Tri-State AI Venture Sprint

Building the next Frontier in AI

From Frontier Research to Working Systems

Format	Intensive Seminar (3 on-site weekends + remote phases)
Duration	March – May 2026
Credit	Participation certificate; ECTS recognition possible (university-specific)
Workload	6–10 hours/week + 3 intensive weekends
Program Lead	Dr. Bastian Burger (Director TUM Venture Labs)
Partners	MCML Munich, Tübingen AI Center, ETH AI Center
Network	ELiAS Alliance
Version	January 2026

Nerds with a Mission

You didn't get into AI research because you wanted to optimize ad clicks. You got into it because you saw something fascinating, maybe even important, and you wanted to understand it. Maybe build something with it.

This program is for researchers who want to find out if their work can matter outside the lab. Not in the vague "broader impacts" section of a grant proposal, but in a concrete way: does anyone actually need this? Can we build something that works? What happens when real people try to use it?

We're not going to teach you how to write business plans or pitch to VCs. We're going to help you run experiments on a different kind of hypothesis: that your research can solve a real problem for real people. The methodology is similar to what you already know—form a hypothesis, design an experiment, gather evidence, update your beliefs—but the experiments involve users instead of datasets, and the success metrics are different.

Your "mission" might be ecological, social, scientific, or economic. What matters is that you have one and that you're willing to test whether your technical capabilities can actually serve it.

Why Now: The SPRIND Opportunity

Europe has exceptional AI talent spread across excellent institutions. What it lacks are teams that can take frontier research and turn it into deployed systems. The US and China have dozens of labs doing this. Europe has... fewer.

SPRIND—Germany's Federal Agency for Breakthrough Innovation—is trying to change this. Their Next Frontier AI Challenge is a €125 million bet that European researchers can build frontier AI labs. Starting July 2026, they'll fund ten teams over 24 months with compute, infrastructure, and support. The goal isn't to catch up on the current LLM race; it's to leapfrog to whatever comes next.

This matters for two reasons. First, the timeline: teams completing our program in May 2026 will be well-positioned to apply before the May 31 deadline. Second, the philosophy: SPRIND explicitly wants "disruptive model architectures, modalities, training paradigms, or multi-agent systems." They want technical depth and originality, not polished pitch decks.

Even if SPRIND isn't your path, the skills transfer: how to validate that your research addresses a real need, how to build systems that work outside controlled environments, how to communicate technical work to people who might fund or use it.

Who This Is For

You should apply if:

- You're a PhD candidate or researcher in ML/AI, robotics, or adjacent fields at one of the partner institutions
- You have technical work (published or not) that you suspect could be useful outside academia
- You're comfortable writing code and building systems, not just running experiments
- You're genuinely curious whether your research can solve real problems—and willing to find out it might not
- You can commit roughly 6–10 hours per week plus three intensive weekends

You probably shouldn't apply if:

- You're looking primarily for entrepreneurship training or networking
- You already have a clear startup idea and just need resources to execute it
- You're not willing to talk to potential users and hear uncomfortable feedback
- Your main goal is adding a line to your CV

Prerequisites

We assume you can code. Python, PyTorch/TensorFlow, whatever—the point is you can build things, not just analyze them. We also assume you can read a paper and understand what's actually new versus what's marketing. English is the working language.

Admission is by application. We're looking for technical depth, a plausible direction for impact, and genuine curiosity about whether your work can matter. Prior startup experience is irrelevant. Publications, preprints, and open-source work are all useful signals.

What We Actually Do

Three months. Three cities. One question: can your research become a working system that someone actually uses?

We're not building startups here. We're building technical demonstrators—systems that prove a capability works in practice, not just on a benchmark. Some of these might become startups later. Some might become open-source tools. Some might become the basis for a grant proposal or a collaboration with an industry partner. And some will turn out to be dead ends, which is also valuable information.

The structure: three intensive on-site weekends (Munich, Tübingen, Zürich), with structured remote work in between. Each phase has a focus:

- Phase I (Zurich): What problem are you actually solving? Who has it? How do you know?
- Phase II (Tübingen): Can you build something that works? What are its limitations? How do you evaluate it?
- Phase III (Munich): Can people outside your lab use it? What happens when they do?

Between on-sites, you're running experiments: talking to potential users, building prototypes, gathering evidence. We have weekly interview targets (10–15 conversations) because the most common failure mode is building something nobody wants. The methodology is borrowed from lean startup and design thinking, but adapted for deep-tech contexts where users often can't articulate what they need because they don't know what's possible.

What You'll Learn

By the end of the program, you should be able to:

1. Identify whether your research addresses a real problem, using evidence from actual potential users rather than your assumptions about what they need.
2. Build working prototypes quickly using modern tools (LLM integration, low-code platforms, rapid deployment), even if the underlying research is complex.
3. Evaluate your system honestly: document capabilities, limitations, and failure modes in a model card format that would survive peer review.
4. Design experiments to test whether your system works in practice, not just whether the underlying algorithm performs well on standard benchmarks.
5. Communicate technical work to non-researchers: potential users, funders, collaborators. This is harder than it sounds.
6. Navigate the landscape of support options: SPRIND, ERC Proof of Concept, EXIST, institutional programs, incubators, industry partnerships.
7. Decide whether to continue, pivot, or stop—based on evidence rather than sunk costs or enthusiasm.

Readings

We're not assigning a textbook. But these are worth your time:

Books

- Fitzpatrick, *The Mom Test* — How to talk to users without leading them. Short, practical, surprisingly hard to actually do.
- Ries, *The Lean Startup* — The original formulation of build-measure-learn. Some parts are dated; the core methodology isn't.
- Blank, *The Four Steps to the Epiphany* — More systematic than Ries. Useful framework for thinking about discovery vs. validation.
- Torres, *Continuous Discovery Habits* — How to make user research an ongoing practice rather than a one-time exercise.

Papers & Essays

- Mitchell et al., "Model Cards for Model Reporting" (FAT* '19) — How to document ML systems responsibly. We'll use this format.
- Graham, "Do Things That Don't Scale" — Why early-stage work looks different from scaled work. paulgraham.com
- SPRIND Next Frontier AI documentation — next-frontier.ai. Know what you might be applying to.

Schedule

The Rhythm

Three phases, each starting with an intensive on-site weekend and followed by 4 weeks of remote work. The on-sites run Thursday evening through Saturday afternoon. Between them, expect 6–10 hours per week: user conversations, prototyping, team coordination.

Throughout the program:

- Weekly user conversations: aim for 10–15 per week across your team. Keep a log.
- Bi-weekly check-ins: structured 30-minute calls with a mentor to review progress and blockers
- Shared Miro board: cross-team visibility so you can learn from each other's experiments

Phase I: Problem Discovery (Munich, Late February)

On-Site

- Research-to-problem mapping: where does your technical capability meet a real need?
- User discovery for deep tech: how do you interview people about problems they don't know are solvable?
- Hypothesis formation: what would have to be true for your research to matter here?
- Team formation and cross-institutional collaboration
- Context-setting on SPRIND and the broader European AI landscape

Remote Phase Deliverables

- 20+ documented user conversations with key insights extracted
- Primary problem hypothesis with supporting evidence
- Technical feasibility sketch: what would you actually build?
- Success criteria: what evidence would convince you this is worth pursuing?

Phase II: Technical Validation (Zürich, Late March)

On-Site

- Rapid prototyping: building working systems in days, not months
- Architecture decisions: what trade-offs are you making and why?
- Model cards: documenting capabilities, limitations, and failure modes
- Evaluation design: how do you know if it actually works?
- Guest session: researcher who successfully bridged lab-to-deployment

Remote Phase Deliverables

- Working technical demonstrator accessible to people outside your team
- Model card draft: capabilities, limitations, intended use, failure modes
- Evaluation protocol: how are you measuring performance in realistic conditions?
- Structured feedback from 10+ domain experts or potential users

Phase III: Deployment & Evidence (Tübingen, Late April)

On-Site

- From prototype to deployed system: what breaks when real people use it?
- Usage patterns: what can you learn from how people actually interact with your system?
- Communicating technical work: demos, narratives, documentation
- SPRIND application workshop: what makes a compelling submission
- Alternative pathways: grants, incubators, industry collaboration, open-source
- Demo Day preparation

Final Deliverables (End of May)

- Deployed system with documented usage (target: 20–100 users or 3+ formal collaborations)
- Complete model card with evaluation results
- Evidence pack: conversation logs, experiment tracker, usage data
- Demo and presentation suitable for SPRIND, grant applications, or similar
- 90-day roadmap: what happens next, regardless of which path you choose

What We Expect

This isn't a graded course. There's no exam, no rubric, no points. What there is: a set of milestones that help you make progress, and a cohort of peers and mentors who will notice if you're not doing the work.

Milestones

These aren't arbitrary checkboxes. They're the minimum evidence that you're actually running experiments rather than just thinking about it.

- 50+ user/expert conversations documented across the program
- Working technical demonstrator that someone outside your team can use
- Model card that honestly documents what your system can and can't do
- Evidence pack: what did you learn, how did you learn it, what changed as a result
- Final demo that communicates your work to a technical but non-specialist audience

What Success Looks Like

Success isn't "we built a unicorn." Success is: you ran real experiments, gathered real evidence, and now you know something you didn't know before about whether your research can address a real problem.

For some teams, that will mean: yes, there's something here, and we're applying to SPRIND / joining an incubator / pursuing a grant / talking to industry partners.

For other teams, it will mean: we learned that this particular application doesn't work, and here's why, and here's what we might try instead.

Both of these are good outcomes. The bad outcome is: we built something in isolation for three months and we still don't know if anyone needs it.

How We Measure the Program

We hold ourselves accountable too:

- 10–18 teams formed across the three institutions
- 80%+ of teams deploy a working technical demonstrator by program end
- Strong teams achieve 50–200 users or multiple formal collaboration commitments
- 3–5 projects positioned for SPRIND applications or equivalent follow-on support
- Open-source artifacts: code repositories, model cards, evaluation datasets shared publicly
- Lasting cross-institutional connections—the network is part of the value

Track Record

This format has been developed and refined over three years at TUM, TopMath, CDTM, and HHL, where it consistently receives the highest ratings among entrepreneurship-adjacent programs. About 20% of past participants have gone on to found companies. The tri-state version aims to raise that bar by connecting researchers across Europe's strongest AI institutions.

Partner Institutions

MCML — Munich Center for Machine Learning

Joint center of LMU and TUM. One of Europe's largest concentrations of ML researchers. Strengths in deep learning theory, computer vision, NLP, and AI for science.

Tübingen AI Center

Home of the Cluster of Excellence "Machine Learning: New Perspectives for Science" and the Cyber Valley ecosystem. Known for work on causality, robustness, interpretability, and AI safety.

ETH AI Center

Coordinates AI research across ETH Zürich. Particular strengths in robotics, autonomous systems, computational biology, and sustainable AI.

ELiAS Alliance

European Lighthouse for AI Startups—the broader network connecting deep-tech entrepreneurship efforts across Europe.

Apply

Applications open: January 2026

Deadline: February 6, 2026

Questions: bastian.burger@unternehmertum.de

Program site: venture-sprint.umso.co

SPRIND Challenge: next-frontier.ai

Ship, don't just publish.