

AI Safety, Strategy & Community Leadership

Hispanic AI Safety Institute | College learning program | Participant edition 1.1, September 2026

The question we will work on

How can we help increasingly capable AI benefit people while preserving human control, accountability, and the ability to recover when systems fail?

This 25-hour program equips students to evaluate claims about advanced AI, trace how harm could occur, compare technical and institutional safeguards, and propose a practical contribution. Participants apply these skills to a fictional community-serving institution and then to a problem of their choice.

Who it serves

College students, including community-college and transfer students, across computing, engineering, health sciences, social sciences, education, and other disciplines. No programming or advanced mathematics required. The pilot is designed for adults aged 18 and above. English and Spanish editions have the same learning objectives and assessment standards. Participants may write and discuss in either language.

HASI's outreach centers Hispanic communities while recognizing differences in background, language, and experience. Participation does not require agreement with a particular AI forecast or policy position.

What participants leave with

A claim-and-evidence record; a capability assessment; a causal risk map; a layered defense proposal; and a short final decision brief with a 30-day action plan.

Delivery

Five weeks or five days. Each module includes 180 minutes of preparation and independent work plus a 120-minute facilitated session, including a five-minute break. Total scheduled workload: 25 hours, including 25 minutes of breaks across the course. Optional enrichment is extra. Recommended group: 8 participants, with one facilitator; split larger cohorts into groups of 6–10.

Developed by HASI, based on the AGI Strategy curriculum and shared facilitation resources of BlueDot Impact. This is an independent HASI adaptation. It does not confer BlueDot certification or university credit.

Learning pathway and schedules

Five modules, five observable outcomes

1. Futures and decisions. Distinguish evidence, forecasts, and values; map at least four actors and their incentives. Deliverable: W1 decision note.
2. Capabilities and uncertainty. Explain compute, data, algorithms, and deployment conditions; identify two limits of a capability claim. Deliverable: W2 evaluation card.
3. Pathways to harm. Construct a causal chain with at least four steps, label assumptions, and identify a point of intervention. Deliverable: W3 risk map.
4. Defenses and governance. Compare three safeguards with owners, failure modes, tradeoffs, and a stopping rule. Deliverable: W4 defense plan.
5. Contributing responsibly. Defend a feasible decision using evidence, an alternative, and an explicit uncertainty. Deliverable: W5 final brief and action plan.

Five-week option

Weeks 1–5 cover modules 1–5 in order. Complete the 180-minute preparation block before each weekly meeting. Example meeting: 17:00–19:00 in the host's announced time zone. Submit W1–W4 at the session close; revise the final brief during session 5. Set actual dates with the host before inviting participants.

Five-day intensive option

Days 1–5 cover modules 1–5 in order. Each day: 09:00–10:30 independent work; 10:30–10:45 additional break; 10:45–12:15 independent work; 12:15–13:00 lunch; 13:00–15:00 facilitated session. Daily elapsed time is six hours, with five scheduled course hours. No previous-evening homework is necessary. Independent blocks can be supervised study rooms.

The 180-minute preparation block

30 min: read the HASI module note, glossary terms, and retrieve key ideas from memory. 30 min: source lab using the assigned source or the in-package bilingual alternative. 60 min: complete the module worksheet and its extension prompts. 45 min: build or revise the cumulative final brief. 15 min: check understanding and record a question. These are time budgets for active work, not estimates of reading length. Stop at the budget and mark unfinished items.

Core and enrichment

The workbook contains all core explanations, fictional evidence cards, and tasks in both languages. The International AI Safety Report has an official Spanish summary. Other English source links are optional enrichment for Spanish learners; the equivalent HASI task always remains available. Do not add a second full BlueDot reading load to these 25 hours.

Participation, completion and host responsibilities

How completion works

Submit all five worksheets, with a reasonable make-up route for missed sessions. Participate in at least four sessions; replace a missed session with a 300-word or three-minute oral analysis and a 15-minute facilitator check-in. Complete the post-course questions. The final brief must earn at least 8/12 on the published rubric, with no criterion at zero. One revision is available within seven days. Criteria assess reasoning, not agreement with HASI. Completion rules are HASI's proposed pilot standard.

A HASI completion certificate may state the program title, participant name, dates, and 25 scheduled hours. It must not imply professional licensure, academic credit, employment placement, or a BlueDot credential.

Working together

Challenge claims respectfully; distinguish a person's experience from a general claim. Share speaking time. A participant may pass, write, or respond orally. No one must disclose immigration status, trauma, family income, or personal encounters with harm. Use fictional examples for sensitive topics. AI tools are optional: identify their contribution, check cited sources, and explain your own reasoning. Never upload private participant or host data.

Host setup

HASI prepares the materials, facilitation and feedback. The host identifies a program contact, recruits participants, confirms accessibility and meeting arrangements, and agrees dates and completion procedures with HASI.

The proposed pilot is free for participants and host institutions. Delivery capacity and dates are confirmed with each host. No paid AI account or programming environment is required. Printed materials or offline files can be used.

Attribution and sources

BlueDot permits independent reuse and adaptation of its own materials with clear independent branding and credit. Linked third-party works retain their own terms. This package links to those works and supplies original HASI explanations and exercises, rather than reproducing complete readings. For an official BlueDot cohort or internal institutional arrangement, consult BlueDot's guidance and team.

[BlueDot reuse guidance](#) • [AGI Strategy](#) • [International AI Safety Report, bilingual summary](#)