
Uncertainty Quantification in Deep Learning

Platform Python 3.6 Implementation Pytorch

This repo contains literature survey and implementation of baselines for predictive uncertainty estimation in deep learning.

Literature survey

Basic background for uncertainty estimation

- B. Efron and R. Tibshirani. “Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy.” Statistical science, 1986. [Link]
- R. Barber, E. J. Candes, A. Ramdas, and R. J. Tibshirani. “Predictive inference with the jackknife+.” arXiv, 2019. [Link]
- B. Efron. “Jackknife-after-bootstrap standard errors and influence functions.” Journal of the Royal Statistical Society: Series B (Methodological), 1992. [Link]
- J. Robins and A. Van Der Vaart. “Adaptive nonparametric confidence sets.” The Annals of Statistics, 2006. [Link]
- V. Vovk, et al., “Cross-conformal predictive distributions.” JMLR, 2018. [Link]
- M. H Quenouille., “Approximate tests of correlation in time-series.” Journal of the Royal Statistical Society, 1949. [Link]
- M. H Quenouille. “Notes on bias in estimation.” Biometrika, 1956. [Link]
- J. Tukey. “Bias and confidence in not quite large samples.” Ann. Math. Statist, 1958.
- R. G. Miller. “The jackknife—a review.” Biometrika, 1974. [Link]
- B. Efron. “Bootstrap methods: Another look at the jackknife.” Ann. Statist., 1979. [Link]
- R. A Stine. “Bootstrap prediction intervals for regression.” Journal of the American Statistical Association, 1985. [Link]
- R. F. Barber, E. J. Candes, A. Ramdas, and R. J. Tibshirani. “Conformal prediction under covariate shift.” arXiv preprint arXiv:1904.06019, 2019. [Link]
- R. F. Barber, E. J. Candes, A. Ramdas, and R. J. Tibshirani. “The limits of distribution-free conditional predictive inference.” arXiv preprint arXiv:1903.04684, 2019b. [Link]
- J. Lei, M. G’Sell, A. Rinaldo, R. J. Tibshirani, and L. Wasserman. “Distribution-free predictive inference for regression.” Journal of the American Statistical Association, 2018. [Link]

-
- R. Giordano, M. I. Jordan, and T. Broderick. “A Higher-Order Swiss Army Infinitesimal Jackknife.” arXiv, 2019. [Link]
 - P. W. Koh, K. Ang, H. H. K. Teo, and P. Liang. “On the Accuracy of Influence Functions for Measuring Group Effects.” arXiv, 2019. [Link]
 - D. H. Wolpert. “Stacked generalization.” Neural networks, 1992. [Link]
 - R. D. Cook, and S. Weisberg. “Residuals and influence in regression.” New York: Chapman and Hall, 1982. [Link]
 - R. Giordano, W. Stephenson, R. Liu, M. I. Jordan, and T. Broderick. “A Swiss Army Infinitesimal Jackknife.” arXiv preprint arXiv:1806.00550, 2018. [Link]
 - P. W. Koh, and P. Liang. “Understanding black-box predictions via influence functions.” ICML, 2017. [Link]
 - S. Wager and S. Athey. “Estimation and inference of heterogeneous treatment effects using random forests.” Journal of the American Statistical Association, 2018. [Link]
 - J. F. Lawless, and M. Fredette. “Frequentist prediction intervals and predictive distributions.” Biometrika, 2005. [Link]
 - F. R. Hampel, E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel. “Robust statistics: the approach based on influence functions.” John Wiley and Sons, 2011. [Link]
 - P. J. Huber and E. M. Ronchetti. “Robust Statistics.” John Wiley and Sons, 1981.
 - Y. Romano, R. F. Barber, C. Sabatti, E. J. Candès. “With Malice Towards None: Assessing Uncertainty via Equalized Coverage.” arXiv, 2019. [Link]
 - H. R. Kunsch. “The Jackknife and the Bootstrap for General Stationary Observations.” The annals of Statistics, 1989. [Link]

Predictive uncertainty for general machine learning models

- A. Malinin, L. Prokhorenkova, A. Ustimenko. “Uncertainty in Gradient Boosting via Ensembles.” ICLR, 2021. [Link]
- S. Feldman, S. Bates, Y. Romano. “Improving Conditional Coverage via Orthogonal Quantile Regression.” arXiv preprint, 2021. [Link]
- S. Bates, A. Angelopoulos, L. Lei, J. Malik, and M. I. Jordan. “Distribution-Free, Risk-Controlling Prediction Sets.” arXiv preprint, 2021. [Link]
- S. Wager, T. Hastie, and B. Efron. “Confidence intervals for random forests: The jackknife and the infinitesimal jackknife.” The Journal of Machine Learning Research, 2014. [Link]

-
- L. Mentch and G. Hooker. “Quantifying uncertainty in random forests via confidence intervals and hypothesis tests.” *The Journal of Machine Learning Research*, 2016. [Link]
 - J. Platt. “Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods.” *Advances in large margin classifiers*, 1999. [Link]
 - A. Abadie, S. Athey, G. Imbens. “Sampling-based vs. design-based uncertainty in regression analysis.” *arXiv preprint (arXiv:1706.01778)*, 2017. [Link]
 - T. Duan, A. Avati, D. Y. Ding, S. Basu, Andrew Y. Ng, and A. Schuler. “NGBoost: Natural Gradient Boosting for Probabilistic Prediction.” *arXiv preprint*, 2019. [Link]
 - V. Franc, and D. Prusa. “On Discriminative Learning of Prediction Uncertainty.” *ICML*, 2019. [Link]
 - Y. Romano, M. Sesia, and E. J. Candès. “Classification with Valid and Adaptive Coverage.” *arXiv preprint*, 2020. [Link]

Predictive uncertainty for deep learning

- I. Osband, Z. Wen, M. Asghari, M. Ibrahimi, X. Lu, and B. Van Roy “Epistemic Neural Networks.” *arXiv*, 2021. [Link]
- Abdar, Moloud, et al. “A review of uncertainty quantification in deep learning: Techniques, applications and challenges.” *Information Fusion*, 2021. [Link]
- Gawlikowski, Jakob, et al. “A Survey of Uncertainty in Deep Neural Networks.” *arXiv preprint*, 2021. [Link]
- P. Morales-Alvarez, D. Hernández-Lobato, R. Molina, J. M. Hernández-Lobato. “Activation-level uncertainty in deep neural networks.” *ICLR*, 2021. [Link]
- A. Angelopoulos, S. Bates, J. Malik, and M. I. Jordan. “Uncertainty Sets for Image Classifiers using Conformal Prediction.” *ICLR*, 2021. [Link]
- K. Patel, W. H. Beluch, B. Yang, M. Pfeiffer, D. Zhang. “Multi-Class Uncertainty Calibration via Mutual Information Maximization-based Binning.” *ICLR 2021*. [Link]
- B. Adlam, J. Lee, L. Xiao, J. Pennington, J. Snoek. “Exploring the Uncertainty Properties of Neural Networks’ Implicit Priors in the Infinite-Width Limit.” *ICLR 2021*. [Link]
- A. Harakeh, S. L. Waslander. “Estimating and Evaluating Regression Predictive Uncertainty in Deep Object Detectors.” *ICLR 2021*. [Link]
- J. Antoran, U. Bhatt, T. Adel, A. Weller, J. M. Hernández-Lobato. “Getting a CLUE: A Method for Explaining Uncertainty Estimates.” *ICLR 2021*. [Link]

-
- A.-K. Kopetzki, B. Charpentier, D. Zügner, S. Giri, S. Günnemann. “Evaluating Robustness of Predictive Uncertainty Estimation: Are Dirichlet-based Models Reliable?” ICML 2021. [Link]
 - A. Zhou and S. Levine. “Amortized Conditional Normalized Maximum Likelihood: Reliable Out of Distribution Uncertainty Estimation.” ICML, 2021. [Link]
 - M. Havasi, R. Jenatton, S. Fort, J. Z. Liu, J. Snoek, B. Lakshminarayanan, A. M. Dai, and D. Tran. “Training independent subnetworks for robust prediction.” ICLR, 2021. [Link]
 - B. Adlam, J. Lee, L. Xiao, J. Pennington, J. Snoek. “Exploring the Uncertainty Properties of Neural Networks’ Implicit Priors in the Infinite-Width Limit”. ICLR, 2021. [Link]
 - A. N. Angelopoulos, S. Bates, T. Zrnic, M. I. Jordan. “Private Prediction Sets.” arXiv, 2021. [Link]
 - B. Charpentier, D. Zügner, S. Günnemann. “Posterior Network: Uncertainty Estimation without OOD Samples via Density-Based Pseudo-Counts.” NeurIPS, 2020. [Link]
 - L. Meronen, C. Irwanto, A. Solin. “Stationary Activations for Uncertainty Calibration in Deep Learning.” NeurIPS, 2020. [Link]
 - F. Wenzel, J. Snoek, D. Tran, R. Jenatton. “Hyperparameter Ensembles for Robustness and Uncertainty Quantification.” NeurIPS, 2020. [Link]
 - J. Liu, Z. Lin, S. Padhy, D. Tran, T. Bedrax Weiss, B. Lakshminarayanan. “Simple and Principled Uncertainty Estimation with Deterministic Deep Learning via Distance Awareness.” NeurIPS, 2020. [Link]
 - J. Lindinger, D. Reeb, C. Lippert, B. Rakitsch. “Beyond the Mean-Field: Structured Deep Gaussian Processes Improve the Predictive Uncertainties.” NeurIPS, 2020. [Link]
 - J. Antoran, J. Allingham, J. M. Hernández-Lobato. “Depth Uncertainty in Neural Networks.” NeurIPS, 2020. [Link]
 - M. Monteiro, L. Le Folgoc, D. C. de Castro, N. Pawłowski, B. Marques, K. Kamnitsas, M. van der Wilk, B. Glocker. “Stochastic Segmentation Networks: Modelling Spatially Correlated Aleatoric Uncertainty.” NeurIPS, 2020. [Link]
 - W. Shi, X. Zhao, F. Chen, Q. Yu. “Multifaceted Uncertainty Estimation for Label-Efficient Deep Learning.” NeurIPS, 2020. [Link]
 - R. Krishnan, O. Tickoo. “Improving model calibration with accuracy versus uncertainty optimization.” NeurIPS, 2020. [Link]
 - J. A. Leonard, M. A. Kramer, and L. H. Ungar. “A neural network architecture that computes its own reliability.” Computers & chemical engineering, 1992. [Link]
 - C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. “Weight uncertainty in neural networks.” ICML, 2015. [Link]

-
- B. Lakshminarayanan, A. Pritzel, and C. Blundell. “Simple and scalable predictive uncertainty estimation using deep ensembles.” NeurIPS, 2017. [Link]
 - Y. Gal and Z. Ghahramani. “Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning.” ICML, 2016. [Link]
 - V. Kuleshov, N. Fenner, and S. Ermon. “Accurate Uncertainties for Deep Learning Using Calibrated Regression.” ICML, 2018. [Link]
 - J. Hernández-Lobato and R. Adams. “Probabilistic backpropagation for scalable learning of bayesian neural networks.” ICML, 2015. [Link]
 - S. Liang, Y. Li, and R. Srikant. “Enhancing The Reliability of Out-of-distribution Image Detection in Neural Networks.” ICLR, 2018. [Link]
 - K. Lee, H. Lee, K. Lee, and J. Shin. “Training Confidence-calibrated classifiers for detecting out-of-distribution samples.” ICLR, 2018. [Link]
 - P. Schulam and S. Saria “Can You Trust This Prediction? Auditing Pointwise Reliability After Learning.” AISTATS, 2019. [Link]
 - A. Malinin and M. Gales. “Predictive uncertainty estimation via prior networks.” NeurIPS, 2018. [Link]
 - D. Hendrycks, M. Mazeika, and T. G. Dietterich. “Deep anomaly detection with outlier exposure.” arXiv preprint arXiv:1812.04606, 2018. [Link]
 - A-A. Papadopoulos, M. R. Rajati, N. Shaikh, and J. Wang. “Outlier exposure with confidence control for out-of-distribution detection.” arXiv preprint arXiv:1906.03509, 2019. [Link]
 - D. Madras, J. Atwood, A. D’Amour, “Detecting Extrapolation with Influence Functions.” ICML Workshop on Uncertainty and Robustness in Deep Learning, 2019. [Link]
 - M. Sensoy, L. Kaplan, and M. Kandemir. “Evidential deep learning to quantify classification uncertainty.” NeurIPS, 2018. [Link]
 - W. Maddox, T. Garipov, P. Izmailov, D. Vetrov, and A. G. Wilson. “A simple baseline for bayesian uncertainty in deep learning.” arXiv preprint arXiv:1902.02476, 2019. [Link]
 - Y. Ovadia, et al. “Can You Trust Your Model’s Uncertainty? Evaluating Predictive Uncertainty Under Dataset Shift.” arXiv preprint arXiv:1906.02530, 2019. [Link]
 - D. Hendrycks, et al. “Using Self-Supervised Learning Can Improve Model Robustness and Uncertainty.” arXiv preprint arXiv:1906.12340, 2019. [Link]
 - A. Kumar, P. Liang, T. Ma. “Verified Uncertainty Calibration.” arXiv preprint, 2019. [Link]
 - I. Osband, C. Blundell, A. Pritzel, and B. Van Roy. “Deep Exploration via Bootstrapped DQN.” NeurIPS, 2016. [Link]

-
- I. Osband. “Risk versus Uncertainty in Deep Learning: Bayes, Bootstrap and the Dangers of Dropout.” NeurIPS Workshop, 2016. [Link]
 - J. Postels et al. “Sampling-free Epistemic Uncertainty Estimation Using Approximated Variance Propagation.” ICCV, 2019. [Link]
 - A. Kendall and Y. Gal. “What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?” NeurIPS, 2017. [Link]
 - N. Tagasovska and D. Lopez-Paz. “Single-Model Uncertainties for Deep Learning.” NeurIPS, 2019. [Link]
 - A. Der Kiureghian and O. Ditlevsen. “Aleatory or Epistemic? Does it Matter?.” Structural Safety, 2009. [Link]
 - D. Hafner, D. Tran, A. Irpan, T. Lillicrap, and J. Davidson. “Reliable uncertainty estimates in deep neural networks using noise contrastive priors.” arXiv, 2018. [Link]
 - S. Depeweg, J. M. Hernández-Lobato, F. Doshi-Velez, and S. Udluft. “Decomposition of uncertainty in Bayesian deep learning for efficient and risk-sensitive learning.” ICML, 2018. [Link]
 - L. Smith and Y. Gal, “Understanding Measures of Uncertainty for Adversarial Example Detection.” UAI, 2018. [Link]
 - L. Zhu and N. Laptev. “Deep and Confident Prediction for Time series at Uber.” IEEE International Conference on Data Mining Workshops, 2017. [Link]
 - M. W. Dusenberry, G. Jerfel, Y. Wen, Yi-an Ma, J. Snoek, K. Heller, B. Lakshminarayanan, D. Tran. “Efficient and Scalable Bayesian Neural Nets with Rank-1 Factors.” arXiv, 2020. [Link]
 - J. van Amersfoort, L. Smith, Y. W. Teh, and Y. Gal. “Uncertainty Estimation Using a Single Deep Deterministic Neural Network.” ICML, 2020. [Link]
 - E. Begoli, T. Bhattacharya and D. Kusnezov. “The need for uncertainty quantification in machine-assisted medical decision making.” Nature Machine Intelligence, 2019. [Link]
 - T. S. Salem, H. Langseth, and H. Ramampiaro. “Prediction Intervals: Split Normal Mixture from Quality-Driven Deep Ensembles.” UAI, 2020. [Link]
 - K. Posch, and J. Pilz, “Correlated Parameters to Accurately Measure Uncertainty in Deep Neural Networks.” IEEE Transactions on Neural Networks and Learning Systems, 2020. [Link]
 - B. Kompa, J. Snoek, and A. Beam. “Empirical Frequentist Coverage of Deep Learning Uncertainty Quantification Procedures.” arXiv, 2020. [Link]

Predictive uncertainty in sequential models

-
- A. Malinin, M. Gales. “Uncertainty Estimation in Autoregressive Structured Prediction.” ICLR, 2021. [Link]
 - C. Wang, C. Lawrence, M. Niepert. “Uncertainty Estimation and Calibration with Finite-State Probabilistic RNNs.” ICLR, 2021. [Link]
 - S. Sankaranarayanan, Y. Chou, E. Goubault, S. Putot. “Reasoning about Uncertainties in Discrete-Time Dynamical Systems using Polynomial Forms.” NeurIPS, 2020. [Link]
 - R. Wen, K. Torkkola, B. Narayanaswamy, and D. Madeka. “A Multi-horizon Quantile Recurrent Forecaster.” arXiv, 2017. [Link]
 - D. T. Mirikitani and N. Nikolaev. “Recursive bayesian recurrent neural networks for time-series modeling.” IEEE Transactions on Neural Networks, 2009. [Link]
 - M. Fortunato, C. Blundell and O. Vinyals. “Bayesian Recurrent Neural Networks.” arXiv, 2019. [Link]
 - P. L. McDermott, C. K. Wikle. “Bayesian Recurrent Neural Network Models for Forecasting and Quantifying Uncertainty in Spatial-Temporal Data.” Entropy, 2019. [Link]
 - Y. Gal, Z. Ghahramani. “A theoretically grounded application of dropout in recurrent neural networks.” NeurIPS, 2016. [Link]