

PROJECT PROPOSAL

Assessing Taiwan's Global Position for Transformative AI

Taiwan manufactures the world's most advanced AI chips and sits at the center of US-China tensions – yet Taiwan is consistently perceived as a bystander rather than an actor in AI safety. This project produces a first baseline demystifying Taiwan's role in AI governance and identifying the specific ways it can leverage its position for global AI safety, written to reach the policymakers who can act on it.

Project Type: Policy research, governance analysis, landscape mapping

Duration: 3 months (September–December 2026)

Mentors: David Sánchez García and Kevin Chen

Ideal Mentees: Background in public policy, international relations, or East Asian or security studies; interest in transformative AI and AI governance; strong analytical writers; Mandarin reading ability for some modules.

Project Overview

Any scenario in which AI proves to be transformative would be heavily influenced by decisions made in Taiwan, yet this area of research is almost completely neglected. Taiwan produces ~90% of the most advanced logic chips in the AI supply chain, and sits at the center of the US-China technology conflict.

There's a small but growing community of AI safety in Taiwan (e.g. NTU AI Safety), Taiwan's DSET has done some adjacent research – though mostly focused on the US-China dynamic and China's semiconductor industry, and the Singapore AI Safety Hub is doing work on compute verification mechanisms that could be interesting for Taiwan. One of the mentors has published a few pieces on Taiwan's AI law, AI sovereignty, and the position of Taiwan in relation to other countries in East Asia. Still, the field is clearly neglected compared to its potential impact, and this project intends to fill this gap.

Decisions made in Taiwan, whether it be on energy, export policy, chip governance, and participation in any future verification regime, shape the global compute landscape whether or not anyone in Taiwan is thinking in those terms. Right now very few people are. A baseline assessment gives international researchers and policymakers something to build on when Taiwan enters AGI-relevant conversations. Hardware verification proposals run physically through the bottleneck of Taiwanese fabs, so where Taiwan can engage with them depends on an analytical capacity that is massively unexplored in AI policy. A Taiwan able to participate expands the option space for US-China coordination on AI, and this project also aims to grow a research pipeline in a strategically underserved jurisdiction.

Structure and Research Questions

Each mentee owns one of four modules end to end. Each module tries to answer a similar analytical template: current posture and key actors, what the posture implicitly assumes about the pace and scale of AI development, how it holds up if significantly capable AI arrives faster than current timelines, and the concrete gaps and levers that follow.

Expect an 8–12 hour weekly engagement (but we are flexible!); the mentors will stay close to the details of each module and frequently engage with your drafts as well as through scheduled meetings while you own the module yourself.

- **Policy situation:** What is the current state of AI in Taiwan? What institutions and initiatives exist? How is AI safety conceptualized? How much funding exists and who is providing it? Some relevant institutions: MoDA, NSTC, the AI Evaluation Center, Taiwan AI Labs, TAIDE, Taiwan's AI Basic Act, DSET.
- **Defense and cognitive security:** Understanding information and disinformation campaigns, how substantially more capable AI could change the cross-strait threat model, and how societies can prepare for this epistemic risk.
- **Compute and verification:** Taiwan's compute infrastructure and energy constraints, and its potential role in hardware-enabled verification mechanisms. What is Taiwan's semiconductor industry's appetite for verification? What levers could the Taiwanese government pull to increase its role in semiconductor policy? Is compliance a relevant angle?
- **Elite discourse mapping:** Structuring a map of the individuals talking about AI in Taiwan, their institutional affiliations, and what they say about advanced AI – mapped through legislative statements, documents, and media sources (C2 Chinese proficiency heavily favored). Potential interviews.

Additionally, mentees would gain access to the mentors' network of AI policy researchers.

Why This Matters

Taiwan is often described as a chokepoint, bottleneck, and one of the most consequential jurisdictions in the AI supply chain. Its dynamics with China and the United States are among the largest geopolitical wildcards in how transformative AI unfolds. Currently, Taiwan has a fledgling, but lagging, indigenous capacity focused on transformative AI-strategy. Its AI policy apparatus is built for incremental AI and the international preparedness conversation figures Taiwan as a bystander at best, or a footnote at worst.

That gap itself is a global risk, and the project addresses this in three ways. First, a baseline assessment gives international researchers and policy conveners something to build on when Taiwan enters transformative-AI conversations. Today there is no document that specifically addresses this issue. Second, hardware-enabled verification proposals run physically through Taiwanese fabs; whether Taiwan can engage with them depends on local policy capacity that does not currently exist, and a Taiwan able to participate expands the option space for global coordination on AI. Third, field-building: the project is designed to

recruit and train Taiwan-based or Taiwan-interested researchers, growing a pipeline in a strategically critical, badly underserved jurisdiction.

The Traditional Chinese executive summary extends this from analysis to outreach by creating a document Taiwanese policymakers and researchers can act on and cite.

Deliverables

- **Main deliverable:** a two-page executive summary in Traditional Chinese and English written for Taiwanese policymakers – deliberately more direct in register than the hedged research underneath it, because the point is a document stakeholders in Taipei will actually read.
- A public report (~8,000 words: modules plus synthesis) for the broader AI policy ecosystem.
- **[Optional]** Serialized publication on Substack or the EA Forum. One of the mentors has a Substack focused on Taiwanese AI governance, Chips & States, and would be happy to crosspost there if mentees are interested.

Timeline

- **Weeks 1–3:** orientation, template refinement, preliminary research to limit scope, source inventories.
- **Weeks 4–8:** module drafting, peer-review pairs, mid-point comparison session.
- **Weeks 9–12:** revisions, comparative synthesis, Chinese language (Traditional) and English executive summary, Demo Day materials.

Projected Outcomes

- A first baseline of Taiwan's transformative-AI posture that international researchers can build on and cite – a document that does not currently exist.
- A Traditional Chinese executive summary written to reach, and be acted on by, Taiwanese policymakers.
- A trained cohort of Taiwan-based or Taiwan-focused researchers, seeding a pipeline in a strategically underserved jurisdiction.
- For each mentee, a bylined research module and a strong analytical writing sample in AI governance.