. Our hope is for Europe to regain this role for transformative technologies like AI and beyond.

The tree of transformative AI is still young — and where its roots take hold, which branches are cultivated, and whose soil its flowers bloom in are not yet decided. That is the choice in front of Europe: to tend this tree deliberately, for its agency, its resilience, and its flourishing, or to let others grow it in its place. Europe can still choose its future.



Read the essay online



Arq Foundation

www.arq.foundation