Combining Aleatoric and Epistemic Uncertainties for Robust Healthcare Decision-Making

Zhiyin Lin, Sam Saarinen, Michael L. Littman

BROWN Computer Science

Introduction

Establishing appropriate trust between healthcare practitioners and machine-learning models is critical to enabling automation of medical processes by models with human-level or super-human performance, thus improving outcomes and lowering costs [5]. Potential issues:

- Under-trust: the healthcare practitioners may unduly reject the prediction provided by the machine-learning model.
- Over-trust: medical practitioners may default to an automated prediction that turns out to be imperfect.

In both cases, a more appropriate level of trust may be established by accurately identifying and communicating a case-specific confidence in the model's prediction [5].

"I predict the value μ"

⬇

"I predict the value μ ± 2σ"

All problems and techniques discussed below are in the context of regression problems and for arbitrary models

Evaluation Metrics

How can a model that outputs means *and* variances be evaluated?

Inspiration from Bayesian modeling: maximize the likelihood of a test set under a model $M : X \to \mathbb{P}[Y]$

⬇ Outputs are independent of each other

$$\text{argmax}_M \prod_{i=1}^{n} M(x_i)[y_i]$$

⬇ Notational & computational reasons

$$\text{argmax}_M \sum_{i=1}^{n} \ln M(x_i)[y_i]$$

Uncertainty Modeling

Aleatoric Risk (Chance-based) + Epistemic Uncertainty (Knowledge-based)

⬇

Total Uncertainty

Method

Fitting Random Priors: Estimate Epistemic Uncertainty [1]

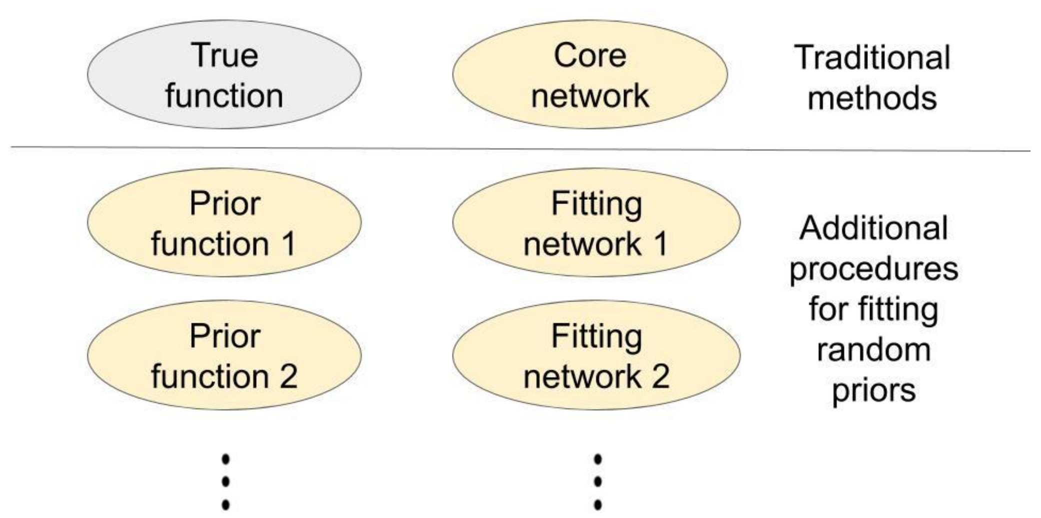

Fitting random priors creates simulated learning problems using randomly sampled prior functions to generate artificial datasets from the same inputs. The performance should correlate with the true problem.

Regress to Variance: Aleatoric Risk Target [3]

- (Predicted - True)² - Epistemic Uncertainty
- Clip negative values at zero

Uncertainty Calibration: Improve Accuracy and Generality of Aleatoric Risk Estimates [6]

- Isotonic Regression: Predicted variance → True error

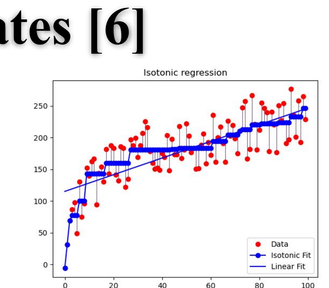

- Demonstration: UCI Wisconsin Breast Cancer (Diagnostic) Dataset [7]

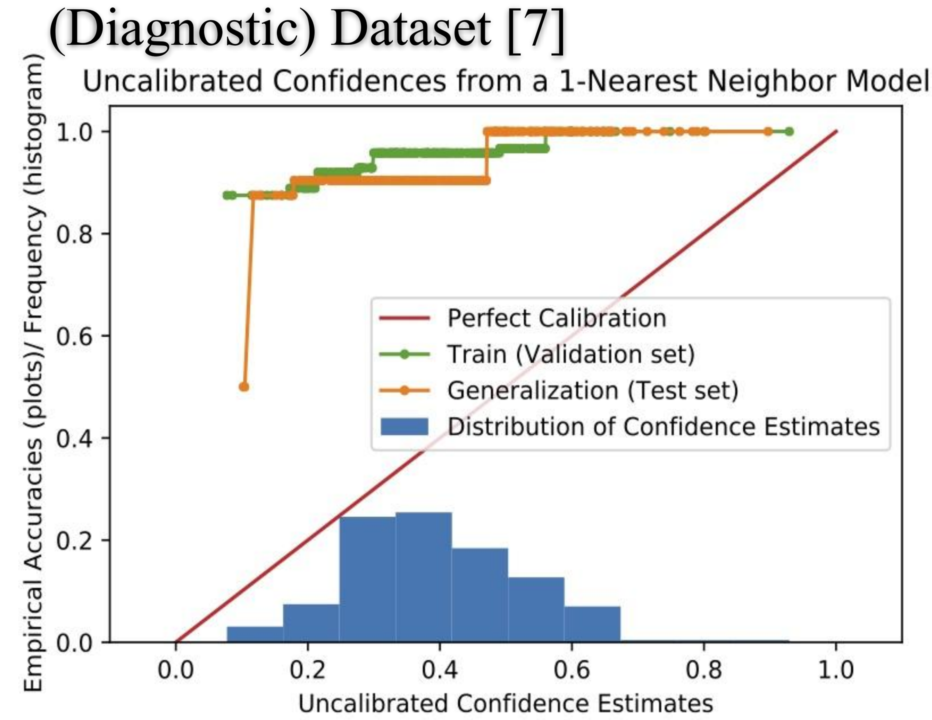

Calibration shifts the estimated uncertainty (x-values) so they agree with empirical errors (y-values).

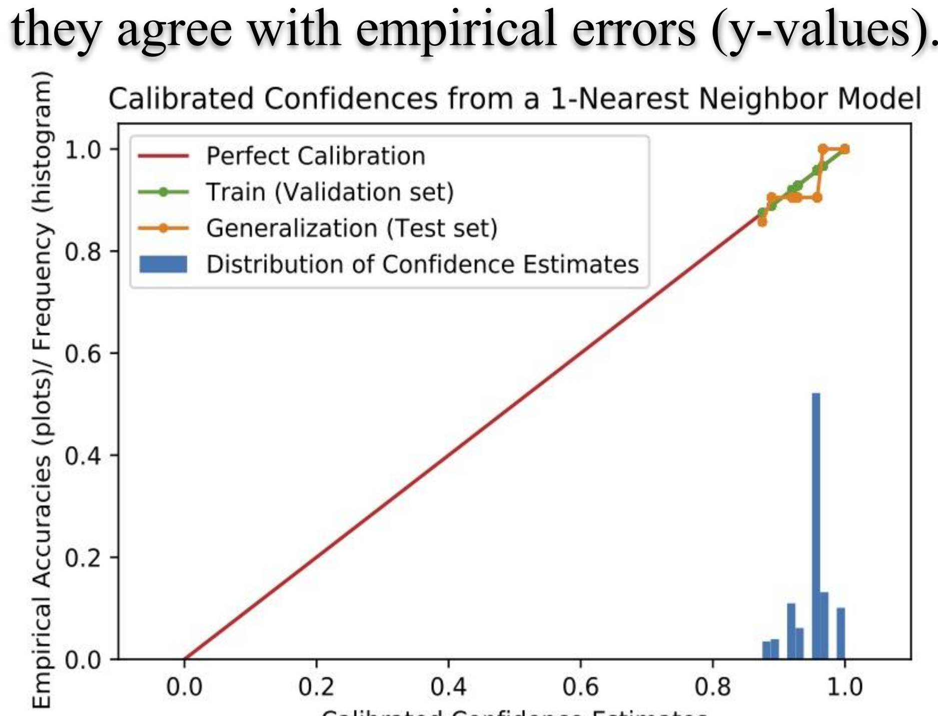

Experiments

Diabetes Progression Dataset [2]: a regression problem, 442 instances, 10 features

- Ablation Study on Log-Likelihood Performance

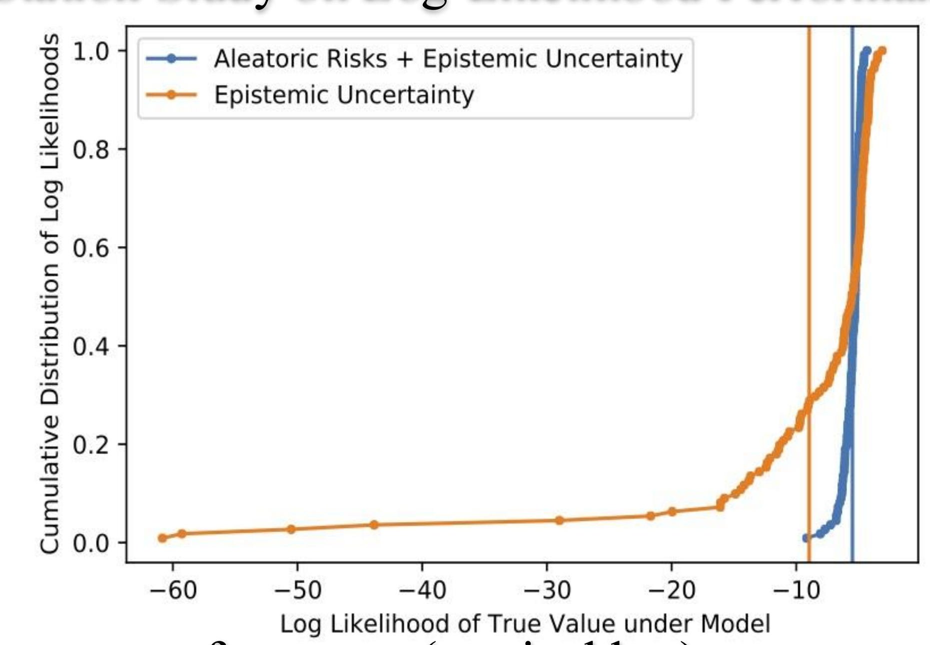

The mean performance (vertical bar) on our evaluation metric shows that combined method is more accurate than either. Aleatoric-Only data could not be plotted due to some 0-uncertainty estimates producing infinitely negative performance values.

- Out-of-Distribution Experiment

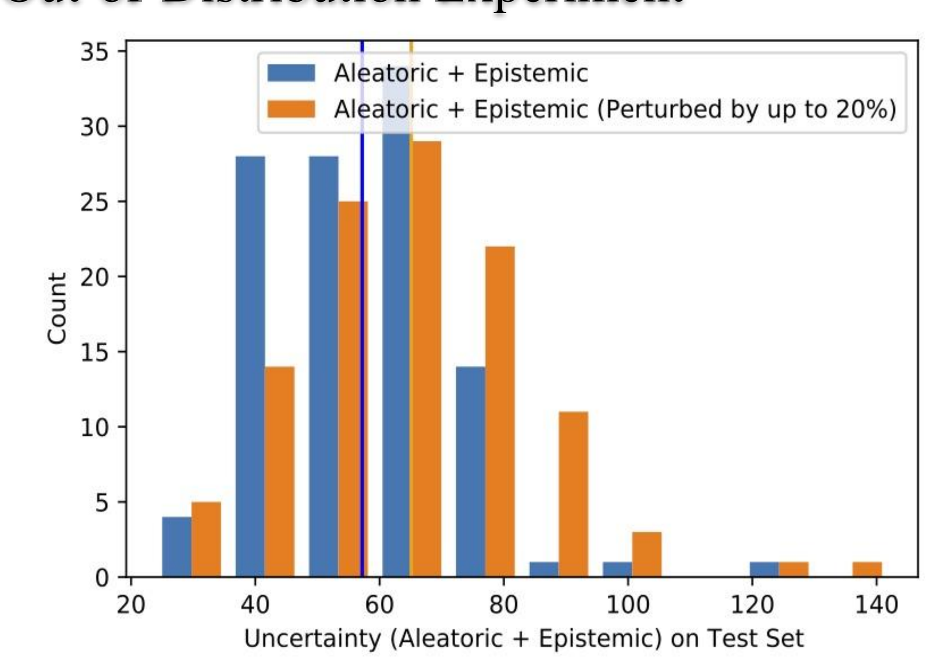

The method produces higher uncertainty on data that didn't come from the original dataset.

- Distinguish Aleatoric Risk and Epistemic Uncertainty

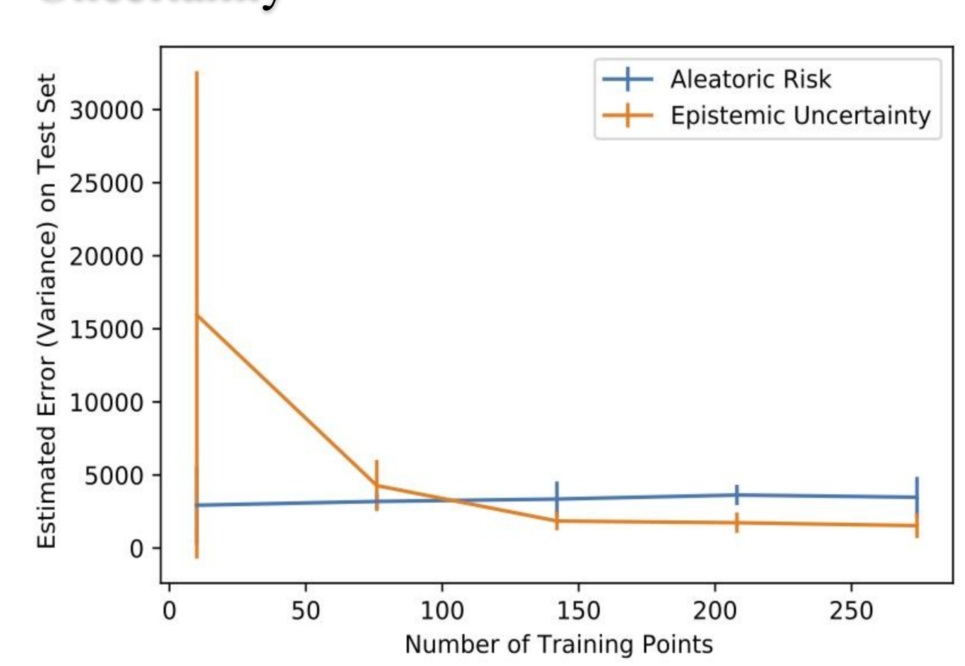

As we'd hope, the epistemic uncertainty goes down as the number of training points increases (getting more information), but the aleatoric risk remains roughly constant, indicating the model is able to differentiate the two sources of errors.

Clinical Applicability

Why distinguish between aleatoric risk and epistemic uncertainty?

Patient 1: 175 ± 91, with 64% of the uncertainty being epistemic.

This is a point where the mean estimate could be wrong by a large margin, mostly because the model simply does not have enough data on similar patients. In this case, the clinician may defer to their own experience and insights or seek more definitive testing and data collection.

Patient 2: 213 ± 121, with 98% of the uncertainty being aleatoric.

For this patient, there have been many similar patients in the data, so the mean value is quite accurate, but the outcome for the patient is still highly uncertain. In this case, it is best that the clinician use the mean estimate provided by the model and communicate with the patient what additional factors (such as behavioral choices) may affect their outcome.

Patient 3: 113 ± 48, with 77% of the uncertainty being epistemic.

The model is already much more confident than average, and additional data on similar patients might further improve the quality of the prediction.

Conclusion

Core contributions:

1. A novel combination of fitted random priors, regression to variance, and uncertainty calibration using isotonic regression
2. Experimental validation of the usefulness of this technique in accurately assessing uncertainty
3. Experimental evidence that the technique correctly distinguishes between aleatoric risk and epistemic uncertainty

Open questions include comparing to other methods for modeling uncertainty, such as ensemble methods, and adapting to classification problems.

References

[1] Ciosek, K.; Fortuin, V.; Tomioka, R.; Hofmann, K.; and Turner, R. 2020. Conservative uncertainty estimation by fitting prior networks. In *International Conference on Learning Representations*.
[2] Efron, B.; Hastie, T.; Johnstone, I.; Tibshirani, R.; et al. 2004. Least angle regression. *The Annals of statistics* 32(2):407–499.
[3] Fan, J., and Yao, Q. 1998. Efficient estimation of conditional variance functions in stochastic regression. *Biometrika* 85(3):645–660.
[4] He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, 1026–1034.
[5] Jorritsma, W.; Cnossen, F.; and van Ooijen, P. M. 2015. Improving the radiologist–cad interaction: designing for appropriate trust. *Clinical radiology* 70(2):115–122.
[6] Kuleshov, V.; Fenner, N.; and Ermon, S. 2018. Accurate uncertainties for deep learning using calibrated regression. *arXiv preprint arXiv:1807.00263*.
[7] Street, W. N.; Wolberg, W. H.; and Mangasarian, O. L. 1993. Nuclear feature extraction for breast tumor diagnosis. In *Biomedical image processing and biomedical visualization*, volume 1905, 861–870. International Society for Optics and Photonics.