

# Knowledge from Uncertainty in Evidential Deep Learning

Cai Davies  
Cardiff University, UK

Marc Roig Vilamala  
Cardiff University, UK

Alun D. Preece  
Cardiff University, UK

Federico Cerutti  
University of Brescia, Italy

Lance M. Kaplan  
DEVCOM Army Research Lab, USA

Supriyo Chakraborty  
IBM, USA

## ABSTRACT

This work reveals an *evidential signal* that emerges from the uncertainty value in Evidential Deep Learning (EDL). EDL is one example of a class of uncertainty-aware deep learning approaches designed to provide confidence (or epistemic uncertainty) about the current test sample. In particular for computer vision and bidirectional encoder large language models, the ‘evidential signal’ arising from the Dirichlet strength in EDL can, in some cases, discriminate between classes, which is particularly strong when using large language models. We hypothesise that the KL regularisation term causes EDL to couple aleatoric and epistemic uncertainty. In this paper, we empirically investigate the correlations between misclassification and evaluated uncertainty, and show that EDL’s ‘evidential signal’ is due to misclassification bias. We critically evaluate EDL with other Dirichlet-based approaches, namely Generative Evidential Neural Networks (EDL-GEN) and Prior Networks, and show theoretically and empirically the differences between these loss functions. We conclude that EDL’s coupling of uncertainty arises from these differences due to the use (or lack) of out-of-distribution samples during training.

## CCS CONCEPTS

• **Computing methodologies** → **Reasoning about belief and knowledge**; *Natural language processing*; *Computer vision*.

## KEYWORDS

evidential deep learning, epistemic uncertainty, uncertainty estimation

### ACM Reference Format:

Cai Davies, Marc Roig Vilamala, Alun D. Preece, Federico Cerutti, Lance M. Kaplan, and Supriyo Chakraborty. 2023. Knowledge from Uncertainty in Evidential Deep Learning. In *Proceedings of 2nd KDD Workshop on Uncertainty Reasoning and Quantification in Decision Making (KDD ’23)*. ACM, New York, NY, USA, 5 pages.

## 1 INTRODUCTION

This work investigates an *evidential signal* that seems to emerge from the evidential deep learning (EDL) approach. EDL is part of a class techniques to realize uncertainty-aware deep models that are able to faithfully represent decision confidence.

When an artificial intelligence (AI) system assists a human operator with predictions, the human has to develop insights (*i.e.* a mental model) of when to trust the AI system with its recommendations [2]. If the human mistakenly relies upon the AI system in regions where it is likely to err, catastrophic failures may occur

[1, 11]. To identify such regions where the AI system is likely to err, we need to distinguish between (at least) two different sources of uncertainty: *aleatoric* (or *aleatory*) and *epistemic* uncertainty [7, 8]. Aleatoric uncertainty refers to the variability in the outcome of an experiment due to inherently random effects (e.g. flipping a fair coin). Epistemic uncertainty refers to the epistemic state of the agent using the model; hence its lack of knowledge that – in principle – can be reduced based on additional data samples.

Estimating model (epistemic) uncertainty is crucial for trust [24]. Approaches to this problem include Ensemble Distribution Distillation [17], Posterior Networks [4], Prior Networks [16], Evidential Deep Learning (EDL) [21], and Evidential Deep Learning with GANs (EDL-GEN) [22], all of which aim at separating aleatoric and epistemic uncertainty. We focus our attention on the last three mentioned approaches. EDL (Section 2.1), in particular, transforms a classification problem into a regression one, with the goal of computing pseudo-counts of evidence in favour of different classes, which then can be mapped into a Dirichlet distribution. Given a classification problem, Prior Networks (Section 2.2), also, returns a Dirichlet distribution, which is computed considering an additional dataset of out-of-distribution (OOD) samples. Finally, EDL-GEN (Section 2.3) expands the idea of Prior Networks by automatically synthesising out-of-distribution samples.

The goal of this paper is to investigate a peculiarity in the Dirichlet strength output by a model trained using EDL which we call an ‘evidential signal’, which we observe to be strongly coupled with the ground truth class, and in some cases carries sufficient information to discriminate between classes on this signal alone. This is seen to be particularly strong when using large language models for text classification, where the discriminatory power of the ‘evidential signal’ matches the accuracy of the model’s classifications. We hypothesise that this is caused by misclassification biases, which are intrinsic in the datasets, indicating that aleatoric and epistemic uncertainty are unintentionally closely coupled. We theoretically and experimentally explore the differences between the EDL, EDL-GEN and Prior Networks approaches to quantifying epistemic uncertainty, and show that these latter two single pass methods do not exhibit this same signal, indicating a better biasing of the Dirichlet strength to represent epistemic uncertainty. We conclude this is due to these methods being heavily dependent on using out-of-distribution samples during training, where EDL is dependent on a KL regularisation term that is high when a misclassification occurs.

**The main contributions of this paper are therefore three-fold:** (1) a critical account of the differences between EDL’s, EDL-GEN’s, and Prior Networks’ loss functions which highlights how the Dirichlet strength of the distribution outputted by EDL is heavily

dependent on misclassification bias; (2) an empirical evaluation confirming such a hypothesis and revealing how the Dirichlet strength is coupled with the class, in some cases carrying enough information for class discrimination matching the accuracy of the model; and (3) a confirmation that such a bias is missing or very weak in Prior Networks and EDL-GEN training with out-of-distribution samples.

## 2 QUANTIFYING ALEATORIC AND EPISTEMIC UNCERTAINTIES: NEURAL APPROACHES

We focus on three approaches to equip neural networks with the ability to express aleatoric and epistemic uncertainty in classification problems. **Evidential Deep Learning** (EDL) [21] we summarise in Section 2.1. **Prior Networks** [16], an approach that makes use of an explicit out-of-distribution dataset, we summarise in Section 2.2. **Evidential Deep Learning with GANs** [22], which builds on the previous two and avoids the need for a pre-chosen out-of-distribution data set while relying on a generative network to synthesise noisy data samples, we summarise in Section 2.3.

### 2.1 Evidential Deep Learning

Evidential Deep Learning (EDL) [21] replaces the classical use of a softmax as the final layer for classification with an estimation of the parameters of a Dirichlet distribution  $D(\boldsymbol{\pi}|\boldsymbol{\alpha})$  over the possible classes.<sup>1</sup> During training, the model can discover patterns in the data and generate evidence for specific class labels such that the overall loss is minimised. However, these features may also be present in counter-examples giving rise to misleading evidence. Reducing the magnitude of generated evidence may increase the overall loss, despite reducing the loss contributed by these counter-examples. To combat this, a regularisation term is included, which incorporates a Kullback-Leibler (KL) divergence term between a uniform Dirichlet distribution and  $\tilde{\alpha}$ , where  $\tilde{\alpha}$  is the parameters of the output Dirichlet distribution  $\alpha$  after removing the non-misleading evidence from  $f(x_i|\Theta)$ , such that a correctly classified sample with no evidence for other classes will generate  $\tilde{\alpha}$  as a uniform Dirichlet distribution.

### 2.2 Prior Networks

Prior Networks [16]<sup>2</sup> differs from EDL in that it uses an explicit out-of-distribution dataset alongside the in-distribution dataset the model is trained on. A Prior Network [16] for classification parameterises a distribution over a simplex, typically a Dirichlet distribution due to its tractable analytic properties. A Prior Network that parameterises a Dirichlet will be referred to as a *Dirichlet Prior Network* (DPN). Similarly to an EDL, a DPN will generate the concentration parameters  $\boldsymbol{\alpha}$  of the Dirichlet distribution.

DPNs for multi-class classification are trained to minimise the KL divergence between the model and a sharp Dirichlet distribution focused on the appropriate class for in-distribution data, and between the model and a flat Dirichlet distribution for out-of-distribution data, chosen in accordance with the *principle of insufficient reason*.

<sup>1</sup><https://github.com/dougbrion/pytorch-classification-uncertainty> (on May 26th 2023, available under MIT license)

<sup>2</sup><https://github.com/KaosEngineer/PriorNetworks> (on May 26th 2023).

Therefore, it is necessary to define the in-distribution targets  $\hat{\boldsymbol{\alpha}}$  and out-of-distribution targets  $\tilde{\boldsymbol{\alpha}}$  with a flat (uniform) Dirichlet distribution. As the true out-of-domain distribution is unknown, and samples are unavailable, a standard solution is to use a different, real dataset as a set of samples from the out-of-domain distribution.

### 2.3 EDL-GEN

To avoid the need to identify a dataset from an out-of-domain distribution, EDL-GEN<sup>3</sup> revisits the idea of Prior Networks by automatically generating out-of-distribution samples adding noise to a latent representation of the in-distribution samples derived using a variational autoencoder. EDL-GEN approximates the log density ratio  $\log(P_k(\mathbf{x})/P_{out}(\mathbf{x}))$  as the logit output of a binary classifier, which is trained to discriminate between the samples from  $P_k$  and  $P_{out}$ . To train such a network, in [22], the authors use the Bernoulli (logarithmic) loss. Similarly to the other two proposals, EDL-GEN also considers a regularisation term that helps achieve a classifier to be totally uncertain in its misclassification, except near decision boundaries. Finally, one of the main contributions of EDL-GEN concerns the generation of out-of-distribution samples by using the latent space of a Variational Auto-Encoder [10] as a proxy for the semantic similarity between samples in input space.

## 3 EXPERIMENTATION METHODOLOGY

### 3.1 Hypotheses

**EDL is heavily dependent on misclassification bias (H1)** as the KL regularisation reduces all evidence near decision boundaries where training data classes overlap. This conflates epistemic (Dirichlet strength) and aleatoric uncertainty. For this, we explore the relationship between recall, which is a measure of misclassification, and the calculated uncertainty outputted by the model. We hypothesise a high correlation for EDL and weaker or no correlation for EDL-GEN and Prior Networks, which would indicate a coupling of aleatoric and epistemic uncertainty in the EDL loss function.

**The Dirichlet strength computed by EDL is coupled with class (H2).** Given the coupling of the Dirichlet strength of EDL and misclassifications, and that aleatoric uncertainty is inherent in the dataset, we hypothesise that the Dirichlet strength will be closely coupled with the ground truth class, to the extent that in some cases where the misclassification bias is particularly strong, we can discriminate between classes based on the Dirichlet strength alone. To quantify this, we train a set of simple models on the Dirichlet strength and record the accuracy across a test set to measure the separability of the classes, which would indicate a coupling to the Dirichlet strength and differences in misclassifications between classes. Given our hypothesis of weak or no correlation with EDL-GEN and Prior Networks, we assume that this separability will be much weaker or non-existent using these approaches.

Finally, **the evidential signal does not appear in approaches that train with out-of-distribution samples (H3)**, namely EDL-GEN and Prior Networks, which follows from the observation that the above two hypotheses do not apply to these methods, or that the observation is much weaker. These methods learn to reduce

<sup>3</sup><https://muratsensoy.github.io/uncertainty.html> (on May 26th 2023)

**Table 1: Datasets used**

DATASET	CLASSES	LABEL
AMAZON [18]	2	SENTIMENT
BLOG [20]	2	GENDER
IMDB [15]	2	SENTIMENT
MNIST [13]	10	DIGIT
NEWSGROUPS [12]	20	TOPIC

the Dirichlet strength (or increase the epistemic uncertainty) over the out-of-distribution data.

### 3.2 Implementation Details

Models are implemented in PyTorch [19] using the PyTorch Lightning<sup>4</sup> framework. For the natural language tasks, the BERT [6] architecture is used with bert-base-uncased pretrained weights, where the model is then fine-tuned to the dataset. For the computer vision tasks, an implementation of LeNet [13] and VGG16 [23] are used, where VGG16 was used as both fine-tuning using pretrained weights and full training with randomly initialised weights.

AdamW [14] was used as an optimiser. It is similar to the Adam optimiser [9] used in [21] but shown to yield better training loss and improve generalisation. The learning rate usually performed well at  $1 \times 10^{-3}$ . The language models are fine-tuned for 30 epochs, while the full training of the computer vision models is 500 epochs.

### 3.3 Datasets

We use a selection of publicly available and widely used datasets in both the natural language and computer vision domains (Table 1). The diversity of datasets includes variations between natural language and computer vision, as well as number of classes and classification types, e.g., sentiment, topic, gender etc.


## 4 EXPERIMENTAL VALIDATION OF HYPOTHESES

### 4.1 EDL is heavily dependent on misclassification bias

To investigate empirically hypothesis 1 (H1), we look to correlations between recall and the Dirichlet strength, or ‘evidential signal’. In Figure 1 we show the results of 90 LeNet models trained on the MNIST dataset, as defined in Section 3.2, and using a random seed for training initialisations; on the x-axis we plot the recall, and y-axis we plot the uncertainty (or Dirichlet strength), and colour-code by ground truth class. There is a strong correlation between the recall and Dirichlet strength independent of the training seed, but also we notice a separation in the classes, in that digit 0 tends to have high recall-low Dirichlet strength and in opposite to digit 9, which indicates evidence towards H2.

To quantify this correlation, we show the Pearson’s and Spearman’s rank correlation coefficients in Table 2. H1 hypothesises that this correlation will not be seen in EDL-GEN and Prior Networks, so we also show the results for these by running the experiment

<sup>4</sup><https://github.com/PyTorchLightning/pytorch-lightning> (on May 10th 2023, Apache-2.0 license)



**Figure 1: Top: Recall against evidential signal plot across models trained on varying annealing\_step values to highlight the correlation. Bottom: Pearson’s and Spearman’s correlation between recall and uncertainty.**

**Table 2: Correlation between recall and evidential signal**

Approach	Pearson’s	Spearman’s
EDL	-0.9975	-0.8517
EDL-GEN	-0.4899	-0.4321
Prior Networks	-0.4030	-0.4320

in the same manner for these two approaches. These results show EDL having a much higher correlation, close to a perfectly negative linear relationship for the Pearson’s coefficient, where EDL-GEN and Prior Networks exhibit much weaker correlations.

### 4.2 The Dirichlet Strength computed by EDL is coupled with class

As observed in Figure 1, across multiple runs, classes in MNIST appear to show a coupling with the Dirichlet strength and recall. To investigate this further, we first look to individual models, by creating a cumulative distribution function from the computed Dirichlet strength, or ‘evidential signal’, of the samples in the test set and split by ground truth class. We show in Figure 2 the results of this for MNIST and the IMDB dataset, which is trained as described in Section 3.2, to show that this coupling is not exclusive to MNIST or computer vision tasks.

With IMDB, we see a very clear separation in the classes from the Dirichlet strength, where each appears to have ‘collapsed’ into a single value; the crossover seen is due to samples being misclassified, which we can remove by separating by predicted class rather than ground truth. For MNIST, we see separation but of a much lower magnitude, with the Dirichlet strength being more distributed.

To quantify the class discrimination, and so the magnitude of the coupling between Dirichlet strength and class, we use three simple machine learning methods – SVM, Decision Tree (DT) classifier [3] and XGBoost [5] – to train on the ‘evidential signal’ alone; we do this for each dataset from Table 1 in accordance with Section 3.2 for 50 runs. From the test set of each dataset, we compute the Dirichlet




Figure 2: Cumulative evidential signal of samples split by the ground-truth label for IMDB (top) and MNIST (bottom).

Table 3: Separability of evidential signal across simple models by test accuracy

Dataset	SVM	DT	XGBoost	Orig
IMDB	<b>0.9426</b>	0.9396	0.9382	0.9400
Blog	<b>0.7746</b>	0.7684	0.7670	0.7702
Amazon	<b>0.9274</b>	0.9230	0.9264	0.9262
Newsgroups	0.2734	0.7326	<b>0.7624</b>	0.8007
MNIST	0.1424	0.2657	<b>0.3789</b>	0.9813

strength and split this into a sub-train and sub-test set with an 80:20 ratio, which is used to train and test each of the simple ML methods keeping the original labels. A high accuracy would indicate strong separability; we show these results in Table 3.

As in Section 4.1, the natural language tasks exhibit the clearest separability, with most of the datasets (with the exception of the multi-class Newsgroups) showing separability in the Dirichlet strength on par with the model’s accuracy. From Figure 2 **bottom** we see the Dirichlet strength being more distributed than the natural language tasks, which is reflected in the poor separability scores with a maximum of 0.38; however this is above the random-guessing threshold.




Figure 3: EDL-GEN implementation with MNIST and computed uncertainty plotted as a cumulative distribution function.




Figure 4: Prior Networks implementation with MNIST and computed uncertainty plotted as a cumulative distribution function.

### 4.3 Evidential signal does not appear in OOD approaches

In Section 4.1 we show the weaker correlation between the recall and Dirichlet strength for the out-of-distribution methods. To explore whether the ‘evidential signal’ observed in EDL is seen in these other approaches (H3), we run the same experiment as in Section 4.2 with implementations of EDL-GEN [22] and Prior Networks [16]. Given that Prior Networks uses an explicit out-of-distribution dataset (see Section 2.2), we choose pairs from our existing datasets for NLP classifications, for example, we take IMDB as the in-distribution and Amazon Books as the out-of-distribution. For MNIST, we use EMNIST handwritten letters as the out-of-distribution dataset. Figures 4 and 3 show these results for MNIST using EDL-GEN and Prior Networks, respectively. Due to the method for generating out-of-distribution samples seen in EDL-GEN, the application of this method to the NLP datasets is not trivial, and so these experiments were omitted.

Visually, the separability of classes, or class discrimination from the Dirichlet strength, appears non-existent. To quantify this, as

**Table 4: Best separability of evidential signal across Dirichlet approaches**

Approach	IMDB	MNIST
EDL	0.9426	0.3789
EDL-GEN	N/A	0.1465
Prior Networks	0.55	0.1515

before, we use the same set of simple machine learning methods trained on the Dirichlet strength and take the mean accuracy of the sub-test set across 50 runs to indicate separability, which is shown in Table 4. To simplify, we show the best mean accuracy of the three methods. As in Table 3, EDL shows very good separation for the natural language tasks (IMDB), and poor but above random-guessing for MNIST. For EDL-GEN and Prior Networks, we see scores that are close to random-guessing, confirming the visual observation that the separability is weak or non-existent, which is in-line with H3 that this ‘evidential signal’ does not appear in approaches that train using out-of-distribution samples.

## 5 CONCLUSION AND FUTURE WORK

In this paper, we have critically explored differences between EDL, EDL-GEN and Prior Networks both theoretically (Section 3.1) and empirically (Section 4). We explored the relationship between recall and the Dirichlet strength of EDL, showing a strong correlation not seen in the other two approaches, which was further shown by the discriminatory power of the evidential signal in EDL only. Since these observations are due to misclassification bias, we have shown that the output of EDL is coupled with aleatoric uncertainty, which is inherent in the dataset. As such, the computed uncertainty differs to these other approaches, given the loss functions of EDL-GEN and Prior Networks heavily depends on out-of-distribution samples to calculate the Dirichlet strength.

Future work will address the limitations of the present study. We will explore additional choices of loss function from [21] (preliminary analysis is consistent with our observation in Section 4.2). We will expand the investigation of the relationship between the evidential signal and misclassification bias to determine the cause of the NLP models exhibiting much stronger discrimination even in the multi-class case (i.e., Newsgroups). Finally, we will look for the presence of the evidential signal in other approaches for estimating epistemic uncertainty beyond EDL and Prior Networks.

## REFERENCES

- [1] Gagan Bansal, Besmira Nushi, Ece Kamar, Walter Lasecki, Dan Weld, and Eric Horvitz. 2019. Beyond Accuracy: The Role of Mental Models in Human-AI Team Performance. In *Proceedings of HCOMP2019*. AAAI, 2–11.
- [2] Gagan Bansal, Besmira Nushi, Ece Kamar, Daniel S Weld, Walter S Lasecki, and Eric Horvitz. 2019. Updates in Human-AI Teams: Understanding and Addressing the Performance/Compatibility Tradeoff. In *AAAI*. 2429–2437.
- [3] Christopher M Bishop. 2006. *Pattern Recognition and Machine Learning*. Springer-Verlag, Berlin, Heidelberg. Citation Key: Bishop2006.
- [4] Bertrand Charpentier, Daniel Zügner, and Stephan Günnemann. 2020. Posterior network: Uncertainty estimation without ood samples via density-based pseudo-counts. *Advances in Neural Information Processing Systems* 33 (2020), 1356–1367.
- [5] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System, booktitle = Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (*KDD '16*). ACM, New York, NY, USA, 785–794. <https://doi.org/10.1145/2939672.2939785>
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv:1810.04805 [cs.CL]
- [7] Stephen C Hora. 1996. Aleatory and Epistemic Uncertainty in Probability Elicitation with an Example from Hazardous Waste Management. *Reliability Engineering & System Safety* 54, 2 (1996), 217–223.
- [8] Eyke Hüllermeier and Willem Waegeman. 2021. Aleatoric and Epistemic Uncertainty in Machine Learning: An Introduction to Concepts and Methods. *Machine Learning* 110, 3 (March 2021), 457–506. <https://doi.org/10.1007/s10994-021-05946-3> arXiv:1910.09457
- [9] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. <https://doi.org/10.48550/ARXIV.1412.6980>
- [10] Diederik P Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *ICLR2014*. arXiv:1312.6114
- [11] Rafal Kocielnik, Saleema Amershi, and Paul N. Bennett. 2019. Will You Accept an Imperfect AI? Exploring Designs for Adjusting End-user Expectations of AI Systems. In *Conference on Human Factors in Computing Systems - Proceedings*. Association for Computing Machinery, New York, New York, USA, 1–14. <https://doi.org/10.1145/3290605.3300641>
- [12] Ken Lang. 1995. Newsweeder: Learning to filter netnews. In *Proceedings of the Twelfth International Conference on Machine Learning*, 331–339.
- [13] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [14] Ilya Loshchilov and Frank Hutter. 2017. Decoupled Weight Decay Regularization. <https://doi.org/10.48550/ARXIV.1711.05101>
- [15] Andrew Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*. 142–150.
- [16] Andrey Malinin and Mark Gales. 2018. Predictive Uncertainty Estimation via Prior Networks. In *Advances in Neural Information Processing Systems*, Vol. 31. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2018/file/3ea2db50e62ceefca70a9d9a56a6f4-Paper.pdf>
- [17] Andrey Malinin, Bruno Mlodozeniec, and Mark Gales. 2019. Ensemble distribution distillation. *arXiv preprint arXiv:1905.00076* (2019).
- [18] Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 188–197.
- [19] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. <https://doi.org/10.48550/ARXIV.1912.01703>
- [20] Jonathan Schler, Moshe Koppel, Shlomo Argamon, and James W Pennebaker. 2006. Effects of age and gender on blogging. In *AAAI spring symposium: Computational approaches to analyzing weblogs*, Vol. 6. 199–205.
- [21] Murat Sensoy, Lance Kaplan, and Melih Kandemir. 2018. Evidential Deep Learning to Quantify Classification Uncertainty. In *Advances in Neural Information Processing Systems*, Vol. 31. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2018/file/a981f2b708044d6f4a71a1463242520-Paper.pdf>
- [22] Murat Sensoy, Lance M. Kaplan, Federico Cerutti, and Maryam Saleki. 2020. Uncertainty-Aware Deep Classifiers Using Generative Models. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. 5620–5627.
- [23] Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. <https://doi.org/10.48550/ARXIV.1409.1556>
- [24] Richard Tomsett, Alun Preece, Dave Braines, Federico Cerutti, Supriyo Chakraborty, Mani Srivastava, Gavin Pearson, and Lance Kaplan. 2020. Rapid trust calibration through interpretable and uncertainty-aware AI. *Patterns* 1, 4 (2020), 100049.