

Clinical Measurements with Calibrated Instance-Dependent Confidence Intervals

Rotem Nizhar¹

ROTEMNIZHAR10@GMAIL.COM

¹ *Computer Science Dept. Tel-Aviv University*

Lior Frenkel²

LIORFRENKEL1992@GMAIL.COM

Jacob Goldberger²

JACOB.GOLDBERGER@BIU.AC.IL

² *Engineering Faculty, Bar-Ilan University*

Editors: Under Review for MIDL 2025

Abstract

Reporting meaningful confidence intervals for the predictions of a regression neural network is critical in medical imaging applications since clinical decisions rely on network predictions. We expect to obtain larger intervals for difficult examples and smaller ones for easier examples to predict. We demonstrate that training a Gaussian regression network, followed by a non-parametric conformal prediction technique to scale the estimated variance, is the most effective way to achieve a small confidence interval with a coverage guarantee. Through extensive experiments on various medical imaging datasets and network architectures, we show that this combined training and calibration procedure produces improved results compared to previous methods

Keywords: confidence interval, uncertainty calibration, conformal prediction

1. introduction

Regression neural networks, which predict continuous quantities, have been a major focus of research in many areas including medical diagnosis (Leibig et al., 2017), weather forecasting (Scher and Messori, 2018) and autonomous driving (Carvalho et al., 2015). Regression neural networks are applied in medical image analysis to measure the size of pathological lesions and the size of anatomical parts and the distance between them. One example is ultrasound-based automatic estimation of fetal biometry which is used to assess the growth and well-being of the fetus (Avisdris et al., 2022). Another example is estimating the bone age of pediatric patients based on radiographs of their hand (Halabi et al., 2019).

The performance of neural network systems has improved dramatically in recent years. However, for safety-critical embodied applications, accurate prediction alone is not sufficient. Uncertainty estimates are important in a wide range of applications, and reporting the confidence of a prediction is essential for reliable and interpretable models. One widely adopted approach to conveying uncertainty is confidence intervals, which enclose the “true value” with a specified probability. The size of these intervals is expected to be small and linked to the case’s complexity.

Standard regression networks are trained by minimizing the Mean Squared Error (MSE). These networks provides prediction intervals that have the same length for all test examples, and thus potentially cannot directly report an instance-based confidence interval. However, it is much more informative to provide larger confidence intervals for difficult examples and

smaller ones for easier examples to predict. An alternative to the MSE approach is to predict the mean and variance simultaneously. The training loss is then formed by the negative Gaussian log-likelihood. During the testing phase, assuming a Gaussian distribution, the predicted mean and variance values can be translated into an instance-based confidence interval.

Consider a regression model that reports confidence intervals with a claimed coverage of $1 - \alpha$. If $1 - \alpha$ of the intervals indeed contain the true value, the model is called calibrated. Deep networks are not well-calibrated and are known to produce unreliable confidence information (Guo et al., 2017). Several recent methods are available for measuring the calibration of a regression network such as the Expected Normalized Calibration Error (ENCE) (Levi et al., 2022) and the Uncertainty Calibration Error (UCE) (Laves et al., 2020; Küppers et al., 2022). Several studies (Laves et al., 2020; Levi et al., 2022; Frenkel and Goldberger, 2023) have proposed a simple calibration method that scales the predictive variance by optimizing a likelihood criterion on a validation set. However, these calibration methods explicitly assume that the conditional density of the correct value given the image is Gaussian, which is not always the case.

Conformal Prediction (CP) (Vovk et al., 2005; Angelopoulos and Bates, 2023) is a general non-parametric calibration method which, given a confidence value, aims to build a confidence interval such that the probability that the correct value is within this set, is indeed the given value. The Conformalized Quantile Regression (CQR) algorithm (Romano et al., 2019) is a calibrated regression method that directly finds a confidence interval without any parametric assumption of the prediction distribution. It consists of a Quantile Regression (QR) (Koenker and Bassett Jr, 1978) followed by a conformalization step. On one hand, CQR is more adaptive to heteroscedasticity and outliers than a Gaussian regression. On the other hand, the pinball loss which is used to train the Quantile regression is less reliable and stable than the Gaussian loss.

CQR is considered the state-of-the-art method to obtain a calibrated instance-based confidence interval of a prediction obtained by a neural network. However, we are not aware of any comparative research, either in medical or non-medical data, (see e.g., a discussion in a recent review paper (Kato et al., 2023)). In the current study, we first analyze and demonstrate the advantages and disadvantages of the current methods. We then propose a calibration strategy based on training a Gaussian network followed by a parameter-free CP calibration of the computed variance. We show that this strategy improves the confidence calibration procedure. We report extensive experiments on several medical imaging regression tasks and network architectures that support this combined training and calibration procedure.

2. Calibrating regression networks

In this section we review the current parametric and non-parametric methods for instance-dependent calibration of a regression network. Consider a regression network that outputs mean $\hat{y} = \mu(x)$ and variance $\sigma^2(x)$ for each input image x . The mean represents the value predicted by the network, while the variance is the level of uncertainty in the prediction. The network output can be viewed as a Gaussian distribution in the form of $y|x \sim \mathcal{N}(\mu(x), \sigma^2(x))$. Given labeled training data $(x_1, y_1), \dots, (x_n, y_n)$, the network is

trained by minimizing the loss function:

$$L(\theta) = - \sum_{t=1}^n \log \mathcal{N}(y_t; \mu_\theta(x_t), \sigma_\theta^2(x_t)) \quad (1)$$

such that θ is the network parameter set. From the distribution $y \sim \mathcal{N}(\mu(x), \sigma^2(x))$ we can extract a confidence interval. Define, $\varphi(a) = \int_{-a}^a f(z)dz$ such that $z \sim \mathcal{N}(0, 1)$ and $f(z)$ is its density function. For example, a 90% confidence interval for the prediction $\mu(x)$ is defined by:

$$\{y \mid c|y - \mu(x)|/\sigma(x) \leq \varphi^{-1}(0.9)\}. \quad (2)$$

Since the variance $\sigma^2(x)$ is predicted by the neural network, it may not be well-calibrated and may be either underestimated or overestimated. The training loss function (1) assumes Gaussian distribution, which may not be correct so that the conversion from variance to confidence interval (2) can be wrong. Gaussian Variance Scaling (Gaussian-VS) (Laves et al., 2020; Levi et al., 2022; Frenkel and Goldberger, 2023) is a method to calibrate the variance $\sigma(x)$ that yields a meaningful confidence interval. This method is based on a Gaussian assumption of the conditional density $f(y|x)$. Gaussian-Vs computes a scalar r , which scales the variance predicted by the network: $\sigma(x) \rightarrow r \cdot \sigma(x)$. Given labeled validation data $(x_1, y_1), \dots, (x_n, y_n)$, we look for a scalar r that minimizes the loss:

$$L(r) = - \sum_{t=1}^n \log \mathcal{N}(y_t; \mu_\theta(x_t), r^2 \sigma_\theta^2(x_t)). \quad (3)$$

It is easy to verify that the optimal r is:

$$\hat{r}^2 = \frac{1}{n} \sum_{t=1}^n s_t^2 \quad \text{s.t.}, \quad s_t = \frac{|y_t - \mu(x_t)|}{\sigma(x_t)}. \quad (4)$$

Given a test image x , the calibrated confidence interval with coverage $1 - \alpha$ is:

$$[\mu(x) - \varphi^{-1}(1 - \alpha) \cdot \hat{r} \cdot \sigma(x), \mu(x) + \varphi^{-1}(1 - \alpha) \cdot \hat{r} \cdot \sigma(x)]. \quad (5)$$

Neural networks tend to be very robust to model mismatch and inaccurate ground truth measurements due to their high non-linearity and over-parameterization. In contrast, Gaussian-VS calibration is linear and consists of a single parameter. Thus if $f(y|x)$ is uni-modal but not Gaussian, while network training works well, Gaussian-VS fails to produce an accurate confidence interval.

The CQR calibration approach (Romano et al., 2019) consists of a non-parametric Quantile Regression (QR) (Koenker and Bassett Jr, 1978) followed by a conformalization step. Define the γ -quantile (pinball) loss:

$$L_\gamma(y, t) = 1_{\{t < y\}}(y - t)\gamma + 1_{\{t > y\}}(t - y)(1 - \gamma).$$

Given a training set $(x_1, y_1), \dots, (x_n, y_n)$ the QR algorithm trains a γ -quantile estimation $\hat{t}_\gamma(x)$ using the pinball loss:

$$L_\gamma(\theta) = \frac{1}{n} \sum_{i=1}^n L_\gamma(y_i, \hat{t}_\gamma(x_i, \theta)),$$

such that θ is the network parameter set. Given a miss-coverage rate α , we train a QR network with two heads $\hat{t}_{\alpha/2}(x)$ and $\hat{t}_{1-\alpha/2}(x)$ to obtain an instance-dependent confidence interval: $[\hat{t}_{\alpha/2}(x), \hat{t}_{1-\alpha/2}(x)]$. The CQR algorithm applies a CP procedure to ensure that the coverage is indeed $1 - \alpha$. Define the following conformal score:

$$s(x, y) = \max\{\hat{t}_{\alpha/2}(x) - y, y - \hat{t}_{1-\alpha/2}(x)\}.$$

The score $s(x, y)$ is the minimum number q such that $y \in [\hat{t}_{\alpha/2}(x) - q, \hat{t}_{1-\alpha/2}(x) + q]$. Let s_1, \dots, s_n be the conformal scores of a given validation set $(x_1, y_1), \dots, (x_n, y_n)$ and let \hat{q} be the $(1 - \alpha)$ quantile of s_1, \dots, s_n . The calibrated confidence interval is:

$$C_{\hat{q}}(x) = [\hat{t}_{\alpha/2}(x) - \hat{q}, \hat{t}_{1-\alpha/2}(x) + \hat{q}].$$

Unlike Gaussian-VS, the interval obtained by CPR has a coverage guarantee. The CP theory (Vovk et al., 2005) guarantees that:

$$1 - \alpha \leq p(y \in C_{\hat{q}}(x)) \leq 1 - \alpha + \frac{1}{n - 1}$$

where y is the (unknown) true label. QR is much more difficult to train than a Gaussian regression network and when QR produces poor interval estimations, the performance of CQR is also affected since it tries to cover the guaranteed validity by sacrificing efficiency (Chung et al., 2021; Kato et al., 2023).

3. CP-based Variance Scaling

In this section we present a method that combines the benefits of the two methods described above. We first train a parametric Gaussian network to predict the target and its variance and then apply a non-parametric CP to calibrate the estimated variance. Assume we trained a regression network that outputs mean $\hat{y} = \mu(x)$ and variance $\sigma^2(x)$ for each input image x as described in Section 2. Given a threshold $1 - \alpha$, we can apply CP to scale the variance and find a confidence interval around $\mu(x)$ in the form of:

$$C_q(x) = [\mu(x) - q \cdot \sigma(x), \mu(x) + q \cdot \sigma(x)] \quad (6)$$

such that the true value y is within this interval with probability $1 - \alpha$, i.e., $p(y \in C_q(x)) = 1 - \alpha$. The calibration parameter q is found in the following way. For each labeled data (x, y) define the conformal score:

$$s = |y - \mu(x)| / \sigma(x).$$

It is easy to verify that:

$$C_s(x) = [\mu(x) - |y - \mu(x)|, \mu(x) + |y - \mu(x)|]$$

and $y \in C_q(x)$ if and only if $q \geq s$. In other words, $C_s(x)$ is the minimal interval centered at $\mu(x)$ which contains the true value y . Let s_1, \dots, s_n be the conformal scores of the validation set $(x_1, y_1), \dots, (x_n, y_n)$ respectively. The calibration value q computed by the CP algorithm

Algorithm 1 Conformal-Prediction based Variance Scaling (CP-VS)

input: A labeled dataset divided into training and validation subsets and a confidence level $1 - \alpha$.

- Train a regression network $x \rightarrow (\mu_\theta(x_t), \sigma_\theta^2(x))$ by minimizing loss (1).
- Compute the conformal scores on the validation set:

$$s_t = |y_t - \mu(x_t)| / \sigma(x_t), \quad t = 1, \dots, n$$

- Sort the scores $s_1 \leq s_2, \dots, \leq s_n$ and set $q = s_{\lceil (n+1)(1-\alpha) \rceil}$.
 - The confidence interval of a new test point x is: $C(x) = [\mu(x) - q \cdot \sigma(x), \mu(x) + q \cdot \sigma(x)]$.
 - There is a coverage guarantee: $p(y \in C(x)) \geq 1 - \alpha$.
-

is the $\frac{\lceil (n+1)(1-\alpha) \rceil}{n}$ quantile of s_1, \dots, s_n . In other words, q is the minimal value for which the true value lies within the confidence interval defined by q in the $(1 - \alpha)$ portion of the validation set. The CP theory (Vovk et al., 2005) guarantees that regardless of the data distribution, for test data (x, y) , the value q found by the CP algorithm satisfies:

$$1 - \alpha \leq p(y \in C_q(x)) \leq 1 - \alpha + \frac{1}{n - 1} \quad (7)$$

such that n is the size of the validation set.

Note that both Gaussian-VS and CP-VS calibrate by scaling the estimated standard deviation $\sigma(x)$ using a scalar value that is learned from the same conformal scores s_1, \dots, s_n (4) obtained from the validation set. Gaussian-VS is based on the scores' average while CP-VS is based on a quantile of the scores. In case the conditional density of the target value given the input image is indeed Gaussian, the two calibration methods asymptotically coincide. In case the conditional density is not Gaussian, the quantile is a more effective measure than the mean when constructing a confidence interval. The CP-VS algorithm is summarized in Algorithm Box 1.

The CP-VS has several benefits. Firstly, network training is conducted using the robust Gaussian loss (unlike CQR which uses the pinball loss). Secondly, CP-VS achieves calibration through the CP procedure which has a parameter-free theoretical coverage guarantee (unlike Gaussian-VS which has neither a theoretical nor practical coverage guarantee).

4. Experimental Results

In this section, we empirically compare the performance of the confidence intervals computed by CP-VS to those computed by Gaussian-VS and CQR in terms of both interval length and coverage.

Datasets. We implemented the proposed calibration methods on several medical imaging regression tasks to evaluate their performance. The experimental setup follows the one used in (Laves et al., 2020) and includes the following medical datasets:

- BoneAge - Hand CT age regression from the RSNA pediatric bone age dataset (Halabi et al., 2019). The task here is to infer a person's age in months from CT scans of the hand. This dataset is the largest used in this study and has 12,811 images, from which we used 6811/2000/4000 images for training/validation/testing.

Table 1: Calibration results for $1 - \alpha = 0.90$ (top) and $1 - \alpha = 0.95$ (bottom), measured by average confidence interval length and coverage (%). The method which reports the minimal length among those satisfying the coverage requirement, is shown in bold.

$1 - \alpha = 0.9$		Gaussian-VS		CQR		CP-VS	
Dataset	Architecture	length ↓	coverage	length ↓	coverage	length ↓	coverage
BoneAge	DenseNet-201	0.184 ± 0.003	91.10 ± 0.71	0.411 ± 0.005	90.02 ± 0.78	0.176 ± 0.003	89.92 ± 0.85
	EfficientNet-B4	0.190 ± 0.003	90.19 ± 0.56	0.602 ± 0.007	90.02 ± 0.83	0.189 ± 0.003	89.95 ± 0.72
OCT	DenseNet-201	0.128 ± 0.001	99.94 ± 0.06	0.106 ± 0.002	90.11 ± 0.99	0.048 ± 0.001	90.47 ± 1.32
	EfficientNet-B4	0.127 ± 0.001	99.93 ± 0.06	0.157 ± 0.003	89.97 ± 1.50	0.050 ± 0.001	90.07 ± 1.80
DLS1	DenseNet201	0.185 ± 0.001	91.67 ± 0.42	0.278 ± 0.003	89.68 ± 1.01	0.180 ± 0.001	90.35 ± 0.69
	EfficientNet-B4	0.175 ± 0.001	89.34 ± 0.54	0.350 ± 0.004	90.37 ± 1.16	0.177 ± 0.001	90.18 ± 0.72
DLS2	DenseNet201	0.121 ± 0.001	87.44 ± 0.64	0.316 ± 0.002	90.36 ± 0.89	0.129 ± 0.001	90.08 ± 0.84
	EfficientNet-B4	0.143 ± 0.001	92.70 ± 0.60	0.299 ± 0.002	89.93 ± 0.78	0.132 ± 0.002	89.76 ± 1.09
DLS3	DenseNet201	0.150 ± 0.001	94.35 ± 0.42	0.178 ± 0.002	90.18 ± 0.90	0.134 ± 0.002	90.33 ± 0.87
	EfficientNet-B4	0.113 ± 0.001	90.63 ± 0.54	0.299 ± 0.002	89.96 ± 0.78	0.113 ± 0.002	90.43 ± 0.81
DLS4	DenseNet201	0.147 ± 0.001	88.41 ± 0.59	0.220 ± 0.002	89.92 ± 0.99	0.154 ± 0.002	90.36 ± 0.87
	EfficientNet-B4	0.150 ± 0.001	89.04 ± 0.62	0.328 ± 0.005	90.21 ± 0.93	0.154 ± 0.001	90.28 ± 0.86
DLS5	DenseNet201	0.225 ± 0.001	90.78 ± 0.59	0.227 ± 0.003	90.19 ± 0.89	0.223 ± 0.002	90.25 ± 0.88
	EfficientNet-B4	0.269 ± 0.001	96.83 ± 0.38	0.474 ± 0.011	89.75 ± 1.14	0.225 ± 0.002	90.27 ± 0.87

$1 - \alpha = 0.95$		Gaussian-VS		CQR		CP-VS	
Dataset	Architecture	length ↓	coverage	length ↓	coverage	length ↓	coverage
BoneAge	DenseNet-201	0.219 ± 0.004	95.40 ± 0.43	0.539 ± 0.011	94.97 ± 0.87	0.224 ± 0.004	94.91 ± 0.53
	EfficientNet-B4	0.227 ± 0.003	94.71 ± 0.43	0.489 ± 0.008	95.08 ± 0.58	0.233 ± 0.004	95.26 ± 0.47
OCT	DenseNet-201	0.153 ± 0.001	99.94 ± 0.06	0.146 ± 0.005	95.07 ± 1.20	0.057 ± 0.001	94.96 ± 0.87
	EfficientNet-B4	0.152 ± 0.002	99.98 ± 0.02	0.235 ± 0.006	95.12 ± 0.69	0.059 ± 0.001	94.99 ± 1.06
DLS1	DenseNet201	0.221 ± 0.001	97.27 ± 0.20	0.362 ± 0.004	95.09 ± 0.67	0.204 ± 0.003	95.17 ± 0.60
	EfficientNet-B4	0.208 ± 0.001	95.82 ± 0.33	0.425 ± 0.004	95.20 ± 0.65	0.202 ± 0.003	95.12 ± 0.67
DLS2	DenseNet201	0.144 ± 0.001	93.76 ± 0.59	0.352 ± 0.003	95.14 ± 0.71	0.151 ± 0.003	94.90 ± 0.84
	EfficientNet-B4	0.170 ± 0.001	96.98 ± 0.35	0.379 ± 0.003	94.88 ± 0.58	0.156 ± 0.002	94.96 ± 0.68
DLS3	DenseNet201	0.179 ± 0.001	97.96 ± 0.21	0.172 ± 0.003	95.17 ± 0.78	0.156 ± 0.003	95.37 ± 0.62
	EfficientNet-B4	0.135 ± 0.001	95.33 ± 0.45	0.432 ± 0.006	95.06 ± 0.47	0.136 ± 0.003	95.38 ± 0.70
DLS4	DenseNet201	0.175 ± 0.001	94.39 ± 0.45	0.251 ± 0.003	95.05 ± 0.72	0.180 ± 0.003	95.16 ± 0.71
	EfficientNet-B4	0.178 ± 0.001	95.06 ± 0.46	0.363 ± 0.008	94.97 ± 0.63	0.179 ± 0.003	95.16 ± 0.69
DLS5	DenseNet201	0.269 ± 0.001	96.93 ± 0.34	0.283 ± 0.003	95.21 ± 0.59	0.252 ± 0.003	95.20 ± 0.67
	EfficientNet-B4	0.320 ± 0.001	99.32 ± 0.14	0.552 ± 0.012	94.70 ± 0.64	0.253 ± 0.003	95.10 ± 0.65

- OCT - Six degrees of freedom (6DoF) needle pose estimation on optical coherence tomography (OCT). This dataset contains 5,000 3D-OCT scans with the accompanying needle pose $y \in [0, 1]^6$, from which we use 3300/850/850 for training/validation/testing (Laves et al., 2020).
- DLS - This dataset is designed to facilitate the detection and classification of degenerative lumbar spine (DLS) conditions using MRI images. Each image includes annotations for the (x, y) positions of five vertebrae. For each vertebra, the dataset is divided into 60% for training, 20% for validation, and 20% for testing, with approximately 10,000 images per vertebra (Richards et al., 2024).




Figure 1: Histograms of the normalized network prediction values computed on the validation sets using DenseNet-201.

Implementation details. The network architectures used were EfficientNet-B4 (Tan and Le, 2019) and DenseNet-201 (He et al., 2016; Guo et al., 2017), which are the state-of-the-art for deep models. The last linear layer of all networks was replaced by two linear layers predicting the mean and log-variance. The networks were trained until no further decrease of the loss on the validation set could be observed. To implement the CQR algorithm, we used the code from the CQR project GitHub¹. CQR is trained to minimize the average length while keeping the coverage valid. For reproducibility, we have made Gaussian-VS and CP-VS code available².

Evaluation measures. The standard direct way to evaluate the performance of confidence-interval estimators on a given test set $(x_1, y_1), \dots, (x_n, y_n)$ is by computing the degree of coverage and the average length of the prediction intervals. Smaller average interval widths indicate higher precision. The length and coverage are formally defined as:

$$\text{length} = \frac{1}{n} \sum_i |C(x_i)|, \quad \text{coverage} = \frac{1}{n} \sum_i \mathbf{1}(y_i \in C(x_i))$$

such that $C(x_i)$ is the confidence interval of x_i and $|C(x_i)|$ is its length. The best algorithm is the one that reports the minimal average interval length among those that satisfy the coverage requirement.

Results. Table 1 shows the comparative calibration results (length and coverage) for the three methods (Gaussian-VS, CQR and CP-VS) on the test set. The results were averaged over 20 random splits of the data into validation and test sets. Both CP-VS and CQR, which apply a CP procedure, obtained the exact required coverage, as guaranteed by the CP theorem. However, the average length reported by the CQR was much larger due to the non-robust training of the QR algorithm. The coverage rate of Gaussian-VS was inconsistent. In some cases, it was below the required coverage $(1 - \alpha)$; in other cases, it was above (resulting in a large average interval). Note that Gaussian-VS has no theoretical coverage guarantee.

Visual examples. We next illustrate the proposed method on several examples. Figure 2 shows examples from the BoneAge dataset. The figures illustrate the same trend that

1. <https://github.com/yromano/cqr>

2. <https://anonymous.4open.science/r/Calibrated-Instance-Dependent-95F9>

was reported in Tables 1 and 2, namely that the CP-VS yields confidence intervals with the smallest size.

Normality check. We next analyzed whether the output distribution of the Gaussian network was indeed Gaussian for each dataset. For each image x in the validation set, we computed the scalar $(y - \mu(x))/\sigma(x)$ such that y was the correct value and $\mu(x)$ and $\sigma(x)$ were predicted by the network. Note that $y \sim \mathcal{N}(\mu(x), \sigma^2(x))$ implies that $(y - \mu(x))/\sigma(x) \sim \mathcal{N}(0, 1)$. The histograms of the three datasets (BoneAge, OCT and DLS) are shown in Fig. 1. We also applied the Kolmogorov–Smirnov test to check the data’s normality and obtained 0.018 (BoneAge), 1.42e-60 (OCT), and 0.001 (DLS1). Hence, the BoneAge task histogram was the only one that resembled a Normal distribution and passed the Gaussianity test. We can see in Table 1 and Table 2 that when the normal assumption is valid, Gaussian-VS works well and produces an effective confidence interval. However, when the normal assumption fails Gaussian-VS has inconsistent behaviors. In some cases it doesn’t satisfy the coverage requirement and in other cases it yields large confidence intervals.




Figure 2: Samples from the BoneAge test set. (a) Target is 0.47, intervals are - Gaussian-VS: [0.37,0.57], CP-VS: [0.38,0.56], CQR: [0.20,0.83]. (b) Target is 0.76, intervals are - Gaussian-VS: [0.67,0.79], CP-VS: [0.68,0.78], CQR: [0.23,0.81]. (c) Target is 0.21, intervals are - Gaussian-VS: [0.17,0.31], CP-VS: [0.17,0.31], CQR: [0.24,0.90].

5. Conclusions

To conclude, in this study, we focused on reporting a reliable prediction uncertainty of a regression network in the form of a confidence interval that varies in size across the images and reflects instance-specific uncertainty. We showed that Gaussian-VS cannot be used in medical applications since its confidence interval has no statistical meaning. We also showed that CQR achieves poor results in terms of average interval length. We demonstrated that CP-VS, which is based on training a network using a Gaussian loss followed by a distribution-free calibration based on CP, achieved the best results. We focused here on medical imaging applications but the conclusions are general and relevant for calibration of any regression network.

References

- Anastasios N. Angelopoulos and Stephen Bates. Conformal prediction: A gentle introduction. *Foundations and Trends in Machine Learning*, 16(4):494–591, 2023.
- Netanell Avisdris, Leo Joskowicz, Brian Dromey, Anna L David, Donald M Peebles, Danail Stoyanov, Dafna Ben Bashat, and Sophia Bano. Biometrynet: Landmark-based fetal biometry estimation from standard ultrasound planes. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2022.
- Ashwin Carvalho, Stéphanie Lefèvre, Georg Schildbach, Jason Kong, and Francesco Borrelli. Automated driving: The role of forecasts and uncertainty—a control perspective. *European Journal of Control*, 24:14–32, 2015.
- Youngseog Chung, Willie Neiswanger, Ian Char, and Jeff Schneider. Beyond pinball loss: Quantile methods for calibrated uncertainty quantification. *Advances in Neural Information Processing Systems (NeurIPs)*, 2021.
- Lior Frenkel and Jacob Goldberger. Calibration of a regression network based on the predictive variance with applications to medical images. In *The IEEE International Symposium on Biomedical Imaging (ISBI)*, 2023.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International Conference on Machine Learning (ICML)*, 2017.
- Safwan S Halabi et al. The RSNA pediatric bone age machine learning challenge. *Radiology*, 290(2):498, 2019.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Yuko Kato, David MJ Tax, and Marco Loog. A review of nonconformity measures for conformal prediction in regression. *Conformal and Probabilistic Prediction with Applications*, pages 369–383, 2023.
- Roger Koenker and Gilbert Bassett Jr. Regression quantiles. *Econometrica: journal of the Econometric Society*, pages 33–50, 1978.
- Fabian Küppers, Jonas Schneider, and Anselm Haselhoff. Parametric and multivariate uncertainty calibration for regression and object detection. *arXiv preprint arXiv:2207.01242*, 2022.
- Max-Heinrich Laves, Sontje Ihler, Jacob F Fast, Lüder A Kahrs, and Tobias Ortmaier. Well-calibrated regression uncertainty in medical imaging with deep learning. In *Medical Imaging with Deep Learning*, 2020.
- Christian Leibig, Vaneeda Allken, Murat Seçkin Ayhan, Philipp Berens, and Siegfried Wahl. Leveraging uncertainty information from deep neural networks for disease detection. *Scientific Reports*, 7(1):1–14, 2017.

- Dan Levi, Liran Gispan, Niv Giladi, and Ethan Fetaya. Evaluating and calibrating uncertainty prediction in regression tasks. *Sensors*, 22(15):5540, 2022.
- Tyler Richards, Jason Talbott, Robyn Ball, Errol Colak, Adam Flanders, Felipe Kitamura, John Mongan, Luciano Prevedello, and Maryam Vazirabad. Rsna 2024 lumbar spine degenerative classification. <https://kaggle.com/competitions/rsna-2024-lumbar-spine-degenerative-classification>, 2024. Kaggle.
- Yaniv Romano, Evan Patterson, and Emmanuel Candes. Conformalized quantile regression. *Advances in Neural Information Processing Systems (NeurIPs)*, 2019.
- Sebastian Scher and Gabriele Messori. Predicting weather forecast uncertainty with machine learning. *Quarterly Journal of the Royal Meteorological Society*, 144(717):2830–2841, 2018.
- Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning (ICML)*, 2019.
- Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic learning in a random world*, volume 29. Springer, 2005.