

Trustworthy and Practical AI for Healthcare: A Guided Deferral System with Large Language Models

Joshua Strong, Qianhui Men*, J. Alison Noble

Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, OX3 7DQ, United Kingdom
 {joshua.strong, alison.noble}@eng.ox.ac.uk, qianhui.men@bristol.ac.uk

Abstract

Large language models (LLMs) offer a valuable technology for various applications in healthcare. However, their tendency to hallucinate and the existing reliance on proprietary systems pose challenges in environments concerning critical decision-making and strict data privacy regulations, such as healthcare, where the trust in such systems is paramount. Through combining the strengths and discounting the weaknesses of humans and AI, the field of Human-AI Collaboration (HAIC) presents one front for tackling these challenges and hence improving trust. This paper presents a novel HAIC guided deferral system that can simultaneously parse medical reports for disorder classification, and defer uncertain predictions with intelligent guidance to humans. We develop methodology which builds efficient, effective and open-source LLMs for this purpose, for the real-world deployment in healthcare. We conduct a pilot study which showcases the effectiveness of our proposed system in practice. Additionally, we highlight drawbacks of standard calibration metrics in imbalanced data scenarios commonly found in healthcare, and suggest a simple yet effective solution: the Imbalanced Expected Calibration Error.

Code —

https://github.com/josh-strong/Guided_Defer_LLM

1 Introduction

The practical implementation of artificial intelligence (AI) in decision-sensitive environments such as healthcare is often a delicate trade-off between the benefits of autonomy, risk, and the costly consequences of errors. Human-AI Collaboration (HAIC), whereby humans and AI work in tandem to produce outcomes superior to independent efforts, seeks to find this balance (Wang et al. 2020). One method of achieving HAIC is through utilising *deferral systems* (Madras, Pitassi, and Zemel 2018; Liu, Gallego, and Barberi 2022), which aim to identify cases that can be rapidly and accurately processed by the AI and then defer the remaining challenging cases to humans. This collaboration reduces decision-makers’ time and fatigue, enabling them to focus on a small set of uncertain cases identified by the deferral

system, with confidence that the AI will efficiently and accurately process the majority of cases. By ensuring that AI handles routine tasks while humans manage the complex ones, the system not only optimises the balance between benefits and risks, but crucially can also enhance human trust in the AI’s role.

Clinicians often make decisions using their expertise in addition to the *intelligent guidance* from colleagues. The colleagues will then respond with their interpretation and the reasoning behind it. However, existing deferral systems lack this guidance (Banerjee, Teso, and Passerini 2023), isolating humans in this critical decision-making. In the spirit of HAIC, we posit that effective deferral systems should additionally offer intelligent guidance to human decision-makers in deferred cases. In this paper, we novelly explore utilising large language models (LLMs) for this purpose, specifically in the task of clinical text parsing for disorder classification.

The main challenges in deploying such systems in practice include the computational cost of LLMs and compliance with strict data privacy regulations (Freyer et al. 2024). Proprietary LLMs offer advantages such as state-of-the-art performance and easy implementation without requiring high-performance hardware, but are impractical for data-sensitive applications such as healthcare due to lack of internal state access and privacy concerns (Ullah et al. 2024). Open-source LLMs can perform well with relatively large amounts of parameters, but are slow and require high-performance hardware, not readily available to all practitioners. Small-scale LLMs are faster and can be run on widely available hardware, but are not effective. A closely-related study (Zöller et al. 2024) that successfully applied LLM for medical record diagnosis has ensembled multiple LLMs and human physicians. Rather than deferring cases to humans, it focuses on a hybrid collaborative approach which may cause cognitive overload. In this paper, we propose methodology for developing *efficient, effective* and *open-source* LLMs capable of guided deferral, suitable for healthcare applications and deployable on modest computational resources.

The performance of deferral systems can be attributed, in part, to the model’s classification performance and its *calibration*—the ability for the model to accurately reflect its confidence in its predictions. Appropriate calibration improves the model’s ability to choose the optimal hierarchy of cases to be deferred. Recent literature has highlighted the

*Work undertaken whilst employed at the University of Oxford. Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

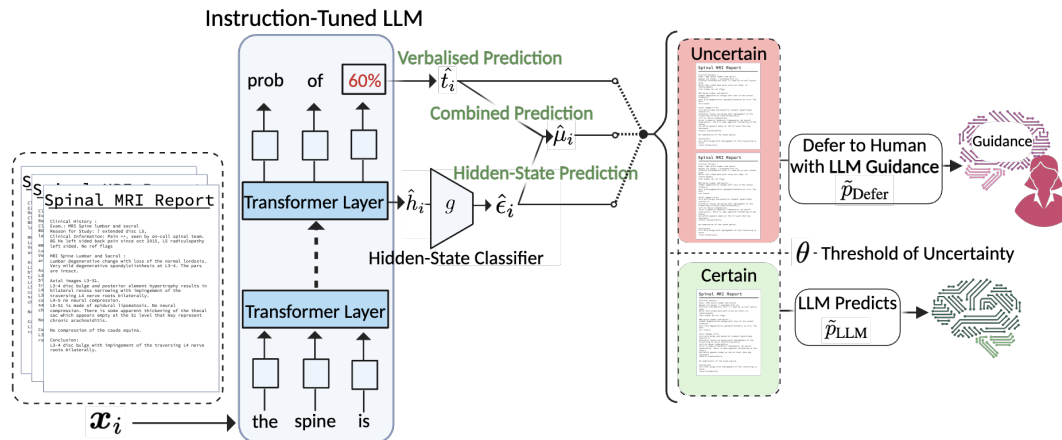


Figure 1: Our *guided deferral* system. Reports are parsed by an instruction-tuned LLM for clinical disorders. From the text output, we extract a *verbalised prediction* \hat{t} . We calculate a *hidden-state* \hat{e} prediction from the final hidden-layer of the LLM, and its combination with \hat{t} through their mean $\hat{\mu}$. Uncertain predictions, determined by either \hat{t} , \hat{e} , or $\hat{\mu}$, are deferred to humans with guidance. Certain predictions are autonomously handled by the LLM. Created in BioRender. Strong, J. (2025).

limitations (Nixon et al. 2020) of the *Expected Calibration Error* (ECE) (Guo et al. 2017), a commonly used metric for calibration. However, the proposed solutions are often inadequate in settings with class imbalance, which is frequently observed in healthcare data. Through analysis on real-world healthcare examples, we propose a new metric: the *Imbalance Expected Calibration Error* (ECE_{Imb}), which adjusts the calibration error by appropriately re-weighting the original calibration error weights in calculating ECE to account for data imbalance. We find that this metric is more suitable for evaluating the calibration of a model in these situations.

In summary, the contributions of this paper are:

- We propose a novel deferral system, *guided deferral*, for large language models (LLMs) in computer-aided clinical diagnosis. This system not only defers cases to human decision-makers, but also provides *intelligent guidance*. We detail its practical application and demonstrate its efficacy through a pilot study.
- We highlight the challenges of calibrating a deferral system with imbalanced data, common in clinical scenarios between diagnosis and negative cases. A new metric, ECE_{Imb} , is proposed as a simple yet effective measurement of calibration in such situations.
- We address the social impact challenges associated with deploying such systems in practice, by developing methodology for creating efficient, effective, and open-source LLMs capable of guided deferral. It is also deployable on modest computational resources and accessible to the community.
- We position our research within the social impact domain of healthcare, addressing challenges of low-count, imbalanced data and data privacy. We propose a human-AI collaboration system crucial for making key decisions and building trust in sensitive fields.

2 Related Work

Deferral Systems for Healthcare. AI systems can defer appropriately to clinicians, which improves accuracy and efficiency in the human-AI collaborative environment for medical support (Dvijotham et al. (2023)). With a focus on system fairness and safety, foundations for advanced deferral algorithms such as Learning-to-Defer (Madras, Pitassi, and Zemel 2018; Liu, Gallego, and Barbieri 2022) have also proposed solutions with uncertainty to assess the confidence level in each prediction.

Human-AI Collaboration with LLMs. Few studies explore the use of LLMs in HAIC. Wiegrefe et al. (2022) first examined LLMs for explaining classification decisions using a human-in-the-loop approach to train a filter assessing explanation quality. Rastogi et al. (2023) used HAIC to audit error-prone LLMs with other LLMs. The closest work is Zöller et al. (2024) who research a hybrid collective intelligence system that combines human expertise with the capabilities of LLMs to improve decision-making in medical diagnostics. Our system differs by using LLMs for guided deferrals. Banerjee, Teso, and Passerini (2023) used LLMs to provide textual guidance for clinicians in decision-making on clinical imaging tasks. The system does not defer cases, which still burdens decision-makers with the time and fatigue issues. Our work combines LLMs in deferral systems with valuable guidance for decision-makers on deferred cases.

Selective Prediction of LLMs. Existing research on LLMs in selective prediction includes methods to measure uncertainty in model responses after generation (Varshney and Baral 2023). Chen et al. (2023) proposed improving selective prediction by incorporating self-assessment. Ren et al. (2022) explored detecting out-of-distribution instances in summarisation and translation tasks. Our paper applies selective prediction to deferral systems in clinical parsing.

We show that combining the model’s internal state with its generated prediction enhances selective prediction without post-generation methods. We provide a comprehensive evaluation of selective prediction performance in clinical classification using real-world data, applied in deferral systems with in-distribution data.

Instruction-Tuning of LLMs. Whilst there exists literature demonstrating improved zero-shot performance of instruction-tuned (IT) LLMs on unseen tasks (Wei et al. 2022), our work studies this improvement specifically on an in-domain task through the use of *guardrails*. Additionally, there exists works on the use of IT in the medical domain (Liu et al. 2023), but these focus on report summarisation. No works have researched the applications of IT LLMs for the use of deferral systems. Zhang et al. (2023) provide a recent survey for additional insight in this area.

On the Calibration of LLMs. There is a wealth of literature on the calibration of LLMs (Geng et al. 2023). Lin, Hilton, and Evans (2022) illustrated the capacity of LLMs to convey the uncertainty of their model answers through verbalisation, specifically in their textual output. Additionally, Xiong et al. (2023) assessed the calibration of LLMs in relation to their verbalised probabilities—when prompted to output predictions through the generated text. In this paper, we discuss the pitfalls of LLM calibration under imbalanced classes of samples in healthcare, and propose an unbiased metric to evaluate model calibration.

The Role of Cognitive Science in HAIC. In the collaboration between humans and AI to achieve a specific goal, the distinct strengths of each enable *complementarity*, a performance surpassing that of either alone and measuring the HAIC system’s success (Bansal et al. 2021). Humans adapt to AI by developing *mental models*—cognitive frameworks aiding understanding and interaction. Andrews et al. (2022) comprehensively reviews mental models, including their definition, measurement, and related concepts. Kulesza et al. (2012) empirically showed that users with better-developed mental models achieved more satisfactory AI results. Research in cognitive science is crucial for developing effective human-AI collaborative systems.

3 Methods

3.1 Sources of Predictions

Utilising LLMs in clinical text parsing for disorder classification presents a unique challenge in determining the classification approach. We focus our study on binary classification utilising the top-performing methods of two distinct sources of classifications; one from the internal-states of the LLM and another from the generated textual output. Additionally, we experiment with a third through combining these sources. Next, we formally define these three sources. See Figure 1 for a visual supplement.

Verbalised Prediction Source. The *verbalised probability* is the probability of the positive class extracted from the generated text of the LLM. We denote this probability for

input x_i as \hat{t}_i . For example, the verbalised probability for the example guidance generated in Figure 2 is $\hat{t}_i = 0.6$.

Hidden-State Prediction Source. The *hidden-state* probability is defined based on the hidden representations of LLM to implicitly detect classifications. Inspired by Ren et al. (2022), which utilises the final-layer hidden embedding of LLMs for out-of-distribution detection, the output embedding $\hat{h}_i \in \mathcal{H}$ for input x_i is computed as the average of the decoder’s final-layer hidden-state vectors $\mathbf{g}_{ik} \in \mathbb{R}^d$ over all K output tokens with a hidden dimension of $d = 5120$ for the small-scale LLM:

$$\hat{h}_i = \frac{1}{K} \sum_{k=1}^K \mathbf{g}_{ik}^n.$$

Then, an MLP $g : \mathcal{H} \mapsto \mathbb{R}$ is trained as a hidden-state classifier to learn the probability of the disorder from the LLM hidden representation, the *hidden-state prediction probabilities*, denoted \hat{e}_i , i.e. $\hat{e}_i = g(\hat{h}_i)$. We experimented with different models for the hidden-state classifier in ambitions of fully utilising the information embedded in \hat{h} (see Technical Appendix for details).

Combined Prediction. We combine \hat{t}_i and \hat{e}_i through a weighted mean to form the *combined prediction*, of which we denote $\hat{\mu}_i$:

$$\hat{\mu}_i = \alpha \hat{e}_i + (1 - \alpha) \hat{t}_i,$$

where $\alpha \in [0, 1]$ can be learned on a validation data split. For simplicity, we fix $\alpha = 0.5$ in our experiments.

3.2 Instruction-Tuning Methodology

In generating well-calibrated verbalised probabilities, we use the “*verb. IS top-k*” prompting strategy (Tian et al. 2023), prompting the LLM to provide the top k guesses and their probabilities in a single response. Adapting this to $k = 1$ for a binary setting, we prompt the LLM to return the probability of the positive class. Additionally, we prompt for the top reason the disorder *might* and *might not* be present, using *dialectic reasoning* (Hegel 2018) to provide intelligent guidance for decision-makers. This technique has proven effective in decision support (Jarupathirun and Zahedi 2007). An example output is in Figure 2.

Instruction-Tuning Data Generation. We use 4-bit quantized 13B and 70B versions of the open-source LLM *Tulu V2 DPO* (Iverson et al. 2023), a Llama 2 (Touvron et al. 2023) derivative, as our *small* and *large*-scale LLMs respectively. As classification, calibration and deferral performance was seen to improve with model size (see Table ??), we instruction-tune the small-scale models on *guard-railed* answers generated by the large-scale model. We experiment instruction-tuning using multiple generated answers for each example in the instruction-data through enabling the stochastic decoding of the LLM. We *guard-rail* the instruction-tuning data generation process such that answers containing hallucinations and “bad” behaviours were rejected. Hallucinations are identified when the generated answers did not match the underlying annotation, and “bad”

Example Generated Guidance

TOP REASON FOR: The MRI report indicates that there is narrowing of the right exit foramen at the L5-S1 level, which is causing compression of the exiting right L5 nerve root. This suggests that there is foraminal stenosis (FS) at this level.

TOP REASON AGAINST: There is no mention of foraminal stenosis specifically at the L5-S1 level.

CONCLUSION: Based on the information provided, there is a possibility of foraminal stenosis at the L5-S1 level due to the narrowing of the right exit foramen and compression of the exiting right L5 nerve root. However, the report does not explicitly mention foraminal stenosis at this level.

PROBABILITY OF FS PRESENT AT L5-S1:
60%

Figure 2: Example guidance based on a spinal MRI report. The instruction-tuned LLM is able to intelligently infer a diagnosis with sound logic without explicit textual diagnosis.

behaviours when the LLM did not output English language, concise answers, or logical statements. The full implementation of guard-rails and the final base-prompt used to generate these examples can be seen in the Technical Appendix.

Instruction-Tuning. We instruction-tune the small-scale LLM using QLoRA (Dettmers et al. 2023) solely on completions. QLoRA enables fine-tuning of LLMs with much less memory requirements whilst maintaining full 16-bit fine-tuning task performance, aligning with our ambitions of a computationally-efficient system. See Technical Appendix on experimental details for full reproducibility.

3.3 Calibration Methods

In improving the deferral performance of our LLM, we must also focus on the calibration of the model in tandem with the classification performance. Note that a highly predictive model can still be poorly calibrated (Wang 2023). We motivate and detail our proposed metric *Imbalanced ECE* in this subsection.

Model Calibration. Intuitively, calibration measures a model’s ability to accurately quantify its confidence in its predictions through its estimated class probabilities. Let x be a realisation of data $X \in \mathcal{X}$ with corresponding label $y \in Y = \{0, 1\}$. A model $h(X) : X \rightarrow [0, 1]$ producing confidence probabilities \hat{p} of belonging to the positive class is said to be perfectly calibrated (Guo et al. 2017) iff

$$P(Y = 1 | h(X) = \hat{p}) = \hat{p}, \quad \forall \hat{p} \in [0, 1].$$

Metrics of Calibration. A common metric for quantifying the overall calibration of a model is the *Expected Calibration Error* (ECE) (Naeini, Cooper, and Hauskrecht 2015). Model confidence probabilities \hat{p} are binned into

M equal-sized interval bins within the domain $[0, 1]$. Let B_m be the set of \hat{p} which fall within the bin bounded by $I_m = (\frac{m-1}{M}, \frac{m}{M}]$, $m \in \{1, \dots, M\}$. Formally, with a dataset $\{x_i, y_i\}_{i=1}^n$,

$$\text{ECE} = \sum_{m=1}^M \frac{|B_m|}{n} |\bar{h}_m - \bar{y}_m|,$$

where \bar{h}_m is the average confidence probability of bin m and \bar{y}_m the average label in bin m . The ECE has been shown to have an inherent statistical bias, particularly for extremely well-calibrated models (Roelofs et al. 2022), in addition to sensitivity to hyperparameters, such as the number of bins, type of norm used, and binning technique (Nixon et al. 2020). We fix these hyperparameters across experiments in this paper, with the number of bins set to $M = 10$ and utilising the ℓ_1 norm in all calibration metrics.

The *Adaptive Calibration Error* (ACE) was proposed by Nixon et al. (2020) as a solution to the problem presented by the ECE, in that the methodology of selecting equally-sized interval bins could result in bins with low-count predictions, and hence, a poor estimate in overall calibration error. Instead, ACE chooses bins B_{m^*} in a way such that each of the M bins contains an equal number of $\frac{1}{M}$ confidence probabilities and bin B_{m^*} contains the m^{*th} set of sorted predictions:

$$\text{ACE} = \sum_{m^*=1}^M \frac{1}{M} |\bar{h}_{m^*} - \bar{y}_{m^*}|.$$

Imbalanced ECE. For imbalanced datasets with a majority negative class, a large proportion of the confidence probabilities of the negative class will inherently lie within the first few bins. In its computation of a weighted average of errors across bins, ECE assigns minimal weights to subsequent bins with higher probability predictions and fewer data points. As a result, the calibration of the majority class overwhelmingly influences the metric, potentially resulting in a misleading interpretation of total calibration through the ECE. The ACE was designed to combat this, but it also suffers in imbalanced situations, as generating equally spaced bins leads to the majority of bins inherently covering only the negative class examples. This problem is further exacerbated by the sharpness of model predictions, a prominent characteristic of neural networks (Gneiting, Balabdaoui, and Raftery 2007), pushing confidence predictions to their extremities of 0 or 1.

In addressing this issue, Guilbert et al. (2024) introduced $\text{ECE}@k\%$, which assesses the ECE of the top $k\%$ most confident predictions by partitioning the probability prediction domain $[\min(\hat{p}), 1]$ into $\log_{10}(k\%)$ bins. While valuable in scenarios that depend on the high calibration of high-confidence predictions, this metric falls short in evaluating the calibration of all predictions—a critical aspect, especially in deferral system use cases.

Hence, we introduce a simple yet effective metric *Imbalanced Expected Calibration Error*, ECE_{Imb} , which weights the calibration error by incorporating a blend of both the bin

sample proportion and a uniform distribution. Formally,

$$\text{ECE}_{\text{Imb}} = \sum_{m=1}^M \Gamma_m |\bar{h}_m - \bar{y}_m|,$$

where $\Gamma_m = (1 - \gamma) \frac{|B_m|}{n} + \gamma \frac{1}{M}$, $\gamma \in [0, 1]$.

When $\gamma = 0$, ECE_{Imb} reduces to the classic ECE. As $\gamma \rightarrow 1$, Γ_m morphs towards a uniform distribution. γ can be heuristically chosen on a validation dataset. In selecting a value for γ , it is recommended to choose a conservative value which does not draw the weights too much from their original bin-weighted distribution. We further address the issues with ECE, ACE, and the ECE_{Imb} solution, as well as the logic behind our choice of $\gamma = 0.3$ for our experiments, in the Technical Appendix.

3.4 Deferral Mechanism

Our deferral mechanism is based on the confidence of predictions, determined by their distance from the decision boundary of 0.5. For a given set of n predictions $\hat{p} \in \{\hat{t}, \hat{e}, \hat{\mu}\}$, we transform these into sorted relative confidence probabilities \tilde{p} :

$$\tilde{p} = \text{sort}(\{2|\hat{p}_i - 0.5|\}_{i=1}^n) = \{\tilde{p}_1, \tilde{p}_2, \dots, \tilde{p}_n\}$$

This transformation allows for an equal comparison between positive and negative classes. The resulting set \tilde{p} measures uncertainty, where \tilde{p}_i has higher deferral priority over \tilde{p}_j if $\tilde{p}_i < \tilde{p}_j$ for all $i, j \in \{1, \dots, n\}$ with $i \neq j$.

Deferral performance is evaluated by iterating through this hierarchy of \tilde{p} , deferring each prediction without replacement, and measuring the classification performance of the LLM on the remaining cases with the respective \hat{p} predictions. Note that \hat{p} and \tilde{p} do not have to be derived from the same set of predictions $\hat{t}, \hat{e},$ or $\hat{\mu}$; classification and deferral can use separate sources of probabilities. This procedure describes the AUARC (Area Under the Accuracy-Rejection Curve) metric (Nadeem, Zucker, and Hanczar 2009) when measuring accuracy. Good deferral behaviour displays a strictly monotonic increasing accuracy with an increasing deferral rate. In practice, the proportion of cases deferred is subject to a threshold of uncertainty $\theta \in [0, 1]$, dependent on operational constraints and/or risk appetite. Deferred cases \tilde{p}_{Defer} are those where $\tilde{p}_i < \theta$, and cases automatically parsed by the LLM \tilde{p}_{LLM} are those where $\tilde{p}_i \geq \theta$, formally:

$$\tilde{p}_{\text{Defer}} = \{\tilde{p}_i \in \tilde{p} \mid \tilde{p}_i < \theta\} \text{ and } \tilde{p}_{\text{LLM}} = \{\tilde{p}_i \in \tilde{p} \mid \tilde{p}_i \geq \theta\},$$

such that $\tilde{p}_{\text{Defer}} \cap \tilde{p}_{\text{LLM}} = \emptyset$ and $\tilde{p}_{\text{Defer}} \cup \tilde{p}_{\text{LLM}} = \tilde{p}$.

3.5 Pilot Study

To evaluate the effectiveness of our LLM guidance in human decision-making, we conduct a pilot study with 20 participants in the scenario of deferring 30 ($\approx 5\%$) of the most uncertain test predictions, such that $\theta = \tilde{p}_{31}$. The participants are researchers with biomedical backgrounds but not specialised in spinal disorders. The 30 medical reports are from OSCLMRIC dataset (introduced in Sect. 4) with ethical approval obtained.

Study Design. All participants received background information, including clinical details, examples of prediction outcomes with associated MRI reports and the LLM’s performance on a validation set to help participants develop a sufficient mental model of the LLM, of which has been shown to be important in HAIC systems (Kulesza et al. 2012; Bansal et al. 2019). The ordering of the 30 questions were randomised for each participant to reduce order bias. Participants moved to the next question if their prediction matched the LLM’s. If it differed, they received guidance and could either change their prediction or keep it based on the guidance and their understanding of the LLM. This process allowed us to assess human performance with and without guidance. Effective guidance should help participants recognise both the accuracy and inaccuracy of their judgements. Full details are in the Technical Appendix.

4 Experiments

Data. We utilise two datasets in our experiments: *OSCLMRIC* and *MIMIC-500*. OSCLMRIC (Oxford Secondary Care Lumbar MRI Cohort), contains 100 professionally radiological reports. The dataset is highly imbalanced, with approximately 95% of labels negative. Each report is parsed to detect the binary presence of foraminal stenosis (FS), spinal canal stenosis (SCS), and lateral recess stenosis (LRS)—at six lumbar spine levels, resulting in 1,800 examples. We randomly split the data into 30% for generating an instruction-tuning dataset, 20% for training the hidden-state classifier, 20% for validation, and 30% for testing. Due to the complexity of these reports, we use this dataset as our primary dataset in evaluating the computational efficacy of our method in addition to the pilot study in assessing the effectiveness of our proposed LLM’s guidance. Due to privacy reasons, we cannot use this dataset to compare the performance of our methodology against proprietary and state-of-the-art (SOTA) LLMs. Instead, we utilise Tulu V2 70B as the SOTA baseline model for our experiments against the OSCLMRIC dataset, of which was the highest-performing open-source LLM against several benchmarks at the time these experiments were conducted (Beeching et al. 2023)

For performance comparison against proprietary and SOTA LLMs, we utilise MIMIC-500 (Gu et al. 2024), a publicly available expert-annotated set of 500 chest X-ray reports (split into impressions and findings sections) of which there exists the SOTA LLM classification benchmarks of GPT-3.5 (Brown et al. 2020) and GPT-4 (OpenAI 2024). We focus on the *impressions* section of the reports, as SOTA LLMs performed best against this section. Each report is parsed for 10 classifications, each with high imbalance. To ensure full reproducibility, we disable stochastic processing in the model’s decoding stage during experiments with this dataset. MIMIC-500 is a subset of MIMIC-CXR-JPG (Johnson et al. 2024), which includes noisy labels from the predictions of CheXpert (Irvin et al. 2019), an open-source rule-based tool. In instruction-tuning our LLM for MIMIC-500, we greatly restrict our access to data in simulating low-count environments commonly found in healthcare, and utilise just .6% (1,500) of MIMIC-CXR-JPG (independent of MIMIC-500), with 500 each for instruction-tuning data generation,

SETUP	CLASSIFICATION PERF.			CALIBRATION PERF.			DEFERRAL PERF.
	RECALL \uparrow	PRECISION \uparrow	F1-SCORE \uparrow	ECE \downarrow	ACE \downarrow	ECE _{Imb} \downarrow	AUARC \uparrow
(1) $\hat{t}_{\text{BASE-13B}}$	0.85 \pm 0.02	0.61 \pm 0.01	0.61 \pm 0.01	0.26 \pm 0.01	0.32 \pm 0.01	0.30 \pm 0.01	0.897460 \pm 0.0063
(2) $\hat{e}_{\text{BASE-13B}}$	0.73 \pm 0.02	0.94 \pm 0.03	0.80 \pm 0.02	0.03 \pm 0.01	0.04 \pm 0.01	0.09 \pm 0.01	0.994715 \pm 0.0006
(3) $\hat{\mu}_{\text{BASE-13B}}$	0.78 \pm 0.03	0.87 \pm 0.04	0.82 \pm 0.03	0.13 \pm 0.01	0.18 \pm 0.01	0.13 \pm 0.01	0.993318 \pm 0.0010
(4) $\hat{t}_{\text{BASE-70B}}$	0.90 \pm 0.02	0.78 \pm 0.01	0.83 \pm 0.01	0.27 \pm 0.01	0.33 \pm 0.01	0.25 \pm 0.01	0.973529 \pm 0.0052
(5) $\hat{e}_{\text{BASE-70B}}$	0.87 \pm 0.01	0.97 \pm 0.01	0.91 \pm 0.01	0.02 \pm 0.01	0.05 \pm 0.01	0.08 \pm 0.01	0.997191 \pm 0.0002
(6) $\hat{\mu}_{\text{BASE-70B}}$	0.90 \pm 0.01	0.95 \pm 0.02	0.92 \pm 0.01	0.14 \pm 0.01	0.19 \pm 0.01	0.14 \pm 0.01	0.997260 \pm 0.0003
(7) $\hat{t}_{\text{IT-13B}}$	0.85 \pm 0.02	0.92 \pm 0.02	0.88 \pm 0.02	0.27 \pm 0.01	0.32 \pm 0.01	0.23 \pm 0.01	0.987501 \pm 0.0024
(8) $\hat{e}_{\text{IT-13B}}$	0.89 \pm 0.01	0.98 \pm 0.01	0.93 \pm 0.01	0.01 \pm 0.01	0.05 \pm 0.01	0.08 \pm 0.01	0.997409 \pm 0.0002
(9) $\hat{\mu}_{\text{IT-13B}}$	0.91 \pm 0.01	0.98 \pm 0.01	0.94 \pm 0.01	0.14 \pm 0.01	0.19 \pm 0.01	0.16 \pm 0.01	0.996978 \pm 0.0013

Table 1: Mean and standard deviation classification, calibration, and deferral performance against test split (N=540) on the OSCLMRIC dataset test split.




Figure 3: Visual comparison of the effectiveness of the calibration metrics (a) ECE, (b) ACE and (c) ECE_{Imb} of $\hat{e}_{\text{IT-13B}}$ predictions against the OSCLMRIC dataset test split.

SETUP	LLM EFFICIENCY		
	REL. S./GEN. \downarrow	MEM. \downarrow	E.R. \downarrow
BASE-13B	0.08 \pm 0.02	6.02	0.29
BASE-70B	1.00	32.08	0.04
IT-13B	0.08 \pm 0.02	6.02	0.0002

Table 2: LLM efficiency metrics for selected setups from the OSCLMRIC dataset test split.

validation and training the hidden-state classifier. In generating instruction-tuning data, we guard-rail with the noisy labels of CheXpert.

4.1 Computational Results

Abbreviations and Evaluation Metrics. We abbreviate IT and base (non-IT) models as subscripts “INSTRUCT” and “BASE”, respectively. Tests for statistical significance (SS) use a one-sided non-parametric Mann-Whitney U test (Mann and Whitney 1947), unless stated otherwise. The Area Under Accuracy-Rejection Curve (AUARC) and F1-score measure deferral and classification performance, respectively. Calibration is evaluated with ECE, ACE, and ECE_{Imb}. LLM efficiency is assessed based on *Rel. s/Gen.* (seconds/generation relative to the baseline), *Mem.* (GPU VRAM in GB), and *E.R.* (error rate: unsuccessful to suc-

cessful generations).

Imbalanced ECE is an Effective Measure of Calibration.

Figure 3 shows the limitations of ECE and ACE in imbalanced situations using $\hat{e}_{\text{IT-13B}}$ predictions. ECE concentrates most bin weight in the first bin, neglecting other bins’ calibration errors. ACE improves on ECE but still mainly represents the negative class. Imbalanced ECE resolves this by modifying ECE to reflect calibration errors more globally.

Instruction-Tuned Models Are More Accurate, Equally Calibrated, Better Deferral Systems, and More Efficient.

We find a SS improvement in F1-Score for $\hat{t}_{\text{IT-13B}}$ compared to $\hat{t}_{\text{BASE-70B}}$, demonstrating the success of our instruction-tuning methodology [Table 1, Rows (1) & (4)]. $\hat{e}_{\text{IT-13B}}$ shows the best deferral performance, surpassing $\hat{e}_{\text{BASE-70B}}$ [Table 1, Rows (8) & (5)]. $\hat{\mu}_{\text{IT-13B}}$ achieves the highest classification performance [Table 1, Row (9)], outperforming $\hat{e}_{\text{IT-13B}}$ and $\hat{t}_{\text{IT-13B}}$ ($p < 0.01$). This improvement is consistent across all setups, highlighting the distinct value of verbalised and hidden-state sources for classification. There is a positive correlation of $\rho = 0.53 \pm 0.07$ between these predictions. When $\hat{\mu}$ is correct and \hat{t} and \hat{e} differ, a correct \hat{e} prediction is associated with an incorrect \hat{t} prediction $75 \pm 0.14\%$ of the time, and vice versa $75 \pm 0.25\%$ of the time. This trend is also seen in MIMIC-500 experiments, where $\hat{\mu}_{\text{IT-13B}}$ achieves the highest macro-averaged F1-score (Table 3). Additionally, the IT model significantly improves across all

CONDITION	PREDICTION SOURCE			
	$\hat{t}_{\text{BASE-70B}}$	GPT-3.5	$\hat{\mu}_{\text{IT-13B}}$	GPT-4
ALECTASIS	95.96	95.77	<u>98.02</u>	98.52
CONSOLIDATION	85.22	85.50	<u>98.04</u>	99.01
EFFUSION	87.32	90.43	<u>96.69</u>	97.69
FRACTURE	76.19	93.88	80.00	88.89
LUNG OPACITY	79.88	87.16	88.17	93.06
NODULE	72.34	80.95	86.36	85.71
PLEURAL LESION	44.44	95.08	77.78	90.00
PNEUMOTHORAX	81.63	60.24	97.78	<u>95.45</u>
PULMONARY EDEMA	87.13	67.24	92.68	85.47
*W.M.S	86.24	73.97	<u>89.72</u>	94.12
MICRO-AVG.	84.94	92.68	<u>93.17</u>	93.87
MACRO-AVG.	79.71	89.64	<u>90.52</u>	92.79
EFFICIENT, EFFECTIVE AND OPEN-SOURCE?	✗	✗	✓	✗

Table 3: **Our proposed guided deferral LLM is competitive with state-of-the-art proprietary LLMs in classification.** F-scores for classification of chest X-ray conditions on the MIMIC-500 dataset for proprietary LLMs (GPT-3.5 and GPT-4) against our guided deferral IT LLM on the combined probability predictions $\hat{\mu}_{\text{BASE-13B}}$. We additionally include $\hat{t}_{\text{BASE-70B}}$ for inference on the effectiveness of our instruction-tuning methodology. Highest and second-highest performing metrics are in **boldface** and underlined, respectively. SOTA LLM baselines are highlighted in grey. *W.M.S = Widened Mediastinal Silhouette.

LLM efficiency metrics [Table 2], producing the fewest errors, requiring the least memory, and taking the shortest time to generate.

The Combined Prediction Has The Greatest Deferral Performance in Mostly-Autonomous Systems. We find a SS ($p < 0.1$) improvement in deferring on $\hat{e}_{\text{IT-13B}}$ and classifying on $\hat{\mu}_{\text{IT-13B}}$ (i.e. taking $\hat{p} = \hat{e}$ and $\hat{p} = \hat{\mu}$) compared to deferring and classifying solely on $\hat{e}_{\text{IT-13B}}$ when deferral is $< 5\%$. This suggests that the combined prediction is a better predictor in high-uncertainty situations. For deferral rates $\geq 5\%$, there is no significant difference between the two strategies, but both significantly outperform all other setups.

Our Guided Deferral LLM is Competitive With Proprietary SOTA LLMs. Table 3 shows that our guided deferral LLM is competitive with proprietary SOTA LLMs. Despite using only 0.6% of the training data, it outperforms both in classifying 3 of the conditions. Overall, $\hat{\mu}_{\text{BASE-13B}}$ exceeds the micro and macro-averaged F-score of GPT-3.5 and falls just short of GPT-4, despite significantly smaller numbers of parameters, and with data privacy intact.

4.2 Pilot Study Results

Following the computational deferral results in high uncertainty, we choose $\theta = \tilde{p}_{31}$ with $\tilde{p} = \hat{e}$ and $\hat{p} = \hat{\mu}$.

Guidance Enhances Human Decision-Making. Figure 4 displays the accuracy of participant decision-making with and without guidance, compared to LLM performance.




Figure 4: Accuracy of prediction methods from pilot study. On average, guided humans (right box plot) outperform both unguided humans (left box plot) and the LLM (red dashed line) alone.

All participants' accuracy improved with guidance. We rejected the null hypothesis that performance without guidance equals performance with guidance (paired one-sided Wilcoxon signed-rank test, $p < 0.01$). When confronted with AI disagreement, the provided guidance proved effective in assisting participants to arrive at the correct final decision. This was true not only when the participant was incorrect, but importantly also when the LLM was incorrect. We rejected the null hypothesis that the proportion of humans changing their prediction is independent of the AI's prediction correctness (χ^2 -test, $p < 0.01$).

5 Discussion and Conclusion

Successful AI implementation in healthcare requires combining human and AI strengths, designing systems that are accessible and beneficial to all. Our approach develops a guided deferral framework using LLMs for clinical decisions. We demonstrate that efficient and open-source LLMs can achieve effective deferral and classification performance, competitive with SOTA proprietary systems, with data privacy intact and without the requirement of high-performance compute. We address issues with calibration metrics (ECE and ACE) in imbalanced data commonly found in healthcare tasks, proposing imbalanced ECE as a solution. The effectiveness of our guided deferral system is validated through a pilot study, of which demonstrates complementarity. Our guided deferral system is theoretically adaptable to any data modality or pathology in healthcare, as well as other high-risk domains. However, its practical application requires rigorous ethical evaluation to address the potential risks associated with deploying AI in these contexts. To prevent unexpected errors during implementation, clinicians need thorough education on the system's capabilities. This includes addressing biases such as anchoring and confirmation bias, and informing them about the system's training data distribution, including data drift and out-of-distribution instances.

Ethical Statement

The OSCLMRIC dataset was collected by Professor Jeremy Fairbank with HRA approval (IRAS Project ID: 207858, Protocol Number:12139). The use of AI in decision-critical settings, such as healthcare, raises significant ethical concerns. Our work employs a LLM to provide textual guidance to humans, using instruction-tuning to minimise hallucinations, as validated by our results. However, ethical considerations must address the risk of presenting users with misleading or incorrect information. The environmental impact of LLMs also warrants attention, as these models are computationally intensive and demand substantial energy. While we take steps toward more sustainable applications, current AI practices still carry a notable carbon footprint. Future research should prioritise energy-efficient LLMs and alternative methods to reduce environmental impact. Finally, LLMs pose privacy risks, as sensitive information can potentially be extracted from training data (Hu et al. 2022). In healthcare, where training data may include patient information, strict data sanitisation, privacy-preserving training, and cautious data use are essential to mitigate data leakage.

Acknowledgements

JS is supported by the EPSRC Center for Doctoral Training in Health Data Science (EP/S02428X/1). JS gratefully acknowledges the Kellogg College Research Support Grant of which supported this research. The authors acknowledge UKRI grant reference EP/X040186/1 (Turing AI Fellowship: Ultra Sound Multi-Modal Video-based Human-Machine Collaboration). The authors extend their gratitude to Professor. Andrew Zisserman and Dr. Amir Jamaludin from the Visual Geometry Group, University of Oxford, for the provision of this data. A special thank-you to all participants of the pilot study.

References

- Andrews, R. W.; Lilly, J. M.; Srivastava, D.; and Feigh, K. M. 2022. The role of shared mental models in human-AI teams: a theoretical review. *Theoretical Issues in Ergonomics Science*, 24(2): 129–175.
- Banerjee, D.; Teso, S.; and Passerini, A. 2023. Learning to Guide Human Experts via Personalized Large Language Models. arXiv:2308.06039.
- Bansal, G.; Nushi, B.; Kamar, E.; Lasecki, W. S.; Weld, D. S.; and Horvitz, E. 2019. Beyond accuracy: The role of mental models in human-AI team performance. In *Proceedings of the AAAI conference on human computation and crowdsourcing*, volume 7, 2–11.
- Bansal, G.; Wu, T.; Zhou, J.; Fok, R.; Nushi, B.; Kamar, E.; Ribeiro, M. T.; and Weld, D. 2021. Does the whole exceed its parts? the effect of ai explanations on complementary team performance. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, 1–16.
- Beeching, E.; Fourrier, C.; Habib, N.; Han, S.; Lambert, N.; Rajani, N.; Sanseviero, O.; Tunstall, L.; and Wolf, T. 2023. Open LLM Leaderboard. https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard.
- 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.
- Chen, J.; Yoon, J.; Ebrahimi, S.; Arik, S. O.; Pfister, T.; and Jha, S. 2023. Adaptation with Self-Evaluation to Improve Selective Prediction in LLMs. arXiv:2310.11689.
- Dettmers, T.; Pagnoni, A.; Holtzman, A.; and Zettlemoyer, L. 2023. QLoRA: Efficient Finetuning of Quantized LLMs. arXiv:2305.14314.
- Dvijotham, K. D.; Winkens, J.; Barsbey, M.; Ghaisas, S.; Stanforth, R.; Pawlowski, N.; Strachan, P.; Ahmed, Z.; Azizi, S.; Bachrach, Y.; Culp, L.; Daswani, M.; Freyberg, J.; Kelly, C.; Kiraly, A.; Kohlberger, T.; McKinney, S.; Mustafa, B.; Natarajan, V.; Geras, K.; Witowski, J.; Qin, Z. Z.; Creswell, J.; Shetty, S.; Sieniek, M.; Spitz, T.; Corrado, G.; Kohli, P.; Cemgil, T.; and Karthikesalingam, A. 2023. Enhancing the reliability and accuracy of AI-enabled diagnosis via complementarity-driven deferral to clinicians. *Nature Medicine*, 29(7): 1814–1820.
- Freyer, O.; Wiest, I. C.; Kather, J. N.; and Gilbert, S. 2024. A future role for health applications of large language models depends on regulators enforcing safety standards. *The Lancet Digital Health*, 6(9): e662–e672.
- Geng, J.; Cai, F.; Wang, Y.; Koepl, H.; Nakov, P.; and Gurevych, I. 2023. A Survey of Language Model Confidence Estimation and Calibration. arXiv:2311.08298.
- Gneiting, T.; Balabdaoui, F.; and Raftery, A. E. 2007. Probabilistic forecasts, calibration and sharpness. *J. R. Stat. Soc. Series B Stat. Methodol.*, 69(2): 243–268.
- Gu, J.; Cho, H.-C.; Kim, J.; You, K.; Hong, E. K.; and Roh, B. 2024. CheX-GPT: Harnessing Large Language Models for Enhanced Chest X-ray Report Labeling. arXiv:2401.11505.
- Guilbert, T.; Caelen, O.; Chirita, A.; and Saerens, M. 2024. Calibration methods in imbalanced binary classification. *Annals of Mathematics and Artificial Intelligence*, 1–34.
- Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On calibration of modern neural networks. In *International conference on machine learning*, 1321–1330.
- Hegel, G. W. F. 1818. *Georg Wilhelm Friedrich Hegel: The Phenomenology of Spirit*. Cambridge Hegel Translations. Cambridge University Press.
- Hu, H.; Salicic, Z.; Sun, L.; Dobbie, G.; Yu, P. S.; and Zhang, X. 2022. Membership Inference Attacks on Machine Learning: A Survey. arXiv:2103.07853.
- Irvin, J.; Rajpurkar, P.; Ko, M.; Yu, Y.; Ciurea-Ilcus, S.; Chute, C.; Marklund, H.; Haghighi, B.; Ball, R.; Shpankaya, K.; et al. 2019. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 590–597.
- Iverson, H.; Wang, Y.; Pyatkin, V.; Lambert, N.; Peters, M.; Dasigi, P.; Jang, J.; Wadden, D.; Smith, N. A.; Beltagy, I.; and Hajishirzi, H. 2023. Camels in a Changing Climate: Enhancing LM Adaptation with Tulu 2. arXiv:2311.10702.

- Jarupathirun, S.; and Zahedi, F. M. 2007. Dialectic decision support systems: System design and empirical evaluation. *Decision Support Systems*, 43(4): 1553–1570.
- Johnson, A.; Lungren, M.; Peng, Y.; Lu, Z.; Mark, R.; Berkowitz, S.; and Horng, S. 2024. MIMIC-CXR-JPG - chest radiographs with structured labels.
- Kulesza, T.; Stumpf, S.; Burnett, M.; and Kwan, I. 2012. Tell me more? The effects of mental model soundness on personalizing an intelligent agent. In *Proceedings of the sigchi conference on human factors in computing systems*, 1–10.
- Lin, S.; Hilton, J.; and Evans, O. 2022. Teaching Models to Express Their Uncertainty in Words. arXiv:2205.14334.
- Liu, J.; Gallego, B.; and Barbieri, S. 2022. Incorporating uncertainty in learning to defer algorithms for safe computer-aided diagnosis. *Scientific reports*, 12(1): 1762.
- Liu, Z.; Li, Y.; Shu, P.; Zhong, A.; Yang, L.; Ju, C.; Wu, Z.; Ma, C.; Luo, J.; Chen, C.; et al. 2023. Radiology-llama2: Best-in-class large language model for radiology. *arXiv preprint arXiv:2309.06419*.
- Madras, D.; Pitassi, T.; and Zemel, R. 2018. Predict responsibly: improving fairness and accuracy by learning to defer. *Advances in neural information processing systems*, 31.
- Mann, H. B.; and Whitney, D. R. 1947. On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics*, 18(1): 50–60.
- Nadeem, M. S. A.; Zucker, J.-D.; and Hanczar, B. 2009. Accuracy-Rejection Curves (ARCs) for Comparing Classification Methods with a Reject Option. In Džeroski, S.; Guerts, P.; and Rousu, J., eds., *Proceedings of the third International Workshop on Machine Learning in Systems Biology*, volume 8 of *Proceedings of Machine Learning Research*, 65–81. Ljubljana, Slovenia: PMLR.
- Naeini, M. P.; Cooper, G. F.; and Hauskrecht, M. 2015. Obtaining well calibrated probabilities using Bayesian Binning. *Proc. Conf. AAAI Artif. Intell.*, 2015: 2901–2907.
- Nixon, J.; Dusenberry, M.; Jerfel, G.; Nguyen, T.; Liu, J.; Zhang, L.; and Tran, D. 2020. Measuring Calibration in Deep Learning. arXiv:1904.01685.
- OpenAI. 2024. GPT-4 Technical Report. arXiv:2303.08774.
- Rastogi, C.; Tulio Ribeiro, M.; King, N.; Nori, H.; and Amershi, S. 2023. Supporting Human-AI Collaboration in Auditing LLMs with LLMs. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23. ACM.
- Ren, J.; Luo, J.; Zhao, Y.; Krishna, K.; Saleh, M.; Lakshminarayanan, B.; and Liu, P. J. 2022. Out-of-distribution detection and selective generation for conditional language models. In *The Eleventh International Conference on Learning Representations*.
- Roelofs, R.; Cain, N.; Shlens, J.; and Mozer, M. C. 2022. Mitigating bias in calibration error estimation. In *International Conference on Artificial Intelligence and Statistics*, 4036–4054.
- Tian, K.; Mitchell, E.; Zhou, A.; Sharma, A.; Rafailov, R.; Yao, H.; Finn, C.; and Manning, C. D. 2023. Just Ask for Calibration: Strategies for Eliciting Calibrated Confidence Scores from Language Models Fine-Tuned with Human Feedback. arXiv:2305.14975.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; Bikel, D.; Blecher, L.; Ferrer, C. C.; Chen, M.; Cucurull, G.; Esiobu, D.; Fernandes, J.; Fu, J.; Fu, W.; Fuller, B.; Gao, C.; Goswami, V.; Goyal, N.; Hartshorn, A.; Hosseini, S.; Hou, R.; Inan, H.; Kardas, M.; Kerkez, V.; Khabsa, M.; Kloumann, I.; Korenev, A.; Koura, P. S.; Lachaux, M.-A.; Lavril, T.; Lee, J.; Liskovich, D.; Lu, Y.; Mao, Y.; Martinet, X.; Mihaylov, T.; Mishra, P.; Molybog, I.; Nie, Y.; Poulton, A.; Reizenstein, J.; Rungta, R.; Saladi, K.; Schelten, A.; Silva, R.; Smith, E. M.; Subramanian, R.; Tan, X. E.; Tang, B.; Taylor, R.; Williams, A.; Kuan, J. X.; Xu, P.; Yan, Z.; Zarov, I.; Zhang, Y.; Fan, A.; Kambadur, M.; Narang, S.; Rodriguez, A.; Stojnic, R.; Edunov, S.; and Scialom, T. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. arXiv:2307.09288.
- Ullah, E.; Parwani, A.; Baig, M. M.; and Singh, R. 2024. Challenges and barriers of using large language models (LLM) such as ChatGPT for diagnostic medicine with a focus on digital pathology—a recent scoping review. *Diagnostic pathology*, 19(1): 43.
- Varshney, N.; and Baral, C. 2023. Post-Abstention: Towards Reliably Re-Attempting the Abstained Instances in QA. arXiv:2305.01812.
- Wang, C. 2023. Calibration in Deep Learning: A Survey of the State-of-the-Art. arXiv:2308.01222.
- Wang, D.; Churchill, E.; Maes, P.; Fan, X.; Shneiderman, B.; Shi, Y.; and Wang, Q. 2020. From human-human collaboration to Human-AI collaboration: Designing AI systems that can work together with people. In *Extended abstracts of the 2020 CHI conference on human factors in computing systems*, 1–6.
- Wei, J.; Bosma, M.; Zhao, V. Y.; Guu, K.; Yu, A. W.; Lester, B.; Du, N.; Dai, A. M.; and Le, Q. V. 2022. Finetuned Language Models Are Zero-Shot Learners. arXiv:2109.01652.
- Wiegrefe, S.; Hessel, J.; Swayamdipta, S.; Riedl, M.; and Choi, Y. 2022. Reframing Human-AI Collaboration for Generating Free-Text Explanations. arXiv:2112.08674.
- Xiong, M.; Hu, Z.; Lu, X.; Li, Y.; Fu, J.; He, J.; and Hooi, B. 2023. Can LLMs Express Their Uncertainty? An Empirical Evaluation of Confidence Elicitation in LLMs. arXiv:2306.13063.
- Zhang, S.; Dong, L.; Li, X.; Zhang, S.; Sun, X.; Wang, S.; Li, J.; Hu, R.; Zhang, T.; Wu, F.; et al. 2023. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792*.
- Zöller, N.; Berger, J.; Lin, I.; Fu, N.; Komarneni, J.; Barabucci, G.; Laskowski, K.; Shia, V.; Harack, B.; Chu, E. A.; Trianni, V.; Kurvers, R. H. J. M.; and Herzog, S. M. 2024. Human-AI collectives produce the most accurate differential diagnoses. arXiv:2406.14981.