

Role of Human-AI Interaction in Selective Prediction

Elizabeth Bondi,^{1*} Raphael Koster,² Hannah Sheahan,² Martin Chadwick,² Yoram Bachrach,² Taylan Cemgil,² Ulrich Paquet,² Krishnamurthy (Dj) Dvijotham^{3*}

¹Harvard University

²DeepMind

³Google Brain

ebondi@g.harvard.edu, {rkoster, hsheahan, mjchadwick, yorambac, taylancemgil, upaq}@deepmind.com, dvij@google.com

Abstract

Recent work has shown the potential benefit of selective prediction systems that can learn to defer to a human when the predictions of the AI are unreliable, particularly to improve the reliability of AI systems in high-stakes applications like healthcare or conservation. However, most prior work assumes that human behavior remains unchanged when they solve a prediction task as part of a human-AI team as opposed to by themselves. We show that this is not the case by performing experiments to quantify human-AI interaction in the context of selective prediction. In particular, we study the impact of communicating different types of information to humans about the AI system’s decision to defer. Using real-world conservation data and a selective prediction system that improves expected accuracy over that of the human or AI system working individually, we show that this messaging has a significant impact on the accuracy of human judgements. Our results study two components of the messaging strategy: 1) Whether humans are informed about the prediction of the AI system and 2) Whether they are informed about the decision of the selective prediction system to defer. By manipulating these messaging components, we show that it is possible to significantly boost human performance by informing the human of the decision to defer, but not revealing the prediction of the AI. We therefore show that it is vital to consider how the decision to defer is communicated to a human when designing selective prediction systems, and that the composite accuracy of a human-AI team must be carefully evaluated using a human-in-the-loop framework.

Introduction

Despite significant progress in machine learning-based AI systems, applications of AI to high stakes domains remains challenging. One of the main challenges is in assessing the reliability of predictions made by a trained machine learning system, particularly when there is a distribution shift between the data the system was trained with and data encountered at deployment. In such situations, communicating the uncertainty associated with ML predictions appropriately is critical (Kompa, Snoek, and Beam 2021; Grote 2021).

Given the difficulty of communicating probabilities to human users (Galesic and Garcia-Retamero 2010; Zhang

and Maloney 2012), a pragmatic alternative is to determine whether an AI system is more likely to make an erroneous prediction than a human, and defer to a human in such cases. A number of such settings have been studied in the literature, including selective prediction (Wiener and El-Yaniv 2013), learning to defer (Mozannar and Sontag 2020) and classification with a reject option (Chow 1970). While there are nuances that differentiate these works, in this paper, we will collectively refer to this body of literature as selective prediction and only emphasize the differences where relevant to our work. A related line of work considers human-AI teams in which humans receive AI assistance but make the final decision (Bansal et al. 2021), and systems in which the AI makes the final decision with human input (Wilder, Horvitz, and Kamar 2020). These prior works either: a) Assume that the human behaves identically even when they know that they are part of a human-AI team or b) Assume a utility-maximizing model for a human decision maker.

However, it has been documented that human-AI interaction may be more complex due to a range of cognitive phenomena. For example, humans have been shown to rely excessively on AI predictions (anchoring bias) (Rastogi et al. 2022; Logg, Minson, and Moore 2019), or even distrust AI predictions after observing AI mistakes (Dietvorst, Simmons, and Massey 2015). Some work has begun to investigate solutions to these issues (Bućinca, Malaya, and Gajos 2021; Rastogi et al. 2022), however they have not focused on selective prediction systems, and context is critical.

In this work, we focus on binary classification tasks and study selective prediction systems (Fig. 1) that determine whether to rely on the outputs of an AI or defer to a human. To evaluate the overall performance of such a selective prediction system, it is important to model how the messaging, hereafter referred to as selective prediction messaging (SPM), that communicates the decision to defer impacts human accuracy. We run experiments with human subjects solving a challenging binary prediction task (that of detecting whether an animal is present in a camera trap image) and study the impact of different choices for communicating the deferring AI system’s decision. We then perform statistical analysis on the human responses under various choices of SPM, and show that the choice of SPM significantly impacts human performance. Our results isolate two ingredients for a statistically effective communication strategy, that is, com-

*Work done at DeepMind

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

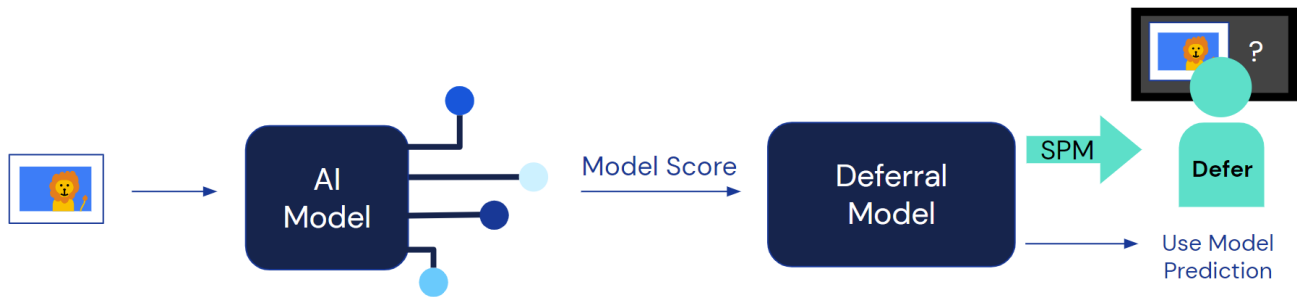


Figure 1: Deferral workflow: (1) obtain AI model prediction for given input; (2) use model score as input to deferral model to decide whether to defer; (3) defer to humans if necessary, using SPM based on the prediction and deferral status.

municating that an AI system deferred (deferral status) and the AI system’s predictions. Manipulating these leads to a boost in overall accuracy. We provide some plausible explanations for this phenomenon and suggest avenues for further work. Our contributions are therefore as follows:

- We develop and implement a balanced experimental design that can be used to measure the impact of SPM on the accuracy of selective prediction systems.
- We investigate the consequences of SPM on joint human-AI performance in a conservation prediction task, as opposed to prior work which assumes human behavior remains static during deferral.
- We discover two separable SPM ingredients, deferral status and prediction, which have distinct, significant effects on human performance, and demonstrate that manipulating these ingredients leads to improved human classifications in a human-AI team, implying that the setup and the information given to humans during such tasks has a large impact on the performance of a human-AI team.
- We suggest that our results may relate to a more general property of naturalistic datasets, that in conditions that are ambiguous, sharing AI model predictions with humans can be detrimental.

Related Work

We start by describing prior work considering different potential roles in human-AI teams, namely, decision aids and systems in which an AI model defers to a human on challenging cases only. We then discuss supporting human-AI decision-making and selective prediction algorithms.

Decision Aids: Potential deployment scenarios for human-AI teams that have been discussed in the literature vary greatly depending on the application. One increasingly common scenario, particularly in high-risk domains, is that of AI systems serving as decision aids to humans making a final decision. For example, Green and Chen (2019) explore the scenario of human decision-makers using a risk assessment model as a decision aid in financial lending and criminal justice (specifically, pretrial detention) settings. They found that humans were biased and that they failed to properly evaluate or take model performance into account, across different messaging conditions communicating the

model’s predicted risk. A decision aid for risk assessment is similarly studied in De-Arteaga, Fogliato, and Choudhchova (2020), particularly in the real-world domain of child maltreatment hotline screening. Humans indeed changed their behavior based on the risk assessment tool, but they were able to identify model mistakes in many cases. Gaube et al. (2021) conduct experiments to measure the interaction between medical AI systems and clinicians. Radiologists who participated reported AI advice as lower quality than human advice, though all advice truly came from humans. Furthermore, clinician diagnostic accuracy was reduced when they were given incorrect predictions from the AI system. Even with humans making the final decision in each case, there is a great deal of variability between these AI systems, their impacts, and their application domains.

Deferral: AI systems have also played a slightly different role by making decisions in straightforward cases and deferring to human decision makers otherwise, which is most similar to our scenario. Wilder, Horvitz, and Kamar (2020) defer to an expert on cases that are best suited for human decision-making compared to model-based decision-making (determined using end-to-end learning), yet the final system is evaluated on historical human data, meaning humans did not know they were part of a human-AI team while labeling. Keswani, Lease, and Kenthapadi (2021) propose deferral to multiple experts, including a classifier, by learning about the experts from their decisions. While Amazon Mechanical Turk was used to collect labels from participants to train and evaluate this model, it is similar to using historical human labels, as the humans again completed the task without knowing the deferral status or model prediction. Such deferral models are not yet widely deployed, nor have their impacts on humans been studied.

Supporting Human-AI Decision-Making: To mitigate some of the known negative impacts of AI on human decision-making, it may be beneficial to present humans with further information, such as uncertainty. Bhatt et al. (2021) find that humans may unreliably interpret uncertainty estimates, but that the estimates may increase transparency and thereby performance. In a system used by people without a background in statistics, for example, presenting categories (such as deferral status) may make it easier to in-

terpret. Bhatt et al. (2021) advocate for testing this with humans, including of different skill levels and in different domains. Further suggestions for positive human-AI interaction are presented in Amershi et al. (2019), including to “Show contextually relevant information,” and “Scope services when in doubt.” We are interested in finding the best way to implement these ideas with SPM.

Selective Prediction: Selective prediction can be traced back to the seminal work of Chow (1970), where theoretical properties of classifiers that are allowed to “reject” (refrain from making a prediction) and ideal rejection strategies for simple classifiers are investigated. In these settings, the main metrics are the accuracy of the classifier on the non-rejected inputs and the rate of rejection, and the natural trade-off between these. A recent survey of theoretical work in this area can be found in Wiener and El-Yaniv (2013). We omit an extensive review of selective prediction literature, but acknowledge there is additional work in this area.

More relevant to our work is the work on learning to defer (Madras, Pitassi, and Zemel 2017; Mozannar and Sontag 2020). Madras, Pitassi, and Zemel (2017) propose to defer to a human decision maker selectively in order to improve accuracy and fairness of a base classifier. Mozannar and Sontag (2020) develop a statistically consistent loss function to learn a model that both predicts and defers. Geifman and El-Yaniv (2017) develop deferral strategies purely based on the confidence estimate of an underlying classifier.

None of these works model the impact of deferral (or its communication to a human decision maker) on the accuracy of a human decision maker. Since this is the primary object of study in our work, we do not include a full selective prediction literature review. We use a simple confidence-based deferral strategy inspired by the work of Geifman and El-Yaniv (2017), but our experimental design is compatible with any selective prediction system.

Background

The primary goal of our work is to evaluate the impact of SPM on human performance in selective prediction systems. We design an experiment to evaluate this impact and base both our design and the questions we study on the psychological literature on joint decision-making.

Psychology Literature on Human Decision-Making:

The psychology literature has extensive studies on decision-making in human teams. Three psychological phenomena stand out as being relevant to our work: 1) Humans are sensitive to the specific way that a task is framed (Tversky and Kahneman 1981), 2) Humans are capable of flexibly deploying greater attentional resources in response to changing task demands and motivation (Kool and Botvinick 2018), and 3) When deciding how best to integrate the decisions of others, humans take into account their own decision confidence as well as the inferred or stated confidence of other decision makers (Boorman et al. 2013; Bahrami et al. 2010).

This suggests that SPM may impact how humans perceive their task, how much they trust the AI system, and ultimately how accurate their final prediction is, which is directly related to the composite performance of a selective prediction



Figure 2: Example camera trap image with distant (circled) animals.

system. We consequently propose an experimental design where we take four natural choices for SPM and estimate their impact on the accuracy of human labelers.

Dataset: The dataset we use in this work is composed of images from camera traps, which are cameras triggered to capture images when there is nearby motion. These can be used to capture images of animals to understand animal population characteristics and even animal behavior, both of which are useful for conservation planning purposes. The volume of images generated in this manner is too high for manual inspection by rangers and scientists directly involved in conservation and monitoring efforts.

To alleviate this burden, the Snapshot Serengeti¹ project was set up to allow volunteers to apply rich labels to camera trap images². These labels are publicly available³, and ground truth comes from label consensus from multiple individuals (Swanson et al. 2015). Given these labels, AI models have been developed that automatically classify and/or detect animals in camera trap images (Norouzzadeh et al. 2018; Tabak et al. 2019; Beery et al. 2020).

Whether relying on an AI model or volunteers, this processing is still difficult. Roughly seven out of ten images contain no animals, as they are the result of false triggers, e.g., due to heat and/or wind. However, it can be challenging to determine which images contain an animal at all, let alone the species, because of challenges like animal camouflage, distance to the camera, or partial visibility in the camera’s field of view. An example camera trap image with animals on the horizon is shown in Fig. 2.

We consider a binary task where the bulk of blank images are first removed before images are uploaded to be labeled by volunteers or a species-identifying AI model. We investigate the use of a selective prediction model to filter out blank images, while prioritizing images for human review for which the blank/animal AI model is uncertain. This is similar to Willi et al. (2019); Norouzzadeh et al. (2021), as they also involve human-AI teams and remove blank images before species identification. However, 1) The human-

¹<https://www.zooniverse.org/projects/zooiniverse/snapshot-serengeti>

²<https://lila.science/datasets/snapshot-serengeti>

³https://lilablobssc.blob.core.windows.net/snapshotserengeti-v-2-0/zooiniverse_metadata/raw_data_for_dryad.csv.zip

AI teams differ, e.g., Norouzzadeh et al. (2021) does not involve humans to remove blank images, but uses active learning for species identification, and 2) They do not focus on human behavior. To begin our workflow (Fig. 1), we obtain model scores from an ensemble of AI models that filters out blank images, then develop a deferral mechanism.

Deferral Mechanism: We use a selective prediction algorithm which finds the optimal threshold(s) for AI model scores to defer, as illustrated in Fig. 1. Concretely, consider a binary classification task with inputs x and labels $y \in \{0, 1\}$. We assume that we are given a pretrained ensemble of AI models m and have access to the predictions made by a human h as well. Given inputs x (for example, pixels of an image), we obtain a continuous score $m(x) \in [0, 1]$ that represents the confidence of the ensemble that the label corresponding to x is 1.

The deferral mechanism we use is a simple rule-based system that identifies ranges of the model score where the model is less likely to be accurate than a human and defers on these. In particular, we use a deferral model that identifies one continuous interval in the model score space to defer on.

The deferral model is parameterized by two real numbers $0 \leq \theta_1 \leq \theta_2$ and is defined as

$$\text{defer}(x; \theta) = \begin{cases} 1 & \text{if } m(x) \in [\theta_1, \theta_2] \\ 0 & \text{otherwise} \end{cases}$$

$\text{defer}(x) = 1$ represents the decision to defer on input x and $\text{defer}(x) = 0$ represents the decision to predict. Given θ , the AI model prediction x and the prediction made by a human $h(x)$, the selective prediction system is given by

$$\text{sp}(x; \theta) = \begin{cases} h(x) & \text{if } \text{defer}(x; \theta) \\ m(x) & \text{otherwise} \end{cases}$$

Given a dataset \mathcal{D} of inputs, corresponding model scores, ground truth labels, and human labels, we choose θ by solving the following optimization problem:

$$\max_{\theta} \text{Accuracy}(\mathcal{D}; \theta) \text{ subject to } \text{DeferralRate}(\mathcal{D}; \theta) \leq r$$

where $\text{Accuracy}(\mathcal{D}; \theta)$ refers to the accuracy of the selective prediction system sp with parameters θ on the dataset, $\text{DeferralRate}(\mathcal{D}; \theta)$ refers to the fraction of points in the dataset for which $\text{defer}(x; \theta) = 1$, and r is a bound on the deferral rate, reflecting the acceptable level of human effort or budget constraints on hiring human decision makers.

This optimization problem can be solved in a brute force manner by considering a discrete grid on the $[0, 1]$ interval and going over all possible choices for the two thresholds θ .

Choosing a Deferral Model: We now describe how we choose a deferral model on the Serengeti dataset. Our goal is to find a model such that the accuracy of the sp classifier is higher than the human h or the AI ensemble m . A large fraction of the camera trap images are false positives, where the camera trap was triggered by a stimulus that was not an actual animal. In order to create a balanced dataset for tuning the deferral model, we subsample these empty images with no animal present. We implement a penalty for




Figure 3: Tradeoff of expected accuracy and deferral rate. More deferral improves performance, but we still gain performance even when deferring less at the circled point.

deferral for the sp classifier, which leads to varied performance at different deferral rates, as seen in Fig. 3. In Fig. 3, individual human accuracy (as opposed to consensus accuracy, which is 1.0) is 0.961, and AI model accuracy is 0.972 (based on choosing one operating point to turn the continuous model scores into binary predictions that maximizes accuracy of the AI-only classifier). We also include the ideal performance if we had a perfect oracle to decide, for each image, whether to defer to a human or rely on the model (given historical human labels). This perfect selective prediction would achieve an accuracy of 0.994.

Given these results, we choose a deferral rate of 1% as an acceptable level of withholding (see Appendix for details) that still improves expected accuracy by a significant margin relative to AI-only or human-only classifiers.

Experiment Design

We hypothesize that the accuracy of a human decision maker in a human-AI team is affected by the SPM in the last step of our workflow (Fig. 1). We specifically consider presenting information about the AI's **prediction** and **deferral status**. The AI prediction refers to the class returned by the AI model, e.g., animal or no animal. The deferral status refers to the result of selective prediction, in which we threshold the model score to determine whether to ask a human to review an image (defer), or rely on the AI prediction. We therefore design a human participant experiment with all possible combinations of these two details: 1) Neither message (NM), 2) Deferral status only (DO), 3) Prediction only (PO), and 4) Both messages (BM), as shown in Fig. 4. We create a survey to host this experiment, consisting of the following sections: 1) Information and consent, 2) Camera trap training and explanation, in which participants are introduced to camera trap images and given an example with an explanation for the best label, 3) Adding AI assistance, in which participants are told about adding an AI model and deferral to assist in the task of sorting camera trap images, 4) AI assist

User would see one of the 4 conditions shown here:

Image 1

Image 40: AI model deferred.

Image 6: AI model predicts no animal present.

Image 37: AI model deferred, but predicts no animal present.



1 2 3 4 5

Definitely no animal present ☐ ☐ ☐ ☐ ☒ Definitely animal present

Figure 4: The four possible SPM conditions in our experiment, along with a challenging example image. The animal is circled in the image here for the reader’s convenience, but in our experiment the circle was not present.

practice and explanations, which consists of 10 examples (9 correct, 1 incorrect) drawn from the Serengeti validation set with AI model assistance, 5) Post training questions asking participants to describe the model strengths and weaknesses, 6) Labeling, and 7) Post dataset questions. Several of these design choices align with the guidelines from Amershi et al. (2019), including describing the AI performance, and providing examples and explanations.

In the labeling section, we display 80 model-deferred images (like Fig. 4) under the four SPM conditions (yielding 20 images per condition, which we believe balances a reasonable number of examples with a manageable amount of participants’ time). Each includes a request for labels of animals present or not, with a Likert scale as in Fig. 4. The images are randomly allocated across the four communication conditions. **We did not inform participants that all images were deferred, we only relied on the different SPM conditions.** Each participant judged images across all four conditions, and no single image was presented more than once to the same participant. The set of 80 images are sampled so as to balance the number of true positive, false positive, true negative, and false negative model classifications, and therefore additionally ensure a balance in the number of images across classes. To ensure that there are no effects due to the specific order or allocation to a condition of each image, four separately seeded random allocations are carried out and each participant is randomly assigned to one of them. To test if the effects of the experimental condi-

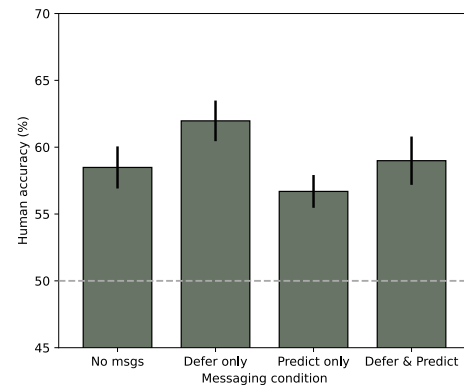


Figure 5: Accuracy of human participants on deferred images, across different SPM conditions. Each bar shows the participant classification accuracies across the entire dataset, errorbars show 95% confidence intervals on the mean. Participants’ responses are more accurate when the images are accompanied by the context that the images are deferred (DO and BM vs. NM and PO). Showing the model’s prediction of the label has a negative effect on accuracy. The horizontal dashed grey line indicates chance performance (50%).

tions on accuracy exceed variation expected by chance, the data are analysed in a 2x2x2 within-subject repeated measures ANOVA with the factors “deferral status” (with the levels: “shown” and “not shown”), “prediction” (with the levels: “shown” and “not shown”) and “model accuracy” (with the levels: “model correct” and “model incorrect”). We received approval from an internal ethics review board, and then recruited 198 participants from Prolific to take part in the experiment. Responses from all 198 participants were included in the ANOVA. Aggregated data are available at https://github.com/deepmind/HAI_selective_prediction/.

Experiment Results and Analysis

SPM Affects Human Accuracy

As can be seen in Fig. 5, the information provided to human participants about the model affected participants’ accuracy. The human-AI communication method that yields the highest human performance is DO. Accuracy in this condition is significantly greater than either humans classifying images by themselves without any information about the model (mean of DO: 61.9%, mean of NM: 58.4%, $p < 0.001$), or the model operating alone (mean of DO: 61.9%, mean of model alone: 50%, $p < 0.001$). Model performance is 50% since the images presented to humans are subsampled from the set of model-deferred images. Furthermore, across all responses, participants are significantly more accurate when the deferral status is shown (conditions DO and BM) than when it is not (conditions PO and NM) (mean of deferral status: 60.4%, mean of no deferral status: 57.4%, $p < 0.001$). We believe this effect of deferral status may be driven by participants inferring that the images are likely to be quite difficult, and therefore concentrating harder. By contrast, partic-




Figure 6: Conformity is the relative increase in agreement when the model predictions are present. For each image, we show how much the agreement between the set of human predictions and model prediction changes from the NM to the PO conditions. Conformity is significantly higher when humans are less confident, indicating that humans are influenced by model predictions more when they are less certain about their judgement.

Participants are significantly less accurate when the prediction is shown (mean of prediction: 57.8%, mean of no prediction: 60.2%, $p = 0.003$). Overall, therefore, we find that providing the deferral status, while avoiding the provision of model predictions, leads to the highest accuracy in human decision-making in this context.

Model Predictions Influence Human Decisions

While the preceding results clearly suggest that participants use the model predictions at least some of the time, the evidence is indirect. We therefore conduct a more direct analysis targeting this question. Specifically, for each image, we compute the proportion of raters who agree with the model prediction under 1) NM and 2) PO. As each image is presented under both conditions, we subtract these two scores to assess how much humans increase their agreement with the model when the model prediction is present. We refer to this as the “conformity” score, and plot it in Fig. 6. Across the set of images, we find an average conformity score of 0.08 (i.e., raters are influenced on 8% of trials), which is significantly greater than zero ($p < 0.0001$), demonstrating that raters use the model information when present. We further test a hypothesis suggested by Bahrami et al. (2010); Boorman et al. (2013), that people are more likely to use model information when they are less confident in their own decisions (as measured by Likert ratings in this case). We separately compute the conformity score for low and high confidence human decisions, and find that, as expected, conformity is higher when rater confidence is low (mean of low confidence: 0.116, mean of high confidence: 0.045, $p = 0.014$).




Figure 7: Accuracy of human participants on deferred images, split by whether the model correctly or incorrectly classified the associated images. (Left) Results split by whether the model correctly classified the image. Notably, images where the model is incorrect has lower participants’ accuracy across all conditions, generally at chance (dashed grey line). Crucially, participants are significantly below chance in the PO condition, for which the model prediction misleads humans. (Right) Human accuracy for each image in the NM and DO conditions, split by whether the model labels the image correctly. Participants’ accuracy is significantly reduced (to around chance) for the subset of images which the model fails to label correctly. Together, both show a congruence between what images the model and humans find difficult to label correctly.

Model Accuracy Affects Human Accuracy

While the preceding analysis demonstrates that people do indeed use the model predictions, this in itself does not explain why we find a decrease in human accuracy when model predictions are available. This suggests a potential bias, such that people tend to use the model information more when it is actually incorrect. To explore this possibility, we focus on the effect of the provided SPM only on the images that the model classifies incorrectly (Fig. 7). We observe that participants’ accuracy in the PO condition is significantly reduced. While participants perform at chance in the other conditions, participants perform significantly below chance in the PO condition (mean of other conditions: 50.6%, mean of PO: 41.9%, $p < 0.001$). In this condition, the label provided by the model is most salient, and critically it provides wrong information. Participants appear to integrate this information as they perform 8.7% worse in the PO condition compared to the NM condition, on the images that are misclassified by the model. In contrast, on images that are correctly classified by the model, and therefore the model can provide a correct prediction message, participants gain 5.1% of accuracy between the NM and PO conditions. This asymmetry, that incorrect model prediction messages are integrated more by the participants than correct model prediction messages, appears to lie at the heart of why we observe an overall negative effect of the prediction message. This pattern is consistent with the observation that participants and the model tend to err on the same images, as shown in the violin plots of Fig. 7. Specifically, participants are correct 66.3% of the time on images shown in the NM and DO conditions and that the model correctly classifies, but only 50.5% on

images in these same conditions that the model classifies incorrectly ($p < 0.05$). Additionally, there is a positive correlation between average human Likert scores in the NM and DO conditions and model scores on the same images, suggesting that humans and models learn similar semantic task dimensions (Pearson's $r = 0.27$, $p = 0.021$).

Asymmetry in Human-AI Agreement

Our results demonstrate a differential change in human accuracy in the presence of model information, based on whether the model is correct or incorrect. When we subdivide the data based on whether the model is correct or incorrect, we see that human accuracy tends to be lower when the model is incorrect (Fig. 7, left). Additionally, exposing humans to incorrect predictions of an AI makes them even more likely to be incorrect. We confirm that human accuracy tends to be lower when the model is incorrect by directly computing the image-wise agreement between human and model ratings on images where the model is correct vs. incorrect (Fig. 7, right). This analysis is performed only on the trials under the NM condition, as we want to investigate independent agreement. As expected, we find that agreement is significantly higher for the correctly labelled images (mean agreement of model correct: 69.6%, mean of model incorrect: 44.9%, $p = 0.007$). This difference in prior agreement over images has potentially important repercussions when introducing model predictions. Specifically, as the subset of images where the model is correct already has high prior agreement with the human raters, there is much less potential for the model to influence human judgments, and hence increase accuracy. By contrast, because the images that are misclassified by the model have a lower prior agreement level with human raters, there is greater potential for model influence, which in this case will decrease accuracy. This asymmetry in potential model influence between correct and incorrect trials is likely to account for the overall drop in human accuracy we see when model predictions are provided to human raters. This is an important result, as it demonstrates that the specific pattern of covariance between model and human decisions can lead to significant downstream effects on joint decision-making when humans have access to the model predictions.

Discussion

To summarize, we find that there are significant effects in the way deferred images are presented to humans in a selective prediction workflow. In particular, in this context, presenting deferral status is helpful, while presenting the uncertain prediction, even when accompanied by a deferral status, can be harmful, especially when the model is incorrect. From this, it is clear that performance will not necessarily be what is estimated from historical human labels. Together, these findings illustrate the importance of considering the human-AI team while designing selective prediction systems, as the SPM can have a significant impact on performance.

We believe these are important findings to direct future research, and have several suggestions for open questions to explore. First, this experiment could be expanded. While we

chose to leave this static, it would be helpful to determine the amount and type of training that is most useful for participants. We also chose to use two categories for deferral status, either defer or not defer. However, it is possible that finer-grained information about model uncertainty could be helpful (Bhatt et al. 2021). We additionally focused on understanding why prediction hurt, but encourage collecting further information, such as timing, to better understand why deferral status helped. This may help inform further research into designing selective prediction algorithms based on human-AI teams, for example by exploring bounded rationality for training improved selective prediction models.

There are also questions about generalizability of these results. These specific results (i.e., that deferral status helps while prediction hurts performance) are not likely to be robust across datasets, different human-AI use scenarios (e.g., decision aids), or even participant expertise levels. For example, in this study, we asked two domain experts working with the Serengeti dataset to go through the same survey provided to Prolific participants. We similarly find that each individual has different performance in the four conditions, and that the deferral status leads to improved performance in both cases. However, the interactions are slightly different. It is necessary to search for generalizable trends across these cases in future work.

Finally, though we worked with humans in this study, it is extremely important to consider specific deployment challenges in these contexts, such as how selective prediction may change existing processes, e.g., in healthcare (Yang, Steinfeld, and Zimmerman 2019), or how to best modify the workflow and instructions in the case where there are multiple human experts we could rely on. In all cases, we stress the importance of remembering the human component of human-AI teams.

Acknowledgments

We would like to thank the Prolific participants, Exavery Kigosi from the Grumeti Fund, DeepMind HuBREC, and the reviewers for their helpful feedback.

Ethics Statement

We received approval for this work from DeepMind's ethics review board, HuBREC 21 008. We did not collect personally identifiable information, and we paid a flat rate of £7 for the task, which took about 20-30 minutes. We believe it's unlikely that any of the participants were exposed to inappropriate imagery during the study. However, the camera trap datasets are made up of animals, both adult and young, in the wild, and thus could potentially include scenes of predator-prey interactions that some raters may find distressing. We included warnings about the species, behaviors, and possible human presence in the information sheet. Furthermore, we ensured participants understood that they were under no obligation to consent to participate, and that if they declined to participate, there would be no negative repercussions.

References

- Amershi, S.; Weld, D.; Vorvoreanu, M.; Fourney, A.; Nushi, B.; Collisson, P.; Suh, J.; Iqbal, S.; Bennett, P. N.; Inkpen, K.; et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*, 1–13.
- Bahrami, B.; Olsen, K.; Latham, P. E.; Roepstorff, A.; Rees, G.; and Frith, C. D. 2010. Optimally interacting minds. *Science*, 329(5995): 1081–1085.
- Bansal, G.; Nushi, B.; Kamar, E.; Horvitz, E.; and Weld, D. S. 2021. Is the Most Accurate AI the Best Teammate? Optimizing AI for Teamwork. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 11405–11414.
- Beery, S.; Wu, G.; Rathod, V.; Votel, R.; and Huang, J. 2020. Context r-cnn: Long term temporal context for per-camera object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13075–13085.
- Bhatt, U.; Antorán, J.; Zhang, Y.; Liao, Q. V.; Sattigeri, P.; Fogliato, R.; Melançon, G.; Krishnan, R.; Stanley, J.; Tickoo, O.; et al. 2021. Uncertainty as a form of transparency: Measuring, communicating, and using uncertainty. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 401–413.
- Boorman, E.; O’Doherty, J.; Adolphs, R.; and Rangel, A. 2013. The Behavioral and Neural Mechanisms Underlying the Tracking of Expertise. *Neuron*, 80(6): 1558–1571.
- Buçinca, Z.; Malaya, M. B.; and Gajos, K. Z. 2021. To trust or to think: cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1): 1–21.
- Chow, C. 1970. On optimum recognition error and reject tradeoff. *IEEE Transactions on information theory*, 16(1): 41–46.
- De-Arteaga, M.; Fogliato, R.; and Chouldechova, A. 2020. A case for humans-in-the-loop: Decisions in the presence of erroneous algorithmic scores. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–12.
- Dietvorst, B. J.; Simmons, J. P.; and Massey, C. 2015. AI-algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1): 114.
- Galesic, M.; and Garcia-Retamero, R. 2010. Statistical numeracy for health: a cross-cultural comparison with probabilistic national samples. *Archives of internal medicine*, 170(5): 462–468.
- Gaube, S.; Suresh, H.; Raue, M.; Merritt, A.; Berkowitz, S. J.; Lermer, E.; Coughlin, J. F.; Gutttag, J. V.; Colak, E.; and Ghassemi, M. 2021. Do as AI say: susceptibility in deployment of clinical decision-aids. *NPJ digital medicine*, 4(1): 1–8.
- Geifman, Y.; and El-Yaniv, R. 2017. Selective classification for deep neural networks. *arXiv preprint arXiv:1705.08500*.
- Green, B.; and Chen, Y. 2019. The principles and limits of algorithm-in-the-loop decision making. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW): 1–24.
- Grote, T. 2021. Trustworthy medical AI systems need to know when they don’t know. *Journal of Medical Ethics*, 47(5): 337–338.
- Keswani, V.; Lease, M.; and Kenthapadi, K. 2021. Towards Unbiased and Accurate Deferral to Multiple Experts. *arXiv preprint arXiv:2102.13004*.
- Kompa, B.; Snoek, J.; and Beam, A. L. 2021. Second opinion needed: communicating uncertainty in medical machine learning. *NPJ Digital Medicine*, 4(1): 1–6.
- Kool, W.; and Botvinick, M. 2018. Mental labour. *Nature Human Behaviour*, 2(5995): 899–908.
- Logg, J. M.; Minson, J. A.; and Moore, D. A. 2019. AI-algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151: 90–103.
- Madras, D.; Pitassi, T.; and Zemel, R. 2017. Predict responsibly: improving fairness and accuracy by learning to defer. *arXiv preprint arXiv:1711.06664*.
- Mozannar, H.; and Sontag, D. 2020. Consistent estimators for learning to defer to an expert. In *International Conference on Machine Learning*, 7076–7087. PMLR.
- Norouzzadeh, M. S.; Morris, D.; Beery, S.; Joshi, N.; Jojic, N.; and Clune, J. 2021. A deep active learning system for species identification and counting in camera trap images. *Methods in Ecology and Evolution*, 12(1): 150–161.
- Norouzzadeh, M. S.; Nguyen, A.; Kosmala, M.; Swanson, A.; Palmer, M. S.; Packer, C.; and Clune, J. 2018. Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proceedings of the National Academy of Sciences*, 115(25): E5716–E5725.
- Rastogi, C.; Zhang, Y.; Wei, D.; Varshney, K. R.; Dhurandhar, A.; and Tomsett, R. 2022. Deciding Fast and Slow: The Role of Cognitive Biases in AI-assisted Decision-making. In *CSCW*.
- Swanson, A.; Kosmala, M.; Lintott, C.; Simpson, R.; Smith, A.; and Packer, C. 2015. Snapshot Serengeti, high-frequency annotated camera trap images of 40 mammalian species in an African savanna. *Scientific data*, 2(1): 1–14.
- Tabak, M. A.; Norouzzadeh, M. S.; Wolfson, D. W.; Sweeney, S. J.; VerCauteren, K. C.; Snow, N. P.; Halseth, J. M.; Di Salvo, P. A.; Lewis, J. S.; White, M. D.; et al. 2019. Machine learning to classify animal species in camera trap images: Applications in ecology. *Methods in Ecology and Evolution*, 10(4): 585–590.
- Tversky, A.; and Kahneman, D. 1981. The framing of decisions and the psychology of choice. *Science*, 211(4481): 453–458.
- Wiener, Y.; and El-Yaniv, R. 2013. *Theoretical foundations of selective prediction*. Ph.D. thesis, Computer Science Department, Technion.
- Wilder, B.; Horvitz, E.; and Kamar, E. 2020. Learning to complement humans. *arXiv preprint arXiv:2005.00582*.

- Willi, M.; Pitman, R. T.; Cardoso, A. W.; Locke, C.; Swanson, A.; Boyer, A.; Veldthuis, M.; and Fortson, L. 2019. Identifying animal species in camera trap images using deep learning and citizen science. *Methods in Ecology and Evolution*, 10(1): 80–91.
- Yang, Q.; Steinfeld, A.; and Zimmerman, J. 2019. Unremarkable ai: Fitting intelligent decision support into critical, clinical decision-making processes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–11.
- Zhang, H.; and Maloney, L. T. 2012. Ubiquitous log odds: a common representation of probability and frequency distortion in perception, action, and cognition. *Frontiers in neuroscience*, 6: 1.