



UNICA

UNIVERSITÀ
DEGLI STUDI
DI CAGLIARI



Università di Cagliari

UNICA IRIS Institutional Research Information System

This is the Author's *accepted* manuscript version of the following contribution:

Emanuele Ledda, Giorgio Fumera, Fabio Roli, ***Dropout Injection at Test Time for Post Hoc Uncertainty Quantification in Neural Networks***, in *Information Sciences*, vol. 645 (2023), article number 119356.

The publisher's version is available at:

<https://doi.org/10.1016/j.ins.2023.119356>

When citing, please refer to the published version.

© 2023. This manuscript version is made available under the CC-BY-NC-ND 4.0 license <https://creativecommons.org/licenses/by-nc-nd/4.0/>

This full text was downloaded from UNICA IRIS <https://iris.unica.it/>

Dropout Injection at Test Time for Post Hoc Uncertainty Quantification in Neural Networks

Emanuele Ledda^{a,b,*}, Giorgio Fumera^c, Fabio Roli^b

^a*Department of Computer, Control and Management Engineering, Sapienza University of Rome, Via Ariosto 25, Rome, 00185, Italy*

^b*Department of Informatics, Bioengineering, Robotics, and Systems Engineering, University of Genova, Via Dodecaneso 35, Genova, 16146, Italy*

^c*Department of Electric and Electronic Engineering, University of Cagliari, Via Marengo 3, Cagliari, 09100, Italy*

Abstract

Among Bayesian methods, Monte Carlo dropout provides principled tools for evaluating the *epistemic* uncertainty of neural networks. Its popularity recently led to seminal works that proposed activating the dropout layers only during inference for evaluating epistemic uncertainty. This approach, which we call dropout *injection*, provides clear benefits over its traditional counterpart (which we call *embedded* dropout) since it allows one to obtain a post hoc uncertainty measure for any existing network previously trained without dropout, avoiding an additional, time-consuming training process. Unfortunately, no previous work thoroughly analyzed injected dropout and compared it with embedded dropout; therefore, we provide a first comprehensive investigation, focusing on regression problems. We show that the effectiveness of dropout injection strongly relies on a suitable scaling of the corresponding uncertainty measure, and propose an alternative method to implement it. We also considered the trade-off between negative log-likelihood and calibration error as a function of the scale factor. Experimental results on benchmark data sets from several regression tasks, including crowd counting, support our claim that dropout injection can effectively behave as a competitive post hoc alternative to embedded dropout.

*Corresponding author.

Email addresses: `emanuele.ledda@uniroma1.it` (Emanuele Ledda), `fumera@unica.it` (Giorgio Fumera), `fabio.roli@unige.it` (Fabio Roli)

Keywords:

Crowd Counting, Epistemic Uncertainty, Monte Carlo Dropout,
Trustworthy AI, Uncertainty Quantification

1. Introduction

Systems based on Artificial Intelligence (AI) are nowadays widespread in industry, workplaces and everyday life. They often operate in critical contexts where errors can cause considerable harm to humans (e.g., self-driving cars), where transparency and accountability are required (e.g., medical and legal domains), and where end users need to be aware of the uncertainty or reliability of AI predictions (e.g., intelligent video surveillance). In particular, given that AI moved out of research laboratories to real-world scenarios, the need for trustworthiness guarantees led to the growth of an entire research field devoted to quantifying *uncertainty* of machine learning models [1, 2].

In this field, it is common to distinguish between two possible sources of uncertainty: one caused by intrinsic data noise (*aleatoric* uncertainty, or data uncertainty) and one caused by the lack of knowledge on the “correct” prediction model (*epistemic* uncertainty, or model uncertainty) [3]. In particular, most of the work on deep neural networks uses their Bayesian extension for quantifying *uncertainty*, modeling the network weights with a probability distribution [4, 5]. In this context, one of the leading and most widely used techniques for evaluating *epistemic* uncertainty of neural networks is the simple and effective Monte Carlo dropout [6]. It is based on *dropout*, a stochastic regularization technique designed for neural networks, consisting of randomly dropping out some neural units and their associated connections during training, with a predefined probability [7]. The main idea behind Monte Carlo dropout is to exploit the weight randomization induced by the dropout layers by keeping them active during inference; this allows one to approximate the distribution of the neural network weights with respect to training data. However, whereas dropout has been originally designed as a stochastic regularization technique to be used only during training, when it is used for uncertainty quantification its purpose is to estimate the weight distribution, and it is no more conceived as a regularizer. This suggests activating dropout layers during training may not be necessary if dropout is used only for uncertainty quantification. Accordingly, a recent work by Loquercio et al. [8] proposed to activate dropout layers for uncertainty quantification

only during inference, instead of using it also during training. The latter approach, which we call dropout *injection*, appears very interesting, as it would allow the *integration* of a measure of epistemic uncertainty into *any* network that was previously trained with no dropout, without the computational burden of a further training procedure required by the original Monte Carlo dropout, which we call *embedded* dropout, nor the need to optimize the related hyper-parameters.

Loquercio et al. [8] provided empirical evidence that injected dropout, combined with assumed density filtering [9] for aleatoric uncertainty evaluation, can be an effective tool for evaluating the *predictive* uncertainty of a neural network. However, no previous work, including the one by Loquercio et al. [8], provided a thorough analysis of injected dropout as an epistemic uncertainty evaluation technique, nor a comparison with classical embedded dropout.

Based on the above motivations, in this work we carry out a deep investigation of injected dropout, focusing on regression problems. First, our analysis shows that its effectiveness critically relies on a suitable scaling of the corresponding uncertainty measure, differently from embedded dropout. To this aim, we extend the original formulation by Loquercio et al. [8] to introduce and compute a suitable scaling factor, exploiting a technique originally devised for calibration [10]; our formulation is summarized in Fig. 1 (we refer the reader to Sect. 3 for the details). We also address the resulting issue of achieving a trade-off between different aspects related to the quality of the corresponding uncertainty measure, i.e., negative log-likelihood (NLL) and calibration error, as a function of the scale factor. We then carry out experiments on eight data sets from the University of California at Irvine (UCI) repository as well as five benchmarks for crowd counting, which is a challenging computer vision task, to thoroughly evaluate the behavior of injected dropout in terms of prediction accuracy and the quality of the corresponding uncertainty measure, providing experimental evidence of its effectiveness over embedded dropout.

Our results shed light on the practical operation of injected dropout and show that it can effectively evaluate epistemic uncertainty at test time without a time-consuming training process with dropout layers.

To sum up, our contributions are the following:

1. we conduct a thorough theoretical analysis of Monte Carlo dropout injection for regression problems;

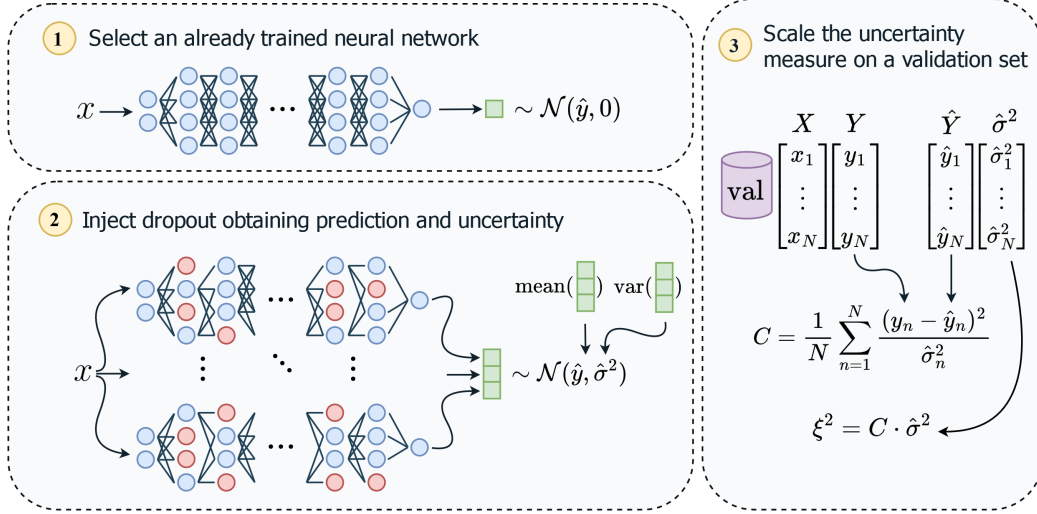


Figure 1: A schematic diagram of the three main steps that we propose to effectively implement injected dropout: 1) select a neural network trained for a regression problem without dropout (network output: $\hat{y} \in \mathcal{R}$); 2) add dropout layers to the network; 3) compute the dropout rate Φ and the scaling factor C of the uncertainty measure by solving the optimization problem in (9) for obtaining the scaled measure ξ^2 .

2. we reveal the existence of the scaling problem affecting such technique and propose a solution built upon sigma scaling [10] for its mitigation;
3. we conduct an experimental analysis confirming our theoretical results and investigating the differences between embedded dropout and the proposed implementation of injected dropout on several UCI datasets;
4. we verify the effectiveness of the proposed method by performing experiments on five benchmark crowd counting and density estimation data sets.

In the following, we present background concepts and describe related work on Bayesian methods for uncertainty evaluation, including Monte Carlo dropout, in Sect. 2; our theoretical analysis of injected dropout and the method we propose to implement it are presented in Sect. 3; experimental results are reported in Sect. 4; discussion and conclusions are given in Sect. 5, where future work on this topic is also discussed.

2. Background and Related Work

We first summarize the main concepts of uncertainty quantification in neural networks based on Bayesian methods; then, we focus on Monte Carlo dropout for epistemic uncertainty quantification. Finally, we review existing state-of-the-art methods for post hoc uncertainty quantification.

2.1. Uncertainty Evaluation through Bayesian Inference

To obtain a network capable of providing predictions with an associated level of uncertainty, one possible approach is a Bayesian extension. In particular, for a regression problem, a traditional neural network implements a predictor $\hat{y} = f(x; w) \in \mathbb{R}$, where w denotes the connection weights. Bayesian neural networks output a full distribution instead of a point prediction \hat{y} , which is usually modeled with a Gaussian distribution $\mathcal{N}(\hat{\mu}, \hat{\sigma}^2)$, albeit other choices are possible (e.g., Gaussian Mixture distributions or Generalized Gaussian distributions [11, 12]). In this case, the mean $\hat{\mu}$ is interpreted as the point prediction \hat{y} , whereas the variance $\hat{\sigma}^2$ represents the desired uncertainty measure; ideally, one wants $\hat{\mu}$ and $\hat{\sigma}^2$ to approximate well the *true* target distribution $\mathcal{N}(\mu, \sigma^2)$.

Under a Bayesian framework, the full probability distribution of the target variable y for an input query x is also conditioned on the training set \mathcal{D} [13]:

$$p(y|x, \mathcal{D}) = \int_w p(y|x, w)p(w|\mathcal{D})dw . \quad (1)$$

This is equivalent to taking the expectation under a posterior on the weights, $p(w|\mathcal{D})$, that provides the desired predictive distribution. However, since analytically obtaining $p(w|\mathcal{D})$ is intractable for most neural networks, an approximation $q(w)$ is usually considered. By substituting the true distribution with its approximation, the posterior can be computed [13] as:

$$p(y|x, \mathcal{D}) \approx \int_w p(y|x, w)q(w)dw . \quad (2)$$

A suitable approximation $q(w)$ is sought among a set of possible candidates by minimizing its Kullback-Leibler (KL) divergence [14] with the real distribution, $KL[p(w|\mathcal{D})||q(w)]$. From variational inference, it is known that this can be attained by minimizing the negative log-likelihood (NLL), that for regression problems is approximated [3] on a given data set of N samples as:

$$\frac{1}{N} \sum_{n=1}^N \frac{1}{2} \frac{(y_n - \hat{y}_n)^2}{\hat{\sigma}_n^2} + \frac{1}{2} \log(\hat{\sigma}_n^2) . \quad (3)$$

2.2. Monte Carlo Dropout as a Bayesian Approximation

To model an approximated distribution $q(w)$, one may consider incorporating a stochastic component inside a network. One possible way to introduce randomness on the network weights to estimate a full probability distribution over them is to use Stochastic Regularization Techniques [1], originally proposed for network regularisation. In particular, it has been shown both theoretically [6] and practically [6, 15, 3] that dropout can be used to this aim.

The standard dropout [7] works by the deactivation of each neuron belonging to a specific network layer with a probability, named *dropout rate*, that follows the Bernoulli distribution $\mathcal{B}(\Phi)$, where the parameter Φ corresponds to the dropout rate. The dropout layers are kept active during training and are deactivated at inference time, which produces a regularisation effect.

Unlike standard dropout, Monte Carlo dropout [6] keeps the dropout layers active at inference time; this allows one to derive a full probability distribution $q(w|\Phi)$ approximating $p(w|\mathcal{D})$. The optimal dropout rate Φ is defined as the one that minimizes the divergence between the real and the approximated distributions; in practice, a suitable Φ is found [6] by minimizing an approximation of NLL computed, e.g., on a validation set of N samples:

$$\Phi = \arg \min_{\varphi} \frac{1}{N} \sum_{n=1}^N \frac{1}{2} \frac{(y_n - \hat{y}_n(\varphi))^2}{\hat{\sigma}_n^2(\varphi)} + \frac{1}{2} \log(\hat{\sigma}_n^2(\varphi)) . \quad (4)$$

Once such a value of Φ and the trained network are obtained, one can use the approximate distribution $q(w|\Phi)$ for computing a prediction and associated predictive uncertainty, modeled as the mean and the variance of the posterior distribution [1]:

$$p(y|x, \mathcal{D}) \approx \frac{1}{T} \sum_{t=1}^T p(y|x, w_t) \sim \mathcal{N}(\hat{\mu}; \hat{y}, \hat{\sigma}^2) , \quad (5)$$

where T denotes the number of Monte Carlo samples obtained by querying the network multiple times with the same sample x . We call this technique *embedded dropout*.

Two widely used techniques for solving the optimization problem (4) are grid search [6] and concrete dropout [15]. In embedded dropout, after solving (4), one can use the obtained dropout rate in the inference stage. To

improve performance, a different dropout rate can be used for each network layer [15]. On the contrary, in the injected dropout approach [8] dropout layers are added to an already-trained network, and the dropout rate Φ is sought by minimizing the NLL (4) of the trained network, e.g., on a validation set.

2.3. *Post Hoc Uncertainty Quantification Methods*

The method we analyze in this paper belongs to the category of *post hoc* uncertainty quantification methods, meaning it can be applied to already trained neural networks.

In recent years several post hoc uncertainty quantification approaches have been proposed, such as new Laplace Approximations, surrogate models, and meta-learning approaches.

Among them we can find Laplace Approximations (LAs) [4] which is a family of techniques that are inherently post hoc. Indeed, these techniques start from a set of weights obtained by minimizing the NLL on a training set, which should correspond to a distribution mode. Then, one performs approximate inference by constructing a Gaussian “around” this set of weights, employing the Hessian as the approximated precision matrix. Such approximation methods have been used for Bayesian Inference since the 1990s [4] and recently gained popularity thanks to many computational advancements in second-order optimization [16]. Indeed, in a recent work by Ritter et al. [17], a Kroneken Factored Laplace Approximation has been employed. Another recent work by Kristiadi et al. [18] proposes to perform an “uncertainty training” procedure by adding additional neurons to the already trained network; such units are added in such a way that they change the geometry of the loss surface while preserving the prediction of the original model. This procedure allows training only a supplementary portion on top of an already trained network for obtaining a well-calibrated Laplace Approximation. Beyond Laplace Approximations, others proposed different strategies for obtaining a post-hoc uncertainty measure. Upadhyay et al. [19] proposed a procedure for obtaining a Bayesian model starting from a deterministic network. This technique, called BayesCap, consists of taking an already trained neural network and using their predictions to train a Bayesian surrogate model: the latter will learn to represent a generalized Gaussian distribution whose mean mimics the predictions of the original model. Other proposed methods use meta-learners [20, 21]: such methods work as model-agnostic predictors that try to associate an uncertainty measure by applying a learning policy on top of

the black-box predictions. Although these methods are conceived for classification and are not explicitly designed for neural networks, they are worth mentioning because they fall under the post hoc category for uncertainty quantification.

The injected dropout technique we focus on in this work is a further approach to post hoc uncertainty estimation. It was originally proposed by Loquercio et al. [8] for estimating the epistemic uncertainty component of an already-trained network. In that work, dropout injection was combined with Assumed Density Filtering (ADS) [22], which estimates the aleatoric uncertainty on the predictions, to obtain a measure of predictive uncertainty. Dropout injection has some advantages compared to many post hoc uncertainty quantification techniques. Unlike BayesCap and meta-learning approaches, it does not necessarily need further optimization. Moreover, beside a set of Monte Carlo forward predictions, it does not need to compute higher-order derivatives to obtain the Hessian (which is required for LAs). The limitation of this technique coincides with the ones of embedded dropout: the uncertainty quantification measure comes at the cost of computing multiple Monte Carlo forward passes at inference time. However, considering the overall state of the art of post hoc uncertainty quantification, dropout injection is an interesting technique that requires a relatively lower computational effort.

3. The Proposed Method for Monte Carlo Dropout Injection

A notable difference between embedded and injected dropout is that the former also performs a regularization, whereas the latter does not. This can lead to differences in the prediction error of the corresponding networks, as well as in their uncertainty measures, which we shall investigate in this section. In particular, we first analyze, theoretically and empirically, how prediction error and quality of the uncertainty measure behave as a function of the dropout rate Φ , for both embedded and injected dropout (Sect. 3.1). The results of this analysis will point out an issue of injected dropout, that we propose to mitigate by suitably *scaling* the uncertainty measure; to this aim, we propose an extension of the original formulation in Eq. (4), exploiting the σ -scaling technique [10]; we also investigate how the scaling factor affects the calibration of the uncertainty measure (Sect. 3.2).

For simplicity, we shall consider a fixed dropout rate across all network layers, but all our results apply to the most general case when different

dropout rates are allowed.

3.1. Analysis of Monte Carlo Dropout Injection

In Sect. 2.2 we mentioned that the dropout rate Φ is usually chosen by minimizing the NLL, both in embedded and injected dropout. Now, let us consider the weights vector w of a neural network: one can obtain w by using a standard optimization process that may or may not involve the activation of dropout layers during *training* for regularizing the network weights (i.e., embedded or injected dropout, respectively). Dropout layers are then activated at *inference* time, both when performing embedded and injected dropout, with a given dropout rate φ , in order to obtain a so-called *Monte Carlo iteration*, i.e. a prediction computed using the network with stochastic weights deactivation. This is equivalent to applying a binary mask $m(\varphi) = \text{diag}(z_1(\varphi) \dots z_L(\varphi))$ – where $z_\ell(\varphi) \sim \mathcal{B}(\varphi)$ and L denotes the size of w – to the network weights. Applying T stochastic iterations is equivalent to applying a series of T masks, each one denoted by $m_t(\varphi)$, to obtain a series of instances of the network weights, $w_1(\varphi) \dots w_T(\varphi)$ with $w_t(\varphi) = w \cdot m_t(\varphi)$, where \cdot denotes the matrix product. It is easy to see that $w_t(\varphi)$ always depends upon φ through the binary mask $m_t(\varphi)$, both for embedded and injected dropout. We can now rewrite Eq. (5) for modelling the prediction \hat{y} and the corresponding uncertainty $\hat{\sigma}^2$ of a given instance x when using a neural network with dropout rate φ :

$$\frac{1}{T} \sum_{t=1}^T p(y|x, w_t(\varphi)) = \frac{1}{T} \sum_{t=1}^T f(x; w_t(\varphi)) \sim \mathcal{N}(\hat{\mu}; \hat{y}, \hat{\sigma}^2) , \quad (6)$$

where $f(x; w_t(\varphi))$ denotes the prediction of the network parametrized with $w_t(\varphi)$ on the instance x .

With abuse of notation, let us denote with x , y , \hat{y} and $\hat{\sigma}^2$ the *vectors* of the instances, their ground truths, and the corresponding predictions and uncertainty measures, respectively, of a data set \mathcal{D} of size N . It is easy to see that, as the dropout rate changes, the prediction vector $\hat{y} = [\hat{y}_1 \dots \hat{y}_N]$ can change, and as a consequence also the quadratic error vector $\epsilon^2 = (y - \hat{y})^2$ and the uncertainty vector $\hat{\sigma}^2 = [\hat{\sigma}_1^2 \dots \hat{\sigma}_N^2]$ will change. In particular, choosing the dropout rate Φ by solving the optimization problem (4) amounts to pursuing two goals: (i) minimizing the quadratic error $\epsilon_n^2 = (y_n - \hat{y}_n)^2$; (ii) minimizing the absolute difference between the prediction’s quadratic error and the variance, $|\epsilon_n^2 - \hat{\sigma}_n^2|$. It is easy to show that the “best” trade-off

between them according to Eq. (4), i.e., the minimum NLL, is achieved when the quadratic error equals the variance for each sample of the set (a simple proof is reported in Appendix A):

$$\hat{\sigma}_n^2 = (y_n - \hat{y}_n)^2 = \epsilon_n^2 \quad \forall n \in \mathcal{D} . \quad (7)$$

A significant difference can now be noticed in the behavior of ϵ^2 and $\hat{\sigma}^2$, as a function of the dropout rate φ , between embedded and injected dropout. As we already pointed out, when injecting dropout the choice of φ does not affect the network weights w ; in embedded dropout, instead, changing φ results in explicitly modifying the vector w . Under the assumption of not having an under-parametrization, the influence of φ will lead to a regularization effect on w when using embedded dropout, but not when using injected dropout: since a network with dropout regularization is usually more robust to dropping network weights, it can be expected that a lower quadratic error is attained when using embedded dropout. Moreover, when using injected dropout we argue that, as φ increases, prediction quality quickly deteriorates, i.e., the quadratic error ϵ^2 increases. This behavior penalizes injected dropout, limiting to a tiny range $[0, \varphi_{\max}]$, for some φ_{\max} , the values of dropout rates φ that lead to acceptably small values of the quadratic error ϵ^2 . However, for very small values of dropout rate, the variability induced on the network predictions, and thus $\hat{\sigma}^2$, are so small that the resulting NLL is very high, which in turn is an indicator of poor quality of the uncertainty measure. Therefore, the range of “useful” values of φ further reduces to $[\varphi_{\min}, \varphi_{\max}]$, for some value $\varphi_{\min} < \varphi_{\max}$. We shall provide experimental evidence of this behavior in Sect. 4.2. In the next section, we propose to address the above issue of injected dropout through a suitable rescaling of its uncertainty measure.

3.2. The Proposed Dropout Injection Method

The problem exposed in Sect. 3.1 limits the dropout rate choice in an interval $[\varphi_{\min}, \varphi_{\max}]$, which can potentially lead to an impractical uncertainty measure due to its small magnitude. Nevertheless, such a degenerate measure can eventually exhibit desired behaviors: indeed, it can still associate high uncertainty with inaccurate predictions and low uncertainty with correct ones. The solution we propose to mitigate this issue is summarized in Fig. 1 and consists in properly *rescaling* the uncertainty measure $\hat{\sigma}^2$.

Rescaling the uncertainty measure – Let us consider a rescaled uncertainty measure $\xi^2(\Phi) = C \cdot \hat{\sigma}^2(\Phi)$, where $C \in (0, +\infty)$. Suitable values of

C and Φ can be sought by jointly optimizing the NLL as a function of both, through the following modification of the optimization problem (4):

$$\Phi, C = \arg \min_{\varphi, c} \frac{1}{N} \sum_{n=1}^N \frac{1}{2} \frac{(y_n - \hat{y}_n(\varphi))^2}{c \cdot \hat{\sigma}_n^2(\varphi)} + \frac{1}{2} \log(c \cdot \hat{\sigma}_n^2(\varphi)) . \quad (8)$$

We argue that the joint optimization in Eq. 8 can lead to a better NLL value than the original formulation (4), which may also correspond to a different optimal dropout rate Φ . An example is reported in Fig. 2: the proposed formulation (10) potentially allows one to explore regions of the NLL surface (considered as a function of *both* the dropout rate and the scale factor), that would otherwise be inaccessible by the original formulation (4), leading to a better NLL; we point out that the joint optimization in Eq. 8 can also provide a better solution (a lower NLL) than a distinct, two-stage optimization, i.e., first finding Φ by solving problem (10) and then finding C by solving problem (8) with respect only to c , keeping fixed Φ . Interestingly, we found that the optimization problem (8) can be solved as efficiently as (4), despite it involves the additional variable C , by exploiting the results by Laves et

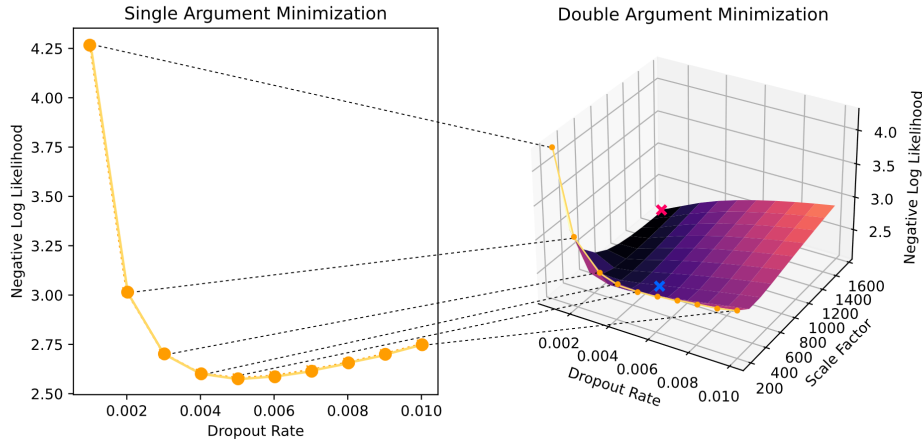


Figure 2: These plots (taken from experiments in Sect. 4.2) point out the difference between the scale-agnostic optimization process of Eq. (4) (left) and the scale-aware one of Eq. (8) (right). The latter can lead to different dropout rates with a lower NLL than the former, thanks to joint optimization (red cross \times in the surface on the right) instead of first solving Eq. (4) and then scaling the resulting measure (blue cross \times in the surface on the right).

al. [10]. Indeed, they found that it is possible to compute *analytically* the optimal scalar that, when multiplied with the uncertainty measure, minimizes the NLL of a predefined data set. However, the technique they designed (named σ -scaling [10]) was aimed at a different purpose, i.e., separately computing a scale factor to calibrate the uncertainty measure, after finding the optimal dropout rate by solving the original embedded dropout optimization problem (4). Our goal is instead to solve the optimization problem of Eq. (8); nevertheless, we found that their original formulation can be adapted to our objective. In this context, given a trained regressor providing a prediction \hat{y} and a corresponding uncertainty measure $\hat{\sigma}^2$, one can further reduce the NLL with respect to the solution of (3) by rescaling $\hat{\sigma}^2$ as $C \cdot \hat{\sigma}^2$. It can be shown that the optimal scale factor C , for any *given* quadratic error vector $\epsilon^2 = (y - \hat{y})^2$ and uncertainty vector $\hat{\sigma}^2$, is analytically given by:

$$C = \frac{1}{N} \sum_{n=1}^N \frac{(y_n - \hat{y}_n)^2}{\hat{\sigma}_n^2} = \frac{1}{N} \sum_{n=1}^N \frac{\epsilon_n^2}{\hat{\sigma}_n^2}. \quad (9)$$

The above result can be derived from the original formulation proposed by Laves et al. [10]; the only difference is to compute analytically a scale factor for the variance instead of the standard deviation. The proof is provided in Appendix B.

Accordingly, the optimal scale factor can be interpreted as an intrinsic indicator of how much the magnitude of the uncertainty vector is close to the magnitude of the quadratic error: if $C > 1$, the uncertainty measure $\hat{\sigma}_n$ is smaller, on average (over the N samples), than the quadratic error ϵ_n ; on the contrary, if $C < 1$, then $\hat{\sigma}_n$ is on average larger than ϵ_n . Finally, since the scale factor is analytically known for a given uncertainty measure, Eq. (9) can be exploited to rewrite C as a function of Φ in the optimization problem (8), which can therefore be simplified as:

$$\Phi = \arg \min_{\varphi} \frac{1}{N} \sum_{n=1}^N \frac{1}{2} \frac{(y_n - \hat{y}_n(\varphi))^2}{C(\varphi) \cdot \hat{\sigma}_n^2(\varphi)} + \frac{1}{2} \log(C(\varphi) \cdot \hat{\sigma}_n^2(\varphi)) . \quad (10)$$

Accordingly, both the dropout rate and the scale factor can be found by optimizing the NLL only with respect to the dropout rate, similarly to the original formulation (4), and at the *same* processing cost.

When introducing a scale factor, the values of the unscaled uncertainty measure $\hat{\sigma}_n^2$ do not need anymore to be close to the corresponding quadratic

error ϵ_n^2 to keep the NLL small, and are therefore not penalized as in the original formulation (4). We point out that, although rescaling can also be applied to embedded dropout, it is much more advantageous for injected dropout to mitigate the issue described in Sect. 3.1, which instead does not affect embedded dropout to a significant extent. In particular, such mitigation is expected to be more beneficial for large data sets, complex tasks, and networks with a large number of weights: we shall provide experimental evidence of this fact in Sect. 4.2.

Calibration of the uncertainty measure as a function of the scale factor – So far, we have focused on NLL minimization, since this allows one to minimize the distance between the actual prior distribution of the weights, $p(w)$, and the approximated one, $q(w)$. However, the quality of an uncertainty measure also depends on its *calibration*, and minimizing NLL does not guarantee to obtain a well-calibrated measure. Given the relevance of the scaling factor to injected dropout, it is interesting to investigate how it affects the trade-off between minimizing NLL and obtaining a well-calibrated measure. To this aim, we first summarise the notion of calibration of an uncertainty measure for regression problems. As mentioned in Sect. 2, the distribution $p(y|x, \mathcal{D})$ is assumed to be Normal, $\mathcal{N}(\mu, \sigma^2)$, as well as its approximation in Eq. (5), $\mathcal{N}(\hat{\mu}, \hat{\sigma}^2)$; to make the latter as close as possible to the former, one minimizes the NLL on a validation set. Considering μ and σ^2 as the point prediction and the corresponding uncertainty for a given instance (x, y) , one can construct a confidence interval (*CI*) with a desired probability α for a generic sample (x, y) :¹

$$CI_\alpha = [\mu - z_\alpha \cdot \sigma, \mu + z_\alpha \cdot \sigma] , \quad (11)$$

where z_α , known as *z-score*, is determined by:

$$p(y \in CI_\alpha) = \alpha, \quad \forall \alpha \in [0, 1] . \quad (12)$$

One would like that Eq. (12) holds for any instance (x, y) also with respect to a confidence interval \hat{CI}_α computed from the estimated distribution:

$$\hat{CI}_\alpha = [\hat{\mu} - z_\alpha \cdot \hat{\sigma}, \hat{\mu} + z_\alpha \cdot \hat{\sigma}] . \quad (13)$$

¹Note that, by definition, the confidence interval is defined using the standard deviation, not the variance.

If so, the uncertainty measure $\hat{\sigma}^2$ is said to be perfectly calibrated. This is, however, not guaranteed in practice.

To empirically evaluate the calibration of an uncertainty measure, the probability $p(y \in CI_\alpha)$ can be estimated on a data set (referred to as validation [23] or calibration set [24]), as the fraction of instances belonging to the confidence interval \hat{CI}_α , for several values of α . We call this fraction the *observed frequency*, and denote it by $\pi(\alpha)$; ideally, $\pi(\alpha)$ should be equal to α , which we call the *expected frequency*. Calibration quality can therefore be evaluated through a *calibration curve*, i.e., a plot of the observed frequency as a function of the expected one, as in the example shown in Fig. 3. Note that a perfect calibration corresponds to a calibration curve equal to the bisector of the X-Y axes. When the observed frequency $\pi(\alpha)$ is below the bisector, the uncertainty measure is *overconfident* for the corresponding values of α (because the expected frequency is greater than the observed one); on the contrary, where $\pi(\alpha)$ is above the bisector, the uncertainty measure is *under-confident*.²

Calibration quality can now be quantitatively evaluated as a scalar measure, in terms of “how far” the calibration curve is from the ideal one. In this work, we employ the area subtended by the calibration curve with respect to the bisector, known in the literature by the name of Miscalibration Area (MA) [25], which corresponds to the shaded regions in Fig. 3:

$$\text{MA} = \int_0^1 |\pi(\alpha) - \alpha| d\alpha . \quad (14)$$

In practice, the MA can be empirically estimated from a validation set as the average difference in absolute value between the expected and the observed frequency, over a set of M values of α :

$$\text{MA} = \frac{1}{M} \sum_{m=1}^M |\pi(\alpha_m) - \alpha_m| . \quad (15)$$

As mentioned above, the scale factor C obtained by solving problem (10) does not guarantee that the corresponding MA is null, i.e., perfect calibration. To analyze the trade-off between MA and NLL that can be obtained as

²Note that the calibration curve can cross the bisector for zero, one or more points, besides the ones corresponding to $\alpha = 0$ and $\alpha = 1$.

a function of the scale factor, we devised an iterative procedure that, starting from the obtained scale factor C , modifies it to reduce MA as much as possible, by increasing or reducing the confidence of each prediction of the calibration set, depending on whether the uncertainty measure is over- or under-confident on that prediction, respectively.

Our procedure (reported as Algorithm 1) aims at reducing to zero the average difference between the expected and the observed frequency (which we refer to as *Balance*) as a proxy for minimizing the MA. In such a setting, the value of *Balance* depends upon the chosen scale factor c ; we denote as $Balance_c$ such dependence, which is also reflected on the observed frequency $\pi_c(\alpha)$. Since $Balance_c$ is a monotonically non-decreasing function of c , one can empirically construct a suitable interval $[C_{\text{low}}, C_{\text{high}}]$ around the scaling factor C obtained by solving problem (10), such that $Balance_{C_{\text{low}}}$ is negative, and $Balance_{C_{\text{high}}}$ is positive; then, our algorithm uses the bisection method to find a “relaxed” scaling factor C_r such that $Balance_{C_r} = 0$ (in practice, the objective is defined as $Balance_{C_r} < \tau$, where τ is a given tolerance threshold).

By analyzing the “path” of the scale factor between C and C_r , i.e., the corresponding values of NLL and MA, it is possible to explore the trade-off between these two measures. Although acting on the scale factor does not guarantee to reduce MA to arbitrarily small values (i.e., to approach perfect calibration), empirical results reported in Sect. 4 show that small enough MA values can nevertheless be obtained. Interestingly, our results also show that the corresponding increase of NLL can be very limited: therefore, although standard post-hoc techniques can be applied (e.g., Kuleshov calibration [24]) to enhance the calibration quality, the above result suggests that, when using injected dropout, a good trade-off between minimizing NLL and obtaining a well-calibrated uncertainty measure can also be achieved by first computing the dropout rate and the corresponding scale factor through problem (10), and then suitably adjusting the latter.

4. Experiments

We carried out two groups of experiments. We first compared the behavior of embedded and injected dropout on eight data sets related to regression problems taken from the UCI repository³, using a multi-layer feed-forward

³<https://archive.ics.uci.edu/ml/datasets.php>

Algorithm 1 Modifying the scale factor to improve calibration

Require: vectors of uncertainty measures $\hat{\sigma}^2$, network predictions \hat{y} and ground truth y for a data set of N instances; a suitable interval $[C_{\text{low}}, C_{\text{high}}]$ around the initial scaling factor C ; M values $\alpha_1, \dots, \alpha_M$; tolerance threshold $\tau > 0$

Ensure: the relaxed scaling factor C_r

do

$C_r \leftarrow (C_{\text{low}} + C_{\text{high}})/2$

$\text{Balance}_{C_r} \leftarrow \frac{1}{M} \sum_{m=1}^M (\pi_{C_r}(\alpha_m) - \alpha_m)$

if $\text{Balance}_{C_r} < 0$ **then**

$C_{\text{low}} \leftarrow C_r$

else if $\text{Balance}_{C_r} > 0$ **then**

$C_{\text{high}} \leftarrow C_r$

end if

while $\text{Balance}_{C_r} \geq \tau$

network as the regression model. In the second group of experiments, we evaluated injected dropout on five benchmark data sets for crowd counting [26], which is a challenging computer vision application using a state-of-the-art Convolutional Neural Network (CNN) architecture for this task. In both experiments, we implemented injected dropout by setting the dropout rate and scaling factor according to Eq. (10), using a validation set; in the experiments on crowd counting we also considered the relaxed scaling factor obtained as described in Sect. 3.2 on the same validation set.

We evaluated both embedded and injected dropout, as a function of the dropout rate, in terms of the three metrics considered in Sect. 3: root mean squared error (RMSE), NLL, and MA. In particular, we evaluated NLL and MA using both the unscaled values $\hat{\sigma}^2(\varphi)$ and the scaled ones $\xi^2(\varphi)$; for reference, we also computed the “ideal” uncertainty measure of Eq. (7), defined as the error vector ϵ^2 computed on the *validation set*

4.1. Experiments on UCI Data Sets

We selected eight data sets related to several multivariate regression problems that have been used in the literature to evaluate uncertainty measures: Protein-Tertiary-Structure, Boston-Housing, Concrete, Energy, Kin8nm, Power-Plant, Wine-Quality-Red, and Yacht. We adopted the same set-up proposed by Hernandez-Lobato et al. [27], described in the following. We also open-




Figure 3: Example of a calibration curve, which plots the observed frequency of instances belonging to the confidence interval CI_α with probability α as a function of the respective expected frequency. In this example, an under-confidence and lower-confidence regions are present (see the text for their meaning).

sourced the code for the UCI experiment reproducibility at the following link https://github.com/EmanueleLedda97/Dropout_Injection.

Experimental set-up – We first distinguished between the challenging Protein-Tertiary-Structure data set and the remaining simpler ones. We used a one-hidden-layer network with 50 hidden units for the latter and 100 for the former. We randomly split each data set into a training set made up of 90% of the original instances and a testing set containing the remaining 10%. This procedure is repeated five times for Protein-Tertiary-Structure and 20 times for the other data sets. In each split, we used 20% of the training partition as a validation set, which is crucial not only for proper hyper-parameter tuning but also – and primarily – for scaling the uncertainty measure.

Differently from the original set-up [27], we used a mini-batch stochastic gradient descent; we empirically tuned the batch size, the learning rate, and the number of epochs on the validation set. We considered 15 possible dropout rates within the interval $[0.001, 0.5]$. We chose these rather extreme values to properly investigate the issue mentioned in Sect. 3.2, where we hypothesized a significant deterioration of the NLL of the unscaled measure for such values of dropout rate.

Experimental results – Results are reported in Fig. 4.




Figure 4: RMSE and NLL as a function of dropout rate on UCI data sets, for injected (blue) and embedded (red) dropout. Two NLL plots are reported for each data set, related to different ranges of dropout rate: $[0.001, 0.5]$ (top plot, same range as the RMSE plot) and $[0.05, 0.5]$ (bottom plot), to better visualize the Y-axis scale.

We can first notice that the RMSE of embedded dropout tends to deteriorate as the dropout rate increases. This behavior suggests that the network’s capacity is very low and that the regularization effect of dropout leads to under-fitting. This over-penalization makes injected dropout more accurate on average, in terms of average RMSE and standard deviation, although a different behavior can be observed on different data sets. Another phenomenon we can observe is that, for the smallest dropout rates, the RMSE is low, but the unscaled NLL is almost always severely affected. This trend reveals that – as we hypothesized in Sect. 3.2 – unscaled uncertainty measures suffer from the correlation between the dropout rate and the corresponding network perturbation’s magnitude. When the dropout rate is extremely low,

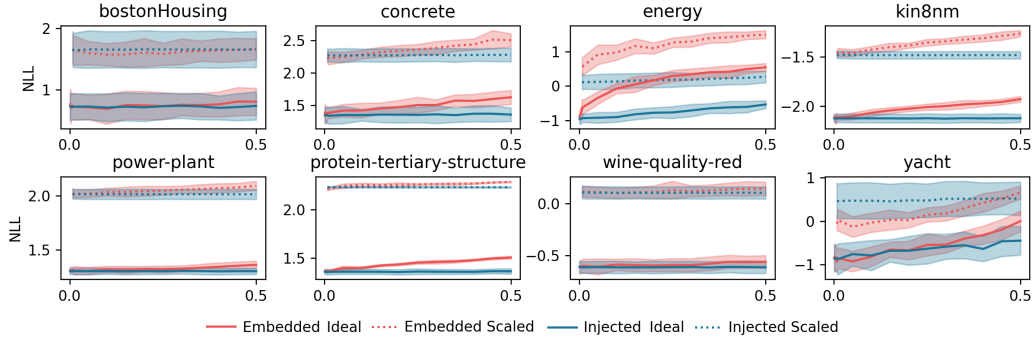


Figure 5: NLL as a function of dropout rate on UCI data sets, for the scaled (dotted line) and ideal (solid line) uncertainty measures, both for injected (blue) and embedded (red) dropout.

the final perturbation (and therefore the uncertainty measure) is so small that it completely disagrees with the quadratic prediction error, worsening the NLL.

Fig. 5 shows that the scaled measures exhibit significantly different behavior instead. In this case, it is evident that even for extremely low dropout rates, the NLL does not deteriorate as much as the unscaled one, in line with our hypothesis in Sect. 3.2. Although the observed trends are always in favor of the injected dropout, it may be possible that this behavior was due to the use of a multi-layer network with a hidden layer of 50 or 100 units, and – as we said before – there is clear evidence that using the dropout at training time in this set-up leads to under-fitting the considered data sets. This deterioration is also visible by looking at the ideal NLL: the solid red line in Fig. 5 tends to be higher than the solid blue line, which means that the ideal uncertainty measure (which is equal to the prediction error, for each sample) is already penalized. Therefore, the observed phenomena are entirely in line with our theoretical considerations in Sect. 3.2.

Looking at the calibration error in Fig. 6, we can see that dropout injection can provide well-calibrated uncertainty measurements. Analyzing the trends of the calibration error in detail, we can notice the following behaviors: first of all, the unscaled measures are clearly more sensitive to changes in dropout rate compared to their scaled counterparts. However, although the scaled measures seem much more stable, there are cases where the unscaled measure of embedded dropout performs comparably to or even outperforms the scaled one; this phenomenon occurs more sporadically – and much less

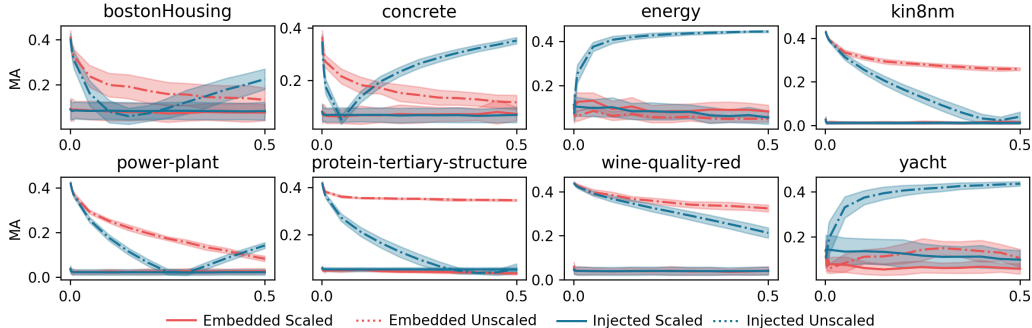


Figure 6: Calibration error as a function of dropout rate on UCI data sets, using the unscaled (dotted line) and scaled (solid line) uncertainty measures, both for injected (blue) and embedded (red) dropout.

clearly – for injected dropout. This is a crucial issue if we consider that, for the corresponding dropout rates, a trade-off can be observed between RMSE, NLL, and MA; this suggests that the embedded dropout inherently achieves a trade-off between these metrics. A further investigation of this issue is an interesting direction for future work.

4.2. Experiments on Crowd Counting

The experiments done on UCI data sets provide useful insights about the behavior of injected vs. embedded dropout and evidence that also the former can be an effective approach to epistemic uncertainty quantification. In the next experiments, we focused on a much more challenging computer vision problem, i.e., crowd counting. State-of-the-art CNN models solve it by first estimating the crowd density map, which is an image-to-image regression problem.

Experimental set-up – We used three benchmark crowd counting data sets, University of California San Diego (UCSD) dataset [28], Mall [29], and the Pedestrian Evaluation on Tracking and Surveillance (PETS) dataset [30]. For the latter, we used three different partitions corresponding to different crowd scenes, named PETSview1, PETSview2, and PETSview3 [31]. We used the Multi-Column Neural Network (MCNN) architecture [32], a well-known state-of-the-art CNN for crowd counting. We implemented dropout injection right before each convolutional layer, except for the first one, following the suggestion of previous works on Monte Carlo dropout [15]. Unlike the previous experiments, we found an exponential increase in the RMSE




Figure 7: RMSE, NLL, and calibration error of the injected dropout as a function of dropout rate for the crowd counting data sets. Two plots are reported for NLL, corresponding to the unscaled (top plot) and scaled (bottom plot) uncertainty measure. The plots of scaled NLL and calibration error report also a comparison with the ideal scale (the former) and with the unscaled measure (the latter)

using the same range of dropout rates. Therefore, we followed the suggestion by Loquercio et al. [8] to use very small dropout rates, which leads to a quasi-deterministic network. In particular, we considered ten dropout rates in the interval $[0.001, 0.01]$.

Experimental results – Fig. 7 reports the RMSE, NLL (scaled and unscaled), and MA as a function of dropout rate. Since there is a significant difference in the magnitude of the NLL for the unscaled measures and the scaled ones, we reported them into two different plots.

We can observe an increase in RMSE in three of five data sets, while for two of them (PETSview1 and PETSview3), the RMSE decreased. This apparently different behavior is due to the very small dropout rates used in these experiments; we observed indeed that for higher dropout rates, the RMSE started to increase also for PETSview1 and PETSview3.

Regarding the NLL, a large difference between the scaled and the unscaled measures can be observed, even in this experiment: similarly to the UCI data sets, the smallest dropout rates lead to a poor unscaled measure, whereas

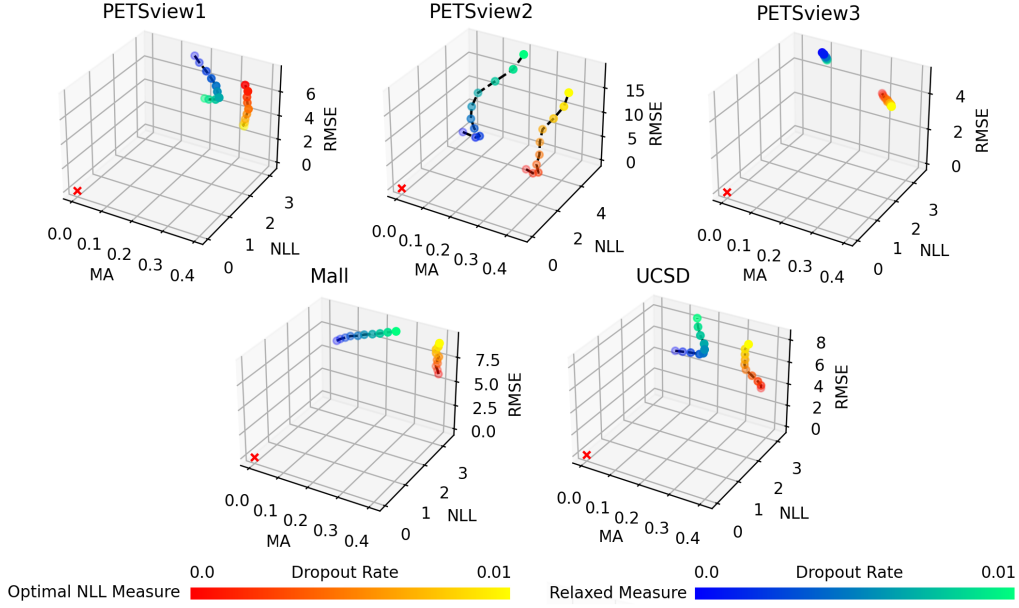


Figure 8: Trade-off between RMSE, NLL, and MA of the injected dropout, as a function of dropout rate, on the crowd counting data sets, for the optimal scaling (warm color gradient) and the relaxed one (cold color gradient).

the scaled measure attained considerably better results. It is also possible to see how the best dropout rate values differ between scaled and unscaled measures: for example, in UCSD, the minimum NLL is around $\Phi = 0.04$ for the scaled measure whereas, for the unscaled measure, it is around $\Phi = 0.01$. Similar considerations can be made for Mall. We can also notice that the scaled measure is much more stable than the unscaled one, even with the smallest dropout rates: this can be advantageous when using injected dropout, as it allows one to make minimal changes to the network to obtain suitable measures of uncertainty. Although re-scaling performs well in terms of the NLL, the same does not hold for the calibration error: in this case, sub-optimal performances can be observed, with – in most cases – a miscalibration area greater than 0.3, which is an indicator of poor calibration. Therefore, in this case, a suitable trade-off between NLL and calibration error should be sought, e.g., by tuning the scaling factor obtained from Eq. (10) through the relaxation procedure presented in Sect. 3.2.

Using a relaxed version of the scaling factor allows one to attain a trade-




Figure 9: Calibration curves of injected dropout on the crowd counting data sets, attained using the unscaled measure (blue), the optimal scaling factor C (red), the relaxation C_r (orange). For reference, the perfect calibration line is also shown (dotted line).

off between RMSE, NLL, and MA, instead of minimizing only the NLL, by manually tuning the scaling factor; Fig. 8 shows the value of the three analyzed metrics as a function of dropout when using both C (cold-colored gradient) and C_r (warm-colored gradient). It is easy to notice that when using C_r , the NLL is not far from the values obtained using C : this suggests that the relaxed measure behaves as a reasonable compromise between NLL and MA.

Finally, the calibration curves in Fig. 9 clearly show that the proposed relaxation procedure effectively reduces the calibration error. In all cases, the relaxed curve is much closer to the axes' bisector, supporting our hypothesis in Sect. 3.2.

Practical utility of uncertainty quantification by dropout injection in video surveillance applications – We first show how the operator of a video surveillance system can benefit from updating its crowd density

estimator software by integrating an uncertainty estimate through dropout injection (i.e., without requiring further training), and then how uncertainty evaluation could be automatically exploited to improve the accuracy of the estimated density maps and of the corresponding people count.

Crowd counting by density estimation uses density maps as ground truth: the standard method for obtaining such maps starts by annotating each image location where one person’s head is present, then the density map is obtained by adding a kernel (e.g., Gaussian) with a unit area centered on each head location. Therefore, the people count corresponds to the sum of each pixel of the density map. In this context, a crowd counting predictor will output an estimated density map reconstructing the original one. The image-to-image nature of this task, similar to what happens in Monocular Depth Estimation and Semantic Segmentation, implies that also the uncertainty estimate consists of a 2D array, associating a level of uncertainty to each pixel-wise predicted density (i.e., its variance). In turn, the pixel-wise uncertainty on the predicted density allows the construction of a confidence interval on the people count, as seen in the examples in Fig. 10, similarly to [26], but in our case, in a post hoc fashion. Having a notion of the confidence level associated with system predictions makes video surveillance operators aware of their reliability, allowing them to make more informed decisions and increasing their trust in the prediction system compared to a point estimator.

An in-depth analysis of our experimental results on crowd counting benchmarks reveals that uncertainty evaluation could also be helpful to detect some inaccuracies in the estimated density maps that correspond to “false positive” pedestrian localization, i.e., image regions where the predicted density is not zero (thus contributing to the pedestrian count) despite the absence of pedestrians. As seen from the examples in Fig. 10, we observed that false positive localizations are often characterized by isolated regions with relatively low estimated density and high uncertainty. This suggests that they could be detected and automatically corrected (by setting to zero the corresponding values in the density map and disregarding their contribution to the people count), thus improving the accuracy of both the density map and the people count. As an example, a simple policy for automatic rejection of such regions may consist in rejecting isolated regions with high pixel-level uncertainty/prediction ratio, but other solutions are conceivable. Regardless of the policy, it is intriguing to see that such automatic rejection mechanisms are theoretically possible and practically applicable by integrating dropout injection in the toolbox of video surveillance software. A further investigation

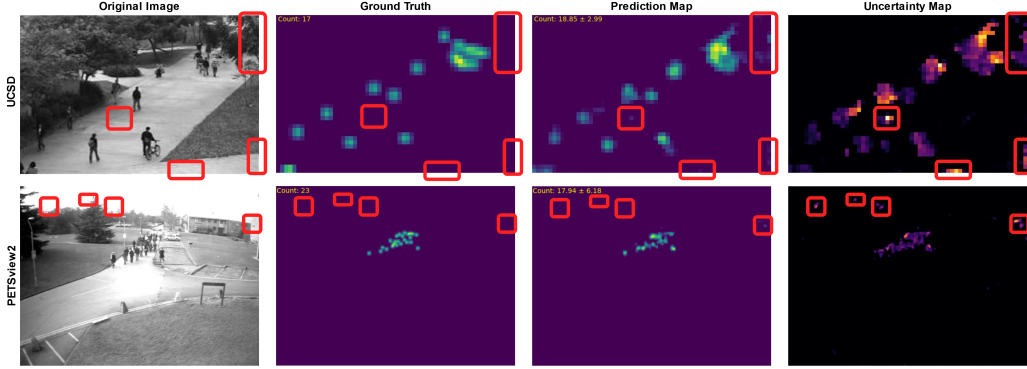


Figure 10: Two examples of “false positive” detections enabled by uncertainty estimation on crowd-counting images (top: UCSD, bottom: PETSview2). From left to right: the original frame, ground truth, predicted density map, and uncertainty map. Each density map is normalized: darker pixels on the gradient palette (blue for ground truth and predicted maps, red for the uncertainty maps) correspond to lower values. “False positive” regions, corresponding to low predicted density and high uncertainty are highlighted with red boxes.

of this issue is an interesting topic for future work.

5. Conclusions

This work provides the first investigation of injected dropout for epistemic uncertainty evaluation in order to assess if and how this method can be a practical alternative to the current use of the traditional Monte Carlo (embedded) dropout, which needs time-consuming uncertainty-aware training (i.e., it should be designed at training time). Our analysis focused on regression problems and encompassed two different dimensions related to the behavior of injected dropout, namely prediction accuracy, measured in terms of RMSE, and quality of the uncertainty measure, evaluated in terms of NLL and MA (i.e., calibration).

We analyzed, theoretically and empirically, how prediction error and quality of the uncertainty measure behave as a function of the dropout rate and proposed a novel injection method grounded on the results of our analysis.

Experimental results on several benchmark data sets, including a challenging computer vision task, provided clear evidence that dropout injection can be an effective and efficient alternative to embedded dropout, also considering that it does not require a computationally costly re-training procedure

of already-trained networks. Our results also confirm that the effectiveness of injected dropout critically depends on a suitable re-scaling of its uncertainty measure.

We point out two limitations of our analysis, that can be addressed in future work: (i) We considered only regression problems; an obvious extension is to analyze dropout injection for classification problems. (ii) We considered the same dropout rate across different network layers; however, our analysis can be easily extended to the case of different dropout rates across the layers, which is known to be potentially beneficial to Monte Carlo dropout [15].

Finally, another interesting issue for future work is to investigate the effectiveness of injected dropout for detecting out-of-distribution samples, which is potentially useful under domain shifts.

Acknowledgments

This work was partially supported by the projects: “Law Enforcement agencies human factor methods and Toolkit for the Security and protection of CROWDs in mass gatherings” (LETSCROWD), EU Horizon 2020 programme, grant agreement No. 740466; “IMaging MAnagement Guidelines and Informatics Network for law enforcement Agencies” (IMMAGINA), European Space Agency, ARTES Integrated Applications Promotion Programme, contract No. 4000133110/20/NL/AF; “Science and engineering Of Security of Artificial Intelligence” (S.O.S. AI) included in the Spoke 3 - Attacks and Defences of the Research and Innovation Program PE00000014, “SEcurity and RIghts in the CybeRspace (SERICS)”, under the National Recovery and Resilience Plan, Mission 4 ”Education and Research” - Component 2 ”From Research to Enterprise” - Investment 1.3, funded by the European Union - NextGenerationEU.

Emanuele Ledda is affiliated with the Italian National PhD in Artificial Intelligence, Sapienza University of Rome. He also acknowledges the cooperation with and the support from the Pattern Recognition and Applications Lab of the University of Cagliari.

References

- [1] Y. Gal, Uncertainty in deep learning (2016).
- [2] E. Hüllermeier, W. Waegeman, Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods, Vol. 110,

Springer US, 2021. doi:10.1007/s10994-021-05946-3.
URL <https://doi.org/10.1007/s10994-021-05946-3>

- [3] A. Kendall, Y. Gal, What uncertainties do we need in bayesian deep learning for computer vision?, in: I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, December 4-9, 2017, Long Beach, CA, USA, 2017, pp. 5574–5584.
URL <https://proceedings.neurips.cc/paper/2017/hash/2650d6089a6d640c5e85b2b88265dc2b-Abstract.html>
- [4] D. J. C. MacKay, *A practical bayesian framework for backpropagation networks*, *Neural Computation* (1992).
- [5] R. M. Neal, *Bayesian learning for neural networks* (1995).
- [6] Y. Gal, Z. Ghahramani, Dropout as a bayesian approximation: Representing model uncertainty in deep learning, in: M. Balcan, K. Q. Weinberger (Eds.), *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016*, New York City, NY, USA, June 19-24, 2016, Vol. 48 of *JMLR Workshop and Conference Proceedings*, JMLR.org, 2016, pp. 1050–1059.
URL <http://proceedings.mlr.press/v48/gal16.html>
- [7] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *Journal of Machine Learning Research* 15 (56) (2014) 1929–1958.
URL <http://jmlr.org/papers/v15/srivastava14a.html>
- [8] A. Loquercio, M. Segu, D. Scaramuzza, A general framework for uncertainty estimation in deep learning, *IEEE Robotics and Automation Letters* 5 (2020) 3153–3160. doi:10.1109/LRA.2020.2974682.
- [9] X. Boyen, D. Koller, Tractable inference for complex stochastic processes, in: G. F. Cooper, S. Moral (Eds.), *UAI '98: Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, University of Wisconsin Business School, Madison, Wisconsin, USA, July 24-26, 1998, Morgan Kaufmann, 1998, pp. 33–42.

- [10] M. Laves, S. Ihler, J. F. Fast, L. A. Kahrs, T. Ortmaier, Well-calibrated regression uncertainty in medical imaging with deep learning, in: T. Arbel, I. B. Ayed, M. de Bruijne, M. Descoteaux, H. Lombaert, C. Pal (Eds.), International Conference on Medical Imaging with Deep Learning, MIDL 2020, 6-8 July 2020, Montréal, QC, Canada, Vol. 121 of Proceedings of Machine Learning Research, PMLR, 2020, pp. 393–412. URL <http://proceedings.mlr.press/v121/laves20a.html>
- [11] C. Blundell, J. Cornebise, K. Kavukcuoglu, D. Wierstra, Weight uncertainty in neural networks, 32nd International Conference on Machine Learning, ICML 2015 2 (2015) 1613–1622.
- [12] U. Upadhyay, S. Karthik, Y. Chen, M. Mancini, Z. Akata, Bayescap: Bayesian identity cap for calibrated uncertainty in frozen neural networks, in: S. Avidan, G. J. Brostow, M. Cissé, G. M. Farinella, T. Hassner (Eds.), Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XII, Vol. 13672 of Lecture Notes in Computer Science, Springer, 2022, pp. 299–317. doi:10.1007/978-3-031-19775-8_18. URL https://doi.org/10.1007/978-3-031-19775-8_18
- [13] M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. W. Fieguth, X. Cao, A. Khosravi, U. R. Acharya, V. Makarenkov, S. Nahavandi, A review of uncertainty quantification in deep learning: Techniques, applications and challenges, Inf. Fusion 76 (2021) 243–297. doi:10.1016/j.inffus.2021.05.008. URL <https://doi.org/10.1016/j.inffus.2021.05.008>
- [14] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, L. K. Saul, An introduction to variational methods for graphical models, Mach. Learn. 37 (2) (1999) 183–233. doi:10.1023/A:1007665907178.
- [15] Y. Gal, J. Hron, A. Kendall, Concrete dropout, in: I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, R. Garnett (Eds.), Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA, 2017, pp. 3581–3590. URL <https://proceedings.neurips.cc/paper/2017/hash/84ddfb34126fc3a48ee38d7044e87276-Abstract.html>

- [16] J. Martens, R. B. Grosse, Optimizing neural networks with kronecker-factored approximate curvature, in: F. R. Bach, D. M. Blei (Eds.), Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015, Vol. 37 of JMLR Workshop and Conference Proceedings, JMLR.org, 2015, pp. 2408–2417.
URL <http://proceedings.mlr.press/v37/martens15.html>
- [17] H. Ritter, A. Botev, D. Barber, A scalable laplace approximation for neural networks, in: 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings, OpenReview.net, 2018.
URL <https://openreview.net/forum?id=Skdvd2xAZ>
- [18] A. Kristiadi, M. Hein, P. Hennig, Learnable uncertainty under laplace approximations, in: C. P. de Campos, M. H. Maathuis, E. Quaeghebeur (Eds.), Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI 2021, Virtual Event, 27-30 July 2021, Vol. 161 of Proceedings of Machine Learning Research, AUAI Press, 2021, pp. 344–353.
URL <https://proceedings.mlr.press/v161/kristiadi21a.html>
- [19] U. Upadhyay, S. Karthik, Y. Chen, M. Mancini, Z. Akata, Bayescap: Bayesian identity cap for calibrated uncertainty in frozen neural networks, in: S. Avidan, G. J. Brostow, M. Cissé, G. M. Farinella, T. Hassner (Eds.), Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XII, Vol. 13672 of Lecture Notes in Computer Science, Springer, 2022, pp. 299–317. doi:10.1007/978-3-031-19775-8_18.
URL https://doi.org/10.1007/978-3-031-19775-8_18
- [20] P. Lahoti, P. K. Gummadi, G. Weikum, Responsible model deployment via model-agnostic uncertainty learning, Mach. Learn. 112 (3) (2023) 939–970. doi:10.1007/s10994-022-06248-y.
URL <https://doi.org/10.1007/s10994-022-06248-y>
- [21] H. Jiang, B. Kim, M. Y. Guan, M. R. Gupta, To trust or not to trust A classifier, in: S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information

- Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada, 2018, pp. 5546–5557.
 URL <https://proceedings.neurips.cc/paper/2018/hash/7180cffd6a8e829dacfc2a31b3f72ece-Abstract.html>
- [22] X. Boyen, D. Koller, Tractable inference for complex stochastic processes, in: G. F. Cooper, S. Moral (Eds.), UAI '98: Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, University of Wisconsin Business School, Madison, Wisconsin, USA, July 24-26, 1998, Morgan Kaufmann, 1998, pp. 33–42.
 URL https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article_id=230&proceeding_id=14
- [23] C. Guo, G. Pleiss, Y. Sun, K. Q. Weinberger, On calibration of modern neural networks, in: D. Precup, Y. W. Teh (Eds.), Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017, Vol. 70 of Proceedings of Machine Learning Research, PMLR, 2017, pp. 1321–1330.
 URL <http://proceedings.mlr.press/v70/guo17a.html>
- [24] V. Kuleshov, N. Fenner, S. Ermon, Accurate uncertainties for deep learning using calibrated regression, in: J. G. Dy, A. Krause (Eds.), Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018, Vol. 80 of Proceedings of Machine Learning Research, PMLR, 2018, pp. 2801–2809.
 URL <http://proceedings.mlr.press/v80/kuleshov18a.html>
- [25] K. Tran, W. Neiswanger, J. Yoon, Q. Zhang, E. P. Xing, Z. W. Ulissi, Methods for comparing uncertainty quantifications for material property predictions, *Mach. Learn. Sci. Technol.* 1 (2) (2020) 25006. doi:10.1088/2632-2153/ab7e1a.
 URL <https://doi.org/10.1088/2632-2153/ab7e1a>
- [26] M. H. Oh, P. Olsen, K. N. Ramamurthy, Crowd counting with decomposed uncertainty, *AAAI 2020 - 34th AAAI Conference on Artificial Intelligence* (2020) 11799–11806doi:10.1609/aaai.v34i07.6852.
- [27] J. M. Hernández-Lobato, R. P. Adams, Probabilistic backpropagation for scalable learning of bayesian neural networks, in: Proceedings of the

- 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015, Vol. 37 of JMLR Workshop and Conference Proceedings, JMLR.org, 2015, pp. 1861–1869.
URL <http://arxiv.org/abs/1502.05336>
- [28] A. B. Chan, Z. J. Liang, N. Vasconcelos, Privacy preserving crowd monitoring: Counting people without people models or tracking, in: 2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2008), 24-26 June 2008, Anchorage, Alaska, USA, IEEE Computer Society, 2008. doi:10.1109/CVPR.2008.4587569.
URL <https://doi.org/10.1109/CVPR.2008.4587569>
 - [29] K. Chen, C. C. Loy, S. Gong, T. Xiang, Feature mining for localised crowd counting, in: R. Bowden, J. P. Collomosse, K. Mikolajczyk (Eds.), British Machine Vision Conference, BMVC 2012, Surrey, UK, September 3-7, 2012, BMVA Press, 2012, pp. 1–11. doi:10.5244/C.26.21.
URL <https://doi.org/10.5244/C.26.21>
 - [30] J. M. Ferryman, A. Shahrokni, Pets2009: Dataset and challenge, 2009 Twelfth IEEE International Workshop on Performance Evaluation of Tracking and Surveillance (2009) 1–6.
 - [31] R. Delussu, L. Putzu, G. Fumera, Investigating synthetic data sets for crowd counting in cross-scene scenarios, in: G. M. Farinella, P. Radeva, J. Braz (Eds.), Proceedings of the 15th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, VISIGRAPP 2020, Volume 4: VISAPP, Valletta, Malta, February 27-29, 2020, SCITEPRESS, 2020, pp. 365–372. doi:10.5220/0008981803650372.
URL <https://doi.org/10.5220/0008981803650372>
 - [32] Y. Zhang, D. Zhou, S. Chen, S. Gao, Y. Ma, Single-image crowd counting via multi-column convolutional neural network, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016, IEEE Computer Society, 2016, pp. 589–597. doi:10.1109/CVPR.2016.70.
URL <https://doi.org/10.1109/CVPR.2016.70>

Appendix A. Proof of Eq. (7)

In the following we prove that, given a data set $\mathcal{D} = \{(x_n, y_n)\}_{n=1}^N$ for a regression problem, and denoting the corresponding predictions of a given regressor by \hat{y}_n , the ideal uncertainty measure σ^{*2} , i.e., the one that minimizes the NLL, is obtained when the following condition holds:

$$\sigma_n^{*2} = (y_n - \hat{y}_n)^2, \quad n = 1, \dots, n. \quad (\text{A.1})$$

Consider any instance $x_n \in \mathcal{D}$. The derivative of the NLL with respect to the corresponding σ_n^2 is:

$$\frac{\partial \left(\frac{1}{2} \cdot \frac{(y_n - \hat{y}_n)^2}{\sigma_n^2} + \frac{1}{2} \cdot \log(\sigma_n^2) \right)}{\partial \sigma_n^2} = \quad (\text{A.2})$$

$$\frac{1}{2} \cdot \left((y_n - \hat{y}_n)^2 \cdot \frac{\partial \left(\frac{1}{\sigma_n^2} \right)}{\partial \sigma_n^2} + \frac{\partial (\log(\sigma_n^2))}{\partial \sigma_n^2} \right) = \quad (\text{A.3})$$

$$\frac{1}{2} \cdot \left(\frac{-(y_n - \hat{y}_n)^2}{(\sigma_n^2)^2} + \frac{1}{\sigma_n^2} \right) = \quad (\text{A.4})$$

$$\frac{\sigma_n^2 - (y_n - \hat{y}_n)^2}{2(\sigma_n^2)^2} \quad (\text{A.5})$$

Setting the above derivative to 0 and solving for σ_n^2 we obtain:

$$\frac{\sigma_n^2 - (y_n - \hat{y}_n)^2}{2(\sigma_n^2)^2} = 0, \quad (\text{A.6})$$

$$\sigma_n^2 = (y_n - \hat{y}_n)^2, \quad \sigma_n^2 \neq 0 \quad (\text{A.7})$$

Appendix B. Proof of Eq. (9)

In Laves et al. [10] the sigma scaling has been computed with respect to the standard deviation. It is straightforward to compute the resulting scaling factor with respect to the variance.

Let C be the optimal scale factor when using the variance as an uncertainty measure (namely the optimal variance scaler), and C' be the one

optimal one obtained by using the standard deviation (namely the optimal standard deviation scaler). First, the value of C' can be obtained as [10]:

$$C' = \sqrt{\frac{1}{N} \sum_{n=1}^N \frac{(y_n - \hat{y}_n)^2}{\hat{\sigma}_n^2}} = \sqrt{\frac{1}{N} \sum_{n=1}^N \frac{\epsilon_n^2}{\hat{\sigma}_n^2}} \quad (\text{B.1})$$

We remind the reader that the optimal variance measure ξ^2 is obtained by multiplying the variance σ^2 with the optimal variance scaler C : $\xi^2 = C \cdot \sigma^2$. We can obtain, instead, the optimal standard deviation measure ξ by multiplying the standard deviation σ with the optimal standard deviation scaler C' : $\xi = C' \cdot \sigma$. From these identities we can derive the relation between the optimal variance scaler C and the optimal standard deviation scaler C' :

$$\xi^2 = C \cdot \sigma^2 \quad (\text{B.2})$$

$$\xi = \sqrt{C \cdot \sigma^2} \quad (\text{B.3})$$

$$C' \cdot \sigma = \sqrt{C \cdot \sigma^2} \quad (\text{B.4})$$

$$C' \cdot \sigma = \sigma \cdot \sqrt{C} \quad (\text{B.5})$$

$$C' = \sqrt{C} \quad (\text{B.6})$$

From B.1 we can now deduce that the optimal variance scaler C is given by:

$$C = \frac{1}{N} \sum_{n=1}^N \frac{(y_n - \hat{y}_n)^2}{\hat{\sigma}_n^2} = \frac{1}{N} \sum_{n=1}^N \frac{\epsilon_n^2}{\hat{\sigma}_n^2} \quad (\text{B.7})$$