

Build AI That Helps Communities Decide Better

Graduate Track (Masters & PhD) | AI for Systems & Society



What You're Building

An AI-powered decision-support system that helps institutions analyze complex societal challenges, simulate policy outcomes, and make more informed, transparent, and long-term decisions for communities.



Build window: June 14–21, 2026 | Teams: 2–5 students | Format: Fully virtual, global

The Problem

Communities and institutions make high-stakes decisions about housing, health, workforce, and public investment — often with **incomplete data, delayed feedback, and no tools to model long-term consequences**. Without better decision-support systems, resources are misallocated, interventions happen too late, and systems remain reactive instead of proactive. The problem isn't lack of intent — it's the absence of tools that make complex system tradeoffs legible and evidence-based.

Your Mission

Design an AI-powered system that helps institutions analyze complex societal challenges, simulate outcomes, and support better decision-making. Choose ONE direction below:

Direction A: The Cost of Doing Nothing Simulator

Model the long-term social and economic cost of delayed or absent intervention.

- Help decision-makers explore: what happens if intervention is delayed?
- Show how risks compound over time and quantify the cost of inaction
- Support scenario planning: 'if we invest X now vs. in 5 years, what is the projected difference?'
- **Possible domains:** Homelessness, workforce development, early childhood education, public health intervention
- Partner inspiration: St. John Center for Homeless Men · Depaul International · Community Solutions (Built for Zero)

Direction B: Community AI Readiness Blueprint

Help communities assess their readiness for the AI era and build a strategic adoption roadmap.

- Identify AI readiness levels across sectors: education, workforce, healthcare, government, nonprofits
- Surface gaps in digital access, AI literacy, and infrastructure
- Recommend priority investments for long-term resilience
- **Possible outputs:** AI readiness scorecards, sector dashboards, strategic roadmaps, policy recommendations
- Partner inspiration: OECD · World Bank · Brookings Institution · Cincy AI Catalyst · African Institute for Development Policy

✚ Optional Extensions (do not affect your base score — they deepen one dimension of your core solution):

- Gender Data Gap & Health Equity — identifying bias in datasets, modeling underrepresentation
 - Diagnostic Delay & System Bottlenecks — mapping delays, identifying systemic inefficiencies
- Choosing an extension does not help or hurt your evaluation. It allows deeper specialization.

Who You're Building For

Your users: city and regional policymakers, economic development organizations, nonprofit coalitions, workforce development agencies, public sector technology leaders.

Guiding question: How might AI help communities understand complex system tradeoffs and develop coordinated strategies for long-term, responsible decision-making?

What a Strong Solution Looks Like

- Clearly defines a **system-level problem and scope** — with explicit non-goals
- Identifies **key stakeholders and decision points** that the system informs
- Uses AI for **modeling, simulation, or structured analysis** — not surface reporting
- Shows architecture: data → model → insight → decision support
- Represents **time-based or system-level dynamics** (trends, scenarios, compounding effects)
- Addresses **uncertainty, tradeoffs, and model limitations** explicitly
- Demonstrates **lifecycle awareness** beyond the demo: drift detection, updates, governance

⚠ Common Mistakes — Judges will mark down:

- 'Predictive analytics' without specifying inputs, model type, training signal, and evaluation metrics
- Single-point predictions without scenario ranges or confidence intervals — policy models must represent uncertainty
- No bypass conditions — when should the system NOT be used?
- No model evaluation or lifecycle mention — what happens after the demo when the model drifts?

Data & Technical Approach

You may use:

- U.S. Census data · HUD housing datasets · public health datasets
- Workforce and labor statistics · World Bank Open Data · Urban Institute research
- Local government open data portals
- **Synthetic population or simulation data** — reasonable assumptions with documented methods

⊖ Avoid: using sensitive personal data, presenting outputs as deterministic predictions.

Responsible AI Requirements (Graduate Level)

1. Risk Identification — e.g. biased/incomplete datasets, misinterpretation of outputs, overconfidence in predictions, policy misuse

2. Concrete Mitigation — e.g. scenario ranges instead of point estimates, transparent assumptions documentation, sensitivity analysis, clear limitations statement
3. Human-in-the-Loop Design — define: (a) what the AI informs vs. what humans decide, (b) decisions requiring two review cycles, (c) bypass conditions where the system should not be used

AI Capabilities You May Use

- Predictive modeling and simulation systems
- Optimization models and causal inference
- Systems dynamics modeling
- Data integration pipelines
- Explainable AI (XAI) and scenario modeling

Real-World Partner Inspiration

You may draw inspiration from:

- St. John Center for Homeless Men · Urban Institute · Community Solutions (Built for Zero) · World Economic Forum · U.S. Department of Housing and Urban Development (HUD)

Additional inspiration:

- OECD · World Bank · Local government innovation labs · Brookings Institution · African Institute for Development Policy

What to Submit on Devpost

Devpost Field	What to Include
Qualifier Approval Code	Enter your 8-character code sent via email after passing the AI Qualifier (e.g., XX26-Q7H9K2M1). Submissions without a valid code will be disqualified.
Project Description	A complete description of your solution, the problem you solved, and who it helps.
Track & Challenge Selection	Select your track (High School / Undergraduate / Graduate) and your challenge number.
AI Architecture Explanation	Describe: (1) Inputs — what data goes in, (2) AI capability used (e.g., chatbot, classification, recommendation), (3) What processing happens, (4) Outputs — what the user receives.
Human-in-Loop Design	What is one decision your AI does NOT make? Explain why a human must remain involved at this point.
Responsible AI Guardrail	What is one risk in your system and how did you reduce it? Examples: bias, misinformation, misuse, over-reliance, exclusion.
Tools Used	List all AI tools, models, APIs, and any AI coding assistance. Indicate which are free vs. paid.
Data Disclosure	What data did you use? Include datasets, APIs, or synthetic data. If synthetic, explain how it was created.
3–5 Minute Pitch Video	A video covering: (1) the problem and user, (2) how your AI works, (3) a live or recorded walkthrough, (4) your responsible AI choice.

Devpost Field	What to Include
Working Demo or Walkthrough	A functional prototype, chatbot, workflow demo, interactive mockup, simulation, or decision-support model.

 Your video should show: (1) the institutional problem, (2) your system architecture, (3) a walkthrough of a key decision scenario, (4) your responsible AI and uncertainty handling.

How You'll Be Evaluated

Dimension	Weight	What Judges Look For
Problem Understanding	20%	Clear assumptions, explicit scope and non-goals, stakeholders correctly identified
AI Reasoning	35%	Architecture tradeoffs, evaluation strategy, justified model choices
Solution Design	20%	Modular system, clear component boundaries, deployment-aware choices
Impact & Insight	15%	Connects model outputs to real institutional decisions and stakeholder needs
Responsible AI	10%	Bias, monitoring, misuse, model drift, lifecycle beyond demo — all addressed

 Judge's Lens: The differentiator at grad level is infrastructure thinking. Did you design for what happens after the demo? A system that acknowledges its own failure modes is more credible than one that doesn't.

Impact

The strongest projects will:

- Help institutions move from **reactive** → **proactive decision-making**
- Improve how resources are allocated across communities
- Make complex systems more understandable and transparent
- Demonstrate how AI can responsibly support policy and societal outcomes

Frequently Asked Questions

Do I need coding experience?	No. You are encouraged to use no-code and low-code AI tools. What matters is your thinking, not your technical stack.
Can my team be from different countries?	Yes! Teams of 2–5 can span different schools, cities, and countries. Register together on Devpost.
What counts as a prototype?	Anything that demonstrates your AI working: a functional app, a chatbot walkthrough, a workflow demo, an interactive mockup, or a clearly narrated simulation. It does not need to be production-ready.
What should my video show?	Cover four things: (1) the problem and user, (2) how your AI works, (3) a live or recorded walkthrough of your solution, and (4) your key responsible AI choice.
Do paid AI tools give an advantage?	No. Judging rewards reasoning and design thinking, not technical resources. Open-source and free tools are equally valid.
What happens after submission?	Submissions go through AI screening, then judge review (Round 2), then a final panel (Round 3) for the top 10 teams per track. All teams receive results by email.
Can I use AI tools to help build my project?	Yes. You must disclose which AI tools you used in your submission. Using AI to help build does not disqualify you — transparency does.
When will I know if I qualified?	Qualifier results are released June 11–13. You will be notified by email and can also check your status on the hackathon platform.
Direction A vs B — which is more technical?	Direction A (Cost of Doing Nothing) typically requires simulation or modeling skills. Direction B (Readiness Blueprint) can be built with data integration, scoring, and visualization. Both are equally valued.
What is a 'non-goal' and why do I need to state one?	A non-goal is something your system explicitly does NOT do. Example: 'This system does not make allocation decisions — it informs them.' Stating non-goals shows you understand system boundaries.
Can I apply this to a country outside the U.S.?	Yes — and encouraged. Use real public data from your country or region. Judges value global examples. Just document your data sources clearly.

Questions? Contact us at aihackathon@usaii.org or join our Discord: discord.gg/ePjenJnyh4
aihackathon.usaii.org