Human Interaction and Emerging Technologies (IHIET-AI 2025), Vol. 161, 2025, 367–377
https://doi.org/10.54941/ahfe1005929

Diversity of Perception in Human-AI Collaboration

Mohamed Quafafou

LIS, CNRS/Aix-Marseille University/UTLN, Marseille, France

ABSTRACT

Two key approaches to building AI systems are Model-Centric AI (MC-AI) and Data-Centric AI (DC-AI). When AI systems are deployed in real-world environments, they become part of a socio-technical ecosystem, interacting with humans, processes, and other systems. This interaction often occurs in hybrid teams, where humans and AI collaborate to achieve shared objectives. However, human influences, at any stage, can lead to suboptimal outcomes, such as model drift or reduced performance. In fact, human introduces variability, as personal experience, biases, and decision-making approaches can significantly impact outcomes. Changing one human in the process can alter the results dramatically. This paper review processes involved into building, deploying, monitoring, and maintaining AI-systems and discusses human influences at each step, the potential risks that may arise and the main skills necessary to avoid human's negative influences. By incorporating perception diversity and tolerating ambiguity, the computing-with-perception framework enhances human-AI collaboration, enabling systems to manage complexity and ambiguity in human-AI collaboration considering real-world problems.

Keywords: Human-AI, Collaboration, Perception, Diversity, Decision-making

INTRODUCTION

Artificial Intelligence (AI) is transforming human life and is shaping the present and future in profound ways. Two key approaches to building AI systems are model-centric AI (MC-AI) and data-centric AI (DC-AI). In DC-AI (Zha et al., 2025), the emphasis lies on enhancing the quality, quantity, and diversity of training data, while MC-AI (Zhu-Hua, 2021) focuses on improving the performance of machine learning models themselves. AI systems are typically developed by humans without direct collaboration between AI and users during the design phase. Once deployed in real-world environments, these systems become integral to socio-technical ecosystems, where they interact with humans, organizational processes, and other technologies. In many cases, AI operates within hybrid teams, where humans and AI collaborate to achieve shared objectives. These interactions influence decision-making, workflow dynamics, and overall system performance, requiring careful integration to ensure effective human-AI collaboration in real-world applications.

Human-AI collaboration (Fragiadakis et al., 2024) is essential not only during the deployment and use of AI systems but also throughout their development. It plays a crucial role in shaping AI models, particularly in Data-Centric AI (DC-AI) and Model-Centric AI (MC-AI) approaches. Integrating human expertise at every stage helps enhance data quality, improve model performance, and ensure AI systems align with real-world needs and ethical considerations. In fact, human choices, and decisions at any stage—such as poorly defined objectives, irrelevant feedback, or lack of trust in AI—can lead to suboptimal outcomes, such as model drift or reduced performance. Hence, when AI agents work interdependently toward a common goal alongside human agents (Li et al., 2024), human influence introduces variability, as personal experience, biases, and decision-making approaches can significantly impact outcomes. Changing one human in the process can alter the results dramatically.

To face the challenges that human-AI systems (Wang et al., 2020; Yang et al., 2020) face, we consider that AI-systems do not depend only on data, methods, algorithms, heuristics, infrastructure and technologies, but they are mainly depending on human with whom they interact and collaborate. This assumption is important because, beyond humans objective decisions, their subjectivity and perception diversity play a central role in problem-solving and decision-making. Humans interpret objects, concepts, and events through the lens of their unique experiences, often tolerating contradictions and ambiguity (Reinecke et al., 2025). The main question is how to account for human perception diversity in artificial intelligence. Dealing with this question will allow us to consider the impact of each human in hybrid-teams and its effect on the collaboration. Then, we can manage alternatives solutions corresponding to the different humans' perceptions. To address the diversity of human choices, decisions, and actions in Human-AI collaboration, we incorporate a perception-driven approach. This enables AI systems to adapt to individual and contextual variations across different stages, including data processing, model development, deployment, monitoring, and maintenance. By integrating perception, we enhance AI's ability to understand and respond to human behaviors, ensuring more effective and adaptive collaboration. This approach is embodied in the framework we propose, which is based on a new notion of sets, where a set is not completely characterized by its members but rely essentially on its observers (Quafafou, 2020).

HUMAN IN DATA-CENTRIC AI

Data-centric AI (DC-AI) posits that improving the quality of the data—such as addressing issues like noise, biases, missing values, or inaccuracies—can be as effective, or even more effective, than refining the models themselves. The key idea is that good data can enhance the performance of simpler models, while poor data can limit even the most advanced models: enhancing data quality leads to more reliable and robust AI models. However, human influences key stages like defining the problem, selecting data sources, labeling data, addressing biases, and pre-processing datasets, etc. (see Table[1] 1.)

At the beginning of the DC-AI process, we must clearly define the problem and determine the role of data in solving it, e.g., identifying the task, specifying the target variables and features required for the task, and defining the data requirements, including types, volume, and quality.

After, we must acquire data that is relevant, diverse, and representative of the problem domain, e.g., gathering data from multiple sources, ensuring the data aligns with the problem requirements. addressing diversity by collecting data representing various conditions, populations, or scenarios. In the context of supervised learning, we must create high-quality labeled datasets, i.e., using manual annotation by domain experts or crowd-sourced annotators, employing semi-automated or fully automated labeling tools, if applicable, validating labels, ensuring consistency in the annotation process, and handling edge cases or ambiguous data carefully with domain expertise.

The data, representing the raw material, is now available and must be cleaned, preprocessed, and transformed into a usable format for model training. Hence, we apply different methods for: (1) handling missing values, duplicates, outliers, (2) normalizing or standardize numerical features, (3) encoding categorical variables, (4) tokenizing or pre-processing text data, and (5) resizing or normalizing image data, if applicable, etc.

Table 1. Some human's influences in data centric AI.

Step	Human Influence	Potential Risks	Required Skills
Data Collection & Acquisition	- Selecting relevant data sources and ensuring data diversity. - Defining data collection protocols to avoid biases. - Ensuring legal and ethical compliance	- Biased or unrepresentative data leading to model discrimination. - Data privacy violations. - Poor-quality data affecting downstream performance.	- Domain expertise to understand data relevance. - Knowledge of data privacy laws and ethical AI principles. - Data engineering skills to manage data pipelines.
Data Cleaning & Preprocessing	- Identifying and handling missing, duplicate, or inconsistent data. - Defining normalization, standardization, and feature engineering techniques. - Applying de-biasing strategies to ensure fairness.	- Incorrect cleaning may remove important information. - Introducing unintended biases while filtering or augmenting data. - Poor preprocessing leading to model instability.	- Proficiency in data preprocessing tools. - Understanding of statistical methods to detect anomalies. - Ethical AI knowledge to minimize biases.
Data Annotation & Labeling	- Establishing high-quality annotation guidelines. - Managing human labelers to ensure consistency. - Detecting and correcting labeling errors or biases.	- Subjective labeling leading to inconsistencies. - Annotation errors impacting model accuracy. - Scalability issues when relying solely on human annotators.	- Experience in annotation tools. - Understanding of inter-annotator agreement techniques. - Knowledge of active learning to optimize labeling efforts.
Data Augmentation & Enrichment	- Designing synthetic data generation strategies. - Augmenting datasets to improve model generalization. - Ensuring added data maintains real-world relevance.	- Artificially generated data may not reflect real-world distributions. - Over-augmentation leading to data redundancy. - Introducing artifacts that distort model learning.	- Knowledge of data augmentation techniques. - Statistical skills to validate augmentation effectiveness. - Domain expertise to maintain data authenticity.

[1] Such tables summarize the influence of human, but they are not exhaustive because the limited number of pages.

Before analyzing the data, it is crucial to validate its accuracy, quality, and relevance for the intended task. This involves conducting statistical analyses, identifying potential biases and imbalances, ensuring data integrity, and visualizing the data to detect anomalies or inconsistencies. These steps help establish a reliable dataset, reducing errors and enhancing the effectiveness of subsequent analysis and model development. In some situation, we need to augment the data to increase the diversity of the dataset without collecting additional data.

In short, human oversight is critical in data-centric AI, as poor data quality directly impacts model fairness, reliability, and performance. To avoid negative influences, professionals must combine domain expertise, data engineering proficiency, statistical knowledge, and ethical AI principles to ensure responsible data handling. After all these tedious steps, we develop a model using the model-centric AI process.

HUMAN IN MODEL-CENTRIC AI

Model-Centric AI (MC-AI) primarily focuses on enhancing machine learning algorithms and models to achieve better performance. By refining the model's architecture, algorithms, and parameters, MC-AI aims to deliver superior results. In MC-AI process, the journey from problem definition to model optimization is methodical and precise, but human influence is crucial (see Table 2).

In Model-Centric AI, the problem definition phase is a critical step where humans shape the project's direction. This involves clearly identifying the problem, ensuring AI is a suitable solution, and aligning objectives—such as improving accuracy, reducing costs, or predicting behavior—with business needs. A key consideration is whether the problem can be effectively addressed using available data, a process influenced by Data-Centric AI. Care must be taken to avoid embedding biases or unethical objectives. Selecting appropriate evaluation metrics (e.g., accuracy, precision, recall, F1-score) is crucial to ensure the model's success. However, misaligned objectives or poorly chosen metrics can lead to ineffective or even harmful models. Overfitting to biased or irrelevant data further compounds these risks. Success in this phase requires domain expertise to grasp problem complexities, strong communication to align stakeholders, and analytical thinking to establish clear, measurable success criteria.

In the model development phase, humans play a crucial role in designing and training machine learning models. This includes selecting appropriate algorithms (e.g., random forests, neural networks) and fine-tuning hyperparameters (e.g., learning rates, batch sizes) to optimize performance. Rigorous experimentation is conducted to improve robustness and prevent biases. However, human intervention introduces risks such as overfitting, which harms generalization, and ethical concerns like fairness and transparency. To mitigate these challenges, expertise in machine learning frameworks, mathematics, and ethical AI principles is essential, ensuring models are both effective and responsible in real-world applications.

Table 2. Some human's influences in model centric AI.

Step	Human Activities	Potential Risks	Required Skills
Problem definition	- Clearly articulate the goal of the model - Assessing whether the problem is solvable with available data and techniques - Ensuring the problem definition does not inadvertently encode bias or unethical goals - Selecting evaluation metrics	- Misaligned objectives or metrics - Overfitting goals to irrelevant or biased aspects of the data	- Good domain expertise to understand the problem and its constraints - Strong communication skills - Analytical thinking
Model development	- Model selection - Hyperparameter Tuning - Experimentation - Incorporating Ethical Considerations	- Overfitting to training data - Ignoring ethical implications	- Proficiency in machine learning frameworks - Strong mathematical foundation - Knowledge of fairness metrics - Knowledge of ethical AI principles.
Model Evaluation	- Defining Evaluation Metrics - Bias Testing - Robustness Testing - Model Interpretability	- Over-reliance on a single metric - Insufficient bias testing	- Knowledge of evaluation metrics - Familiarity with tools for fairness - Ability to interpret and explain model behavior
Model Optimization	- Hyperparameter Tuning - Model Simplification - Use regularization - Model Compression - Improving Generalization - Optimize the model to scale efficiently and making it interpretable	- Excessive optimization - Focusing solely on efficiency - Amplifying biases present in the data - Overlooking hardware or software constraints - Ignoring ethical implications	- Advanced knowledge of machine learning algorithms and their optimization techniques - Expertise in hyperparameter tuning methods - Familiarity with regularization techniques - Proficiency in model interpretability tools - Ability to analyze trade-offs - Understanding of metrics beyond accuracy - Working closely with domain experts - Knowledge of AI ethics and fairness

Once the model is trained, its performance is evaluated using appropriate metrics. This involves selecting evaluation metrics based on task requirements and deciding on validation methods, such as k-fold cross-validation. Performance metrics are interpreted to identify areas for improvement.

Finally, the model optimization step begins to enhance model performance by addressing identified weaknesses. This involves fine-tuning hyperparameters and model architecture, and potentially incorporating advanced techniques like transfer learning or distillation. In this phase, humans play a crucial role in testing the model for performance, robustness, and fairness. They define evaluation metrics by selecting those that reflect real-world performance, such as precision for fraud detection and recall for medical diagnosis. Additionally, they assess the model for potential biases across different groups, test it on edge cases and adversarial scenarios, and ensure it is interpretable and explainable to stakeholders. The main risks of human intervention include over-reliance on a single metric, which can lead to skewed results, and insufficient bias testing, which may result in unfair or discriminatory outcomes. To mitigate these risks, several competencies are essential: knowledge of evaluation metrics and their alignment with business goals, familiarity with fairness testing tools, and the ability to interpret and explain model behavior.

Risks include over-optimization, which may harm generalization, create overly complex models, or degrade fairness. Ignoring deployment constraints or ethical considerations can also lead to impractical or harmful outcomes. Success requires expertise in machine learning, ethical AI principles, and domain knowledge to make informed, balanced decisions.

After the development of the model, we deploy it and becomes part of a socio-technical ecosystem where it interacts with humans, processes, and other systems.

HUMAN IN DEPLOYMENT, MONITORING AND MAINTENANCE

Humans play a critical role in every phase of the deployment, monitoring, and maintenance of an AI system. Their expertise, decisions, and interventions influence the success and ethical implementation of these systems. Below will briefly introduce the main steps and we discuss how humans can influence each one among them, the potential risks they require, and the required skills needed to avoid negative consequences (see Table 3.)

Let us start with the deployment phase that involves integrating the machine learning model into production environments. It consists of several steps, for example infrastructure setup, model integration, deployment pipeline, scalability and optimization.

In the infrastructure setup step, human influences concern the selection of appropriate hardware and cloud infrastructure (e.g., GPUs, TPUs, or edge devices) and setting up software environments. This creates potential risks like a poor resource planning may lead to performance bottlenecks or excessive costs. To avoid negative human influence, the person in the loop must have several skills like knowledge of cloud platforms, he must also be familiar with hardware requirements for AI, and understands DevOps practices. Moreover, human can also influence model integration as he decides how to deploy the model, choosing deployment tools, and designing the API. Hence a potential risk arises consisting of an inefficient deployment strategy, which may lead to misalignment between the model's output and the application requirements. Thus, several skills are required like to have a good knowledge of APIs and frameworks for model serving, software development and system integration skills and the ability to understand business needs to align technical decisions.

Table 3. Some human's influences during the deployment, monitoring and maintenance phases.

Step	Human Influence	Potential Risks	Required Skills
Deployment & Integration	- Model Selection & Readiness Check - Infrastructure Setup - Scalability & Performance Optimization - Integration with Business Systems: - Security & Compliance Validation	- Deploying an untested or biased model leading to failures. - Performance bottlenecks due to poor infrastructure choices. - Security vulnerabilities exposing sensitive data.	- Proficiency in cloud platforms (AWS, Azure, GCP) and containerization (Docker, Kubernetes) - Knowledge of model optimization (quantization, pruning) for efficient deployment. - Understanding of security best practices and compliance standards.
Monitoring & Performance Tracking	- Defining Key Performance Indicators (KPIs) - Real-time Model Performance Analysis - Error Analysis & Failure Investigation - Anomaly Detection	- Ignoring model drift leads to outdated and unreliable predictions. - Failure to detect bias shifts can reinforce discrimination over time. - Inadequate logging makes it difficult to troubleshoot issues.	- Expertise in monitoring tools. - Statistical knowledge for drift detection and root cause analysis. - Experience with logging frameworks and automated alerting.
Continuous Model Updates & Retraining	- Data Refresh & Labeling Strategies - Model Retraining Pipelines - Fine-tuning for New Use Cases - Ethical & Bias Audits	- Retraining on biased or noisy data leading to performance decline. - Overfitting to recent data, reducing generalization. - Ignoring fairness metrics, causing ethical concerns.	- Proficiency in data versioning tools (DVC, Delta Lake). - Experience with automated ML pipelines (Kubeflow, Airflow). - Understanding of fairness metrics and AI ethics.
Infrastructure Maintenance & Scalability	- Resource Optimization - Hardware & Software Upgrades - Disaster Recovery & Backup Planning - Energy Efficiency & Sustainability	- Downtime or system failures due to outdated infrastructure. - High operational costs from inefficient resource allocation. - Environmental impact due to excessive energy consumption.	- Experience in cloud cost management and autoscaling strategies. - Knowledge of sustainable AI practices and energy-efficient computing. - Expertise in DevOps tools (Terraform, Ansible) for infrastructure automation.

Humans play a crucial role in designing CI/CD pipelines, which automate the deployment and updating of AI systems. Poor pipeline design can cause deployment delays, errors, and production bugs, compromising system reliability and performance. To mitigate these risks, expertise in CI/CD tools, scripting, automation, and MLOps principles is essential. A well-designed pipeline ensures smooth integration, efficient model updates, and automated lifecycle management, reducing manual intervention and minimizing errors. By implementing robust automation strategies, teams can enhance the reliability, scalability, and maintainability of AI systems in production environments.

Finally, human influence the model optimization for inference and scaling deployments. In this case the potential risk is Under- or over-scaling infrastructure, leading to poor user experiences or wasted resources. Hence, good knowledge of scaling systems and model optimization techniques for deployment are necessary for good human-AI collaboration.

When AI-systems are developed and deployed, the following main question arise: can they have abilities of human to live in the real world, e.g., physical world?

COMPUTING WITH PERCEPTION

Humans vs. Machines

Humans achieve their daily goals using their ability to think, e.g., various cognitive abilities, including reasoning and decision-making. However, we often overlook the influence of the real, physical world in human-AI collaboration, despite human's perceptions playing a crucial role in shaping human thought processes, choices, and decisions. Understanding this interaction is essential for accurately modeling human cognition and behavior in human-AI collaboration.

As pointed by John McCarthy, today's machines are limited because their relation to the world is almost non-existent: "What the robot believes about the world in general doesn't arise for the limited robots of today because the languages they are programmed to use can't express assertions about the world in general" (McCarthy 2008). On the contrary, humans perceive the world and manage all aspects of everyday life by allowing or tolerating contradictions and ambiguities, whereas machines perform computation under consistency (no contradiction) and completeness (no ambiguity) constraints.

Human perception extends beyond the five senses, being shaped by factors such as education, culture, and past experiences. As AI continues to evolve, perception will play an increasingly vital role in modern life. Advancements in transportation, social media, and global communication have made the world more interconnected, enabling rapid information exchange. This fosters diverse viewpoints, emotions, and expressions, but also facilitates the spread of misinformation and fake news. In this hyper-connected environment, humans and machines interact extensively, supported by technologies that collect, store, and process massive amounts of data. As a result, individuals may interpret the same reality differently, shaping personalized perspectives based on their unique perceptions: we are in the same world, but each one lives in his own world. To enhance Human-AI collaboration, we integrate perception-aware AI systems, ensuring more efficient, robust, and meaningful interactions between humans and machines.

Perceptions and Sets

Epistemologists have proposed various theories of what perception is and how we perceive reality, i.e., the outside world. The three main perception schools (Lewis, 1946) are: (1) In naive realism we directly perceive the world as it is; i.e. things are what they seem, (2) Representative realism is an alternative view, developed by John Locke (Uzgalis, 2007), where we are actively involved in perception, and (3) Idealism, which is defended by George

Berkeley who is persuaded by the thought that we only have direct access to our experiences of the world, but not to the world itself: to be is to be perceived. Beyond these three introduced perception theories, on one hand, our knowledge is related to our perception (Russel, 1967) while, on the other hand, our knowledge, or at least common knowledge, can be mathematically formalized using set theory.

Beyond the three perception theories introduced, knowledge and perception are deeply interconnected. On one hand, our understanding of the world is shaped by our perception. On the other hand, knowledge—particularly common knowledge—can be mathematically formalized using set theory. This formalization provides a structured way to represent and analyze shared understanding, enhancing our ability to model perception-driven reasoning. Thus, using fuzzy sets, Zadeh has introduced, in his paper (Zadeh, 2001), a computational theory of perception considering that perceptions are intrinsically imprecise and stressed the need of "a methodology in which the objects of computation are perceptions - perceptions of time, distance, form, direction, color, shape, truth, likelihood, intent, and other attributes of physical and mental objects".

Diversity of Perceptions and oSets

We have introduced accessible sets or oSets, which are sets that depend on their observers. Let's consider that the external world is represented by the universe U, where each concept c of the real world has its own representative element $X_c \in 2^U$ or $X \in 2^U$ in short. In this context, each observer $i \in I$ perceives the concept X through his own perception function f_i and the perception of X by i is $f_i(X)$. So, $\forall i \in I, \exists!\ f_i: 2^U \to 2^U$, such that $f_i(X)$ is the perception of X by the observer i.

We say that X is accessible for the observer i if and only if $f_i(X)=X$. The accessibility notion is related to the perception and can be best summarized as follows "to be accessible is to be perceived", which is weaker than the Berkeley's idealism "to be is to be perceived".

We can generalize this definition considering several observers simultaneously. In this context and object o can belong to set X for a given observer i, but not for another observer j, e.g., $o \in fi(X)$, but $o \notin fj(X)$. Thus, we introduce a new ternary membership relation, denoted $\in i$, where $x \in i\ X$ means that "x is perceived, by the observer i, to be a member of X". Doing so, we do not exclude the variability of perception of concepts assuming their multiplicity. In fact, the observer i has his own space $Ui=(U,\{i\})$. Hence, each observer constructs his universe according to his perception of both concepts and use it for reasoning, decision making and collaboration with AI-systems. The constructed space may be identical to the real world; more or less different or completely different.

eHealth Example: Blood Pressure and Weight

The description of the computing with perception framework is beyond the scope of this article, but let's take a simple example to show how oSets can be used and to underline their methodological impact on human-AI collaboration. Consider the following e-health problem, where a doctor is collaborating with an AI-system to resolve the flowing problem: "predict a patient's blood pressure from their weight". In machine learning (ML), this problem known as "simple linear regression". Several alternative solutions are possible, including: (1) the AI-system has a prior knowledge of ML methods of the literature, (2) the doctor interact with a large language model (LLM), etc. Each solution its advantage and limits, for example, it's a hard task to update the knowledge of ML methods, and LLM can hallucinate leading to false solutions. Using computing with perception framework, the AI-systems start analyzing its previous interaction with other persons, which they have treated this regression problem, and which methods they have used.

This analysis results into 4 users u_1, u_2, u_3 and u_4 that they have used 6 methods, which are Linear Regression (LiR), Ridge Regression (RR), Lasso Regression (LaR), Elastic Net Regression (ENR), Probabilistic Approaches (PA), and Neural Network (NN). For example, the first and fourth users used only one method, whereas the second and third used two methods, e.g., u_1={LiR}, u_2={LiR, ENR, NN}, u_3={PA, NN }and u_4={RR}. The question is then which method to use?

This question is crucial because there is no consensus between the four users: $u_1 \cap u_2 \cap u_3 \cap u_4 = \emptyset$. Hence, we try to rank method according to the of their users. The result is: LiR and NN are ranked 1 as they are used by two users, ENR, PA, the other methods are ranked 2 because they are used by only 1 user, and LaR is not ranked as it is not used. Hence, can we decide to use LiR and/or NN? Here, we are in a situation of total ignorance because two users have used these two methods, but two others have not used them. In this case where the decision is blocked, we consider that methods to use to resolve the regression problem cannot be represented by a set, but it will be by an oSet. How to compute this oSet?

We proposed a hypergraph-based algorithm to compute "minimal admissible sets", (see Quafafou, 2020; 2016). The resulted oSet is X = {{LiR, ENR, RR}, {PA, LiR, RR}}, which is a set of sets representing the diversity of perception by the four users of the set of ML methods to use to resolve the doctor's problem, e.g., predict a patient's blood pressure from their weight. The interpretation is as follows: there are not just one method or a single set of methods to apply, but two distinct sets of methods—Set {LiR, ENR, RR} and Set {PA, LiR, RR}—based on the solutions provided by the four users who have already resolved the problem. Hence, the doctor will use an ensemble learning approach based on methods in {LiR, ENR, RR} or {PA, LiR, RR}. Using this result, the doctor follow DC-AI and MC-AI processes and deployment, monitoring, and maintenance phases previously introduced in this article.

CONCLUSION

Humans play a pivotal and crucial role during the development of AI-systems and their usage. Their choices, and decisions at any stage can lead to suboptimal outcomes, making the collaboration toxic. In this paper we have analyzed both the construction and the usage of AI-Systems showing the important role of humans' perceptions and introduction of computing with perception based on oSets.

REFERENCES

B. Russell, An Inquiry into Meaning and Truth, Penguin Books Publisher, Pelican book Series, 1967, 332 p.

Daochen Zha, Zaid Pervaiz Bhat, Kwei-Herng Lai, Fan Yang, Zhimeng Jiang, Shaochen Zhong, and Xia Hu. (2025). Data-centric Artificial Intelligence: A Survey. ACM Comput. Surv. 57, 5, Article 129 (May 2025), 42 pages.

Fragiadakis, George, Christos Diou, George Kousiouris, and Mara Nikolaidou. "Evaluating human-ai collaboration: A review and methodological framework." arXiv preprint arXiv:2407.19098 (2024).

Lewis White Beck, Secondary Quality, The Journal of Philosophy, Volume 43, Issue 22, October 1946, pp. 599–610.

Li, You, Yi Li, Qian Chen, and Yaping Chang. "Humans as teammates: The signal of human–AI teaming enhances consumer acceptance of chatbots." International Journal of Information Management 76 (2024): 102771.

McCarthy, John, The Philosophy of AI and the AI of Philosophy, Philosophy of Information, 12, 2008, pp. 711–740.

Mohamed Quafafou, Rough Sets Are ε-Accessible. IJCRS 2016*: 77–86*.

Mohamed Quafafou. Perception in Human-Computer Symbiosis. HCI (48) 2020: 83–89

Reinecke, M. G., Kappes, A., Porsdam Mann, S., Savulescu, J. and Earp, B. D., 2025. The need for an empirical research program regarding human–AI relational norms. AI and Ethics, pp. 1–10.

Uzgalis, William L., Locke's Essay Concerning Human Understanding—A Reader's Guide, London: Continuum, 2007.

Wang, Dakuo, Elizabeth Churchill, Pattie Maes, Xiangmin Fan, Ben Shneiderman, Yuanchun Shi, and Qianying Wang. "From human-human collaboration to Human-AI collaboration: Designing AI systems that can work together with people." In Extended abstracts of the 2020 CHI conference on human factors in computing systems, pp. 1–6. 2020.

Yang, Qian, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. "Re-examining whether, why, and how human-AI interaction is uniquely difficult to design." In Proceedings of the 2020 chi conference on human factors in computing systems, pp. 1–13. 2020.

Zadeh, Lotfi A. "A new direction in AI: Toward a computational theory of perceptions." AI magazine 22, no. 1 (2001): 73–73.

Zhou, Zhi-Hua. Machine learning. Springer nature, 2021.