

Research Director: Formal Verification and AI

Location: Remote

Travel: likely 10-20% travel

Salary range: \$180-230k/yr + a stipend for benefits

Note:

- Text in dark green is specific to this role
- This is a more senior version of our Research Lead role. The differences are small but meaningful changes to salary, expectations, and prior experience.

Key Summary

Advancements in AI bring both opportunities and serious risks, which simultaneously enable and necessitate more powerful approaches to high-assurance systems. Atlas Computing is a nonprofit working to ensure robust democratic oversight and control of critical infrastructure and AI. We are building an ecosystem for AI systems with provable properties.

Building an ecosystem is a massive group effort: we work with the formal methods, AI, cyber, and biosecurity communities across academia, government, and industry. We contribute through both technical work (targeted research, product development, technology transfer) and community organizing (education, government outreach, conferences, coordinating funding, incubating startups).

We believe that machine learning capabilities will continue to grow rapidly and should be directed at advancing formal verification. Atlas Computing is hiring a research director with demonstrated experience directing formal verification and AI research. This role is responsible for paving the way to dramatically accelerate formal verification through research that influences researchers, developers, funders, and policymakers. This involves pursuing a research agenda at the overlap of formal methods and AI safety communities, building a world-class research team, and collaborating with applications engineers to translate research into real-world adoption.

As the FV+AI Research Director, success in the first year looks like:

(In order of priority from highest to lowest)

- You designed a portfolio of research directions at the intersection of AI and formal methods and are helping funders design research agendas.
- You've helped source and convince funders to grow your team.
- You've advanced the state-of-the-art in using machine learning to accelerate formal verification at scale
 - You've already worked with our software lead to integrate one of these into a useful tool.
- Your technical insights and results have played a key role in Atlas's influence on high-impact decision-makers outside the organization (for instance, policymakers or regulators).

- You've enabled conversations with important stakeholders given your track record.
- Perhaps you've identified an amazing hire for Product Lead (if we haven't hired one already).

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

- experience running a lab or building a team that has advanced the state of the art in formal methods.
- extensive knowledge of the latest research in AI
- a PhD (or equivalent experience) and have published research using ML for formal methods in peer reviewed journals.
- a track record of sourcing your own research mentors and collaborators.
- comfort picking research directions that bet that capabilities of frontier models continue to grow.
- 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 and the ability to personally make and externally drive progress with little oversight.
- experience coordinating and aligning groups to work toward a common goal.
- an ability to think across disciplines and research-development timescales.
- strong analytical, problem solving, and communication skills.
- interest in shaping our culture, as we're still a very small organization.

Bonus points if you have...

- fluency with Lean or Coq.
- expertise managing high-performance machine learning compute.
- experience with high assurance engineering in real-world systems.
- experience building and applying ML models to solve complex problems.
- expertise in cybersecurity.
- familiarity with SCADA systems.
- experience shaping the research agendas of science funders.
- seriously considered starting your own company, but the opportunity or push hasn't come up yet.
- experience managing research consortia and knowledge of how research labs are run and funded.
- been the person in the conversation who is always recommending books.
- a deep commitment to improving the human condition with technology... and understanding how different people might differently define "improving", "human", "condition", and "technology" in that phrase.

Here's a supplemental living document if you have more questions: [More information for potential hires](#)