Human-AI Collaboration in Decision-Making: Beyond Learning to Defer

Diogo Leitão [1 2 3] **Pedro Saleiro** [1] **Mário A. T. Figueiredo** [2 3] **Pedro Bizarro** [1]

Abstract

Human-AI collaboration (HAIC) in decision-making aims to create synergistic teaming between human decision-makers and AI systems. *Learning to defer* (L2D) has been presented as a promising framework to determine who among humans and AI should make which decisions in order to optimize the performance and fairness of the combined system. Nevertheless, L2D entails several often unfeasible requirements, such as the availability of predictions from humans for every instance or ground-truth labels that are independent from said humans. Furthermore, neither L2D nor alternative approaches tackle fundamental issues of deploying HAIC systems in real-world settings, such as capacity management or dealing with dynamic environments. In this paper, we aim to identify and review these and other limitations, pointing to where opportunities for future research in HAIC may lie.

1. Introduction

With the advent of machine learning (ML), artificial intelligence (AI) is now fast, scalable, and, in several applications, very accurate. As a result, AI is increasingly present in decision-making processes previously conducted solely by humans, such as in healthcare (Esteva et al., 2017), or in criminal justice (Kirchner & Larson, 2017). The AI may take on a pure advisory role, providing the human decision-maker with a prediction, a score, or explanations. However, in several applications, the volume of cases is so large that it is impractical or even economically unfeasible to solely rely on human decision-making. This is the case in fraud detection (Yousefi et al., 2019), and credit scoring (Burrell, 2016). Because AI models are able to process large volumes of information, they offer superior operational efficiency and statistical precision to that of humans, who are known to

struggle in this regard (Meehl, 1954). With that said, humans are capable of swiftly learning from few data points, and can leverage information that the AI cannot grasp, either because it cannot be converted to symbolic representation, or because it originates from accrued experience. As such, often the ideal system consists neither of AI alone, nor of humans alone, but rather of a collaboration between the two.

A key challenge in *human-AI collaboration* (HAIC) is to determine who makes which decisions. An ideal HAIC system is capable of identifying the strengths and weaknesses of humans and AI in order to optimally assign instances to decision-makers. This encompasses modelling not only predictive performance, but also **fairness**. ML models are known to be capable of producing biased decisions against protected groups (Angwin et al., 2016; Buolamwini & Gebru, 2018). Humans, too, can be prejudiced, especially when information correlated with protected attributes is available to them. HAIC systems have the opportunity to reduce unfairness by optimizing assignments to mitigate group-wise disparities.

The current state-of-the-art approach for performance- and fairness-aware assignments in HAIC is *learning to defer* (L2D) (Madras et al., 2018; Mozannar & Sontag, 2020; Keswani et al., 2021). L2D frames the assignment problem as a classification task, providing a supervised method to find the best decision-maker for each given instance. Nevertheless, L2D exhibits structural limitations: it **requires predictions from every human decision-maker for every instance** in the training set, and trades off **robustness** for specialization. Additionally, **capacity management** in teams of humans is disregarded: L2D aims to find the optimal assignment given an individual instance, instead of aiming to find the optimal set of assignments for a set of instances given a team of humans with limited capacity.

Settings where decisions influence outcomes are also usually not considered. A particularly ubiquitous problem is **selective labels** (Lakkaraju et al., 2017). For instance, in criminal justice, rearrest can only take place if bail is first granted; in credit scoring, failing to repay a loan can only happen if a loan is granted. Lastly, a key issue still unaddressed in the literature is dealing with **dynamic environments**, which may entail addressing concept drift (Gama et al., 2014) and feedback loops (Dalvi et al., 2004; Perdomo et al., 2020).

---

[1]Feedzai [2]Instituto Superior Técnico, Universidade de Lisboa [3]Instituto de Telecomunicações. Correspondence to: Diogo Leitão <diogo.leitao@feedzai.com>.

To avoid performance degradation, HAIC systems must account for non-stationary factors, or at least be updated with new data. However, L2D offers no solution to learn from partial, non-random data, and neither do alternative contributions.

In short, HAIC systems have the opportunity to optimize performance and fairness beyond that of individual contributors. Although L2D provides a valid framework to this end, it entails several structural limitations that inhibit its adoption in real-world use-cases and disregards other key challenges. The goal of this paper is to clearly define these challenges and limitations, while shedding light on possible avenues of future research to bring us closer to performant, fair, and practical human-AI collaboration.

2. Learning to Defer

The simplest approach to manage assignments in HAIC is to defer based on model confidence. This approach is drawn from learning with a reject option, a framework first studied by Chow (1970), where the goal is to optimize the performance of the non-rejected predictions by rejecting to predict in cases of high uncertainty. The baseline for uncertainty estimation is to take the predicted probability score(s) as a measure of confidence (Hendrycks & Gimpel, 2016), while more sophisticated techniques range from ensemble methods (Gal & Ghahramani, 2016; Lakshminarayanan et al., 2017) to auxiliary models (Corbière et al., 2019). In HAIC systems, these uncertainty estimation techniques can be leveraged to defer low model confidence cases to humans.

Madras et al. (2018) introduced learning to defer (L2D) as an improvement upon confidence-based deferral in the context of HAIC systems with a single human expert. The authors argue that confidence-based deferral is suboptimal, as it fails to consider the downstream human decision-maker. In some instances of high model uncertainty, they may be just as inaccurate as the model; as such, it may be preferable to defer to them other lower uncertainty instances where they are able to outperform the model. To model these differences, Madras et al. (2018) expand upon the framework of learning with a reject option. In particular, they adapt the work of Cortes et al. (2016), who incorporated the reject option into the learning algorithm, allowing the classifier to reject to predict, incurring instead in a constant, label-independent cost. To account for the performance of the human decision-maker, L2D adds a variable cost of deferral to this setup:

$$\mathcal{L}_{system}(\hat{Y}_M, \hat{Y}_H, s) = \sum_i [(1 - s_i)\mathcal{L}_{CE}(Y_i, \hat{Y}_{M,i})$$
$$+ s_i \mathcal{L}_{0\text{-}1}(Y_i, \hat{Y}_{H,i}) + s_i \gamma_{defer}]$$

The system loss depends on the choice to defer ($s = 1$) or not to defer ($s = 0$) to the human decision-maker H. If the human decision-maker is chosen, the system loss is the 0-1 loss between ground-truth label Y and human prediction $\hat{Y}_H$, plus a constant penalty term γ_{defer}. Conversely, if the decision is not deferred, the system loss corresponds to the cross-entropy loss of model M's prediction $\hat{Y}_M$.

Madras et al. (2018) propose a learning scheme consisting of two separate neural networks that are jointly trained to minimize the system loss. The main model focuses on the classification task at hand; the deferral model decides whether to assign the decision to the model or to the human decision-maker. During training, the output of the deferral model is taken as a Bernoulli probability, with both networks backpropagating on the expected value of the system loss. Consequently, the deferral model learns to choose the decision-maker with best expected performance, and the main classifier is allowed to specialize on the cases that are not deferred, disregarding the rest.

An additional concern of Madras et al. (2018) is that both the human decision-maker and the ML model may be biased against certain protected groups. To allow the L2D system to mitigate these biases, the authors propose a regularized loss with an additional penalty term for error-rate disparities between protected groups. We will further develop on the topic of fairness in HAIC in Section 3.6.

Others authors have since joined the effort of expanding and improving learning to defer. Mozannar & Sontag (2020) generalize L2D to multi-class settings and introduce a consistent surrogate loss to be used in a single classifier capable of deferral. Verma & Nalisnick (2022) elaborate upon their work, proposing a loss calibrated for human correctness. Keswani et al. (2021) expand L2D to model multiple human decision-makers. Lastly, some authors propose different methodologies (Wilder et al., 2020; Gao et al., 2021), or different end-goals (Bansal et al., 2021). In the next section, we will shed light on the limitations of L2D, as well as unaddressed challenges by this and alternative frameworks.

3. Challenges and Limitations

To help illustrate the shortcomings of L2D, we will take the domain of online merchant fraud detection as a running example. In this use-case, ML models and human fraud analysts work in tandem to identify and stop fraud in streams of thousands of daily purchases. As such, the AI is used to process the majority of purchases, while humans, due to limited capacity, only intervene on a small fraction (e.g. 5%) of the most difficult instances.

3.1. Learning from data without human predictions

Learning to defer entails a fundamental, structural limitation: it requires predictions from human decision-makers

for every training instance. This implies an additional effort that goes much beyond what is strictly necessary in regular operations. One cannot expect a team of fraud analysts hired to intervene only on the most difficult cases to be able to review all the instances necessary to train an ML model.

Mozannar & Sontag (2020) consider this problem, and suggest an imputation scheme where the behavior of the human decision-makers is modelled from the complete data and replicated on the incomplete data. However, that method fails when the complete and incomplete data are not identically distributed, as the basis for generalization falters. This will be the case whenever the collected data stems from a non-random assignment system. In real-world use-cases, this is to be expected, as aforementioned techniques such as deferring low-confidence instances to the human decision-makers already bring substantial performance gains. In fraud detection, analysts do not review random cases; they review cases where the ML model is less confident.

As such, L2D cannot be implemented as an improvement to an existing deferral system. In fact, if a L2D system needs to be updated to keep up with changing trends in the data, the framework cannot be reused to perform this update, as non-deferred instances will lack a human prediction, and will be distributed differently from deferred instances. We will revisit this point in Section 3.7, elaborating on why model updating will, in fact, be necessary in most settings.

Future research should focus on being able to learn from data with and without human predictions, as that is its natural state in HAIC settings. An interesting approach would be to resort to reinforcement learning techniques, where partial and biased information is the norm. Choices of decision-maker could be framed as actions, with the environment being the features and any other relevant information.

3.2. Joint learning: losing robustness and advisory

One of the proposed benefits of L2D is that, by jointly training the main classifier with the deferral model, the former is allowed to specialize on non-deferred instances, in detriment of those that are deferred (Madras et al., 2018; Mozannar & Sontag, 2020). However, by design, specialization renders the main classifier useless in the instances likely to be deferred, as gradients are stopped from backpropagating into it. As such, **learning to defer is unsuitable in domains where the AI advises the human decision-makers**. In such use-cases, the ML model may provide the human decision-maker with a tentative prediction, a score, or a set of explanations, that they will then combine with their own accrued experience and analysis of the instance at hand. In fraud detection, the ML model score is one of the pieces of information considered by analysts when making a prediction (Jesus et al., 2021). This would not be possible if the classifier specialized away from the deferred instances.

Furthermore, **specialization makes the system brittle**: by trading off generalization for specialization, the AI is not robust to post-deployment changes on human capacity for review. If fraud analysts become temporarily unavailable, the AI will not be capable of deciding on a fraction of instances, seriously harming performance.

A promising avenue of research could be to develop an alternative approach to L2D that follows a sequential learning setup instead of joint learning. This option is briefly considered by Madras et al. (2018), when the authors suggest stopping system error from backpropagating to the main classifier; however, results of this alternative setup are not shown, and it has not been considered in posterior contributions. A HAIC assignment system trained sequentially would solely focus on finding the best assignment policy given the pre-trained classifier and the set of human decision-makers. This approach would not only avoid the aforementioned drawbacks of altering the main classifier, but also eliminate the need for human predictions for every instance, which, as mentioned in Section 3.1, is often an unfeasible demand.

3.3. Deferral to multiple decision-makers

Madras et al. (2018) introduce learning to defer in the context of **deferral to a single human decision-maker**. However, in several applications, teams may be composed of several people, and they may have different biases and types of expertise. For example, fraud analysts may be specialized in specific fraud patterns.

Keswani et al. (2021) propose a L2D method that accounts for the diversity of human decision-makers, modelling each available expert individually. However, to this end, they require not only human predictions for every instance — as in single-human L2D — but also predictions from every considered decision-maker. This data will rarely be available in real-world use-cases, as it is inefficient for individuals in teams of human experts to review each other's decisions. If every fraud analyst in a team with n elements starts reviewing every deferred instance, the capacity for reviews decreases by a factor of n, harming the system's performance. Keswani et al. (2021) posit that missing predictions may be imputed; however, the robustness of that method is not discussed. If the data reflects non-random assignments, then generalization may not be possible.

Future work should focus on modelling several decision-makers without requiring concurrent predictions, as those will rarely be available without additional effort. As mentioned in Section 3.1, this could be accomplished with reinforcement learning techniques, where partial information is the norm.

3.4. Capacity management

Learning to defer frames deferral as a classification task: given a specific instance, the goal is to find which decision-maker is expected to be the most accurate and less biased. However, in the presence of capacity constraints, the best decision-maker may not be the optimal choice. For example, if a fraud analyst is universally better than the rest of the team, then, ideally, they would decide on only the hardest cases where others are most likely to err.

Current L2D methods would simply assign them to every instance, disregarding capacity constraints. While this may be acceptable in use-cases where the hiring process is fully flexible, such as with Amazon Mechanical Turk, in human-AI teams with stricter arrangements, such extreme assignments are unfeasible. As a result, for an optimal deferral solution, capacity must be managed explicitly. We believe this to be a blind spot and an opportunity for future research.

3.5. Selective labels

Human-AI decision-making teams often operate in environments where actions can influence outcomes. In fraud detection, a purchase may only be found to be fraudulent if it is not declined. Lakkaraju et al. (2017) refer to this as the selective labels problem. L2D is incapable of dealing with selective labels, as it requires ground-truth labels to evaluate human decision-makers on.

Gao et al. (2021) propose an alternative method anchored not on predictions and ground-truth labels, but rather actions and rewards. For instance, in fraud detection, the reward of past declined purchases is always zero, regardless of the (unobservable) ground-truth label, as no revenue or loss occurs. On the other hand, predicted negatives either have a negative reward — the cost of accepting a fraudulent purchase —, or a positive reward — the revenue of a legitimate purchase. (Gao et al., 2021) propose using this bandit feedback to learn both a decision-making model and a routing model, leveraging counterfactual risk minimization (Swaminathan & Joachims, 2015). Despite modelling rewards instead of labels, their approach still requires human predictions for every training instance.

Alternatively, selective labels can also be directly addressed. Imputation techniques may be used, leveraging expert consensus when available (De-Arteaga et al., 2018; Keswani et al., 2022). Furthermore, selective labels do not completely eliminate relevant signal, as predicted negatives will still result in either true or false negatives. Consequently, performance can be compared on the basis of the trade-off between predicted positive prevalence and false negative prevalence (Lakkaraju et al., 2017). A HAIC system based on these signals could still be beneficial, without requiring additional data.

3.6. Fairness

As mentioned in Section 1, both humans and ML models are capable of bias against protected groups (Angwin et al., 2016; Buolamwini & Gebru, 2018). When introducing learning to defer, Madras et al. (2018) were keen on building a framework that optimizes not only for performance, but also for fairness. Considering the ML model's biases, as well as those of humans, allows the deferring system to mitigate existing biases. On the contrary, introducing customized deferral systems without considering fairness may result in aggravating existing biases, as demonstrated by Jones et al. (2020) in the context of learning with rejection.

Fairness in HAIC represents not only an opportunity, but also a threat, and should thus always be considered. While most contributions on L2D do consider fairness (Madras et al., 2018; Keswani et al., 2021), other contributions inside and outside L2D do not (Mozannar & Sontag, 2020; Gao et al., 2021; Bansal et al., 2021). With AI being increasingly used to support decision-making in key societal sectors, future research should aim to maintain fairness at the forefront of human-AI collaboration.

3.7. Dynamic environments

Lastly, a key unaddressed issue is dealing with dynamic environments. The underlying data distribution may change due to new trends, resulting in concept drift (Gama et al., 2014). Furthermore, in performative prediction settings (Perdomo et al., 2020), the actions of the system may influence the environment. In adversarial classification settings (Dalvi et al., 2004), it may even elicit a response (e.g. fraudsters changing strategies as a response to fraud detection systems). The system's decision-makers may also change their behavior, as they learn new concepts, update beliefs, or are influenced by exogenous factors. HAIC systems must account for all these non-stationary factors, or at least be updated with new data to avoid performance degradation.

L2D does not endogenously model any of the aforementioned non-stationarity factors, and it cannot be updated with new data naturally obtained from the HAIC environment (see Section 3.1). Future research should focus on developing HAIC systems that do not have this limitation.

4. Conclusion

In this paper, we shed light on the structural limitations and unaddressed challenges of the state-of-the-art method for managing assignments in human-AI collaborative decision-making — learning to defer. Our goal is to motivate research moving towards a holistic human-AI collaboration system that learns to optimize performance and fairness from the available data, while managing existing capacity constraints, and keeping up with dynamic environments.

References

Angwin, J., Larson, J., Mattu, S., and Kirchner, L. Machine bias risk assessments in criminal sentencing. *ProPublica, May*, 23, 2016.

Bansal, G., Nushi, B., Kamar, E., Horvitz, E., and Weld, D. S. Is the most accurate ai the best teammate? optimizing ai for teamwork. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 11405–11414, 2021.

Buolamwini, J. and Gebru, T. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pp. 77–91. PMLR, 2018.

Burrell, J. How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society*, 3 (1):2053951715622512, 2016.

Chow, C. K. On optimum recognition error and reject tradeoff. *IEEE Trans. Inf. Theory*, 16:41–46, 1970.

Corbière, C., Thome, N., Bar-Hen, A., Cord, M., and Pérez, P. Addressing failure prediction by learning model confidence. *Advances in Neural Information Processing Systems*, 32, 2019.

Cortes, C., DeSalvo, G., and Mohri, M. Learning with rejection. 2016.

Dalvi, N., Domingos, P., Sanghai, S., and Verma, D. Adversarial classification. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 99–108, 2004.

De-Arteaga, M., Dubrawski, A., and Chouldechova, A. Learning under selective labels in the presence of expert consistency. *arXiv preprint arXiv:1807.00905*, 2018.

Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., and Thrun, S. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639):115–118, 2017.

Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1050–1059, New York, New York, USA, 20–22 Jun 2016. PMLR.

Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., and Bouchachia, A. A survey on concept drift adaptation. *ACM computing surveys (CSUR)*, 46(4):1–37, 2014.

Gao, R., Saar-Tschansky, M., De-Arteaga, M., Han, L., Lee, M. K., and Lease, M. Human-AI collaboration with bandit feedback. *arXiv preprint arXiv:2105.10614*, 2021.

Hendrycks, D. and Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.

Jesus, S. M., Belém, C., Balayan, V., Bento, J., Saleiro, P., Bizarro, P., and Gama, J. How can I choose an explainer? An application-grounded evaluation of post-hoc explanations. *CoRR*, abs/2101.08758, 2021.

Jones, E., Sagawa, S., Koh, P. W., Kumar, A., and Liang, P. Selective classification can magnify disparities across groups. *arXiv preprint arXiv:2010.14134*, 2020.

Keswani, V., Lease, M., and Kenthapadi, K. Towards unbiased and accurate deferral to multiple experts. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 154–165, 2021.

Keswani, V., Lease, M., and Kenthapadi, K. Designing closed human-in-the-loop deferral pipelines. *arXiv preprint arXiv:2202.04718*, 2022.

Kirchner, L. and Larson, J. How we analyzed the compas recidivism algorithm. 2017. URL https : / / www.propublica.org / article / how - we - analyzed - the - compas - recidivism-algorithm.

Lakkaraju, H., Kleinberg, J., Leskovec, J., Ludwig, J., and Mullainathan, S. The selective labels problem: Evaluating algorithmic predictions in the presence of unobservables. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 275–284, 2017.

Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.

Madras, D., Pitassi, T., and Zemel, R. Predict responsibly: Improving fairness and accuracy by learning to defer. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, pp. 6150–6160, Red Hook, NY, USA, 2018. Curran Associates Inc.

Meehl, P. E. Clinical versus statistical prediction: A theoretical analysis and a review of the evidence. 1954.

Mozannar, H. and Sontag, D. Consistent estimators for learning to defer to an expert. In *Proceedings of the 37th International Conference on Machine Learning*, ICML'20. JMLR.org, 2020.

Perdomo, J., Zrnic, T., Mendler-Dünner, C., and Hardt, M. Performative prediction. In *International Conference on Machine Learning*, pp. 7599–7609. PMLR, 2020.

Swaminathan, A. and Joachims, T. Counterfactual risk minimization: Learning from logged bandit feedback. In *International Conference on Machine Learning*, pp. 814–823. PMLR, 2015.

Verma, R. and Nalisnick, E. Calibrated learning to defer with one-vs-all classifiers. *arXiv preprint arXiv:2202.03673*, 2022.

Wilder, B., Horvitz, E., and Kamar, E. Learning to complement humans. *arXiv preprint arXiv:2005.00582*, 2020.

Yousefi, N., Alaghband, M., and Garibay, I. A comprehensive survey on machine learning techniques and user authentication approaches for credit card fraud detection. *arXiv preprint arXiv:1912.02629*, 2019.