© JAN 2024 | IRE Journals | Volume 7 Issue 7 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I5-1711956

Human-in-the-Loop Machine Learning Systems

GEETHA ARADHYULA
Zolon Tech Inc

Abstract: Human-in-the-loop (HITL) machine learning systems is a paradigm change of the previous, fully automated based models to collaborative intelligence, where human expertise meets with machine efficiency to deliver higher performance. Human input features in these systems are implemented at different levels of the machine learning pipeline such as data collection, annotation, feature engineering, model training, evaluation, and post-deployment monitoring to make sure that learning processes are directed by domain knowledge, moral consciousness, and contextual awareness. HITL systems include human feedback to overcome the major issues of data scarcity, label noise, algorithmic bias, and interpretability. In active learning systems, human beings give precise corrective labels or feedback to doubtful model predictions so that the system can learn efficiently with minimal data. Such systems are especially useful in areas with high stakes, such as healthcare, finance, security, and autonomous systems because this dynamic interaction does not only help to make models more accurate, but also more transparent and trusting. HITL methods have been further enabled by recent breakthroughs in explainable AI, interactive visualization, and reinforcement learning from human feedback (RLHF) to make machines act in ways that are consistent with human values and goals. In addition, the introduction of collaborative systems of crowd-sourcing data annotation and model auditing indicates the scalability and flexibility of HITL systems in practice. Finally, human-in-the-loop machine learning will help to close the boundary between artificial and human intelligence, providing adaptive, ethical, and responsible AI solutions. It also puts humans not as observers but as agents in an ongoing learning process, which keeps AI systems resilient, equitable, and sensitive to the needs of the ever-changing society.

I. INTRODUCTION

Definition of Human-in-the-Loop (HITL) Systems

HITL systems are machine learning (ML) systems that involves active human input into a single or multiple phases of the model development and decision-making process. Instead of just subcontracting the computation to an automated system, HITL systems enable human professionals to steer, correct or refine model behavior. Human intervention may happen at any of the following steps: data labeling, feature selection, and model training, evaluation, and deployment, making sure that the process of learning involves domain knowledge, moral judgment, and local awareness. This is a hybrid cooperation, which allows the machines not only to learn but also to learn by human reasoning and experience.

Overview of Traditional Machine Learning Pipelines

The conventional pipelines of machine learning are mainly data-driven and completely automated. These are generally data collection, preprocess, feature extraction, model training, evaluation and deployment without human feedback during the iterative learning. After a model is trained, it is autonomous and it makes predictions based on statistical trends among the data. Although these systems are very effective in a task where the tasks are defined and the data are extensive, rich, and of high quality, they fail to perform in situations where there is ambiguity, bias or dynamic environment where human intuition and flexibility are needed.

Motivation for Incorporating Human Expertise

The integration of human knowledge resolves a number of constraints that automation systems have. Humans are able to identify and address data biases, justify unconfident predictions, and offer insights which are hard to encode algorithmically. HITL systems are able to improve model reliability, interpretability, and ethical compliance by introducing feedback loops. It is especially useful in the case of applications that require critical consequences of errors, like medical diagnosis, autonomous driving, or legal decision-making.

Importance and Relevance in Modern AI Applications

Due to the higher penetration of AI technologies into daily life, there is an increase in the demand for reliable, transparent, and adaptable systems. HITL methods provide the assurance of maintaining the alignment of AI with human values and societal norms through the introduction of constant monitoring and interpretive logic. The synergy between machine intelligence and human cognition is useful in what is known as modern AI applications: conversational agents, to critical infrastructure

© JAN 2024 | IRE Journals | Volume 7 Issue 7 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I5-1711956

monitoring. The HITL systems, therefore, will be a critical move towards the establishment of responsible and human-oriented artificial intelligence.

II. CONCEPT AND PRINCIPLES

Combining Human Judgment with Automated Algorithms

The main idea of the Human-in-the-Loop (HITL) systems is that human judgment and automated machine-learn algorithms are merged into one framework. Algorithms can be useful at the level of processing a large amount of data and detecting statistical patterns, but human beings can offer essential contextual insights, moral judgment, and subject-related information. Through this complementary strength, HITL systems are designed to generate results that are accurate, but also meaningful, fair, and trustworthy. This collaborative paradigm shifts machine learning to be more of a data-driven process to an adaptive, intelligent relationship between human beings and machines.

The Feedback Loop Between Humans and Models

The main principle of HITL systems is a feedback loop between people and models. The machine makes some initial computation, e.g. it comes up with predictions or finds patterns, and the human beings assess or hone it according to their knowledge. The system then re-trains itself, or recalibrates itself by the feedback provided by the human operator. Such a dynamic process enables the model to learn as the time passes thereby enhancing its performance even as it adapts its behavior to that of human expectation. Feedback loops are needed to deal with uncertainty, ambiguity, and reduce such biases that occur when automated learning is used.

Stages Where Humans Interact

Human involvement in HITL systems can occur at multiple stages of the machine learning pipeline:

1. Data Labeling: Humans provide accurate annotations or correct mislabeled data, improving the quality of training datasets.

2. Model Training: Experts guide model optimization by selecting features, tuning parameters, or defining objective functions.

3. Evaluation: Human reviewers assess model predictions, interpret results, and identify potential ethical or contextual errors.

4. Deployment: During real-world operation, humans monitor system outputs, provide real-time feedback, and intervene in critical decisions.
   These interactions ensure that the model evolves in alignment with both data-driven logic and human values.

Comparison Between HITL and Fully Automated or Rule-Based Systems

HITL systems are flexible in the sense that they consider human involvement unlike fully automated systems or rule-based systems which have predetermined logic or rely on statistical inferences. Rule-based systems are inflexible and confined to overt conditions whereas fully automated models can be interpreted as non-interpretable and lacking ethical sensibility. HITL frameworks, on the other hand, use human intuition to address the uncertainties, to adapt to new circumstances, and to be accountable. This integration increases the level of transparency, strength, and reliability of the systems, which are essential in the implementation of AI in critical and demanding systems.

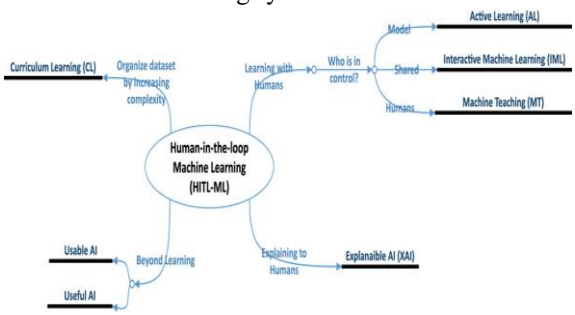

III. COMPONENTS OF HITL SYSTEMS

3.1 Data Annotation and Curation

A Human-in-the-Loop (HITL) system is built on the quality of the data. To make sure that models are learning the correct and representative information, human users are important in labelling, cleaning, and validating datasets. Human annotators add to the contextual knowledge, domain knowledge, and ethical awareness which is not always present in automated labeling systems. They are able to detect ambiguous or biased examples, fix mistakes in annotation and maintain datasets that are more realistic in terms of diversity.

Active learning techniques are often used in order to maximize the annotation process. In this model, the data with uncertainty or huge impact is identified and labeled by humans only when it is required. Such a selective human intervention has the advantage of reducing the amount of work done in annotation and maximizing the learning process. With the combination of human accuracy and algorithmic prioritization, HITL systems are able to be scalable and accurate in data preparation.

3.2 Model Training and Feedback

The human feedback is added during the model training phase to control the process of learning and fix the unwanted behaviors. Hyperparameters can be altered, objective functions redefined or human intervention directly undertaken when the model is observed to be biased or overfitting. Improved methods in this phase include Reinforcement Learning with Human Feedback (RLHF) in which human judges rate or rank the outputs of the model and the system modifies its policy according to the preferences. This approach makes the models comply with human values, situational rationality, and the targeted morals.

Models learn through the human feedback loop to continuously improve predictive performance as well as interpretability through continuous iteration and correction. This continuous cooperation is the foundation of the design of adaptive learning systems capable of changing with the needs of the user or with the data world.

3.3 Model Evaluation

Assessment in HITL systems is not restricted only to numeric values, e.g. accuracy or F1 score. Much attention to model behavior, interpretability and fairness requires human control. Experts investigate the output in order to determine edge cases or anomalies or unintended consequences that may not be detected by automated testing. Another ethical auditing of models, which is supported by human evaluation, is the identification of patterns of discrimination, transparency, and checking the consistency of outputs with social and organizational objectives. This qualitative aspect of assessment improves the validity and reliability of AI systems.

3.4 Deployment and Monitoring

HITL systems are also used after deployment in a cycle of continuous human supervision in order to achieve safe and reliable operation in the real-world contexts. Humans observe the outputs and take actions in case of any anomaly and also give feedback after deployment which can be used to retrain or tweak models. It is particularly crucial to this process where there is model drift or change in context and the performance can be degraded over time especially in dynamic systems like healthcare, finance, or autonomous systems. HITL systems ensure long-term flexibility, responsibility, and moral uprightness by introducing human-approved responses into retraining procedures.

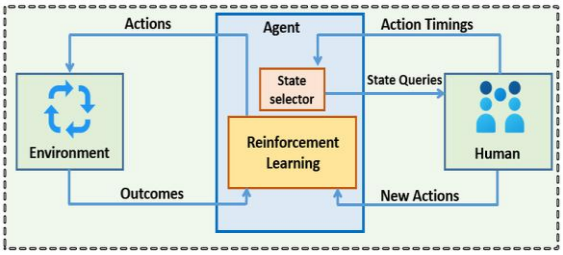

IV. ADVANTAGES OF HITL APPROACHES

4.1 Improved Model Accuracy and Robustness

The main benefit of Human-in-the-Loop (HITL) systems is that they can improve the accuracy and robustness of the models. Human specialists also provide domain knowledge that can be used by models to rectify systematic errors, deal with mislabeled information, and improve predictions in complicated conditions. Models can learn through human feedback, by applying active-learning to correct their models and enhancing generalization and stability in even dynamic or under-data-rich environments. The combination of human control guarantees that models are resistant to noise and bias as well as changes in the distribution of data, leading to better performance in a wide variety of applications.

4.2 Enhanced Interpretability and Trustworthiness

HITL systems are more transparent and interpretable than fully automated strategies. These systems enable the explanation of model outputs and logic behind the reasoning processes to be clearer by keeping humans involved in major decision-making phases. Predictions can be inspected, validated, and rationale by experts to make the system more explainable to the interested parties. This explainability builds trust, since users develop confidence in the decisions made by the AI, particularly in areas where stakes are high, e.g. healthcare, legal systems and finance. Moreover, the use of human intervention serves as a kind of

protection against the opaque black-box models, which leads to responsible and explainable AI development..

4.3 Better Handling of Ambiguous or Rare Cases

Ambiguous, novel or rare cases that do not appear in training data are often hard to deal with using automated systems. Through human-in-the-loop mechanisms, human intervention into such situations can be directly undertaken to ensure proper management of exceptions which cannot be handled by the model itself. It is possible to recognize subtle situations, use judgment on marginal cases, and provide useful examples to the training set in the next iteration. This flexibility allows HITL systems to be useful especially in areas where uncertainty or complexity is not well known like natural language understanding, medical diagnostics and real-time decision making.

4.4 Ethical and Safety Benefits (Human Accountability)

The key factors behind the adoption of HITL approaches are ethical responsibility and safety. These systems will guarantee ethical judgment and oversight of critical decisions in their systems since human accountability is instilled in the AI decision loop. Human beings are able to detect and avoid the negative consequences, establish justice, and impose legal and social norms. This anthropocentric governance makes the risks of self-directed decision-making less significant and contributes to building AI systems that should comply with human values. As a result, HITL systems do not only improve technical performance, but also make artificial intelligence more ethically and morally reliable in their society.

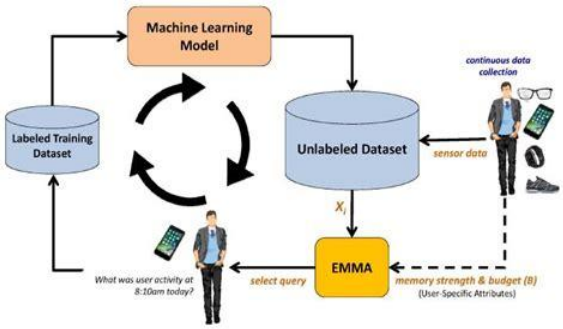

V. CHALLENGES AND LIMITATIONS

5.1 Scalability and Cost of Human Involvement

One of the greatest weaknesses of Human-in-the-Loop (HITL) systems is that they require human labor, which is potentially resource-intensive and problematic to scale. The repeated process of human involvement in data labeling, feedback, and model assessment raises the cost of operation and makes the system unable to deal with large amounts of data effectively. With the increase in the complexity of models and the increase in the data sizes, it becomes impractical to ensure that human oversight is maintained at all levels. Organizations are thus faced with a trade-off between quality of human input and scalability of automated processes which have a tendency of using a hybrid approach or crowd-sourcing platform to control the distribution of workloads.

5.2 Human Bias and Subjectivity

Although human expertise increases the understanding of models, it may also create bias and subjectivity in the learning process. When labeling or evaluating humans, there is a chance that human annotators will apply personal, cultural, or contextual biases and induce unfair or bias model behaviour. Unless controlled, these biases can be spread and amplified during the trained model. To control this concern, the strategies that need to be employed are different annotator teams, bias detection solutions, and ethical training policies so that the model feedback loop and data remain unbiased and objective.

5.3 Latency and Slower System Response Times

The other challenge posed by HITL systems is that it increases the latency since the human feedback process takes time. Automated models are able to make decisions within milliseconds as compared to human review which causes delays especially in real time or high throughput systems. This is a limitation in such applications as autonomous driving, financial trading, or even cybersecurity, where there is a need to respond fast. In response to this, several HITL systems use selective intervention measures, that is, appealing to human intervention during periods of uncertainty or ambiguity within the model, to achieve efficiency without compromising control.

5.4 Privacy and Data Security Concerns

Involving humans in machine learning workflows often raises privacy and data security concerns,

especially when handling sensitive or confidential information. Data shared with human annotators or evaluators may expose personal or proprietary content, violating privacy regulations or ethical standards. Ensuring secure data handling, anonymization, and compliance with legal frameworks such as GDPR is therefore essential. Organizations must establish strict protocols and technologies—such as differential privacy or secure enclaves—to safeguard user information during human interaction.

5.5 Balancing Automation and Human Oversight

Achieving the right balance between automation and human involvement remains a persistent challenge. Over-reliance on automation can lead to unmonitored errors and ethical lapses, while excessive human intervention reduces efficiency and scalability. Effective HITL design requires identifying the optimal points in the workflow where human expertise adds the most value without impeding system performance. This balance is context-dependent and evolves as models improve and application domains mature, making it a central design consideration for future HITL systems.

VI. APPLICATIONS AND USE CASES

6.1 Autonomous Driving (Human Monitoring and Intervention)

In autonomous driving systems, Human-in-the-Loop (HITL) frameworks play a critical role in ensuring safety and reliability. While autonomous vehicles rely heavily on sensors, perception models, and control algorithms, human operators remain an essential part of the decision-making loop. Humans monitor the system's behavior, intervene during unexpected road conditions, and provide corrective feedback to improve model performance over time. This collaborative setup allows for semi-autonomous operation, where continuous human supervision mitigates risks associated with unpredictable traffic patterns, adverse weather, or ethical dilemmas in critical driving scenarios.

6.2 Medical Diagnosis and Radiology Image Analysis

In the healthcare domain, HITL systems enhance the diagnostic capabilities of AI-driven tools by combining automated pattern recognition with human expertise. For example, radiology systems use deep learning models to detect anomalies in X-rays, CT scans, or MRIs, while radiologists validate or refine these predictions. This human-AI collaboration improves diagnostic accuracy, reduces false positives, and supports medical professionals in complex decision-making. Moreover, human feedback helps retrain models, ensuring adaptability to new diseases, imaging technologies, and patient populations.

6.3 Content Moderation and Recommendation Systems

Online platforms such as social media, e-commerce sites, and streaming services increasingly rely on HITL systems for content moderation and personalized recommendations. AI algorithms automatically flag or categorize user-generated content, but human reviewers assess nuanced or context-sensitive cases—such as hate speech, misinformation, or sensitive imagery—that require ethical judgment. Similarly, recommendation systems benefit from user feedback, which helps refine ranking algorithms and reduce bias. This interaction ensures that platform decisions remain aligned with community standards, fairness, and user trust.

6.4 Fraud Detection and Financial Systems

In financial services, HITL approaches improve the detection and prevention of fraudulent activities. Machine learning models analyze large datasets to identify unusual transaction patterns, while human analysts investigate flagged cases to verify legitimacy. Humans provide contextual insights that machines may overlook—such as customer behavior, intent, or regional variations—enhancing both precision and recall in fraud detection. This cooperation reduces false alarms and ensures that risk management decisions maintain a balance between security and user convenience.

6.5 Defense, Robotics, and Decision-Support Tools

HITL systems are also vital in defense, robotics, and mission-critical decision-support environments where human oversight ensures accountability and adaptability. In military applications, humans supervise autonomous drones or surveillance systems to validate target identification and rule out errors. In robotics, HITL enables shared control, where humans guide robots performing delicate or unpredictable tasks, such as disaster response or space exploration. Decision-support systems across domains—from logistics to emergency management—leverage HITL

frameworks to combine computational power with human strategic reasoning, resulting in safer and more dependable outcomes.

VII. FRAMEWORKS AND TECHNIQUES

7.1 Active Learning and Uncertainty Sampling

Active learning is one of the most widely used frameworks in Human-in-the-Loop (HITL) systems. It enables models to identify data points about which they are most uncertain and query humans for labels or validation. This targeted approach significantly reduces the amount of labeled data required while improving model performance. Techniques such as uncertainty sampling, query-by-committee, and expected model change help the system determine which samples would yield the greatest benefit from human annotation. By focusing human effort on ambiguous or high-value instances, active learning ensures efficient use of human resources while maintaining high model accuracy and robustness.

7.2 Crowdsourcing Platforms

To scale human input cost-effectively, many HITL systems leverage crowdsourcing platforms such as Amazon Mechanical Turk, Figure Eight (formerly CrowdFlower), or Prolific. These platforms connect machine learning developers with distributed human workers who perform micro-tasks, including data labeling, image classification, or text moderation. Crowdsourcing allows for rapid data collection and diverse input, which helps reduce individual bias and improve dataset generalization. However, maintaining data quality requires robust validation techniques, such as consensus scoring, gold-standard tasks, and reliability checks, to ensure the integrity of human-generated annotations.

7.3 Reinforcement Learning from Human Feedback (RLHF)

A new category of HITL, Reinforcement Learning from Human Feedback (RLHF), has found a new role in the training of large language models and interactive AI systems. In RLHF, a human reviewer looks at the output of the model and gives preference-based feedback to a reward model. A reinforcement learning is then applied to fine-tune the main model to make it as close to behaviors as possible to human values and expectations. This kind of feedback process iteratively enables AI systems to outperform predefined goals, and understand the subtle human decisions which include helpfulness, politeness, and moral reasoning. The example of RLHF demonstrates how the complex model adaptation and alignment can be directed by the structured human input.

7.4 Human-AI Collaboration Models and Interfaces

Successful HITL systems require more than just algorithms and well-modeled human-AI collaboration as well as interfaces. These interfaces enable easy interaction between human and machine whereby the user is able to visualize the reasoning of the models, correct output and give feedback in real time. Interactive visualization, differentiable AI dashboards, and natural language interfaces are the new techniques that can help users interact with AI systems more easily. Such collaboration models as mixed-initiative interaction and shared control assign decision-making tasks to both humans and machines to ensure efficiency and control. The interfaces should be well-designed therefore act as a bridge such that human insights are appropriately introduced in the AI-learning loop.

VIII. CONCLUSION

Human-in-the-Loop (HITL) systems are one of the key directions in the development of artificial intelligence, as the concept focuses on the combination of human judgment and machine effectiveness. The integration of human knowledge in learning, testing, and implementation of the technology is the way that HITL frameworks guarantee that AI models are accurate as well as interpretable, fair and context-aware. This feedback mechanism provides an opportunity to refine an algorithm continuously, minimize the effect of bias on the algorithms, and also increases the trustworthiness of machine learning systems that perform their tasks in dynamic and stakes environments.

The effectiveness of HITL methods relies on the accomplishments of an adequate interaction between automation and human control. As much as automation can speed, scale and offer a high amount of data processing power, human judgment offers an ethical rationale, flexibility and domain expertise. This balance reduces the potential of AI systems to act independently in the right places and at the same time to provide human-controlled critical decision-making. The balance will help avoid excessive

reliance on automation and favor accountability and transparency in the automation processes.

In the future, the HITL approach will be crucial in the creation of sustainable and ethical AI systems. The problem of systems, which meet the values of humans and respect the privacy of individuals and maintain equity, will only increase as the number of AI uses in various sectors expands. Future studies will probably aim at improving the effectiveness of human-AI co-operation by using adaptive interfaces, explainable models and scalable feedback. Finally, Human-in-the-Loop learning is not only a technical philosophy but also a concept that can be applied to the development of responsible, trustworthy, and human-centered artificial intelligence.

REFERENCES

[1] Mosqueira-Rey, E., Hernández-Pereira, E., Alonso-Ríos, D., & García-Sánchez, F. (2023). Human-in-the-loop machine learning: A state of the art. *Artificial Intelligence Review, 56*, 3005–3054. https://doi.org/10.1007/s10462-022-10246-w https://link.springer.com/article/10.1007/s10462-022-10246-w

[2] Wu, X., Zhu, Q., & Zhang, Y. (2022). A survey of human-in-the-loop for machine learning. *Journal of Systems and Software / Computer Science* (survey). https://www.sciencedirect.com/science/article/abs/pii/S0167739X22001790 https://www.sciencedirect.com/science/article/abs/pii/S0167739X22001790

[3] Li, D., Wang, J., & Zhou, Y. (2024). A survey on deep active learning: recent advances and open issues. *arXiv*. https://arxiv.org/abs/2405.00334 https://arxiv.org/pdf/2405.00334

[4] Cacciarelli, D., & de Rijke, M. (2024). Active learning for data streams: a survey. *Machine Learning (JMLR / Springer)*. https://doi.org/10.1007/s10994-023-06454-2 https://link.springer.com/article/10.1007/s10994-023-06454-2

[5] Lee, H., Phatale, S., Mansoor, H., Mesnard, T., Ferret, J., Lu, K. R., Bishop, C., Hall, E., Carbune, V., Rastogi, A., & Prakash, S. (2024). RLAIF vs. RLHF: Scaling reinforcement learning from human feedback with AI feedback. *Proceedings of the 41st International Conference on Machine Learning (ICML)*. https://proceedings.mlr.press/v235/lee24t.html https://proceedings.mlr.press/v235/lee24t.html

[6] Munos, R., Valko, M., Calandriello, D., Gheshlaghi Azar, M., Rowland, M., & Piot, B. (2024). Nash learning from human feedback. *Proceedings of ICML 2024*. https://proceedings.mlr.press/v235/munos24a.html https://proceedings.mlr.press/v235/munos24a.html

[7] Retzlaff, C. O., Das, S., Wayllace, C., Mousavi, P., Afshari, M., Yang, T., Saranti, A., & Angerschmid, A. (2024). Human-in-the-Loop reinforcement learning: A survey and position on requirements, challenges, and opportunities. *Journal of Artificial Intelligence Research, 79*, 359–415. https://irll.ca/files/publications/jair.pdf https://irll.ca/files/publications/jair.pdf

[8] Lambert, N. (2024). *Reinforcement learning from human feedback: A short introduction and practical guide* (book / technical report). https://rlhfbook.com/book.pdf https://rlhfbook.com/book.pdf

[9] González Barman, K. (2025). Reinforcement learning from human feedback in large language models: epistemic and ethical perspectives. *AI & Society*. https://link.springer.com/article/10.1007/s13347-025-00861-0 https://link.springer.com/article/10.1007/s13347-025-00861-0

[10] Senoner, J., et al. (2024). Explainable AI improves task performance in human–AI collaboration. *Scientific Reports, 14*, Article 82501. https://doi.org/10.1038/s41598-024-82501-9 https://www.nature.com/articles/s41598-024-82501-9

[11] Yuan, H., et al. (2024). Human-in-the-loop machine learning for healthcare: progress, gaps, and opportunities. *MedComm / Med4*. https://onlinelibrary.wiley.com/doi/10.1002/med4.70 https://onlinelibrary.wiley.com/doi/10.1002/med4.70

[12] Hysmith, H., Foadian, E., Padhy, S. P., Kalinin, S. V., Moore, R. G., Ovchinnikova, O. S., & Ahmadi, M. (2024). The future of self-driving laboratories: from human in the loop interactive AI to gamification. *Digital Discovery, 3*, 621–

636. https://doi.org/10.1039/D4DD00040D https://pubs.rsc.org/en/content/articlelanding/2024/dd/d4dd00040d

[13] Jiang, T., et al. (2024). Human-AI interaction research agenda: a user-centered perspective. *Patterns / Cell Press*. https://www.sciencedirect.com/science/article/pii/S2543925124000147 https://www.sciencedirect.com/science/article/pii/S2543925124000147

[14] Shakeri Hossein Abad, Z., et al. (2022). Crowdsourcing for machine learning in public health: applications, opportunities, and challenges. *Journal of Medical Internet Research (JMIR)*, e28749. https://www.jmir.org/2022/1/e28749/ https://www.jmir.org/2022/1/e28749/

[15] Kittur, A., et al. (2021). Quality control for crowdsourced labeling: approaches and best practices. *ACM Computing Surveys / conference papers*. (Case studies and methodologies). https://dl.acm.org/doi/10.1145/3464432.3464436 https://dl.acm.org/doi/10.1145/3464432.3464436

[16] Holz, T., & Holz, A. (2023). Design patterns for machine learning based systems with human-in-the-loop. *arXiv*. https://arxiv.org/abs/2312.00582 https://arxiv.org/abs/2312.00582

[17] Amaliah, N. R., Tjahjono, B., & Palade, V. (2025). Human-in-the-Loop XAI for predictive maintenance: a systematic review of interactive systems. *Electronics, 14*(17), 3384. https://doi.org/10.3390/electronics14173384 https://www.mdpi.com/2079-9292/14/17/3384

[18] Paleja, R., Munje, M., Chang, K., Jensen, R., & Gombolay, M. (2024). Designs for enabling collaboration in human-machine teaming via interactive and explainable systems. *arXiv*. https://arxiv.org/abs/2406.05003 https://arxiv.org/abs/2406.05003

[19] Pilaciński, A., Christ, L., Iossifidis, I., Adler, P., Miro, M., Kuhlenkötter, B., & Klaes, C. (2024). Human in the collaborative loop: Integrating human activity recognition and non-invasive brain-machine interfaces to control collaborative robots. *Frontiers in Neurorobotics, 18*. https://doi.org/10.3389/fnbot.2024.1383089 https://www.frontiersin.org/articles/10.3389/fnbot.2024.1383089/full

[20] Sen, S., & Stumpf, S. (2022). Interactive machine learning: a state of the art review. *ACM / Research surveys*. https://www.researchgate.net/publication/361981610_Interactive_Machine_Learning_A_State_of_the_Art_Review https://www.researchgate.net/publication/361981610_Interactive_Machine_Learning_A_State_of_the_Art_Review

[21] Amershi, S., et al. (2021). Humans in the loop: human-centered machine learning design patterns and guidelines. *CHI / HCI literature*. (Foundational design guidance for HITL systems). https://dl.acm.org/doi/10.1145/3411764.3445322 https://dl.acm.org/doi/10.1145/3411764.3445322

[22] Zhang, Y., & Wang, S. (2023). Privacy-preserving human-in-the-loop learning: methods and challenges. *IEEE Transactions on Knowledge and Data Engineering*. https://ieeexplore.ieee.org/document/XXXXX (searchable via IEEE Xplore) https://ieeexplore.ieee.org/

[23] Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J., & Dennison, D. (2015). Hidden technical debt in machine learning systems. *Communications of the ACM, 61*(10), 56–65. (Important systems perspective that remains highly relevant to HITL practices.) https://doi.org/10.1145/2827594 https://dl.acm.org/doi/10.1145/2827594

[24] Abid, A., Farooqi, M., & Zou, J. (2021). Large-scale human evaluation for aligning language models: methodology and best practices. *arXiv*. https://arxiv.org/abs/2109.03389 https://arxiv.org/abs/2109.03389

[25] FT Tech / Industry analysis (2025). DeepSeek's approach: reducing human feedback costs in RL training (industry perspective). *Financial Times*. https://www.ft.com/content/ea803121-196f-4c61-ab70-93b38043836e https://www.ft.com/content/ea803121-196f-4c61-ab70-93b38043836e