

Effective Generative AI: The Human-Algorithm Centaur

Soroush Saghaian^{1,2,3,4,5,6,7,8}, Lihi Idan⁹

¹ Associate Professor, Harvard University

² Visiting Scholar, MIT

³ Faculty Affiliate, Harvard Data Science Initiative

⁴ Founder and Director, Harvard's Public Impact Analytics Science Lab (PIAS-Lab)

⁵ Core Faculty, Harvard Center for Health Decision Science

⁶ Faculty Affiliate, Harvard Mossavar-Rahmani Center for Business and Government

⁷ Faculty Affiliate, Belfer Center for Science & International Affairs

⁸ Faculty Affiliate, Center for Public Leadership

⁹ Postdoctoral Fellow, Harvard University

ABSTRACT. Advanced analytics science methods have enabled combining the power of artificial and human intelligence, creating *centaurs* that allow superior decision-making. Centaurs are hybrid human-algorithm models that combine both formal analytics and human intuition in a symbiotic manner within their learning and reasoning process. We argue that the future of AI development and use in many domains needs to focus more on centaurs as opposed to other AI approaches. This paradigm shift towards centaur-based AI methods raises some fundamental questions: How are centaurs different from other human-in-the-loop methods? What are the most effective methods for creating centaurs? When should centaurs be used, and when should the lead be given to pure AI models? Doesn't the incorporation of human intuition—which at times can be misleading—in centaurs' decision-making process degrade its performance compared to pure AI methods? This work aims to address these fundamental questions, focusing on recent advancements in generative AI, and especially in Large Language Models (LLMs), as a main case study to illustrate centaurs' critical essentiality to future AI endeavors.

Keywords: Centaurs, Generative AI, Large Language Models, Human-AI Algorithms

1. INTRODUCTION

Developing analytics science methods that can enable combining the power of artificial and human intelligence has brought the concept of *centaurs* from myth to reality. In the Greek mythology, centaurs are half-human and half-horse creatures. In modern analytics science, centaurs represent

strong merges between humans and algorithms, referring to hybrid human-algorithm models that combine both formal analytics and human intuition in a symbiotic manner. One of the main users in the U.S. has been the Defense Department, which has been working with tech companies to combine the power of algorithms with the capabilities of humans (Times, 2016b). The concept has attracted the attention of the U.S. military, both in research programs at the Defense Advanced Research Projects Agency and the Pentagon’s third-offset strategy for military advantage (PARC, 2023). Robert O. Work, for example, who was the deputy secretary of defense under Presidents Trump and Barack Obama, advocated for the idea of centaur weapons systems, which would require human control, instead of pure Artificial Intelligence (AI) systems, and could combine the power of AI with the capabilities of humans (Times, 2016a).

The concept of centaurs is not that new, but it received spotlight attention within the analytics science domain because of its success in applications like playing free-style chess. Specifically, prominent advocates of free-style chess like Gary Kasparov repeatedly argued that humans paired with algorithms can do better than just the single strongest computer program in chess (Kasparov, 2017). As the chess legend put it:

“Weak human plus machine plus better process was superior to a strong computer alone and, more remarkably, superior to a strong human plus machine plus inferior process.” (Kasparov, 2010)

Beyond free-style chess, the centaur model is being widely used in a variety of applications of analytics science. In clinical decision-making related to rehabilitation assessment, for example, algorithms provide therapists with detailed analysis on patient’s status, where the collaboration with therapist and such algorithm is shown to improve the practices of rehabilitation assessment (Lee et al., 2021).

Research in the first author’s lab at Harvard, which was conducted in collaboration with the Mayo Clinic, showed very promising results for a centaur model developed to enhance decision-making and reduce readmission risks for a large number of patients who underwent transplantation. The research showed that combining human experts’ intuition with the power of a strong machine-learning algorithm through a human-algorithm centaur model can outperform both the best algorithm and the best human experts (Orfanoudaki et al., 2022).

Other examples of using the centaur model to create public impact include systems for spotting anomalies and preventing cyber-attacks, improving design components in manufacturing systems, and assisting officers balance their workloads and helping them to better ensure public safety (PARC, 2023). And the potential for developing and making use of centaurs is endless. Thus, it is reasonable to expect most data-driven organizations to take advantage of them in the near future. A department of human services, for example, can use algorithms to help predict which child welfare cases are likely to lead to child fatalities and raise a red flag for high-risk cases. Such cases are then reviewed by human experts and the results are shared with frontline staff, who then might choose remedies designed to lower risk and improve outcomes (Deloitte, 2020). The algorithm then can be augmented by using human intuition related to specific cases, creating a human-algorithm centaur.

In this article, we focus on the role of centaurs in recent advancements in Generative AI, and especially in Large Language Models (LLMs). We first present a framework that allows an understanding of the core characteristics of centaurs. We argue that the incorporation of *human intuition* and *symbiotic learning*—a specific learning method based on a human-algorithm symbiosis that we will define and theorize—are key characteristics of centaurs that distinguish them from other AI

models. Using these core characteristics, we also present a few specific methods of creating centaurs. We then argue that the growth and success of LLMs are to a great extent due to the fact that they are moved from pure AI algorithms to human-algorithm centaurs. We present various pieces of evidence to demonstrate this, including the advantages of the so-called “fine-tuning” approaches such as the Reinforcement Learning with Human Feedback (RLHF) method used in various LLMs.

We also discuss evidence showing that symbiotic learning on human intuition data can turn Generative AI tools into cognitive models capable of representing human behavior. We also discuss models in which only some of the conditions for creating centaurs are fulfilled, including In-Context Learning or Chain-of-Thought reasoning methods. We conclude by discussing two main points: (1) recent advancements in creating centaurs have moved us closer to reaching the goals that the founding fathers of AI—John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon—stated in 1955 as part of their proposed 2-month, 10-man study of AI to be held at Dartmouth; and (2) the future of AI development and use in many domains will most likely need to focus on centaurs as opposed to other AI approaches.

2. KEY IDEAS

While the concept of centaurs might sound similar to other hybrid paradigms facilitating collaboration between humans and machines such as Human-In-The-Loop-based (HITP) methods, centaurs propose an inherently different Human-Machine Interaction (HMI) paradigm that creates a *merge* between humans and algorithms as opposed to just a simple *partnership*. Similar to the fact that two businesses might either create a partnership or merge, humans and algorithms can either work alongside each other through a simple partnership (while being two separate entities that may or may not listen to one another) as suggested by HITP-type methods, or combine their powers through a strong merge, as suggested by centaurs. The unique characteristics of centaurs that allow them to facilitate such strong merge-based collaboration is the use of *symbiotic learning* and *human intuition data*. That is, their strength is in both *how* they learn, and *what* they learn, and the existence of *both* components is critical for creating centaurs. Section 5 provides an elaborated definition of centaurs and symbiotic learning, and presents a mathematical framework of centaurs.

Given that the two core components of centaurs are symbiotic learning and learning from human intuition data, a key question that should be posed is whether intuition is at all useful for learning meaningful generalizations. After all, intuition might seem the opposite of rigorous, or logical, or formal (Otte, 1990). The consensus of most works is that while statistical rules indeed lead to more accurate results on unconstrained decision-making problems, intuition can surprisingly yield better results in some constrained decision-making problems, where typical constraints include incomplete data, uncertainty, and limited decision time, among others (L. Huang, 2019; Kahneman, 2011; Khatri & Ng, 2000; Otte, 1990). In data-oriented terms, such constraints may refer to insufficient training instances or the lack of representative training cases, missing values for some of the features in the test case, and limited learning time or inference period. Centaurs, in essence, provide a strong merge that benefits from the *complementary* effect of intuition-based human decision-making and algorithmic decision-making (Orfanoudaki et al., 2022; Saghafian, 2023), providing a single entity that performs well in both constrained and unconstrained situations. Another scenario where human intuition, and specifically, training a model on human intuition data can be useful is when the loss function used prioritizes decisions that coincide with a human-like decision-making mechanism,

even if that does not coincide with the optimal decision to take in the environment. We elaborate on specific cases in which the imitation of human decision-making can be beneficial in Section 7.

Understanding that benefiting from intuition can indeed be useful, the next question is, when can it be trusted for making high-quality decisions? Simon and Kahneman define intuition succinctly as unconscious pattern recognition (Kahneman, 2011; Simon, 1992). Kahneman describes the recognition-primed decision model as follows: When faced with a problem in a domain of expertise, we initially do not try to come up with several possible courses of action and then evaluate which one is the best. Instead, our associative memory automatically draws on a collection of patterns in our subconscious mind and comes up with the closest one (Kahneman, 2011). Based on this definition, Hogarth (Hogarth, 2001), followed by Kahneman and Klein (Kahneman & Klein, 2009; G. Klein, 2015) suggested that intuitive decisions can only be trusted if they reflect a rich and repeated pattern-matching-like experience in environments with reliable feedback. That is, the environment should be stable enough so that past patterns can still be applicable to new ones. Furthermore, meaningful feedback should be constantly provided, and a rich enough set of stored patterns as well as efficient pattern-matching methods should be readily available.

Through symbiotic learning, centaurs ensure the validity of the use of intuition, as the above works suggest. By choosing a model trained on a large set of patterns and with good-enough pattern recognition abilities, the continuous mutual learning process between the human and the machine provides meaningful feedback to both sides. And the fact that centaurs limit the extent to which intuition can influence the model’s weights, mean that they effectively utilize human intuition only to the extent that is beneficial.

The former becomes even stronger when the model used for pattern-matching is a generative model, including the modern architectures of LLMs as we thoroughly elaborate in Section 6. Traditionally, discriminative models were considered preferable to generative models on pattern recognition tasks (Chen, 2015). While generative models were still considered preferable in cases where training sets are small, fast and accurate inference algorithms for complex generative models have been the bottleneck for the application of such models to large-scale instances (Chen, 2015). A major step towards more effective inference in generative modeling has been the Transformer architecture (Vaswani et al., 2017) and its application to LLMs. Notably, recent works have demonstrated impressive results obtained by LLMs on a variety of pattern-recognition problems (S. Huang et al., 2023; Mirchandani et al., 2023). Apart from the high effectiveness of LLMs for pattern recognition tasks, an additional benefit that makes LLMs highly suitable to be used as centaurs is the use of methods such as Reinforcement Learning with Human Feedback (RLHF) (Christiano et al., 2017). These methods allow reliable human-based feedback in the form of shaped rewards while guaranteeing the environment remains stable enough by constraining the differences between two consecutive policies.

While modern Generative AI constructs such as LLMs incorporate a built-in collaboration between humans and algorithms in the form of prompting and In-Context Learning, we note that it is unclear to what extent such collaborations have a true effect on model learning. In fact, existing literature (Min et al., 2022; Webson & Pavlick, 2021) suggests that prompting does not result in real learning, but rather functions merely as a sophisticated “query normalization” method for building better encoding of user queries. That is in contrast to one of symbiotic learning’s inherent traits, the human input having a direct influence on the model’s weights. We stress that this is only one view of prompting, and understanding the exact influence of prompting on LLMs’ learning and inference is still a matter of debate.

Newer models, such as the newest release of OpenAI’s o1 (OpenAI, 2024) take one step forward and incorporate human-machine collaboration not only via ad-hoc prompting, but also as an inherent part of the model’s architecture in the form of human-produced Chains-of-Thoughts (CoTs). While CoTs have been commonly used in an In-Context Learning setting for enriching prompts (CoT Prompting), o1 incorporates such human-produced CoTs within its learning and reasoning process using Reinforcement Learning (RL), where the overall objective is for the model to understand not only what a good outcome is, but also what a good “reasoning process” is. The incorporation of human-produced CoTs within the model’s learning and reasoning process can certainly be seen as one step forward toward creating centaur models; however, we argue that such models are still different than centaurs and, hence, do not provide the strong human-algorithm merge we introduced earlier. The key difference lies in the definition of intuition as suggested by (Kahneman, 2011; Simon, 1992), and specifically on its inherent trait of being “unconscious.” According to (G. Klein, 2015), this fact can be attributed to the reliance of intuition on *tacit knowledge*: knowledge gained through experience. Tacit knowledge lacks structured order, is difficult to explain, and often unavailable to consciousness (Kahneman, 2011; G. Klein, 2015; Simon, 1992). CoT prompting or learning represents the opposite: it is a structured, conscious, and purposely explainable human input that aims at conveying purely logical arguments. Models trained on CoT human input alone, thus, might not be considered full centaurs, since they lack human intuition.

In closing this section, we note that our aim in this paper is mainly (a) starting to raise important questions about human-algorithm centaurs, and (b) providing the first steps for encouraging a wide-ranging debate about their potential in making future AI models more effective. The definitions and arguments we provide, thus, are not black-and-white. Furthermore, they might change over time, especially since it is up to future scholars to create highly effective centaur-based AI models.

The rest of this document is organized as follows. Section 3 presents the motivation for creating and using centaurs. Section 4 characterizes common existing methods for human-machine collaboration, clarifying why centaurs are different. Section 5 presents the notion of symbiotic learning and provides a concrete mathematical framework for creating centaurs. Section 6 discusses why LLMs are highly suitable as base models for creating centaurs, and presents evidence suggesting that centaur-based LLMs can function, at least partially, as cognitive models. Section 7 discusses the advantages and disadvantages of centaur models, their potential use cases, and some ethical considerations. Finally, Section 8 concludes by discussing why future use of AI in various applications will likely be centaur-based.

3. MOTIVATION

Humans often face complex decisions, and it seems that their intuition is not always helpful. When facing critical life-changing decisions such as quitting a job or ending a relationship, we tend to be happier with the outcomes later on, when a coin toss tells us to make a change than when it promotes maintaining the status quo (Levitt, 2021).

Nonetheless, human intuition is often very powerful, especially when we want to make quick decisions. Put it differently, while intuition often misfires when we are dealing with complex problems that require careful analytics (e.g., in predicting whether a particular startup will be successful and is worth large investment (Ang et al., 2022), finding ways to reduce incidents of diabetes for organ transplanted patients (Munshi et al., 2020; Saghafian, 2024), deciding upon cell formation and layout design for a cellular manufacturing system (Saghafian & Akbari Jokar, 2009), or finding

most effective ways of saving lives in emergency rooms (Saghafian et al., 2012, 2024; Traub et al., 2016)), it can be very useful when using data, models, and careful analytics is not an option. Malcolm Gladwell’s popular book “Blink: The Power of Thinking Without Thinking” provides various examples of this, including when police officers need to quickly decide whether to shoot a suspect (Gladwell, 2006). Good intuitive decision-making also helps fire fighters when they face a burning building (G. A. Klein, 2017).

A central question, however, that might arise is this: when and why should one incorporate human intuition in AI-based decision making processes? Wouldn’t it degrade the overall performance? After all, aren’t AI models better than human intuition in many tasks?

While relying on intuition in handling complex problems can be misleading, combining intuition with the most useful analytics approaches can often be better than just relying on analytics. To better understand this, it is useful to see how our own system of thinking works. Daniel Kahneman—the famed psychologist, known for his groundbreaking work on the psychology of judgment and decision-making as well as behavioral economics, who won the Nobel Prize in Economics in 2002—highlighted in his book “Thinking Fast and Slow” that our brain has two modes of thinking: System 1 and System 2. System 1 is fast and instinctive, but System 2 is slower, more deliberative, and more logical. What is perhaps more interesting is that these systems greatly complement each other. Our body somehow knows that we need both systems to be able to make good decisions in different situations.

Similar to how System 1 and 2 complement each other, intuition and analytics can help each other as well. And this is where human-machine centaurs can play a vital role. We—humans—can use our intuition in many ways while developing analytics methods and taking advantage of computers to run them. For starters, intuition often allows us to develop better models, or considering the George Box’s aphorism “all models are wrong, but some models are useful,” more “useful” ones. Intuition also allows us to verify the results obtained from models and make sure that the assumptions made in the model are not problematic. Analytics scientists often use this simple technique as a feedback loop: when the results obtained from a model are not sensible and can be related to a wrong assumption in the model, they modify their model to obtain a better one. And if they are working with a cloud of models to address the curse of ambiguity (Saghafian, 2018; Saghafian & Tomlin, 2016), they can replace the models that might be causing preposterous results. Preposterous results could also be related to abnormalities and/or outliers in data that need to be removed before feeding them to models, and intuition is often very helpful here as well.

The complementary functionalities of formal analytics and human intuition are the key to centaurs’ potential of working better than the best algorithms, even in tasks in which human intuition is relatively weak. In the first author’s experiments with the Mayo Clinic, for example, the ranking in performance (from best to worst) was as follows: human-algorithm centaur, algorithm, and human experts. Notably, the best human expert’s performance was significantly below that of an AI model that was not benefiting from human intuition. Yet, feeding human intuition to it could create a desirable boost in performance (Orfanoudaki et al., 2022; Saghafian, 2023).

Realizing that intuition alone can be misleading in understanding and analyzing complex systems, and that human-machine integration is needed to harness the full power of both advanced analytics and human intuition, has introduced new methods of creating human-algorithm centaurs capable of symbiotic learning, as will be discussed in Section 5.

3.1. Main Advantages of Centaurs Compared to Other AI Models. There are various reasons why developing and implementing centaurs should be on top of mind in different sectors:

1. **Enhanced interpretability:** While some models such as decision trees are inherently interpretable, other common models such as neural networks lack interpretability. Incorporating human intuition into the learning process of models may allow humans to better understand their predictions as well as the inference process that led to those predictions, thus increasing the interpretability level of naturally uninterpretable models. This observation was empirically validated by (Lage et al., 2018), which shows that explicitly including human feedback in the algorithm’s learning process improves the algorithm’s interpretability on some datasets even to a larger extent than methods imposing direct interpretability constraints.

2. **Reduced algorithm aversion:** Algorithm aversion has been widely reported as a main hindrance in creating impactful AI in applications in which the final decision-maker is a human (Burton et al., 2020; Dietvorst et al., 2015, 2018). Through experiments, we have found that utilizing centaurs can reduce algorithm aversion, mainly because recommendations from centaurs are closer to human way of thinking since centaurs incorporate human intuition (Orfanoudaki et al., 2022). For example, using the “weight on advice measure,” we found that physicians are more willing to consider recommendations that come from centaurs than those from pure AI models (Orfanoudaki et al., 2022). Furthermore, the recommendations coming out of centaurs represent less “human aversion” in that, not surprisingly, they better represent human behavior and preferences.

3. **Reduced causation aversion:** Centaurs can help AI developers overcome what we might call “causation aversion.” This refers to the fact that most of the focus of the current AI algorithms has been on the association layer of the “Ladder of Causation” (Pearl & Mackenzie, 2018). What decision-makers in various sectors need are algorithms that can work beyond the association between various input and output variables. Prescriptive analytics, where the intention is to recommend decisions that can cause improvements, can benefit a lot from moving algorithms to human-algorithm centaurs. Given various ongoing research on causal ML (see, e.g., (Saghafian, 2024) and the references therein), one can expect to see more efforts in developing centaurs capable of causal and counterfactual reasoning.

4. **Enhanced adaptability to behavioral tasks:** Centaurs achieve better results on behavioral tasks—tasks that require the model’s output to coincide with human behaviors such as human-study replication (Ziems et al., 2024). Likewise, centaurs show better results on tasks evaluated using behavioral evaluation metrics (Ouyang et al., 2022; Stiennon et al., 2020). For instance, chatbots are often evaluated according to the extent to which human users found them “helpful”. We further elaborate on the use of centaurs for behavioral tasks in Section 7.

5. **Better specification of hard-to-define objectives:** Traditionally, AI models have been applied to tasks for which it is quite easy to come up with well-specified objective functions. However, as these models have been widely adopted in numerous areas and for a wide array of diverse tasks beyond the classification of well-structured data, many of such newer, more complex tasks involve goals that are hard to define, and thus it is not immediately clear how to design a cost function for such tasks. Instead of explicitly defining an objective function for such tasks, the objective can be dynamically learned using human intuition in the form of preferences or demonstrations. The model then follows the human-guided, learned objective function (Christiano et al., 2017; Ibarz et al., 2018).

6. **Enhanced performance on low-quality datasets:** A key trait of AI models is the ability to generalize from one dataset, a training set, to other unseen datasets. However, training datasets

may contain flaws that inhibit the model’s generalization ability. For instance, the training set at hand can be unrepresentative of the general population. Alternatively, the training dataset may be representative of the general population at *train time*, t_{train} , but not at *deployment time*, $t_{test} = t_{train} + t'$, (*i.e.* concept drift). Human intuition can be used to mitigate such generalization caveats by adding more nuanced yet general reasoning into the model, enabling it to generalize to unseen, potentially novel, examples (Agarwal et al., 2024; Grønsund & Aanestad, 2020; Idan, 2022; Knoche & Rigoll, 2023; Langlotz, 2019).

4. CENTAUR VERSUS HUMAN-IN-THE-LOOP APPROACHES

In 1960, J.C.R. Licklider published a paper entitled “Man-Computer Symbiosis” in which he popularized the idea of man-computer symbiosis “as an expected development in cooperative interaction between men and electronic computers” (Licklider, 1960). In his words, the aim was to “enable men and computers to cooperate in making decisions and controlling complex situations without inflexible dependence on predetermined programs.” In today’s analytics science, we can think of symbiotic learning and the incorporation of human intuition as two important characteristics of centaurs. These two characteristics distinguish centaurs from other related designs in AI such as those under the “human-in-the-loop” umbrella. We discussed the advantages of incorporating human intuition in the previous section. In the next section, we will discuss the other distinctive feature of centaurs: symbiotic learning. Here, we start by reviewing human-in-the-loop approaches so as to distinguish them from centaurs.

Active-learning approaches. Some works (see, e.g., (Settles, 2011)) model the human as an *oracle*: given a query on a subset of unlabeled data records, issued by the AI model, the human responds with the labels of the data records in the query. The human in this case is used as a “data annotator.” After the query’s response is submitted back to the AI model, the annotated records are used to improve the model’s performance by conducting either full or partial retraining of the model using the annotated records. Active learning approaches are particularly useful when the vast majority of the available data is unlabeled and the cost of annotations is high. Thus, it is a popular approach in domains such as healthcare and finance (Carcillo et al., 2018; Gal et al., 2017; Lorenz et al., 2020; Warmuth et al., 2003).

Interactive-learning approaches. Some other works (see, e.g., (Ware et al., 2001)) also model the human as an oracle, but this time the human itself can actively control *which* queries are submitted to her and on which ones she will respond with annotations. Unlike an active-learning process which may or may not be iterative, interactive-learning algorithms are both iterative and incremental in nature (Amershi et al., 2014; Fails & Olsen Jr, 2003); the human is actively participating not only in the annotation process but also in the query creation process and/or in the model evaluation process aimed at assessing whether more queries are needed (Berg et al., 2019; Heimerl et al., 2012).

Workload partitioning approaches. In this set of approaches, one typically divides the prediction workload between the human and the machine based on pre-selected criteria. That is, each record is processed either by the machine, the human, or both the machine and the human where the goal is to minimize the number of records that humans are required to process. This approach is oftentimes used in domains in which, on the one hand, experts’ annotation time is expensive, yet, on the other hand, the degree of error tolerance is low. In such case, the minority of “complex” cases will be processed by the human whereas the majority of cases will be processed

by the machine. The allocation of records either to the human expert or to the model can be done either based on manually selected criteria such as the risk associated with the record, or based on automatic criteria in which the model itself can decide who should be assigned with the record (e.g., based on its confidence level on the record). Workload partitioning schemes are thus oftentimes deployed in high-risk domains, such as applications related to security and health. While in some works humans and machines share the same workload (Cranor, 2008; Reverberi et al., 2022; Zhalechian et al., 2024), in most workload-partitioning works the majority of the workload is put on the AI algorithm, and only the minority of harder or riskier tasks are assigned to humans (Cohen et al., 2023; Hanczar & Dougherty, 2008).

Machine teaching approaches. The main human functionality in machine-teaching schemes is typically in designing an optimal training set for the AI algorithm (Goldman & Kearns, 1995; Simard et al., 2017). Importantly, machine-teaching approaches often assume that the human already knows the “true” decision boundary, and its goal is merely to convey that boundary to the model via the selection of a proper training set. Of course, the human can also use its control on the training set to guide the model towards a different decision boundary than the true one. Indeed, most applications of the machine-teaching approach are either in the security domain, which uses machine teaching to construct both attacks and defense schemes on automated learning agents (Mei & Zhu, 2015a, 2015b) or in the education domain, which uses machine teaching to build models that are better at conveying information to human learners (Patil et al., 2014; Zhu, 2015).

Curriculum learning approaches. These set of approaches are similar to machine teaching approaches in that the main human functionality is organizing the training set used by the AI model (Bengio et al., 2009). Unlike machine teaching, the human in curriculum-learning schemes is not assumed to know the true decision boundary. Instead, curriculum learning assumes the human to have a close familiarity with the prediction task at hand. Using that familiarity, and given a fixed set of training data, the human’s goal is designing an optimal curriculum based on the fixed training data. The idea behind curriculum learning is that the learning process can be quicker if the examples are not randomly presented but rather organized in a meaningful order, where the common wisdom is that “easier” examples should be presented before “harder” examples (Bengio et al., 2009). Defining a proper “difficulty” measure is thus the main challenge when designing a curriculum-learning scheme, and different works use different difficulty measures. Common difficulty metrics in the Natural Language Processing (NLP) domain include sentence length (Platanios et al., 2019), parse tree depth (Tsvetkov et al., 2016), and level of word rarity (Kocmi & Bojar, 2017). Common difficulty metrics in the vision domain include number of objects (Y. Wei et al., 2016); shape variability (Bengio et al., 2009); and inter-image clustering density (Guo et al., 2018).

Meta-Learning approaches. In meta-learning approaches, human feedback is neither used for preprocessing actions such as training set organization nor is combined in the learning process. Instead, human feedback is used for different meta-learning tasks such as model evaluation, hyperparameter tuning, and model calibration, among others. Importantly, such feedback is not used to further train the model, but merely to further improve the model without directly combining the feedback in the training process. For example, (Idan, 2024) builds an anomaly-detection ensemble and uses human experts’ opinions on a subset of query records for tuning the ensemble’s hyperparameters: the weight assigned to each model in the ensemble. Likewise, (Corvelo Benz & Rodriguez, 2024) presents a human-preference-based calibration approach, suggesting calibration with respect to the human’s confidence in her own predictions produces calibrated estimates that lead to optimal decisions. Finally, (Chang et al., 2009; Denton et al., 2015) use humans as model

evaluators, where humans observe the learned model’s outputs on a representative set of inputs and evaluate its output quality by ranking to what extent it aligns with their subjective preferences.

Other approaches. More recent methods for human-machine collaboration such as Human-Guided Rewards, Human-based Prompt Engineering, Chain-of-Thought Reasoning, Supervised Fine-Tuning on Human Intuition Data, and more, are discussed in Section 5 and Section 6.

In almost all of the approaches discussed above, the human functions mainly as a “service provider:” she is either treated as an oracle, providing labeling services, or as a teacher, providing preprocessing services such as training set design. Notably, the human is not treated as an *equal* of the learning model who can directly incorporate her own intuition—in the observed form of preferences, actions, decisions, or opinions—as part of the learning process. Such observed signals of human intuition result from the human’s unique judgment, which is subjective in nature, as opposed to the human’s perception as an oracle where the human input is assumed to always coincide with a certain objective truth (e.g. ground-truth labels). In those human-in-the-loop approaches where humans are not merely used as service providers but are rather queried for their own subjective input, such inputs are not combined directly in the learning algorithm. Instead, they are used *orthogonally* to the AI model, either by working on a different set of records than the machine (workload partitioning approaches) or by providing subjective feedback that is not incorporated directly in the learning process (meta-learning approaches).

As we will discuss in the next section, symbiotic learning which is a core element of centaurs makes them different from these human-in-the-loop approaches by taking a broader approach to the cooperation between humans and machines, treating both parties as equals rather than a-priori giving the lead to one party. Such a holistic approach resembles the concept of Human-Intelligence Collectives (HIC) (Jennings et al., 2014; Muller, 2022), which researches the creation of human-agent teams and the best cooperating and control methods for achieving a predefined set of goals that may defer among team members. Importantly, HIC considers each member of the team to be an autonomous stakeholder, having its own set of goals which may not coincide with other members’ goal. Thus, HIC mainly focuses on the research of new incentive engineering and reward modeling techniques, rather than on the research of mutual human-AI learning methods under commonly shared goals as is the case with centaurs.

5. SYMBIOTIC LEARNING: A CORE CHARACTERISTIC OF CENTAURS

As mentioned in previous sections, the distinguishing characteristic of centaurs from human-in-the-loop approaches is symbiotic learning. Symbiotic learning departs from the methods discussed in Section 4 in that it implements a mutual human-AI learning process where the human is not treated merely as an oracle or a preprocessor, but rather as an equal partner in the learning process by incorporating her unique input into the learning process. Human input includes observed signals of human intuition such as preferences, behaviors, or opinions, and unlike ground-truth labels, such inputs are inherently subjective. The core differences between human-in-the-loop approaches as described in Section 4 and symbiotic learning are three-fold:

1. Symbiotic learning facilitates multi-input learning, in which learning is performed using multiple datasets corresponding to either the same set of observations or different sets of observations. At least one dataset represents human input, and includes observed signals of human intuition such as subjective preferences, behaviors, opinions, or decisions.

2. The human input is directly incorporated in the learning process of the downstream model’s parameters (e.g., weights), which we denote by $\theta_{\text{symbiotic}}$.

3. The influence of human input data on the final model’s parameters is constrained to ensure appropriate reliance on human intuition given the specific traits of the environment, the degree of human expertise, and feedback quality. The constraints are optimized either as a separate learning task or jointly with $\theta_{\text{symbiotic}}$.

Symbiotic learning defines a new paradigm, where humans do not function merely as service providers supporting the AI model, but rather as full partners to the model in the learning process. To account for a wide array of tasks and applications, the nature of such partnership is learned on a case-by-case basis using flexible constraints. A high-level diagram of symbiotic learning is given in Figure 1. In what follows, we present a mathematical framework that further details this representation.

We frame symbiotic learning as a general constrained optimization problem, abstracting away specific symbiotic learning techniques using different loss functions, models, and constraints. A parametric example of our symbiotic learning framework is:

$$(5.1) \quad \hat{\theta}_M = \arg \min_{\theta \in \Theta} \mathbb{E}_{z \sim d_{\text{data}}} l_{\text{data}}(z; \theta)$$

$$(5.2) \quad \hat{\theta}_H = \arg \min_{\theta \in \Theta} \mathbb{E}_{z \sim d_{\text{human}}} l_{\text{human}}(z; \theta)$$

$$(5.3) \quad \hat{\theta}_{\text{symbiotic}} = \gamma_{\text{symbiotic}}(\hat{\theta}_M, \hat{\theta}_H) = \arg \min_{\gamma \in \Gamma, \mathcal{C}(\gamma, \hat{\theta}_M) \leq c_1, \mathcal{C}(\gamma, \hat{\theta}_H) \leq c_2} \mathbb{E}_{z \sim d_{\text{human}}, d_{\text{data}}} [l_{\text{comb}}(z; \gamma; \hat{\theta}_M; \hat{\theta}_H)]$$

$$(5.4) \quad h_{\theta_M} : \mathcal{X}_M \rightarrow \mathcal{Y}_M, h_{\theta_H} : \mathcal{X}_H \rightarrow \mathcal{Y}_H, h_{\theta_{\text{symbiotic}}} : \mathcal{X}_S \rightarrow \mathcal{Y}_S$$

Given a data distribution p_{data} and a corresponding empirical sample d_{data} , we train a base model parametrized by $\hat{\theta}_M$ to minimize an arbitrary loss function, l_{data} , on d_{data} . Given a human-preference distribution p_{human} and a corresponding empirical sample d_{human} , we train a human preference model parametrized by $\hat{\theta}_H$ to minimize an arbitrary loss function, l_{human} , on d_{human} . We denote the space of symbiotic-model parameters by Γ , and model symbiotic learning as minimizing an arbitrary combination loss l_{comb} on either d_{data} or d_{human} w.r.t symbiotic learning parameters $\gamma \in \Gamma$ and subject to constraints, \mathcal{C} , on either $\hat{\theta}_M$, $\hat{\theta}_H$, or both. Finally, the learned parameters can be used to learn the functions denoted by h in the last line, where a mapping between the space of inputs to that of outputs occurs. In this setting, loss functions are quite general and include various aspects such as regularizations that is often used in high-dimensional settings to avoid over-fitting.

In practice, different symbiotic learning techniques use different instantiations of the definition in Equations (5.1)-(5.4). For instance, in human-preference-based supervised fine-tuning approaches no θ_H is trained, and thus l_{comb} is optimized solely on d_{human} given constraints on $\hat{\theta}_M$. In preference-based augmented-covariate-space approaches, no θ_M is trained and thus l_{comb} is solely optimized on d_{data} given constraints on $\hat{\theta}_H$. In preference-constrained cost function approaches, neither θ_M nor θ_H are trained, and thus the symbiotic learning framework collapses into a simple constrained loss function using a tunable hyperparameter that controls the importance given to human preferences within the loss. In human-machine ensemble approaches, both θ_H and θ_M are trained, and l_{comb} is optimized on d_{data} .

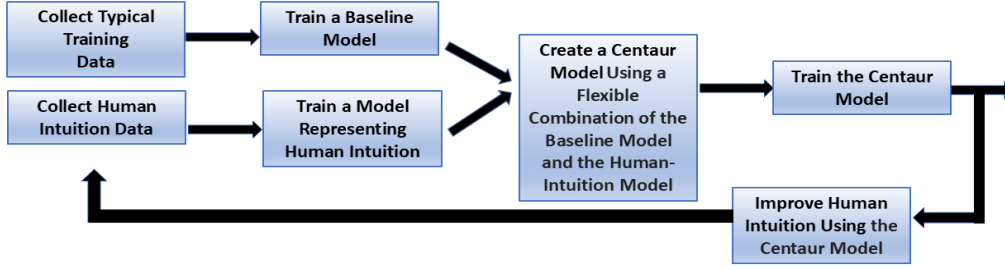


Figure 1. A high-level representation of symbiotic learning

Distributions d_{data} and d_{human} can also take different forms in different symbiotic learning techniques. Generally speaking, in supervised learning applications of symbiotic learning such as the one discussed above we assume d_{data} to contain pairs of the following form: $z_i = (x_i, y_i) \in \mathcal{X}_{data} \times \mathcal{Y}_{data}$ and d_{human} to contain pairs of the following form: $z_i = (x_i^h, y_i^h) \in \mathcal{X}_{human} \times \mathcal{Y}_{human}$. In some symbiotic learning techniques, both d_{data} and d_{human} must represent the same underlying distribution, while in other techniques they may represent two different distributions. For instance, an underlying assumption of preference-based augmented-covariate-space approaches (see below) is that d_{data} and d_{human} represent the same underlying distribution, where $\mathcal{X}_{data} = \mathcal{X}_{human}$, $\mathcal{Y}_{data} = \mathcal{Y}_{human}$, whereas in human-preference-based supervised fine-tuning approaches (see below) d_{data} and d_{human} can take different distributions. In fact, in the latter case d_{human} is often labeled, where y_i^h corresponds to human feedback, whereas d_{data} is unlabeled, and thus $\mathcal{Y}_{data} = \emptyset$.

We now discuss multiple concrete techniques for the implementation of symbiotic learning. These techniques are specific ways of utilizing the general framework presented via Equations (5.1)-(5.4).

Preference-based augmented covariate space. In this approach, γ refers to the learned parameters (e.g., weights) of a model, $h_\gamma: \mathcal{X}_S \rightarrow \mathcal{Y}_S$, where \mathcal{X}_S is an augmented version of \mathcal{X}_{data} which contains for each record, i , its corresponding human feedback y_i^h extracted from d_{human} so that h_γ is fed, for each record, $i \in d_{data}$, the following pair: $(x_i \cup \{y_i^h\}, y_i)$. The new set of features for a record, i , contains its original set of features, x_i , as well as the human feedback on that record, y_i^h .

A practical caveat of that approach is that human feedback is needed on each record in d_{data} , both in training and in production, and hence the following must hold $|d_{human}| \geq |d_{data}|$, which is impractical in most scenarios. An alternative approach is to collect human preferences on a different, smaller dataset, so that $|d_{human}| \ll |d_{data}|$ while d_{human} is highly representative of d_{data} . To this end, for for each $i \in d_{data}$ one can search for its k-nearest neighbors in d_{human} , and use a combination of their assigned feedback (y_i^h), such as a majority vote, as an additional feature for i in h_γ . However, it is not immediately clear how to choose the best distance metric between samples in d_{data} and samples in d_{human} , and the results might be sensitive to the choice made for this metric.

A generalization of the nearest-neighbor approach presented in (Orfanoudaki et al., 2022) is to implicitly learn such similarity relations by building a human-preference model and use its prediction as an additional feature to h_γ , as opposed to using the raw preferences y_i^h . That is, one can first learn a human preference model h_{θ_H} using d_{human} :

$$(5.5) \quad \hat{\theta}_H = \arg \min_{\theta \in \Theta} \mathbb{E}_{z \sim d_{human}} l_{human}(z; \theta).$$

For each $i \in d_{data}$, one can then use the value $h_{\theta_H}(x_i)$ as an additional feature for learning a model h_γ , trained to minimize the loss over d_{data} . In this setting, the set of constraints would be $\mathcal{C}(\gamma, h_{\theta_H})$ which denotes constraints on the importance assigned to feature $h_{\theta_H}(x_i)$ during the learning process compared to other features in \mathcal{X}_{data} .

Human-guided rewards. As in the preference-based augmented covariate space approach, this approach starts by learning a human preference model, h_{θ_H} , using the human preference data, d_{human} :

$$(5.6) \quad \hat{\theta}_H = \arg \min_{\theta \in \Theta} \mathbb{E}_{z \sim d_{human}} l_{human}(z; \theta).$$

One then uses the outputs of $\hat{\theta}_H$ on samples in d_{data} for estimating a newly-learned symbiotic learning model, h_γ with newly-learned parameters, γ . However, unlike the preference-based augmented covariate space approach, the outputs of h_{θ_H} are not embedded within h_γ 's feature space, \mathcal{X}_S , as additional features, but rather used either consistently or sporadically to guide h_γ towards solution regions compatible with the human input in d_{data} . Moreover, unlike prior symbiotic learning schemes where the exact form of $\hat{\theta}_H$, as well l_{human} used for training it was rather flexible, here the human preference model, $\hat{\theta}_H$, also referred to as the reward model, must have a specific structure. First, a reward model must take as input pairs (x, y) where x denotes a context and y denotes a possible output. Second, the output of the reward model must be a scalar, denoting the predicted human-preference score given to output y created in the context of an input x . Thus, $\mathcal{X}_H = \mathcal{X}_S \times \mathcal{Y}_S$ and $\mathcal{Y}_H = \mathbb{R}$. $\hat{\theta}_H$ is trained to maximize the rewards of outputs that are highly preferred by humans while minimizing the reward of outputs that are less preferred by humans. Specifically, l_{human} often takes the following general form:

$$(5.7) \quad l_{human}(\hat{\theta}_H) = \mathbb{E}_{(x, y_1, y_2) \sim d_{human}} [f_1(f_2(\hat{\theta}_H(x, y_1) - \hat{\theta}_H(x, y_2)))],$$

where the training set d_{human} is composed of triplets (x, y_1, y_2) : x denotes a context (an input to h_γ), y_1 denotes the output that is more preferred by humans in the context of x , and y_2 denotes the output that is less preferred by humans in the context of x . f_1, f_2 are task-specific and may correspond, for instance, to the $\log()$ and $\text{sigmoid}()$ functions, respectively.

In some configurations, $\hat{\theta}_H$ is learnt first, and then used for learning h_γ . In other configurations, both $\hat{\theta}_H$ and h_γ are jointly learned, where the output of h_{θ_H} is used as rewards to guide h_γ 's learning process and optimizing l_{data} while at the same time outputs of h_γ are used to further tune $\hat{\theta}_H$ using l_{human} . While human-guided rewards can be used for aligning different types of AI schemes, its most common use is for replacing or augmenting the reward function in reinforcement learning, creating Reinforcement Learning with Human Feedback (RLHF) schemes (Christiano et al., 2017; Ouyang et al., 2022). As we will discuss later, RLHF in turn has played an important role in moving various generative AI models, including a vast majority of LLMs, from a pure AI model to a centaur-based system by incorporating human intuition and symbiotic learning—the two distinguishing characteristics of centaurs.

Using human-guided rewards as opposed to a fixed reward function is especially useful for optimizing tasks in which it is easier to recognize the desired behaviors than explicitly defining them.

In an RLHF-based setting, l_{comb} , which denotes the policy’s loss, takes the following general form:

$$(5.8) \quad l_{comb}(\gamma) = \mathbb{E}_{(x,y) \sim D_{h_\gamma}} [\hat{\theta}_H(x, y) - \mathcal{C}(h_\gamma(y|x), h_{\hat{\theta}_I}(y|x))],$$

where $x \in \mathcal{X}_{data}$, $y = h_\gamma(x)$, and \mathcal{C} denotes a constraint applied to the symbiotic model, aimed at preventing it from “overfitting” to human preferences by limiting its divergence from an initialization model, $\hat{\theta}_I$. We will elaborate on the use of human-guided rewards for LLM training later in this section where we use GPT-4 training as a specific example.

Human-preference-based supervised fine-tuning. While the term “fine-tuning” is nowadays mostly used in the context of LLMs, it is more broadly a re-branding attempt of a collection of well-known, well-established transfer-learning-based methods (Torrey & Shavlik, 2010). Nonetheless, some newer techniques are designed specifically for applications in LLMs as we will discuss in Section 6. A typical supervised fine-tuning task consists of taking an existing model trained on a given labeled or unlabeled dataset, and adapting it to another task using another labeled dataset. Fine-tuning a model with human feedback can, thus, be thought of as first training a model h_{θ_M} using d_{data} :

$$(5.9) \quad \hat{\theta}_M = \arg \min_{\theta \in \Theta} \mathbb{E}_{z \sim d_{data}} l_{data}(z; \theta),$$

and then adapting h_{θ_M} to human preferences using another dataset, d_{human} :

$$(5.10) \quad \hat{\theta}_{symbiotic} = \arg \min_{\gamma \in \Gamma, \mathcal{C}(\gamma, \hat{\theta}_M) \leq c_1} \mathbb{E}_{z \sim d_{human}} [l_{human}(z; \gamma; \hat{\theta}_M)].$$

In the supervised case, γ corresponds to a small set of parameters *within* the same model h_{θ_M} so that most parameters are frozen using the initialization of θ_M and only few parameters, γ , can change during fine-tuning using d_{human} subject to the constraints enforced by \mathcal{C} . For example, when using adapter layers (Houlsby et al., 2019; Rebuffi et al., 2017), due to the high training cost only parameters corresponding to the last layers will be tuned while all other parameters are frozen according to θ_M ’s initialization. Details on supervised fine-tuning can be found in Section 6.

Human-machine ensembles. Ensemble-based approaches view both the human and the machine as experts and aim at optimizing a joint model of the two types of experts (Cha & Lee, 2021; Wang et al., 2016). Specifically, we are given N heterogeneous machine learning models, indexed by $n = 1, 2, \dots, N$ and trained using d_{data} :

$$(5.11) \quad \hat{\theta}_M^n = \arg \min_{\theta \in \Theta_n} \mathbb{E}_{z \sim d_{data}} l_{data}(z; \theta),$$

as well as K heterogeneous human-preference models indexed by $k = 1, 2, \dots, K$ and trained using d_{human} :

$$(5.12) \quad \hat{\theta}_H^k = \arg \min_{\theta \in \Theta_k} \mathbb{E}_{z \sim d_{human}} l_{human}(z; \theta).$$

The goal is to construct an ensemble model $h(\hat{\theta}_M^1, \dots, \hat{\theta}_M^N, \hat{\theta}_H^1, \dots, \hat{\theta}_H^K)$ such that h effectively combines $\hat{\theta}_{M_{n \in 1, \dots, N}}^n$ and $\hat{\theta}_{H_{k \in 1, \dots, K}}^k$ where h is either a fixed combination function or a learnable model h_γ with learnable weights, γ .

In the simplest form, h_γ functions as a meta-learner using a stacking approach that learns the optimal weights of $\hat{\theta}_M$ and $\hat{\theta}_H$ in the final prediction, or further learns how to combine $\hat{\theta}_M$ ’s and $\hat{\theta}_H$ ’s predictions. In more complex configurations, the contribution of the human-preference model to h_γ is further constrained. An example of such a constrained ensemble setting is the GPT-4 architecture

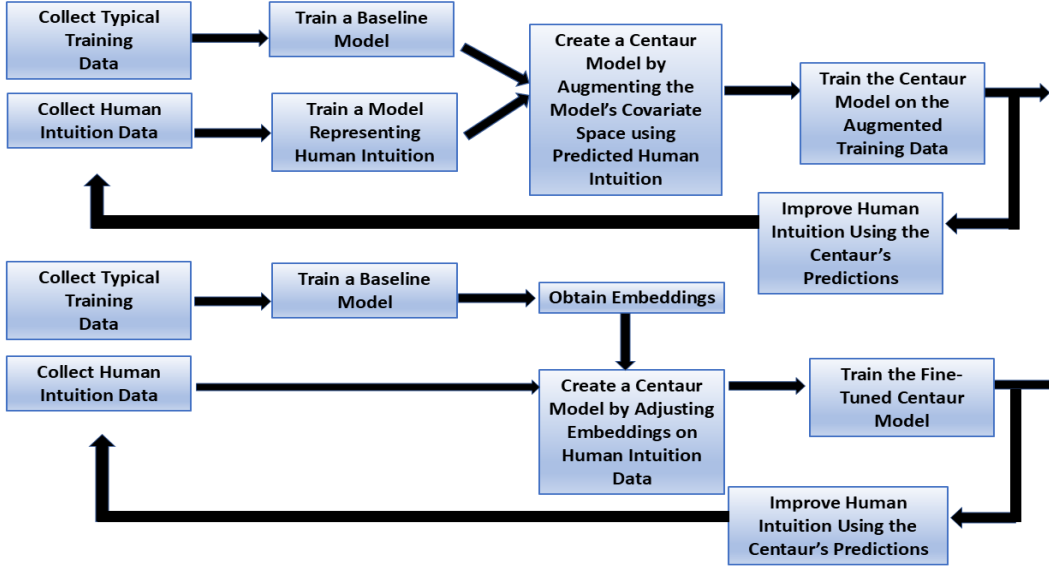


Figure 2. Two modern methods of creating centaurs. Top: Preference-based augmented covariate space. Bottom: Human-preference-based supervised fine-tuning

(OpenAI, 2023), in which $\hat{\theta}_M$ corresponds to a pre-trained language model, $\hat{\theta}_H$ corresponds to a learned human-based reward model, and γ are the parameters of a fine-tuned model, h_γ . In GPT-4, γ is learned using the pre-trained language model and the reward model while constraining the current policy, resulting from rewards provided by $\hat{\theta}_H$, not to diverge too much from $\hat{\theta}_M$. This later step of constraining the current policy is operationalized using a KL-divergence penalty applied on a per-token basis between h_γ and $\hat{\theta}_M$.

Another example of such a constrained ensemble setting is reward-model ensembles (Zhang et al., 2024), an emergent solution to the challenge of reward over-optimization often seen on out-of-distribution datasets. In such a setting, there exists K reward models, representing diverse sources of human intuition-based feedback. Each reward model starts from a pre-trained initialized model on d_{data} , $\hat{\theta}_M$, and then fine-tunes it on its own d_{human} dataset to a different extent, which allows capturing various degrees of tradeoff between machine-based and human-based intuition. Here, γ is the parameters of each model's Low Rank (LORA) adapters (Hu et al., 2021) which effectively combine all the fine-tuned reward models $\{\hat{\theta}_H^k(\hat{\theta}_M)\}_{k=1,\dots,K}$ into a single ensemble model, h_γ .

Preference-constrained cost functions. Unlike the former approaches in which the human feedback is implicitly combined in $\theta_{symbiotic}$ using a human-preferences model, when using direct cost-function adaptation methods the human feedback is explicitly combined in the model's cost function (Correia & Lecue, 2019). Hence, cost-function adaptation can be seen as a flattened version of Equation (5.3) where there exist no θ_H, θ_M and thus:

$$(5.13) \quad \hat{\theta}_{symbiotic} = \arg \min_{\gamma \in \Gamma} \mathbb{E}_{z \sim d_{human}, d_{data}} [l_{comb}(z; \gamma)],$$

which, in its simplest form, translates into:

$$(5.14) \quad l_{comb_i}(x_i, y_i; \gamma) = \arg \min_{\gamma \in \Gamma} l_1(y_i, h(x_i, \gamma)) + \lambda \times l_2(f_1(y_i^h), f_2(h(x_i, \gamma))).$$

In Equation (5.14), the constraints $\mathcal{C}(\gamma, \hat{\theta}_M) \leq c_1$, $\mathcal{C}(\gamma, \hat{\theta}_H) \leq c_2$ introduced earlier have been translated into a single hyperparameter, λ , that explicitly controls the influence of human preferences on the overall learning process.

A simple application of Equation (5.14) is when $l_1 = l_2$, and $f_1 = f_2 = \mathcal{I}$ where \mathcal{I} represents the exact form of the loss function. In such a case, data points y_i^h denote human-preference-based labels, which are used to ensure that the predictions align with those preferences. Such an approach will work well if, for instance, there exists multiple, equally-performing decision boundaries, and we choose the one that also highly aligns with human preferences. However, by parametrizing l_{comb_i} using both l_1 and l_2 , as well as both f_1 and f_2 , the loss can express a richer set of relations between d_{data} and d_{human} . An example is when human preference input is not provided on the actual records' labels, but rather on indirect components of the prediction process such as recommended features. In such a case, f_2 denotes the important features extracted from a learned model h_γ , l_1 denotes a standard loss function such as cross entropy while l_2 denotes a distance metric between the two sets of features.

6. RECENT ADVANCEMENTS IN LLMs AND THEIR COGNITIVE ABILITIES: EVIDENCE ON THE POWER OF CENTAURS

In recent years we have seen a fast growth in the development and use of Generative AI. Notable examples include efforts by OpenAI with GPT (Generative Pre-trained Transformer) and DALL-E (“Dali” and “Eve,” from artist Salvador Dali and the character Eve from the Pixar movie WALL-E), Meta with Llama (Large Language Model Meta AI), Google with Bard that uses LaMDA (Language Model for Dialog Applications) and more recently PaLM (Pathways Language Model), Microsoft with Bing Chat, Stability AI with Stable Diffusion, Github (and OpenAI) with GitHub Copilot, and Anthropic with Claude.

In 2022, for example, LLMs saw a major advancement: ChatGPT—a large language model chatbot developed by OpenAI based on GPT-3.5 with the ability to provide conversational responses that can appear surprisingly human. Like other language models, the main idea behind ChatGPT is simple: to predict the next word in a sentence or phrase based on the context of previous words, using a model trained on a large number of instances. GPT-3 had about 175 billion machine learning parameters, and some estimates showed that it consumed about 936 Mwh to train (Register, 2024)—equivalent of about 30k American households power usage in a day. While recent improvements have focused on making GPT-3 more efficient by reducing these numbers, other works focused on further improving its reasoning abilities. For instance, the “chain of thought prompting” method, first presented by Google Brain (J. Wei & Zhou, 2022), enabled language models of sufficient scale (e.g., models with 100 billion parameters) to solve complex reasoning problems that are not solvable with standard prompting methods. In 2023, OpenAI introduced GPT-4, which is not only multimodal (e.g., accepts images), but according to the creator “exhibits human-level performance on various professional and academic benchmarks” (OpenAI, 2023).

Though LLMs have demonstrated extraordinary performance, the exact mechanism behind those impressive achievements is not entirely clear and has been the subject of an ongoing debate. Another, topic sparking ongoing debate is whether, and under which terms, LLMs can be used as *cognitive models*. A cognitive model is defined as “an approximation of one or more cognitive processes in humans or other animals for the purposes of comprehension and prediction” (Sun, 2008). Thus, a model’s cognitive abilities can be captured either internally, by the extent to which the

model’s internal learning and inference processes—captured via various neural activities—resemble those of humans; or externally, by the extent to which the model’s output follows human behaviors, decisions, and preferences. We claim that (a) modern LLMs already possess traits that highly resemble important human cognitive processes, and (b) a main reason for this is that they have moved from pure AI models to centaurs by benefiting from direct incorporation of human intuition as well as symbiotic learning—the two main characteristics of centaurs. Building on these traits and enhancing them via more advanced symbiotic learning approaches or further incorporation of human intuition is an important building block for future endeavors in turning LLMs into much stronger cognitive models, capable of further representing human behavior.

The first cognitive-like mechanism underlying LLMs is *attention*. Attention is a fundamental mechanism of cognition needed to perform complex reasoning tasks. Modern cognitive research defines attention as the capacity to select and enhance certain aspects of currently processed information while suppressing the remaining aspects (Allport, 1989). When one attends to an object, she implicitly assigns a higher importance score to that object compared to other objects in the environment. The higher-scored objects receive preferential processing, leading to an increased neural response which results in improved functionalities on the attended object such as better memory storage (Desimone & Duncan, 1995).

Relative importance-score assignment is also the core objective of the attention mechanism used as the key novelty in the Transformer architecture—the neural-network (NN) architecture underlying the vast majority of modern LLMs (Vaswani et al., 2017). The core idea of artificial attention—a set of comparisons to relevant items in some context, a normalization of those scores to provide a probability distribution, followed by a weighted contextualization based on those scores—closely resembles the human cognitive attention process, in which humans assign importance scores to objects in the environment, where those scores are both relative and dynamically learned by the human.

Furthermore, the usage of attention for creating contextualized representations of objects based on prior relevant ones resembles another cognitive process, *priming*, occurring when the exposure to one stimulus influences a response to a subsequent stimulus. Specifically, the Transformer’s use of attention weights for the creation of contextualized representations resembles a specific type of priming, Contextual Priming (Stanovich & West, 1983) in which humans use a context to speed up processing for stimuli that are likely to occur in that context. Indeed, studies in human language confirm that while reading, people will look back at previous sections of text in order to clarify what they are currently reading (Coltheart, 2016; Myachykov & Posner, 2005).

While the attention mechanism as used in modern LLM architectures resembles human cognitive processes in many respects, artificial attention mechanisms as described above and cognitive attention vary on a number of dimensions. For instance, while cognitive attention is manifested via both top-down and bottom-up attention mechanisms, artificial attention mechanisms such as the one used in the Transformer architecture are inherently goal-driven and are thus solely top-down based. Moreover, the attention model as currently used in most NN architectures computes an attention score for a given query for every possible key, whereas cognitive attention is “dynamically selective:” the human’s choice of which objects to attend depends on the human’s current state (Spencer, 1870). One possible way to bridge those gaps between artificial attention and human attention is a *symbiotic learning-based attention mechanism*. In such a symbiotic learning process, human attention preferences as extracted from real-world observational data on tasks such as reading and text comprehension are incorporated in the learned attention matrices as either hard

or soft constraints, creating a centaur Transformer model. Such human attention preferences can both incorporate bottom-up insights that current attention mechanisms in the NLP domain are not able to explicitly express due to the intrinsic guided learning process, as well as provide attention insights that dynamically depend on the human’s current state.

When GPT-3 was first launched, its creators noticed a rather interesting behavior: the pre-trained model is able to adapt itself to a variety of downstream tasks without any further training—*i.e.*, without changing the model’s weights using task-specific data (Brown et al., 2020). Instead, all that was needed in order to obtain a high degree of accuracy on the downstream task is a well-designed prompt—a natural language description of the task—coupled with a very small number of demonstrations of object-label examples (“few shot learning”). Interestingly, the idea of in-context learning in LLMs strongly reassembles the idea of Concept Learning, a critical aspect of cognition. Concept learning was first defined by (Jerome S. Bruner, 1967) as “the search for and listing of attributes that can be used to distinguish exemplars from non exemplars of various categories.” Concept learning resembles LLM’s in-context learning in two key points: the ability to obtain abstract concepts from examples, and the ability to learn from very few examples. As shown in prior work, humans are able to acquire the definition of a general category from given sample positive and negative training examples of the category (Lake et al., 2015; Loewenstein & Gentner, 2005). As further shown in prior work, humans can learn such general-category definitions using only a few examples (Quinn et al., 1993; Smith et al., 2002).

The learning process in the classic in-context-learning setting is a one-way learning process in which the human “teaches” the machine by designing prompts and demonstrations, fed to the LLM which outputs the required text according to the human-designed prompt. While such one-way in-context learning schemes were shown to yield high performance, recent work has shown that in-context learning does not result in true model learning, but is rather merely a sophisticated “query normalization” method for building better encoding of user queries (Min et al., 2022; Webson & Pavlick, 2021). An interesting question that should be explored is whether symbiotic learning can be used to improve existing in-context learning approaches, facilitating true learning by enhancing existing one-way in-context learning schemes into two-way schemes. Two-way in-context learning schemes result in centaur models in which both the human and the machine learn from each other’s inputs, instead of only the machine learning from the human. A concrete example of such a two-way learning process for in-context learning is symbiotic prompt learning as illustrated below: instead of the human manually crafting prompts and examples, a human can use a machine-learning model, h_1 , that, given a general task description and goals inputted by the human, d_{human} , outputs a prompt and a set of demonstrations, d_{prompt} designed to be consumed by a specific type of LLM, h_2 :

$$\text{Human} (\rightarrow d_{human}) \rightarrow \text{Machine} (\rightarrow d_{prompt} = h_1(d_{human})) \rightarrow \text{Machine} (\rightarrow h_2(h_1(d_{human}))).$$

The interactions between the human and the machine can become even more flexible by enabling the human to give feedback (\mathcal{H}) on the prompts and demonstrations outputted by the machine before inputting them to the LLM:

$$\text{Human} (\rightarrow d_{human}) \rightarrow \text{Machine} (\rightarrow d_{prompt} = h_1(d_{human})) \rightarrow \text{Human} (\rightarrow \mathcal{H}(h_1(d_{human}))) \rightarrow \text{Machine} (\rightarrow h_2(\mathcal{H}(h_1(d_{human}))))).$$

When faced with complex problems, one reasoning strategy employed by humans' cognition is "divide-and-conquer:" decomposing the problem into manageable, intermediate steps and solving each one before giving the final answer (J. Wei et al., 2022). Such intermediate sub-problems can be seen as a series of logical deductions that sequentially lead to a conclusive answer. The benefits of such sequential decomposition reasoning in humans raised the question of whether such "intermediate thinking" patterns can improve LLMs' reasoning abilities, leading to the development of Chain-of-Thought reasoning (J. Wei et al., 2022). In its simplest form, Chain-of-Thought (CoT) denotes a series of intermediate natural language reasoning steps that lead to a final output (Zhou et al., 2022).

Until recently, CoT has been incorporated within LLMs primarily using CoT Prompting: instead of demonstrations solely composed of a (query, answer) pair, CoT Prompting additionally includes a rationale for each example, encouraging the model to verbalize the intermediate reasoning steps for solving the task (Zhou et al., 2022). Newer models, such as the latest release of OpenAI's o1 (OpenAI, 2024), take one step forward and incorporate CoTs within the model's learning and reasoning process. In this way, human input in the form of CoTs is explicitly combined within the model's architecture rather than being queried in an ad-hoc fashion by the end user as part of the prompt. The o1 model is trained to follow human-produced CoTs by using an RL scheme designed to improve the model's alignment not only on final answers but also on their corresponding CoTs—the reasoning process that led to the final answers (Valmeekam et al., 2024; Zhong et al., 2024).

While CoT reasoning can certainly be seen as one step forward towards creating centaur models, as mentioned earlier, such models are still different than centaurs. A critical difference is in that they include only *conscious* reasoning patterns, depriving LLMs from any training data representing *unconscious reasoning patterns*. Thus, models trained on CoT human input alone may not be considered as centaurs: they lack true human intuition, at least based on the definition of intuition we discussed earlier.

Indeed, recent works show that while o1 significantly outperforms GPT-4 on problem-solving tasks such as math and coding problems, it performs on par (and sometimes, subpar) with GPT-4 on ambiguous tasks, highly-abstract tasks, tasks requiring flexibility and adaptability to real-time dynamics, or ones that require knowledge gained through experience (OpenAI, 2024; Zhong et al., 2024). Furthermore, the iterative reasoning process of o1 makes it significantly slower than GPT-4, leading to longer output times (Zhong et al., 2024).

Nonetheless, enriching LLMs' built-in, cognitive-like mechanisms such as attention and concept learning with symbiotic learning techniques and high-quality human-preference data has a great potential to further enrich LLMs' cognitive abilities. This can be operationalized, for example, by using two-way in-context learning and human-preference-based attention weights as described above. Below we discuss two specific symbiotic-learning methods that have shown to be the underlying reasons why LLMs have gained cognitive abilities: human-preference-based supervised fine-tuning and human-guided rewards.

Fine-tuning LLMs with human-preference data has been shown by multiple works to yield models that successfully follow human behaviors on a diverse array of applications. All works apply a similar technique to the one illustrated in Section 5, while varying the nature of d_{data} and d_{human} as well as the structure of γ and \mathcal{C} to meet the needs of the specific downstream task and using efficient fine-tuning techniques such as linear probing or LORA adaptation. For example, (Coltheart, 2016; Schrimpf et al., 2021) showed that LLMs can predict behavioral responses in reading tasks by

fine-tuning pre-trained transformers on both behavioral and neural human reading data, fitting a regression model from the model activations' to the corresponding neural/behavioral measurements. Furthermore, (v. Zhang et al., 2024) showed that LLMs can follow human motion behavior by fine-tuning a pre-trained LLM on human motion data using LORA adaptation (Hu et al., 2021), and (Binz & Schulz, 2023a) discussed that LLMs can imitate human's decision-making process by fine-tuning pre-trained LLMs on human-choice datasets, fitting a regression model from the model's activations to the corresponding human choices. By doing so, (Binz & Schulz, 2023a) created a new class of models that they termed CENTaUR. The resulting model showed close to human behavior in two general sets of human decision-making experiments. The first set is known as decisions from descriptions. In this set, a complete, idealized, and abstract set of information about the values and frequencies of potential outcomes from each choice is provided to participants before choices are made (Weiss-Cohen et al., 2018). An example is when participants are told there is "50% chance to win 1000\$; 50% chance to win nothing" ((Kahneman & Tversky, 1979), p. 264). The second set is known as decisions from experience. In this set, participants need to form their own view of the potential outcomes from each choice via feedback provided after each selection is made (Binz & Schulz, 2023a; Weiss-Cohen et al., 2018).

Human-guided rewards is another symbiotic learning approach that has been used to enhance LLMs' cognitive abilities. Human-preference-based reward schemes were shown by multiple works (Stiennon et al., 2020; Ziegler et al., 2019) as crucial for unsupervisedly pre-trained LLMs to create outputs that better align with human's choices. Newer works have also applied human-guided reward schemes to increase LLMs' alignment with collaborative human decision-making mechanisms (Bakker et al., 2022), to create LLMs that both follow and predict humans' implicit and explicit intentions (Ouyang et al., 2022), or to teach LLMs how to align with humans' moral values, such as honesty and fairness (Ji et al., 2024).

Other works such as (Binz & Schulz, 2023b) and (Kiciman et al., 2023) show that LLMs have been able to perform well compared to human subjects in vignette-based tasks, and has also excelled in multi-arm bandit type of decision making experiments that requires balancing exploitative and exploratory actions. Some specific cognitive experiments performed by researchers to assess the cognitive ability of LLMs include:

Base Rate Fallacy (The Cab Problem). This experiment is about judging the color of a cab that has been involved in an accident. Base rate fallacy refers to the fact that participants often fail to take into account the base rate of different cab colors that operate in a city. GPT has been able to avoid this fallacy, providing accurate answers (Binz & Schulz, 2023b).

Causal and Counterfactual Reasoning (Blicket, Pills, and Do Experiments). One of the experiments used in assessing the causal reasoning ability of GPT is the "Blicket" experiment. The goal is to identify whether an object turns on a machine (*i.e.*, is a "blicket"). Two objects are introduced, and the participants are informed that the first object alone can turn on the machine. But for the second object to turn on the machine, it needs to be accommodated with the first one. GPT-3 has been able to correctly identify that the first object is a blicket, but the second one is not (Binz & Schulz, 2023b). A more serious set of investigations includes using the so-called "pills experiments" and "do experiments". In both of these, an underlying causal relationship between different entities (e.g., pills and death) is provided and then the participant is asked some counterfactual questions. GPT-3 without fine-tuning has shown good performance in some of these experiments, but has also raised concerns about its overall causal and counterfactual reasoning. For example, authors in (Binz & Schulz, 2023b) concluded that "GPT-3 has difficulties with causal reasoning in

tasks that go beyond a vignette-based characterization.” Using experiments related to over hundred causal relationships from various domains such as physics, biology, zoology, cognitive science, epidemiology, and soil science, authors in (Kırcıman et al., 2023) concluded that “algorithms based on GPT-3.5 and 4 outperform existing algorithms on a pairwise causal discovery task, counterfactual reasoning task, and actual causality. At the same time, LLMs exhibit unpredictable failure modes.”

Problem Solving and Decision-Making (Multi-Arm Bandits and Prospect Theory Experiments). GPT has also shown strong performance in problem-solving and decision-making tasks (Binz & Schulz, 2023b). As mentioned earlier, these include both decisions from descriptions and decisions from experience. The latter is the more challenging one, and often requires some fine-tuning to make sure human intuition is directly fed to the underlying model. In testing the ability of LLMs to make decisions from experience, some researchers have focused on multi-arm bandit problems where tradeoffs in exploration and exploitation play a key role. Recent work also demonstrates that LLMs might show similar to human behavior biases such as those that were documented by Kahneman and Tversky as part of their celebrated Prospect Theory (Kahneman & Tversky, 1972, 1979). For example, in one set of experiments researchers found that GPT-3 showed three of the six biases in human decision-making that Prospect Theory predicts. Specifically, GPT-3 displayed a “framing effect”, a “certainty effect,” and an “overweighting bias effect” but not a “reflection effect,” an “isolation effect,” or a “magnitude perception effect” (Binz & Schulz, 2023b).

Finally, a particularly interesting work is that of (Hagendorff et al., 2023), which aims at evaluating the cognitive abilities of different LLMs by probing LLMs with the Cognitive Reflection Test and comparing their performance and mistakes. The authors found that older versions of GPT made significantly more cognitive errors compared to GPT-4. The authors discuss potential reasons for such cognitive differences and conclude that the heavier reliance on human-guided rewards (using RLHF) for model training is the most plausible reason for GPT-4’s improved cognitive abilities.

As discussed earlier, these methods are all based on symbiotic learning. Furthermore, they allow the incorporation of human intuition so as to augment performance in various tasks. As such, LLMs that benefit from these techniques have quietly moved to become centaurs. It is not clear whether and how without becoming centaurs, LLMs could reach such cognitive ability.

7. DISCUSSION: ARE CENTAURS ALWAYS SUPERIOR?

In prior sections we have surveyed multiple techniques for creating centaurs. A key question that arises when considering the deployment of centaurs is the following: Does the fact that centaurs incorporate human intuition and symbiotic learning make them always have a better performance vis-a-vis some pure AI models?

The answer to this question is simple but important and requires a holistic view of the specific application as well as the performance metrics used. In some applications, there is a strong positive correlation between the model’s alignment with human preferences and high performance as dictated by the task-specific performance metrics. In some other applications, such a correlation is weak at best, or even negative. Thus, we claim that there exists an inherent tension between the model’s alignment with human preferences and the model’s performance, and represent this tension by evaluating the model using two separate evaluation metrics: a performance metric, ϕ_p , and a behavioral metric, ϕ_b . We see centaurs as a key conceptual mechanism over the next few years to create AI models that perform well in both ϕ_p and ϕ_b dimensions.

However, ensuring that future centaurs perform well in both dimensions needs further research and deeper understandings. To illustrate this, we consider current efforts for increasing LLMs’ cognitive abilities as discussed in Section 6. Turning LLMs into cognitive models means that LLMs will also inherit at least some flaws of human cognition, not just its strengths. Indeed, it has been shown that on some tasks, LLMs’ ability to follow human’s intuitions, behaviors, or decisions negatively correlates with the model’s performance (Acerbi & Stubbersfield, 2023; Aher et al., 2023). That is to be expected as human decision-making process is oftentimes irrational, influenced by noise factors such as emotions and cognitive biases which might result in suboptimal decisions and behaviors (Schwarz, 2000). For instance, it has been shown in (Acerbi & Stubbersfield, 2023), (Liang et al., 2021), and (Schramowski et al., 2022) that LLMs share similar cognitive biases with humans, and that even though those biases lead to behaviors that are more similar to human behaviors, the biased behaviors also lead to deterioration in LLMs’ performance. While the latter has been recently highlighted in the context of LLMs, the cognitive bias problem presents a challenge in the design of any AI model where human biases are combined in the training process either explicitly by combining human preferences in the training algorithm, or implicitly by using biased training sets (Barocas et al., 2023). Another potential threat in developing centaur-type models driven by human intuition is the development of dark-triad-like traits (Paulhus & Williams, 2002) within LLMs. At scale, such traits may pose serious safety risks by driving the model to engage in an array of negative behaviors from systematic deception and content manipulation to power-seeking behaviors, potentially with the collaboration of other centaur-like agents (Markov et al., 2023; Phuong et al., 2024; Shevlane et al., 2023).

Nonetheless, there exist multiple settings in which we can expect centaurs to score well on both ϕ_p and ϕ_b dimensions. This is clearly the case in which these dimensions are strongly and positively correlated, such as developing LLM-based chatbots (Ouyang et al., 2022) in which the performance metric is the extent to which human customers find them “helpful.” Such a performance metric is heavily human-aligned, and thus different from an “objective” performance metric such as the proportion of chatbot’s answers that contained factually correct information. Likewise, the performance metric of many summarizing LLMs (Stiennon et al., 2020; Ziegler et al., 2019) is whether the produced summary is highly-rated by human readers—a purely behavioral metric, as opposed to a factual metric such as the number of grammar mistakes found in the summary.

Another type of tasks in which we can expect centaurs to score well on both ϕ_p and ϕ_b is tasks in which the model’s goal is merely to accurately replicate human behavior. For example, multiple works have examined the extent to which LLMs can be used as a tool for simulating human behaviors in human subject studies (Aher et al., 2023; Argyle et al., 2023; Ziems et al., 2024). It has been shown that certain LLMs can accurately reproduce economic, psycholinguistic, and social psychology studies, and existing research efforts are targeted at generalizing those results by using LLMs for simulating diverse populations, both in terms of their simulated demographic characteristics and their simulated personality traits (Caron & Srivastava, 2022; Jiang et al., 2024).

Other notable examples include using centaurs in tasks for which there exist multiple optimal solutions with respect to ϕ_p . Benefiting from human intuition and symbiotic learning, a centaur model can pick an optimal solution that best aligns with human preferences. A similar situation exists in tasks applied to unrepresentative datasets or datasets with many edge cases.

We consider modern centaur-based symbiotic learning to be a key mechanism for ensuring that LLMs in particular and AI models in general score high on both the performance metric and the behavioral metric for a variety of reasons. A particular reason is that by carefully choosing

γ and \mathcal{C} , the model is able to learn when following human intuition is beneficial, while selectively filtering damaging intuitions resulting from irrational patterns such as cognitive biases and emotions. Alternatively, if the performance metric is merely a behavioral metric, as is the case when using LLMs for behavior simulation in human subject studies, γ and \mathcal{C} can be fine-tuned to solely capture human intuition. Evidence of our hypothesis can already be found in existing work, and as the research of centaurs continues, more sophisticated ways for balancing performance-based evaluation metrics and behavior-based evaluation metrics will emerge, creating “selectively cognitive” models that are able to learn when and how it is best to follow human intuition.

8. CONCLUSION: THE FUTURE OF AI IN VARIOUS APPLICATIONS IS CENTAUR

In this article, we presented the concept of centaurs as AI models that benefit from the paradigm of symbiotic learning as well as the incorporation of human intuition. We argued that centaurs have various advantages over other AI models. We also discussed that various human-level performances observed from recent LLMs including demonstration of cognitive ability stem from the fact that they have moved from pure AI models to centaur-based systems.

We started by presenting a wide array of motivations for the incorporation of human intuition in training AI models. We then proceeded to surveying traditional human-in-the-loop learning techniques and explained how they differ from symbiotic learning techniques which are used in centaurs. We then presented a formal framework for symbiotic learning and discussed five concrete symbiotic learning techniques: preference-based augmented covariate space, human-guided rewards, human-preference-based supervised fine-tuning, human-machine ensembles, and preference-constrained cost functions.

In closing, we note that while AI is still far from reaching “human-level intelligence,” the advancements in Generative AI have shown promising results for making use of centaurs in various applications. Centaurs might indeed be at the center of future AI developments. They might also be our main hope to get closer to reaching the goals that were stated by the founding fathers of AI. To realize this, it is useful to revisit what the founding fathers (John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon) stated in 1955 as part of their proposed 2-month, 10-man study of AI (to be held at Dartmouth):

“The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves” (McCarthy et al., 1955).

Making “machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves” no longer seems an unrealistic goal, given what modern centaur-based AI methods have been able to achieve to this date. Future research can make these methods even more powerful, pushing us further to reach the goals of AI research as envisioned by the founding fathers of AI.

REFERENCES

- Acerbi, A., & Stubbersfield, J. M. (2023). Large language models show human-like content biases in transmission chain experiments. *Proceedings of the National Academy of Sciences*, *120*(44), e2313790120.

- Agarwal, N., Huang, R., Moehring, A., Rajpurkar, P., Salz, T., & Yu, F. (2024). Comparative advantage of humans vs ai in the long tail. *American Economic Association*.
- Aher, G. V., Arriaga, R. I., & Kalai, A. T. (2023). Using large language models to simulate multiple humans and replicate human subject studies. *International Conference on Machine Learning*, 337–371.
- Allport, A. (1989). Visual attention.
- Amershi, S., Cakmak, M., Knox, W. B., & Kulesza, T. (2014). Power to the people: The role of humans in interactive machine learning. *AI magazine*, 35(4), 105–120.
- Ang, Y. Q., Chia, A., & Saghafian, S. (2022). Using machine learning to demystify startups’ funding, post-money valuation, and success. *Springer International Publishing*, 271–296.
- Argyle, L. P., Busby, E. C., Fulda, N., Gubler, J. R., Rytting, C., & Wingate, D. (2023). Out of one, many: Using language models to simulate human samples. *Political Analysis*, 31(3), 337–351.
- Bakker, M., Chadwick, M., Sheahan, H., Tessler, M., Campbell-Gillingham, L., Balaguer, J., McAleese, N., Glaese, A., Aslanides, J., Botvinick, M., et al. (2022). Fine-tuning language models to find agreement among humans with diverse preferences. *Advances in Neural Information Processing Systems*, 35, 38176–38189.
- Barocas, S., Hardt, M., & Narayanan, A. (2023). *Fairness and machine learning: Limitations and opportunities*. MIT Press.
- Bengio, Y., Louradour, J., Collobert, R., & Weston, J. (2009). Curriculum learning. *Proceedings of the 26th annual international conference on machine learning*, 41–48.
- Berg, S., Kutra, D., Kroeger, T., Straehle, C. N., Kausler, B. X., Haubold, C., Schiegg, M., Ales, J., Beier, T., Rudy, M., et al. (2019). Ilastik: Interactive machine learning for (bio) image analysis. *Nature methods*, 16(12), 1226–1232.
- Binz, M., & Schulz, E. (2023a). Turning large language models into cognitive models. *arXiv preprint arXiv:2306.03917*.
- Binz, M., & Schulz, E. (2023b). Using cognitive psychology to understand gpt-3. *Proceedings of the National Academy of Sciences*, 120(6), e2218523120.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33, 1877–1901.
- Burton, J. W., Stein, M.-K., & Jensen, T. B. (2020). A systematic review of algorithm aversion in augmented decision making. *Journal of behavioral decision making*, 33(2), 220–239.
- Carcillo, F., Le Borgne, Y.-A., Caelen, O., & Bontempi, G. (2018). Streaming active learning strategies for real-life credit card fraud detection: Assessment and visualization. *International Journal of Data Science and Analytics*, 5, 285–300.
- Caron, G., & Srivastava, S. (2022). Identifying and manipulating the personality traits of language models. *arXiv preprint arXiv:2212.10276*.
- Cha, Y., & Lee, S. W. (2021). Human uncertainty inference via deterministic ensemble neural networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(7), 5877–5886.
- Chang, J., Gerrish, S., Wang, C., Boyd-Graber, J., & Blei, D. (2009). Reading tea leaves: How humans interpret topic models. *Advances in neural information processing systems*, 22.
- Chen, C. H. (2015). *Handbook of pattern recognition and computer vision*. World scientific.

- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., & Amodei, D. (2017). Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Cohen, I. G., Babic, B., Gerke, S., Xia, Q., Evgeniou, T., & Wertenbroch, K. (2023). How ai can learn from the law: Putting humans in the loop only on appeal. *npj Digital Medicine*, 6(1), 160.
- Coltheart, M. (2016). *Attention and performance xii: The psychology of reading*. Routledge.
- Correia, A. H., & Lecue, F. (2019). Human-in-the-loop feature selection. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01), 2438–2445.
- Corvelo Benz, N., & Rodriguez, M. (2024). Human-aligned calibration for ai-assisted decision making. *Advances in Neural Information Processing Systems*, 36.
- Cranor, L. F. (2008). A framework for reasoning about the human in the loop.
- Deloitte. (2020). How new human-machine collaborations could make government organizations more efficient. *Harvard Business Review*.
- Denton, E. L., Chintala, S., Fergus, R., et al. (2015). Deep generative image models using a laplacian pyramid of adversarial networks. *Advances in neural information processing systems*, 28.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual review of neuroscience*, 18(1), 193–222.
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114.
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management science*, 64(3), 1155–1170.
- Fails, J. A., & Olsen Jr, D. R. (2003). Interactive machine learning. *Proceedings of the 8th international conference on Intelligent user interfaces*, 39–45.
- Gal, Y., Islam, R., & Ghahramani, Z. (2017). Deep bayesian active learning with image data. *International conference on machine learning*, 1183–1192.
- Gladwell, M. (2006). Blink: The power of thinking without thinking.
- Goldman, S. A., & Kearns, M. J. (1995). On the complexity of teaching. *Journal of Computer and System Sciences*, 50(1), 20–31.
- Grønsund, T., & Aanestad, M. (2020). Augmenting the algorithm: Emerging human-in-the-loop work configurations. *The Journal of Strategic Information Systems*, 29(2), 101614.
- Guo, S., Huang, W., Zhang, H., Zhuang, C., Dong, D., Scott, M. R., & Huang, D. (2018). Curriculumnet: Weakly supervised learning from large-scale web images. *Proceedings of the European conference on computer vision (ECCV)*, 135–150.
- Hagendorff, T., Fabi, S., & Kosinski, M. (2023). Human-like intuitive behavior and reasoning biases emerged in large language models but disappeared in chatgpt. *Nature Computational Science*, 3(10), 833–838.
- Hanczar, B., & Dougherty, E. R. (2008). Classification with reject option in gene expression data. *Bioinformatics*, 24(17), 1889–1895.
- Heimerl, F., Koch, S., Bosch, H., & Ertl, T. (2012). Visual classifier training for text document retrieval. *IEEE Transactions on Visualization and Computer Graphics*, 18(12), 2839–2848.
- Hogarth, R. M. (2001). Educating intuition. *The University of Chicago*.

- Houlsby, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., & Gelly, S. (2019). Parameter-efficient transfer learning for nlp. *International conference on machine learning*, 2790–2799.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., & Chen, W. (2021). Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Huang, L. (2019). When it's ok to trust your gut on a big decision. *Harvard Business Review*.
- Huang, S., Dong, L., Wang, W., Hao, Y., Singhal, S., Ma, S., Lv, T., Cui, L., Mohammed, O. K., Patra, B., et al. (2023). Language is not all you need: Aligning perception with language models. *Advances in Neural Information Processing Systems*, 36, 72096–72109.
- Ibarz, B., Leike, J., Pohlen, T., Irving, G., Legg, S., & Amodei, D. (2018). Reward learning from human preferences and demonstrations in atari. *Advances in neural information processing systems*, 31.
- Idan, L. (2022). Beyond purchase intentions: Mining behavioral intentions of social-network users. *International Journal of Human-Computer Interaction*, 40(5), 1111–1132. <https://doi.org/10.1080/10447318.2022.2132195>
- Idan, L. (2024). Towards unsupervised validation of anomaly detection models. *27th European Conference on Artificial Intelligence (ECAI)*. <https://doi.org/10.3233/FAIA240859>
- Jennings, N. R., Moreau, L., Nicholson, D., Ramchurn, S., Roberts, S., Rodden, T., & Rogers, A. (2014). Human-agent collectives. *Communications of the ACM*, 57(12), 80–88.
- Jerome S. Bruner, G. A. A., Jacqueline F. Goodnow. (1967). *A study of thinking*. Wiley.
- Ji, J., Liu, M., Dai, J., Pan, X., Zhang, C., Bian, C., Chen, B., Sun, R., Wang, Y., & Yang, Y. (2024). Beavertails: Towards improved safety alignment of llm via a human-preference dataset. *Advances in Neural Information Processing Systems*, 36.
- Jiang, G., Xu, M., Zhu, S.-C., Han, W., Zhang, C., & Zhu, Y. (2024). Evaluating and inducing personality in pre-trained language models. *Advances in Neural Information Processing Systems*, 36.
- Kahneman, D. (2011). Thinking, fast and slow. *Farrar, Straus and Giroux*.
- Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree. *American psychologist*, 64(6), 515.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive psychology*, 3(3), 430–454.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. 47, 263–292.
- Kasparov, G. (2010). The chess master and the computer. *The New York Review of Books*, 57(2), 16–19.
- Kasparov, G. (2017). Garry kasparov on ai, chess, and the future of creativity.
- Khatri, N., & Ng, H. A. (2000). The role of intuition in strategic decision making. *Human relations*, 53(1), 57–86.
- Kiciman, E., Ness, R., Sharma, A., & Tan, C. (2023). Causal reasoning and large language models: Opening a new frontier for causality. *arXiv preprint arXiv:2305.00050*.
- Klein, G. (2015). A naturalistic decision making perspective on studying intuitive decision making. *Journal of applied research in memory and cognition*, 4(3), 164–168.
- Klein, G. A. (2017). *Sources of power: How people make decisions*. MIT press.

- Knoche, M., & Rigoll, G. (2023). Tackling face verification edge cases: In-depth analysis and human-machine fusion approach. *2023 18th International Conference on Machine Vision and Applications (MVA)*, 1–5.
- Kocmi, T., & Bojar, O. (2017). Curriculum learning and minibatch bucketing in neural machine translation. *arXiv preprint arXiv:1707.09533*.
- Lage, I., Ross, A., Gershman, S. J., Kim, B., & Doshi-Velez, F. (2018). Human-in-the-loop interpretability prior. *Advances in neural information processing systems*, 31.
- Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350(6266), 1332–1338.
- Langlotz, C. P. (2019). Will artificial intelligence replace radiologists? *Radiology: Artificial Intelligence*, 1(3).
- Lee, M. H., Siewiorek, D. P., Smailagic, A., Bernardino, A., & Bermúdez i Badia, S. B. (2021). A human-ai collaborative approach for clinical decision making on rehabilitation assessment. *Proceedings of the 2021 CHI conference on human factors in computing systems*, 1–14.
- Levitt, S. D. (2021). Heads or tails: The impact of a coin toss on major life decisions and subsequent happiness. *The Review of Economic Studies*, 88(1), 378–405.
- Liang, P. P., Wu, C., Morency, L.-P., & Salakhutdinov, R. (2021). Towards understanding and mitigating social biases in language models. *International Conference on Machine Learning*, 6565–6576.
- Licklider, J. C. (1960). Man-computer symbiosis. *IRE transactions on human factors in electronics*, (1), 4–11.
- Loewenstein, J., & Gentner, D. (2005). Relational language and the development of relational mapping. *Cognitive psychology*, 50(4), 315–353.
- Lorenz, J., Silva, M. I., Aparício, D., Ascensão, J. T., & Bizarro, P. (2020). Machine learning methods to detect money laundering in the bitcoin blockchain in the presence of label scarcity. *Proceedings of the first ACM international conference on AI in finance*, 1–8.
- Markov, T., Zhang, C., Agarwal, S., Nekoul, F. E., Lee, T., Adler, S., Jiang, A., & Weng, L. (2023). A holistic approach to undesired content detection in the real world. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(12), 15009–15018.
- McCarthy, J., Minsky, M., Rochester, N., & Shannon, C. (1955). A proposal for the dartmouth summer research project on artificial intelligence.
- Mei, S., & Zhu, X. (2015a). The security of latent dirichlet allocation. *Artificial Intelligence and Statistics*, 681–689.
- Mei, S., & Zhu, X. (2015b). Using machine teaching to identify optimal training-set attacks on machine learners. *Proceedings of the aai conference on artificial intelligence*, 29(1).
- Min, S., Lyu, X., Holtzman, A., Artetxe, M., Lewis, M., Hajishirzi, H., & Zettlemoyer, L. (2022). Re-thinking the role of demonstrations: What makes in-context learning work? *arXiv preprint arXiv:2202.12837*.
- Mirchandani, S., Xia, F., Florence, P., Ichter, B., Driess, D., Arenas, M. G., Rao, K., Sadigh, D., & Zeng, A. (2023). Large language models as general pattern machines. *arXiv preprint arXiv:2307.04721*.
- Muller, E. (2022). How ai-human symbiotes may reinvent innovation and what the new centaurs will mean for cities. *Technology and Investment*, 13(1), 1–19.

- Munshi, V. N., Saghafian, S., Cook, C. B., Werner, K. T., & Chakkerla, H. A. (2020). Comparison of post-transplantation diabetes mellitus incidence and risk factors between kidney and liver transplantation patients. *PloS one*, *15*(1), e0226873.
- Myachykov, A., & Posner, M. I. (2005). Attention in language. In *Neurobiology of attention* (pp. 324–329). Academic Press.
- OpenAI. (2023). Gpt-4.
- OpenAI. (2024). Openai o1 system card.
- Orfanoudaki, A., Saghafian, S., Song, K., Chakkerla, H. A., & Cook, C. (2022). Algorithm, human, or the centaur: How to enhance clinical care?
- Otte, M. (1990). Intuition and logic. *For the learning of mathematics*, *10*(2), 37–43.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. (2022). Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, *35*, 27730–27744.
- PARC. (2023). Half-human, half-computer? meet the modern centaur.
- Patil, K. R., Zhu, J., Kopeć, L., & Love, B. C. (2014). Optimal teaching for limited-capacity human learners. *Advances in neural information processing systems*, *27*.
- Paulhus, D. L., & Williams, K. M. (2002). The dark triad of personality: Narcissism, machiavelianism, and psychopathy. *Journal of research in personality*, *36*(6), 556–563.
- Pearl, J., & Mackenzie, D. (2018). *The book of why: The new science of cause and effect*. Basic books.
- Phuong, M., Aitchison, M., Catt, E., Cogan, S., Kaskasoli, A., Krakovna, V., Lindner, D., Rahtz, M., Assael, Y., Hodkinson, S., et al. (2024). Evaluating frontier models for dangerous capabilities. *arXiv preprint arXiv:2403.13793*.
- Platanios, E. A., Stretcu, O., Neubig, G., Poczos, B., & Mitchell, T. M. (2019). Competence-based curriculum learning for neural machine translation. *arXiv preprint arXiv:1903.09848*.
- Quinn, P. C., Eimas, P. D., & Rosenkrantz, S. L. (1993). Evidence for representations of perceptually similar natural categories by 3-month-old and 4-month-old infants. *Perception*, *22*(4), 463–475.
- Rebuffi, S.-A., Bilen, H., & Vedaldi, A. (2017). Learning multiple visual domains with residual adapters. *Advances in neural information processing systems*, *30*.
- Register, T. (2024). Energy breakthrough needed to build agi, says openai boss altman.
- Reverberi, C., Rigon, T., Solari, A., Hassan, C., Cherubini, P., & Cherubini, A. (2022). Experimental evidence of effective human–ai collaboration in medical decision-making. *Scientific reports*, *12*(1), 14952.
- Saghafian, S. (2018). Ambiguous partially observable markov decision processes: Structural results and applications. *Journal of Economic Theory*, *178*, 1–35.
- Saghafian, S. (2023). The analytics science behind ChatGPT: Human, algorithm, or a human-algorithm centaur? *Public Impact Analytics Science (PIAS) Blog*.
- Saghafian, S. (2024). Ambiguous dynamic treatment regimes: A reinforcement learning approach. *Management Science*, *70*(9), 5667–5690.
- Saghafian, S., & Akbari Jokar, M. R. (2009). Integrative cell formation and layout design in cellular manufacturing systems. *Journal of Industrial and Systems Engineering*, *3*(2), 97–115.
- Saghafian, S., Hopp, W. J., Van Oyen, M. P., Desmond, J. S., & Kronick, S. L. (2012). Patient streaming as a mechanism for improving responsiveness in emergency departments. *Operations Research*, *60*(5), 1080–1097.

- Saghafian, S., Kilinc, D., & Traub, S. (2024). *Dynamic assignment of patients to primary and secondary inpatient units: Is patience a virtue?* Cambridge Handbook on Productivity, Efficiency; Effectiveness in Healthcare, Cambridge University Press.
- Saghafian, S., & Tomlin, B. (2016). The newsvendor under demand ambiguity: Combining data with moment and tail information. *Operations Research*, 64(1), 167–185.
- Schramowski, P., Turan, C., Andersen, N., Rothkopf, C. A., & Kersting, K. (2022). Large pre-trained language models contain human-like biases of what is right and wrong to do. *Nature Machine Intelligence*, 4(3), 258–268.
- Schrimpf, M., Blank, I. A., Tuckute, G., Kauf, C., Hosseini, E. A., Kanwisher, N., Tenenbaum, J. B., & Fedorenko, E. (2021). The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45), e2105646118.
- Schwarz, N. (2000). Emotion, cognition, and decision making. *Cognition & emotion*, 14(4), 433–440.
- Settles, B. (2011). From theories to queries: Active learning in practice. *Active learning and experimental design workshop in conjunction with AISTATS 2010*, 1–18.
- Shevlane, T., Farquhar, S., Garfinkel, B., Phuong, M., Whittlestone, J., Leung, J., Kokotajlo, D., Marchal, N., Anderljung, M., Kolt, N., et al. (2023). Model evaluation for extreme risks. *arXiv preprint arXiv:2305.15324*.
- Simard, P. Y., Amershi, S., Chickering, D. M., Pelton, A. E., Ghorashi, S., Meek, C., Ramos, G., Suh, J., Verwey, J., Wang, M., et al. (2017). Machine teaching: A new paradigm for building machine learning systems. *arXiv preprint arXiv:1707.06742*.
- Simon, H. A. (1992). What is an “explanation” of behavior? *Psychological science*, 3(3), 150–161.
- Smith, L. B., Jones, S. S., Landau, B., Gershkoff-Stowe, L., & Samuelson, L. (2002). Object name learning provides on-the-job training for attention. *Psychological science*, 13(1), 13–19.
- Spencer, H. (1870). *The principles of psychology* (Vol. 1). Williams; Norgate London.
- Stanovich, K. E., & West, R. F. (1983). On priming by a sentence context. *Journal of Experimental Psychology: General*, 112(1), 1.
- Stiennon, N., Ouyang, L., Wu, J., Ziegler, D., Lowe, R., Voss, C., Radford, A., Amodei, D., & Christiano, P. F. (2020). Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33, 3008–3021.
- Sun, R. (2008). *The cambridge handbook of computational psychology*. Cambridge University Press.
- Times, N. Y. (2016a). A case for cooperation between machines and humans.
- Times, N. Y. (2016b). Pentagon turns to silicon valley for edge in artificial intelligence.
- Torrey, L., & Shavlik, J. (2010). Transfer learning. In *Handbook of research on machine learning applications and trends: Algorithms, methods, and techniques* (pp. 242–264). IGI global.
- Traub, S. J., Bartley, A. C., Smith, V. D., Didehban, R., Lipinski, C. A., & Saghafian, S. (2016). Physician in triage versus rotational patient assignment. *The Journal of emergency medicine*, 50(5), 784–790.
- Tsvetkov, Y., Faruqui, M., Ling, W., MacWhinney, B., & Dyer, C. (2016). Learning the curriculum with bayesian optimization for task-specific word representation learning. *arXiv preprint arXiv:1605.03852*.
- Valmeekam, K., Stechly, K., & Kambhampati, S. (2024). Llms still can’t plan; can lrms? a preliminary evaluation of openai’s o1 on planbench. *arXiv preprint arXiv:2409.13373*.

- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, H., Gong, S., Zhu, X., & Xiang, T. (2016). Human-in-the-loop person re-identification. *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, 405–422.
- Ware, M., Frank, E., Holmes, G., Hall, M., & Witten, I. H. (2001). Interactive machine learning: Letting users build classifiers. *International Journal of Human-Computer Studies*, 55(3), 281–292.
- Warmuth, M. K., Liao, J., Rätsch, G., Mathieson, M., Putta, S., & Lemmen, C. (2003). Active learning with support vector machines in the drug discovery process. *Journal of chemical information and computer sciences*, 43(2), 667–673.
- Webson, A., & Pavlick, E. (2021). Do prompt-based models really understand the meaning of their prompts? *arXiv preprint arXiv:2109.01247*.
- Wei, J., & Zhou, D. (2022). Language models perform reasoning via chain of thought.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q. V., Zhou, D., et al. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35, 24824–24837.
- Wei, Y., Liang, X., Chen, Y., Shen, X., Cheng, M.-M., Feng, J., Zhao, Y., & Yan, S. (2016). Stc: A simple to complex framework for weakly-supervised semantic segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(11), 2314–2320.
- Weiss-Cohen, L., Konstantinidis, E., Speekenbrink, M., & Harvey, N. (2018). Task complexity moderates the influence of descriptions in decisions from experience. *Cognition*, 170, 209–227.
- Zhalechian, M., Saghafian, S., & Robles, O. (2024). Harmonizing safety and speed: A human-algorithm approach to enhance the fda’s medical device clearance policy. *Available at SSRN 4863134*.
- Zhang, Chen, Z., Chen, S., Shen, Y., Sun, Z., & Gan, C. (2024). Improving reinforcement learning from human feedback with efficient reward model ensemble. *arXiv preprint arXiv:2401.16635*.
- Zhang, v., Huang, D., Liu, B., Tang, S., Lu, Y., Chen, L., Bai, L., Chu, Q., Yu, N., & Ouyang, W. (2024). Motiongpt: Finetuned llms are general-purpose motion generators. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(7), 7368–7376.
- Zhong, T., Liu, Z., Pan, Y., Zhang, Y., Zhou, Y., Liang, S., Wu, Z., Lyu, Y., Shu, P., Yu, X., et al. (2024). Evaluation of openai o1: Opportunities and challenges of agi. *arXiv preprint arXiv:2409.18486*.
- Zhou, H., Nova, A., Larochelle, H., Courville, A., Neyshabur, B., & Sedghi, H. (2022). Teaching algorithmic reasoning via in-context learning. *arXiv preprint arXiv:2211.09066*.
- Zhu, X. (2015). Machine teaching: An inverse problem to machine learning and an approach toward optimal education. *Proceedings of the AAAI conference on artificial intelligence*, 29(1).
- Ziegler, D. M., Stiennon, N., Wu, J., Brown, T. B., Radford, A., Amodei, D., Christiano, P., & Irving, G. (2019). Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.
- Ziems, C., Held, W., Shaikh, O., Chen, J., Zhang, Z., & Yang, D. (2024). Can large language models transform computational social science? *Computational Linguistics*, 1–55.