

Human-Algorithm Collaboration: Achieving Complementarity and Avoiding Unfairness

Kate Donahue*
kdonahue@cs.cornell.edu
Cornell University
USA

Alexandra Chouldechova
Carnegie Mellon University and
Amazon Web Services
USA

Krishnaram Kenthapadi†
Fiddler AI
USA

ABSTRACT

Much of machine learning research focuses on predictive accuracy: given a task, create a machine learning model (or algorithm) that maximizes accuracy. In many settings, however, the final prediction or decision of a system is under the control of a human, who uses an algorithm's output along with their own personal expertise in order to produce a combined prediction. One ultimate goal of such collaborative systems is *complementarity*: that is, to produce lower loss (equivalently, greater payoff or utility) than either the human or algorithm alone. However, experimental results have shown that even in carefully-designed systems, complementary performance can be elusive. Our work provides three key contributions. First, we provide a theoretical framework for modeling simple human-algorithm systems and demonstrate that multiple prior analyses can be expressed within it. Next, we use this model to prove conditions where complementarity is impossible, and give constructive examples of where complementarity is achievable. Finally, we discuss the implications of our findings, especially with respect to the fairness of a classifier. In sum, these results deepen our understanding of key factors influencing the combined performance of human-algorithm systems, giving insight into how algorithmic tools can best be designed for collaborative environments.

ACM Reference Format:

Kate Donahue, Alexandra Chouldechova, and Krishnaram Kenthapadi. 2022. Human-Algorithm Collaboration: Achieving Complementarity and Avoiding Unfairness. In *FAccT '22: ACM Conference on Fairness, Accountability and Transparency*, June 21–24, 2022, Seoul, South Korea. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3531146.3533221>

1 INTRODUCTION

Consider a prediction task where the goal is to take a set of features about the world as input and predict an outcome of interest. A typical machine learning approach to such a task is to attempt to select a model with low (generalization) loss for the problem at

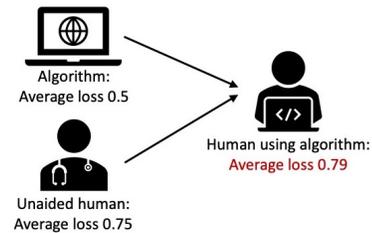
*Work conducted while on an internship at Amazon.

†Work done while at Amazon.

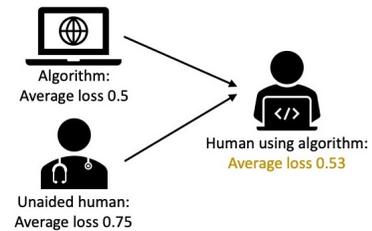
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22), Seoul, Republic of Korea,

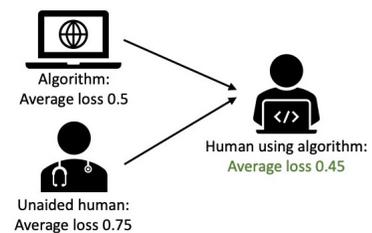
© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-9352-2/22/06...\$15.00
<https://doi.org/10.1145/3531146.3533221>



(a) Scenario 1: Combined system has higher loss than either unaided human or algorithm.



(b) Scenario 2: Combined system has lower loss than unaided human, but higher loss than algorithm.



(c) Scenario 3: Complementary performance: lower error than either unaided human or algorithm.

Figure 1: Three possible scenarios for human-algorithm collaboration, each with the same algorithmic and unaided human loss. However, the loss of the combined system (human using the algorithm) might vary substantially. Section 4 gives a more detailed analysis.

hand. If such a model is applied directly to the prediction task, it will minimize expected loss.

However, this standard approach does not necessarily reflect the way that machine learning tools are actually implemented. Often, algorithmic predictions are presented to humans, who then make a final decision by additionally relying on their own expertise

[1, 17, 35, 38]. For example, consider a doctor looking at a medical record and trying to make a determination of whether disease is present. An algorithmic prediction based on the record may be useful, but it almost certainly will not be the sole factor influencing the doctor's diagnosis. For example, the doctor may have access to different data, such as conversations with the patient. The doctor may also have access to different knowledge, such as distilled expertise from years of practice. The doctor's decision will be a function of the algorithm's prediction, as well as their own inherent belief. Note that the doctor's decision-making may be imperfect, such as relying on their own judgement when the algorithm may have better performance. A successful outcome occurs when the combined system (the doctor using algorithmic output) has low loss, not when the algorithm alone has low loss. Figure 1 illustrates three scenarios where a combined human-algorithm system could have differing levels of loss.

In particular, one especially valuable goal is *complementarity* (or *complementary performance*). Complementarity (originally defined in Bansal et al. [2]) is achieved whenever the combined human-algorithm system has strictly lower expected loss than either the human or the algorithm alone (Figure 1c). Complementarity is not necessary for a combined system to be deemed successful: for example, a combined system that does better than the human alone, but not necessarily better than the algorithm alone, would still reflect an improvement from a human-alone status-quo. However, complementarity creates the strongest incentive for adoption of a combined human-algorithm system, which is why it is the focus of our analysis.

Contributions: At a high level, we address the following problems: (i) *How do we formally and tractably model human-algorithm collaborative systems?* (ii) *When can human-algorithm collaborative systems produce higher accuracy than either the human or algorithm alone?* (iii) *What are the fairness implications of such collaborative systems?*

The contributions of this work are three-fold. First, in Section 3, we introduce a simple theoretical framework for analyzing human-algorithm collaboration, and demonstrate the richness of this framework by showing that it can encapsulate models from previous works analyzing human decision-making. In Section 4, we provide a simple, concrete motivating example using this framework that illustrates the core results of this paper.

Next, in Section 5, we use this approach to analyze complementarity. First, we present several impossibility results that characterize regimes in which human-algorithm collaboration can never achieve complementarity. We then give concrete conditions for when complementarity can be achieved. In particular, our results suggest that complementarity is easier to achieve when loss rates are highly variable: when the unaided human (or algorithm) has very low loss on some inputs and very high loss on inputs. Disparate levels of loss raises issues of fairness, which we turn to next.

In Section 6 we conclude our analysis by examining the fairness impacts of complementarity. The variability in loss rates implied by our results has implications for fairness, since types of inputs with very high error rates may correspond to protected attributes, such as race, gender, or ethnicity. To investigate this concern, we propose and analyze three types of fairness relating to human-algorithm systems, giving conditions for when they can and cannot

be achieved. One of main results shows that when complementarity is achieved, at least one group does worse in the combined system than under the human-only status quo. Additionally, we give a simple condition where the combined human-algorithm system will guarantee that loss disparity between different protected groups will not increase.

2 RELATED WORK

2.1 Human-Algorithm Collaboration

A series of papers have explored issues related to human-algorithm collaboration. For example, Poursabzi-Sangdeh et al. [28] and Yin et al. [38] analyze how explainability and accuracy, respectively, influence how humans use algorithmic predictions. Similarly, Dietvorst and Bharti [8] hypothesize that humans may prefer algorithmic predictions that are more variable in their loss rates and Dietvorst et al. [9] suggests that allowing algorithmic predictions to be modified may make humans more likely to use them.

Other papers center more on in-depth qualitative assessments of how professionals incorporate tailor-made algorithmic tools into their workflow. For example, Lebovitz et al. [21, 22] studies how doctors in major US hospitals use AI predictions in their daily work. Similarly, Okolo et al. [27] studies how community healthcare workers in India believe AI tools could influence their work. Finally, Yang et al. [37] studies how UX designers work with machine learning tools and the data scientists who create them.

Some research teams that develop tools for human-in-the-loop settings have run experiments analyzing how their tools perform with human collaboration. For example, Beede et al. [3], Raghu et al. [29] both study how an AI tool for predicting diabetic retinopathy fits in with a broader ecosystem (human doctors and the overall healthcare system). Similarly, De-Arteaga et al. [6] studies how child welfare call screeners incorporate algorithmic predictions in their risk assessments. Tan et al. [33] studies how human and algorithmic distributions of loss rates differ for recidivism predictions for the COMPAS dataset (but not in ways that allowed for complementary performance by combined systems). Similarly, Geirhos et al. [13] compares the similarity (consistency) of loss in predictions made by humans and a deep learning algorithm.

Some computer science papers specifically analyze models of human-algorithm interaction, such as [1, 4, 17, 30, 31, 35]. Of these, Bansal et al. [2] is especially relevant because it is framed through the goal of complementarity. Bansal et al. [1] also highlights the fact that the optimizing for the algorithm's error may not minimize the loss of the combined system. In Section 3, we show how human decision-making processes rules inspired by analyses Bansal et al. [1] and Vodrahalli et al. [35] can be represented in our model. Some papers show how to build models optimized for a human-algorithm *deferral* system, where the final decision is made by either the human or the algorithm [5, 25, 26]. Straitouri et al. [32] studies a variant of this problem for classification where the algorithm presents a subset of possible labels to the human, who selects the final decision from among them.

2.2 Fairness in Human-Algorithm Collaboration

Some papers specifically consider the fairness implications of combined human-algorithm systems. For example, Madras et al. [23] studies fairness and accuracy in deferring to a human expert, while Keswani et al. [18, 19] extends this analysis to deferring to multiple different human experts. For example, Gillis et al. [14] takes a theoretical approach towards modeling the human-algorithm system and gives conditions where adding a biased (unfair) human can change the fairness properties of the overall system. Valera et al. [34] studies a system with multiple biased “experts”, where assigning the correct expert to each task can improve accuracy while still satisfying fairness requirements.

2.3 Related Papers From Other Areas

Finally, some papers in seemingly unrelated areas end up being relevant to our analysis. For example, ensemble learning studies how to incorporate predictions from multiple algorithms into a unified (more accurate) system [20]. Ensemble learning differs from our analysis in that each expert (predictor) is assumed to be an algorithm, and predictions are assumed to be combined by some additional algorithm under our control (rather than a human decision-maker we cannot control). However, certain factors identified in the ensemble learning literature as affecting overall performance, such as diversity, are relevant for our analysis [7]. Additionally, multiple works study how fairness properties of predictors change when they are composed [10, 11, 36]. These works are relevant to our analysis in Section 6 of the fairness of a combined human-algorithm system, but differ somewhat from ours: in general, these other papers tend to study fairness of allocating or achieving some desired prediction, while our analysis describes fairness as equal loss across groups. Finally, Meehl [24] compares statistical and clinical methods of reasoning, a framing that parallels to our analysis of algorithmic versus human prediction methods.

3 MODEL AND ASSUMPTIONS

3.1 Model

Our model considers a prediction task: given some element $x \in \mathcal{X}$, make a prediction $y \in \mathcal{Y}$ that minimizes some loss function \mathcal{L} , with loss bounded ≥ 0 . This loss could reflect any error rates for any type of learning problem—for example, regression and classification tasks could both be represented by this loss function. We model the input space \mathcal{X} as being made up of N discrete regimes: all inputs within the same regime are identical from the perspective of algorithmic and human loss. This is without loss of generality, given that N could be arbitrarily large. We are *not* assuming that either the human or algorithm has knowledge of these regimes, simply that they exist. We will denote the probability of seeing regime i is given by p_i , with $\sum_{i=1}^N p_i = 1$.

The human-algorithm system consists of three components:

- (1) An *algorithm*, which for each regime in the input space $x_i \in \mathcal{X}$ makes a prediction \hat{y}_i^a with some loss rate a_i . The average loss is given by $\sum_{i=1}^N p_i \cdot a_i = A$. We can write $a_i = A + \delta_{ai}$, with $\sum_{i=1}^N p_i \cdot \delta_{ai} = 0$. The term δ_{ai} represents how much a_i varies (differs from the average loss A).

- (2) A *unaided human*, which similarly for each regime in the input space $x_i \in \mathcal{X}$ makes some prediction \hat{y}_i^h . The average loss of the human is given by $\sum_{i=1}^N p_i \cdot h_i = H$. Similarly, we write $h_i = H + \delta_{hi}$, with $\sum_{i=1}^N p_i \cdot \delta_{hi} = 0$.
- (3) Finally, some *combiner* (a human using algorithmic input) $g(\hat{y}_i^a, \hat{y}_i^h)$, which takes predictions given by the algorithm and unaided human and returns a combined prediction, \hat{y}_i^c . The combining function reflects human decision-making: it could select the algorithm’s prediction, the unaided human’s prediction, or interpolate between the two of them. We could also view this as a (loss) *combining function* $c(a_i, h_i)$ that takes the algorithmic loss and human loss on a particular instance and returns some combined loss.

In general, we may not have control over all (or even any) of these components. For example, the combining function reflects human judgement, which typically can’t be directly manipulated. A primary goal of our analyses is to determine when a human-algorithm system displays complementarity, defined in Definition 1 below.

DEFINITION 1 (FROM BANSAL ET AL. [2]). *A human-algorithm system displays **complementary performance** when the combined system has (strictly) lower loss than either the human or algorithm:*

$$\sum_{i=1}^N p_i \cdot c(a_i, h_i) < \min \left(\sum_{i=1}^N p_i \cdot a_i, \sum_{i=1}^N p_i \cdot h_i \right) = \min(A, H)$$

3.2 Assumptions

The combining function models the key question in human-algorithm collaboration: how do humans incorporate algorithmic predictions with their own expertise? In this work, we will make two main assumptions about how such combination occurs. First, throughout this paper, we will find it useful to work in the space of combining *losses*, rather than combining predictions. Specifically, we will use the $c(a_i, h_i)$ loss combining function. Assumption 1 describes the assumption this implies.

ASSUMPTION 1. *The loss of a combined human-algorithm system can be modeled by a combining rule relying only on the loss rates of the unaided human and algorithm in a particular regime: $c(a_i, h_i)$. That is, regimes with identical (unaided human, algorithm) pairs of loss rates are treated identically.*

This assumption reflects the case where the level of accuracy in the algorithm and (unaided) human is the only feature influencing the accuracy of the combined system. An example of a situation that might *violate* this assumption is if regime 1 and 2 both have loss of 3% for the unaided human and 5% for the algorithm, but the human using the algorithm (combined system) has loss of 3% for regime 1 and 4% for regime 2. Considering the case where this assumption is relaxed could be an interesting avenue for future work: however, it would likely result in much more complicated analysis.

Next, Assumption 2 below, describes the assumption that the combining rule’s outputs are bounded.

ASSUMPTION 2. For each regime, the loss of the combined system is bounded between the loss of the human and algorithm:

$$\min(a_i, h_i) \leq c(a_i, h_i) \leq \max(a_i, h_i)$$

This assumption reflects a case where the combiner operates by interpolating between the predictions made by the human or algorithm. An example of a situation that might *violate* this assumption is if the combined system has loss 3% in a certain regime, where the human has loss 4% and the algorithm has loss 6% in that regime. A bounded combining rule makes modeling human-algorithm collaboration more realistic: complementarity is trivial to achieve if the combining rule's loss can be arbitrarily disconnected from the loss of the human and algorithm.

It's worth considering why a system may satisfy bounded inputs (Assumption 2) and still exhibit complementarity. Assumption 2 refers to bounds at the level of each *regime*, while complementarity refers to *overall average loss*. For example, Scenario 3 (Table 3) in Section 4 obeys the bound in Assumption 2, and yet also achieves complementarity.

3.3 Weighting function

In this work, we will find it helpful to think about the combined human-algorithm system as involving a *weighting function* $0 \leq w_h(a_i, h_i) \leq 1$ controlling how much the human influences the final prediction:

$$c(a_i, h_i) = (1 - w_h(a_i, h_i)) \cdot a_i + w_h(a_i, h_i) \cdot h_i \quad (1)$$

Lemma 1, below, shows that using a weighting function requires no new assumptions.

LEMMA 1. Any combining rule relying only on loss rates (Assumption 1) with bounded output (Assumption 2) can be written as a combining rule with a weighting function $0 \leq w_h(a_i, h_i) \leq 1$.

One simple (ideal) combining function is given by Example 1: it simply selects whichever of the unaided human or algorithm has lower loss. While this is the best possible combining function (given our assumptions), it is likely not a realistic model of how human decision-makers incorporate algorithmic advice.

EXAMPLE 1 (MIN). The combining function becomes $c(a, h) = \min(a, h)$ is represented by the weighting function:

$$w_h(a, h) = \begin{cases} 1 & h \leq a \\ 0 & \text{otherwise} \end{cases}$$

While our framework is simple, it is also sufficiently flexible to capture models of human-algorithm collaboration studied in multiple previous papers (described in greater detail in Appendix A). Examples 2 and 3 demonstrate this in reference two particular models suggested by prior literature. First, Example 2 selects whichever of the human or algorithm has lower loss rate with probability p_s . For high p_s , this reflects a decision-maker who accurately trusts whichever has lower loss.

EXAMPLE 2 (BANSAL ET AL. [1]). The analysis in Bansal et al. [1] suggests the weighting function:

$$w_h(a, h) = \begin{cases} p_s & h \leq a \\ 1 - p_s & \text{otherwise} \end{cases}$$

Next, in Example 3, the decision-maker first decides whether to consider algorithmic advice at all: it does so only if the loss rate is ϵ lower than the human loss rate. Then, the decision-maker incorporates algorithmic advice with some probability $p_s(\cdot)$ that is a function depending on the gap between human and algorithmic loss rates.

EXAMPLE 3 (VODRAHALLI ET AL. [35]). The two-stage model in Vodrahalli et al. [35] could be written as:

$$w_h(a, h) = \begin{cases} 1 & a \geq h - \epsilon \\ p_s(h - a) & \text{otherwise} \end{cases}$$

3.4 Research Ethics and Social Impact

While our paper is primarily theoretical, its application area prompts a number of ethical considerations. For example, in this work, we are primarily concerned with building better prediction functions. In general, such functions could be used for positive means (helping a doctor correctly diagnose a disease) or negative ones (enabling the identification and repression of minority groups). Additionally, even if a function is being used for positive goals, it could be the case that factors besides the loss rate are ultimately more important. For example, it could be that the process of coming to a prediction, rather than the prediction itself, is more important. This is especially salient for our discussion of fairness, which assumes that the fairness of outcomes (of loss rate disparities) is the relevant factor to consider, rather than fairness of the prediction process. The issue of explanation of algorithmic predictions, which is orthogonal to our main analysis, could be relevant for this consideration.

4 MOTIVATING EXAMPLE

To further motivate the analysis, let us revisit the medical application from the introduction (Figure 1) and introduce further specifics. Consider the medical prediction task of using information from a patient's medical record to predict disease severity (on a scale from 0 to 5, as in [29]). As illustrated in Figure 1, we will assume that doctors relying on their medical training (unaided humans) have an average loss rate of 0.75: they are off by 0.75 grades, on average. A data science team has created a machine learning algorithm that has average loss of 0.5.

Even though the algorithm has lower loss, doctors won't simply rubber-stamp algorithmic suggestions. Because they have specialized training and access to additional information (such as conversations with patients), a doctor might reasonably incorporate algorithm advice only partially, or only sometimes. However, this leaves open a crucial question: what is the combined human-algorithm loss? That is, what is the average loss once doctors start incorporating the new machine-learning algorithm into their decision-making process?

4.1 Three scenarios

To help build intuition before delving into our formal theoretical results, we will consider three different example scenarios for how a human-algorithm collaborative system might look like, given in Tables 1, 2, 3. Each scenario has the same average loss for the unaided human and for the algorithm (0.75 and 0.5, respectively). However, each scenario differs in 1) the way (unaided) human and

	(Unaided) human	Algorithm	Combined (human using algorithm)	Weight (on unaided human)
Regime 1	1	0.35	0.94	0.9
Regime 2	0.5	0.65	0.64	0.1
Average	0.75	0.5	0.79	0.5

Table 1: Scenario 1 Loss rates: An example of a combined human-algorithm system. There are two regimes, each making up equal proportions of the input space ($p_0 = p_1 = 0.5$). Note here that complementarity is *not* satisfied: in fact, the combined system has higher loss than either the human or algorithm alone!

	(Unaided) human	Algorithm	Combined (human using algorithm)	Weight (on unaided human)
Regime 1	1	0.35	0.51	0.25
Regime 2	0.5	0.65	0.54	0.75
Average	0.75	0.5	0.53	0.5

Table 2: Scenario 2 Loss rates: A second example of a combined human-algorithm system, but with different loss distributions. Here, the combined system (human using algorithm) has average loss which is lower than the loss of the unaided human, but higher than the loss of the algorithm alone.

	(Unaided) human	Algorithm	Combined (human using algorithm)	Weight (on unaided human)
Regime 1	1.15	0.2	0.44	0.25
Regime 2	0.35	0.8	0.46	0.75
Average	0.75	0.5	0.45	0.5

Table 3: Scenario 3 Loss rates: A third example of a combined human-algorithm system. Here, the combined system (human using algorithm) displays *complementary performance*: its average loss of 0.45 is lower than the loss of either the unaided human (0.75) or algorithm alone (0.5).

Table 4: Three possible scenarios for human-algorithm collaboration. In each, the algorithm and unaided human have the same average loss. However, the loss of the human using the algorithm varies. Figure 1 gives a visual description of these scenarios.

algorithm loss is distributed and 2) the way the human combines algorithmic advice. Specifically, these three scenarios illustrate the simplified case where patient records (“regimes”) come in one of two types: regime 1 and regime 2, each of which makes up 50% of the total input space. These regimes might differ in multiple ways: say disease progression, data quality, or patient characteristics. We will treat observations from within the same regime as identical: the doctor or algorithm or combined system each has uniform loss for every record within the same regime. (In later sections, we will relax this consider arbitrary numbers of regimes, reflecting arbitrarily complex distributions of loss).

Scenario 1 (combined loss higher than unaided human or algorithm): For example, in Table 1, the unaided human has loss 1 on instances of type 1, but loss 0.5 for instances of type 2. The algorithm’s loss rate distribution differs: 0.35 for type 1, and 0.65 for type 2. Finally, the combined loss (loss of the human using algorithmic input) for a particular regime is a function of the loss of the unaided human and the algorithm: for this example, it’s 0.94 in regime 1, and 0.64 in regime 2. Note that this results in an average loss of 0.79—the human using the algorithm has a strictly *greater* loss rate than either the unaided human or the algorithm! This unfortunate case could result from inappropriately relying on the algorithm: for example, the doctor might mistakenly incorporate algorithmic advice more frequently in regimes where it happens to have higher loss. The fourth column concretely reflects this cause: it calculates the weighting function for each regime (reliance on the unaided human, as opposed to the algorithm). In this case, the regime 1 weighting function is 0.9, meaning the combined system

is relying heavily on the unaided human for this instance, even though it has higher loss than the algorithm. Similarly, in regime 2, the weighting function is 0.1, indicating that the combined system is relying more on the algorithm, even though (for this regime), the algorithm has higher loss than the human. This inappropriate reliance explains why the combined system has higher loss than either the unaided human or algorithm alone.

Scenario 2 (combined system lower error than human, higher error than algorithm): Table 2 presents a slightly more optimistic Scenario 2. Here, *the distribution of loss rates for the unaided human and the algorithm are the same as in Scenario 1, but the combined loss differs*: the average loss of the human using the algorithm is 0.53—lower than the loss of the unaided human (0.75), but higher than the loss of the algorithm alone (0.5). This reflects a common scenario in human-algorithm collaboration: the combined system ends up improving over the human alone, but still falls short of the loss rate achievable by the algorithm alone [2]. In this case, the reason is because *the combined system is doing a better job than in Scenario 1 of appropriately relying on the unaided human or algorithm*. Note that the weighting function is lower in regime 1 (when the unaided human has higher loss) and higher in regime 2 (where the unaided human has lower loss). This more appropriate reliance explains why the combined system has better performance than Scenario 1.

Scenario 3 (complementarity: lower than unaided human or algorithm): Finally, Table 3 presents Scenario 3. In this case, *we keep the weighting function the same as in Scenario 2*, as well as the average loss of the unaided human and the algorithm. However,

we change the distribution of loss rates across the regimes: we make them more variable. For example, the unaided human now has loss 1.15 in regime 1 and 0.35 in regime 2, while the algorithm has loss 0.2 in regime 1 and 0.8 in regime 2. While average (unaided) human and algorithmic loss are the same as in Scenarios 1 and 2, the combined human-algorithmic system ends up having average loss of 0.45, which is strictly lower than either the unaided human or algorithm alone. Complementarity arises here because diversity of errors and appropriate reliance when algorithmic error is lower than the human's. As we show in Section 5, these conditions are necessary for complementarity to arise.

Let's return to our motivating example: a doctor using algorithmic input to make predictions about patients. These three examples illustrate different possible overall loss rates from a doctor and algorithm with the same average loss rates alone. It is clearly important for overall performance (including patient well-being) to determine whether the human-algorithm system will result in something like Scenario 1 (with higher loss than either the doctor or algorithm alone) or Scenario 3 (with lower loss than either alone). In the coming sections we lay out a theoretical framework for analyzing combined human-algorithm systems more generally to understand their implications for fairness and complementary performance. Our analysis formalizes the observations made in the scenarios above by precisely characterizing the role that "appropriate" reliance on the algorithm and variability in performance across regimes plays in complementarity and fairness.

4.2 Complementarity and fairness

Complementarity is the best-case scenario for human-algorithm collaboration, and much of our paper will revolve around proving when it can and cannot exist. For instance, in Section 5 we will give theoretical results describing the kinds of factors necessary for complementarity to be achievable. In particular, we will give conditions on the distributions of loss rates as well as the way predictions are combined. We will show that, all else being equal, complementarity is easier to achieve when distributions of loss for the algorithm and unaided human are highly variable. For example, the algorithm's loss rates are less variable in Tables 1 and 2 (ranging from 0.35 to 0.65) than they are in Table 3 (where they go from 0.2 to 0.8). However, variable loss rates naturally have fairness implications.

Fairness concerns are especially salient if the regimes are correlated with sensitive attributes, such as race, ethnicity, sex, gender, or socioeconomic status. One fairness question could revolve around the "loss disparity" - the difference in loss rates between multiple regimes. Many papers in algorithmic fairness focus on loss disparity (also called accuracy parity or disparate mistreatment [12, 39]), making them a natural focus in our work. In Table 3, there's a loss disparity (difference in loss rates between regimes) of 0.8 for the unaided human and 0.6 for the algorithm. However, the combined system has an loss disparity of 0.02—much lower than either the human or algorithm! Can we guarantee that combined systems will always have lower loss disparities? In Section 6 we will show that, under some conditions, complementarity implies a bounded loss disparity. Another fairness concern could revolve around whether the benefits of incorporating an algorithm are shared among all

groups. For example, in Table 3, regime 2 sees a reduction in loss when the algorithm is incorporated, going from 1.15 (unaided human) to 0.44 (combined human with algorithm). However, regime 1 sees an increase in loss, from 0.35 to 0.46. Ideally, a combined system should benefit all regimes—but (when) is this possible? In Section 6, we will show that, unfortunately, any system exhibiting complementarity can't be one where all regimes see their loss decrease from what it was with the unaided human.

5 COMPLEMENTARITY

In this section, we analyze complementarity: when will a combined human-algorithm system have lower average loss than either the unaided human or algorithm? First, we give general results for when complementarity is impossible to achieve. Secondly, we build on these previous results in order to give constructive examples where complementarity is possible. Finally, we discuss some implications of our findings. All proofs are given in Appendix B.

5.1 Cases Where Complementarity is Impossible

This section gives cases where complementarity is *impossible* to achieve. These results help to narrow the scope of cases that we must consider for future analysis. They could also be helpful for practitioners: if any of their cases is addressed by these lemmas, then they can immediately know that their system can never achieve complementarity.

We begin by presenting two lemmas that concern the distribution of loss rates for the unaided human or the algorithm. Lemma 2 considers a case where the loss of the algorithm and unaided human are constant across regimes. Constant loss is of course unlikely to arise in practice, but is often a setting considered in the computational science literature.

LEMMA 2. *A human-algorithm system where unaided human and algorithm loss rates are constant over regimes can never achieve complementary performance.*

The next result, Lemma 3, considers the setting where one of the components outperforms the other in every regime. Specifically, it says that if one of the unaided human or algorithm always has lower loss than the other, then complementarity is impossible.

LEMMA 3. *Complementarity is impossible if one of the human or algorithm always weakly dominates the loss of the other: that is, if $a_i \leq h_i$ for all i , or $a_i \geq h_i$ for all i .*

This result may have implications for tasks where the algorithm has extremely high performance, achieving lower loss than the human for all types of inputs. While this could mean that the combined system will have lower loss than the human alone, Lemma 3 tells us that it can't achieve lower loss than both the human and algorithm.

Next, we will consider two lemmas that concern properties of the combining rules (the way the human incorporates advice from the unaided human and the algorithm). First, Lemma 4 analyzes a case where the combining function is convex in its arguments.

LEMMA 4. *A combining function $c(a_i, h_i)$ that is convex in a_i, h_i can never achieve complementary performance.*

Recalling that the maximum function, $\max(a, h)$, that returns whichever of the unaided human or algorithm has higher loss, is convex, this result is very intuitive.

A simple but important corollary (Corollary 1) is that complementarity is impossible whenever the weighting function $w_h(a_i, h_i)$ is constant (independent of the algorithm or human’s loss rate). This might reflect a situation where the decision-maker either is ignorant of the loss rates by the human or algorithm, or else decides to ignore them.

COROLLARY 1. *A combining function with a constant weighting function $w_h(a_i, h_i) = w_h$ can never achieve complementarity performance.*

PROOF. Note that $c(a, h) = w_h \cdot h + (1 - w_h) \cdot a$ is convex in both a and h . \square

These results, taken together, show that any system that could potentially achieve complementarity must have weighting function that varies with the inputs, must have human or algorithmic loss that varies, must have a combining function that is not convex, and cannot have either the human or algorithm dominate the loss of the other.

5.2 Cases Where Complementarity is Possible: $N = 2$ regimes

Having shown cases where complementarity is impossible, in this section we give conditions where complementarity is possible. As we described in Section 3, our notation describes the loss of the unassisted human in regime i by $h_i = H + \delta_{hi}$ and the loss of the algorithm by $a_i = A + \delta_{ai}$, where H and A are the average losses of the unassisted human and algorithm, respectively. Additionally, $w_h(a_i, h_i)$ (as defined in Equation 1) is the weighting function—the weight that the human places on themselves in making a final prediction.

We will first build intuition with the $N = 2$ case: there are exactly two regimes, with regime 1 having probability p of occurring. By assumption, $\sum_{i=1}^N p_i \cdot \delta_{ai} = 0$ (and similarly for unaided human loss). This allows us to simplify the $N = 2$ loss distributions into:

$$\text{unaided human loss: } \begin{cases} h_1 = H + \delta_h \\ h_2 = H - \frac{p}{1-p} \delta_h \end{cases} \quad \text{algorithmic loss: } \begin{cases} a_1 = A + \delta_a \\ a_2 = A - \frac{p}{1-p} \delta_a \end{cases}$$

This formulation is without loss of generality because we allow δ_a, δ_h to be positive, negative, or zero. In the case that δ_a, δ_h are the same sign, loss rates for the human and algorithm are correlated. If they are of a different sign, then loss rates are anti-correlated. δ_a, δ_h values of larger magnitude correspond to a more variable distribution of losses.

Lemma 5 gives a condition for when a $N = 2$ example can achieve complementarity. Note that the condition gives a lower bound on the magnitude of $|\delta_a - \delta_h|$, terms which reflect the variability in human and algorithmic loss. This bound depends on A and H , the average loss rates of the algorithm and human. It also depends on $w_h(a_1, h_1)$, $w_h(a_2, h_2)$, the weighting functions for the human in the 1st and 2nd regime, respectively. Recall that Corollary 1 tells us that any system exhibiting complementarity cannot have $w_h(a_1, h_1) = w_h(a_2, h_2)$, because that would imply a constant weighting function.

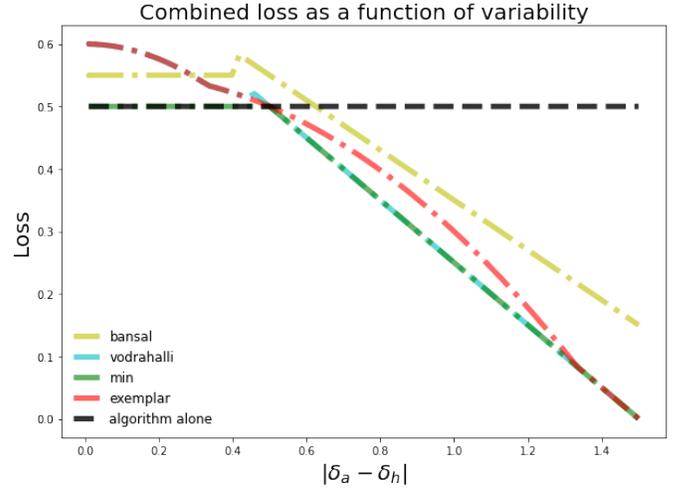


Figure 2: Low combined loss occurs for high variability (large $|\delta_a - \delta_h|$). The plot displays combined loss for a $N = 2$ system, for four combining functions. The black dashed line gives loss of algorithm alone: since for this setting the algorithm has lower error than the unaided human, complementarity occurs below this line. Note that for all four combining functions, complementarity occurs where $|\delta_a - \delta_h|$ is large. For details on the combining functions, including the exemplar function (original to this work), see Appendix A.

LEMMA 5. *Consider the case where $N = 2$, and WLOG assume that $A \leq H$: the algorithm has lower average loss than the human. Then, the combined system exhibits complementarity whenever:*

$$(H - A) \cdot \frac{w_h(a_1, h_1) + \frac{1-p}{p} \cdot w_h(a_2, h_2)}{|w_h(a_2, h_2) - w_h(a_1, h_1)|} < |\delta_a - \delta_h|$$

Figure 2 displays the loss of the combined system for four different possible combining functions (models of how humans incorporate algorithmic advice). For each of them, complementarity occurs when $|\delta_a - \delta_h|$ is high, as Lemma 5 suggests. However, as Lemma 6 states, this does *not* mean that unaided human and algorithm loss need to be anti-correlated.

LEMMA 6. *A system can exhibit complementarity even if the unaided human and the algorithm have correlated loss (both have higher loss in the same regime).*

This is a reassuring result because, otherwise, complementarity would not be possible in settings where a given regime is fundamentally harder than another. For example, one regime might be low-resolution images while another might be high-resolution images, so both humans and algorithms are expected to perform worse in the low-resolution regime. In our notation, if $\delta_a > 0$ and $\delta_h > 0$, then both the unaided human and algorithm have lower loss in regime 2. If $|\delta_a - \delta_h|$ is large, complementarity still may be possible. However, because δ_a, δ_h have the same signs, the variability must be larger in order to make $|\delta_a - \delta_h|$ satisfy the lower bound.

Tables 5 and 6 prove Lemma 6 through an example. Table 5 is a copy of Scenario 3 (Table 3) from the motivating example in Section 4. The values in the table correspond to $\delta_a = 0.4, \delta_h = -0.3$ (giving

	(Unaided) human	Algorithm	Combined (human using algorithm)	Weight (on unaided human)
Regime 1	1.15	0.2	0.44	0.25
Regime 2	0.35	0.8	0.46	0.75
Average	0.75	0.5	0.45	0.5

Table 5: (Reproduced version of Table 3). The unaided human has loss 0.75, the algorithm has loss 0.5, and the combined system has loss 0.45 (complementary performance). In this table, we have $\delta_a = 0.4, \delta_h = -0.3$, for $|\delta_a - \delta_h| = 0.7$. Note that here, losses are anti-correlated: the unaided human has higher loss in regime 1, while the algorithm has higher loss in regime 2.

	(Unaided) human	Algorithm	Combined (human using algorithm)	Weight (on unaided human)
Regime 1	1.48	0.53	0.77	0.25
Regime 2	0.02	0.47	0.13	0.75
Average	0.75	0.5	0.45	0.5

Table 6: Correlated loss: The combined system (human using algorithm) displays complementary performance even though the unaided human and the algorithm both have higher loss for regime 1. In this table, we have $\delta_a = 0.73, \delta_h = 0.03$, for $|\delta_a - \delta_h| = 0.7$.

$|\delta_a - \delta_h| = 0.7$). Here, the losses are anti-correlated: the unaided human has higher loss in regime 1, while the algorithm has higher loss in regime 2.

Table 6 gives an example where the losses are correlated: both the unaided human and algorithm have higher loss for regime 1. The values in the table are given by $\delta_a = 0.03, \delta_h = 0.73$, which again gives $|\delta_a - \delta_h| = 0.7$. Even though losses are correlated, the overall system still displays complementarity because $|\delta_a - \delta_h|$ has remained the same, as Lemma 5 suggests.

5.3 Cases Where Complementarity is Possible: $N > 2$ regimes

Finally, we will consider the general $N > 2$ case.

LEMMA 7. *WLOG, assume that $A \leq H$: the algorithm has lower loss, on average. Then, the condition below gives necessary and sufficient conditions for complementarity of the human-algorithm system:*

$$(H - A) \cdot \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) < \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) \cdot (\delta_{ai} - \delta_{hi})$$

If we view $w_h(a_i, h_i)$ and δ_{ai}, δ_{hi} as random variables over the instance space with probability mass governed the distribution of instances given by $\{p_i\}$, then we can interpret the condition as:

$$(H - A) \cdot \mathbb{E}[w_h(a_i, h_i)] < \text{Cov}(w_h(a_i, h_i), \delta_{ai} - \delta_{hi})$$

where $\text{Cov}(\cdot)$ gives the covariance.

Lemma 7 gives conditions on complementarity, requiring that the weighting function $w_h(a_i, h_i)$ have high covariance with the difference between δ_{ai} and δ_{hi} . Intuitively, this means that when the algorithm is more *above* its typical loss than the unaided human (when $\delta_{ai} > \delta_{hi}$), then the combined loss should rely more heavily on the unaided human ($w_h(a_i, h_i)$ should be large). Conversely, if the algorithm is more *below* its typical loss than the unaided human (when $\delta_{ai} < \delta_{hi}$), then the combined loss should rely more heavily on the algorithm ($w_h(a_i, h_i)$ should be small). The lefthand side of the equation lower bounds how large this covariance must, as a function of the gap between the average loss of the unaided

human and the algorithm, and the expected value of the weighting function.

The results in this section imply that unaided human and algorithmic error rates must be highly variable. One question this raises is about achievability: is it even possible to arrange human and algorithmic loss rates like this? For example, achieving a highly variable loss rate for an algorithm may require retraining, and for certain cases may not be possible. Manipulating the loss rate for a human may be possible through re-assigning human effort: for example, assigning multiple humans to certain portions of the input space in order to reduce loss.

Even if highly variable loss rates are possible, they may not be desirable for other reasons. In particular, the next section discusses the fairness implications of complementarity.

6 FAIRNESS

There are numerous possible notions of fairness that could be relevant for a human-algorithm classifier. In this section, we will analyze three of them and describe how they relate to complementarity. (All proofs are given in Appendix B). In general, the notions of fairness concern disparities in loss across different regimes. Disparities in loss rates could be alarming (if they align or correlate with sensitive attributes, such as race, gender, or socioeconomic status) or innocuous (if they are unrelated to sensitive attributes). In this paper, we will consider two general classes of fairness concerns: *fairness of benefit* and *loss disparity* rates. Fairness of benefit relates to which regimes see their loss rate decrease when the algorithm is incorporated into the decision-making process. Ideally, all regimes would benefit from human-algorithm collaboration. We will show that is unfortunately not possible. Loss disparity rates relate to the gap in loss rates between different regimes: ideally, this gap would be small, which would reflect equal loss rates for all regimes. In this work, we will show that complementarity can sometimes help bound loss disparity.

6.1 Fairness of benefit

First, Definition 2 says that a system has “fairness of benefit” if all regimes experience a lower loss from the combined human-algorithm system than they would experience with the unassisted human.

DEFINITION 2 (FAIRNESS OF BENEFIT). *A human/algorithm system exhibits fairness of benefit if all regimes benefit from the combined system (experience a lower loss than with the unaided human).*

This notion captures the desideratum that the benefits resulting from switching to a combined human-algorithm system are shared by all. However, Lemma 8 shows that this notion of fairness is incompatible with complementarity.

LEMMA 8. *Any system exhibiting fairness of benefits cannot have complementary performance.*

PROOF. This result is largely due to Lemma 3. Note that we have:

$$c(a_i, h_i) = (1 - w_h(a_i, h_i)) \cdot a_i + w_h(a_i, h_i) \cdot h_i$$

for $s_i \in [0, 1]$. Then, if we have fairness of benefits, we know:

$$(1 - w_h(a_i, h_i)) \cdot a_i + w_h(a_i, h_i) \cdot h_i < h_i \quad \forall i \in [N]$$

In order to achieve this, we need $w_h(a_i, h_i) < 1$ and $a_i < h_i$, for all $i \in [N]$. However, this is exactly the condition addressed in Lemma 3: the algorithm always has lower loss than the unassisted human, which means that complementarity is impossible. \square

Note that Definition 2 could be defined symmetrically, defining “benefit” as a reduction in loss as compared to the algorithmic prediction. In this case, an analogous version of Lemma 8 would hold, similarly showing that this notion of fairness is incompatible with complementarity.

Lemma 8 tells us that there is an inherent tension between achieving complementarity and ensuring all people who use the system see their loss rates decrease. For certain application areas, fairness of benefit might be more important than complementarity, so practitioners might consciously choose to prioritize it. On the other hand, a practitioner who opts to achieve complementarity instead of fairness of benefit might wish to consider alternate ways to support any groups that see their loss increase in the combined system.

6.2 Loss disparity

Next, Definition 3 describes fairness as the disparity in loss rates between different regimes. Intuitively, this notion relates to group-based notions of fairness: loss rates should be relatively similar between members of different groups. However, Lemma 9 again describes how this notion of fairness is in tension with complementarity, which puts a lower bound on the level of this kind of unfairness.

DEFINITION 3 (LOSS DISPARITY). *A prediction system exhibits ϵ -loss disparity if the losses in different regimes differ by no more than ϵ . We will use $\epsilon_h, \epsilon_a, \epsilon_c$ to refer to the loss disparity of the unaided human, algorithm alone, and combined human-algorithm system, respectively.*

LEMMA 9. *WLOG assume that $A \leq H$: the algorithm has lower average loss than the human. Then, any system exhibiting complementarity has a lower bound on $\epsilon_a + \epsilon_h$: the combined loss disparity of the unaided human and algorithm.*

$$A - C + (H - A) \cdot \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) < \epsilon_a + \epsilon_h$$

Lemma 9 should match our intuition from Section 5. There, results indicated that complementarity is easiest to achieve when loss rates are highly variable. However, this directly contradicts with the goal of minimizing loss disparity.

Next, we will consider the loss disparity of the combined human-algorithm system. Unfortunately, it is possible for a system exhibiting complementarity to *exacerbate* unfairness. Table 7 gives an example where the combined system has a higher loss disparity than either the unaided human or the algorithm, even though it exhibits complementarity. This result (which parallels results from [10, 11, 36]) means that practitioners must be careful when discussing fairness of combined systems: fairness guarantees from the individual components don’t necessarily transfer to the combined human-algorithm system. However, Lemma 10 gives a condition where the loss disparity of the combined human-algorithm system is upper bounded by the loss disparity of the unaided human or the algorithm.

LEMMA 10. *Define $i+$ as the regime where the combined human-algorithm system has highest loss and $i-$ as the regime where it has lowest loss. Then, the loss disparity of the combined system is upper bounded by the loss disparity of the unaided human or algorithm, so long as neither the unaided human or algorithm dominates the other in both $i+, i-$. That is,*

$$\begin{aligned} \text{If either case is satisfied: } & \begin{cases} h_{i+} \leq a_{i+} \text{ and } h_{i-} \geq a_{i-} \\ h_{i+} \geq a_{i+} \text{ and } h_{i-} \leq a_{i-} \end{cases} \\ \Rightarrow & \epsilon_c \leq \max(\epsilon_a, \epsilon_h) \end{aligned}$$

This last lemma gives our first positive result for fairness: it gives conditions where human-algorithm system exhibiting complementarity at least doesn’t exacerbate unfairness. Specifically, what it requires is that neither the unaided human nor the algorithm dominates the other for both of the most extreme regimes (where the combined system has highest and lowest loss). As we would expect, the scenario in Table 7 violates this: the unaided human dominates in both regime 1 (highest loss) and regime 2 (lowest loss), which allows the combined loss disparity ϵ_c to be greater than $\max(\epsilon_a, \epsilon_h)$. Interestingly, Lemma 10 is quite powerful: it only relies on the losses within two specific regimes and holds regardless of whether the overall system satisfies complementarity. Practitioners could use Lemma 10 to guide their algorithm development: so long as the preconditions are satisfied, they can guarantee that the combined system will never exacerbate unfairness.

As we mentioned previously, these definitions are only a few of the possible fairness concerns we could analyze. However, this analysis highlights the importance of considering fairness, especially ways that it might be in tension with achieving complementarity.

	(Unaided) Human	Algorithm	Combined (human using algorithm)
Regime 1	0.95	0.85	0.895
Regime 2	0.95	0.02	0.05
Regime 3	0.15	0.45	0.255
Average loss	0.68	0.44	0.40

Table 7: This system exhibits complementarity, since average loss is lowest in the combined human-algorithm system. However, note that loss disparity is increased in the combined system: $\epsilon_h = 0.8$, $\epsilon_a = 0.83$, but $\epsilon_c = 0.84$.

7 CONCLUSION AND FUTURE DIRECTIONS

In this work, we introduce a simple theoretical model of human-algorithm collaboration, which we show is flexible enough to encompass models analyzed in prior work. Using this model, we obtain theoretical impossibility results that characterize settings where complementarity is not achievable. We also use this framework to construct cases where complementarity is possible, given certain conditions on the loss distributions. Finally, we consider the implications, especially fairness, of the requirements in order to achieve complementarity.

Our approach admits multiple possible avenues for future work. Our work highlights the importance of variable loss rates: algorithmic loss that is not constant over the input space. However, as mentioned previously, it may not be possible to achieve extremely variable loss rates. Future work could model algorithmic loss rates more explicitly, describing loss distributions that are both achievable and lead to complementary performance. For example, some prior work has demonstrated that, for many algorithms, reducing loss becomes harder as the level of loss decreases, which could make it more difficult to achieve highly variable loss rates [15, 16].

Similarly, future work could relax assumptions made in our work. Relaxing Assumption 2, for example, could involve analyzing combining rules that, on any individual regime, could do better or worse than the human or algorithmic input loss rates. Relaxing Assumption 1 could involve modeling cases where regimes with identical human and algorithmic loss rates might be treated differently by the combiner. Any of these future analyses could allow us to have greater insight into a variety of ways human-algorithm systems perform.

ACKNOWLEDGMENTS

This work was supported in part by NSF grant DGE-1650441. We are grateful to Vijay Keswani, Jon Kleinberg, Pang Wei Koh, Michela Meister, Emma Pierson, Charvi Rastogi, Aaron Roth, Kiran Tomlinson, Aaron Tucker, Qian Yang, Joyce Zhou, James Zou, the AWS Machine Learning team, the AI, Policy, and Practice working group at Cornell, and the attendees at the NeurIPS 2021 Workshops on Human-Centered AI and Human and Machine Decisions for invaluable discussions.

REFERENCES

- [1] Gagan Bansal, Besmira Nushi, Ece Kamar, Eric Horvitz, and Daniel S Weld. 2021. Is the Most Accurate AI the Best Teammate? Optimizing AI for Teamwork. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 11405–11414.
- [2] Gagan Bansal, Tongshuang Wu, Joyce Zhou, Raymond Fok, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, and Daniel Weld. 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.
- [3] Emma Beede, Elizabeth Baylor, Fred Hersch, Anna Iurchenko, Lauren Wilcox, Paisan Ruamviboonsuk, and Laura M Vardoulakis. 2020. A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [4] Bo Cowgill and Megan T Stevenson. 2020. Algorithmic social engineering. In *AEA Papers and Proceedings*, Vol. 110. 96–100.
- [5] Abir De, Nastaran Okati, Ali Zaregade, and Manuel Gomez-Rodriguez. 2021. Classification under human assistance. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 5905–5913.
- [6] Maria De-Arteaga, Riccardo Fogliato, and Alexandra Chouldechova. 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.
- [7] Luca Didaci, Giorgio Fumera, and Fabio Roli. 2013. Diversity in classifier ensembles: Fertile concept or dead end? In *International workshop on multiple classifier systems*. Springer, 37–48.
- [8] Berkeley J Dietvorst and Soham Bharti. 2020. People reject algorithms in uncertain decision domains because they have diminishing sensitivity to forecasting error. *Psychological science* 31, 10 (2020), 1302–1314.
- [9] Berkeley J Dietvorst, Joseph P Simmons, and Cade Massey. 2018. Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science* 64, 3 (2018), 1155–1170.
- [10] Cynthia Dwork and Christina Ilvento. 2018. Fairness Under Composition. In *10th Innovations in Theoretical Computer Science Conference (ITCS 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [11] Cynthia Dwork, Christina Ilvento, and Meena Jagadeesan. 2020. Individual Fairness in Pipelines. In *1st Symposium on Foundations of Responsible Computing (FORC 2020)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [12] Aaron Fraenkel. 2020. Fairness and Algorithmic Decision Making. <https://fraenkel.github.io/fairness-algo-decision>
- [13] Robert Geirhos, Kristof Meding, and Felix A Wichmann. 2020. Beyond accuracy: quantifying trial-by-trial behaviour of CNNs and humans by measuring error consistency. *Advances in Neural Information Processing Systems* 33 (2020), 13890–13902.
- [14] Talia Gillis, Bryce McLaughlin, and Jann Spiess. 2021. On the Fairness of Machine-Assisted Human Decisions. *arXiv preprint arXiv:2110.15310* (2021).
- [15] Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory F. Diamos, Heewoo Jun, Hassan Kianinejad, Md. Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. 2017. Deep Learning Scaling is Predictable, Empirically. *CoRR abs/1712.00409* (2017). [arXiv:1712.00409](http://arxiv.org/abs/1712.00409) <http://arxiv.org/abs/1712.00409>
- [16] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361* (2020).
- [17] Gavin Kerrigan, Padhraic Smyth, and Mark Steyvers. 2021. Combining Human Predictions with Model Probabilities via Confusion Matrices and Calibration. (2021).
- [18] Vijay Keswani, Matthew Lease, and Krishnaram Kenthapadi. 2021. Towards unbiased and accurate deferral to multiple experts. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 154–165.
- [19] Vijay Keswani, Matthew Lease, and Krishnaram Kenthapadi. 2022. Designing Closed Human-in-the-loop Deferral Pipelines. *arXiv preprint arXiv:2202.04718* (2022).
- [20] Ludmila I Kuncheva. 2014. *Combining pattern classifiers: methods and algorithms*. John Wiley & Sons.
- [21] Sarah Lebovitz, Natalia Levina, and Hila Lifshitz-Assaf. 2021. Is AI ground truth really “true”? The dangers of training and evaluating AI tools based on experts’ know-what. *Management Information Systems Quarterly* (2021).
- [22] Sarah Lebovitz, Hila Lifshitz-Assaf, and Natalia Levina. 2020. To incorporate or not to incorporate AI for critical judgments: The importance of ambiguity in professionals’ judgment process. *Collective Intelligence, The Association for Computing Machinery* (2020).
- [23] David Madras, Toni Pitassi, and Richard Zemel. 2018. Predict Responsibly: Improving Fairness and Accuracy by Learning to Defer. In *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2018/file/09d37c08f7b129e96277388757530c72-Paper.pdf>
- [24] Paul E Meehl. 1954. Clinical versus statistical prediction: A theoretical analysis and a review of the evidence. (1954).

- [25] Vahid Balazadeh Meresht, Abir De, Adish Singla, and Manuel Gomez-Rodriguez. 2020. Learning to switch between machines and humans. *ICML Workshop on Human-AI Collaboration in Sequential Decision-Making*. (2020).
- [26] Nastaran Okati, Abir De, and Manuel Rodriguez. 2021. Differentiable learning under triage. *Advances in Neural Information Processing Systems* 34 (2021).
- [27] Chinasa T Okolo, Srujana Kamath, Nicola Dell, and Aditya Vashistha. 2021. “It cannot do all of my work”: Community Health Worker Perceptions of AI-Enabled Mobile Health Applications in Rural India. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [28] Forough Poursabzi-Sangdeh, Daniel G Goldstein, Jake M Hofman, Jennifer Wortman Vaughan, and Hanna Wallach. 2021. Manipulating and measuring model interpretability. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–52.
- [29] Maithra Raghu, Katy Blumer, Greg Corrado, Jon Kleinberg, Ziad Obermeyer, and Sendhil Mullainathan. 2018. The algorithmic automation problem: Prediction, triage, and human effort. *NeurIPS Workshop on Machine Learning for Health (ML4H)* (2018).
- [30] Rory Sayres, Ankur Taly, Ehsan Rahimy, Katy Blumer, David Coz, Naama Hammel, Jonathan Krause, Arunachalam Narayanaswamy, Zahra Rastegar, Derek Wu, Shawn Xu, Scott Barb, Anthony Joseph, Michael Shumski, Jesse Smith, Arjun B. Sood, Greg S. Corrado, Lily Peng, and Dale R. Webster. 2019. Using a deep learning algorithm and integrated gradients explanation to assist grading for diabetic retinopathy. *Ophthalmology* 126, 4 (2019), 552–564.
- [31] Megha Srivastava, Besmira Nushi, Ece Kamar, Shital Shah, and Eric Horvitz. 2020. An empirical analysis of backward compatibility in machine learning systems. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 3272–3280.
- [32] Eleni Straitouri, Lequn Wang, Nastaran Okati, and Manuel Gomez Rodriguez. 2022. Provably Improving Expert Predictions with Conformal Prediction. arXiv:2201.12006 [cs.LG]
- [33] Sarah Tan, Julius Adebayo, Kori Inkpen, and Ece Kamar. 2018. Investigating human+ machine complementarity for recidivism predictions. *NeurIPS 2018 Workshop on Ethical, Social and Governance Issues in AI* (2018).
- [34] I. Valera, A. Singla, and M. Gomez Rodriguez. 2018. Enhancing the Accuracy and Fairness of Human Decision Making. In *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*. Curran Associates, Inc., 1774–1783. <http://papers.nips.cc/paper/7448-enhancing-the-accuracy-and-fairness-of-human-decision-making.pdf>
- [35] Kailas Vodrahalli, Tobias Gerstenberg, and James Zou. 2021. Do Humans Trust Advice More if it Comes from AI? An Analysis of Human-AI Interactions. *arXiv preprint arXiv:2107.07015* (2021).
- [36] Xuezhong Wang, Nithum Thain, Anu Sinha, Flavien Prost, Ed H Chi, Jilin Chen, and Alex Beutel. 2021. Practical Compositional Fairness: Understanding Fairness in Multi-Component Recommender Systems. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. 436–444.
- [37] Qian Yang, Alex Scuito, John Zimmerman, Jodi Forlizzi, and Aaron Steinfeld. 2018. Investigating how experienced UX designers effectively work with machine learning. In *Proceedings of the 2018 Designing Interactive Systems Conference*. 585–596.
- [38] Ming Yin, Jennifer Wortman Vaughan, and Hanna Wallach. 2019. Understanding the effect of accuracy on trust in machine learning models. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [39] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummadi. 2017. Fairness Beyond Disparate Treatment & Disparate Impact: Learning Classification without Disparate Mistreatment. In *Proceedings of the 26th International Conference on World Wide Web (Perth, Australia) (WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1171–1180. <https://doi.org/10.1145/3038912.3052660>

A MODELING COMBINING FUNCTIONS

The definitions presented in Section 3 reflect multiple ways humans could incorporate algorithmic inputs. However, their functional forms in some cases are less clean to analyze. Definition 4 gives an *exemplar* weighting function (original to this work) that we created in order to illustrate common patterns in weighting functions, while also allowing for tractable theoretical analysis. In this function, for $m > 0$, the combining rule is more likely to select the algorithm when $a < h$ (when the algorithm has lower loss rate). For $m < 0$, the reverse is true: the combining rule “mistakenly” goes with the input with higher loss.

DEFINITION 4 (EXEMPLAR WEIGHTING FUNCTION). *The exemplar weighting function is given by:*

$$w_h(a, h) = \begin{cases} b - m \cdot (h - a) & 0 \leq b - m \cdot (h - a) \leq 1 \\ 0 & b - m \cdot (h - a) < 0 \\ 1 & b - m \cdot (h - a) > 1 \end{cases}$$

Figure 3 plots the weighting function for each of the selection rules presented. Note that, in general, as the human gets higher loss, the weight on the human decreases.

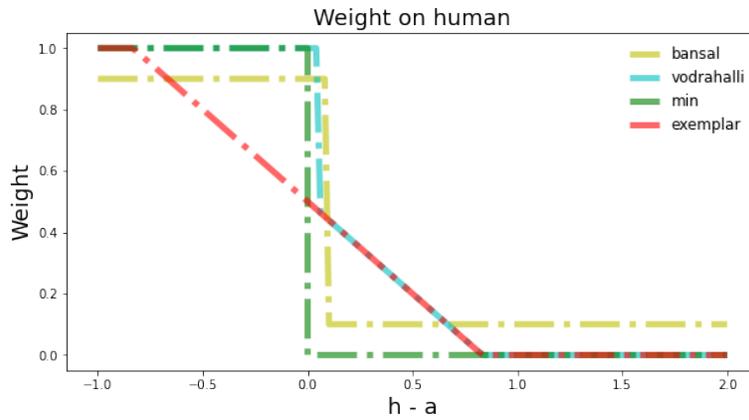


Figure 3: The weight on the human, given a difference $h - a$ between algorithmic and human loss, for multiple weighting functions. Note that, in general, as the human gets higher loss than the algorithm, the weight on the human decreases.

Lemma 11, below, gives an analogous version of Lemma 5 for the exemplar function. Note that it similarly shows that complementarity occurs when $|\delta_a - \delta_h|$ is large (when losses are highly variable).

LEMMA 11. *Consider the exemplar weighting function with $N = 2$ and $w_h(a_0, h_0), w_h(a_1, h_1) < 1$, and where $A \leq H$. Then, the system exhibits complementarity whenever:*

$$\sqrt{H - A} \cdot \sqrt{\frac{1 - p}{p} \cdot \left(\frac{1 - b}{m} - (H - A) \right)} < |\delta_a - \delta_h|$$

B PROOFS

LEMMA 12. *Any combining rule relying only on loss rates (Assumption 1) with bounded output (Assumption 2) can be written as a combining rule with a weighting function $0 \leq w_h(a_i, h_i) \leq 1$.*

PROOF. We will define a weighting function as follows:

$$w_h(a, h) = \begin{cases} \frac{c(a, h) - a}{h - a} & a \neq h \\ \frac{1}{2} & a = h \end{cases}$$

Note that if $c(a, h) = h$, then $w_h(a, h) = 1$ (the weighting function puts all weight on the unaided human), and if $c(a, h) = a$, then $w_h(a, h) = 0$ (the weighting function puts no weight on the unaided human).

First, we'll show that $w_h(a, h) \in [0, 1]$. If $a = h$, this is true by construction (because $w_h(a, h) = \frac{1}{2}$). What we want to show is:

$$0 \leq \frac{c(a, h) - a}{h - a} \leq 1$$

If $h > a$, then this is equivalent to requiring:

$$0 \leq c(a, h) - a \leq h - a$$

The lefthand inequality is satisfied because $c(a, h) \geq \min(a, h) = a$ (in this case). The righthand inequality is satisfied because $c(a, h) \leq \max(a, h) = h$ (in this case).

On the other hand, if $h < a$, then the inequality we want to show is:

$$0 \geq c(a, h) - a \geq h - a$$

Again, the lefthand inequality is satisfied because $c(a, h) \leq \max(a, h) = a$ (in this case). The righthand inequality is satisfied because $c(a, h) \geq \min(a, h) = h$ (in this case).

Next, we'll show that this weighting function is correct: that is, that:

$$c(a, h) = (1 - w_h(a, h)) \cdot a + w_h(a, h) \cdot h$$

Note that Assumption 2 means that if $a = h$, then $c(a, h) = a = h$, so any $w_h(a, h) \in [0, 1]$ would result in a correct weighting function. For $a \neq h$, we can write:

$$\begin{aligned} & (1 - w_h(a, h)) \cdot a + w_h(a, h) \cdot h \\ = & a \cdot \left(1 - \frac{c(a, h) - a}{h - a}\right) + h \cdot \frac{c(a, h) - a}{h - a} \\ = & a \cdot \frac{h - a - c(a, h) + a}{h - a} + h \cdot \frac{c(a, h) - a}{h - a} \\ = & a \cdot \frac{h - c(a, h)}{h - a} + h \cdot \frac{c(a, h) - a}{h - a} \\ = & \frac{a \cdot h - a \cdot c(a, h) + h \cdot c(a, h) - a \cdot h}{h - a} \\ = & \frac{c(a, h) \cdot (h - a)}{h - a} \\ = & c(a, h) \end{aligned}$$

as desired. □

LEMMA 13. *A human-algorithm system where unaided human and algorithm loss rates are constant over regimes can never achieve complementary performance.*

PROOF. Constant loss rates means that $a_i = A$ and $h_i = H$ for all $i \in [N]$. The combined system has loss:

$$\sum_{i \in [N]} p_i \cdot c(a_i, h_i) = \sum_{i \in [N]} p_i \cdot c(A, H) = c(A, H) \geq \min(A, H)$$

□

LEMMA 14. *Complementarity is impossible if one of the human or algorithm always weakly dominates the loss of the other: that is, if $a_i \leq h_i$ for all i , or $a_i \geq h_i$ for all i .*

PROOF. WLOG, we will assume that $A \leq H$. Then, complementarity occurs when:

$$\begin{aligned} & \sum_{i=1}^N p_i \cdot c(a_i, h_i) < A \\ \sum_{i=1}^N p_i \cdot ((1 - w_h(a_i, h_i)) \cdot a_i + w_h(a_i, h_i) \cdot h_i) & < \sum_{i=1}^N p_i \cdot a_i \\ \sum_{i=1}^N p_i \cdot (-w_h(a_i, h_i) a_i + w_h(a_i, h_i) \cdot h_i) & < 0 \\ \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) \cdot (h_i - a_i) & < 0 \end{aligned}$$

If the algorithm weakly dominates the unaided human ($h_i \geq a_i$ for all i), then this inequality can never be satisfied because the entire lefthand side is positive or 0. If the unaided human weakly dominates the algorithm ($a_i \geq h_i$ for all i) then $A \geq H$. From our previous assumption that $A \leq H$, which means that we must have that $A = H$: both have identical average losses.

We will now show that further analysis means in the case that $A = H$, complementarity is still impossible. If $A = H$, we could have equivalently written the complementarity condition as:

$$\begin{aligned} \sum_{i=1}^N p_i \cdot c(a_i, h_i) &< H = A \\ \sum_{i=1}^N p_i \cdot ((1 - w_h(a_i, h_i)) \cdot a_i + w_h(a_i, h_i) \cdot h_i) &< \sum_{i=1}^N p_i \cdot h_i \\ \sum_{i=1}^N p_i \cdot ((1 - w_h(a_i, h_i)) \cdot a_i + (1 - w_h(a_i, h_i)) \cdot h_i) &< 0 \\ \sum_{i=1}^N p_i \cdot (1 - w_h(a_i, h_i)) \cdot (a_i - h_i) &< 0 \end{aligned}$$

In this case, we've assumed that $a_i \geq h_i$. However, this means that the lefthand side of the above inequality is positive, which means complementarity isn't satisfied. \square

LEMMA 15. *A combining function $c(a_i, h_i)$ that is convex in a_i, h_i can never achieve complementary performance.*

PROOF. The first inequality in this proof is by Jensen's inequality and the last inequality is due to our construction of the combining function:

$$\sum_{i \in [N]} p_i \cdot c(a_i, h_i) \geq c\left(\sum_{i \in [N]} p_i \cdot a_i, \sum_{i \in [N]} p_i \cdot h_i\right) = c(A, H) \geq \min(A, H)$$

\square

LEMMA 16. *Consider the case where $N = 2$, and WLOG assume that $A \leq H$: the algorithm has lower average loss than the human. Then, the combined system exhibits complementarity whenever:*

$$(H - A) \cdot \frac{w_h(a_1, h_1) + \frac{1-p}{p} \cdot w_h(a_2, h_2)}{|w_h(a_2, h_2) - w_h(a_1, h_1)|} < |\delta_a - \delta_h|$$

PROOF. If we assume that $A \leq H$, then complementarity occurs whenever $C < A$, or: The system exhibits complementarity when:

$$\begin{aligned} p \cdot c(a_1, h_1) + (1-p) \cdot c(a_2, h_2) &< A \\ p \cdot ((1 - w_h(a_1, h_1)) \cdot a_1 + w_h(a_1, h_1) \cdot h_1) + (1-p) \cdot ((1 - w_h(a_2, h_2)) \cdot a_2 + w_h(a_2, h_2) \cdot h_2) &< p \cdot a_1 + (1-p) \cdot a_2 \\ p \cdot w_h(a_1, h_1) \cdot (h_1 - a_1) + (1-p) \cdot w_h(a_2, h_2) \cdot (h_2 - a_2) &< 0 \end{aligned}$$

We insert the values of a_i, h_i , to get:

$$\begin{aligned} p \cdot w_h(a_1, h_1) \cdot (H + \delta_h - A - \delta_a) + (1-p) \cdot w_h(a_2, h_2) \cdot \left(H - \frac{p}{1-p} \cdot \delta_h - A + \frac{p}{1-p} \cdot \delta_a\right) &< 0 \\ (H - A) \cdot (p \cdot w_h(a_1, h_1) + (1-p) \cdot w_h(a_2, h_2)) + p \cdot w_h(a_1, h_1) \cdot (\delta_h - \delta_a) - p \cdot w_h(a_2, h_2) \cdot (\delta_h - \delta_a) &< 0 \\ (H - A) \cdot (p \cdot w_h(a_1, h_1) + (1-p) \cdot w_h(a_2, h_2)) - p \cdot (\delta_a - \delta_h) \cdot (w_h(a_1, h_1) - w_h(a_2, h_2)) &< 0 \end{aligned}$$

Note that we have assumed $H \geq A$, so in order for this inequality to hold, we must have that the other term (involving δ_a, δ_h) be positive. Note that by Lemma 2 we cannot have $w_h(a_1, h_1) = w_h(a_2, h_2)$ if it exhibits complementarity. In order to have the inequality satisfied, one of two conditions must hold:

- **Case 1:** $\delta_a > \delta_h$ and $w_h(a_1, h_1) > w_h(a_2, h_2)$ (the unaided human is weighted more heavily in regime 1). Note that in this case, we must have $h_2 > a_2$, because:

$$\begin{aligned} h_2 &> a_2 \\ H - \frac{p}{1-p} \cdot \delta_h &> A - \frac{p}{1-p} \cdot \delta_a \\ H - A - \frac{p}{1-p} \cdot (\delta_h - \delta_a) &> 0 \\ H - A + \frac{p}{1-p} \cdot (\delta_a - \delta_h) &> 0 \end{aligned}$$

which is satisfied because $H - A \geq 0$ and $\delta_a > \delta_h$. By Lemma 3, we know that $h_2 > a_2$ implies that $h_1 \leq a_1$. Taken together with $w_h(a_1, h_1) > w_h(a_2, h_2)$, this means that the combined system must weight the human more heavily in instance 1, where it has lower loss than the human.

- **Case 2:** $\delta_a < \delta_h$ and $w_h(a_1, h_1) < w_h(a_2, h_2)$ (the unaided human is weighted more heavily in regime 2). In this case, we must have $h_1 > a_1$, because:

$$\begin{aligned} H + \delta_h &> A + \delta_a \\ H - A + \delta_h - \delta_a &> 0 \end{aligned}$$

This must be satisfied because $H - A > 0$ and $\delta_h > \delta_a$. By similar reasoning to above, this means that $h_2 \leq a_2$. Taken with $w_h(a_1, h_1) < w_h(a_2, h_2)$, this again means that the system must weight the unaided more heavily in the regime where it has lower loss.

Finally, we can simplify the inequality:

$$(H - A) \cdot (p \cdot w_h(a_1, h_1) + (1 - p) \cdot w_h(a_2, h_2)) < p \cdot (\delta_a - \delta_h) \cdot (w_h(a_1, h_1) - w_h(a_2, h_2))$$

If $\delta_a > \delta_h$, then we know that $w_h(a_1, h_1) > w_h(a_2, h_2)$, so we can rewrite this as:

$$(H - A) \cdot \frac{w_h(a_1, h_1) + \frac{1-p}{p} \cdot w_h(a_2, h_2)}{w_h(a_1, h_1) - w_h(a_2, h_2)} \leq \delta_a - \delta_h$$

On the other hand, if $\delta_a < \delta_h$, we know that $w_h(a_1, h_1) < w_h(a_2, h_2)$, so we can rewrite this as:

$$\begin{aligned} (H - A) \cdot (p \cdot w_h(a_1, h_1) + (1 - p) \cdot w_h(a_2, h_2)) &< p \cdot (\delta_h - \delta_a) \cdot (w_h(a_2, h_2) - w_h(a_1, h_1)) \\ (H - A) \cdot \frac{w_h(a_1, h_1) + \frac{1-p}{p} \cdot w_h(a_2, h_2)}{(w_h(a_2, h_2) - w_h(a_1, h_1))} &< \delta_a - \delta_h \end{aligned}$$

Either way simplifies to:

$$(H - A) \cdot \frac{w_h(a_1, h_1) + \frac{1-p}{p} \cdot w_h(a_2, h_2)}{|w_h(a_2, h_2) - w_h(a_1, h_1)|} < |(\delta_a - \delta_h)|$$

□

LEMMA 17. *WLOG, assume that $A \leq H$: the algorithm has lower loss, on average. Then, the condition below gives necessary and sufficient conditions for complementarity of the human-algorithm system:*

$$(H - A) \cdot \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) < \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) \cdot (\delta_{ai} - \delta_{hi})$$

If we view $w_h(a_i, h_i)$ and δ_{ai}, δ_{hi} as random variables over the instance space with probability mass governed the distribution of instances given by $\{p_i\}$, then we can interpret the condition as:

$$(H - A) \cdot \mathbb{E}[w_h(a_i, h_i)] < \text{Cov}(w_h(a_i, h_i), \delta_{ai} - \delta_{hi})$$

where $\text{Cov}(\cdot)$ gives the covariance.

PROOF. If we assume that $A \leq H$, then complementarity occurs whenever $C < A$, or:

$$\begin{aligned} \sum_{i=1}^N p_i \cdot c(a_i, h_i) &\leq \sum_{i=1}^N p_i \cdot a_i \\ \sum_{i=1}^N p_i \cdot ((1 - w_h(a_i, h_i)) \cdot a_i + w_h(a_i, h_i) \cdot h_i) &\leq \sum_{i=1}^N p_i \cdot a_i \\ \sum_{i=1}^N p_i \cdot (-w_h(a_i, h_i) \cdot a_i + w_h(a_i, h_i) \cdot h_i) &\leq 0 \\ \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) \cdot (h_i - a_i) &\leq 0 \end{aligned}$$

Plugging in for the values gives:

$$\begin{aligned} \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) \cdot (H + \delta_{hi} - A - \delta_{ai}) &\leq 0 \\ (H - A) \cdot \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) &< \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) \cdot (\delta_{ai} - \delta_{hi}) \end{aligned}$$

as desired. \square

LEMMA 18. *WLOG assume that $A \leq H$: the algorithm has lower average loss than the human. Then, any system exhibiting complementarity has a lower bound on $\epsilon_a + \epsilon_h$: the combined loss disparity of the unaided human and algorithm.*

$$A - C + (H - A) \cdot \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) < \epsilon_a + \epsilon_h$$

PROOF. First, we will use the complementarity result from partway through the proof of Lemma 7. We calculate the difference between the average algorithmic loss and the average combined human-algorithmic loss (which gives us the average benefit of collaboration):

$$\begin{aligned} A - C &= \sum_{i=1}^N p_i \cdot a_i - \sum_{i=1}^N p_i \cdot c(a_i, h_i) \\ &= \sum_{i=1}^N p_i \cdot a_i - \sum_{i=1}^N p_i \cdot ((1 - w_h(a_i, h_i)) \cdot a_i + w_h(a_i, h_i) \cdot h_i) \\ &= \sum_{i=1}^N p_i \cdot (a_i - a_i + w_h(a_i, h_i) \cdot a_i - w_h(a_i, h_i) \cdot h_i) \\ &= \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) \cdot (a_i - h_i) \\ &= \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) \cdot (A + \delta_{ai} - H - \delta_{hi}) \\ &= (A - H) \cdot \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) + \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) \cdot (\delta_{ai} - \delta_{hi}) \end{aligned}$$

Rearranging gives:

$$A - C + (H - A) \cdot \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) = \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) \cdot (\delta_{ai} - \delta_{hi})$$

Next, we will define:

$$\begin{aligned} \epsilon_a &= a_{a^+} - a_{a^-} = A + \delta_{a^+} - A - \delta_{a^-} = \delta_{a^+} - \delta_{a^-} \\ \epsilon_h &= h_{h^+} - h_{h^-} = H + \delta_{h^+} - H - \delta_{h^-} = \delta_{h^+} - \delta_{h^-} \end{aligned}$$

where a^+, a^-, h^+, h^- are the indices of the maximum and minimum loss for the algorithm and unaided human, respectively. Note that we require $\sum_{i=1}^N p_i \cdot \delta_{ai} = \sum_{i=1}^N p_i \cdot \delta_{hi} = 0$, which implies that $\delta_{a^+}, \delta_{h^+} \geq 0$ and $\delta_{a^-}, \delta_{h^-} \leq 0$. Then, we know that:

$$\epsilon_a + \epsilon_h = \delta_{a^+} - \delta_{a^-} + \delta_{h^+} - \delta_{h^-} > \delta_{a^+} - \delta_{h^-}$$

Define $\mathcal{P} = \{i \mid \delta_{ai} \geq \delta_{hi}\}$ and $\mathcal{N} = \{i \mid \delta_{ai} < \delta_{hi}\}$. By definition, $\delta_{a^+} \geq \delta_{ai}$ and $\delta_{h^-} \leq \delta_{hi}$ for all $i \in [N]$. Then, we know that:

$$\delta_{a^+} - \delta_{h^-} \geq \sum_{i \in \mathcal{P}} p_i \cdot (\delta_{hi} - \delta_{ai}) \geq \sum_{i \in \mathcal{P}} p_i \cdot w_h(a_i, h_i) \cdot (\delta_{hi} - \delta_{ai})$$

where we have used the fact that $w_h(a_i, h_i) \leq 1$. Finally, we know that:

$$\sum_{i \in \mathcal{P}} p_i \cdot w_h(a_i, h_i) \cdot (\delta_{hi} - \delta_{ai}) \geq \sum_{i \in \mathcal{P}} p_i \cdot w_h(a_i, h_i) \cdot (\delta_{hi} - \delta_{ai}) + \sum_{i \in \mathcal{N}} p_i \cdot w_h(a_i, h_i) \cdot (\delta_{hi} - \delta_{ai}) = \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) \cdot (\delta_{hi} - \delta_{ai})$$

where the first inequality comes by the fact that $\delta_{hi} - \delta_{ai}$ for $i \in \mathcal{N}$. Finally, we can combine this analysis with our previous analysis on the gap in loss rate between the algorithm and the combined system:

$$A - C + (H - A) \cdot \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) = \sum_{i=1}^N p_i \cdot w_h(a_i, h_i) \cdot (\delta_{a_i} - \delta_{h_i}) \leq \delta_{a_+} - \delta_{h_-} < \epsilon_a + \epsilon_h$$

as desired. \square

LEMMA 19. Define i_+ as the regime where the combined human-algorithm system has highest loss and i_- as the regime where it has lowest loss. Then, the loss disparity of the combined system is upper bounded by the loss disparity of the unaided human or algorithm, so long as neither the unaided human or algorithm dominates the other in both i_+ , i_- . That is,

$$\text{If either case is satisfied: } \begin{cases} h_{i_+} \leq a_{i_+} \text{ and } h_{i_-} \geq a_{i_-} \\ h_{i_+} \geq a_{i_+} \text{ and } h_{i_-} \leq a_{i_-} \end{cases} \Rightarrow \epsilon_c \leq \max(\epsilon_a, \epsilon_h)$$

PROOF. We wish to upper bound ϵ_c , which is given by:

$$\epsilon_c = c(a_{i_+}, h_{i_+}) - c(a_{i_-}, h_{i_-})$$

such that:

$$i_+ = \operatorname{argmax}_{i \in [N]} c(a_i, h_i) \quad i_- = \operatorname{argmin}_{i \in [N]} c(a_i, h_i)$$

Note that by Assumption 2, we must have:

$$c(a_{i_+}, h_{i_+}) \leq \max(a_{i_+}, h_{i_+}) \text{ and } c(a_{i_-}, h_{i_-}) \geq \min(a_{i_-}, h_{i_-})$$

The statement of this lemma gives two cases, which we will consider in turn.

In Case 1, we assume that:

$$h_{i_+} \leq a_{i_+} \text{ and } h_{i_-} \geq a_{i_-}$$

In this case, we know that:

$$c(a_{i_+}, h_{i_+}) - c(a_{i_-}, h_{i_-}) \leq a_{i_+} - a_{i_-} \leq \epsilon_a \leq \max(\epsilon_a, \epsilon_h)$$

In Case 2, we assume that:

$$h_{i_+} \geq a_{i_+} \text{ and } h_{i_-} \leq a_{i_-}$$

In this case, we know that:

$$c(a_{i_+}, h_{i_+}) - c(a_{i_-}, h_{i_-}) \leq h_{i_+} - h_{i_-} \leq \epsilon_h \leq \max(\epsilon_a, \epsilon_h)$$

In either case, $\epsilon_c \leq \max(\epsilon_a, \epsilon_h)$ \square

LEMMA 20. Consider the exemplar weighting function with $N = 2$ and $w_h(a_0, h_0), w_h(a_1, h_1) < 1$, and where $A \leq H$. Then, the system exhibits complementarity whenever:

$$\sqrt{H - A} \cdot \sqrt{\frac{1-p}{p} \cdot \left(\frac{1-b}{m} - (H - A) \right)} < |\delta_a - \delta_h|$$

PROOF. This lemma is a more specific version of Lemma 5, so we start with an intermediate result from that proof. Complementarity occurs whenever:

$$p \cdot w_h(a_1, h_1) \cdot (h_1 - a_1) + (1-p) \cdot w_h(a_2, h_2) \cdot (h_2 - a_2) < 0$$

For the exemplar combining rule, we have that:

$$w_h(a_i, h_i) = b - m \cdot (h_i - a_i)$$

Plugging in for the values of $w_h(a_1, h_1), w_h(a_2, h_2)$ gives:

$$\begin{aligned} & p \cdot (b - m \cdot (h_1 - a_1)) \cdot (h_1 - a_1) + (1-p) \cdot (b - m \cdot (h_2 - a_2)) \cdot (h_2 - a_2) < 0 \\ & b \cdot (p \cdot (h_1 - a_1) + (1-p) \cdot (h_2 - a_2)) - m(p \cdot (h_1 - a_1)^2 + (1-p) \cdot (h_2 - a_2)^2) < 0 \\ & b \cdot (p \cdot h_1 - p \cdot a_1 + (1-p) \cdot h_2 - (1-p) \cdot a_2) - m(p \cdot (h_1 - a_1)^2 + (1-p) \cdot (h_2 - a_2)^2) < 0 \\ & b \cdot (H - A) - m(p \cdot (h_1 - a_1)^2 + (1-p) \cdot (h_2 - a_2)^2) < 0 \end{aligned}$$

We can analyze the term with the m coefficient by plugging in for values of a_i, h_i .

$$\begin{aligned} & p \cdot (h_1 - a_1)^2 + (1-p) \cdot (h_2 - a_2)^2 \\ & = p \cdot (H + \delta_h - A - \delta_a)^2 + (1-p) \cdot \left(H - \frac{p}{1-p} \cdot \delta_h - A + \frac{p}{1-p} \cdot \delta_a \right)^2 \end{aligned}$$

We expand out each to get:

$$\begin{aligned} &= p \cdot (H - A)^2 + p \cdot (\delta_h - \delta_a)^2 + 2p \cdot (H - A) \cdot (\delta_h - \delta_a) + (1 - p) \cdot (H - A)^2 \\ &\quad + (1 - p) \cdot \frac{p^2}{(1 - p)^2} \cdot (\delta_a - \delta_h)^2 + 2 \cdot (1 - p) \cdot (H - A) \cdot \frac{p}{1 - p} \cdot (\delta_a - \delta_h) \end{aligned}$$

We note that two of the terms cancel:

$$\begin{aligned} &2p \cdot (H - A) \cdot (\delta_h - \delta_a) + 2 \cdot (1 - p) \cdot (H - A) \cdot \frac{p}{1 - p} \cdot (\delta_a - \delta_h) \\ &= 2p \cdot (H - A) \cdot (\delta_h - \delta_a) + 2 \cdot p \cdot (H - A) \cdot (\delta_a - \delta_h) \\ &= 0 \end{aligned}$$

Next, we can simplify the other terms:

$$\begin{aligned} &= p \cdot (H - A)^2 + p \cdot (\delta_h - \delta_a)^2 + (1 - p) \cdot (H - A)^2 + (1 - p) \cdot \frac{p^2}{(1 - p)^2} \cdot (\delta_a - \delta_h)^2 \\ &= (H - A)^2 + p(\delta_h - \delta_a)^2 + \frac{p^2}{1 - p} \cdot (\delta_a - \delta_h)^2 \\ &= (H - A)^2 + (\delta_h - \delta_a)^2 \cdot \left(p + \frac{p^2}{1 - p} \right) \\ &= (H - A)^2 + (\delta_h - \delta_a)^2 \cdot \frac{p}{1 - p} \end{aligned}$$

where we have used that $p + \frac{p^2}{1 - p} = \frac{p - p^2 + p^2}{1 - p} = \frac{p}{1 - p}$. We can combine this with the inequality we were analyzing earlier to get:

$$\begin{aligned} (1 - b) \cdot (H - A) - m \cdot \left((H - A)^2 + (\delta_h - \delta_a)^2 \cdot \frac{p}{1 - p} \right) &< 0 \\ (1 - b) \cdot (H - A) - m \cdot (H - A)^2 &< m \cdot (\delta_h - \delta_a)^2 \cdot \frac{p}{1 - p} \\ \frac{1 - p}{p} \left(\frac{1 - b}{m} \cdot (H - A) - (H - A)^2 \right) &< (\delta_h - \delta_a)^2 \end{aligned}$$

where we have used the assumption that $0 < p < 1$ and $m > 0$. If the lefthand side of the inequality is positive, we can take the square root of both sides to get:

$$\sqrt{H - A} \cdot \sqrt{\frac{1 - p}{p} \left(\frac{1 - b}{m} - (H - A) \right)} < |\delta_h - \delta_a|$$

Finally, we will show that the term under the square root, give the assumptions of this lemma. The term under the square root is *negative* if:

$$\frac{1 - b}{m} - (H - A) < 0$$

or:

$$1 < b - m \cdot (H - A) \tag{2}$$

We will show that, if this happens, we must have $w_h(a_1, h_1) > 1$ or $w_h(a_2, h_2) > 0$ (either of which violate the assumptions of this lemma). For this combining rule,

$$w_h(a_1, h_1) = b - m \cdot (h_1 - a_1) = b - m \cdot (H + \delta_h - A - \delta_a) = b - m \cdot (H - A) + m \cdot (\delta_h - \delta_a)$$

In the event that $\delta_h > \delta_a$, then the above equation is greater than Equation 2. Therefore, if Equation 2 is greater than 1, then $w_h(a_1, h_1)$ is also greater than 1, which violates the assumptions of the lemma. Similarly, we can write:

$$w_h(a_2, h_2) = b - m \cdot (h_2 - a_2) = b - m \cdot \left(H - \delta_h \cdot \frac{p}{1 - p} - A + \delta_a \cdot \frac{p}{1 - p} \right) = b - m \cdot (H - A) + m \cdot \frac{p}{1 - p} \cdot (\delta_a - \delta_h)$$

In the event that $\delta_h < \delta_a$, then the above equation is greater than Equation 2. Therefore, if Equation 2 is greater than 1, then $w_h(a_2, h_2)$ is also greater than 1, which violates the assumptions of the lemma. \square