AI to Learn (AI2L): Guidelines and Practice for Human-Centered AI Utilization as a Learning Support Tool—Four Pillars of Black-Box Elimination, Accountability, Information Protection, and Energy Efficiency

Seine A. Shintani*

1) Department of Biomedical Sciences, College of Life and Health Sciences,
Chubu University, Kasugai, Aichi, 487-8501, Japan

2) Center for Mathematical Science and Artificial Intelligence,
Chubu University, Kasugai, Aichi, 487-8501, Japan

3) Institute for Advanced Research, Nagoya University,
Furo-cho, Chikusa-ku, Nagoya-shi, 464-8601, Japan

Abstract

Contemporary generative AI—especially large language models (LLMs)—is rapidly permeating diverse domains such as research, education, and healthcare owing to its remarkable efficiency and expressive power. Conversely, AI systems bring serious challenges: their black-box nature, the risk of privacy leakage from input data, ethical concerns arising from outputs whose rationale is opaque, and the substantial energy consumption and environmental burden associated with large-scale deployment. This paper proposes AI to Learn (AI2L), a set of guidelines that deliberately limits AI to a learning-support role for humans and eliminates any black-box components from the final deliverables. AI2L rests on four principles: (1) humans retain ultimate decision-making authority; (2) human verification ensures accountability for AI outputs; (3) the risk of information leakage is rigorously minimized; and (4) AI usage is managed for energy efficiency and long-term sustainability. We examine several concrete implementations of AI2L—including Grad CAM–based image interpretation, the discovery

*Corresponding author. E-mail: s-shintani@fsc.chubu.ac.jp

of novel insights via symbolic regression, the development of AI-generated yet human-auditable code, and reversible anonymization for data protection—and analyze them from both practical and theoretical perspectives. Recent studies showing that foundation models fail to grasp underlying physical laws, despite high predictive accuracy, further underscore the necessity of AI2L's approach. By acknowledging AI's limitations and hazards while harnessing its strengths, AI2L provides a robust framework for ethical, sustainable, and human-centered integration of AI into society.

1 Introduction

Advances in artificial intelligence technology—most notably the rise of large-scale generative AI—are transforming the ways in which society conducts research, education, and creative work [1, 2]. New applications in natural language processing, image recognition, and programming assistance are emerging almost daily, fundamentally reshaping how knowledge is produced and disseminated [1]. Yet the very convenience that allows users to feed vast amounts of unpublished or sensitive data into these systems conceals two intrinsic dangers: (i) once submitted, the information may effectively reside in the cloud beyond institutional oversight, and (ii) the opaque decision processes of AI models make it difficult for humans to detect misinformation or bias embedded in the outputs [3, 4]. Such risks are unacceptable in high-stakes areas such as medicine and research, where incidents of data leakage and ethical violations have already been reported worldwide [3, 4]. Moreover, training and deploying LLMs and other large models demand enormous computational resources and electricity; recent studies document a steep rise in the carbon footprint of deep learning research as a whole [5, 6]. Against this backdrop, Schwartz et al. introduced the concept of "Green AI," arguing that research evaluation should prioritize efficiency (energy per unit of computation) alongside conventional performance metrics [6]. In other words, real-world adoption of AI must balance accuracy with environmental responsibility and other societal constraints.

Rather than advocating the unrealistic option of "living without AI," the present study proposes the AI to Learn (AI2L) paradigm: AI is leveraged as a learning companion under human control, its use is minimized to conserve computational resources, and all traces of black-box components are removed from the final deliverables. We examine both the theoretical foundations and practical value of this approach, demonstrating how AI2L reconciles the power of generative models with the demands of transparency, privacy, and sustainability.

2 Risks and Limitations of Conventional Generative-AI Use

2.1 Data Leakage and Security Concerns

Submitting confidential information directly to an Internet-connected generative AI service—such as ChatGPT—exposes corporate secrets and personal data, previously guarded under strict internal controls, to third-party cloud providers [2]. Indeed, several organizations have already restricted in-house use of such services, and legal debates have arisen under the GDPR and Japan's Act on the Protection of Personal Information [7]. Because model operators may reuse user inputs for continual training or service improvement, the data are irreversibly absorbed into the model's parameter space, making subsequent leakage or misuse difficult to detect or remediate [2]. Persisting in the use of generative AI while ignoring this fundamental risk is indefensible both from an information governance perspective and from the standpoint of broader social responsibility [3].

2.2 Black-Box Opacity and the Accountability Gap

Most contemporary AI systems, including those based on deep learning, cannot explain why they produce a given output. Techniques such as Grad CAM [8] and SHAP [9] offer partial avenues toward explainable AI (XAI), yet many scholars argue that AI models whose inner workings elude human oversight should not be entrusted with high-stakes decisions [10]. Recent work on foundation models reinforces this caution: Vafa et al. showed that a model able to predict planetary orbits with 99.99% accuracy nevertheless failed to recover the underlying Newtonian laws of motion [11]. In other words, the system mastered pattern recognition but could not extract or generalize the governing principles. Such findings highlight the peril of adopting AI outputs at face value and underscore the need for rigorous human vetting before those outputs are incorporated into any final product.

2.3 Bias, Ethical Pitfalls, and Environmental Burden

Generative AI presents a two-fold challenge: it can reproduce the historical and societal biases embedded in its training data while simultaneously amplifying global energy consumption and environmental impact [3]. Because such systems readily inherit prejudices from their data sets, they may generate discriminatory language or make unfair decisions; in high-stakes domains, opaque outputs of this kind can lead to serious ethical violations [10]. Strubell et al. warned that training a BERT-scale language model once—augmented with neural architecture search—can emit roughly 284 metric tons of CO_2, exacerbating the resource gap between academia and industry [5]. In response, Schwartz et al. introduced the notion of "Green AI," arguing that, given comparable performance, researchers should favor models with superior energy efficiency [6].

AI2L operationalizes this principle by (i) restricting AI usage to the learning-support phase, (ii) distilling the results into lightweight, human-readable code, and (iii) eliminating cloud-based inference, thereby slashing power consumption during deployment. For example, the seating-chart workflow described in this paper relies on Python code generated once by ChatGPT; thereafter the task runs locally with zero GPU resources and no further calls to generative AI services. Thus, AI2L constitutes a practical framework that embodies the ideals of Green AI while maintaining transparency and accountability.

2.4 Human–AI Collaboration for Discovery and Accelerated Learning

Although large language models achieve strikingly accurate predictions, their inability to internalize deep physical laws dampens hopes that AI can operate as a fully autonomous scientist. From the AI2L perspective, however, this limitation merely clarifies the optimal division of labor between machines and humans. Vafa et al.'s "inductive bias probe" demonstrated that a foundation model trained on orbital data could not generalize Newtonian mechanics [11]. In the symbolic regression case study presented in Section 4.2, an AI system proposed several nonlinear candidate equations [12, 13]; a human researcher scrutinized these suggestions and ultimately refined them into the homeochaos concept for cardiac rhythms [14, 15]. Here, AI rapidly canvassed the hypothesis space, while theory building and generalization remained a human responsibility—thereby offsetting AI's weaknesses and accelerating scientific discovery. Within the AI2L framework, AI serves as a "hypothesis generator" or "pattern extractor," allowing investigators to focus on *why* the patterns arise. The complementary strengths of machine prediction and human interpretation thus combine to improve both the speed and the quality of law-seeking research.

3 The AI to Learn (AI2L) Framework: Concept and Positioning

3.1 Definition and Philosophy

AI to Learn (AI2L) treats artificial intelligence as a temporary assistant that augments human learning and creativity; its guiding rule is that no black-box AI component remains in the final deliverables—whether algorithms, research papers, teaching materials, or production code [10]. Here, removal means that the finished system or knowledge artifact no longer depends on hidden model weights or external cloud APIs: humans must be able to understand and explain every structural element and rationale.

Traditional Human-in-the-Loop (HIL) design inserts people into the training or inference loop for labeling and feedback, improving performance and adding a safety valve; yet in practice the final decision or execution often rests with the AI system [16]. AI2L diverges fundamentally by limiting AI to a catalyst for thought. The machine rapidly explores vast hypothesis spaces and proposes candidate formulas or designs, but humans perform the truth testing, theory building, and social accountability—ultimately excising the AI black-box from the finished work. This approach strengthens accountability while minimizing both data leakage risk and environmental impact [2].

AI2L is anchored in four pillars:

1. Human sovereignty over final decisions.
2. Mandatory explainability of all retained components.
3. Rigorous prevention of information leakage.
4. Energy efficiency and sustainability.

These pillars operationalize the principles of human oversight, accountability, and sustainability emphasized in the EU AI Act and the NIST AI Risk Management Framework [17]. The fourth pillar is particularly significant: by distilling outputs into lightweight, human-readable programs and avoiding large model inference at deployment, AI2L offers a concrete, Green AI solution that sharply reduces computational resources and CO_2 emissions.

3.2 AI2L Guidelines (Four Principles)

1. **Learning support only**: AI is confined to tasks that assist human learning; it is never allowed to complete deliverables autonomously.
2. **Human verification and accountability**: All AI outputs undergo visualization, logical checking, and reproducibility testing; humans bear ultimate responsibility for validity and safety.
3. **Privacy and security first**: Both inputs and outputs are anonymized or processed locally; confidential data are never uploaded to cloud-based AI services.
4. **Energy-aware sustainability**: Large models are used solely during the learning-support phase; the operational stage relies on human-readable, lightweight code.

4 AI2L Case Studies

4.1 A Grad-CAM Success Story: Detecting Under-polished Regions on Titanium Plates

We built a convolutional neural network (CNN) to classify the surface condition of titanium plates—adequately versus inadequately polished—and used Grad CAM to visualize the evidence for each decision (Fig. 1) [8]. The heat map overlays consistently highlighted microscopic dark regions of residual titanium oxide in images labeled "under polished," precisely the areas that experienced inspectors had been focusing on tacitly. This enabled us to codify a clear quality control rule—"presence of dark TiO_2 residue $\rightarrow$ insufficient polishing"—and to develop a short training module through which non-experts could master the skill in minutes.

From the AI2L standpoint, the AI system handled only two tasks during the learning-support phase: (i) training the polishing quality classifier and (ii) externalizing tacit knowledge via Grad CAM. Routine inspection was then shifted to a manual procedure in which humans directly look for the highlighted dark regions; the AI model was removed from the operational loop, eliminating black-box dependence. What remains on the shop floor is the human understanding—"black oxide spots signal under polishing"—and a rapid manual check, while the AI served merely as a short-term accelerator. This case neatly illustrates the AI2L principle of leveraging AI's visualization power to transform tacit expertise into explicit, shareable knowledge and then reverting to a lightweight, human-centered workflow.

4.2 New Insights Uncovered via Symbolic Regression

Cardiac sarcomeres exhibit self-sustained oscillations—termed HSOs (Hyperthermal Sarcomeric Oscillations)—when locally warmed; their frequency remains stable at $\approx 7\,\mathrm{Hz}$, close to the heart rate, yet their amplitude waveforms display puzzling fluctuations [18, 19]. To probe this unresolved behavior, we fed length-trace data from multiple consecutive sarcomeres, obtained with SL nanometry, into physics-oriented symbolic regression pipelines. We employed both the free-form equation search algorithm of Schmidt & Lipson [12] and AI Feynman [13], which couples neural networks with analytical heuristics, to harvest candidate formulas.

Although none of the AI-generated expressions constituted an explicit chaos model, several terms simultaneously hinted at amplitude self-amplification and inhibitory coupling between neighboring sarcomeres. Guided by this structure, we conducted additional recurrence plot analyses and computed the maximal Lyapunov exponent, revealing a two-tiered dynamics: each sarcomere possesses a positive Lyapunov exponent (chaotic instability) while mutual interactions enforce a stable oscillation period [14]. This phenomenon was documented as Contraction Rhythm Homeostasis in [14]. Subsequent experiments and model refinements led us to propose the broader concept of Chaordic Homeodynamics—the cooperation of

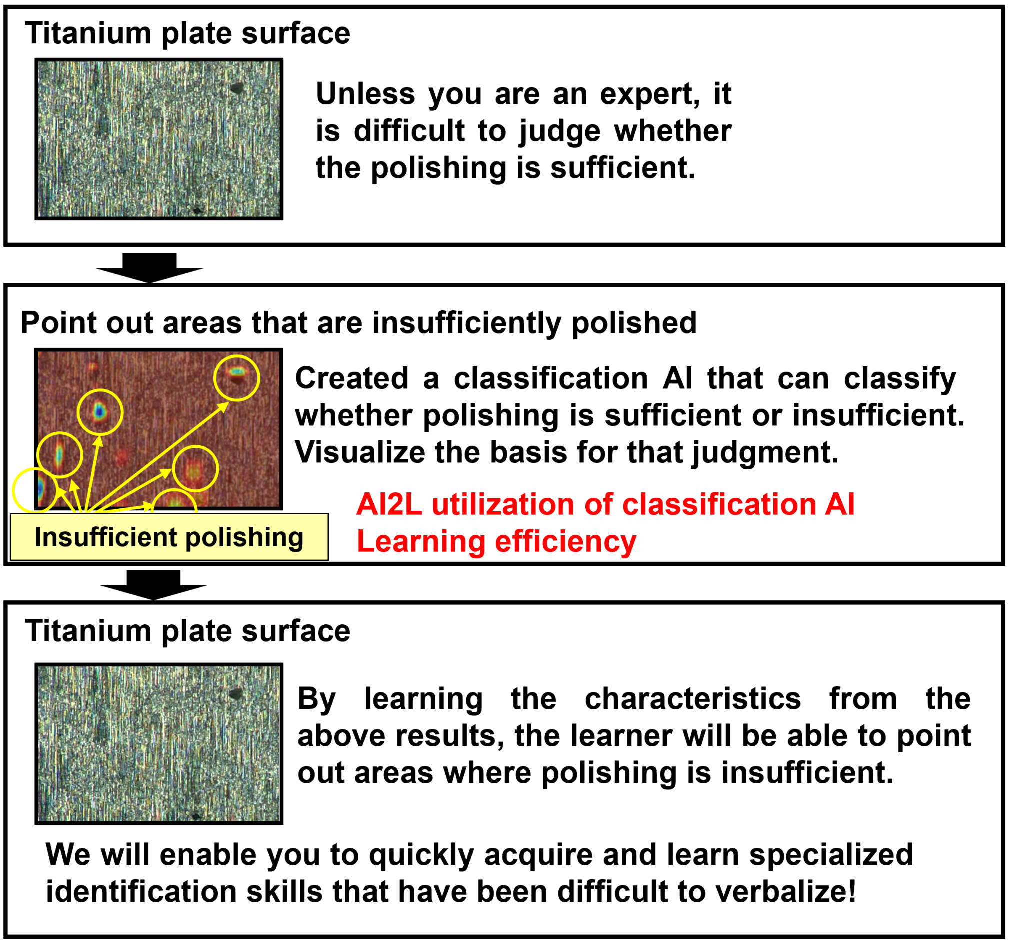

Figure 1: Grad-CAM visualization of the titanium-polish classifier. Top: CNN prediction alongside the Grad-CAM heat map; the yellow circle marks the dark oxide region that triggered the "under polished" label. Bottom: When shown as instructional material, the visualization allows learners to grasp in a short time the expert cue—dark regions corresponding to residual TiO_2—that was previously difficult to articulate.

periodicity and chaos in maintaining homeostasis—detailed in [15].

Crucially, in line with AI2L principles, the AI's role was confined to "suggesting candidate equations," whereas the final validation and theorizing were performed by humans. Because no black-box component remains in the deliverable, the readable equations narrowed the exploratory search space and dramatically shortened the experiment–analysis cycle. By interrogating—rather than blindly adopting—the AI-proposed terms, we exposed a previously hidden chaotic property within HSOs. Ongoing work is leveraging this discovery for ultra-early diagnostics of heart failure and for next-generation pacemaker control strategies, demonstrating how AI2L can act as a knowledge accelerator from basic science to clinical application.

4.3 AI-Generated Code, Human-Owned Asset: a Seating-Chart Workflow Fully Aligned with AI2L

The seemingly mundane task of assigning more than 120 student IDs to an eight-column seating chart typically requires at least thirty minutes and is prone to copy-and-paste errors. Applying AI2L principles, we compressed the entire workflow to 17 seconds without uploading any confidential data to a generative AI service. The process comprised four stages (Fig. 2) [3].

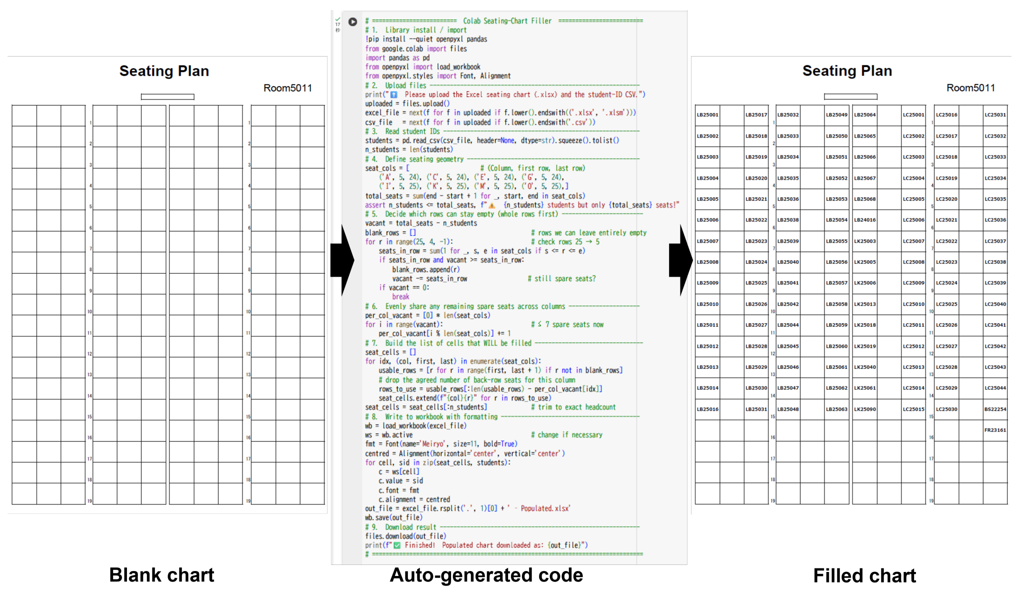

Figure 2: AI-assisted yet human-owned seating-chart workflow. Left: blank Excel template (Room 5011). Centre: excerpt of the Python script auto-generated by ChatGPT from a single prompt. Right: completed chart produced in 17 s; the algorithm leaves the rear rows empty by evenly distributing surplus seats.

1. **One-shot code generation.** A single English prompt (centre panel, Fig. 2) was

submitted to ChatGPT o3 pro. Within thirty seconds the model produced an open-source Python script that (i) reads a headerless CSV of student IDs, (ii) maps the designated Excel seat ranges, (iii) evenly distributes surplus seats to the back rows, and (iv) formats every entry in Meiryo 11 pt, bold, and center aligned.

2. **Human review and hardening.** The script was downloaded, variable names were refactored, and explicit `assert` statements plus exception handling were added. This audit erased the last vestige of black-box dependence, turning the code into a fully transparent asset.

3. **Execution in a controlled cloud sandbox.** For convenience, the hardened script was executed on Google Colab, but only the blank Excel template and an anonymized ID list—information acceptable for external storage—were uploaded. No generative AI inference occurs at runtime; Colab functions solely as a commodity Python interpreter, so the model weights of ChatGPT are never exposed to the data and associated privacy risks are eliminated [3].

4. **Download and local reuse.** The populated seating chart (right panel, Fig. 2) was downloaded and archived. The identical script can subsequently be run on an on-premise PC—without internet connectivity or GPU acceleration—incurring negligible energy cost and zero future dependence on ChatGPT or any other large model [5].

This workflow epitomizes the balanced ethos of AI2L: a generative model is leveraged once to accelerate development, after which day-to-day operations proceed independently of large-scale AI. The outcome is a lightweight, auditable, and easily modifiable script that fosters user autonomy while minimizing cloud exposure and carbon footprint [6].

4.4 Reconciling Data Anonymisation with Generative-AI Use—An AI2L Three-Step Protocol

First and foremost, many educational and research tasks do not require uploading real data at all. If the inputs supplied to a generative AI service are replaced with structurally similar dummy data, model prototyping and code generation proceed perfectly well. For instance, the seating-chart script described above was fully validated with a CSV containing random IDs; no confidential file ever left the local environment.

There are, however, situations—such as grading student essays with highly variable formatting—where one might wish to process genuine submissions through AI. At this point, AI2L introduces reversible anonymization. Student numbers and names are replaced en bloc with either Fernet-based symmetric encryption or tokenization, ensuring that the file transmitted to the external service contains nothing that a human reader could trace back to an individual. The encryption key and mapping table remain in an offline lab computer; scores or feedback returned by the AI are decrypted locally. In this way, automated assessment is achieved without compromising personal privacy.

The safe level of anonymization is context dependent. Even after removing direct identifiers, free-text answers may still reveal personal details [20]. Therefore, before any upload the following checks are mandatory:

1. Encrypt all direct identifiers.
2. Generalize quasi-identifiers (e.g., replace "Department of Bioengineering, Year 3" with "STEM, Upper division").
3. Scan free-text fields for residual re-identification risks.

Only when residual risk falls below the operational threshold—and a human explicitly confirms this—may the data be sent. AI2L emphasizes that ultimate responsibility lies with humans, including the choice of anonymization depth.

In summary, AI2L prescribes a three-layer guardrail:

1. Avoid attachments altogether whenever possible.
2. Substitute with dummy data for development and testing.
3. When real data are unavoidable, apply reversible anonymization and keep decryption strictly local.

This protocol allows practitioners to benefit from generative AI while preventing data leakage and simultaneously cultivating higher data governance literacy—a concrete, safe, and sustainable pathway for AI adoption that AI2L explicitly endorses.

5 Theoretical and Societal Significance

5.1 True Understanding versus Pattern Recognition—AI2L's Collaborative Model

The impressive predictive power of large models is often misconstrued as proof that "AI understands the world." Yet Yildirim et al.'s probe shows decisively that a system able to predict orbital positions with 99.99% accuracy still fails to recover Newtonian mechanics; high accuracy prediction and deep understanding are not the same [11]. Bridging that gap requires human critical thinking and theory building [10]. AI2L assigns AI the role of scanning the vast hypothesis space, while humans verify, interpret, and generalize the results into universal laws. In this scheme, pattern recognition—the "Keplerian" level of knowledge—serves as a springboard for humans to leap toward explanatory principles—the "Newtonian" level. This division of labor captures AI2L's scholarly value.

5.2 A Unified Demand for AI Ethics, Governance, and Energy Efficiency

International frameworks such as the EU AI Act and the NIST AI RMF designate human oversight, accountability, and sustainability as indispensable [17]. In practice, however, operational models that simultaneously suppress (a) black-box dependence, (b) personal data leakage, and (c) massive power consumption are rare. AI2L offers a three-in-one protocol:

1. **Explainability**: All AI outputs are distilled into human-readable formulas or code; no external service remains in the production pipeline [10].
2. **Information protection**: Dummy data and reversible anonymization keep re-identification risk to a minimum [3].
3. **Green AI implementation**: Generative models are used only during the learning-support phase, while routine tasks run locally in lightweight form [5, 6].

By meeting these three goals at once, AI2L uniquely realizes an AI-specific triple bottom line, integrating ethical, legal, and environmental imperatives into a single, practical methodology.

6 Challenges and Future Directions for AI2L Adoption

Because AI2L deliberately limits AI autonomy and assigns ultimate responsibility to humans, it imposes short-term costs—such as (i) the human effort required for code audits and formula verification, and (ii) the relinquishment of some automated pipeline features offered by AI APIs. Yet these costs yield long-term returns in the form of quality assurance, risk reduction, and enhanced researcher skill sets. Key practical challenges and prospects include:

1. **Feedback loops from real-world practice**. Domain-specific artifacts—such as anonymization templates for medical records or grading rubrics for education—should be shared in open repositories so that AI2L workflows can evolve through continuous community feedback.
2. **Standardizing evaluation metrics**. Quantitative benchmarks are needed: XAI scores based on Grad CAM or SHAP [8, 9], energy indicators such as watt hours per inference, and other measures that allow the "degree of AI2L compliance" to be assessed objectively.
3. **Horizontal expansion to other sectors**. While this paper focused on materials inspection, biophysics, and educational administration, the benefits of AI2L are even greater in fields where human accountability is paramount—government document review, legal support, automotive maintenance, and more. Cross-disciplinary workshops

should foster the co-development of common modules such as anonymization toolkits and lightweight inference frameworks.

4. **Integration into policy and education**. University curricula and corporate training programs should incorporate "AI2L practicum" modules, enabling researchers and engineers to master AI while managing its limits in a systematic way.

In building a sustainable AI society, AI2L offers a healthy counterbalance to the myth of "fully autonomous AI." By uniting black-box elimination, accountability, data protection, and energy conservation, the framework is poised to become the new default in near-future AI governance and education.

Acknowledgements

All experiments and manuscript writing were solely performed by the author. This work was supported by JSPS KAKENHI Grant Number JP25K00269 (Grant-in-Aid for Scientific Research (C), Project Title: "Elucidation of Myosin Molecular Dynamics Associated with Sarcomere Morphological Changes in the Intracellular Environment").

References

[1] R. Bommasani, D. A. Hudson, E. Adeli, *et al.*, "On the Opportunities and Risks of Foundation Models," *arXiv preprint* arXiv:2108.07258, 2021.

[2] OpenAI, "GPT-4 Technical Report," *arXiv preprint* arXiv:2303.08774, 2023.

[3] N. Carlini, F. Tramèr, E. Wallace, *et al.*, "Extracting Training Data from Large Language Models," in *Proc. 30th USENIX Security Symp.*, 2021, pp. 2633–2650.

[4] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" in *Proc. ACM Conf. Fairness, Accountability, and Transparency (FAccT)*, 2021, pp. 610–623.

[5] E. Strubell, A. Ganesh, and A. McCallum, "Energy and Policy Considerations for Deep Learning in NLP," in *Proc. 57th Annu. Meeting Assoc. Comput. Ling. (ACL)*, 2019, pp. 3645–3650.

[6] R. Schwartz, J. Dodge, N. Smith, and O. Etzioni, "Green AI," *Communications of the ACM*, vol. 63, no. 12, pp. 54–63, 2020.

[7] T. Kuru, "Lawfulness of the Mass Processing of Publicly Accessible Online Data to Train Large Language Models," *International Data Privacy Law*, Oct. 2024.

[8] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," in *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, 2017, pp. 618–626.

[9] S. M. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," in *Advances in Neural Information Processing Systems 30 (NeurIPS)*, 2017, pp. 4765–4774.

[10] C. Rudin, "Stop Explaining Black Box Machine Learning Models for High-Stakes Decisions and Use Interpretable Models Instead," *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206–215, 2019.

[11] K. Vafa, P. G. Chang, A. Rambachan, and S. Mullainathan, "What Has a Foundation Model Found? Using Inductive Bias to Probe for World Models," *arXiv preprint* arXiv:2507.06952, Jul. 2025, to appear in *Proc. 42nd Int. Conf. Machine Learning (ICML)*.

[12] M. Schmidt and H. Lipson, "Distilling Free-Form Natural Laws from Experimental Data," *Science*, vol. 324, no. 5923, pp. 81–85, Apr. 2009.

[13] S.-M. Udrescu and M. Tegmark, "AI Feynman: A Physics-Inspired Method for Symbolic Regression," *Science Advances*, vol. 6, no. 16, eaay2631, Apr. 2020.

[14] S. A. Shintani, "Hyperthermal Sarcomeric Oscillations Generated in Warmed Cardiomyocytes Control Amplitudes with Chaotic Properties While Keeping Cycles Constant," *Biochem. Biophys. Res. Commun.*, vol. 611, pp. 8–13, 2022.

[15] S. A. Shintani, "Chaordic Homeodynamics: The Periodic Chaos Phenomenon Observed at the Sarcomere Level and its Physiological Significance," *Biochem. Biophys. Res. Commun.*, vol. 760, 151712, 2025.

[16] S. Amershi, M. Cakmak, W. B. Knox, and T. Kulesza, "Power to the People: The Role of Humans in Interactive Machine Learning," *AI Magazine*, vol. 35, no. 4, pp. 105–120, 2014.

[17] E. Tabassi, "Artificial Intelligence Risk Management Framework (AI RMF) 1.0," National Institute of Standards and Technology, Gaithersburg, MD, USA, NIST AI 100-1, Jan. 2023.

[18] S. A. Shintani, K. Oyama, N. Fukuda, and S. Ishiwata, "High-Frequency Sarcomeric Auto-Oscillations Induced by Heating in Living Neonatal Cardiomyocytes of the Rat," *Biochem. Biophys. Res. Commun.*, vol. 457, no. 2, pp. 165–170, 2015.

[19] S. A. Shintani, T. Washio, and H. Higuchi, "Mechanism of Contraction Rhythm Homeostasis for Hyperthermal Sarcomeric Oscillations of Neonatal Cardiomyocytes," *Scientific Reports*, vol. 10, 20468, 2020.

[20] L. Sweeney, "k-Anonymity: A Model for Protecting Privacy," *Int. J. Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 5, pp. 557–570, 2002.