Mathematical Modeling and Algorithm Application
ISSN: 3006-0842 | Vol. 2, No. 3, 2024

Nondeterministic Features in Deep neural network design, training and inference

Jackob Antonion[1], Annie Wang [1, *], Maziar Raissi[2], Ribana Joshie[2]

[1]Department of Technology of Computers and Communications, University of Extremadura, 10071 Caceres, Spain

[2]Department of Applied Mathematics, The University of Sydney, Sydney, NSW 2006, Australia

Abstract: Neural networks are increasingly integral to scientific modeling and a wide range of real-world applications. However, standard neural networks often lack reliable certainty and confidence, and can be poorly calibrated, leading to nondeterministic results. To address these challenges, researchers have focused on understanding and quantifying uncertainty in neural network predictions. This work offers a comprehensive discussion of uncertainty estimation in neural networks, discussing recent advances, current challenges, and future research opportunities. This work explains key sources of uncertainty, challenges from model uncertainty and irreducible data uncertainty. We explore various methods for modeling these uncertainties, such as deterministic neural networks, Bayesian neural networks (BNNs), ensembles, and test-time data augmentation approaches. Additionally, the paper provides practical examples from fields like medical data analysis, robotics, and autonomous driving, illustrating the challenges and requirements associated with uncertainty in real-world applications. We also examine the limitations of current uncertainty quantification methods in safety-critical applications, and provide an outlook on future developments aimed at broader adoption of these methods in diverse domains. This paper serves as a valuable resource for both newcomers and those experienced in the field of uncertainty estimation in neural networks.

Keywords: Neural Networks, Machine Learning, Real-World Applications, Data Uncertainty, Deterministic Neural Networks.

1. Introduction

Over the past decade, deep neural networks (DNNs) have made significant strides, leading to their adoption across diverse research fields, particularly in the modeling and understanding of complex systems such as language model[1, 2], medical image analysis, robotics[3] and topology optimization[4]. For the security-critical tasks, the neutral networks have been used in pathological assessments in Hepatic steatosis[5-10], which influences transplant outcomes while the reliability of data source remains to be improved[11]. Despite their appeal in high-risk areas like medical imaging and autonomous vehicle control[12], DNNs remain underutilized in safety-critical real-world applications. This limited deployment is primarily due to:

Challenges in providing reliable uncertainty estimates for decisions[13].

Vulnerability to adversarial attacks, which exposes them to potential sabotage.

A lack of expressiveness and transparency in their inference models[14].

Inability to distinguish between in-domain and out-of-domain data.

These limitations often stem from uncertainties inherent in the data or gaps in the neural network's knowledge. Developing reliable uncertainty estimates is crucial, as these can help identify uncertain predictions, allowing them to be either disregarded or escalated to human experts. Uncertainty estimation is essential not only for safe decision-making in high-risk domains but also in fields with heterogeneous data sources and limited labeled data, such as sensing mm-Wave wireless signal[15-19] and traffic user equilibrium[12, 20] and autonomous driving[12, 21, 22]. Moreover, it plays a critical role in areas like active learning and reinforcement learning, where managing uncertainty is a fundamental aspect of the learning process. In recent popular transformer model plus its variants like detection Transformer[23] and Domain-Switch Learning[24], this uncertainty plays a great role in cross-attention and self-attention layers with various decode model that needs to be precisely estimated in feature extractor[25] and real-world applications like person re-identification[26-29], recognition of human actions[30-34], sentiment analysis[35, 36] , large language model[37, 38] and physics-informed reinforcement learning[39].

In recent years, there has been growing interest in estimating uncertainty in DNNs. The most common approach involves separately modeling two types of uncertainty: model uncertainty, which arises from the model itself, and data uncertainty (aleatoric uncertainty), which is inherent in the data. While model uncertainty can be reduced by improving the learning process, data uncertainty is irreducible. Key methods for modeling these uncertainties include Bayesian inference, ensemble techniques, test-time augmentation, and single deterministic networks with explicit components for representing both model and data uncertainty[40]. However, estimating predictive uncertainty alone is insufficient for safe decision-making; ensuring that these estimates are reliable is equally important. Consequently, the calibration properties of DNNs have been extensively studied, and re-calibration methods have been developed to enhance the reliability of uncertainty estimates.

Several studies have provided introductions and overviews of uncertainty in statistical modeling, though these works often lack a specific focus on neural network applications. For instance, researchers have introduced general frameworks and conceptual descriptions of Bayesian neural networks (BNNs)[41], along with updated overviews of Bayesian methods for uncertainty quantification in neural networks, particularly in recommender systems, topic modeling, and control systems. Evaluations of uncertainty quantification methods in deep learning have also been presented, comparing techniques such as SoftMax output ensemble networks, BNNs, and autoencoders on datasets like MNIST[40]. These studies highlight the practical challenges and considerations for applying uncertainty quantification methods in real-world, mission- and safety-critical applications.

Our aim is to guide the reader through the entire process—from identifying sources of uncertainty to understanding where uncertainty estimations are crucial. Additionally, we highlight the limitations of current methodologies and discuss future challenges. This paper offers a broad introduction and comparison of various approaches and foundational concepts, serving as a valuable resource for those familiar with deep learning who wish to integrate uncertainty estimation into their models, as well as for those already knowledgeable in the field.

2. Neural Networks with Nondeterministic Sources

A neural network is a non-linear function f_θ parameterized by model parameters θ (i.e. the network weights) that maps from a measurable input set A to a measurable output set B:

$$f_\theta: A \to B \text{ or } f_\theta(A) = B \qquad (1)$$

We consider four distinct steps from raw environmental information to a neural network's prediction with quantified uncertainties, namely:

Data Acquisition Process: The occurrence of some information in the environment (e.g., a bird’s singing) and a measured observation of this information (e.g., an audio recording).

DNN Building Process: The design and training of a neural network.

Applied Inference Model: The model applied for inference.

Prediction’s Uncertainty Model: The modeling of uncertainties caused by the neural network and/or the data.

Above highlights the nondeterministic features related to these steps and explain uncertainties are propagated through the process. Finally, a model for the uncertainty of a neural network’s prediction and discuss the main types of uncertainty considered in neural networks is introduced. The goal is to provide an accountable understanding of uncertainties in neural networks. For simplicity, the mathematical properties relevant to understanding the approaches and applying the methodology in different fields will be the focus for future research.

3. Challenges in Design and Training

The design of a DNN involves the explicit modeling of the neural network and its stochastic training process. The assumptions about the problem structure induced by the design and training of the neural network are referred to as inductive bias[42]. We summarize all decisions of the modeler regarding the network’s structure and training process in a structure configuration s. The defined network structure introduces a third factor of uncertainty in neural network predictions. For a given network structure s and a training dataset D, the training of a neural network is a stochastic process, and therefore, the resulting neural network f_Θ is based on a random variable.

The process is stochastic due to random decisions, such as the order of the data, random initialization, or random regularization (e.g., augmentation or dropout). The loss landscape of a neural network is highly non-linear, and the randomness in the training process generally leads to different local optima, resulting in different models. Parameters such as batch size, learning rate, and the number of training epochs also affect the training, resulting in different models. Depending on the underlying task, these models can significantly differ in their predictions for individual samples, even leading to differences in overall model performance. This sensitivity to the training process introduces a fourth factor for uncertainties in neural network predictions. The inference describes the prediction of an output for a new data sample by the neural network. At this time, the network is trained for a specific task. Thus, samples that are not inputs for this task cause errors and are therefore also a source of uncertainty.

In general, the distribution is unknown and can only be estimated based on the given data in D. For this estimation, neural networks form a very powerful tool for many tasks and applications. Neural networks form a powerful tool for many tasks and applications, but the prediction of a neural network is subject to both model-dependent and input data-dependent errors as illustrated in Figure1. Consequently, the predictive uncertainty associated with $y*$ is typically separated into data uncertainty (also known as statistical or aleatoric uncertainty) and model uncertainty. Depending on the approach, additional explicit modeling of distributional uncertainty may be used to account for uncertainty caused by examples from regions not covered by the training data.

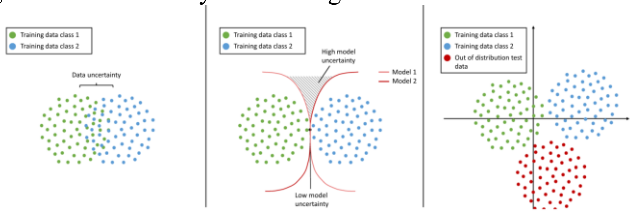

Figure 1 Visualization of the data, the model, and the distributional uncertainty for classification and regression models[34].

4. How to evaluation and address Nondeterministic Features

As described in Section 2, several factors may cause model and data uncertainty, affecting a DNN’s predictions. This variety of sources of uncertainty makes the complete exclusion of uncertainties in a neural network nearly impossible for most applications. Particularly in practical applications using real-world data, the training data is only a subset of all possible input data, which means that a mismatch between the DNN domain and the unknown actual data domain is often unavoidable. However, it is also impossible to compute an exact representation of a DNN prediction's uncertainty since the different uncertainties cannot be modeled accurately and are often even unknown. Therefore, methods for estimating uncertainty in DNN predictions are a popular and vital field of research.

The data uncertainty component is typically represented in the prediction (e.g., in the SoftMax output of a classification network or in the explicit prediction of a standard deviation in a regression network)[43]. In contrast, several different approaches that model the model uncertainty and seek to separate it from the data uncertainty have been introduced to obtain an accurate representation of the data uncertainty. Methods for estimating uncertainty can generally be split into four different types based on the number (single or multiple) and the nature (deterministic or stochastic) of the DNNs used:

Single Deterministic Methods: These methods provide predictions based on a single forward pass within a deterministic network. The uncertainty quantification is either derived using additional (external) methods or directly predicted by the network[44].

Bayesian Methods: These cover all kinds of stochastic DNNs, where two forward passes of the same sample generally lead to different results.

Ensemble Methods: These combine predictions from several different deterministic networks during inference.

Test-Time Augmentation Methods: These methods provide predictions based on a single deterministic network but augment the input data at test-time to generate several predictions that are used to evaluate the certainty of the prediction.

Many internal uncertainty quantification approaches follow the idea of predicting the parameters of a distribution over the predictions rather than a direct pointwise maximum-a-posteriori estimation. The loss function of such networks often takes the expected divergence between the true and predicted distributions into account. The distribution over the outputs can be interpreted as a quantification of the model uncertainty, attempting to emulate the behavior of Bayesian modeling of the network parameters. The prediction is then given as the expected value of the predicted distribution.

For regression tasks, Oala[45] introduced an uncertainty score based on the lower and upper bound output of an interval neural network. This network has the same structure as the underlying deterministic neural network and is initialized with the deterministic network's weights. Unlike Gaussian representations of uncertainty (given by a standard deviation), this approach can produce non-symmetric values of uncertainty and is more robust in the presence of noise. A simultaneous quantile regression loss function was introduced to generate well-calibrated prediction intervals for data uncertainty. Model uncertainty is quantified based on a mapping from the training data to zero using Orthonormal Certificates, ensuring that out-of-distribution samples, where the model is uncertain, are mapped to non-zero values and can thus be recognized.

Compared to other principles, single deterministic methods are computationally efficient in training and evaluation. Only one network needs to be trained, and the approaches can often be applied to pre-trained networks. Depending on the approach, only a single or at most two forward passes are required for evaluation. The underlying networks may contain more complex loss functions, which can slow down the training process, or external components that need to be trained and evaluated additionally. However, this is generally more efficient than the number of predictions required for ensemble-based methods, Bayesian methods, and test-time data augmentation methods. A drawback of single deterministic neural network approaches is their reliance on a single opinion, which can make them sensitive to the underlying network architecture, training procedure, and training data as shown in Figure1.

Various methods for modeling and predicting different types of uncertainty in neural networks will be presented. To assess these approaches, specific measures must be applied to the derived uncertainties. In the following section, we outline different methods for quantifying the predicted types of uncertainty. It is important to note that the accuracy and reliability of these uncertainties are not inherently guaranteed. Evaluating the quality of uncertainty estimates poses several challenges:

Method Dependency: The quality of uncertainty estimation is highly dependent on the underlying method used. Different approximations of Bayesian inference, for example, can lead to varying levels of accuracy in uncertainty estimates.

Absence of Ground Truth: There is a notable lack of ground truth uncertainty estimates, making it difficult to validate predictions. Defining ground truth uncertainty is itself a complex task. For example, if ground truth is defined based on human subjectivity, questions such as "How many subjects are needed?" and "How should these subjects be selected?" must be addressed.

Lack of Unified Metrics: There is no standardized quantitative evaluation metric across machine learning tasks. Uncertainty is defined differently depending on the task, such as classification, segmentation, or regression. For instance, prediction intervals or standard deviations are typically used to represent uncertainty in regression tasks, while measures like entropy are employed to capture uncertainty in classification and segmentation tasks.

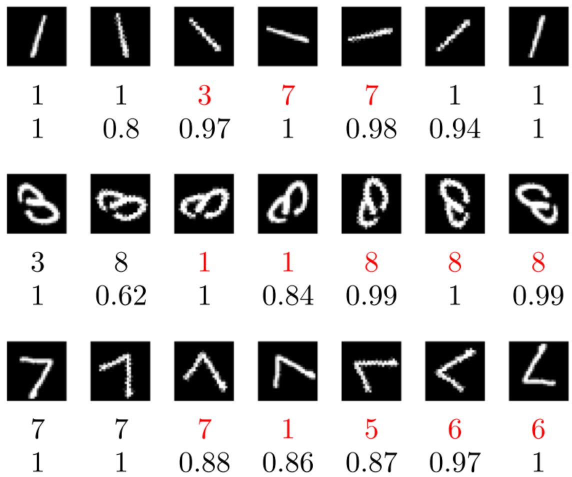

Figure2. Predictions received from a LeNet network trained on MNIST's handwritten digits from 0 to 9 and evaluated on different rotations of test samples[34].

Post-processing (or post-hoc) methods are applied after the training process and are designed to learn a re-calibration function. In this approach, a subset of the training data is held out during the training process and used as a calibration set. The re-calibration function is then applied to the network's outputs (e.g., the logit vector) to improve calibration based on the held-out calibration set. Zhang et al. [46] identified three key requirements that post-hoc calibration methods should meet:

Preserve Accuracy: The method should not compromise the predictor's performance.

Data Efficiency: Only a small fraction of the training dataset should be required for calibration.

Correct Re-Calibration Mapping: The method should be able to approximate the correct re-calibration map, provided there is sufficient data for calibration.

However, they also noted that no existing approach fully satisfies all three requirements.

For classification tasks, one of the simplest but effective post-hoc calibration techniques are temperature scaling. This involves optimizing the temperature T > 0 of the SoftMax function to achieve better calibration in below equation (2).

$$\mathrm{softmax}(z_i) = \frac{\exp^{z_i/T}}{\sum_{j=1}^{K} \exp^{z_j/T}} \tag{2}$$

This distinction is evident as these methods quantify model and data uncertainty separately, focusing on reducing model uncertainty in predictions. In addition to approaches that enhance calibration by addressing model uncertainty, a significant and expanding body of literature has explored methods specifically aimed at reducing calibration errors.

5. Conclusion and outlook

Despite significant advances in uncertainty quantification in neural networks over recent years, their adoption in practical mission- and safety-critical applications remains limited. While DNNs have become the main standard for solving numerous computer vision and medical image processing tasks, most existing models are not capable of appropriately quantifying the uncertainty inherent in their inferences, particularly in real-world applications. This is primarily because baseline models are often developed using standard datasets, such as ImageNet, or well-known regression datasets specific to particular use cases, and are therefore not readily applicable to complex real-world environments, such as low-resolution satellite data or other data sources affected by noise.

Although many researchers from other fields apply uncertainty quantification, a broad and structured evaluation of existing methods based on different real-world applications is not yet available. Existing methods for evaluating estimated uncertainty are better suited for comparing uncertainty quantification methods based on measurable quantities, such as calibration or performance on out-of-distribution (OOD) detection. However, a clear, standardized protocol of tests for uncertainty quantification methods is still lacking. This absence makes it difficult for researchers from other domains to identify state-of-the-art methods relevant to their field and complicates the direct comparison of the latest approaches. Moreover, the limited acceptance and adoption of currently existing methods for uncertainty quantification further highlight this issue.

Existing measures for evaluating estimated uncertainty, such as the expected calibration error, are based on the entire testing dataset. This means that, similar to classification tasks on unbalanced datasets, the uncertainty associated with single samples or small groups of samples may be biased towards the performance on the rest of the dataset. However, for practical applications, assessing the reliability of predicted confidence on a pointwise basis could offer more possibilities than aggregated reliability based on some testing data, independent of the current situation. This pointwise evaluation is particularly important for mission- and safety-critical applications. Current methods are empirically evaluated, with their performance underlined by reasonable and explainable values of uncertainty. Ground truth uncertainty, however, is generally unavailable for validation. Even though existing methods are calibrated on given datasets, these results cannot simply be transferred to other datasets, as shifts in data distribution must be considered, and many fields can only cover a small portion of the actual data environment. Preparing a large amount of training data is challenging and expensive, so synthetic data is often used to train models. For this, artificial uncertainties in labels and data should be considered to gain a better understanding of uncertainty quantification performance. The gap between real and synthetic data, or estimated and real uncertainty, further limits the adoption of current methods for uncertainty quantification.

Existing methods for neural network uncertainty quantification provide predictions of certainty without any indication of what causes the uncertainties. Although these certainty values often appear reasonable to a human observer, it is unclear whether the uncertainties are actually based on the same observations the human observer made. Without understanding the reasons and motivations behind single uncertainty estimations, it is much harder to guarantee performance when transferring from one dataset to another or when dealing with domain shifts. In safety-critical real-life applications, the lack of explainability makes it significantly harder to apply available methods. Besides the explainability of neural network decisions, existing methods for uncertainty quantification are not well understood at a higher level.

For instance, explaining the behavior of single deterministic approaches, ensembles, or Bayesian methods is a current research direction that remains difficult to fully grasp[47]. However, understanding how these methods operate and capture uncertainty is crucial for identifying pathways for refinement and detecting and characterizing uncertainty failures and shortcomings.

References

[1] J. Yao, B. Yuan, Research on the Application and Optimization Strategies of Deep Learning in Large Language Models, Journal of Theory and Practice of Engineering Science 4(05) (2024) 88-94.

[2] J. Guo, L. Du, H. Liu, M. Zhou, X. He, S. Han, Gpt4graph: Can large language models understand graph structured data? an empirical evaluation and benchmarking, arXiv preprint arXiv:2305.15066 (2023).

[3] S. Cao, J. Xiao, On Efficient and Flexible Autonomous Robotic Insertion Assembly in the Presence of Uncertainty, IEEE Robotics and Automation Letters (2024).

[4] J. Huo, Y. Wang, N. Wang, W. Gao, J. Zhou, Y. Cao, Data-driven design and optimization of ultra-tunable acoustic metamaterials, Smart Materials and Structures 32(5) (2023) 05LT01.

[5] H. Guo, A.B. Tikhomirov, A. Mitchell, I.P.J. Alwayn, H. Zeng, K.C. Hewitt, Real-time assessment of liver fat content using a filter-based Raman system operating under ambient light through lock-in amplification, Biomedical Optics Express 13(10) (2022) 5231-5245.

[6] H. Guo, A.E. Stueck, J.B. Doppenberg, Y.S. Chae, A.B. Tikhomirov, H. Zeng, M.A. Engelse, B.L. Gala-Lopez, A. Mahadevan-Jansen, I.P. Alwayn, Evaluation of minimum-to-severe global and macrovesicular steatosis in human liver specimens: a portable ambient light-compatible spectroscopic probe, medRxiv (2023) 2023.12. 04.23299259.

[7] D. Petroff, V. Blank, P.N. Newsome, C.S. Voican, M. Thiele, V. de Lédinghen, S. Baumeler, W.K. Chan, G. Perlemuter, A.-C. Cardoso, Assessment of hepatic steatosis by controlled attenuation parameter using the M and XL probes: an individual patient data meta-analysis, The Lancet Gastroenterology & Hepatology 6(3) (2021) 185-198.

[8] H. Guo, A.E. Stueck, J.B. Doppenberg, Y.S. Chae, A.B. Tikhomirov, H. Zeng, B.L. Gala-Lopez, A. Mahadevan-Jansen, M.A. Engelse, I.P. Alwayn, Assessment of liver steatosis using an ambient light-compatible Raman system: enhancing specificity with supplementary reflectance information, Biomedical Vibrational Spectroscopy 2024: Advances in Research and Industry, SPIE, 2024, p. PC128390B.

[9] H. Guo, A.E. Stueck, A.B. Tikhomirov, H. Zeng, I.P. Alwayn, B.L. Gala-Lopez, A. Mahadevan-Jansen, A.K. Locke, K.C. Hewitt, Evaluation of Steatosis in Human Liver Specimens Using an Ambient Light-compatible Raman Spectroscopy Approach, Bio-Optics: Design and Application, Optica Publishing Group, 2023, p. JTu4B. 26.

[10] H. Guo, V.S. Zions, B.A. Law, K. Hewitt, Potential of Raman-reflectance combination to quantify liver steatosis and fat droplet size: evidence from Monte Carlo simulations and phantom

studies, Authorea Preprints (2024).
[11] H. Guo, B.L. Gala-Lopez, I.P. Alwayn, K.C. Hewitt, Liver discard rate due to conservative estimations of steatosis: an inference-based approach, medRxiv (2023) 2023.12. 04.23299406.
[12] K. Huang, X. Chen, X. Di, Q. Du, Dynamic driving and routing games for autonomous vehicles on networks: A mean field game approach, Transportation Research Part C: Emerging Technologies 128 (2021) 103189.
[13] A.G. Wilson, P. Izmailov, Bayesian deep learning and a probabilistic perspective of generalization, Advances in neural information processing systems 33 (2020) 4697-4708.
[14] A.G. Roy, S. Conjeti, N. Navab, C. Wachinger, A.s.D.N. Initiative, Bayesian QuickNAT: Model uncertainty in deep whole-brain segmentation for structure-wise quality control, NeuroImage 195 (2019) 11-22.
[15] M. Yin, Millimeter Wave Wireless Assisted Indoor Robot Navigation, New York University Tandon School of Engineering, 2024.
[16] K. Pfeiffer, Y. Jia, M. Yin, A.K. Veldanda, Y. Hu, A. Trivedi, J. Zhang, S. Garg, E. Erkip, S. Rangan, Path planning under uncertainty to localize mmWave sources, 2023 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2023, pp. 3461-3467.
[17] M. Yin, A.K. Veldanda, A. Trivedi, J. Zhang, K. Pfeiffer, Y. Hu, S. Garg, E. Erkip, L. Righetti, S. Rangan, Millimeter wave wireless assisted robot navigation with link state classification, IEEE Open Journal of the Communications Society 3 (2022) 493-507.
[18] Y. Hu, M. Yin, W. Xia, S. Rangan, M. Mezzavilla, Multi-frequency channel modeling for millimeter wave and thz wireless communication via generative adversarial networks, 2022 56th Asilomar Conference on Signals, Systems, and Computers, IEEE, 2022, pp. 670-676.
[19] H. Abedi, S. Luo, V. Mazumdar, M.M. Riad, G. Shaker, AI-powered in-vehicle passenger monitoring using low-cost mm-wave radar, IEEE Access 10 (2021) 18998-19012.
[20] X. Chen, X. Di, Ridesharing user equilibrium with nodal matching cost and its implications for congestion tolling and platform pricing, Transportation Research Part C: Emerging Technologies 129 (2021) 103233.
[21] X. Chen, X. Di, Z. Li, Social Learning for Sequential Driving Dilemmas, Games 14(3) (2023) 41.
[22] X. Chen, Z. Li, X. Di, Social learning in Markov games: Empowering autonomous driving, 2022 IEEE Intelligent Vehicles Symposium (IV), IEEE, 2022, pp. 478-483.
[23] Z. Hu, Y. Sun, J. Wang, Y. Yang, DAC-DETR: Divide the attention layers and conquer, Advances in Neural Information Processing Systems 36 (2024).
[24] Z. Hu, Y. Sun, Y. Yang, Switch to generalize: Domain-switch learning for cross-domain few-shot classification, International Conference on Learning Representations, 2022.
[25] Z. Hu, Y. Sun, Y. Yang, Suppressing the heterogeneity: A strong feature extractor for few-shot segmentation, The Eleventh International Conference on Learning Representations, 2023.
[26] X. Chen, X. Liu, W. Liu, X.-P. Zhang, Y. Zhang, T. Mei, Explainable person re-identification with attribute-guided metric distillation, Proceedings of the IEEE/CVF international conference on computer vision, 2021, pp. 11813-11822.
[27] Z. Hu, Y. Sun, Y. Yang, J. Zhou, Divide-and-regroup clustering for domain adaptive person re-identification, Proceedings of the AAAI Conference on Artificial Intelligence, 2022, pp. 980-988.
[28] X. Chen, W. Liu, X. Liu, Y. Zhang, T. Mei, A cross-modality and progressive person search system, Proceedings of the 28th ACM International Conference on Multimedia, 2020, pp. 4550-4552.
[29] E. Ning, C. Wang, H. Zhang, X. Ning, P. Tiwari, Occluded person re-identification with deep learning: a survey and perspectives, Expert Systems with Applications (2023) 122419.
[30] X. Chen, W. Liu, X. Liu, Y. Zhang, J. Han, T. Mei, MAPLE: Masked pseudo-labeling autoencoder for semi-supervised point cloud action recognition, Proceedings of the 30th ACM International Conference on Multimedia, 2022, pp. 708-718.
[31] X. Chen, X. Liu, K. Liu, W. Liu, T. Mei, A baseline framework for part-level action parsing and action recognition, arXiv preprint arXiv:2110.03368 (2021).
[32] M.A. Khan, K. Javed, S.A. Khan, T. Saba, U. Habib, J.A. Khan, A.A. Abbasi, Human action recognition using fusion of multiview and deep features: an application to video surveillance, Multimedia tools and applications 83(5) (2024) 14885-14911.
[33] X. Chen, X. Liu, W. Liu, K. Liu, D. Wu, Y. Zhang, T. Mei, Part-level Action Parsing via a Pose-guided Coarse-to-Fine Framework, 2022 IEEE International Symposium on Circuits and Systems (ISCAS), IEEE, 2022, pp. 419-423.
[34] X. Chen, W. Liu, Q. Bao, X. Liu, Q. Yang, R. Dai, T. Mei, Motion Capture from Inertial and Vision Sensors, arXiv preprint arXiv:2407.16341 (2024).
[35] X. Chen, H. Zeng, H. Xu, X. Di, Sentiment analysis of autonomous vehicles after extreme events using social media data, 2021 IEEE International Intelligent Transportation Systems Conference (ITSC), IEEE, 2021, pp. 1211-1216.
[36] X. Chen, Z. Wang, X. Di, Sentiment analysis on multimodal transportation during the covid-19 using social media data, Information 14(2) (2023) 113.
[37] B. Wang, H. Duan, Y. Feng, X. Chen, Y. Fu, Z. Mo, X. Di, Can LLMs Understand Social Norms in Autonomous Driving Games?, arXiv preprint arXiv:2408.12680 (2024).
[38] M. Qu, X. Chen, W. Liu, A. Li, Y. Zhao, ChatVTG: Video Temporal Grounding via Chat with Video Dialogue Large Language Models, Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 1847-1856.
[39] M. Yin, T. Li, H. Lei, Y. Hu, S. Rangan, Q. Zhu, Zero-Shot Wireless Indoor Navigation through Physics-Informed Reinforcement Learning, 2024 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2024, pp. 5111-5118.
[40] J. Gawlikowski, C.R.N. Tassi, M. Ali, J. Lee, M. Humt, J. Feng, A. Kruspe, R. Triebel, P. Jung, R. Roscher, A survey of uncertainty in deep neural networks, Artificial Intelligence Review 56(Suppl 1) (2023) 1513-1589.
[41] L.V. Jospin, H. Laga, F. Boussaid, W. Buntine, M. Bennamoun, Hands-on Bayesian neural networks—A tutorial for deep learning users, IEEE Computational Intelligence Magazine 17(2) (2022) 29-48.
[42] P.W. Battaglia, J.B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, Relational inductive biases, deep learning, and graph networks, arXiv preprint arXiv:1806.01261 (2018).
[43] A. Kendall, Y. Gal, What uncertainties do we need in bayesian deep learning for computer vision?, Advances in neural information processing systems 30 (2017).
[44] J. Lee, G. AlRegib, Gradients as a measure of uncertainty in neural networks, 2020 IEEE International Conference on Image Processing (ICIP), IEEE, 2020, pp. 2416-2420.
[45] L. Oala, C. Heiß, J. Macdonald, M. März, W. Samek, G. Kutyniok, Interval neural networks: Uncertainty scores, arXiv preprint arXiv:2003.11566 (2020).
[46] J. Zhang, B. Kailkhura, T.Y.-J. Han, Mix-n-match: Ensemble and compositional methods for uncertainty calibration in deep learning, International conference on machine learning, PMLR, 2020, pp. 11117-11128.
[47] S. Fort, H. Hu, B. Lakshminarayanan, Deep ensembles: A loss landscape perspective, arXiv preprint arXiv:1912.02757 (2019).