



## OPEN ACCESS

## EDITED BY

Frank Emmert-Streib,  
Tampere University, Finland

## REVIEWED BY

Adrian J. Green,  
Sciome LLC, United States

## \*CORRESPONDENCE

Thomas Hartung  
✉ Thomas.Hartung@uni-konstanz.de

RECEIVED 21 May 2025

ACCEPTED 27 August 2025

PUBLISHED 10 September 2025

## CITATION

Sillé FCM and Hartung T (2025) AI: the Apollo  
guidance computer of the Exposome  
moonshot.  
*Front. Artif. Intell.* 8:1632520.  
doi: 10.3389/frai.2025.1632520

## COPYRIGHT

© 2025 Sillé and Hartung. This is an  
open-access article distributed under the  
terms of the [Creative Commons Attribution  
License \(CC BY\)](#). The use, distribution or  
reproduction in other forums is permitted,  
provided the original author(s) and the  
copyright owner(s) are credited and that the  
original publication in this journal is cited, in  
accordance with accepted academic  
practice. No use, distribution or reproduction  
is permitted which does not comply with  
these terms.

# AI: the Apollo guidance computer of the Exposome moonshot

Fenna C. M. Sillé<sup>1</sup> and Thomas Hartung<sup>1,2,3\*</sup>

<sup>1</sup>Center for Alternatives to Animal Testing (CAAT), Department of Environmental Health and Engineering, Bloomberg School of Public Health and Whiting School of Engineering, Johns Hopkins University, Baltimore, MD, United States, <sup>2</sup>Doerenkamp-Zbinden-Chair for Evidence-Based Toxicology, Department of Environmental Health and Engineering, Bloomberg School of Public Health and Whiting School of Engineering, Johns Hopkins University, Baltimore, MD, United States, <sup>3</sup>CAAT-Europe, Department of Biology, University of Konstanz, Konstanz, Germany

The Exposome—the totality of environmental exposures across a lifetime—remains one of the most significant challenges in understanding and preventing human disease. Translating its vast, heterogeneous data streams into actionable knowledge requires artificial intelligence (AI) integrated with human-relevant experimental systems. We propose a unifying vision in which Microphysiological Systems (MPS) and multi-omics platforms generate high-quality, context-specific data that iteratively calibrate AI models, enabling the creation of digital twins of organs, individuals, and ultimately populations. This “Exposome Moonshot” parallels the Apollo program in ambition, with MPS as the rocket, multi-omics as the lunar module, and AI as the guidance computer. Early applications demonstrate that deep learning can already outperform canonical animal tests for several toxicological endpoints, while reducing cost and time to decision. Realizing the full potential of Exposome intelligence will require expanding the applicability domain of models, implementing robust data security, and prioritizing transparent, interpretable algorithms. By linking predictive AI with experimental feedback, we can move toward a prevention-driven, personalized paradigm for human health and regulatory science.

## KEYWORDS

artificial intelligence, Exposome, digital twin, microphysiological systems, multi-omics integration6, human Exposome project, predictive toxicology, systems toxicology

## 1 Ignition: why exposomics needs an AI engine

Genomics taught us that mapping static code is only half the story; the Exposome—the moving target of everything we breathe, eat, touch, or worry about—drives most of the remaining disease burden (Wild, 2005; Hartung, 2023c). Complex human *in vitro* systems, also known as Microphysiological Systems (MPS), allow modeling these exposures. Yet exposure effects arrive as an unruly torrent of heterogeneous, high-dimensional data. Turning that torrent into knowledge requires artificial intelligence (AI) in the same way the Apollo spaceflight program needed an onboard guidance computer: “*Mass-spectroscopy is our telescope, MPS our lunar module, and AI the guidance computer that stitches the trajectory.*” Extending the analogy, Microphysiological Systems (MPS) can be viewed as the rocket delivering our mission payload. At the same time, multi-omics technologies function as the lunar module enabling precision landing on specific biological questions.

Recently, we have already fused the terms into Exposome Intelligence (EI = Exposome + AI), calling it the “central tool for making sense of ~omics big data” (Sillé et al., 2024; Hartung, 2025). EI is no longer aspirational: industrial-scale machine learning now integrates untargeted mass-spectrometry, wearables, satellite feeds, and electronic health records in near-real time.

## 2 Pattern-finding at planetary scale

Early proof-points show what happens when deep learning meets safety science. Neural networks trained on 600,000 chemicals already outperform the canonical animal tests for skin sensitization, acute toxicity, mutagenicity, and skin and eye hazards, and screen *thousands* of structures in hours (Luechtefeld et al., 2018; Golden et al., 2021; Walter et al., 2024; Duy and Srisongkram, 2025). In the Implementation Moonshot Project for Alternative Chemical Testing (IMPACT) (Sillé et al., 2024), we combine those predictors with evidence-to-decision frameworks<sup>1</sup> so that regulators can rank hazards and benefits on the same probabilistic scale. Let the algorithms sweat the data so that scientists can sweat the hypotheses.

While AI models show promising performance across diverse toxicological endpoints, their predictive accuracy is ultimately constrained by the applicability domain—defined by the chemical structures, exposure scenarios, and biological contexts represented in their training data. Applicability domain is a concept we introduced earlier (Hartung et al., 2004) for *in vitro* systems, borrowing from the Quantitative Structure-Activity Relationship (QSAR) literature. Still, it is now equally applicable to AI-facilitated New Approach Methods (NAMs), also known as alternatives to animal testing. Current coverage of both chemical space and human-relevant biological responses remains incomplete, particularly for complex mixtures, low-abundance environmental contaminants, and underrepresented population groups. Extrapolation beyond these domains can lead to overconfident or biased predictions (Hartung et al., 2025b). To address these limitations, AI development in exposomics must be coupled to an iterative feedback loop with experimental platforms such as Microphysiological Systems (MPS). In this approach, AI models guide targeted MPS experiments to fill gaps in chemical–biological coverage, while new experimental data are used to recalibrate and extend model applicability. This bidirectional exchange not only improves model robustness and generalizability but also ensures that predictions remain anchored in human-relevant biology, thereby increasing regulatory and clinical confidence in their use.

## 3 Microphysiological systems and their digital twins—the test bed for human digital twins with destination personalization

Fueled by stem cell and sensor technologies (Young et al., 2019), MPS platforms have evolved, which do not only keep single cell types alive and measure cell death, but also replicate aspects of native tissue architecture—such as multi-cellular organization, 3D structure, and barrier function—and physiological functionality, including electrophysiological activity, hormone secretion, and metabolic processing (Roth and Berlin, 2019; Marx et al., 2025; Hartung and Smirnova, 2025). They can act both as human-relevant testbeds and as embodied simulators that refine *in-silico* models on the fly. So, by creating a digital twin of an MPS, running virtual experiments and based on this refining our twin (Smirnova et al., 2018), we learn how

to build the twins for entire humans and then populations. Like this, we are teaching our computers to model organ and whole body responses, so that regulators can finally press ‘quit’ on obsolete animal tests.

In a fully integrated workflow, MPS platforms act as dynamic, human-relevant testbeds that both inform and are informed by AI models. Initial *in silico* predictions, generated from existing chemical–biological data, can be used to prioritize compounds, exposure scenarios, or biological pathways for targeted MPS experimentation. These experiments generate high-content, mechanistically anchored datasets—spanning molecular, cellular, and functional endpoints—which are then fed back into the AI pipeline to refine parameters, extend the applicability domain, and reduce prediction uncertainty. This iterative calibration cycle not only enhances model robustness but also guides the design of subsequent experiments, ensuring that each new data generation step strategically fills knowledge gaps identified by the computational models. Such bidirectional learning aligns with the “systems toxicology” vision, in which experimental and computational tools evolve in concert to progressively approximate human biology while minimizing reliance on animal testing (Smirnova et al., 2018).

The Exposome moonshot ultimately aims at full human digital twins—virtual replicas that integrate genomics, exposomics and clinical trajectories to forecast individual risk and therapy response (De Domenico et al., 2025; Trevena et al., 2024; Gangwal and Lavecchia, 2025). Building such twins depends on advanced modelling plus constant data assimilation, a textbook task for adaptive AI. Key components have been sketched<sup>2</sup>—data fusion, generative modelling, iterative validation. The Exposome Moonshot aims to scale these approaches. If the genome was Apollo 11, the Exposome is Artemis—same audacity, bigger destination.

## 4 AI-driven knowledge creation: faster, cheaper, fairer

Chemical safety testing is traditionally slow, costly and biased toward animal biology. A Human Exposome Project platform uses AI to transform this paradigm (Hartung, 2023a, 2023b; Kleinstreuer and Hartung, 2024). It accelerates data interpretation, slashes costs by automating analysis and reducing animal use, and opens the door to more equitable science by leveraging diverse, real-world datasets—including those from underrepresented populations. An AI-powered Human Exposome Project platform can democratize access to knowledge creation, enabling even low-resource labs and countries to contribute insights. As we expand from chemical safety to exposure sciences and human biomonitoring, AI enables high-throughput, real-time synthesis of data that mirrors how people actually live. Calibrating these outputs against human disease etiology, has the potential to upend medicine as we know it—shifting the gravity from symptom-treatment to prevention and personalization.

<sup>1</sup> Available online at: <https://www.ebtox.org>.

<sup>2</sup> Available online at: <https://digital-strategy.ec.europa.eu/en/policies/virtual-human-twins>.

In practice, implementing this vision requires integration of: (i) high-resolution environmental exposure data (e.g., air pollutants, dietary profiles, occupational hazards); (ii) biological readouts from minimally invasive sampling (e.g., blood, saliva, hair metabolomics); and (iii) contextual data from wearables and geospatial mapping. MPS platforms can incorporate these datasets by simulating relevant exposure mixtures, using donor-specific induced pluripotent stem cell (iPSC) lines to enable personalized risk modeling. Stem cell-derived MPSs allow individual genetic and epigenetic backgrounds to be represented in exposure-response testing.

## 5 Trust, transparency and the hallucination hang-over

Sceptics fear the “black-box” nature of large models. Encouragingly, hallucination rates in leading large language models have dropped from ~9% to 1–3% in the past year alone, the latest release of Chat-GPT-5 claims 0.7%, and explainability toolkits mature alongside. As Eliezer Yudkowsky warned, “The greatest danger of AI is that people conclude too early that they understand it.” Our task as a community is therefore four-fold:

1. Make models legible—through open weights, provenance metadata and causal audits.
2. Make data FAIR by design—so that learning systems evolve under robust version control (Hartung et al., 2025a).
3. Make the scientific community more AI literate to create a workforce for the Human Exposome and beyond.
4. Using the simplest model when possible ensures highest data security. Where model performance is comparable, simpler and more interpretable algorithms should be favored to facilitate regulatory uptake. Moreover, given the sensitive nature of Exposome and personal health data, the highest standards of data security and privacy must be maintained to ensure public trust and foster data sharing.

## 6 A collective flight-plan

Moonshots are team sports. The Exposome Moonshot Forum (Washington DC, 12–15 May 2025) convened AI researchers, exposure and environmental health scientists, regulators, policy-makers, ethicists and the public to co-author a “Declaration toward a Human Exposome Project.” Working groups shall tackled now critical priorities—from exposomics reporting standards and model cards for safety AI to avenues for public–private co-funding. Regulators aren’t just kicking the tires of new approach methods anymore—they are asking for the keys.

## 7 Conclusions: lighting the next-gen engines

AI has progressed to the point where dismissing it as futuristic is no longer tenable. While some AI domains—such as large language models—are currently doubling in specific performance benchmarks every few months, domains such as exposure science, genomics, and chemistry advance at different rates due to data generation constraints. By the time this reaches print, today’s algorithms will already look

vintage. What matters is how quickly we embed them in transparent, ethical and human-centric frameworks.

“You will never see an AI as bad as today’s—tomorrow’s will be twice as smart and half as hallucinogenic,” we told journalists recently at an Science Media Centre (SMC) London media briefing. The same exponential curve that propels model capability can, if we steer wisely, propel the Exposome moonshot from concept to clinical and regulatory reality. So let us strap in, check the boosters, and light the engines—Exposome intelligence is ready for lift-off.

## Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author/s.

## Author contributions

FS: Writing – review & editing. TH: Writing – original draft.

## Funding

The author(s) declare that financial support was received for the research and/or publication of this article. This work is supported by the NEXUS (Network for EXposomics in the United States), a Center for Exposome Research Coordination (CERC), supported by the National Institutes of Environmental Health Sciences (NIEHS) under grant number U24ES036819 is gratefully appreciated.

## Acknowledgments

We also want to acknowledge a 2024 Johns Hopkins Bloomberg Center Nexus Convening award which funded the Exposome Moonshot Forum in May 2025.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The author(s) declared that they were an editorial board member of Frontiers, at the time of submission. This had no impact on the peer review process and the final decision.

## Generative AI statement

The authors declare that no Gen AI was used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated

organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

## References

- De Domenico, M., Allegrì, L., Caldarelli, G., d'Andrea, V., Di Camillo, B., Rocha, L. M., et al. (2025). Challenges and opportunities for digital twins in precision medicine from a complex systems perspective. *NPJ Digit. Med.* 8:37. doi: 10.1038/s41746-024-01402-3
- Duy, H. A., and Srisongkram, T. (2025). Bidirectional long short-term memory (BiLSTM) neural networks with conjoint fingerprints: application in predicting skin-sensitizing agents in natural compounds. *J. Chem. Inf. Model.* 65, 3035–3047. doi: 10.1021/acs.jcim.5c00032
- Gangwal, A., and Lavecchia, A. (2025). Artificial intelligence in preclinical research: enhancing digital twins and organ-on-chip to reduce animal testing. *Drug Discov. Today* 30:104360. doi: 10.1016/j.drudis.2025.104360
- Golden, E., Macmillan, D. S., Dameron, G., Kern, P., Hartung, T., and Maertens, A. (2021). Evaluation of the global performance of eight in silico skin sensitization models using human data. *ALTEX* 38, 33–48. doi: 10.14573/altex.1911261
- Hartung, T. (2023a). A call for a human Exposome project. *ALTEX* 40, 4–33. doi: 10.14573/altex.2301061
- Hartung, T. (2023b). AI as the new frontier in chemical risk assessment. *Front. AI. Sec. Med. Pub. Health* 6:1269932. doi: 10.3389/frai.2023.1269932
- Hartung, T. (2023c). ToxAIcology - the evolving role of artificial intelligence in advancing toxicology and modernizing regulatory science. *ALTEX* 40, 559–570. doi: 10.14573/altex.2309191
- Hartung, T. (2025). How AI can deliver the human Exposome project. *Nat. Med.* 31:1738. doi: 10.1038/s41591-025-03749-w
- Hartung, T., Bremer, S., Casati, S., Coecke, S., Corvi, R., Fortaner, S., et al. (2004). A modular approach to the ECVAM principles on test validity. *Altern. Lab. Anim* 32, 467–472. doi: 10.1177/026119290403200503
- Hartung, T., and Smirnova, L. (2025). A path forward advancing microphysiological systems. *ALTEX* 42, 183–203. doi: 10.14573/altex.2504091
- Hartung, T., Whelan, Tong, W., and Califf, R. M. (2025a). Is regulatory science ready for artificial intelligence? *NPJ Digit. Med.* 8:200. doi: 10.1038/s41746-025-01596-0
- Hartung, T., Hoffmann, S., and Whaley, P. (2025b). Assessing risk of bias in toxicological studies in the era of artificial intelligence. *Archives in Toxicology* 99, 3065–3090. doi: 10.1007/s00204-025-03978-5
- Kleinstreuer, N., and Hartung, T. (2024). Artificial intelligence (AI) – it's the end of the tox as we know it (and I feel fine) - AI for predictive toxicology. *Arch. Toxicol.* 98, 735–754. doi: 10.1007/s00204-023-03666-2
- Luechtefeld, T., Marsh, D., Rowlands, C., and Hartung, T. (2018). Machine learning of toxicological big data enables read-across structure activity relationships (RASAR) outperforming animal test reproducibility. *Toxicological Sciences* 165, 198–212. doi: 10.1093/toxsci/kfy152
- Marx, U., Beken, S., Chen, Z., Dehne, E.-M., Doherty, A., Ewart, L., et al. (2025). Biology-inspired dynamic microphysiological system approaches to revolutionize basic research, Healthcare and Animal Welfare. *ALTEX* 42, 204–223. doi: 10.14573/altex.2410112
- Roth, A., and Berlin, M. P. S.-W. S. (2019). Human microphysiological systems for drug development. *Science* 373, 1304–1306. doi: 10.1126/science.abc3734
- Sillé, F. C. M., Busquet, F., Fitzpatrick, S., Herrmann, K., Leenhouts-Martin, L., Luechtefeld, T., et al. (2024). The implementation Moonshot project for alternative chemical testing (IMPACT) toward a human Exposome project. *ALTEX* 41, 344–362. doi: 10.14573/altex.2407081
- Smirnova, L., Kleinstreuer, N., Corvi, R., Levchenko, A., Fitzpatrick, S. C., and Hartung, T. (2018). 3S – systematic, systemic, and systems biology and toxicology. *ALTEX* 35, 139–162. doi: 10.14573/altex.1804051
- Trevena, W., Zhong, X., Lal, A., Rovati, L., Cubro, E., Dong, Y., et al. (2024). Model-driven engineering for digital twins: a graph model-based patient simulation application. *Front. Physiol.* 15:1424931. doi: 10.3389/fphys.2024.1424931
- Walter, M., Webb, S. J., and Gillet, V. J. (2024). Interpreting neural network models for toxicity prediction by extracting learned chemical features. *J. Chem. Inf. Model.* 64, 3670–3688. doi: 10.1021/acs.jcim.4c00127
- Wild, C. P. (2005). Complementing the genome with an "Exposome": the outstanding challenge of environmental exposure measurement in molecular epidemiology. *Cancer Epidemiol. Biomarkers Prev.* 14, 1847–1850. doi: 10.1158/1055-9965.EPI-05-0456
- Young, A. T., Rivera, K. R., Erb, P. D., and Daniele, M. A. (2019). Monitoring of microphysiological systems: integrating sensors and real-time data analysis toward autonomous decision-making. *ACS Sens.* 4, 1454–1464. doi: 10.1021/acssensors.8b01549