Uncertainty as a Form of Transparency: Measuring, Communicating, and Using Uncertainty

Umang Bhatt^{1,2}, Javier Antorán², Yunfeng Zhang³, Q. Vera Liao³, Prasanna Sattigeri³,
Riccardo Fogliato^{1,4}, Gabrielle Melançon⁵, Ranganath Krishnan⁶, Jason Stanley⁵, Omesh Tickoo⁶,
Lama Nachman⁶, Rumi Chunara⁷, Madhu Srikumar¹, Adrian Weller^{2,8}, Alice Xiang^{1,9}

¹Partnership on AI, ²University of Cambridge, ³IBM Research, ⁴Carnegie Mellon University, ⁵Element AI, ⁶Intel Labs,
⁷New York University, ⁸The Alan Turing Institute, ⁹Sony AI

ABSTRACT

Algorithmic transparency entails exposing system properties to various stakeholders for purposes that include understanding, improving, and contesting predictions. Until now, most research into algorithmic transparency has predominantly focused on explainability. Explainability attempts to provide reasons for a machine learning model's behavior to stakeholders. However, understanding a model's specific behavior alone might not be enough for stakeholders to gauge whether the model is wrong or lacks sufficient knowledge to solve the task at hand. In this paper, we argue for considering a complementary form of transparency by estimating and communicating the uncertainty associated with model predictions. First, we discuss methods for assessing uncertainty. Then, we characterize how uncertainty can be used to mitigate model unfairness, augment decision-making, and build trustworthy systems. Finally, we outline methods for displaying uncertainty to stakeholders and recommend how to collect information required for incorporating uncertainty into existing ML pipelines. This work constitutes an interdisciplinary review drawn from literature spanning machine learning, visualization/HCI, design, decision-making, and fairness. We aim to encourage researchers and practitioners to measure, communicate, and use uncertainty as a form of transparency.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; *Uncertainty quantification*; • **Human-centered computing** → *Visualization*;

KEYWORDS

uncertainty; transparency; machine learning; visualization

ACM Reference Format:

Umang Bhatt^{1,2}, Javier Antorán², Yunfeng Zhang³, Q. Vera Liao³, Prasanna Sattigeri³, Riccardo Fogliato^{1,4}, Gabrielle Melançon⁵, Ranganath Krishnan⁶, Jason Stanley⁵, Omesh Tickoo⁶, Lama Nachman⁶, Rumi Chunara⁷, Madhu Srikumar¹, Adrian Weller^{2,8}, Alice Xiang^{1,9}. 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 (AIES '21)*, May 19–21, 2021, Virtual Event, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3461702.3462571>

This work is licensed under a Creative Commons Attribution International 4.0 License.

AIES '21, May 19–21, 2021, Virtual Event, USA.
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8473-5/21/05.
<https://doi.org/10.1145/3461702.3462571>

1 INTRODUCTION

Transparency in machine learning (ML) encompasses a wide variety of efforts to provide stakeholders, such as model developers and end users, with relevant information about how a ML model works [8, 96, 128]. One form of transparency is procedural transparency, which provides information about model development (e.g., code release, model cards, dataset details) [3, 39, 86, 104]. Another form is algorithmic transparency, which exposes information about a model's behavior to various stakeholders [67, 108, 120]. The ML community has mostly considered explainability, which attempts to provide reasoning for a model's behavior to stakeholders, as a proxy for algorithmic transparency [78]. With this work, we seek to encourage researchers to study uncertainty as an alternative form of algorithmic transparency and practitioners to communicate uncertainty estimates to stakeholders. Uncertainty is crucial yet often overlooked in the context of ML-assisted, or automated, decision-making [66, 111]. If well-calibrated and effectively communicated, uncertainty can help stakeholders understand when they should trust model predictions and help developers address fairness issues in their models [123, 138].

Uncertainty refers to our lack of knowledge about some outcome. As such, uncertainty will be characterized differently depending on the task at hand. In regression tasks, uncertainty is often expressed in terms of error bars, also known as confidence intervals. For example, when predicting the number of crimes in a given city, we could report that the number of predicted crimes is 943 ± 10 , where “ ± 10 ” represents a 95% *confidence interval* (capturing two standard deviations on either side of the central, mean estimate). The smaller the interval, the more certain the model. In classification tasks, probabilities are often used to capture how confident a model is in a specific prediction. For example, a classification model may decide that a person is at a high risk for developing diabetes given a prediction of 85% chance of diabetes. Broadly, uncertainty in data-driven decision-making systems may stem from different sources and thus communicate different information to stakeholders [36, 52]. Aleatoric uncertainty is induced by inherent randomness (or noise) in the quantity to predict given inputs. Epistemic uncertainty can arise due to lack of sufficient data to learn our model precisely.

Why do we care?

We posit uncertainty can be useful for obtaining fairer models, improving decision-making, and building trust in automated systems. Throughout this work, we will use the following cancer diagnostics scenario for illustrative purposes: Suppose we are tasked with diagnosing individuals as having breast cancer or not, as in [25, 31].

Given categorical and continuous characteristics about an individual (medical test results, family medical history, etc.), we estimate the probability of an individual having breast cancer. We can then apply a threshold to classify them into high- or low-risk groups. Specifically, we have been tasked with building ML-powered tools to help three distinct audiences: doctors, who will be assisted in making diagnoses; patients, who will be helped to understand their diagnoses; and review board members, who will be aided in reviewing doctors' decisions across many hospitals. Throughout the paper, we will refer back to the scenario above to discuss how uncertainty may arise in the design of an ML model and why, when well-calibrated and well-communicated, uncertainty can act as a useful form of transparency for stakeholders. We now briefly examine how ignoring uncertainty can be detrimental to transparency in three different use cases with the help of our scenario.

Fairness: Uncertainty, if not properly quantified and considered in model development, can endanger efforts to assess model fairness. Developers often aim to assess fairness to mitigate or prevent unwanted biases. Breast cancer is more common for older patients, potentially leading to datasets where young age groups are under-represented. If our breast cancer diagnostic tool is trained on data presenting such a bias, the model might be under-specified and present larger error rates for younger patients. Such dataset bias will manifest itself as epistemic uncertainty. In Section 3, we detail ways uncertainty interacts with bias in the data collection/modeling stages and how such biases can be mitigated by ML practitioners.

Decision-making: Treating all model predictions the same, independent of their uncertainty, can lead decision-makers to over-rely on their models in cases where they produce spurious outputs or to under-rely on them when their predictions are accurate. As such, a doctor could make better use of an automated decision-making system by observing its uncertainty estimates before leveraging the model's output in making a diagnosis. In Section 3, we draw upon the literature of judgment and decision-making (JDM) to discuss the potential implications of showing uncertainty estimates to end users of ML models.

Trust in Automation: Well-calibrated and well-communicated uncertainty can be seen as a sign of model trustworthiness, in turn improving model adoption and user experience. However, when communicated inadequately, uncertainty estimates can be incomprehensible to stakeholders and perceived negatively, thus spawning confusion and impairing trust formation. Suppose our model's predictions of a patients' breast cancer stage are accompanied by 95% confidence intervals. Due to the pervasiveness of breast cancer, a large number of healthy patients might fall within these errorbars, suggesting to doctors that they could have early stage breast cancer. In this situation, the doctors may choose to always override the model's output, the model's seeming imprecision resulting in an erosion of the doctors' trust. For this task, errors bars representing the predictive interquartile range might have been more appropriate. The optimal way to communicate uncertainty will depend on the task at hand. In Section 3, we review prior work on how users form trust in automated systems and discuss the potential impact of uncertainty on this trust.

This work is structured as follows. In Section 2, we review possible sources of uncertainty and methods for uncertainty quantification. In Section 3, we describe how uncertainty can be leveraged in

Figure 1: Top: The same prediction (no cancer) is made in two different hypothetical cancer diagnosis scenarios. However, our model is much more confident in the first. This is reflected in the predictive entropy for each case (dashed red line denotes the maximum entropy for 3-way classification). Bottom: In a regression task, a predictive distribution contains rich information about our model's predictions (modes, tails, etc). We summarise predictive distributions with means and error-bars (here standard deviations).

each use case mentioned above. Finally, in Section 4, we discuss how to communicate uncertainty effectively. Therein, we also discuss how to take a user-centered approach to collecting requirements for uncertainty quantification and communication.

2 MEASURING UNCERTAINTY

In ML, we use the term uncertainty to refer to our lack of knowledge about some outcome of interest. We use the tools of probability to reason about and quantify uncertainty. The Bayesian school of thought interprets probabilities as subjective degrees of belief in an outcome of interest occurring [81]. For frequentists, probabilities reflect how often we would observe the outcome if we were to repeat our observation multiple times [10, 99]. Fortunately for end-users, uncertainty from Bayesian and frequentist methods conveys similar information in practice [131], and can almost always be treated interchangeably in downstream tasks.

2.1 Metrics for Uncertainty

The metrics used to communicate uncertainty vary between research communities and application domains. Predictive distributions, shown in Figure 1, tell us about our models' degree of belief in every possible prediction. Despite containing a lot of information (prediction modes, tails, etc.), a full predictive distribution may not always be desirable. In our cancer diagnostic scenario, we may want our automated system to abstain from making a prediction and instead request the assistance of a medical professional when its uncertainty is above a certain threshold. When deferring to a human expert, it may not matter to us whether the system believes the patient has cancer or not, just how uncertain it is. For this reason, **summary statistics** of the predictive distribution are often

Table 1: Commonly used metrics for the quantification and communication of uncertainty.

	Full Information	Summary Statistics
Regression	Predictive Density	Predictive Variance, Percentile (or quantile) Confidence Intervals
Classification	Predictive Probabilities	Predictive Entropy, Expected Entropy, Mutual Information, Variation Ratio

used to convey information about uncertainty. For classification, the predictive distribution is composed of class probabilities. These intuitively communicate our degree of belief in an outcome. On the other hand, predictive entropy decouples our predictions from their uncertainty, only telling us about the latter. For regression, the predictive distribution is often summarized by a predictive mean together with error bars (written $\pm\sigma$). These commonly reflect the standard deviation or some percentiles of the predictive distribution. In Section 4.3, we discuss how the best choice of uncertainty metric is use case dependent. We show common summary statistics for uncertainty in Table 1 and elaborate in the supplementary material.

2.2 The Different Sources of Uncertainty

While there can be many sources of uncertainty [127], herein we focus on those that we can quantify in ML models: *aleatoric uncertainty* (also known as indirect uncertainty) and *epistemic uncertainty* (also known as direct uncertainty) [26, 27, 36].

Aleatoric uncertainty stems from noise, or class overlap, in our data. Noise in the data is a consequence of unaccounted-for factors that introduce variability in the inputs or targets. Examples of this could be background noise in a signal detection scenario or the imperfect reliability of a medical test in our cancer diagnosis scenario. Aleatoric uncertainty is also known as irreducible uncertainty: it cannot be decreased by observing more data. If we wish to reduce aleatoric uncertainty, we may need to leverage different sources of data, e.g., switching to a more reliable clinical test. In practice, most ML models account for aleatoric uncertainty through the specification of a noise model or likelihood function. A homoscedastic noise model makes the assumption that all of the input space is equally noisy, $y = f(x) + \epsilon$; $\epsilon \sim p(\epsilon)$. However, this may not always be true. Returning to our medical scenario, consider using results from a clinical test which produces few false negatives but many false positives as an input to our model. A heteroscedastic noise assumption allows us to express aleatoric uncertainty as a function of our inputs $y = f(x) + \epsilon$; $\epsilon \sim p(\epsilon|x)$. Perhaps the most commonly used heteroscedastic noise models in ML are those induced by the sigmoid or softmax output layers. These enable almost any classification model to express aleatoric uncertainty: see the supplementary material for more details.

Epistemic uncertainty stems from a lack of knowledge about which function best explains the data we have observed. There are two reasons why epistemic uncertainty may arise. Consider a scenario in which we employ a very complex model relative to the amount of training data available. We will be unable to properly constrain our model's parameters. This means that, out of all the possible functions that our model can represent, we are unsure of which ones to choose. In this work, we refer to uncertainty about

a model's parameters as *model uncertainty*. We might also be uncertain of whether we picked the correct model class in the first place. Perhaps we are using a linear predictor but the phenomenon we are trying to predict is non-linear. We will refer to this as *model specification uncertainty* or *architecture uncertainty*. Epistemic uncertainty can be reduced by collecting more data in input regions where the training dataset was sparse. It is less common for ML models to capture epistemic uncertainty. Often, those that do are referred to as probabilistic models.

Given a probabilistic predictive model, aleatoric and epistemic uncertainties can be quantified separately, as described in the supplementary material. We depict them separately in Figure 2. Being aware of which regions of the input space present large aleatoric uncertainty can help ML practitioners identify issues in their data collection process. On the other hand, epistemic uncertainty tells us about which regions of input space we have yet to learn about. Thus, epistemic uncertainty is used to detect dataset shift [97], or adversarial inputs [134]. It is also used to guide methods that require exploration like active learning [53], continual learning [92], Bayesian optimisation [47], and reinforcement learning [58].

2.3 Methods to Quantify Uncertainty

Most ML approaches involve a noise model, thus capturing aleatoric uncertainty. However, few are able to express epistemic uncertainty. When we say that a method is able to quantify uncertainty, we are implicitly referring to those that capture both epistemic and aleatoric uncertainty. These methods can be broadly classified into two categories: Bayesian approaches [40, 90, 106] and Frequentist approaches [14, 70, 113].

Bayesian methods explicitly define a hypothesis space of plausible models *a priori* (before observing any data) and use deductive logic to update these priors given the observed data. In parametric models, like Bayesian Neural Networks (BNNs) [80, 91], this is most often done by treating model weights as random variables instead of single values, and assigning them a prior distribution $p(w)$. Given some observed data $D = \{y, x\}$, the conditional likelihood $p(y|x, w)$ tells us how well each weight setting w explains our observations. The likelihood is used to update the prior, yielding the posterior distribution over the weights $p(w|D)$:

$$p(w|D) = \frac{p(y|x, w) p(w)}{\int p(y|x, w) p(w) dw} \quad (1)$$

Prediction for a test point x^* is made via marginalization: all possible weight configurations are considered with each configuration's prediction being weighed by that weights' posterior density. The disagreement among predictions from different plausible weight settings induces model (epistemic) uncertainty. The predictive posterior distribution:

$$p(y|x^*) = \int p(y|x^*, w) p(w|D) dw \quad (2)$$

captures both epistemic and aleatoric uncertainty.

Recently, the ML community has moved towards favoring NNs as their choice of model due to their flexibility and scalability to large amounts of data. Unfortunately, the more complicated the model, the more difficult it is to compute the exact posterior $p(w|D)$ and predictive $p(y|x^*)$ distributions. For NNs, it is analytically and

Figure 2: Uncertainty quantification and evaluation. Left three plots: A 15-element deep ensemble provides increased model (epistemic uncertainty) in data sparse regions. A homoscedastic Gaussian noise model provides aleatoric uncertainty matching the noise in the data. Both combine to produce a predictive distribution. Right: calibration plot for classification task. Each bar corresponds to a bin in which predictions are grouped. Their height corresponds to their proportion of correct predictions.

computationally intractable [46]. However, various approximations have been proposed. Among the most popular are variational inference [12, 37, 48] and stochastic gradient MCMC [21, 129, 136]. Methods that provide more faithful approximations, and thus more calibrated uncertainty estimates, tend to be more computationally intensive and scale worse to larger models. As a result, the best method will vary depending on the use case. Also worth mentioning are Bayesian non-parametrics, such as Gaussian Processes [106]. However, for brevity, we discuss these, with a broader range of Bayesian methods, in the supplementary material.

Frequentist methods do not specify a prior distribution over hypothesis. They exclusively consider how well the distribution over observations implied by each hypothesis matches the data. Here, uncertainty stems from how we expect our chosen hypothesis to change if we were to repeatedly sample different sets of data. Perhaps the most salient Frequentist technique is ensembling [28, 76]. This consists of training multiple models in different ways to obtain multiple plausible fits. At test time, the disagreement between ensemble elements' predictions yields model uncertainty, as shown in Figure 2. Currently, deep ensembles [70] are one of the best performing uncertainty quantification approaches for NNs [5], retaining calibration even under dataset shift [97]. Unfortunately, the computational cost involved with running multiple models at both train and test time also make ensembles one of the most expensive methods. There is a vast heterogeneous literature on frequentist uncertainty quantification, some of which is covered in the supplementary material.

2.4 Uncertainty Evaluation

Calibration is a form of quality assurance for uncertainty estimates. It is not enough to provide larger error bars when our model is more likely to make a mistake. Our predictive distribution must reflect the true distribution of our targets. Recall our cancer diagnosis scenario, where the system declines to make a prediction when uncertainty is above a threshold, and a doctor is queried instead. Due to the doctor's time being limited, we might design our system such that it only declines to make a prediction if it estimates there is a probability greater than 0.05 of the prediction being wrong. If instead of being well-calibrated, our system is underconfident, we

would over-query the doctor in situations where the AI's prediction is correct. Overconfidence would result in taking action on unreliable predictions: delivering unnecessary treatment or abstaining from providing necessary treatment.

Calibration is orthogonal to accuracy. A model with a predictive distribution that matches the marginal distribution of the targets $p(y|x) = p(y) \forall x$ would be perfectly calibrated but would not provide any useful predictions. Thus, calibration is usually measured in tandem with accuracy through either a general fidelity metric (most often chosen to be a proper scoring rule [41]) which subsumes both objectives, or using two separate metrics. The most common metrics of the former category are *negative log-likelihood* (NLL) and *Brier score* [15]. Of the latter type, *Expected calibration error* (ECE) [88], illustrated in Figure 2, is popularly used in classification scenarios. We refer to the supplementary material for a detailed discussion of calibration metrics.

Recently, Antoran et al. [2] introduced CLUE, a method to help identify which input features are responsible for models' uncertainty. This class of approach opens avenues for stakeholders to qualitatively evaluate if their models' reasons for uncertainty align with their own intuitions [73]. A transparent ML model requires both well-calibrated uncertainty estimates and an effective way to communicate them to stakeholders. If uncertainty is not well-calibrated, our model cannot be transparent since its uncertainty estimates provide false information. Thus calibration is a precursor to using uncertainty as a form of transparency.

Takeaways:

- (1) Well-calibrated uncertainty is a key factor in transparent data-driven decision-making.
- (2) Uncertainty can stem from both the difficulty of the problem we are trying to solve and the modelling choices we are using to solve it.
- (3) The more complex the problem, the more costly it is to obtain calibrated uncertainty estimates.

3 USING UNCERTAINTY

In this section, we discuss the use of uncertainty based on the three motivations presented in the Introduction: fairness, decision

making and trust in automation. These are not mutually exclusive; as such, they could all be leveraged simultaneously depending on the use case. In Section 4.3, we discuss the importance of gathering requirements on how stakeholders, both internal and external, will use uncertainty estimates.

3.1 Uncertainty and Fairness

An unfair ML model is one that exhibits unwanted bias towards or against one or more target classes. Here, we discuss possible ways in which bias can appear as a consequence of unaccounted-for uncertainty and potential approaches to mitigate it. The supplementary material contains further discussion on the interplay of model specification uncertainty and fairness.

Measurement bias, also known as feature noise, is a case of aleatoric uncertainty (defined in Section 2). It arises when one or more of the features in the data only represent a proxy for the features that would have ideally been measured and can be mitigated by a properly specified noise model. We describe the potential effects of noise on inputs and targets.

Commencing with the former, we discuss noise in the sensitive attribute. In contexts such as medical diagnosis, information on race and ethnicity of patients may not be collected [20]. Alternatively, in some contexts, survey participants may have incentives to misreport characteristics such as their religious or political affiliations. The experimental results of Gupta et al. [44] have shown that enforcing fairness constraints on a noisy sensitive attribute, without assumptions on the structure of the noise, is not guaranteed to lead to any improvement in the fairness properties of our model. Successive papers have explored which assumptions are needed in order to obtain such guarantees.

The “mutually contaminated learning model” assumes the noise only depends on the true unobserved value of the sensitive attribute [112]. Here, measures of demographic parity and equalized odds computed on the observed data are equal to the true metrics up to a scaling factor, which is proportional to the value of the noise [71]; if the noise rates are known, then the true metrics can be directly estimated. When information on the sensitive attribute is unavailable (e.g., information on gender was not collected), but can be predicted from an auxiliary dataset, disparity measures are generally unidentifiable [20, 63].

Second, we discuss noise in the targets (i.e., aleatoric uncertainty in the labels). This source of bias has attracted less attention in the fairness community despite being similarly pervasive. Obermeyer et al. [94] found that when medical expenses were used as a proxy for illness, their algorithm severely underpredicted the prevalence of the illness for the population of Black patients.

Similarly to noise in the sensitive attribute, Blum and Stangl [11], Jiang and Nachum [59] show that fairness constraints are guaranteed to improve the properties of a predictor only under appropriate assumptions on the label noise. Indeed, Fogliato et al. [35] show that even a small amount of label noise can greatly impact the assessment of the fairness properties of the model. There is a small but growing body of work on appropriate noise model specification for bias mitigation. For example, the problem of auditing algorithms for fairness in the positive and unlabeled learning setting (i.e., noise is one-sided) has been studied by Fogliato et al. [35].

Representation bias stems from how we define the population under study and how we sample our observations from said population [121]. Representation bias is epistemic in nature and thus may be reduced by collecting additional, potentially more diverse, data. Models trained in the presence of representation bias could exhibit unwanted bias towards an under-represented group. For example, the differential performance of gender classification systems across racial groups may be due to under-representation of Black individuals in the sampled population [18]. Similarly, historical over-policing of some communities has unavoidable impacts on the sampling of the data used to train models for predictive policing [32, 79].

In our cancer diagnostics scenario, the data used to train a model for a specific hospital should closely reflect the demographic statistics of that hospital’s patients, i.e. ensuring proper representation in terms of age groups, cancer stages, etc. As such, ideally, the data used to train our hospital’s models, would stem from its own patients, and not those of another hospital where population statistics may be different [95].

In order to tackle representation bias, we must first detect that such bias exists. This can be done by leveraging uncertainty as a form of transparency. ML practitioners building a probabilistic model could check for representation bias by ensuring that their validation datasets closely matches the distribution expected at test time (in deployment). Their model presenting large epistemic uncertainty on this validation set would indicate the existence of representation bias in the training data. The practitioner could then identify which subgroups are unequally represented between their training and validation sets [2]. Finally, they could leverage this knowledge to improve the data collection procedure. Alternatively, we could consider designing a system that automatically checks for epistemic inputs in deployment, thus raising an alert if the statistics of the population to which it is exposed change over time.

Unfortunately, sample size may still represent an issue. It is not always simple to collect more diverse data. Additionally, in many domains, the existing sample size may also not be large enough to assess the existence of biases [33]. We refer to Mehrabi et al. [83] for a more detailed breakdown of the potential sources of representation bias.

3.2 Uncertainty and Decision-making

Depending on the context, ML systems can be used to support human decision-making to various degrees or even substitute it outright. Crucially, in all types of data-driven decision-making, uncertainty plays a key role. Uncertainty enables stakeholders to better understand their decision-making system and thus better combine the systems’ output with their own judgement.

First, we consider fully automated decision-making without human intervention. Decision theory [9] allows us to combine probabilistic predictions $p(y|x)$ about an outcome of interest y given an input x with the cost of taking each action a given each outcome $L(a|y)$, leading us to make an optimal decision a^* under uncertainty:

$$a^* = \arg \min_a \int L(a|y)p(y|x) dy. \quad (3)$$

Recall our medical scenario; our model is tasked with providing breast cancer diagnoses $p(\text{cancer}|x)$. Here, we might quantify the

cost of a false negative $L(\text{report} = \text{healthy}|\text{cancer})$ as 100 times greater than that of a false positive $L(\text{report} = \text{cancer}|\text{healthy})$. Armed with this information, the optimal diagnosis we report to the patient can simply be obtained by plugging into Equation 3. More reasonably, we may decide to have a medical professional intervene in a small number of situations where our system is likely to be wrong. This is known as a “reject option” [87]. In turn, decision theory can be used to appropriately place a rejection threshold.

In both the fully automated and “reject option” systems, uncertainty increases the transparency of the decision-making system. When a decision is made automatically, a model’s uncertainty can be leveraged post-hoc to help elucidate why that specific decision was made. When the reject option is triggered, learning about the models’ uncertainty (in the form of a predictive distribution or some summary statistic) may aid the human expert in her decision.

We now consider situations of the latter sort: a human expert is tasked with making a decision while being aided by an ML system. Here, uncertainty plays a key role, as the end user has to weight how much they should trust the model’s output. This question corresponds to prototypical tasks¹ studied in the Judgment and Decision-Making (JDM) literature, i.e., “action threshold decision” and “multi-option choices” [34]. The ML literature has only just begun to examine how uncertainty estimates affect user interactions with ML models and task performance [4, 138]. However, we highlight relevant conclusions from the JDM literature that pertain to using uncertainty estimates in decision-making. Prospect Theory suggests that uncertainty (or risk) is not considered independently but together with the expected outcome [61, 125]. This relationship is non-linear and asymmetrical. A certain prediction of a large loss is often perceived more negatively than an uncertain prediction of a small loss. When risk is presented positively, e.g., a 95% chance of the model being correct, people tend to be risk-averse; when it is presented negatively they are risk-seeking. Additionally, as the stake of the decision-outcome increases, tolerance for uncertainty seems to decrease at a superlinear rate [126]. Of course, differences among individuals’ tolerances for uncertainty also play an important role [45, 85, 102, 114].

Assessment of risk, however, also depends on how uncertainty is communicated and perceived. Both lay people and experts rely on mental shortcuts, or heuristics, to *interpret* uncertainty [124]. This could lead to biased appraisals of uncertainty even if model outputs are well-calibrated. We discuss this in Section 4. Stowers et al. [117] find that communicating and visualizing uncertainty information to operators of unmanned vehicles helped improve human-AI team performance; however, they note that their findings may not generalize to other tasks and contexts.

To our knowledge, the empirical understanding of how decision-makers make use of aleatoric versus epistemic uncertainty is limited. Furthermore, the JDM literature has mostly focused on discrete outcomes. There is not a good understanding of how people perceive uncertainty over continuous outcomes (e.g., errorbars). We judge these to be important gaps in the human-ML interaction literature.

3.3 Uncertainty and Trust Formation

While trust could be implicit in a decision to rely on a model’s suggestion, the communication of a model’s uncertainty can also affect people’s general trust in an ML system. At a high level, communicating uncertainty is a form of transparency that can help gain people’s trust. The relationship between uncertainty estimates and trust in automation is a relatively unexplored idea. Here, we explore the relevant literature on the underlying construct of trust and the processes of trust formation with the goal of painting a more complex picture of how stakeholders might use uncertainty estimates to form trust in an ML system.

While not limited to ML systems, the HCI and Human Factors communities have a long history of studying trust in automation [50, 68, 72]. These models of trust often build on Mayer et al. [82]’s classic ABI (Ability, Benevolence, Integrity) model of interpersonal trust. Taking into account some fundamental differences between interpersonal trust and trust in automation, Lee and See [72] adapted the ABI model to trust in automated systems. They highlight three underlying dimensions: 1) competence of the system within the targeted domain, 2) intention of developers: the extent to which they are perceived to want to do good to the trustor, and 3) predictability/understandability: the extent to which the system consistently operates according to a set of principles that the trustor finds acceptable.

We speculate that communicating uncertainty estimates could be relevant to all three dimensions. If a model’s uncertainty is perceived to be too large or miscalibrated, it may harm the model’s perceived competence. If uncertainty is not communicated or intentionally mis-communicated, users or stakeholders might form a negative opinion on the intention of developers. If a model shows uncertainty that could not be understood or expected, it will be negatively perceived in predictability.

To anticipate how uncertainty estimates, and ways to communicate them, could impact stakeholder trust, we highlight existing “process models” on how people develop trust. Rooted in information-processing and decision-making theories [19, 60, 101], “process models” differentiate between an analytic (or systematic) process of trust formation and an affective (or heuristic) process of trust formation [72, 84, 118].

The former involves rational evaluation of a trustee’s characteristics; systematic trust formation in an ML system could be facilitated by providing detailed probabilistic uncertainty estimates. The latter process relies on feelings or heuristics to form a quick judgment to trust or not; when lacking either the technical ability or motivation to perform an analytic evaluation, people rely more on the affective or heuristic route [101, 119]. For example, for some users the mere presence of uncertainty information could signal that the ML engineers are being transparent and sincere, enhancing their trust [54]. For others, uncertainty could invoke negative heuristics [127]. Furthermore, prior work suggests that the style in which uncertainty estimates are communicated is highly relevant to how these are perceived [98]. We elaborate on communication methods for uncertainty in Section 4.

¹Note that social scientists often use the terms “risk” and “uncertainty” differently than in this text [65, 105]

Lastly, we highlight a non-trivial point that the goal of presenting uncertainty estimates to stakeholders should support forming *appropriate* trust, rather than blindly enhancing trust. A well-measured and well-communicated uncertainty estimate should not only facilitate the *calibration* of overall trust on a system, but also the *resolution* of trust [23, 72]. The latter referring to how precisely the judgment of trust could differentiate types of model capabilities in decision-making.

Based on relevant work from HCI, JDM, and ML, we argue that leveraging uncertainty as a form of transparency can be helpful for trust formation. However, how uncertainty estimates are processed for trust formation and what kind of affective impact uncertainty could invoke remain open questions and merit future research.

Takeaways:

- (1) Uncertainty can manifest as unfairness in the form of noisy features/targets (aleatoric) or in the sampling procedure (epistemic). Being aware of uncertainty can allow ML practitioners to mitigate these issues.
- (2) Being aware of uncertainty allows stakeholders to make better use of ML-assisted decision-making systems.
- (3) Delivering uncertainty estimates to stakeholders can enhance transparency by ensuring trust formation.

4 COMMUNICATING UNCERTAINTY

Treating uncertainty as a form of transparency requires accurately communicating it to stakeholders. However, even well-calibrated uncertainty estimates could be perceived inaccurately by people because (a) they have varying levels of understanding about probability and statistics, and (b) human perception of uncertainty quantities is often biased by decision-making heuristics. In this section, we will review some of the issues that hinder people's understanding of uncertainty estimates and will discuss how various communication methods may help address these issues. We will first describe how to communicate uncertainty in the form of confidence or prediction probabilities for classification tasks, and then more broadly in the form of ranges, confidence intervals, or full distributions. In the supplementary material, we dive into a case study on the utility of uncertainty communication during the COVID-19 pandemic.

4.1 Issues in Understanding Uncertainty

Many application domains involve communicating uncertainty estimates to the general public to help them make decisions, e.g., weather forecasting, transit information delivery [64], medical diagnosis and interventions [102]. One key issue in these applications is that a great deal of their audience may not have the numeracy skills required to interpret uncertainty correctly. In a survey Galesic and Garcia-Retamero [38] conducted in 2010 on statistical numeracy across the US and Germany, they found that many people do not understand relatively simple statements that involve statistics concepts. For example, 20% of the German and US participants could not say "which of the following numbers represents the biggest risk of getting a disease: 1%, 5%, or 10%," and almost 30% could not answer whether 1 in 10, 1 in 100, or 1 in 1000 represents the largest risk. Another study found that people's numeracy skills significantly affect how well they comprehend risks [139]. Many of

the aforementioned decision-making scenarios involve high-stakes decisions, so it is vital to find alternative ways to communicate uncertainty estimates to people with low numeracy skills.

Besides numeracy skills, research shows that humans in general suffer from a variety of cognitive biases, some of which hinder our understanding of uncertainty [60, 107, 115]. One is called ratio bias, which refers to the phenomenon where people sometimes believe a ratio with a big numerator is larger than an equivalent ratio with a small numerator. For example, people may see 10/100 as a larger odds of having breast cancer than 1/10. This same phenomenon is sometimes manifested as an underweighting of the denominator, e.g. believing 9/11 is smaller than 10/13. This is also called denominator neglect.

In addition to ratio biases, people's perception of probabilities is also distorted in that they tend to underweight high probabilities while overweighting low probabilities. This distortion prevents people from making optimal decisions. Zhang and Maloney [135] showed that when people are asked to estimate probabilities or frequencies of events based on memory or visual observations, their estimates are distorted in a way that follows a log-odds transformation of the true probabilities. Research also found that this bias occurs when people are asked to make decisions under risk and that their decisions imply such distortions [125, 137]. Therefore, when communicating probabilities, we need to be aware that people's perception of high risks may be lower than the actual risk, while that of low risks may be higher than actual.

A different kind of cognitive bias that impacts people's perception of uncertainty is framing [60]. Framing has to do with how information is contextualized. Typically, people prefer options with positive framing (e.g., a 80% likelihood of surviving breast cancer) than an equivalent option with negative framing (e.g., a 20% likelihood of dying from breast cancer). This bias has an effect on how people perceive uncertainty information. A remedy for this bias is to always describe the uncertainty of both positive and negative outcomes, rather than relying on the audience to infer what's left out of the description.

4.2 Communication Methods

Choosing the right communication methods can address some of the above issues. Hullman et al. [55] review methods for evaluating the success of uncertainty visualization. van der Bles et al. [127] categorize the different ways of expressing uncertainty into nine groups with increasing precision, from explicitly denying that uncertainty exists to displaying a full probability distribution. While high-precision communication methods help experts understand the full scale of the uncertainty of the ML models, low precision methods can suffice for lay people, who may have potentially low numeracy skills. We now focus on the pros and cons of the four more precise methods of communicating uncertainty: 1) describing the degree of uncertainty using a predefined categorization, 2) describing a numerical range, 3) showing a summary of a distribution, and 4) showing a full probability distribution. The first two methods can be communicated verbally, while the last two often require visualizations.

A predefined, ordered categorization of uncertainty and risk levels reduces the cognitive effort needed to comprehend uncertainty

Figure 3: Examples of uncertainty visualizations.

estimates, and therefore is particularly likely to help people with low numeracy skills [100]. A great example of how to appropriately use this technique is the GRADE guidelines, which introduce a four-category system, from high to very low, to rate the quality of evidence for medical treatments [7]. GRADE has provided definitions for each category and a detailed description of the aspects of studies to evaluate for constructing quality ratings. Uncertainty ratings are also frequently used by financial agencies to communicate the overall risks associated with an investment instrument [29].

The main drawback of communicating uncertainty via predefined categories is that the audience, especially non-experts, might not be aware of or even misinterpret the threshold criteria of the categories. Many studies have shown that although individuals have internally consistent interpretation of words for describing probabilities (e.g., likely, probably), these interpretations can vary substantially from one person to another [17, 22, 74]. Recently, Budescu et al. [16] investigated how the general public interpret the uncertainty information in the climate change report published by the Intergovernmental Panel on Climate Change (IPCC). They found that people generally interpreted the IPCC's categorical description of probabilities as less likely than the IPCC intended. For example, people took the word "very likely" as indicating a probability of around 60%, whereas the IPCC's guideline specifies that it indicates a greater than 90% probability. To avoid such misinterpretation, categorical and numerical forms of uncertainty can be communicated when possible.

Though numbers and numerical ranges are more precise than categorical scales in communicating uncertainty, as discussed earlier, they are harder to understand for people with low numeracy and can induce ratio biases. However, a few techniques can be used to remediate these problems. First, to overcome the adverse effect of denominator neglect, it is important to present ratios with the same denominator so that they can be compared with just the numerator [116]. Denominators that are powers of 10 are preferred since they are easier to compute. There is so far no conclusive findings on whether frequency format ("Out of every 100 patients, 2 are likely to be misdiagnosed") is easier to understand than ratios/percentages ("Out of every 100 customers, 2% are likely to be misdiagnosed"): people do seem to perceive risk probabilities represented in the frequency format as showing higher risk than those represented in the

percentage format [107]. Therefore, it is helpful to use a consistent format to represent probabilities. If the audience underestimates risk levels, the frequency format may be preferred.

Uncertainty estimates can also be represented with graphics, which have several advantages over verbal communication, such as attracting and holding the audience's attention, revealing trends or patterns in the data, and evoking mental mathematical operations [75]. Commonly used visualizations include pie charts, bar charts, and more recently, icon arrays (Figure 3a). Pie charts are particularly useful for conveying proportions since all possible outcomes are depicted explicitly. However, it is more difficult to make accurate comparisons with pie charts than with bar charts because pie charts use areas to represent probabilities. Icon arrays vividly depict part-to-whole relationship, and because they show the denominator explicitly, they can be used to overcome ratio biases.

So far, what we have discussed pertains mostly to conveying uncertainty of a binary event, which takes the form of a single number (probability), whereas the uncertainty of a continuous variable or model prediction takes the form of a distribution. This latter type of uncertainty estimate can be communicated either as a series of summary statistics about the distribution, or directly as the full distribution. As mentioned in Table 1, some commonly reported summary statistics include mean, median, confidence intervals, standard deviation, and quartiles [1]. These statistics are often depicted graphically as error bars and boxplots for univariate data, and two dimensional error bars and bagplots for bivariate data [109]. We describe these summary statistics and plots in detail in the supplementary material. Error bars only have a few graphical elements and are hence relatively easy to interpret. However, since they have represented a range of different statistics in the past, they are ambiguous if presented without explicit labeling [130]. Error bars may also overly emphasize the range within the bar [24]. Boxplots and bagplots are less popular in the mass media, and generally require some training to understand.

When presenting uncertainty about a single model prediction, it might be better to show the entire posterior predictive distribution, which can avoid over-emphasis of the within-bar range and allow more granular visual inferences. Popular visualizations of distributions are histograms, density plots, and violin plots (Hintze and Nelson [49] show multiple density plots side-by-side), but they

seem to be hard for an uninitiated audience to grasp. They are often mistaken as bivariate case-value plots in which the lines or bars denote values instead of frequencies [13]. More recently, Kay et al. [64] develop quantile dot plots to convey distributions (see Figure 3b for an example). These plots use stacked dots, where each dot represents a group of cases, to approximate the data frequency at particular values. This method translates the abstract concept of probability distribution into a set of discrete outcomes, which are more familiar concepts to people who have not been trained in statistics. Kay et al. [64]’s study showed that people could more accurately derive probability estimates from quantile dot plots than from density plots. Kruschke [69] defines a method that attempts to simultaneously convey aleatoric and epistemic uncertainty in a single plot by showing the predictive densities resulting from various samples the posterior distribution; however, the efficacy of such plots in practice is still unknown.

One very different approach to conveying uncertainty is to individually show random draws from the probability distribution as a series of animation frames called hypothetical outcome plots (HOP) [56]. Similar to the quantile dot plots, HOPs accommodate the *frequency* view of uncertainty very well. In addition, showing events individually does not add any new visual encodings (such as the length of the bar or height of the stacked dots) and thus requires no additional learning from the viewers. Hullman et al. [56] show that this visualization enabled people to make more accurate comparisons of two or three random variables than error bars and violin plots, presumably because statistical inference based on multiple distribution plots require special strategies while HOP does not. The drawbacks of HOP are: (a) it takes more time to show a representative sample of the distribution, and (b) it may incur high cognitive load since viewers need to mentally count and integrate frames. Nevertheless, because this method is easy to understand for people with low numeracy, similarly animated visualizations are frequently used in the mass media [6, 133]. Hofman et al. [51] found that 95% confidence intervals (inferential uncertainty) can be more misleading than 95% prediction intervals (outcome uncertainty): HOPs tend to help reduce error when laypeople are asked to estimate effect size. However, Kale et al. [62] note that simple heuristics may suffice instead of complex uncertainty visualizations: they find that the best visualizations for understanding effect size may not be the best for decision-making.

The above methods are designed to communicate uncertainty around a single quantity, so they need to be extended for visualizing uncertainty around a range of predictions, such as those in time-series forecasting. The simplest form of such visualization is a quantile plot, which uses lines to connect predictions at equal quantiles of the uncertainty distribution across the output range. When used in time-series forecasting, such plots are called cone-of-uncertainty plots (see Figure 3c), in which the cone enlarges over time, indicating increasingly uncertain predictions. Gradient plots, or fan charts (see Figure 3d) in the context of time series forecasting, can be used to show more granular changes in uncertainty, but they require extra visual encoding that may not be easily understood by the viewer. In contrast, spaghetti plots simply represent each model’s predictions as one line, while uncertainty can be inferred from the closeness of the lines. However, they might put too much emphasis on the lines themselves and de-emphasize the range of

the model predictions. Lastly, HOP can also be used to show uncertainty estimates over a range of predictions by showing each model’s predictions in an animation frame [130].

4.3 Uncertainty Requirements

Familiarity with the findings above will be helpful for teams building uncertainty into ML workflows. Yet, none of these findings should be treated as conclusive when it comes to predicting how usable a given expression of uncertainty will be for different types of users facing different kinds of constraints in real-world settings. Instead, findings from the literature should be treated as fertile ground for generating hypotheses in need of testing with real users, ideally engaged in the concrete tasks of their typical workflow, and ideally doing so in real-world settings.

It is important to recognize just how diverse individual users are, and how different their social contexts can be. In our cancer diagnostic scenario, the needs and constraints of a doctor making a time-pressed and high-stakes medical decision using an ML-powered tool will likely be very different from those of a patient attempting to understand their diagnosis, and different again from those of an ML engineer reviewing model output in search of strategies for model improvement. Furthermore, if we zoom in on any one of these user populations, we still typically observe a tremendous diversity in skills, experience, environmental constraints, and so on. For example, among doctors, there can be big differences in terms of statistical literacy, openness to trusting ML-powered tools, time available to consume and decide on model output, and so on. These variations have important implications for designing effective tools.

To design and build an effective expression of uncertainty, we need to begin with an understanding of who the tool will be used by, what goal that user has, and what needs and constraints the user has. Frequently we also need to understand the organizational and social context in which a user is embedded. For example, to understand how an organization calculates and processes risk, which can influence the design of human-in-the-loop processes, automation, where thresholds are set, and so on. This point is not a new one, and it is by no means unique to the field of ML. User-centered design (UCD), human-computer interaction (HCI), user experience (UX), human factors, and related fields have arisen as responses to this challenge across a wide range of product and tool design contexts [42, 43, 103].

UCD and HCI have a firm footing in many software development contexts, yet they remain relatively neglected in the field of ML. Nevertheless, a growing body of research is beginning to demonstrate the importance of user-centered design for work on ML tools [57, 122]. For example, Yang et al. [132] draw on field research with healthcare decision-makers to understand why an ML-powered tool that performed well in laboratory tests was rejected by clinicians in real-world settings. They found that users saw little need for the tool, lacked trust in its output, and faced environmental barriers that made it difficult to use. Narayanan et al. [89] conduct a series of user tests for explainability to uncover which kinds of increases in explanation complexity have the greatest effect on the time it takes for users to achieve certain tasks. Doshi-Velez and Kim [30] propose a framework for evaluation of

explainability that incorporates tests with users engaged in concrete and realistic tasks. From a practitioner’s perspective, Lovejoy [77] describes the user-centered design lessons learned by the team building Google Clips, an AI-enabled camera designed to capture candid photographs of familiar people and animals. One of their key conclusions is that “[m]achine learning won’t figure out what problems to solve. If you aren’t aligned with a human need, you’re just going to build a very powerful system to address a very small — or perhaps nonexistent — problem.”

Research to uncover user goals, needs, and constraints can involve a wide spectrum of methods, including but not limited to in-depth interviews, contextual inquiry, diary studies, card sorting studies, user tests, user journey mapping, and jobs-to-be-done workshops with users [42, 43, 110]. It is helpful to divide user research into two buckets: 1) discovery research, which aims to understand what problem needs to be solved and for which type of user; and 2) evaluative research, which aims to understand how well our attempts to solve the given problem are succeeding with real users. Ideally, discovery research precedes any effort to build a solution, or at least occurs as early in the process as possible. Doing so helps the team focus on the right problem and right user type when considering possible solutions, and can help a team avoid costly investments that create little value for users. Which of the many methods a researcher uses in discovery and evaluative research will depend on many factors, including how easy it is to find relevant participants, how easy it is for the researcher to observe participants in the context of their day-to-day work, how expensive and time-consuming it is for teams to prototype potential solutions for the purposes of user testing, etc. One takeaway is that teams building uncertainty into ML workflows should do user research to understand what problem needs solving and for what type of user.

Takeaways:

- (1) Stakeholders struggle interpreting uncertainty estimates.
- (2) While uncertainty estimates can be represented with a variety of methods, the chosen method should be one that is tested with stakeholders.
- (3) Teams integrating uncertainty into ML workflows should undertake user research to identify the problems they are solving for and cater to different stakeholder types.

5 CONCLUSION

Throughout this paper, we have argued that uncertainty is a form of transparency and is thus pertinent to the machine learning transparency community. We surveyed the machine learning, visualization/HCI, design, decision-making, and fairness literature. We reviewed how to quantify uncertainty and leverage it in three use cases: (1) for developers reducing the unfairness of models, (2) for experts making decisions, and (3) for stakeholders placing their trust in ML models. We then described the methods for and pitfalls of communicating uncertainty, concluding with a discussion on how to collect requirements for leveraging uncertainty in practice. In summary, well-calibrated uncertainty estimates improve transparency for ML models. In addition to calibration, it is important that these estimates are applied coherently and communicated

clearly to various stakeholders considering the use case at hand. Future work could study the interplay between fairness, transparency, and uncertainty. For example, one could explore how communicating uncertainty to a stakeholder affects their perception of a model’s fairness, or one could study how to best measure the calibration of uncertainty in regression settings. We hope this work inspires others to study uncertainty as transparency and to be mindful of uncertainty’s effects on models in deployment.

ACKNOWLEDGMENTS

The authors would like to thank the following individuals for their advice, contributions, and/or support: James Allingham (University of Cambridge), McKane Andrus (Partnership on AI), Kemi Bello (Partnership on AI), Hudson Hongo (Partnership on AI), Eric Horvitz (Microsoft), Jessica Hullman (Northwestern University), Matthew Kay (Northwestern University), Terah Lyons (Partnership on AI), Elena Spitzer (Google), Richard Tomsett (IBM), Kush Varshney (IBM), and Carroll Wainwright (Partnership on AI).

UB acknowledges support from DeepMind and the Leverhulme Trust via the Leverhulme Centre for the Future of Intelligence (CFI), and from the Partnership on AI. JA acknowledges support from Microsoft Research, through its PhD Scholarship Programme, and from the EPSRC. AW acknowledges support from a Turing AI Fellowship under grant EP/V025379/1, The Alan Turing Institute under EPSRC grant EP/N510129/1 and TU/B/000074, and the Leverhulme Trust via CFI.

REFERENCES

- [1] Franklin A. Graybill Alexander McFarlane Mood and Duane C. Boes. 1974. *Introduction to the theory of statistics* (3rd ed. ed.). McGraw-Hill New York.
- [2] Javier Antoran, Umang Bhatt, Tameem Adel, Adrian Weller, and Jose Miguel Hernandez-Lobato. 2021. Getting a CLUE: A Method for Explaining Uncertainty Estimates. In *International Conference on Learning Representations (ICLR)*.
- [3] Matthew Arnold, Rachel KE Bellamy, Michael Hind, Stephanie Houde, Sameep Mehta, A Mojsilović, Ravi Nair, K Natesan Ramamurthy, Alexandra Olteanu, David Piorowski, et al. 2019. FactSheets: Increasing trust in AI services through supplier’s declarations of conformity. *IBM Journal of Research and Development* 63, 4/5 (2019), 6–1.
- [4] Syed Z Arshad, Jianlong Zhou, Constant Bridon, Fang Chen, and Yang Wang. 2015. Investigating user confidence for uncertainty presentation in predictive decision making. In *Proceedings of the Annual Meeting of the Australian Special Interest Group for Computer Human Interaction*. 352–360.
- [5] Arsenii Ashukha, Alexander Lyzhov, Dmitry Molchanov, and Dmitry Vetrov. 2019. Pitfalls of In-Domain Uncertainty Estimation and Ensembling in Deep Learning. In *International Conference on Learning Representations*.
- [6] Emily Badger, Claire Cain Miller, Adam Pearce, and Kevin Quealy. 2018. Income Mobility Charts for Girls, Asian-Americans and Other Groups. Or Make Your Own. (2018).
- [7] Howard Balshem, Mark Helfand, Holger J. Schünemann, Andrew D. Oxman, Regina Kunz, Jan Brozek, Gunn E. Vist, Yingve Falck-Ytter, Joerg Meerpohl, Susan Norris, and Gordon H. Guyatt. 2011. GRADE Guidelines: 3. Rating the Quality of Evidence. 64, 4 (2011), 401–406. <https://doi.org/10/d49b4h>
- [8] Umang Bhatt, Alice Xiang, Shubham Sharma, Adrian Weller, Ankur Taly, Yunhan Jia, Joydeep Ghosh, Ruchir Puri, José MF Moura, and Peter Eckersley. 2020. Explainable machine learning in deployment. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 648–657.
- [9] Christopher M. Bishop. 2006. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg.
- [10] J Martin Bland and Douglas G Altman. 1998. Bayesians and frequentists. *BMJ* 317, 7166 (1998), 1151–1160. <https://doi.org/10.1136/bmj.317.7166.1151> arXiv:<https://www.bmj.com/content/317/7166/1151.1.full.pdf>
- [11] Avrim Blum and Kevin Stangl. 2019. Recovering from biased data: Can fairness constraints improve accuracy? *arXiv preprint arXiv:1912.01094* (2019).
- [12] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. 2015. Weight Uncertainty in Neural Network. In *International Conference on Machine Learning*. 1613–1622.

- [13] Lonneke Boels, Arthur Bakker, Wim Van Dooren, and Paul Drijvers. 2019. Conceptual Difficulties When Interpreting Histograms: A Review. 28 (2019), 100291. <https://doi.org/10/gbhwb6>
- [14] Leo Breiman. 1996. Bagging Predictors. *Machine Learning* 24, 2 (01 Aug 1996), 123–140. <https://doi.org/10.1023/A:1018054314350>
- [15] Glenn W Brier. 1950. Verification of forecasts expressed in terms of probability. *Monthly weather review* 78, 1 (1950), 1–3.
- [16] David V. Budescu, Han-Hui Por, and Stephen B. Broomell. 2012. Effective Communication of Uncertainty in the IPCC Reports. 113, 2 (2012), 181–200. <https://doi.org/10/c93bb5>
- [17] David V. Budescu and Thomas S. Wallsten. 1985. Consistency in Interpretation of Probabilistic Phrases. 36, 3 (1985), 391–405. <https://doi.org/10/b9qss9>
- [18] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. 77–91.
- [19] Shelly Chaiken. 1999. The heuristic—systematic. *Dual-process theories social psychology* 73 (1999).
- [20] Jiahao Chen, Nathan Kallus, Xiaojie Mao, Geoffry Svacha, and Madeleine Udell. 2019. Fairness under unawareness: Assessing disparity when protected class is unobserved. In *Proceedings of the conference on fairness, accountability, and transparency*. 339–348.
- [21] Tianqi Chen, Emily Fox, and Carlos Guestrin. 2014. Stochastic gradient hamiltonian monte carlo. In *International conference on machine learning*. 1683–1691.
- [22] Dominic A. Clark. 1990. Verbal Uncertainty Expressions: A Critical Review of Two Decades of Research. 9, 3 (1990), 203–235. <https://doi.org/10/c4kmb>
- [23] Marvin S Cohen, Raja Parasuraman, and Jared T Freeman. 1998. Trust in decision aids: A model and its training implications. In *in Proc. Command and Control Research and Technology Symp.* Citeseer.
- [24] Michael Correll and Michael Gleicher. 2014. Error Bars Considered Harmful: Exploring Alternate Encodings for Mean and Error. 20, 12 (2014), 2142–2151. <https://doi.org/10/23c>
- [25] Christina Curtis, Sohrab P Shah, Suet-Feung Chin, Gulisa Turashvili, Oscar M Rueda, Mark J Dunning, Doug Speed, Andy G Lynch, Shamith Samarajiva, Yinyin Yuan, et al. 2012. The genomic and transcriptomic architecture of 2,000 breast tumours reveals novel subgroups. *Nature* 486, 7403 (2012), 346–352.
- [26] Stefan Depeweg. 2019. *Modeling Epistemic and Aleatoric Uncertainty with Bayesian Neural Networks and Latent Variables*. Ph.D. Dissertation. Technical University of Munich.
- [27] Armen Der Kiureghian and Ove Ditlevsen. 2009. Aleatory or epistemic? Does it matter? *Structural safety* 31, 2 (2009), 105–112.
- [28] Thomas G. Dietterich. 2000. Ensemble Methods in Machine Learning. In *Proceedings of the First International Workshop on Multiple Classifier Systems (MCS '00)*. Springer-Verlag, Berlin, Heidelberg, 1–15.
- [29] Andreia Dionisio, Rui Menezes, and Diana A Mendes. 2007. Entropy and uncertainty analysis in financial markets. *arXiv preprint arXiv:0709.0668* (2007).
- [30] Finale Doshi-Velez and Been Kim. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608* (2017).
- [31] Dheeru Dua and Casey Graff. 2017. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>
- [32] Danielle Ensign, Sorelle A Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. 2018. Runaway feedback loops in predictive policing. In *Conference on Fairness, Accountability and Transparency*. 160–171.
- [33] Kavin Ethayarajh. 2020. Is Your Classifier Actually Biased? Measuring Fairness under Uncertainty with Bernstein Bounds. *arXiv preprint arXiv:2004.12332* (2020).
- [34] Baruch Fischhoff and Alex I Davis. 2014. Communicating scientific uncertainty. *Proceedings of the National Academy of Sciences* 111, Supplement 4 (2014), 13664–13671.
- [35] Riccardo Fogliato, Max G'Sell, and Alexandra Chouldechova. 2020. Fairness Evaluation in Presence of Biased Noisy Labels. *arXiv preprint arXiv:2003.13808* (2020).
- [36] Yarin Gal. 2016. Uncertainty in deep learning. *University of Cambridge* 1, 3 (2016).
- [37] Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*. 1050–1059.
- [38] Mirta Galesic and Rocio Garcia-Retamero. 2010. Statistical Numeracy for Health: A Cross-Cultural Comparison With Probabilistic National Samples. 170, 5 (2010), 462–468. <https://doi.org/10/fmj7q3>
- [39] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Dauméé III, and Kate Crawford. 2018. Datasheets for datasets. *arXiv preprint arXiv:1803.09010* (2018).
- [40] Zoubin Ghahramani. 2015. Probabilistic machine learning and artificial intelligence. *Nature* 521, 7553 (01 May 2015), 452–459. <https://doi.org/10.1038/nature14541>
- [41] Tilmann Gneiting and Adrian E Raftery. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association* 102, 477 (2007), 359–378.
- [42] Elizabeth Goodman. 2009. Three environmental discourses in human-computer interaction. In *CHI'09 Extended Abstracts on Human Factors in Computing Systems*. 2535–2544.
- [43] Elizabeth Goodman, Mike Kuniavsky, and Andrea Moed. 2012. *Observing the user experience: A practitioner's guide to user research*. Elsevier.
- [44] Maya Gupta, Andrew Cotter, Mahdi Milani Fard, and Serena Wang. 2018. Proxy fairness. *arXiv preprint arXiv:1806.11212* (2018).
- [45] Chip Heath and Amos Tversky. 1991. Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of risk and uncertainty* 4, 1 (1991), 5–28.
- [46] José Miguel Hernández-Lobato and Ryan Adams. 2015. Probabilistic back-propagation for scalable learning of bayesian neural networks. In *International Conference on Machine Learning*. 1861–1869.
- [47] José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. 2014. Predictive entropy search for efficient global optimization of black-box functions. In *Advances in neural information processing systems*. 918–926.
- [48] Geoffrey E. Hinton and Drew van Camp. 1993. Keeping the Neural Networks Simple by Minimizing the Description Length of the Weights. In *Proceedings of the Sixth Annual Conference on Computational Learning Theory (COLT '93)*. ACM, New York, NY, USA, 5–13. <https://doi.org/10.1145/168304.168306>
- [49] Jerry L. Hintze and Ray D. Nelson. 1998. Violin Plots: A Box Plot-Density Trace Synergism. 52, 2 (1998), 181–184. <https://doi.org/10/gf5hpg>
- [50] Kevin Anthony Hoff and Masooda Bashir. 2015. Trust in automation: Integrating empirical evidence on factors that influence trust. *Human factors* 57, 3 (2015), 407–434.
- [51] Jake M Hofman, Daniel G Goldstein, and Jessica Hullman. 2020. How visualizing inferential uncertainty can mislead readers about treatment effects in scientific results. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [52] Stephen C Hora. 1996. Aleatory and epistemic uncertainty in probability elicitation with an example from hazardous waste management. *Reliability Engineering & System Safety* 54, 2-3 (1996), 217–223.
- [53] Neil Houlsby, Ferenc Huszar, Zoubin Ghahramani, and Máté Lengyel. 2011. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745* (2011).
- [54] Carl Iver Hovland, Irving Lester Janis, and Harold H Kelley. 1953. Communication and persuasion. (1953).
- [55] Jessica Hullman, Xiaoli Qiao, Michael Correll, Alex Kale, and Matthew Kay. 2018. In pursuit of error: A survey of uncertainty visualization evaluation. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 903–913.
- [56] Jessica Hullman, Paul Resnick, and Eytan Adar. 2015. Hypothetical Outcome Plots Outperform Error Bars and Violin Plots for Inferences about Reliability of Variable Ordering. 10, 11 (2015), e0142444. <https://doi.org/10/f3tvsd>
- [57] Kori Inkpen, Stevie Chancellor, Munmun De Choudhury, Michael Veale, and Eric PS Baumer. 2019. Where is the human? Bridging the gap between AI and HCI. In *Extended abstracts of the 2019 chi conference on human factors in computing systems*. 1–9.
- [58] David Janz, Jiri Hron, Przemysław Mazur, Katja Hofmann, José Miguel Hernández-Lobato, and Sebastian Tschiatschek. 2019. Successor Uncertainties: Exploration and Uncertainty in Temporal Difference Learning. In *Advances in Neural Information Processing Systems* 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 4507–4516. <http://papers.nips.cc/paper/8700-successor-uncertainties-exploration-and-uncertainty-in-temporal-difference-learning.pdf>
- [59] Heinrich Jiang and Ofir Nachum. 2020. Identifying and correcting label bias in machine learning. In *International Conference on Artificial Intelligence and Statistics*. 702–712.
- [60] Daniel Kahneman. 2011. *Thinking, fast and slow*. Macmillan.
- [61] Daniel Kahneman and Amos Tversky. 2013. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I*. World Scientific, 99–127.
- [62] Alex Kale, Matthew Kay, and Jessica Hullman. 2020. Visual reasoning strategies for effect size judgments and decisions. *IEEE Transactions on Visualization and Computer Graphics* (2020).
- [63] Nathan Kallus, Xiaojie Mao, and Angela Zhou. 2019. Assessing algorithmic fairness with unobserved protected class using data combination. *arXiv preprint arXiv:1906.00285* (2019).
- [64] Matthew Kay, Tara Kola, Jessica R. Hullman, and Sean A. Munson. 2016. When (Is) Is My Bus? User-Centered Visualizations of Uncertainty in Everyday, Mobile Predictive Systems. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016-05-07) (CHI '16). Association for Computing Machinery, 5092–5103. <https://doi.org/10/b2pj>
- [65] Frank Hyneman Knight. 1921. *Risk, uncertainty and profit*. Vol. 31. Houghton Mifflin.
- [66] Mykel J Kochenderfer. 2015. *Decision making under uncertainty: theory and application*. MIT press.
- [67] Pang Wei Koh and Percy Liang. 2017. Understanding black-box predictions via influence functions. In *Proceedings of the 34th International Conference on*

- Machine Learning-Volume 70 (ICML 2017)*. Journal of Machine Learning Research, 1885–1894.
- [68] Moritz Körber. 2018. Theoretical considerations and development of a questionnaire to measure trust in automation. In *Congress of the International Ergonomics Association*. Springer, 13–30.
 - [69] John Kruschke. 2014. Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan. (2014).
 - [70] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in neural information processing systems*. 6402–6413.
 - [71] Alex Lamy, Ziyuan Zhong, Aditya K Menon, and Nakul Verma. 2019. Noise-tolerant fair classification. In *Advances in Neural Information Processing Systems*. 294–306.
 - [72] John D Lee and Katrina A See. 2004. Trust in automation: Designing for appropriate reliance. *Human factors* 46, 1 (2004), 50–80.
 - [73] Dan Ley, Umang Bhatt, and Adrian Weller. 2021. δ -CLUE: Diverse Sets of Explanations for Uncertainty Estimates. *arXiv preprint arXiv:2104.06323* (2021).
 - [74] Sarah Lichtenstein and J. Robert Newman. 1967. Empirical Scaling of Common Verbal Phrases Associated with Numerical Probabilities. 9, 10 (1967), 563–564. <https://doi.org/10/ghbkh9>
 - [75] I. M. Lipkus and J. G. Hollands. [n. d.]. The Visual Communication of Risk. 25 ([n. d.]), 149–163. <https://doi.org/10/gd589v> arXiv:10854471
 - [76] Daniel Hernández Lobato. 2009. Prediction based on averages over automatically induced learners ensemble methods and Bayesian techniques.
 - [77] Josh Lovejoy. 2018. The UX of AI: Using Google Clips to understand how a human-centered design process elevates artificial intelligence. In *2018 AAAI Spring Symposium Series*.
 - [78] Ana Lucic, Madhulika Srikumar, Umang Bhatt, Alice Xiang, Ankur Taly, Q Vera Liao, and Maarten de Rijke. 2021. A Multistakeholder Approach Towards Evaluating AI Transparency Mechanisms. *arXiv preprint arXiv:2103.14976* (2021).
 - [79] Kristian Lum and William Isaac. 2016. To predict and serve? *Significance* 13, 5 (2016), 14–19.
 - [80] David JC MacKay. 1992. A practical Bayesian framework for backpropagation networks. *Neural computation* 4, 3 (1992), 448–472.
 - [81] David JC MacKay. 2003. *Information theory, inference and learning algorithms*. Cambridge university press.
 - [82] Roger C Mayer, James H Davis, and F David Schoorman. 1995. An integrative model of organizational trust. *Academy of management review* 20, 3 (1995), 709–734.
 - [83] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and A. Galstyan. 2019. A Survey on Bias and Fairness in Machine Learning. *ArXiv abs/1908.09635* (2019).
 - [84] Miriam J Metzger and Andrew J Flanagan. 2013. Credibility and trust of information in online environments: The use of cognitive heuristics. *Journal of pragmatics* 59 (2013), 210–220.
 - [85] Suzanne M Miller. 1987. Monitoring and blunting: validation of a questionnaire to assess styles of information seeking under threat. *Journal of personality and social psychology* 52, 2 (1987), 345.
 - [86] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, 220–229.
 - [87] Malik Sajjad Ahmed Nadeem, Jean-Daniel Zucker, and Blaise Hanczar. 2009. Accuracy-Rejection Curves (ARCs) for Comparing Classification Methods with a Reject Option. In *Proceedings of the third International Workshop on Machine Learning in Systems Biology (Proceedings of Machine Learning Research)*, Sašo Džeroski, Pierre Guerts, and Juho Rousu (Eds.), Vol. 8. PMLR, Ljubljana, Slovenia, 65–81. <http://proceedings.mlr.press/v8/nadeem10a.html>
 - [88] Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. 2015. Obtaining well calibrated probabilities using bayesian binning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.
 - [89] Menaka Narayanan, Emily Chen, Jeffrey He, Been Kim, Sam Gershman, and Finae Doshi-Velez. 2018. How do humans understand explanations from machine learning systems? an evaluation of the human-interpretability of explanation. *arXiv preprint arXiv:1802.00682* (2018).
 - [90] Radford M. Neal. 1995. *Bayesian Learning for Neural Networks*. Ph.D. Dissertation. CAN. Advisor(s) Hinton, Geoffrey. AAINN02676.
 - [91] Radford M Neal. 1995. Bayesian Learning for Neural Networks. *PhD thesis, University of Toronto* (1995).
 - [92] Cuong V. Nguyen, Yingzhen Li, Thang D. Bui, and Richard E. Turner. 2018. Variational Continual Learning. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=BkQqOgRb>
 - [93] NHC. 2020. Definition of the NHC Track Forecast Cone. <https://www.nhc.noaa.gov/aboutcone.shtml>
 - [94] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453.
 - [95] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kiciman. 2019. Social data: Biases, methodological pitfalls, and ethical boundaries. *Frontiers in Big Data* 2 (2019), 13.
 - [96] Onora O'Neill. 2018. Linking trust to trustworthiness. *International Journal of Philosophical Studies* 26, 2 (2018), 293–300.
 - [97] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, D Sculley, Sebastian Nowozin, Joshua V Dillon, Balaji Lakshminarayanan, and Jasper Snoek. 2019. Can You Trust Your Model's Uncertainty? Evaluating Predictive Uncertainty Under Dataset Shift. *arXiv preprint arXiv:1906.02530* (2019).
 - [98] Raja Parasuraman and Christopher A Miller. 2004. Trust and etiquette in high-criticality automated systems. *Commun. ACM* 47, 4 (2004), 51–55.
 - [99] Jolynn Pek and Trisha Van Zandt. 2020. Frequentist and Bayesian approaches to data analysis: Evaluation and estimation. *Psychology Learning & Teaching* 19, 1 (2020), 21–35. <https://doi.org/10.1177/1475725719874542> arXiv:https://doi.org/10.1177/1475725719874542
 - [100] Ellen Peters, Judith Hibbard, Paul Slovic, and Nathan Dieckmann. 2007. Numeracy Skill And The Communication, Comprehension, And Use Of Risk-Benefit Information. 26, 3 (2007), 741–748. <https://doi.org/10/c77p38>
 - [101] Richard E Petty and John T Cacioppo. 1986. The elaboration likelihood model of persuasion. In *Communication and persuasion*. Springer, 1–24.
 - [102] Mary C. Politi, Paul K. J. Han, and Nananda F. Col. 2007. Communicating the Uncertainty of Harms and Benefits of Medical Interventions. 27, 5 (2007), 681–695. <https://doi.org/10/c9b8g4>
 - [103] Jennifer Preece, Helen Sharp, and Yvonne Rogers. 2015. *Interaction design: beyond human-computer interaction*. John Wiley & Sons.
 - [104] Inioluwa Deborah Raji and Jingying Yang. 2019. ABOUT ML: Annotation and Benchmarking on Understanding and Transparency of Machine Learning Lifecycles. *arXiv preprint arXiv:1912.06166* (2019).
 - [105] Tim Rakow. 2010. Risk, uncertainty and prophet: The psychological insights of Frank H. Knight. *Judgment and Decision Making* 5, 6 (2010), 458.
 - [106] Carl Edward Rasmussen and Christopher K. I. Williams. 2005. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.
 - [107] Valerie F. Reyna and Charles J. Brainerd. 2008. Numeracy, Ratio Bias, and Denominator Neglect in Judgments of Risk and Probability. 18, 1 (2008), 89–107. <https://doi.org/10/bdqnh5>
 - [108] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why should I trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1135–1144.
 - [109] Peter J. Rousseeuw, Ida Ruts, and John W. Tukey. 1999. The Bagplot: A Bivariate Boxplot. 53, 4 (1999), 382–387. <https://doi.org/10/gg6cp5>
 - [110] Jeff Rubin and Dana Chisnell. 2008. How to plan, design, and conduct effective tests. (2008).
 - [111] David A Schum, Gheorghe Tecuci, Dorin Marcu, and Mihai Boicu. 2014. Toward cognitive assistants for complex decision making under uncertainty. *Intelligent Decision Technologies* 8, 3 (2014), 231–250.
 - [112] Clayton Scott, Gilles Blanchard, and Gregory Handy. 2013. Classification with asymmetric label noise: Consistency and maximal denoising. In *Conference On Learning Theory*. 489–511.
 - [113] Shai Shalev-Shwartz and Shai Ben-David. 2014. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, USA.
 - [114] Richard M Sorrentino, D Ramona Bobocel, Maria Z. Gitta, James M Olson, and Erin C Hewitt. 1988. Uncertainty orientation and persuasion: Individual differences in the effects of personal relevance on social judgments. *Journal of Personality and social Psychology* 55, 3 (1988), 357.
 - [115] David Spiegelhalter. 2017. Risk and Uncertainty Communication. 4, 1 (2017), 31–60. <https://doi.org/10.1146/annurev-statistics-010814-020148>
 - [116] D. Spiegelhalter, M. Pearson, and I. Short. 2011. Visualizing Uncertainty About the Future. 333, 6048 (2011), 1393–1400. <https://doi.org/10/cd4>
 - [117] Kimberly Stowers, Nicholas Kasdaglis, Michael Rupp, Jessie Chen, Daniel Barber, and Michael Barnes. 2017. Insights into human-agent teaming: Intelligent agent transparency and uncertainty. In *Advances in Human Factors in Robots and Unmanned Systems*. Springer, 149–160.
 - [118] S Shyam Sundar. 2008. *The MAIN model: A heuristic approach to understanding technology effects on credibility*. MacArthur Foundation Digital Media and Learning Initiative.
 - [119] S Shyam Sundar and Jinyoung Kim. 2019. Machine heuristic: When we trust computers more than humans with our personal information. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–9.
 - [120] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70 (ICML 2017)*. Journal of Machine Learning Research, 3319–3328.
 - [121] Harini Suresh and John V. Guttag. 2019. A Framework for Understanding Unintended Consequences of Machine Learning. arXiv:cs.LG/1901.10002
 - [122] Anja Thieme, Danielle Belgrave, and Gavin Doherty. 2020. Machine learning in mental health: A systematic review of the HCI literature to support the

- development of effective and implementable ML systems. *ACM Transactions on Computer-Human Interaction (TOCHI)* 27, 5 (2020), 1–53.
- [123] Richard Tomsett, Alun Preece, Dave Braines, Federico Cerutti, Supriyo Chakraborty, Mani Srivastava, Gavin Pearson, and Lance Kaplan. 2020. Rapid trust calibration through interpretable and uncertainty-aware AI. *Patterns* 1, 4 (2020), 100049.
- [124] Amos Tversky and Daniel Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *science* 185, 4157 (1974), 1124–1131.
- [125] Amos Tversky and Daniel Kahneman. 1992. Advances in Prospect Theory: Cumulative Representation of Uncertainty. 5, 4 (1992), 297–323. <https://doi.org/10/cb57hk>
- [126] Anne Marthe Van Der Bles, Sander van der Linden, Alexandra LJ Freeman, and David J Spiegelhalter. 2020. The effects of communicating uncertainty on public trust in facts and numbers. *Proceedings of the National Academy of Sciences* 117, 14 (2020), 7672–7683.
- [127] Anne Marthe van der Bles, Sander van der Linden, Alexandra L. J. Freeman, James Mitchell, Ana B. Galvao, Lisa Zaval, and David J. Spiegelhalter. 2019. Communicating Uncertainty about Facts, Numbers and Science. 6, 5 (2019), 181870. <https://doi.org/10/gf2g9j>
- [128] Adrian Weller. 2019. Transparency: motivations and challenges. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. Springer, 23–40.
- [129] Max Welling and Yee W Teh. 2011. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*. 681–688.
- [130] Claus O. Wilke. 2019. *Fundamentals of Data Visualization: A Primer on Making Informative and Compelling Figures* (1st edition ed.). O'Reilly Media.
- [131] Min-ge Xie and Kesar Singh. 2013. Confidence Distribution, the Frequentist Distribution Estimator of a Parameter: A Review. *International Statistical Review* 81, 1 (2013), 3–39. <https://doi.org/10.1111/insr.12000> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/insr.12000>
- [132] Qian Yang, John Zimmerman, Aaron Steinfeld, Lisa Carey, and James F Antaki. 2016. Investigating the heart pump implant decision process: opportunities for decision support tools to help. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 4477–4488.
- [133] Nathan Yau. 2015. Years You Have Left to Live, Probably. <https://flowingdata.com/2015/09/23/years-you-have-left-to-live-probably/>
- [134] Nanyang Ye and Zhanxing Zhu. 2018. Bayesian Adversarial Learning. In *Advances in Neural Information Processing Systems* 31, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.). Curran Associates, Inc., 6892–6901. <http://papers.nips.cc/paper/7921-bayesian-adversarial-learning.pdf>
- [135] Hang Zhang and Laurence T. Maloney. 2012. Ubiquitous Log Odds: A Common Representation of Probability and Frequency Distortion in Perception, Action, and Cognition. 6 (2012). <https://doi.org/10/fxxssh>
- [136] Ruqi Zhang, Chunyuan Li, Jianyi Zhang, Changyou Chen, and Andrew Gordon Wilson. 2020. Cyclical Stochastic Gradient MCMC for Bayesian Deep Learning. *International Conference on Learning Representations* (2020).
- [137] Yunfeng Zhang, Rachel KE Bellamy, and Wendy A. Kellogg. 2012. Designing Information for Remediating Cognitive Biases in Decision-Making. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (2015). ACM, 2211–2220. <https://doi.org/10.1145/2702123.2702239>
- [138] Yunfeng Zhang, Q Vera Liao, and Rachel KE Bellamy. 2020. Effect of confidence and explanation on accuracy and trust calibration in AI-assisted decision making. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 295–305.
- [139] Brian J. Zikmund-Fisher, Dylan M. Smith, Peter A. Ubel, and Angela Fagerlin. 2007. Validation of the Subjective Numeracy Scale: Effects of Low Numeracy on Comprehension of Risk Communications and Utility Elicitations. 27, 5 (2007), 663–671. <https://doi.org/10/cf2642>