BUILDING A ROADMAP FOR AI-ENABLED HUMAN AND ENVIRONMENTAL HEALTH PROTECTION

Author: OpenAI's ChatGPT and the HESI Global Scientific Staff
Date: June 2025

Executive Summary

On June 4–5, 2025, the Health and Environmental Sciences Institute (HESI Global) convened a two-day in-person workshop "*Building a Roadmap for AI-Enabled Human and Environmental Health Protection*" in Washington, DC, to explore the strategic use of artificial intelligence (AI) in advancing human and environmental health protection within a One Health framework. The workshop gathered over 150 participants from more than 10 countries, representing a wide range of sectors including academia, government, industry, and nonprofit organizations. A novel element of the workshop was the deployment of a customized generative AI model (Workshop GPT), designed to synthesize live inputs and identify key thematic focus areas—referred to as Guiding Pillars—that serve as strategic destinations on the path toward AI-enabled health and environmental systems. Drawing from expert presentations, moderated discussions, and real-time polling, the Workshop GPT produced five Guiding Pillars: Connected Data, Trust and Validation, Real-World Use Cases, Capacity Building, and Responsible Design. These Guiding Pillars define critical areas for the practical and responsible implementation of AI across biological, ecological, and regulatory domains. The report that follows presents these outputs, combining AI-generated content with extensive review and refinement by the HESI Global scientific team.

Introduction

The accelerating complexity of global environmental and public health challenges has underscored the urgent need for integrative, scalable, and adaptive solutions. Artificial Intelligence (AI) and Machine Learning (ML) technologies offer powerful opportunities to enhance the capacity to model, monitor, and manage risks across biological and ecological systems. In recent years, AI has emerged as a transformative force, enabling novel approaches across these scientific and regulatory domains. However, the field is still actively evolving, and there is a critical need to evaluate where AI can be most effectively applied and how to bridge traditionally siloed human, animal, and environmental health domains across a One Health platform through AI-enabled connections.

This report summarizes the outcomes of a two-day experimental workshop organized by HESI Global to identify challenges and opportunities in the application of AI-enabled tools for human and environmental health protection. Held on June 4–5, 2025, at the National Academy of Sciences in Washington, DC, the workshop "*Building a Roadmap for AI-Enabled Human and Environmental Health Protection*" convened 151 registrants, including 139 in-person attendees from 11 countries, and featured 23 invited expert speakers. Participants represented a range of disciplines—including toxicology, computational biology, regulatory science, public health, engineering, ecology, and ethics—drawn from academia, government, industry, and nonprofit sectors. The objective was to explore the role of AI in human, ecological, and environmental health protection systems, and to define areas where progress could be made.

A defining feature of the workshop was the integration of generative AI technologies in the execution and output of the event. A customized Generative Pre-Trained Transformer (GPT) instance of OpenAI's ChatGPT was developed and deployed to support real-time synthesis and documentation. GPTs, developed from large language models, can be pre-trained on existing testing data and have the ability to generate text based on specific inputs and requests. This specialized model was trained on a curated dataset (see methods and appendix) and used throughout the workshop to identify thematic focus areas, called Guiding Pillars. The Guiding Pillars represent core areas where investment is necessary to advance the effective and responsible use of AI for human and environmental health assessment. The workshop agenda was structured using an Outcome Pathway (OP) framework (**Figure 1**) and organized into four thematic sessions: (1) Molecular Initiating Events and Cellular Responses, (2) Organism and Population Responses, (3) Ecosystem and One Health, and (4) Exploring the Last Mile of Implementation.

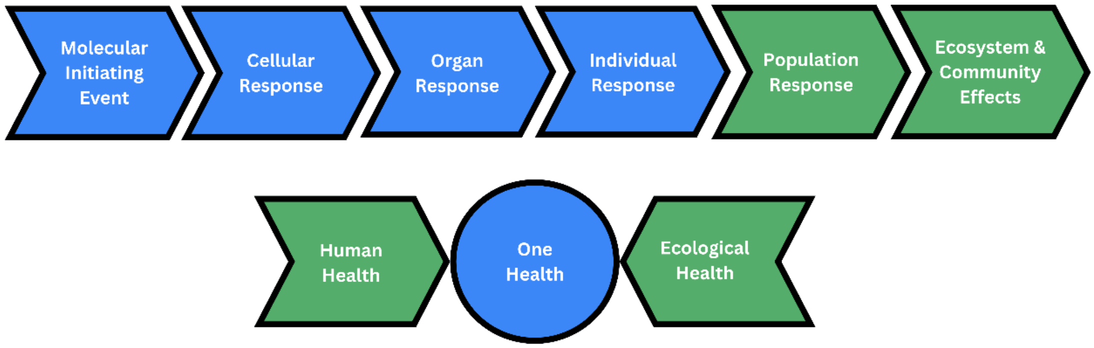

Figure 1: Outcome Pathway Framework upon which workshop sessions were organized, scaling from molecular initiation events to the ecosystem and community. *Figure adapted from the Adverse Outcome Pathway (AOP) framework.*

Each session included expert presentations, moderated discussions, and real-time polling, all of which were used to support the GPT-generated synthesis of Guiding Pillars. The outputs were generated through an iterative process combining AI-driven synthesis with active human guidance, including detailed review, editing, and refinement of the Guiding Pillars by the HESI Global scientific staff to ensure clarity, accuracy, and alignment with workshop content. The Guiding Pillars were developed by analyzing the workshop's aggregated outputs and are expected to be 'major destinations' on the road to AI-enabled human and environmental health protection. Ongoing discussions will help define how, when, and by whom these destinations should be explored. Recognizing that the positioning of these Guiding Pillars on a 'map' will vary by organization, application, and geography, this work and the workshop products are shared freely with the scientific community, with a request for citation and an invitation to adapt and build upon them. We also welcome feedback and examples of how this work is being used (contact: hesi@hesiglobal.org).

This report should be viewed as a prototype—an exploratory demonstration of how AI, paired with expert curation and review, can accelerate the development of strategic perspectives on emerging science. While AI expedited certain elements of the workshop synthesis process, significant time and expertise were required for initial AI training and development phases, technical curation of transcripts, slides, reports, and inputs during the meeting, and the extensive review and revision of the outputs. The final product reflects both AI-generated content and extensive refinement by the HESI Global staff. A detailed peer-reviewed methods paper describing this process is forthcoming.

Methods

GPT Configuration and Experimental Framework

A custom-configured instance of OpenAI's GPT-4o model (Workshop GPT) was deployed to support real-time synthesis of expert insights into structured Guiding Pillars representing thematic focus areas. This model followed a three-stage, prompt-governed workflow developed by HESI Global staff specifically for the context of this workshop.

The Workshop GPT followed a staged output process based on sequential input (**Figure 2**). Before the workshop, it was loaded with foundational materials, including the roadmap title, intended audience, speaker background papers [1–46], the full agenda, themes to guide alignment, and contextual information about HESI Global. During the live event, the model continuously received real-time transcripts from speaker presentations and panel discussions, along with live polling data from our Mentimeter polls and staff notes. With each new input, the

Workshop GPT reassessed and refined the set of Guiding Pillars in an iterative loop, updating the outputs until no additional content was available, and a final version was confirmed.

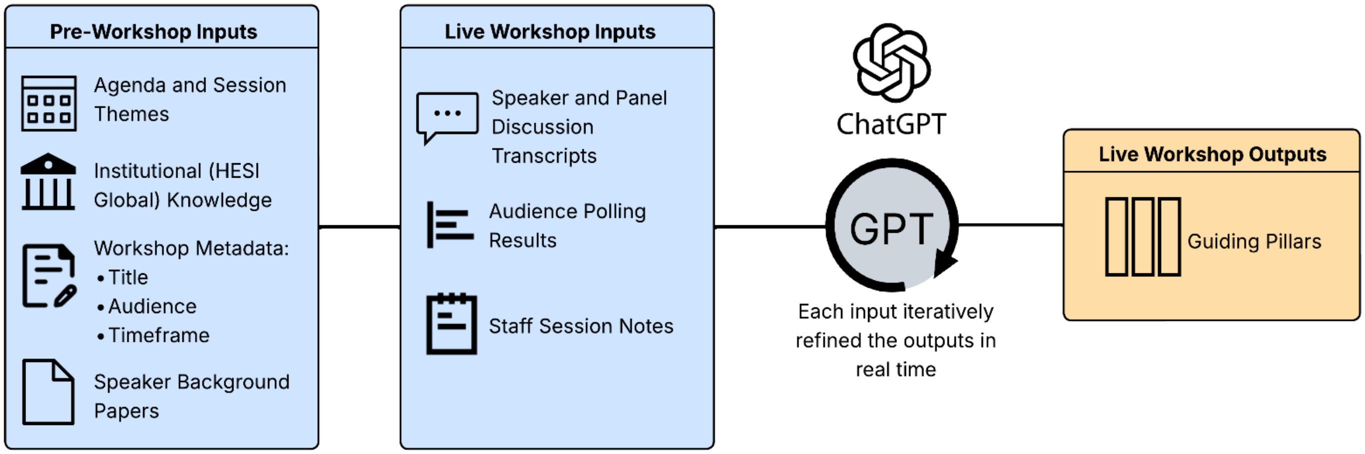

Figure 2: Staged GPT process, highlighting the pre-workshop and live workshop inputs loaded and iteratively refined to produce the live workshop output. *Figure credit – Dr. Alexandra Taraboletti*

Pre-Deployment Testing

Prior to deployment, the Workshop GPT underwent structured stress testing using historical workshop transcripts and simulated input data thematically aligned with One Health, exposure, toxicology, and AI governance. This process included testing roadmap generation against prior transcript sets and hypothetical polling results and introducing partial datasets (e.g., missing speaker biographies or agenda context) to observe model flexibility and structural adherence. Based on these tests, HESI Global staff refined the Workshop GPT's prompt architecture to improve the consistency of output structure and ensure that the model appropriately weighted key content inputs.

Live Workshop Testing

Configuration variations of the Workshop GPT were deployed live to evaluate the contribution of different input types—for example, runs excluding speaker background papers, runs with and without staff notes, and a benchmark run using Anthropic's Claude model with identical prompts. Each variation generated a distinct version of the Guiding Pillars, though with much overlap. The final Guiding Pillars were derived from the combination of run outputs and were presented live during the concluding session of the workshop. This live integration enabled direct audience feedback on the accuracy and relevance of AI-synthesized content, closing the loop between input, analysis, and collective review.

Guiding Pillars

The Guiding Pillars presented here reflect a synthesized collection of outputs generated through AI-enabled analysis and edited by HESI Global Staff (**Figure 3**). Per **Figure 2** above, they were developed by integrating speaker transcripts, panel discussions, polling data, staff notes, and workshop background materials. While care was taken to ground each Guiding Pillar in source content, the resulting framework inevitably mirrors the emphasis and perspectives of participating speakers and attendees. As such, this section should be viewed as a cross-section of collective priorities identified and discussed at the time of the workshop and in pre-meeting materials.

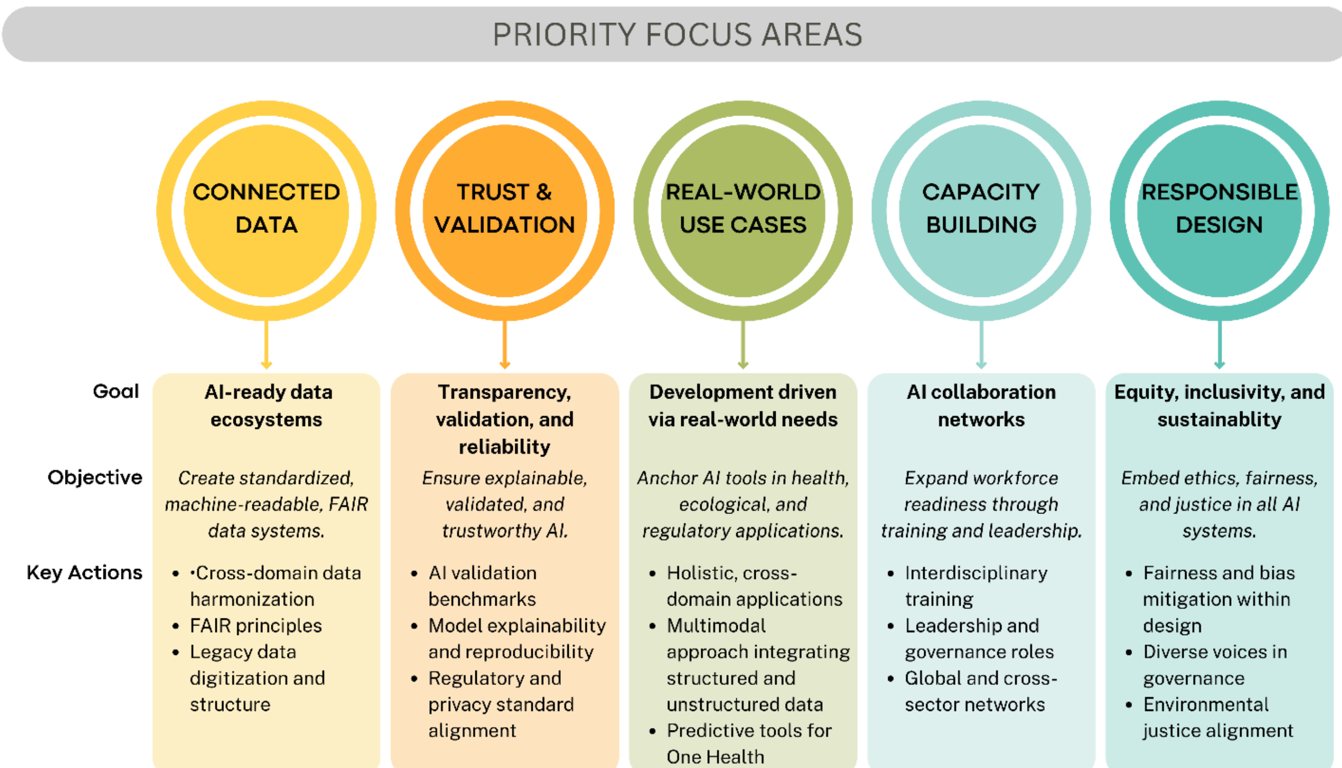

Figure 3: Summary of Guiding Pillars. The Guiding Pillars represent priority focus areas identified during the workshop. *Figure credit –Dr. Julie Krzykwa*

The following pages provide a detailed overview of each of the five Guiding Pillars. These summaries capture the core themes, objectives, and feasibility considerations discussed during the workshop and synthesized through the Workshop GPT and expert review process.

PILLAR 1: Connected Data

Establish comprehensive, interoperable, and context-specific data ecosystems that enable robust real-world AI solutions across health, environmental, and population scales.

AI applications in health and environmental sciences require seamless access to data from diverse domains, including molecular, clinical, environmental, and population levels. Participants across sessions consistently emphasized that fragmented, non-interoperable, and poorly structured data limit the potential of AI systems. The highest ranked polling option across all sessions was the need for better machine-readable and harmonized data. A consistent theme emerged around building an integrated, standardized, and FAIR (findability, accessibility, interoperability, and reusability)-aligned data infrastructure as a foundation for trustworthy and effective AI deployment.

Objectives & Key Actions

Support Machine-Readable Data Development	Accelerate the creation of structured, machine-readable datasets by prioritizing the digitization and annotation of experimental, clinical, and environmental data using ontologies and controlled vocabularies aligned with AI model requirements.
Improve Data Harmonization and Interoperability	Establish cross-sector data standards and interoperable metadata frameworks [e.g., HL7 Fast Healthcare Interoperability Resources (FHIR), Open Biological and Biomedical Ontologies (OBO) Foundry] to enable seamless integration of toxicology, biomedical, environmental, and exposure datasets for model training and reuse.
Implement FAIR Data Principles	Enforce adherence to FAIR (Findable, Accessible, Interoperable, Reusable) principles through mandatory metadata tagging, persistent identifiers, and open-access data portals to ensure AI models can discover and leverage diverse data sources.
Standardized Acquisition Protocols	Develop and disseminate consensus-based protocols for consistent data collection, including time-point sampling, sensor calibration, assay documentation, and version control, to ensure comparability across studies and sites.
Legacy Data Utilization	Implement automated digitization pipelines using natural language processing and image recognition to extract structured data from legacy records (e.g., PDFs, scanned lab notes), enabling retrospective analyses and AI training on previously siloed information.

Feasibility Rating: The following action areas were identified within this Guiding Pillar as high impact and high to moderate feasibility:

- Launch machine-readable data initiatives (top-rated recommendation across all sessions)
- Develop FAIR-aligned infrastructure
- Initiate legacy dataset conversions (deemed high feasibility and of immediate utility)

The following key needs will be challenging to address (i.e. 'low to moderate' feasibility)

- Promoting increased data harmonization standards and certifications
- Supporting the development of consolidated data repositories with cross-sector and geographically diverse participation

PILLAR 2: Trust and Validation

Establish rigorous, transparent, and continuous validation frameworks that build confidence in AI systems among users, regulators, and communities through demonstrated real-world value.

Stakeholders from all sectors highlighted the critical need for transparency, explainability, and model validation. From the challenge of prompt sensitivity in large language models to real-world discrepancies in toxicity predictions, there is a growing demand for models that are not just high-performing but understandable and reproducible. Panelists emphasized trust-building mechanisms like public case studies, regulatory alignment, and stakeholder engagement as critical.

Objectives & Key Actions

Develop AI-Specific Validation Frameworks	Design rigorous validation protocols tailored to AI workflows—including cross-validation, stress testing, and external benchmarking—that assess model robustness, generalizability, and domain-specific performance before deployment.
Promote Explainable and Interpretable AI	Advance algorithmic transparency by integrating techniques such as attention maps, feature importance ranking, or counterfactual explanations to ensure users can understand, audit, and trust AI outputs in high-stakes scientific decisions.
Engage with Regulatory Qualification Processes	Collaborate with regulatory agencies to co-develop and pilot model acceptability criteria—such as reliability thresholds, reproducibility standards, and decision-support boundaries—for preclinical, clinical, and environmental AI applications.
Develop Real-World Use Cases	Generate independent, peer-reviewed demonstration projects that apply AI to real-world problems—such as in silico trials, predictive toxicology, or outbreak forecasting—to establish empirical evidence of benefit, reproducibility, and usability.
Enhance Built-In Security/Privacy	Integrate federated learning, differential privacy, and encryption-by-design into AI systems to ensure compliance with the Health Insurance Portability and Accountability Act (HIPAA), General Data Protection Regulation (GDPR), and other data governance standards, while preserving analytical power and fostering user adoption.

Feasibility Rating: The following action areas were identified within this Guiding Pillar as of high impact and high to moderate feasibility:

- Build AI-specific validation protocols and regulatory pathways
- Launch real-world validation case studies with public release plans
- Develop explainable AI toolkits for toxicology and exposure science

They noted that the following key needs will be challenging to address (i.e. ‘low to moderate feasibility)

- Embedding security/privacy co-design in early-stage AI research and development
- Securing sustained regulatory engagement for iterative qualification processes

PILLAR 3: Real-World Use Cases

Develop and validate AI solutions driven explicitly by real-world health, environmental, and regulatory scenarios to ensure practical impact and stakeholder alignment.

Participants stressed the need to move from abstract model development to real-world use case implementation. Sessions 2 and 3 in particular emphasized examples like digital twins, early warning biosurveillance, and public health simulation tools. The emphasis was on holistic, cross-domain applications of AI—spanning human, animal, and ecosystem health, and multimodal approaches that integrate structured and unstructured data.

Objectives & Key Actions

Develop Multi-Scale Modeling	Develop clinically validated digital twin platforms that integrate multi-scale physiological modeling (from molecular pathways to whole organs) with real-world patient data (e.g., Electronic Healh Records, wearables) to enable in silico simulations for diagnostics, treatment planning, and regulatory submissions.
Improve Cross-Species/Domain Extrapolation	Train AI systems on harmonized mechanistic datasets to accurately predict toxicological and pharmacological responses across species and regulatory sectors, using standardized ontologies and embedded uncertainty metrics to flag low-confidence inferences.
Advance Ecosystem-Level Applications	Deploy AI-driven One Health models that integrate environmental, wildlife, and human surveillance data to forecast health threats from ecosystem disruptions, zoonotic spillover, or environmental exposures, enabling cross-agency preparedness and response.
Deploy Geospatial/Temporal Models	Integrate geospatial mapping, satellite data, and time-series health/environmental records into AI tools that model region-specific exposure risks and support targeted interventions for at-risk populations.
Design Multi-Modal AI Platforms	Build interpretable AI platforms capable of fusing diverse data types—scientific literature, clinical images, genomics, and sensor feeds—into unified models for toxicity prediction and public health decision-making across the human-environment interface.

Feasibility Rating: The following action areas were identified within this Guiding Pillar as of high impact and high to moderate feasibility:

- Develop multiscale models (high feasibility, real-world applications ready, e.g. Digital Twins)
- Begin ecosystem-level modeling for integrated health and environment systems.
- Enable geospatial and temporal data alignment across domains
- Integrate diverse modalities (e.g., -omics, clinical, environmental) into single modeling systems

PILLAR 4: Capacity Building

Expand global and cross-sector capacity to implement, understand, and sustainably leverage AI solutions through networks, training, and interdisciplinary collaboration.

Human capacity emerged as a foundational enabler across all sessions. As AI technologies evolve rapidly, there is an urgent need for scalable and inclusive training infrastructure, leadership frameworks, and communications pathways. From toxicologists to data scientists to regulators, multiple user groups need tailored upskilling and alignment. Cross-sector networks were also highlighted as vital for shared innovation and implementation.

Objectives & Key Actions

Design and Implement Training Programs	Launch curriculum-integrated programs and microcredential tracks that train life and environmental scientists in foundational AI concepts—such as machine learning, data ethics, and model validation—while immersing data scientists in regulatory, toxicological, and biological contexts.
Advance Novel Leadership and Governance Models	Develop institutional frameworks that delineate roles and responsibilities for AI implementation—including chief AI ethics officers, interdisciplinary advisory boards, and designated decision authorities—to ensure accountability and alignment with strategic goals.
Promote Public and Interagency Communication	Design multilingual, audience-specific communication toolkits and digital interfaces that translate AI outputs—like confidence scores or predicted risk levels—into accessible narratives for stakeholders, including the public, policymakers, and regulators.
Support Workforce Development	Establish apprenticeship, fellowship, and reskilling programs in AI for public health, environmental protection, and biomedical research fields, linked to real job placements and industry needs across government, non-profits, and private sectors.
Catalyze Global and Regional Networks	Formalize collaborative hubs and consortia that enable knowledge sharing, joint infrastructure development, and harmonization of standards across continents, leveraging regional strengths while aligning with international AI governance and scientific objectives.

Feasibility Rating: The following action areas were identified within this Guiding Pillar as of high impact and high to moderate feasibility:

- Stand up interdisciplinary training and cross-sector fellowships
- Develop cross-agency leadership structures for AI deployment of governance
- Create shared communication frameworks and public engagement guidance
- Pilot international and multi-agency coordination mechanisms

PILLAR 5: Responsible Design

Implement robust ethical frameworks and governance mechanisms that ensure inclusive, equitable, and responsible AI deployment while addressing fears and building community trust.

Polls and transcripts revealed deep concern about whether AI systems will reinforce existing disparities or create new ones. From model bias to data center energy use, ethical implications span social, environmental, and technical dimensions. This Guiding Pillar consolidates cross-cutting calls for inclusive governance, early-stage ethical design, and community engagement as central tenets of responsible AI.

Objectives & Key Actions

Advance "Ethics by Design"	Embed frameworks such as value-sensitive design and impact assessments into every stage of AI system development—from data selection to model deployment—to proactively ensure fairness, accountability, and social benefit.
Deploy Bias Mitigation and Fairness Tools	Deploy algorithmic auditing protocols and fairness-enhancing interventions (e.g., counterfactual fairness testing, re-weighting techniques) throughout model training and monitoring pipelines to detect and reduce disparities in performance across demographic groups.
Encourage Inclusive Governance	Establish multi-stakeholder advisory boards and participatory review panels that include historically marginalized communities, ensuring their perspectives shape policy, data governance, and resource allocation tied to AI systems.
Expect Community Co-Design	Implement structured co-development processes where affected populations actively contribute to the design, testing, and feedback loops of AI models—especially those deployed in public health, environmental monitoring, and clinical contexts.
Promote More Equitable Distribution of Costs and Benefits	Coordinate with energy regulators and sustainability experts to ensure AI infrastructure (e.g., data centers, edge devices) is equitably distributed, energy-efficient, and aligned with regional goals for environmental health and social equity.

Feasibility Rating: While participants agreed that this Guiding Pillar is a high priority and has high potential impact, they noted that the following key needs will be challenging to address (i.e. 'low to moderate feasibility)

- Establishing enforceable ethical guardrails and cross-border governance mechanisms
- Creating long-term oversight boards, including affected communities and regulators
- Developing metrics for bias, equity, and inclusion in AI deployment

Conclusion and Next Steps

This report represents an exploratory effort to use generative AI in co-creating strategic directions at the intersection of human and environmental health. The Guiding Pillars presented here reflect a synthesized view of shared priorities and insights from the workshop, shaped by AI-enabled processing of speaker content, polling responses, and supporting materials, with significant oversight, revision, and refinement by the HESI Global scientific team. Designed to follow the logic of an Outcome Pathway (**Figure 1**), the workshop addressed a wide range of topics from the understanding of toxicity at the molecular scale, to population environmental stress response, to the protection of ecosystems, and overarching ethical considerations. The breadth of these topics resulted in the development of Guiding Pillars that could be applicable to many evolving scientific arenas. However, this experiment provides an important proof of concept that could also be refined to address more focused questions in future iterations.

All workshop materials, including the Workshop GPT configuration, background papers [1–46], transcripts, polling results, and outputs, are publicly available through the HESI Global website. We encourage others to explore these resources, adapt the process, and develop their own roadmaps to these shared "destination" Guiding Pillars. HESI Global intends to pursue additional work in this space, including the development of a technical methods paper to document the GPT-assisted synthesis process in more detail. We also anticipate exploring where HESI Global's programs and collaborative networks may be most impactful in advancing specific priorities highlighted through this exercise.

Acknowledgements

Workshop Planning, Writing, and Design Team

The authors gratefully acknowledge the contributions of the HESI Global Staff, the Board of Trustees, and those who supported the planning and execution of this project.

Special thanks to Alex Taraboletti and Julie Krzykwa for their leadership in writing and GPT design/operations. We also thank Syril Pettit, Michelle Embry, Raechel Puglisi, Sandrine Deglin, Jennifer Pierson, Cissy Li, and Hannah Richardson for their essential roles in writing, workshop design, and meeting planning. We are especially grateful to Nancy Washton for her guidance on GPT-related strategy and implementation.

We also recognize the HESI Board of Trustees—Carole Yauk, Dorit Donoviel, Gary Minsavage, Justin Teeguarden, Karissa Adkins, Nina Hallmark, Sarah Hughes, Warren Casey, and Yasu Kanda—for their support in workshop planning and moderating. This report would not have been possible without the collaborative effort and dedication of the entire planning team.

Appendices

Appendix A: Glossary of Terms

- **AI (Artificial Intelligence):** Technologies that perform tasks typically requiring human intelligence.
- **FAIR:** Data principles ensuring Findability, Accessibility, Interoperability, and Reusability.
- **GPT (Generative Pre-Trained Transformer):** A large language model developed by OpenAI and adapted for real-time synthesis during the workshop.
- **HESI Global (Health and Environmental Sciences Institute):** The nonprofit scientific organization that hosted the workshop and oversaw the GPT-based roadmap development.
- **ML (Machine Learning):** A subset of AI focused on algorithms that improve automatically through experience and data.
- **OP (Outcome Pathway):** The framework used to organize the workshop agenda, scaling from molecular events to ecosystem impacts.
- **One Health:** An integrated approach that recognizes the interconnection between people, animals, plants, and their shared environment.
- **Workshop GPT:** The specific, customized instance of OpenAI's GPT-4o model used to support the live synthesis of workshop outputs.

Appendix B: Contributing Sources

Workshop contributions that informed the report development included presentations, polling responses, panel discussions, and supporting documents provided by speakers and organizers.

Invited Speakers and Keynotes

- Robert M. Califf – Opening Keynote
- Nicole Kleinstreuer – Session 1 Keynote
- Zhoumeng Lin – Session 1
- Carlie LaLone – Session 1

— Mathieu Vinken – Session 1
— Zhichao Liu – Session 1
— Thomas Steger-Hartmann – Session 2
— Dongying Li – Session 2
— Steve Levine – Session 2
— Lauren Charles – Session 3
— Victoria Baxter – Session 3
— Levente Klein – Session 3
— Julian Heinrich – Session 3
— Marcelo D'Agostino – Session 3 Keynote
— Shaolei Ren – Day 2 Keynote
— Syril Pettit – Welcome and Session 4
— Brian Berridge – Session 4
— Yasu Kanda – Session 4
— Vasiliki N. Rahimzadeh – Session 4

Additional Workshop Materials

The Workshop GPT model can be found at: https://chatgpt.com/g/g-67f5412e765c8191ad204875e542a23f-roadmap-ai-for-human-and-environmental-health

The workshop agenda, speaker slides, poster session details, and workshop reports can be found at: *https://hesiglobal.org/hesi-conference-series-2025-ai-health-protection/*.

Polling Results can be found at:
https://www.mentimeter.com/app/presentation/albjycbzp4qu6uj7i33ku31ye57d6xvm/edit?source=share-modal

Other workshop materials, including transcripts and prompt guides, are available upon request.

Background Source References

A full list of the speaker's background papers, scientific articles, and reference documents used to train the Workshop GPT. These sources served as foundational input to the AI system and helped develop the output.

1. Shimomura, Y., Komukai, S., Kitamura, T., et al. (2023) Identifying the optimal conditioning intensity for stem cell transplantation in patients with myelodysplastic syndrome: a machine learning analysis. *Bone Marrow Transplant*, **58**,186–194. Accessed 6 July 2025.

2. Aghasafari, P., Yang, P.C., Kernik, D.C., Sakamoto, K., Kanda, Y., Kurokawa, J., Vorobyov, I., and Clancy, C.E. (2021) A deep learning algorithm to translate and classify cardiac electrophysiology. *Elife*, **10**,e68335. Accessed 6 July 2025.

3. The IRC. (2025) Follow the Forecast: An Evaluation of the Rapid contingency Planning Approach to Anticipatory Action Through a Case Study in Guatemala. Accessed 6 July 2025.

4. Shabani, M. and Rahimzadeh, V. (2025) Introducing a Fairness Checkpoint for Data Quality and Evidence During Regulatory Review of AI/ML-Enabled Medical Devices. *Houst J Health Law Policy*, **24**,191–218. Accessed 6 July 2025.

5. Steger-Hartmann, T., Sanz, F., Bringezu, F., and Soininen, I. (2025) IHI VICT3R: Developing and Implementing Virtual Control Groups to Reduce Animal Use in Toxicology Research. *Toxicol Pathol*, **53**,230–233. Accessed 6 July 2025.

6. Steger-Hartmann, T., Kreuchwig, A., Vaas, L., et al. (2020) Introducing the concept of virtual control groups into preclinical toxicology animal testing. *ALTEX*, **37**,343–349. Accessed 6 July 2025.

7. Steger-Hartmann, T., Kreuchwig, A., Wang, K., Birzele, F., Draganov, D., Gaudio, S., and Rothfuss, A. (2023) Perspectives of data science in preclinical safety assessment. *Drug Discov Today*, **28**. Accessed 6 July 2025.

8. Steger-Hartmann, T., Duchateau-Nguyen, G., Bringezu, F., Onidi, M., and Stirn, M. (2025) Virtual control groups in non-clinical toxicology – A replicability challenge. *ALTEX - Alternatives to animal experimentation*,. Accessed 6 July 2025.

9. Sanz, F., Pognan, F., Steger-Hartmann, T., et al. (2023) eTRANSAFE: data science to empower translational safety assessment. *Nat Rev Drug Discov*, **22**,605–606. Accessed 6 July 2025.

10. Huang, L., Duan, Q., Liu, Y., et al. (2025) Artificial intelligence: A key fulcrum for addressing complex environmental health issues. *Environ Int*, **198**,109389. Accessed 6 July 2025.

11. Li, P., Yang, J., Wierman, A., and Ren, S. (2024) Towards Environmentally Equitable AI via Geographical Load Balancing,291–307. Accessed 6 July 2025.

12. Cavasotto, C.N. and Scardino, V. (2022) Machine Learning Toxicity Prediction: Latest Advances by Toxicity End Point. *ACS Omega*, **7**,47536–47546. Accessed 6 July 2025.

13. Li, P., Yang, J., Islam, M.A., and Ren, S. (2023) Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. Accessed 6 July 2025.

14. Moyano-Fernández, C., Rueda, J., Delgado, J., and Ausín, T. (2024) May Artificial Intelligence take health and sustainability on a honeymoon? Towards green technologies for multidimensional health and environmental justice. *Global Bioethics*, **35**. Accessed 6 July 2025.

15. Kleinstreuer, N. and Hartung, T. (2024) Artificial intelligence (AI)—it's the end of the tox as we know it (and I feel fine)*. *Arch Toxicol*, **98**,735–754. Accessed 6 July 2025.

16. Vinken, M., Benfenati, E., Busquet, F., et al. (2021) Safer chemicals using less animals: kick-off of the European ONTOX project. *Toxicology*, **458**,152846. Accessed 6 July 2025.

17. Moreno-Torres, M., López-Pascual, E., Rapisarda, A., et al. (2024) Novel clinical phenotypes, drug categorization, and outcome prediction in drug-induced cholestasis: Analysis of a database of 432 patients developed by literature review and machine learning support. *Biomedicine and Pharmacotherapy*, **174**. Accessed 6 July 2025.

18. da Silva, J.B., Lima, N.T., Haddad, A.E., et al. (2024) From national and regional commitments to global impact: artificial intelligence for equitable public health at the G20. *Revista Panamericana de Salud Pública*, **48**,e73. Accessed 6 July 2025.

19. Organization, P.A.H. (2021) Eight Guiding Principles of Digital Transformation of the Health Sector. A Call to Pan American Action. *Digital Transformation Toolkit;Knowledge Tools*,. Accessed 6 July 2025.

20. Pan American Health Organization. (n.d.) Artificial Intelligence in Public Health: Readiness Assessment Toolkit. Accessed 6 July 2025.

21. Pan American Health Organization. (n.d.) CE168/10 - Roadmap for the Digital Transformation of the Health Sector in the Region of the Americas. Accessed 6 July 2025.

22. US FDA. (n.d.) Considerations for the Use of Artificial Intelligence To Support Regulatory Decision-Making for Drug and Biological Products. Accessed 6 July 2025.

23. Wu, P.Y., Chou, W.C., Wu, X., Kamineni, V.N., Kuchimanchi, Y., Tell, L.A., Maunsell, F.P., and Lin, Z. (2024) Development of machine learning-based quantitative structure–activity relationship models for predicting plasma half-lives of drugs in six common food animal species. *Toxicological Sciences*, **203**,52–66. Accessed 6 July 2025.

24. Mi, K., Wu, X., and Lin, Z. (2025) Chemical risk assessment in food animals via physiologically based pharmacokinetic modeling − Part I: Veterinary drugs on human food safety assessment. *Environ Int*, **197**,109339. Accessed 6 July 2025.

25. Chou, W.C. and Lin, Z. (2023) Machine learning and artificial intelligence in physiologically based pharmacokinetic modeling. *Toxicological Sciences*, **191**,1–14. Accessed 6 July 2025.

26. Chou, W.C., Chen, Q., Yuan, L., Cheng, Y.H., He, C., Monteiro-Riviere, N.A., Riviere, J.E., and Lin, Z. (2023) An artificial intelligence-assisted physiologically-based pharmacokinetic model to predict nanoparticle delivery to tumors in mice. *Journal of Controlled Release*, **361**,53–63. Accessed 6 July 2025.

27. Lin, Z. and Chou, W.C. (2022) Machine Learning and Artificial Intelligence in Toxicological Sciences. *Toxicological Sciences*, **189**,7–19. Accessed 6 July 2025.

28. Wu, L., Xu, J., Thakkar, S., Gray, M., Qu, Y., Li, D., and Tong, W. (2024) A framework enabling LLMs into regulatory environment for transparency and trustworthiness and its application to drug labeling document. *Regulatory Toxicology and Pharmacology*, **149**. Accessed 6 July 2025.

29. Morrison, T.M., Stitzel, J.D., and Levine, S.M. (2023) Modeling and Simulation in Biomedical Engineering: Regulatory Science and Innovation for Advancing Public Health. *Ann Biomed Eng*, **51**,1–5. Accessed 6 July 2025.

30. Brian Buntz. (2024) How the Living Heart Project could transform FDA's approach to cardiac device testing. *Research & Development World*,. Accessed 6 July 2025.

31. Aycock, K.I., Battisti, T., Peterson, A., et al. (2024) Toward trustworthy medical device in silico clinical trials: a hierarchical framework for establishing credibility and strategies for overcoming key challenges. *Front Med (Lausanne)*, **11**,1433372. doi:10.3389/FMED.2024.1433372/BIBTEX.

32. Soni, S., Dudhane, A., Debary, H., et al. (2024) EarthDial: Turning Multi-sensory Earth Observations to Interactive Dialogues. Accessed 6 July 2025.

33. Klein, L., Singh, G., Atluri, A.P., et al. (2024) Biodiversity and forest fragmentation assessment based on aboveground biomass prediction produced by fine-tuned geospatial foundation models. Accessed 6 July 2025.

34. PNNL. (2025) MedINT. Accessed 6 July 2025.

35. Keshavamurthy, R., Pazdernik, K.T., Ham, C., Dixon, S., Erwin, S., and Charles, L.E. (2025) Meeting Global Health Needs via Infectious Disease Forecasting: Development of a Reliable Data-Driven Framework. *JMIR Public Health Surveill*, **11**,e59971. Accessed 6 July 2025.

36. Charles-Smith, L.E., Reynolds, T.L., Cameron, M.A., et al. (2015) Using Social Media for Actionable Disease Surveillance and Outbreak Management: A Systematic Literature Review. *PLoS One*, **10**,e0139701. Accessed 6 July 2025.

37. Knight, E., Balzter, H., Breeze, T.D., et al. (2024) Adapting genetic algorithms for multifunctional landscape decisions: A theoretical case study on wild bees and farmers in the UK. *Methods Ecol Evol*, **15**,2153–2167. Accessed 6 July 2025.

38. Knight, E., Balzter, H., Breeze, T.D., et al. (2024) Adapting genetic algorithms for multifunctional landscape decisions: A theoretical case study on wild bees and farmers in the UK. *Methods Ecol Evol*, **15**,2153–2167. Accessed 6 July 2025.

39. Buddendorf, W.B., Wipfler, L., Beltman, W., Baveco, H., Braakhekke, M.C., Bub, S., Gergs, A., and Schad, T. (2023) Aquatic Risks at the Landscape Scale: A Case Study for Pyrethroid Use in Pome Fruit Orchards in Belgium. *Environ Sci Technol*, **57**,15608–15616. Accessed 6 July 2025.

40. LaLone, C.A., Villeneuve, D.L., Lyons, D., Helgen, H.W., Robinson, S.L., Swintek, J.A., Saari, T.W., and Ankley, G.T. (2016) Sequence alignment to predict across species susceptibility (seqapass): A web-based tool for addressing the challenges of cross-species extrapolation of chemical toxicity. *Toxicological Sciences*, **153**,228–245. Accessed 6 July 2025.

41. Doering, J.A., Lee, S., Kristiansen, K., Evenseth, L., Barron, M.G., Sylte, I., and LaLone, C.A. (2018) In Silico Site-Directed Mutagenesis Informs Species-Specific Predictions of Chemical Susceptibility Derived From the Sequence Alignment to Predict Across Species Susceptibility (SeqAPASS) Tool. *Toxicol Sci*, **166**,131–145. Accessed 6 July 2025.

42. Hartung, T., Whelan, M., Tong, W., and Califf, R.M. (2025) Is regulatory science ready for artificial intelligence? *npj Digital Medicine 2025 8:1*, **8**,1–5. Accessed 6 July 2025.

43. Baran, S.W., Bratcher, N., Dennis, J., et al. (2022) Emerging Role of Translational Digital Biomarkers Within Home Cage Monitoring Technologies in Preclinical Drug Discovery and Development. *Front Behav Neurosci*, **15**. Accessed 6 July 2025.

44. Mamoshina, P., Bueno-Orovio, A., and Rodriguez, B. (2020) Dual Transcriptomic and Molecular Machine Learning Predicts all Major Clinical Forms of Drug Cardiotoxicity. *Front Pharmacol*, **11**,504317. Accessed 6 July 2025.

45. Baran, S.W., Bolin, S.E., Gaburro, S., et al. (2024) Validation framework for in vivo digital measures. *Frontiers in Toxicology*, **6**,1484895. Accessed 6 July 2025.

46. Berridge, B.R., Baran, S.W., Kumar, V., Bratcher-Petersen, N., Ellis, M., Liu, C.N., and Robertson, T.L. (2024) Digitalization of toxicology: improving preclinical to clinical translation. *Frontiers in Toxicology*, **6**. Accessed 6 July 2025.