International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

Uncertainty Quantification in Machine Learning Predictions: A Bayesian Statistical Perspective

Patil Sunil Murlidhar
Research Scholar, Department of Statistics, Malwanchal University, Indore
Dr. Shoyeb Ali Sayed
Supervisor, Department of Statistics, Malwanchal University, Indore

Abstract
Machine learning (ML) has achieved remarkable success across domains such as healthcare, finance, climate science, and autonomous systems, yet the absence of reliable mechanisms to communicate predictive confidence remains a critical limitation for safe and trustworthy deployment. Uncertainty quantification (UQ) addresses this gap by enabling models to measure both aleatoric uncertainty, arising from data noise, and epistemic uncertainty, stemming from model limitations. Among various approaches, Bayesian statistics provides a principled and mathematically grounded framework for integrating prior knowledge, modeling probability distributions over parameters, and producing uncertainty-aware predictions. This paper explores the theoretical foundations of Bayesian UQ and reviews key methodologies including Monte Carlo sampling, variational inference, Gaussian processes, and Bayesian neural networks. Practical applications are examined across multiple high-stakes domains, demonstrating how Bayesian UQ enhances decision-making, fosters transparency, and aligns with ethical and regulatory standards. The study further highlights challenges such as computational cost, scalability, and prior selection, while emphasizing recent advances in approximate Bayesian methods and deep learning integration. By bridging the gap between accuracy and interpretability, Bayesian UQ strengthens trust in machine learning systems and paves the way for responsible AI adoption. Future directions are outlined, focusing on scalable frameworks, hybrid techniques, and integration with explainable AI to ensure robustness in increasingly complex environments.

Keywords: Bayesian statistics, uncertainty quantification, machine learning, probabilistic modeling, responsible AI

Introduction

The rapid advancement of machine learning (ML) has revolutionized numerous domains such as healthcare, finance, engineering, climate science, and autonomous systems, where predictive models are increasingly used to support high-stakes decision-making. Despite their growing success, traditional ML models are often criticized for acting as “black boxes” that provide

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

deterministic outputs without communicating the degree of confidence associated with their predictions. This lack of interpretability and reliability poses significant risks, especially in safety-critical applications where wrong decisions can have severe consequences. For instance, in medical diagnosis, an algorithm may correctly classify a disease in most cases, but without knowing the uncertainty level, a clinician cannot assess whether to rely on the prediction for life-altering treatment. Similarly, in autonomous driving or financial forecasting, unquantified uncertainty can lead to catastrophic failures or financial losses. Hence, uncertainty quantification (UQ) has emerged as an essential tool in bridging the gap between high-performing algorithms and their safe, trustworthy deployment. Among the different approaches, Bayesian statistics provides a mathematically principled framework for quantifying uncertainty, enabling ML models to incorporate prior knowledge, model variability, and account for both epistemic (model-based) and aleatoric (data-based) uncertainties.

Bayesian methods offer a structured way to estimate probability distributions over parameters and predictions, rather than single-point estimates, which helps capture the full range of possible outcomes. This probabilistic view allows practitioners to not only achieve accurate predictions but also measure the credibility of those predictions, improving interpretability and decision-making under uncertainty. Bayesian inference, Monte Carlo sampling, variational methods, and Gaussian processes are some of the key techniques used to integrate uncertainty quantification within ML pipelines. While these methods can be computationally intensive, recent advancements in scalable Bayesian approximations and deep learning integration have significantly improved their feasibility for large datasets and complex models. Furthermore, UQ powered by Bayesian approaches fosters robustness, transparency, and accountability in AI systems, aligning with the growing global emphasis on ethical and responsible AI. In this context, the present study examines the role of Bayesian statistics in uncertainty quantification for machine learning predictions, highlighting methodological frameworks, comparative strengths, challenges, and real-world applications. By doing so, it seeks to emphasize not only the importance of UQ in improving predictive reliability but also the broader implications for trust, safety, and adoption of AI technologies across diverse industries.

Background and Motivation

The growth of machine learning (ML) over the last decade has been transformative, enabling intelligent systems to perform tasks such as image recognition, natural language processing, financial forecasting, and autonomous navigation with remarkable accuracy. These advancements have created a shift in how decisions are made in domains ranging from healthcare to transportation and finance. However, despite their performance, conventional ML models often lack the ability to communicate how confident they are in their predictions. Most models provide deterministic outputs—single values or labels—without indicating the

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

underlying level of uncertainty. This absence of transparency presents a significant limitation in real-world applications where wrong decisions may lead to costly or even life-threatening consequences. For instance, a medical diagnosis system suggesting the presence of cancer without an accompanying uncertainty measure may leave physicians unable to judge whether the result should be trusted or verified through additional tests. Similarly, autonomous vehicles must not only detect obstacles but also assess the confidence of those detections to avoid accidents. The motivation for uncertainty quantification (UQ) thus lies in bridging the gap between predictive accuracy and trustworthiness. By incorporating probabilistic reasoning into ML through Bayesian statistics, models can capture both data-driven randomness and uncertainty arising from limited knowledge. This perspective fosters safe deployment, enhances interpretability, and aligns with the growing global demand for responsible and ethical AI applications.

Importance of Uncertainty Quantification in ML

Uncertainty quantification (UQ) in machine learning plays a pivotal role in transforming predictive models from black-box systems into interpretable and trustworthy decision-support tools. In many high-stakes applications, it is not sufficient to simply provide a prediction; rather, it is essential to understand the confidence associated with that prediction. UQ addresses this by distinguishing between two major types of uncertainty: aleatoric uncertainty, which arises from inherent variability or noise in the data, and epistemic uncertainty, which originates from limitations in the model's knowledge or parameters. By quantifying both, practitioners gain deeper insight into model reliability and potential risks. For example, in medical imaging, a model may detect a tumor with high uncertainty, alerting doctors that further evaluation is necessary, whereas low uncertainty predictions may justify direct action. Beyond safety, UQ also contributes to resource optimization by enabling risk-sensitive decision-making in domains such as finance, supply chain management, and energy systems. In machine learning research, UQ serves as a foundation for model calibration, out-of-distribution detection, and active learning, where uncertainty guides data acquisition strategies. Bayesian statistics provides a natural framework for this, offering a probabilistic view of both parameters and predictions. The importance of UQ is also reflected in its role in explainable AI, regulatory compliance, and ethical deployment of intelligent systems, as policymakers and end-users increasingly demand transparency, accountability, and trust in AI-powered decision-making processes.

Scope and Objectives of the Study

The present study focuses on exploring the role of Bayesian statistics in quantifying uncertainty within machine learning predictions, with the aim of providing a comprehensive understanding of both theoretical foundations and practical applications. The scope encompasses an examination of how Bayesian frameworks model uncertainty, the techniques employed—such as

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

Monte Carlo methods, variational inference, Gaussian processes, and Bayesian neural networks—and their applicability across different types of machine learning tasks. The study further addresses the comparative strengths and limitations of these approaches, highlighting trade-offs between computational cost, scalability, and interpretability. Additionally, real-world use cases are considered, spanning healthcare diagnostics, financial forecasting, autonomous systems, and climate modeling, to demonstrate the practical value of Bayesian UQ in diverse domains. The objectives are threefold: first, to critically analyze how Bayesian methods capture epistemic and aleatoric uncertainties; second, to evaluate the performance and reliability of Bayesian UQ techniques against traditional deterministic ML models; and third, to identify ongoing challenges and future directions in this emerging field. By addressing these objectives, the study aims to contribute to the broader discourse on responsible AI, showing how Bayesian UQ can enhance the safety, transparency, and ethical deployment of machine learning systems. This research seeks to provide scholars, practitioners, and policymakers with insights that encourage the adoption of uncertainty-aware models in high-impact decision-making environments.

Need of the Study

Machine learning has become a critical enabler in decision-making across various domains, yet the absence of reliable uncertainty estimates continues to pose challenges for real-world adoption. In high-stakes applications—such as healthcare diagnostics, financial risk assessment, or autonomous navigation—errors caused by overconfident models can lead to severe consequences (Pathiravasan et al., 2020). Traditional predictive models focus primarily on maximizing accuracy or minimizing loss functions, but they often fail to account for the confidence associated with their outputs. This neglect creates a gap between model performance in controlled environments and reliability in dynamic, real-world contexts. Thus, there is an urgent need to develop robust uncertainty quantification (UQ) frameworks that can make ML predictions both interpretable and trustworthy.

Another factor driving the necessity of UQ is the inherent variability and imperfections of data. Real-world datasets are often noisy, incomplete, and subject to distributional shifts. Models trained on such data may offer misleadingly precise predictions, thereby propagating uncertainty into critical decision processes (Ober & Rasmussen, 2020). For instance, in clinical settings, two patients with similar medical histories may respond differently to treatments. Without capturing predictive uncertainty, ML systems cannot effectively convey the risks of false positives or false negatives, which are central to medical decision-making (Antorán et al., 2020). This underscores the importance of methodologies that quantify both aleatoric and epistemic uncertainties, allowing for a more comprehensive evaluation of prediction reliability.

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

The growing demand for transparency and interpretability in AI systems further strengthens the need for uncertainty-aware models. Regulatory bodies and policymakers increasingly emphasize the importance of explainability in AI, particularly in safety-critical industries (Bhatt et al., 2020). Bayesian methods, by producing full predictive distributions rather than point estimates, align with these requirements and foster greater trust in ML systems. Moreover, they enable practitioners to distinguish between areas where the model is confident and scenarios where more data or expert oversight is required. This form of interpretability is vital for risk-sensitive environments, where decisions must balance accuracy with the quantified risk of uncertainty.

From a research perspective, the integration of Bayesian statistics into ML highlights a promising path for bridging the gap between theoretical rigor and practical applicability. Traditional deterministic models often lack mechanisms to adapt to evolving data streams or distributional changes. In contrast, Bayesian frameworks naturally accommodate new data through posterior updates, making them well-suited for applications where knowledge evolves continuously (Liu et al., 2020). This adaptability ensures that predictions remain relevant and calibrated in rapidly changing environments such as financial markets, climate systems, or public health surveillance.

Uncertainty quantification plays a pivotal role in decision-theoretic optimization. In many applications, it is not sufficient to know the most likely outcome; decision-makers must also understand the risks of alternative outcomes. By embedding uncertainty into prediction pipelines, Bayesian ML provides a principled framework for risk-sensitive optimization, active learning, and resource allocation (Abdar et al., 2020). For example, in active learning scenarios, models with well-calibrated uncertainty estimates can identify instances where additional data collection is most beneficial, thereby reducing labeling costs and improving generalization.

The need for this study arises from three critical gaps: (1) the lack of reliable uncertainty measures in traditional ML models, (2) the inability of current methods to address data imperfections and distributional shifts, and (3) the growing demand for interpretability and risk-sensitive decision-making. By focusing on Bayesian statistics as a foundation for uncertainty quantification, this research aims to address these challenges and contribute to the development of safe, trustworthy, and adaptive machine learning systems.

Literature Review

Uncertainty quantification (UQ) in machine learning has become a prominent research focus in recent years, particularly with the growing need for reliable and trustworthy artificial intelligence (AI) systems. Existing literature spans classical Bayesian theory, scalable inference methods, calibration techniques, and application-specific studies. This section synthesizes contributions from foundational works to recent advancements, drawing upon scholarly studies that collectively highlight the role of Bayesian statistics in enhancing predictive reliability.

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

The theoretical basis for UQ in ML originates in Bayesian probability, which treats parameters as random variables with distributions rather than fixed values. Gelman et al. (2013) outlined Bayesian inference as a principled framework for updating beliefs in light of new data, establishing the foundation for probabilistic ML. This perspective contrasts with deterministic approaches that yield only point estimates. Further, Berger (2013) emphasized the decision-theoretic role of Bayesian methods, where actions are guided not only by expected outcomes but also by associated uncertainties.

Murphy (2012) provided a modern introduction to probabilistic modeling, framing Bayesian inference as essential for learning under uncertainty. Neal (2012) expanded this understanding by demonstrating how Bayesian methods can be applied to neural networks, introducing Bayesian neural networks (BNNs) that explicitly model uncertainty in weights. These works collectively positioned Bayesian inference as the theoretical cornerstone for integrating uncertainty into ML systems.

A recurring theme in literature is distinguishing between different types of uncertainty. Der Kiureghian and Ditlevsen (2009) categorized uncertainty into aleatoric (inherent randomness in data) and epistemic (lack of knowledge due to limited data or imperfect models). Hüllermeier and Waegeman (2020) extended this classification to ML, arguing that robust systems must quantify both forms to ensure reliability in real-world tasks.

In applied contexts, Begoli et al. (2019) emphasized the need for UQ in medical decision-making, where overconfident predictions could lead to harmful outcomes. Similarly, Ghosh et al. (2018) demonstrated how ignoring uncertainty in disease detection models can result in unreliable outputs. These studies underscore the critical importance of capturing multiple sources of uncertainty, especially in safety-critical domains.

The challenge of Bayesian inference lies in its computational intractability for high-dimensional models. To address this, Blei et al. (2017) provided a comprehensive review of variational inference (VI), highlighting its ability to approximate posteriors efficiently. Blundell et al. (2015) advanced this area by introducing Bayes by Backprop, an algorithm for training BNNs using stochastic variational methods. These innovations enabled Bayesian approaches to scale with modern deep learning architectures.

Wang and Yeung (2020) surveyed Bayesian deep learning techniques, focusing on computational strategies such as stochastic gradient methods and probabilistic programming. Their analysis demonstrated how approximate inference can make Bayesian ML feasible in practice. Similarly, Tran et al. (2019) explored distributed probabilistic programming, showcasing methods for accelerating inference in large-scale models. Together, these works highlight how Bayesian inference has evolved from theoretical rigor to practical applicability.

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

Recent studies emphasize the necessity of uncertainty-aware deep learning. Lakshminarayanan et al. (2017) proposed deep ensembles as a simple yet effective approach to estimate predictive uncertainty, outperforming many Bayesian approximations in practice. Gal and Ghahramani (2016) introduced dropout as a Bayesian approximation, a breakthrough that connected regularization with uncertainty estimation in deep networks.

Wilson (2020) built upon these advancements by advocating for Bayesian deep learning as a more robust alternative to deterministic neural networks, especially under distributional shifts. Michelmore et al. (2018) evaluated uncertainty in autonomous driving systems, demonstrating that models with uncertainty estimates perform better in safety-critical decision-making. These studies collectively affirm that uncertainty-aware deep learning methods improve both performance and robustness.

Calibration is another central theme in UQ research. Guo et al. (2017) showed that modern deep networks are often poorly calibrated, producing overconfident predictions. To address this, Mukhoti et al. (2020) proposed calibration through focal loss, improving the alignment between predicted probabilities and observed frequencies. Malinin and Gales (2020) explored calibration in natural language processing tasks, emphasizing uncertainty as a key factor in interpretability.

Molnar (2019) contributed to the interpretability literature by framing uncertainty as a pathway to explainability in ML. Bhatt et al. (2020) similarly argued that uncertainty should be viewed as a form of transparency, essential for communicating the reliability of predictions. Together, these contributions underscore how Bayesian UQ enhances both interpretability and trust.

Abdar et al. (2020) and Abdar et al. (2020) explored UQ in medical image analysis, showing that Bayesian ML models can improve diagnostic reliability by communicating uncertainty to clinicians. Antorán et al. (2020) introduced CLUE, a method for explaining uncertainty estimates, specifically tailored to healthcare scenarios. These contributions highlight how UQ supports informed clinical decision-making and risk management.

Amini et al. (2020) presented deep evidential regression, an approach that directly models predictive distributions, enhancing decision-making in navigation and control tasks. Their findings demonstrated that uncertainty-aware models reduce catastrophic errors in autonomous driving.

Kuhn et al. (2019) explored distributionally robust optimization, integrating uncertainty into financial decision-making frameworks. O'Hagan (2019) addressed uncertainty in expert elicitation, providing a scientific basis for incorporating human knowledge into policy-oriented ML systems. These works illustrate how Bayesian UQ extends beyond technical modeling to shape evidence-based governance.

Foong et al. (2019) benchmarked Bayesian deep learning methods, revealing limitations and strengths across inference techniques. Their evaluation demonstrated that Bayesian models,

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

when properly implemented, provide reliable predictive intervals for scientific data analysis. Similarly, Ober and Rasmussen (2020) discussed the promises and pitfalls of deep kernel learning, stressing the importance of UQ in extending Gaussian process models to complex scientific domains.

Applications of Bayesian UQ in Machine Learning

Healthcare and Medical Diagnostics

Bayesian UQ is crucial in healthcare, where decisions directly impact patient safety. Medical imaging models enhanced with Bayesian inference not only detect anomalies like tumors but also indicate prediction confidence, allowing doctors to judge whether additional tests are needed. In personalized medicine, Bayesian methods account for variability in patient responses, ensuring safer treatment recommendations. In critical care, predictive models with quantified uncertainty help clinicians anticipate patient deterioration and intervene early. By incorporating both aleatoric and epistemic uncertainties, Bayesian UQ builds trust in AI-assisted diagnostics, aligning with ethical and regulatory demands for transparency in healthcare systems.

Financial Forecasting and Risk Management

Financial markets are highly uncertain, making reliable forecasts essential. Bayesian UQ strengthens ML applications in credit scoring, stock prediction, and portfolio optimization by offering probabilistic forecasts instead of deterministic outputs. For example, Bayesian neural networks provide predictive intervals for stock prices, helping investors weigh risks. In credit risk, Bayesian methods distinguish between high-confidence and low-confidence evaluations, reducing misclassification. They also support stress testing, where institutions evaluate performance under extreme scenarios. By capturing uncertainty, Bayesian UQ enhances transparency and trust, enabling financial systems to balance profitability with risk control.

Autonomous Systems and Robotics

Autonomous systems operate in unpredictable environments where errors may cause accidents. Bayesian UQ equips robots and self-driving cars with uncertainty-aware decision-making, ensuring safer operation. In perception tasks, such as obstacle detection, uncertainty scores help vehicles slow down or seek human intervention when confidence is low. In robotic manipulation, Bayesian models guide grasping and navigation under uncertain conditions. Similarly, drones use UQ for trajectory planning in dynamic weather. Beyond safety, UQ supports reinforcement learning by balancing exploration and risk. Thus, Bayesian UQ is central to building trustworthy, resilient, and interpretable autonomous systems.

Climate Modeling and Environmental Predictions

Climate science relies on models that inherently involve uncertainty due to incomplete data and complex dynamics. Bayesian UQ addresses this by integrating diverse datasets, such as satellite records and ground sensors, while quantifying both epistemic and aleatoric uncertainties. This

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

results in probabilistic climate forecasts that inform disaster preparedness and infrastructure planning. For instance, Bayesian UQ helps identify high-risk regions for floods or extreme weather events. It also strengthens ensemble modeling by weighting predictions according to uncertainty levels. By making forecasts transparent and interpretable, Bayesian methods enable policymakers to make better-informed environmental decisions.

Industrial and Engineering Applications

In industry, predictive accuracy must be paired with reliability to avoid costly failures. Bayesian UQ supports predictive maintenance by quantifying uncertainty in equipment failure forecasts, allowing timely interventions. In manufacturing, it improves quality control by flagging uncertain defect classifications for further inspection. Engineering design also benefits, as Bayesian methods capture uncertainty in simulations, leading to safer and more efficient structures. Furthermore, UQ supports active learning, reducing data collection costs in fields like aerospace. By embedding probabilistic reasoning, industries gain not only efficiency but also transparency and resilience in critical operations.

Research Problem

Despite remarkable progress in machine learning (ML), a persistent challenge is the lack of reliable uncertainty quantification (UQ) in predictive models. Most conventional ML methods provide point estimates or probability scores without adequately capturing the uncertainty associated with those predictions (Lakshminarayanan et al., 2017). This limitation often results in overconfident outputs, which can be misleading and potentially harmful in high-stakes decision-making contexts such as medical diagnostics, autonomous driving, and financial forecasting.

Uncertainty in ML arises from multiple sources: data noise and variability (aleatoric uncertainty) and model limitations or lack of knowledge (epistemic uncertainty) (Der Kiureghian & Ditlevsen, 2009; Hüllermeier & Waegeman, 2020). Traditional deterministic models struggle to account for these uncertainties, limiting their reliability when applied to dynamic, real-world environments. While ensemble methods and calibration techniques provide partial solutions, they often lack theoretical rigor and scalability (Guo et al., 2017).

Bayesian statistics provides a principled framework for modeling uncertainty by representing parameters as probability distributions and updating beliefs with observed data (Gelman et al., 2013). However, the computational complexity of Bayesian inference in large-scale ML models presents significant challenges (Neal, 2012). Approximate inference methods such as variational inference and Bayesian deep learning (Blei et al., 2017; Blundell et al., 2015) attempt to address these limitations, but questions remain regarding their efficiency, calibration, and interpretability in practical applications.

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

Conclusion

Uncertainty quantification (UQ) has emerged as a vital component in advancing machine learning from high-performing yet opaque models to trustworthy, interpretable, and ethically deployable systems. This study highlights how Bayesian statistics, with its probabilistic framework, provides a principled approach for modeling both aleatoric and epistemic uncertainties, thereby enabling predictions that communicate not only outcomes but also confidence levels. The review of theoretical underpinnings and practical methodologies—ranging from Monte Carlo sampling and variational inference to Bayesian neural networks and Gaussian processes—demonstrates that Bayesian UQ offers robust solutions across diverse application domains, including healthcare, finance, autonomous systems, climate modeling, and industrial engineering. These methods empower practitioners to make risk-aware decisions, optimize resources, and build transparency into AI workflows, thus bridging the critical gap between accuracy and reliability. Despite challenges such as computational complexity, prior selection, and scalability, rapid progress in approximate Bayesian methods and deep learning integration is making UQ increasingly feasible for real-world systems. Furthermore, its alignment with explainable AI and regulatory compliance strengthens the case for broader adoption. Ultimately, Bayesian UQ not only enhances model performance but also promotes accountability and trust in AI-driven decision-making, ensuring that machine learning technologies can be responsibly applied to high-stakes environments. Future research should focus on improving scalability, developing hybrid methods, and integrating Bayesian UQ with interpretable and ethical AI frameworks, thereby solidifying its role as a cornerstone of reliable and human-centered artificial intelligence.

References

1. Abdar, M., et al. (2020). A review of uncertainty quantification in deep learning: Techniques, applications and challenges. Information Fusion, 76, 243–297.
2. Abdar, M., Pourpanah, F., Hussain, S., Rezazadegan, D., Liu, L., Ghavamzadeh, M., … Nahavandi, S. (2020). A review of uncertainty quantification in deep learning: Application to medical image analysis. Computers in Biology and Medicine, 140, 105063.
3. Antorán, J., Allingham, J. U., & Hernández-Lobato, J. M. (2020). Getting a CLUE: A method for explaining uncertainty estimates. arXiv preprint arXiv:2006.06848.
4. Begoli, E., Bhattacharya, T., & Kusnezov, D. (2019). The need for uncertainty quantification in machine-assisted medical decision making. Nature Machine Intelligence, 1(1), 20–23.
5. Berger, J. O. (2013). Statistical decision theory and Bayesian analysis. Springer.

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

6. Bhatt, U., Antorán, J., Zhang, Y., Liao, Q. V., Sattigeri, P., Fogliato, R., … Weller, A. (2020). Uncertainty as a form of transparency: Measuring, communicating, and using uncertainty. Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, 401–413.
7. Blei, D. M., Kucukelbir, A., & McAuliffe, J. D. (2017). Variational inference: A review for statisticians. Journal of the American Statistical Association, 112(518), 859–877.
8. Blundell, C., et al. (2015). Weight uncertainty in neural networks. Proceedings of ICML, 1613–1622.
9. Der Kiureghian, A., & Ditlevsen, O. (2009). Aleatory or epistemic? Does it matter? Structural Safety, 31(2), 105–112.
10. Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation. Proceedings of ICML, 1050–1059.
11. Gelman, A., et al. (2013). Bayesian data analysis (3rd ed.). CRC Press.
12. Gelman, A., et al. (2017). Prior choice recommendations. Bayesian Analysis, 12(3), 515–533.
13. Ghahramani, Z. (2015). Probabilistic machine learning and artificial intelligence. Nature, 521(7553), 452–459.
14. Ghosh, S., et al. (2018). Uncertainty estimation in deep learning for disease detection. arXiv preprint arXiv:1807.01771.
15. Hüllermeier, E., & Waegeman, W. (2020). Aleatoric and epistemic uncertainty in machine learning. Machine Learning, 110(3), 457–506.
16. Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. Science, 349(6245), 255–260.
17. Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. NeurIPS, 6402–6413.
18. Liu, H., Ong, Y. S., Shen, X., & Cai, J. (2020). When Gaussian process meets big data: A review of scalable GPs. IEEE Transactions on Neural Networks and Learning Systems, 31(11), 4405–4423.
19. Michelmore, R., et al. (2018). Evaluating uncertainty in deep learning for autonomous driving. arXiv preprint arXiv:1811.06817.
20. Molnar, C. (2019). Interpretable machine learning. Leanpub.
21. Murphy, K. P. (2012). Machine learning: A probabilistic perspective. MIT Press.
22. Neal, R. M. (2012). Bayesian learning for neural networks. Springer.
23. Ober, S. W., & Rasmussen, C. E. (2020). The promises and pitfalls of deep kernel learning. Journal of Machine Learning Research, 22(1), 1–68.

24. Pathiravasan, C. H., Miratrix, L. W., & Rubin, D. B. (2020). Uncertainty quantification in causal inference and predictive modeling. Annual Review of Statistics and Its Application, 8, 495–517.
25. Wang, H., & Yeung, D. Y. (2020). A survey on Bayesian deep learning. ACM Computing Surveys, 53(5), 1–37.
26. Wilson, A. G. (2020). The case for Bayesian deep learning. arXiv preprint arXiv:2001.10995