Field Strategist

Location: Remote, with many calls in Pacific Time.

Preference for candidates based in the US who are eligible for a security clearance

Travel: likely ~20% travel

Salary range: $140 to ~200k/yr + benefits

Key Summary

Advancements in AI bring both opportunities and serious risks. Atlas Computing takes on society-scale AI risks that no existing institution is built to address — we design the intervention, validate it with domain experts and funders, then recruit and empower the specialist who will own it.

You can think of us as:

- A tiger team that identifies, scopes, and launches interventions that aren't a good fit for existing organizations
- A design studio developing and staffing product requirements docs for coordination and market failures
- A think tank combined with whatever their analog of "tech transfer" would be.

We start from the opportunity or risk, map the context in which the need for interventions arise. Then we sketch possible interventions, iteratively working with experts to ensure feasibility and potential stakeholders to establish sufficiency to have the desired impact. Finally, we will make it as easy as possible for a career professional to take ownership over each interventions (we believe most of these efforts don't require risk-seeking young generalists courted by traditional VCs, but rather are best addressed by career professionals who have spent the last decade or more executing on very similar work).

Atlas Computing is looking to hire a cohort of **Field Strategists,** each of whom will scope multiple potential interventions simultaneously and hand them all off.

While this is a separate role from our Founders (who will scope a single intervention they would then lead themselves), we expect the Field Strategists will work closely in collaboration with our cohort of Founders.

While Atlas is interested in many aspects of AI Resilience, we expect this first cohort will likely focus on some combination of AI infrastructure security (especially these 6 problems) and epistemic security; we will prioritize candidates with relevant experience.

As a Field Strategist, success in the first 6 months looks like:

- The hire has created two or more clear write-ups, each describing:
    - Why advances in AI capabilities pose a specific risk to society
    - What a proposed intervention would look like.
    - Who are the best people in the world to advise and steer the intervention
        - (and their attestation that this solution is plausible).
    - What decision-makers would change their behavior because of this intervention
        - (and their attestation that this intervention plan is sufficient).
    - Who would likely fund it
        - (and what their decision criteria would be).
    - What skills and experience the founder needs (and where are we likely to find that person)
- The hire has worked with a recruiter to pitch potential founders for that organization and has introduced one to potential funders.

- A more junior hire would:
    - focus on 1-2 interventions at a time.
    - benefit from checking in with Atlas leadership multiple times per week.
- A more senior hire would:
    - develop 3-5 organizations simultaneously and on ~3-5 month timescales.
    - benefit from the freedom of little oversight.
    - be "fire and forget", sharing weekly updates as needed to highlight potential opportunities.
    - (perhaps) provide mentorship to junior hires.

Note that we expect this first cohort to focus primarily on starting new efforts in the United States that could scale globally, rather than creating efforts starting outside the US, but this is not a hard constraint.

You may be a fit for this role if you have...

- expertise in navigating and/or modifying complex systems (natural or human-made),
    - This could look like complex systems research in academia, policy research at a think tank, market research in a product role, or analyst research in the intelligence community.
- generalist: able to quickly understand nuances in a variety of topics, including ML architecture, export control laws, and ideological differences within the AI community.
- meaningful experience in at least one of these areas:
    - Cybersecurity - Technical:
        - Software security, vulnerability analysis (or related).
        - Hardware security or hardware supply chain security (or related).
        - Identity & authentication systems (or related).
        - Privacy engineering or privacy threat modeling (or related).
        - Datacenter or cloud infrastructure security (or related).
    - Cybersecurity - Strategic:
        - Cybersecurity strategy, Red Team, or threat intelligence (or related).
        - Cyber-defense coordination or policy (e.g., IC, CISA).
        - Risk analysis in highly complex environments.
    - Epistemic Security - Technical:
        - Media forensics, deepfake detection, or content provenance (or related).
        - Platform trust & safety, content moderation systems, or recommender integrity (or related).
        - Computational propaganda analysis or bot/network detection (or related).
    - Epistemic Security - Strategic:
        - Information operations analysis or counter-influence strategy (e.g., IC, think tank, or allied government).
        - Election security, democratic resilience, or civic information integrity (or related).
        - Institutional epistemic governance or public trust measurement (or related).
- Cross-cutting
    - Strategic technological or product planning.
    - Working in an "ARPA" entity.
    - Technology roadmapping or multi-year strategic planning.
    - a strong ability to interview experts and synthesize across domains.
    - comfort forming and revising hypotheses quickly.

    - strong creative problem-solving skills, love finding pareto improvements to the obvious options.
    - a capability to reason about AI trajectory, safety, and future-state technologies.

At Atlas, we also expect everyone to have...

- a practice of reviewing the state-of-the-art after having an idea, coupled with an urge to prototype, test, and iterate.
- a founder mindset, with a high degree of ownership over outcomes (not process) and the ability to personally make and externally drive progress with little oversight.
- experience coordinating and aligning stakeholders to work toward a common goal.
- strong analytical, problem-solving, and communication skills.
- comfort with uncertainty and interest in shaping our culture, as we're still a very small organization.

Bonus points if you...

- strongly hold a conviction that AI is going to reshape the face of civilization fairly soon, and you want to join a hard-working team focused on making that go as well as possible.
- have deep networks in cybersecurity, biosecurity, or epistemic security
- have a background in product or taking things from zero to one: organizational formation, paired with domain depth.
- have experience acting as a COO-type integrator across complex, multi-expert efforts
- thrive in a high-trust, low-structure environment.
- love talking to people across multiple domains about concrete actions that can address important problems.
- know how to solve problems with money, but also know how to operate on a tight budget.
- seriously considered starting your own company, but the opportunity or push hasn't come up yet.
- find yourself always recommending books, models, and frameworks.
- are deeply committed to improving the human condition with technology...
- and understanding how different people might differently define "improving", "human", "condition", and "technology" in that sentence.
- have work authorization to work in the US (and ideally also would be eligible to hold a US security clearance)

We're excited to be running Navigators in partnership with Coefficient Giving. Read more in this blogpost.