Bundesministerium für Bildung und Forschung | Deutsches Zentrum für Luft- und Raumfahrt (DLR) German Aerospace Center

DeepDoubt

Weiterentwicklung von Unsicherheitsmaßen zur Erhöhung der Erklärbarkeit und Transparenz des Maschinellen Lernens und der Künstlichen Intelligenz

KNIME | IOS Institut für Optische Systeme

The project DeepDoubt aims for the development and practical application of uncertainty measures in deep learning. Up to now the project is about midway the focus has been in the development of a fast query function for active learning based on predictive uncertainty. The development of flexible variational distribution for variational inference based on special normalizing flows. Modeling uncertainty also plays a key role in optical inspection tasks where only error free data is available for training and hence one-class classification techniques need to be applied. Here, we demonstrate how uncertainty measures can be used to improve detection capabilities. In the future, we will also try to incorporate uncertainty for multisensor tasks like the fusion of LiDAR and camera data based on the uncertainties of the respective modalities.

Overview

Timespan:
01.04.2020 until 31.03.2023

Project partners:
IOS Konstanz, Hochschule Konstanz, KNIME

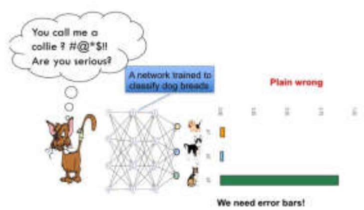

Schematic illustration showcasing the problem with standard neural network that are often overconfident and miscalibrated.

Methodological development

Uncertainty quantification
- Bayesian Modeling complex posteriors for Variational Inference
- Deep ensembles and boostrapping methods
- Last Layer Multivariate Gaussian

Focus applications

Active Learning
- Uncertainty as a proposal function speed up using

Optical Inspection
- Modeling non-defective reference data p(x)
- Detecting outliers p(x)
- Failure detection with outlier score and uncertainty (Figure)

Object recognition
- Detect objects with given uncertainty
- Application fusion from lidar and camera

Improving active learning by using uncertainty as proposal function

In active learning only partial labelled data is available and the model fit is done incrementally by querying an oracle, potentially human.
Several types of querying strategies exist. Most popular are:
Entropy $\tilde{x} = \operatorname{argmax}_{x \in X} \sum_{c=0}^{\text{Classes}} p(y=c \mid x) \log p(y=c \mid x)$
or
Minimum probability $\tilde{x} = \operatorname{argmin}_{x \in X} \operatorname{argmax}_{c \in C} p(y=c \mid x)$.

Quantifying uncertainty
Currently we analyze different standard uncertainty estimation techniques, such as MC Dropout, Deep Ensembles, Swag and VI, and their interaction with different query strategies. For accelerating the MC Dropout method in real-time applications, we applied moment propagation technique without loss of accuracy.

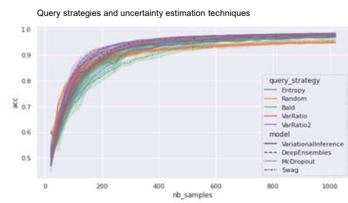

Visualizing the accuracy on MNIST dataset for different query strategies and uncertainty estimation techniques. Across all query strategies we see the following order Ensemble better than MC-dropout uncertainty better than Swag.

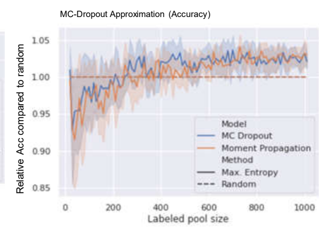

No loss in accuracy when approximating MC-Dropout with Moment Propagation.

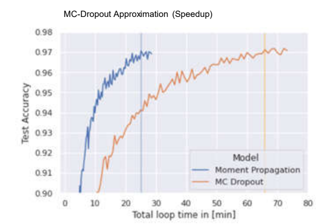

Runtime comparison between moment propagation and MC Dropout (25 forward passes). Moment propagation gives a significant speed up.

Literature / Related Work

Using uncertainty as a query function: Houlsby, N., Huszár, F., Ghahramani, Z., & Lengyel, M. (2011). Bayesian active learning for classification and preference learning. arXiv preprint arXiv:1112.5745
Using MC-Dropout as uncertainty estimation: Yarin Gal, Riashat Islam, Zoubin Ghahramani https://arxiv.org/abs/1703.02910
MC-Dropout approximation with Moment Propagation: Kai Brach, Beate Sick, Oliver Dürr https://arxiv.org/abs/2007.03293
Manuscript in preparation

Flexible Posterior for Variational Inference

Bayesian neural networks are a natural way to incorporate epistemic uncertainty into deep learning. However, the gold standard MCMC is hardly applicable. Variational inference (VI) is a technique to approximate difficult to compute posteriors $p(\theta|D)$ by a variational approximation $q_\lambda(\theta)$ paramerized by λ. This project uses a transformation model using Bernstein polynomials $Be_i(z)$ to construct a complex posterior from a simple distribution. This method fits in the DL framework, optimizing a parameter λ, and allows so to combine DL with Bayesian Modelling.

Main Idea of the transformation model for $f = f_3 \circ f_2 \circ \sigma \circ f_1$

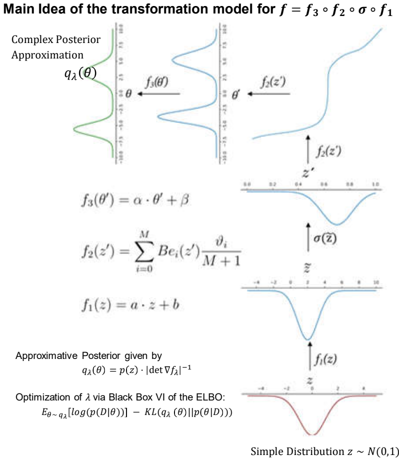

Results: 1-D Example

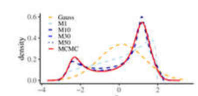

Model (Stan)

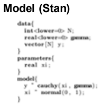

Bayesian Linear Model with 4 Parameters (MCMC in Blue)

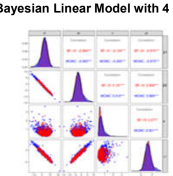

Multivariate Generalization can be done by letting the transformation model for θ_j depend on $\theta_{<j}$ and z_j using a NN (Jacobian ∇f_λ is diagonal). Effective implementation with Masked Autoregressive Flows.

Model (Stan)

Combination CNN of and simple Bayesian models with interpretable parameters

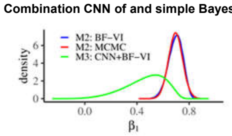

SIIM-ISIC Melanoma Class. Challenge):
Features: images of skin + age information.
Outcome: benign / malignant

Models:
M2 Bayesian Logistic Regression on age β_1
M3 M2 + Images modeled

Literature / Related Work

- First used in statistics: Hothorn, T., Moest, L., and Buehlmann, P. (2018). Most likely transformations. Scandinavian Journal of Statistics, 45(1):110–134 / arxiv.org/abs/1508.06749
- Normalizing Flow are similar concepts in DL. Using Bernstein and NN for conditional outcome distribution: Sick, B., Hothorn, T., and Dürr, O. (2021). Deep transformation models ICPR 2020
- 1-D version and mean-field approximation: Sefan Hörtling, Daniel Dold, Oliver Dürr, Beate Sick https://arxiv.org/abs/2106.00528
- Generalization to more than 1-D manuscript in preparation

Improving one-class optical inspection techniques by incorporating uncertainty

A major problem in optical inspection is the lack of defective examples for training a supervised classificator. One-class classification techniques train a classificator with access only to non-defective examples. The output is a probability estimate of beeing defective $p(\hat{x} = defect \mid \theta, \phi)$.

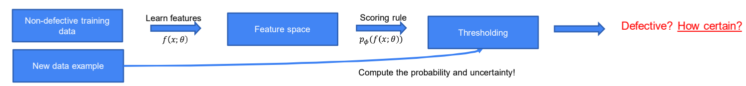

Model
Our current model uses a pooled CNN feature space based on VGG19: $f(x;\theta) = VGG19(x)$.
The scoring rule is a standard multivariate gaussian $p_\phi(\cdot) = MVG(x \mid \Sigma, \mu)$.

Quantifying uncertainty
Estimating the uncertainty of predictions based on posterior of model parameters $p(\Theta \mid \text{Data})$, with $\Theta = \{\theta, \phi\}$.
Posterior is approximated by bootstrapping, which is Bayesian inference approximation.
Using the estimated credibility interval can improve model calibration and detection performance.
Incorporation of flexible posteriors based on variational inference is planned,

	SVM	KDE	VAE	LSA	DSVD	μshift
0	0.988	0.885	**0.998**	0.993	0.980	0.997
1	0.999	0.996	**0.999**	0.999	0.997	0.993
2	0.902	0.710	0.962	0.959	0.917	**0.986**
3	0.950	0.693	0.947	0.966	0.919	**0.979**
4	0.955	0.844	0.965	0.956	0.949	**0.971**
5	0.968	0.776	0.963	0.964	0.885	**0.981**
6	0.978	0.861	0.995	0.994	0.983	**0.995**
7	0.965	0.884	0.974	**0.980**	0.946	0.973
8	0.853	0.669	0.905	0.953	0.939	**0.969**
9	0.955	0.825	0.978	**0.981**	0.965	0.977
	0.951	0.814	0.969	0.975	0.948	**0.982**

AUC values on the MNIST dataset using one-vs-all protocol.

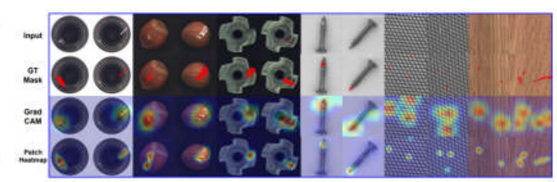

Optical inspection examples based on the MVTec dataset showing defect detection and localization.

Literature / Related Work

Li, Chun-Liang et. Al (2021). CutPaste: Self-Supervised Learning for Anomaly Detection and Localization. https://arxiv.org/pdf/2104.04015v1.pdf
Bergmann, P (2019). MVTec AD – A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 9592–9600
Rippel, O., Mertens, P., & Merhof, D. (2021, January). Modeling the distribution of normal data in pre-trained deep features for anomaly detection. In 2020 25th International Conference on Pattern Recognition (ICPR) (pp. 6726-6733). IEEE.