QUTCC : QUANTILE UNCERTAINTY TRAINING AND CONFORMAL CALIBRATION FOR IMAGING INVERSE PROBLEMS

Anonymous authorsPaper under double-blind review

000

001

003

004 005 006

012 013

014

016

018

021

025

026

027

028

029

031

032

034

035

037

040

041

042

043

044

046

047

048

051

052

ABSTRACT

Deep learning models often hallucinate, producing realistic artifacts that are not truly present in the sample. This can have dire consequences for scientific and medical inverse problems, where accuracy is more important than perceptual quality. Uncertainty quantification techniques, such as conformal prediction, can pinpoint outliers and provide guarantees for image regression tasks, improving reliability. However, existing methods predict fixed quantiles and utilize a linear constant scaling factor to calibrate uncertainty bounds, resulting in larger, less informative bounds. We propose *QUTCC*, a quantile uncertainty training and calibration technique that enables nonlinear, non-uniform scaling of quantile predictions to enable tighter uncertainty estimates. Using a U-Net architecture with a quantile embedding, OUTCC can predict the full conditional distribution of quantiles for each image. After conformal calibration, QUTCC can predict pixel-wise uncertainty intervals that satisfy coverage guarantees and also estimate a pixelwise conditional probability density function. We evaluate our method on image denoising, quantitative phase imaging, and compressive MRI reconstruction. Our method successfully pinpoints hallucinations in image estimates and consistently achieves tighter uncertainty intervals than prior methods while maintaining the same statistical coverage.

1 Introduction

In recent years, deep learning models have become the dominant approach across many inverse problems, favored for their ability to learn powerful and complex priors from an abundance of data (Ongie et al., 2020; Alshardan et al., 2024; Barbastathis et al., 2019; Xue et al., 2019). However, these models are generally limited in their ability to represent uncertainty in their predictions, which has been a significant barrier to their use in scientific and medical applications, where identifying observations as out-of-distribution is crucial and the consequences of model hallucination can be severe (Begoli et al., 2019). While significant progress has been made in estimating uncertainty in deep learning, many methods require substantial computational demands, incorporate strong prior data assumptions, or may not provide formal statistical guarantees (Psaros et al., 2023; Gal & Ghahramani, 2016; Gal et al., 2017; Lakshminarayanan et al., 2017; Sun et al., 2024). These limitations have motivated research into conformal prediction methods for uncertainty quantification, which aim to overcome these challenges (Angelopoulos & Bates, 2021; Angelopoulos et al., 2021; Dewolf et al., 2023; Barber et al., 2023; Romano et al., 2019).

Specifically, recent methods apply the statistical rigor of conformal prediction to inverse problems and image regression tasks through multi-dimensional conformalized quantile regression (Angelopoulos et al., 2022; Romano et al., 2019; Koenker & Bassett Jr, 1978). This is typically done by modifying a neural network to output two additional predictions: a lower-bound and an upper-bound image, which together define a confidence interval for each pixel. A calibration step is then used to scale these bounds so that the interval captures the true intensity values with the desired coverage, resulting in pixel-wise intervals that are statistically guaranteed to contain the true values with a user-specified confidence level (Angelopoulos et al., 2022). Multi-dimensional conformalized quantile regression is simple to implement and computationally inexpensive, unlike ensemble methods, which often require

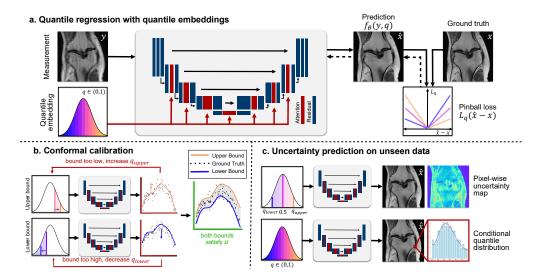


Figure 1: QUTCC Overview. a) Quantile Regression with Quantile Embedding: During training, a neural network with a quantile embedding predicts an image as a function of the measurement and quantile, q. The quantile embedding is randomly sampled ($q \in (0,1)$) and the value of q determines the asymmetry of the pinball loss, enabling the model to learn a range of conditional quantiles. b) Conformal Calibration: During calibration, the predictive bounds (q_{lower}, q_{upper}) are iteratively adjusted on a held-out dataset to satisfy the desired miscoverage level α . c) Uncertainty Prediction on Unseen Data: At test time, the model can be queried to predict the mean image, a pixel-wise uncertainty map, or a conditional probability density function at each pixel.

training multiple models, while still providing formal coverage guarantees that are independent of the model choice or data distribution. This makes it particularly well-suited for scientific and medical imaging, which often has small datasets, and adaptive imaging, where inference speed is crucial (Ye et al., 2025). However, prior methods use a linear pixel-wise scaling during conformal calibration to approximate a complex, nonlinear distribution, which leads to overly loose prediction intervals (Angelopoulos et al., 2022). Furthermore, existing methods learn fixed quantiles, which carry the risk of quantile crossing (He, 1997).

To address these challenges, we present **Quantile Uncertainty Training and Conformal Calibration** (**QUTCC**), a novel method for simultaneous quantile prediction and conformal calibration that enables efficient and accurate uncertainty quantification for imaging inverse problems (Fig 1). Building on past work in multi-quantile estimation (Rodrigues & Pereira, 2020; Tagasovska & Lopez-Paz, 2019), QUTCC uses a single neural network to estimate a distribution of quantiles. During the conformal calibration step, QUTCC applies a non-uniform, nonlinear scaling to the uncertainty bounds, compared to constant scaling used by prior methods. This results in smaller and potentially more informative uncertainty intervals. Additionally, because all quantiles are learned during training, QUTCC can query the full quantile range at inference time to construct a pixel-wise estimate of the underlying probability distribution. We evaluate QUTCC on five imaging inverse problems and show that QUTCC can highlight regions of hallucination and maintain smaller uncertainty interval sizes than prior methods, while satisfying the same statistical guarantees. Additionally, we show that QUTCC effectively approximates pixel-wise probability density functions, providing a richer representation of uncertainty that could enable more informed downstream decision-making in scientific and medical imaging applications.

2 Related Work

2.1 CONFORMAL PREDICTION

Conformal prediction is a robust statistical technique that constructs predictive intervals with a formal guarantee of statistical coverage (Angelopoulos & Bates, 2021; Correia et al., 2024; Romano et al.,

109

110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126 127

128 129

130

131

132

133

134

135

136

137

138

139

140

141

142

143

144

145

146

147 148

149 150

151

152

153

154

155

156

157

158

159

160

161

2019). In general, conformal methods take heuristic or model-derived uncertainty estimates and refine them into statistically valid predictive intervals by leveraging a held-out calibration dataset. These intervals are constructed to satisfy a finite-sample marginal coverage guarantee at a user-specified level $(1-\alpha)$, meaning that, with high probability, the true response lies within the predicted interval at least $(1-\alpha)\%$ of the time. For example, setting $\alpha=0.1$ yields prediction intervals that are guaranteed to contain the ground truth in at least 90% of test cases. Conformal prediction has become increasingly popular due to its simplicity, speed, formal guarantees, and lack of assumptions on model or data distribution (Angelopoulos et al., 2023). It has been applied to a variety of areas, including classification Angelopoulos et al. (2020); Ding et al. (2023); Angelopoulos et al. (2021), language modeling (Quach et al., 2023; Campos et al., 2024), robotics (Lindemann et al., 2023; Lekeufack et al., 2024), protein design (Fannjiang et al., 2022), and time series estimation (Xu & Xie, 2023; Zaffran et al., 2022). Recently, conformal prediction has been proposed for image-to-image regression tasks to enable pixel-wise uncertainty prediction (Angelopoulos et al., 2022), and thereafter applied to adaptive-microscopy (Ye et al., 2025). These methods alter a neural network to predict fixed upper and lower quantile estimates, in addition to the mean image, then uniformly scale these predictions by a constant factor using conformal risk control to ensure valid statistical coverage at test time (Angelopoulos et al., 2022). Instead of predicting fixed quantiles and uniformly scaling them during conformal calibration, we use one network to predict all possible conditional quantiles. This enables non-uniform, asymmetric scaling during conformal calibration and leads to tighter uncertainty intervals for image regression tasks and imaging inverse problems.

2.2 QUANTILE REGRESSION

Quantile regression is a general approach to estimate the conditional quantiles of a target distribution rather than the mean of a response variable (Koenker & Bassett Jr, 1978; Koenker & Hallock, 2001). This is often accomplished by leveraging an asymmetric loss function, called pinball loss (Fig. 1a, Eq. 3), tailored to the specified quantile level (Steinwart & Christmann, 2011). The estimated intervals obtained by quantile regression do not have formal guarantees on their own, but can be paired with conformal prediction to obtain coverage guarantees (Romano et al., 2019). Learning quantiles during neural network training can improve predictive performance by introducing a regularization effect, while enabling uncertainty estimation (Rodrigues & Pereira, 2020). One limitation of quantile-based methods is the tendency for 'quantile crossing' to occur, a phenomenon wherein quantiles trained independently violate their natural ordering, resulting in lower quantiles exceeding higher ones (Das et al., 2019). Training a single network to simultaneously predict multiple quantiles, called 'simultaneous quantile prediction' can mitigate this issue, while also enabling the estimation of the entire conditional distribution (Sangnier et al., 2016; Liu & Wu, 2011; Tagasovska & Lopez-Paz, 2019; Rodrigues & Pereira, 2020). In our work, we leverage a single-network with shared parameters for simultaneous quantile prediction. We embed the quantile level as an explicit input parameter into a U-Net, which is well-suited for a variety of image regression and imaging inverse tasks. Furthermore, we pair our network with conformal prediction to achieve coverage guarantees. We show that our network mitigates the issue of quantile crossing while also maintaining overall prediction accuracy. We are the first to demonstrate that a single network trained for simultaneous quantile prediction can predict conformally calibrated uncertainty intervals for imaging inverse problems.

2.3 Predicting Tighter or More Interpretable Bounds

Achieving smaller interval lengths without sacrificing coverage guarantees reflects increased confidence in the model's predictions, leading to more precise and reliable uncertainty quantification. Producing tighter bounds is a common objective across uncertainty estimation methods, not just conformal prediction (Xie et al., 2024). Several approaches have been proposed to enhance conformal prediction by targeting user-specified properties such as reduced interval length or improved conditional coverage (Xie et al., 2024; Chung et al., 2021); however, to date, none of these techniques have been applied to imaging tasks. On the other hand, several methods aim to improve the interpretability of uncertainty prediction for imaging tasks by moving away from per-pixel uncertainty estimates. These methods leverage principal components, posterior projected distribution, and spatial/topological relationships (Nehme et al., 2023; Yair et al., 2024; Belhasin et al., 2023; Gupta et al., 2023) to predict uncertainty in a more interpretable way. However, without incorporating conformal prediction, these methods lack statistical guarantees. Several methods pair inverse problems with downstream tasks, such as classification, to estimate the uncertainty in a more interpretable way (Cheung et al., 2024;

Wen et al., 2024), and others represent uncertainty in a semantically-meaningful latent space (Sankara-narayanan et al., 2022). While these methods are promising, they are less general and often tied to a specific application. We present a more general method that can predict uncertainty for any imaging inverse problem while achieving smaller uncertainty interval lengths than previous image-to-image regression methods.

3 METHODOLOGY

3.1 PROBLEM OVERVIEW

Given an unseen measurement image \mathbf{y} , our goal is to train a network f_{θ} that can be queried at any quantile level $q \in (0,1)$ to produce an estimate $\hat{\mathbf{x}} = f_{\theta}(\mathbf{y},q)^1$. The image estimate is obtained by querying the network at q = 0.5, which gives us the median image:

$$\hat{\mathbf{x}} = f_{\theta}(\mathbf{y}, q = 0.5). \tag{1}$$

The uncertainty interval is obtained by querying the network at a high quantile value, q_{upper} , and a low quantile value, q_{lower} . The upper quantile prediction serves as the upper bound of the uncertainty interval, and the lower quantile serves as the lower bound. These two quantities can be subtracted to get a notion of pixel-wise uncertainty: $\hat{\mathbf{u}} = f_{\theta}(\mathbf{y}, q_{\text{upper}}) - f_{\theta}(\mathbf{y}, q_{\text{lower}})$.

During conformal calibration, we use a small calibration dataset, $\mathcal{D}_c = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{N_c}$, to search over different values of the input parameters q_{upper} and q_{lower} until we reach the desired coverage. That is, the constructed interval, $\mathcal{C} \in [f_{\theta}(\mathbf{y}, q_{lower}), f_{\theta}(\mathbf{y}, q_{upper})]$, contains at least $1 - \alpha$ of the ground truth pixels with probability $1 - \delta$ (Angelopoulos et al., 2022). Formally, this means that with probability of at least $1 - \delta$:

$$E[\mathbf{x}_{\text{test}} \in C(\mathbf{y}_{\text{test}})] \ge 1 - \alpha \tag{2}$$

where \mathbf{x}_{test} , \mathbf{y}_{test} are unseen test data from the same distribution as the calibration data. After calibration, the network can provide a pixel-wise uncertainty map for new, unseen data or can be queried to predict a conditional distribution for any pixel in the image. Our method QUTCC, is summarized in Fig. 1. We elaborate on our network architecture, training procedure, and the conformal calibration step below.

3.2 Network Architecture

Our proposed network, $f_{\theta}(\mathbf{y},q)$, takes the parameter q as an input to predict the conditional quantile. To do this in practice, we propose to use an attention U-Net (Ronneberger et al., 2015; Oktay et al., 2018) with a quantile-embedding to condition the U-Net on a given q. This is inspired by the architecture of U-Nets used for diffusion models, which include a time-embedding (Ho et al., 2020; Zhang et al., 2023). Specifically, we encode randomly sampled quantiles ($q \in (0,1)$) as high-dimensional vectors during training, allowing the network to learn the distribution of the target variable for each specified quantile. This allows us to query an image prediction at any quantile during a forward pass. The neural network weights are shared across different quantile predictions, limiting quantile crossing. The proposed network is shown in Fig. 1, and the full architecture and training details are described in the Supplement.

3.3 SIMULTANEOUS QUANTILE REGRESSION

In order to train our neural network to predict an arbitrary quantile image, we use pinball loss (L_q) , an asymmetric loss function commonly used in quantile regression (Eq. 3), where \hat{x} denotes the predicted value, x represents the ground truth, and $q \in (0,1)$ is the quantile of interest:

$$L_q(x,\hat{x}) = \begin{cases} q \cdot |x - \hat{x}| & \text{if } x - \hat{x} \ge 0\\ (1 - q) \cdot |x - \hat{x}| & \text{otherwise.} \end{cases}$$
 (3)

When $q \neq 0.5$, the pinball loss introduces asymmetry by penalizing overestimates and underestimates unequally. Specifically, when q > 0.5, the loss assigns a greater penalty to underestimations

¹Please note that we adopt the notation $\hat{x} = f(y)$, which is commonly used in the field of inverse problems instead of $\hat{y} = f(x)$, which is more common in the machine learning literature.

(i.e., when $\hat{x} < x$), encouraging the model to predict higher values. Conversely, for q < 0.5, overestimations (i.e., when $\hat{x} > x$) incur a larger penalty, biasing predictions downward. This asymmetry enables the model to learn conditional quantiles of the target distribution, in contrast to losses like mean squared error (MSE), which are symmetric and designed to estimate the conditional mean. At each training step, the quantile parameter, q is randomly sampled and used to both condition the network and as an input to the loss function. This allows the model to learn the full conditional quantile function, rather than a discrete, fixed quantile value as in prior image-to-image regression methods (Angelopoulos et al., 2022; Ye et al., 2025). The total loss is given by:

$$\mathcal{L}_{\text{total}}(\theta) = \sum_{i=1}^{N_t} \mathcal{L}_q(x_i, f_{\theta}(y_i, q)), \tag{4}$$

where f_{θ} is a neural network with parameters θ , and $f_{\theta}(y_i,q)$ is the output of the neural network given an input measurement, y_i and quantile value q. This loss is minimized with the Adam optimizer (Kingma & Ba, 2014) using backpropagation.

3.4 CONFORMAL CALIBRATION

After training, the neural network can be queried to obtain the uncertainty interval predictions. However, these predictions may not be valid. To ensure the statistical coverage in Eq. 2, a conformal calibration step is necessary. Following the procedure in (Angelopoulos et al., 2022), we use a small, separate calibration dataset, $\mathcal{D}_c = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{N_c}$ to adjust q_{lower} and q_{upper} until the desired coverage, $1-\alpha$, is reached. Our model's risk over the calibration dataset is the number of miscovered pixels for each image normalized by the total number of pixels, K, and the size of the calibration dataset, N_c :

$$\hat{R}(q_{\text{lower}}, q_{\text{upper}}) = \frac{1}{N_c} \sum_{i=1}^{N} \frac{1}{K} \sum_{k=1}^{K} \mathbb{1}\{\mathbf{x}_i(k) \notin \mathcal{C}(\mathbf{y}_i(k), q_{\text{lower}}, q_{\text{upper}})\}, \tag{5}$$

where k is the index of the pixels in the image. This measures the miscoverage as a function of q_{lower} and q_{upper} . We can decompose this total risk as a function of the miscoverage from the lower quantile (the number of ground truth pixels that are lower than the lower quantile) and the miscoverage from the upper bound (the number of ground truth pixels that are higher than the upper quantile):

$$\hat{R}(q_{\text{lower}}, q_{\text{upper}}) = \frac{1}{N_c} \sum_{i=1}^{N} \frac{1}{K} \sum_{k=1}^{K} \left[\mathbb{1} \left\{ \mathbf{x}_i(k) < f_{\theta}(\mathbf{y}_i, q_{\text{lower}}) \right\} + \mathbb{1} \left\{ \mathbf{x}_i(k) > f_{\theta}(\mathbf{y}_i, q_{\text{upper}}) \right\} \right]. \quad (6)$$

During conformal calibration, we calibrate the lower and upper quantile bounds independently. To satisfy a target total miscoverage rate of α , the calibration process allocates half of this error budget to each bound. That is, the lower and upper bounds are each adjusted to capture violations at a rate no greater than $\alpha/2$. Thus, if $\mathbf{x}(k)$ denotes the ground truth at pixel k and $[\hat{\mathbf{x}}_{lower}(k), \hat{\mathbf{x}}_{upper}(k)]$ the predicted interval, the goal is to ensure:

$$\mathbb{P}(\mathbf{x}(k) < \hat{\mathbf{x}}_{lower}(k)) \le \frac{\alpha}{2} \quad \text{and} \quad \mathbb{P}(\mathbf{x}(k) > \hat{\mathbf{x}}_{upper}(k)) \le \frac{\alpha}{2},$$
 (7)

which implies:

$$\mathbb{P}(\hat{\mathbf{x}}_{\text{lower}}(k) \le \mathbf{x} \le \hat{\mathbf{x}}_{\text{upper}}(k)) \le \alpha. \tag{8}$$

By allowing the quantile bounds to vary independently and adaptively during calibration, our method supports non-uniform scaling as a function of image characteristics. Rather than scaling the quantile predictions by a constant, linear factor, as proposed in prior work (Angelopoulos et al., 2022), our quantile predictions can be scaled in a non-uniform, non-linear way as a function of the neural network, $f_{\theta}(\mathbf{y}, q_{\text{lower}}, q_{\text{upper}})$. In practice, this can result in smaller uncertainty intervals.

Pseudocode for this calibration process is provided in **Algorithm 1**. At each step, we compute the miscoverage from the quantile upper and lower bounds over the entire calibration dataset. If the violation rate for a bound exceeds the adjusted α , we relax the corresponding bound; otherwise, we tighten it. This process proceeds via a binary search over the quantile space until the desired coverage is reached. Conducting a binary search over the quantile space assumes that the learned quantile function is monotonic, which we verify in Supplement Tbl. 3. Note that we adjust the error rate, $\alpha' \leftarrow \alpha - \frac{1-\alpha}{N_c}$ to account for the finite calibration dataset size (Angelopoulos & Bates, 2021; Vovk, 2012). At the end of this procedure, we can obtain a constructed interval $\mathcal{C} \in [f_{\theta}(\mathbf{y}, q_{\text{lower}}^*), f_{\theta}(\mathbf{y}, q_{\text{upper}}^*)]$ that satisfies Eq. 2.

Algorithm 1 Calibrating Quantile Bounds q_{lower} , q_{upper}

- 1: Compute adjusted error: $\alpha' \leftarrow \alpha \frac{1-\alpha}{N}$
- 2: Initialize bounds: $q_{\text{lower}} \leftarrow \alpha', \quad q_{\text{upper}} \leftarrow 1 \alpha'$ 3: Define step size Δq for bound updates
- 4: while $R(q_{\text{lower}}, q_{\text{upper}}) > \alpha'$ do

- $\begin{array}{l} (r_{\mathrm{lower}}, r_{\mathrm{upper}}) \leftarrow (R_{\mathrm{lower}}(q_{\mathrm{lower}}), \ R_{\mathrm{upper}}(q_{\mathrm{upper}})) \\ \mathrm{Update} \ q_{\mathrm{lower}} \leftarrow q_{\mathrm{lower}} \pm \Delta q \\ \mathrm{Update} \ q_{\mathrm{upper}} \leftarrow q_{\mathrm{upper}} \pm \Delta q \end{array}$ $\begin{array}{l} \rhd -\Delta q \text{ if } r_{\text{lower}} \leq \alpha'/2, \text{ else } +\Delta q \\ \rhd +\Delta q \text{ if } r_{\text{upper}} \leq \alpha'/2, \text{ else } -\Delta q \end{array}$
- 8: **return** $(q_{lower}^*, q_{upper}^*)$

270

271

272

273 274

275 276

277 278

279

281

282 283

284

285

287

288

289

290

291

292

293

295 296

297

298 299

300 301

303 304

305

306

307

308

310

311 312

313

314 315

316

317 318

319 320 321

322

323

ESTIMATING THE CONDITIONAL DISTRIBUTION

Since our network is trained to predict the full conditional quantile distribution rather than a single fixed quantile, we can recover an estimate of the entire conditional quantile function, $Q_k(q)$, at each pixel k. This is accomplished by querying the network over a range of quantile levels $q \in (0,1)$. Recall that the quantile function $\hat{Q}_k(q)$ is the inverse of the cumulative distribution function for the conditional distribution at pixel k.

By approximating the derivative of the quantile function with respect to q, we can obtain the rate of change of the predicted values across quantile levels. Since the quantile function is the inverse of the cumulative distribution function, the reciprocal of this derivative provides an estimate of the conditional probability density function (PDF) at each pixel. This derivative is numerically approximated using finite differences on the discrete set of quantile predictions:

$$\frac{\partial \hat{\mathbf{x}}}{\partial q}\Big|_{q_i} \approx \nabla_q \hat{\mathbf{x}}\Big|_{q_i} = \frac{\hat{\mathbf{x}}^{(q_{i+1})} - \hat{\mathbf{x}}^{(q_{i-1})}}{q_{i+1} - q_{i-1}} \quad \text{for } i = 2, \dots, n-1, \tag{9}$$

where $\hat{\mathbf{x}}^{(q_i)}$ denotes the image prediction at quantile level q_i . The estimated PDF at each pixel is then given by the reciprocal of this derivative $\hat{p} = \left(\frac{\partial \hat{\mathbf{x}}}{\partial q}\right)^{-1}$. The network is first queried at all desired quantile levels to obtain the conditional quantile function estimates. Next, the derivative of this function is approximated using finite differences, and finally, the pixel-wise PDF is computed as the inverse of the derivative. This approach not only enables accurate pixel-wise PDF estimation but, when combined with conformal calibration of multiple bounds, also provides statistically guaranteed coverage of the estimated density.

To 'conformalize' the pixel-wise PDF estimation, we first specify the desired coverage level (e.g., for 90% coverage, we initialize with quantiles 0.05 and 0.95; for 60%, with 0.20 and 0.80). These initial bounds are then calibrated using Algorithm 1 to guarantee the target coverage. By performing this calibration for multiple quantile levels $\{q_i\}_{i=1}^n$ and their associated coverage rates $\{1-\alpha_i\}_{i=1}^n$, we obtain a set of conformally corrected quantiles,

$$\{\hat{Q}_k^{\text{conf}}(q_i)\}_{i=1}^n,$$
 (10)

from which the pixel-wise PDF is reconstructed via finite differences:

$$\hat{p}_k^{\text{conf}}(\hat{Q}_k^{\text{conf}}(q_i)) = \left(\frac{\hat{Q}_k^{\text{conf}}(q_{i+1}) - \hat{Q}_k^{\text{conf}}(q_{i-1})}{q_{i+1} - q_{i-1}}\right)^{-1}.$$
(11)

By calibrating multiple quantile levels with their corresponding coverage rates, we construct a PDF that provides reliable and conformally calibrated estimates.

RESULTS

To evaluate our proposed approach, we fit and calibrate a separate model f_{θ} to five separate imaging inverse problems: accelerated MRI (Zbontar et al., 2018), quantitative phase imaging (QPI) (Pinkard et al., 2024), and denoising under real-noise, synthetic Poisson, and Gaussian noise (Zhang et al., 2019). We compare against Im2Im-UQ, which is the leading conformal prediction approach for

image-to-image regression. To ensure that our performance improvements come from our uncertainty quantification technique and not network improvements, we upgrade Im2Im-UQ to use the same architecture and depth as QUTCC, which we call Im2Im-Deep. In all cases, we set $\alpha=0.1$ during conformal calibration. Full training details are provided in the supplement. We evaluate predicted interval lengths, empirical risk, and model performance. We also visualize uncertainty and highlight the model's ability to identify hallucinations, which are realistic features not present in the ground truth. Finally, we show that QUTCC infers pixel-wise PDFs without relying on prior distributional assumptions.

4.1 Uncertainty interval length and risk

We compare the predicted uncertainty interval lengths and total risk for each imaging modality in Fig. 2 and Table 1. QUTCC consistently produces smaller prediction intervals than Im2Im-Deep across all five modalities, while keeping the total risk under 0.1. By achieving smaller interval lengths while exhibiting comparable risk, QUTCC demonstrates that its uncertainty quantification is both more precise and well-calibrated without sacrificing coverage.

In Figure 3, we compare QUTCC and Im2Im-Deep for Gaussian denoising with progressively increasing noise, plotting the average uncertainty interval lengths as a function of pixel intensity. When the noise variance is lower, QUTCC produces smaller interval lengths that better match the noise level compared to Im2Im-Deep. In addition, QUTCC consistently estimates narrower uncertainty intervals in regions of high signal intensity.

Metric	Method	MRI	QPI	Gaussian	Poisson	Real Noise
Interval Length	Im2Im-Deep QUTCC	0.109 ± 0.056 0.108 ± 0.057	0.065 ± 0.014 0.062 ± 0.014	0.063 ± 0.049 0.059 ± 0.048	0.047 ± 0.038 0.040 ± 0.029	0.038 ± 0.045 0.036 ± 0.035
Total-Risk	Im2Im-Deep QUTCC	$\begin{array}{c} 0.099 \pm 0.047 \\ 0.097 \pm 0.031 \end{array}$	0.098 ± 0.099 0.100 ± 0.102	0.094 ± 0.065 0.090 ± 0.046	0.049 ± 0.042 0.093 ± 0.091	0.096 ± 0.040 0.098 ± 0.029

Table 1: Interval Length and Total Risk

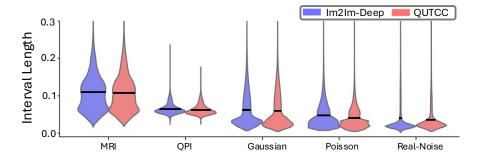


Figure 2: **QUTCC exhibits smaller uncertainty interval sizes:** We compare the predictive interval sizes of Im2Im-Deep and QUTCC across all five inverse problems. QUTCC consistently produces narrower uncertainty intervals. The black bolded line indicates the mean interval length.

4.2 Uncertainty Visualizations

Next, we visualize the predicted pixel-wise uncertainty of Im2Im-Deep and QUTCC for an undersampled MRI and a real noise image (Fig. 4). We compare these against the true error, which we obtain by taking the difference between the model prediction and the ground truth. Both Im2Im-Deep and QUTCC predict regions of high uncertainty that align with areas of high reconstruction error. However, QUTCC produces a more informative uncertainty map by selectively highlighting only the most uncertain regions, whereas Im2Im-Deep tends to assign broad, uniform areas of elevated uncertainty. In the MRI task shown in Fig. 4, arrows in the model prediction, uncertainty, and error maps indicate a hallucinated structure that appears in both the Im2Im-Deep and QUTCC predictions. QUTCC is

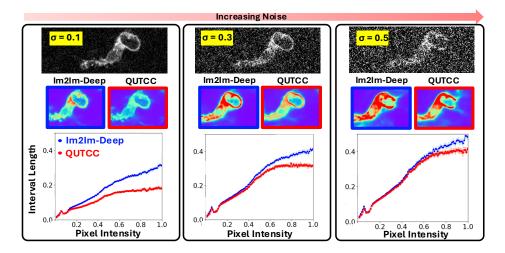


Figure 3: QUTCC exhibits smaller uncertainty intervals in regions of high pixel intensity: In a Gaussian denoising setting, we analyze how uncertainty interval lengths vary with pixel intensity across increasing noise levels ($\sigma=0.1,\ 0.3,\ 0.5$). Under low-noise conditions, QUTCC exhibits substantially narrower uncertainty intervals compared to the baseline. As noise increases to $\sigma=0.5$, this advantage becomes less pronounced overall. However, QUTCC continues to produce shorter uncertainty intervals for high-intensity pixels (intensity > 0.8), even at higher noise levels. All confidence intervals were estimated over a set of 10 samples.

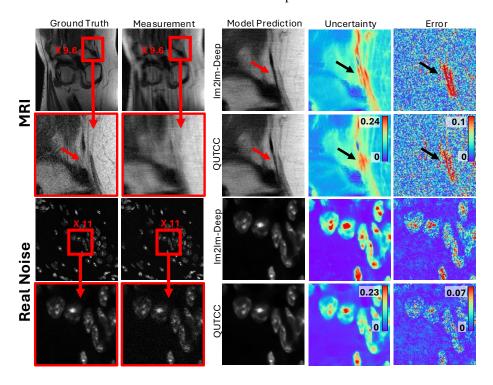


Figure 4: **Hallucination Visualization:** We show pixel-wise QUTCC and Im2Im-Deep uncertainty quantification for MRI and Real Noise tasks. In both imaging scenarios, both models highlight regions of high uncertainty that correspond to regions of high error. In the MRI task, the arrows point to a hallucination that appears in the Im2Im-Deep and QUTCC model predictions that is not present in the ground truth. QUTCC produces tighter uncertainty intervals that can better pinpoint uncertainty and hallucinations compared to Im2Im-Deep, which highlights a larger region.

 able to accurately localize this hallucination with its uncertainty prediction, but Im2Im-Deep exhibits broader, less specific uncertainty across the surrounding structure.

4.3 ESTIMATING THE CONDITIONAL DISTRIBUTION

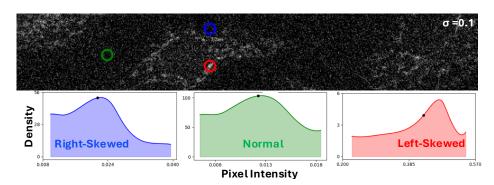


Figure 5: **QUTCC predicts diverse pixel-wise distributions:** By querying multiple quantiles and conformalizing the result, we can predict pixel-wise conditional distributions with statistical coverage guarantees. Here, we show several predicted conditional distributions for regions within an image given a measurement with Gaussian noise. At the blue, green, and red pixels, the model predicts a right-skewed distribution, a Gaussian-like distribution, and left-skewed, respectively.

Finally, we demonstrate QUTCC's ability to construct a conformalized, pixel-wise conditional PDF in Fig. 5. To provide statistical coverage guarantees, we calibrate the model across a range of miscoverage levels α (see Suppl. Fig. 9 for visualization). By systematically varying α (e.g., from 0.1 to 0.9) and recording the corresponding quantile bounds, we obtain a collection of confidence intervals that, when aggregated, approximate the full cumulative distribution function. Differentiating this cumulative distribution function yields a conformalized pixel-wise PDF with formal coverage guarantees at each risk level. In Figure 5, we illustrate QUTCC 's ability to model diverse pixel-wise predictive distributions. The blue, green, and red patches show noticeably right, symmetric, and left-skewed distributions, respectively.

5 LIMITATIONS AND CONCLUSION

We propose QUTCC, a new uncertainty quantification method for imaging inverse problems that can achieve tighter uncertainty estimates than previous methods while maintaining the same statistical coverage. QUTCC accomplishes this by training a U-Net with a quantile embedding simultaneously on $q \in (0,1)$ quantiles and then dynamically adjusting its quantile bound predictions during calibration until the desired risk is satisfied. We validated QUTCC on five imaging inverse problems - undersampled MRI, QPI, denoising under Gaussian, Poisson, and real noise, comparing its performance against prior conformal methods for image regression tasks. Our method exhibited tighter uncertainty intervals, on average, while still pinpointing model hallucinations and regions of high error. This can be attributed to our model learning and applying a nonlinear and asymmetrical scaling to its pixel-wise uncertainty predictions. Furthermore, QUTCC can estimate a conformalized conditional PDF, which previous conformal uncertainty quantification methods for image-to-image regression tasks could not do. While quantifying model uncertainty remains a significant open challenge in the field of deep learning, we believe that QUTCC offers a simple, yet robust method of uncertainty quantification for imaging inverse problems and image-to-image regression tasks. Some limitations are that QUTCC has a need for paired data for both the training and conformal calibration step. Additionally, we do not consider motion or 3D effects, which can be present in real samples. Interesting future work includes considering the effects of sample movement and distribution shifts, as well as uncertainty across multiple measurements instead of considering uncertainty for a single-frame independently.

Ethical Considerations: Using machine learning methods for scientific and medical applications has inherent risk - producing realistic artifacts that are not truly present in the image can be catastrophic for discovery and medical diagnostics. As research in uncertainty quantification matures, we hope that some of these risks can be mitigated to enable more trustworthy imaging. We acknowledge and fully comply with the ICLR Code of Ethics.

Reproducibility: We describe the different imaging inverse tasks in more detail in A.1. In this section, we also include the model epochs that we used to evaluate in this paper. Additional explanation of QUTCC's model architecture can be found in A.2, where we describe the quantile embedding in more detail. Code to train our model and reproduce the experiments shown here can be found at this anonymous code repo: https://anonymous.4open.science/r/QUTCC-246F.

REFERENCES

- Amal Alshardan, Hany Mahgoub, Nuha Alruwais, Abdulbasit A Darem, Wafa Sulaiman Almukadi, and Abdullah Mohamed. Deep learning solutions for inverse problems in advanced biomedical image analysis on disease detection. *Scientific Reports*, 14(1):18478, 2024.
- Anastasios Angelopoulos, Stephen Bates, Jitendra Malik, and Michael I Jordan. Uncertainty sets for image classifiers using conformal prediction. *arXiv preprint arXiv:2009.14193*, 2020.
- Anastasios N Angelopoulos and Stephen Bates. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv* preprint arXiv:2107.07511, 2021.
- Anastasios N Angelopoulos, Stephen Bates, Emmanuel J Candès, Michael I Jordan, and Lihua Lei. Learn then test: Calibrating predictive algorithms to achieve risk control. *arXiv preprint arXiv:2110.01052*, 2021.
- Anastasios N Angelopoulos, Amit Pal Kohli, Stephen Bates, Michael Jordan, Jitendra Malik, Thayer Alshaabi, Srigokul Upadhyayula, and Yaniv Romano. Image-to-image regression with distribution-free uncertainty quantification and applications in imaging. In *International Conference on Machine Learning*, pp. 717–730. PMLR, 2022.
- Anastasios N Angelopoulos, Stephen Bates, et al. Conformal prediction: A gentle introduction. *Foundations and Trends*® *in Machine Learning*, 16(4):494–591, 2023.
- George Barbastathis, Aydogan Ozcan, and Guohai Situ. On the use of deep learning for computational imaging. *Optica*, 6(8):921–943, 2019.
- Rina Foygel Barber, Emmanuel J Candes, Aaditya Ramdas, and Ryan J Tibshirani. Conformal prediction beyond exchangeability. *The Annals of Statistics*, 51(2):816–845, 2023.
- Edmon Begoli, Tanmoy Bhattacharya, and Dimitri Kusnezov. The need for uncertainty quantification in machine-assisted medical decision making. *Nature Machine Intelligence*, 1(1):20–23, 2019.
- Omer Belhasin, Yaniv Romano, Daniel Freedman, Ehud Rivlin, and Michael Elad. Principal uncertainty quantification with spatial correlation for image restoration problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5):3321–3333, 2023.
- Margarida Campos, António Farinhas, Chrysoula Zerva, Mário AT Figueiredo, and André FT Martins. Conformal prediction for natural language processing: A survey. *Transactions of the Association for Computational Linguistics*, 12:1497–1516, 2024.
- Matt Y Cheung, Tucker J Netherton, Laurence E Court, Ashok Veeraraghavan, and Guha Balakrishnan. Metric-guided image reconstruction bounds via conformal prediction. *ArXiv*, pp. arXiv–2404, 2024.
- Youngseog Chung, Willie Neiswanger, Ian Char, and Jeff Schneider. Beyond pinball loss: Quantile methods for calibrated uncertainty quantification. *Advances in Neural Information Processing Systems*, 34:10971–10984, 2021.
- Alvaro HC Correia, Fabio Valerio Massoli, Christos Louizos, and Arash Behboodi. An information theoretic perspective on conformal prediction. *arXiv* preprint arXiv:2405.02140, 2024.
- Kiranmoy Das, Martin Krzywinski, and Naomi Altman. Quantile regression, 2019.
- Nicolas Dewolf, Bernard De Baets, and Willem Waegeman. Valid prediction intervals for regression problems. *Artificial Intelligence Review*, 56(1):577–613, 2023.

- Tiffany Ding, Anastasios Angelopoulos, Stephen Bates, Michael Jordan, and Ryan J Tibshirani. Class-conditional conformal prediction with many classes. *Advances in neural information processing systems*, 36:64555–64576, 2023.
 - Clara Fannjiang, Stephen Bates, Anastasios N Angelopoulos, Jennifer Listgarten, and Michael I Jordan. Conformal prediction under feedback covariate shift for biomolecular design. *Proceedings of the National Academy of Sciences*, 119(43):e2204569119, 2022.
 - Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059. PMLR, 2016.
 - Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International conference on machine learning*, pp. 1183–1192. PMLR, 2017.
 - Saumya Gupta, Yikai Zhang, Xiaoling Hu, Prateek Prasanna, and Chao Chen. Topology-aware uncertainty for image segmentation. *Advances in Neural Information Processing Systems*, 36:8186–8207, 2023.
 - Xuming He. Quantile curves without crossing. The American Statistician, 51(2):186–192, 1997.
 - Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
 - Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
 - Florian Knoll, Jure Zbontar, Anuroop Sriram, Matthew J Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J Geras, Joe Katsnelson, Hersh Chandarana, et al. fastmri: A publicly available raw k-space and dicom dataset of knee images for accelerated mr image reconstruction using machine learning. *Radiology: Artificial Intelligence*, 2(1):e190007, 2020.
 - Roger Koenker and Gilbert Bassett Jr. Regression quantiles. *Econometrica: journal of the Econometric Society*, pp. 33–50, 1978.
 - Roger Koenker and Kevin F Hallock. Quantile regression. *Journal of economic perspectives*, 15(4):143–156, 2001.
 - Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
 - Jordan Lekeufack, Anastasios N Angelopoulos, Andrea Bajcsy, Michael I Jordan, and Jitendra Malik. Conformal decision theory: Safe autonomous decisions from imperfect predictions. In 2024 IEEE International Conference on Robotics and Automation (ICRA), pp. 11668–11675. IEEE, 2024.
 - Lars Lindemann, Matthew Cleaveland, Gihyun Shim, and George J Pappas. Safe planning in dynamic environments using conformal prediction. *IEEE Robotics and Automation Letters*, 8(8):5116–5123, 2023.
 - Yufeng Liu and Yichao Wu. Simultaneous multiple non-crossing quantile regression estimation using kernel constraints. *Journal of nonparametric statistics*, 23(2):415–437, 2011.
 - Elias Nehme, Omer Yair, and Tomer Michaeli. Uncertainty quantification via neural posterior principal components. *Advances in Neural Information Processing Systems*, 36:37128–37141, 2023.
 - Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, et al. Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999*, 2018.
 - Gregory Ongie, Ajil Jalal, Christopher A. Metzler, Richard G. Baraniuk, Alexandros G. Dimakis, and Rebecca Willett. Deep learning techniques for inverse problems in imaging, 2020. URL https://arxiv.org/abs/2005.06001.
 - Henry Pinkard, Cherry Liu, Fanice Nyatigo, Daniel A Fletcher, and Laura Waller. The berkeley single cell computational microscopy (bsccm) dataset. *arXiv preprint arXiv:2402.06191*, 2024.
 - Apostolos F Psaros, Xuhui Meng, Zongren Zou, Ling Guo, and George Em Karniadakis. Uncertainty quantification in scientific machine learning: Methods, metrics, and comparisons. *Journal of Computational Physics*, 477:111902, 2023.
 - Victor Quach, Adam Fisch, Tal Schuster, Adam Yala, Jae Ho Sohn, Tommi S Jaakkola, and Regina Barzilay. Conformal language modeling. *arXiv* preprint arXiv:2306.10193, 2023.

- Filipe Rodrigues and Francisco C Pereira. Beyond expectation: Deep joint mean and quantile regression for spatiotemporal problems. *IEEE transactions on neural networks and learning systems*, 31(12):5377–5389, 2020.
 - Yaniv Romano, Evan Patterson, and Emmanuel Candes. Conformalized quantile regression. *Advances in neural information processing systems*, 32, 2019.
 - Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pp. 234–241. Springer, 2015.
 - Maxime Sangnier, Olivier Fercoq, and Florence d'Alché Buc. Joint quantile regression in vector-valued rkhss. Advances in Neural Information Processing Systems, 29, 2016.
 - Swami Sankaranarayanan, Anastasios Angelopoulos, Stephen Bates, Yaniv Romano, and Phillip Isola. Semantic uncertainty intervals for disentangled latent spaces. In *NeurIPS*, 2022.
 - Ingo Steinwart and Andreas Christmann. Estimating conditional quantiles with the help of the pinball loss. 2011.
 - Yu Sun, Zihui Wu, Yifan Chen, Berthy T Feng, and Katherine L Bouman. Provable probabilistic imaging using score-based generative priors. *IEEE Transactions on Computational Imaging*, 2024.
 - Natasa Tagasovska and David Lopez-Paz. Single-model uncertainties for deep learning. *Advances in neural information processing systems*, 32, 2019.
 - Vladimir Vovk. Conditional validity of inductive conformal predictors. In Asian conference on machine learning, pp. 475–490. PMLR, 2012.
 - Jeffrey Wen, Rizwan Ahmad, and Philip Schniter. Task-driven uncertainty quantification in inverse problems via conformal prediction. In *European Conference on Computer Vision*, pp. 182–199. Springer, 2024.
 - Ran Xie, Rina Barber, and Emmanuel Candes. Boosted conformal prediction intervals. *Advances in Neural Information Processing Systems*, 37:71868–71899, 2024.
 - Chen Xu and Yao Xie. Sequential predictive conformal inference for time series. In *International Conference on Machine Learning*, pp. 38707–38727. PMLR, 2023.
 - Yujia Xue, Shiyi Cheng, Yunzhe Li, and Lei Tian. Reliable deep-learning-based phase imaging with uncertainty quantification. *Optica*, 6(5):618–629, 2019.
 - Omer Yair, Elias Nehme, and Tomer Michaeli. Uncertainty visualization via low-dimensional posterior projections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11041–11051, 2024.
 - Cassandra Tong Ye, Jiashu Han, Kunzan Liu, Anastasios Angelopoulos, Linda Griffith, Kristina Monakhova, and Sixian You. Learned, uncertainty-driven adaptive acquisition for photon-efficient scanning microscopy. *Optics Express*, 33(6):12269–12287, 2025.
 - Margaux Zaffran, Olivier Féron, Yannig Goude, Julie Josse, and Aymeric Dieuleveut. Adaptive conformal predictions for time series. In *International Conference on Machine Learning*, pp. 25834–25866. PMLR, 2022.
 - Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, et al. fastmri: An open dataset and benchmarks for accelerated mri. arXiv preprint arXiv:1811.08839, 2018.
 - Yi Zhang, Xiaoyu Shi, Dasong Li, Xiaogang Wang, Jian Wang, and Hongsheng Li. A unified conditional framework for diffusion-based image restoration. Advances in Neural Information Processing Systems, 36: 49703–49714, 2023.
 - Yide Zhang, Yinhao Zhu, Evan Nichols, Qingfei Wang, Siyuan Zhang, Cody Smith, and Scott Howard. A poisson-gaussian denoising dataset with real fluorescence microscopy images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11710–11718, 2019.

A APPENDIX

These supplementary materials contain additional information to improve the reproducibility of the paper, as well as several additional analyses and experimental results. Code for training and analysis of QUTCC can be found at the following anonymous repo: https://anonymous.4open.science/r/QUTCC-246F. Section A.1 provides more details about the different inverse problems used throughout the paper and their associated datasets. Section A.2 details the network architecture used for QUTCC. Section A.3 provides additional comparisons against the original formulation of Im2Im-UQ without our network modifications and analysis on quantile crossing. Section A.4 includes analysis on the size-stratified risk, more visualizations of our uncertainty predictions across the different inverse problems, and a visualization of how QUTCC produces narrower uncertainty intervals through asymmetry. Finally, Sec. A.5 provides several visualizations of QUTCC's conditional distribution predictions, including an example showing how we conformalize the quantiles and the PDF prediction as a function of noise.

A.1 EXPERIMENT DETAILS

In this section, we provide additional experimental details about our training process and datasets used for our five different tasks: denoising (real, Gaussian, Poisson), MRI, and QPI. Im2Im-Deep and QUTCC were each trained for 50 epochs on their respective datasets using a single NVIDIA RTX PRO 6000. Following training, we conducted a model selection sweep to identify the epoch that yielded the narrowest uncertainty intervals while satisfying the target risk level ($\alpha = 0.1$).

For all models, an initial learning rate of 1e-4 and weight decay of 0 was used. Batch size was adjusted depending on the task, with 4 used for denoising tasks, 12 for MRI, and 72 for QPI for U-Net backbones and 32, 512, and 16 (respectively) for ResNet-18 backbones. Images were normalized before training.

A.1.1 REAL NOISE TASK

For the real noise task, we used the Fluorescence Microscopy Denoising (FMD) dataset (Zhang et al., 2019), which contains experimentally acquired fluorescence microscopy images spanning 12 wide-field, confocal, and two-photon modalities. The model was trained on 10,000 images, with 500 confocal mouse images used for calibration and an additional 500 for validation. The epochs used for Im2Im-Deep and QUTCC are 5 and 10, respectively.

A.1.2 GAUSSIAN AND POISSON NOISE TASK

For both the Gaussian and Poisson tasks, we synthetically introduced varying levels of noise to the FMD ground truth images. The dataset was split into 180 ground truth images for training, 40 for calibration, and 20 for validation. The pseudocode for generating gaussian and poisson noise are shown below in Algorithm 2 and Algorithm 3.

Algorithm 2 Add Gaussian Noise

- 1: **Input:** Clean image x, max noise level σ_{max}
- 2: Output: Noisy image y
- 3: Sample noise std: $\sigma \sim \mathcal{U}(0, \sigma_{\text{max}})$
- 4: Sample Gaussian noise: $\eta \sim \mathcal{N}(0, \sigma^2)$
- 5: Add noise: $\mathbf{y} \leftarrow \mathbf{x} + \boldsymbol{\eta}$
- 6: Clamp: $\mathbf{y} \leftarrow \text{Clamp}(\mathbf{y}, 0, 1)$
- 7: return y

Algorithm 3 Add Poisson Noise

- 1: **Input:** Clean image x, min/max noise levels σ_{\min} , σ_{\max}
- 2: Output: Noisy image y
- 3: Sample scale: $\lambda \sim \mathcal{U}(\sigma_{\min}, \sigma_{\max})$
- 4: Scale image: $\mathbf{x}_{\text{scaled}} \leftarrow \lambda \cdot \mathbf{x}$
- 5: Sample noise: $\eta \sim \text{Poisson}(\mathbf{x}_{\text{scaled}})$
- 6: Clamp: $\mathbf{y} \leftarrow \text{Clamp}(\boldsymbol{\eta}, 0, 1)$
- 7: return y

For Gaussian noise, we set $\sigma_{\rm max}=0.5$. For Poisson noise, the noise level range $(\lambda_{\rm min},\lambda_{\rm max})$ was set to (50,100). In each iteration, with a batch size of 16, we generated 25 random noise levels uniformly sampled between the specified minimum and maximum values. The number of training

epochs for Im2Im-Deep and QUTCC with Gaussian noise were 15 and 20, respectively. For Poisson noise, Im2Im-Deep was trained for 10 epochs, while QUTCC was trained for 35 epochs.

A.1.3 MAGNETIC RESONANCE IMAGING (MRI) TASK

Data used for the MRI task was obtained from the NYU fastMRI initiative database (fastmri.med.nyu.edu) (Knoll et al., 2020; Zbontar et al., 2018). The primary goal of fastMRI is to test whether machine learning can aid in the reconstruction of medical MRI images. To train our models, we split the fastMRI dataset into 700 volumes for training, 200 volumes for calibration, and 200 for validation. It is important to note that a single volume contains multiple MRI images. The epochs used for Im2Im-Deep and QUTCC are 25 and 30, respectively.

To simulate the forward model in MRI, we start with a fully-sampled 3D volume composed of multiple 2D image slices. Each 2D image slice is transformed into its frequency domain representation using the 2D Fourier Transform, producing its k-space data. To simulate undersampled acquisition, we apply a $4\times$ undersampling mask to the k-space. The resulting undersampled k-space is then transformed back into the image domain using the inverse Fourier Transform, yielding an aliased or artifact-corrupted image that serves as the input for the models.

A.1.4 QUANTITATIVE PHASE IMAGING (QPI) TASK

Data used for the QPI task were obtained from the Berkeley Single Cell Computational Microscopy (BSCCM) dataset (Pinkard et al., 2024). The BSCCM dataset contains image samples of individual white blood cells that have been captured with several illumination patterns on an LED array microscope. For this task, we used 289,059 images for training, 82,588 images for calibration, and 41,294 images for validation. The number of training epochs for Im2Im-Deep and QUTCC were 15 and 20, respectively.

In QPI, the goal is to image the structure of transparent cells by recovering the phase delay of the light that passes through the cells. Several intensity-only images are used to computationally estimate the phase of the light, since this cannot be measured directly. The input measurement y is the concatenation of two cell intensity images acquired at different illumination angles. The corresponding ground truth is the quantitative phase image recovered from four or more illumination angles. The model is trained to map these two intensity images into a phase image.

A.2 Model Architecture

Our method, QUTCC, is based on a U-Net backbone augmented with self-attention mechanisms, where quantile embeddings are propagated through the self-attention layers to guide the network's quantile predictions. In our design, the target quantile level is embedded as a continuous scalar, analogous to the time-step embeddings in diffusion models. A core part of this architecture is the integration of self-attention layers within the U-Net, implemented as AttentionBlock modules. These blocks allow the model to capture global dependencies across spatial dimensions, enabling each pixel or feature location to attend to all other locations within its feature map. Specifically, attention layers are incorporated at various downsampling resolutions within both the encoder and decoder paths, as well as at the bottleneck of the U-Net.

The ability to condition the model's output on a given quantile is achieved through the quantile embedding mechanism. A quantile value between (0, 1) is chosen for each image sample at each iteration and then transformed into a high-dimensional vector representation using a positional encoding scheme, employing sinusoidal functions to create generalizable embeddings. This initial embedding is then processed by a small multi-layer perceptron, which is then used throughout to condition the network. By introducing this conditioning, the model learns to generate outputs that are responsive to the entire range of quantile levels. This mechanism works in tandem with the self-attention layers, which are specified using the attention_resolutions parameter.

In our experiments, QUTCC was trained using attention_resolutions configured as [16, 8, 4, 2, 1]. This means for 512×512 images, using the specified configuration results in attention layers within the encoder path (at resolutions 512x512, 256x256, 128x128, 64x64, and 32x32), one central attention layer in the middle block (at 8x8 resolution), and additional attention layers distributed across the decoder path (at resolutions 32x32, 64x64, 128x128, 256x256,

and 512x512). The attention_resolutions parameter dictates the spatial scales at which these attention mechanisms are introduced. Further model configurations can be found in models/model_config.yaml.

A.3 MODEL ANALYSIS

In Fig. 4, we restricted our comparisons to the Im2Im-Deep and QUTCC models. The original Im2Im-UQ model was excluded due to its comparatively shallow architecture, resulting in decreased performance. However, for the subsequent analysis, we reintroduce Im2Im-UQ for completeness. In this section, we assess the model's mean predictive performance and the quantile crossing occurrences.

A.3.1 MEAN PREDICTION PERFORMANCE

Does QUTCC produce tighter intervals because it is simply a better image prediction network? To investigate this, we compare the predictive performance of Im2Im-UQ, Im2Im-Deep and QUTCC. In Table 2, we compare the performance of Im2Im-Deep and QUTCC using standard image reconstruction metrics: MSE, SSIM, PSNR, and LPIPS. For QUTCC, the mean prediction was obtained by setting the quantile level to q=0.5. The results indicate that all models achieve nearly identical performance in terms of MSE, with only minor differences observed in SSIM, PSNR, and LPIPS. These variations are not substantial enough to suggest that QUTCC provides a significantly better mean prediction. These findings suggest that QUTCC 's improved uncertainty quantification predictions are not attributed to better mean prediction performance. Rather, its ability to more effectively characterize uncertainty appears to come from the explicit learning of quantiles during training.

Metric	Model	MRI	QPI	Gaussian	Poisson	Real Noise
MSE (↓)	Im2Im-UQ Im2Im-Deep QUTCC	$\begin{array}{c} 0.003 \pm 0.002 \\ 0.001 \pm 0.002 \\ 0.001 \pm 0.002 \end{array}$	$\begin{array}{c} 0.0006 \pm 0.0005 \\ 0.0004 \pm 0.0003 \\ 0.0004 \pm 0.0003 \end{array}$	$\begin{array}{c} 0.0007 \pm 0.0006 \\ 0.0006 \pm 0.0006 \\ 0.0006 \pm 0.0005 \end{array}$	0.0003 ± 0.0003 0.0003 ± 0.0002 0.0003 ± 0.0003	0.0030 ± 0.0004 0.0004 ± 0.0002 0.0002 ± 0.0001
SSIM (†)	Im2Im-UQ Im2Im-Deep QUTCC	$\begin{array}{c} 0.668 \pm 0.127 \\ 0.707 \pm 0.139 \\ 0.708 \pm 0.139 \end{array}$	$\begin{array}{c} 0.949 \pm 0.017 \\ 0.961 \pm 0.010 \\ 0.959 \pm 0.010 \end{array}$	$\begin{array}{c} 0.852 \pm 0.102 \\ 0.856 \pm 0.107 \\ 0.865 \pm 0.102 \end{array}$	$\begin{array}{c} 0.931 \pm 0.038 \\ 0.937 \pm 0.035 \\ 0.941 \pm 0.036 \end{array}$	$\begin{array}{c} 0.803 \pm 0.016 \\ 0.959 \pm 0.006 \\ 0.957 \pm 0.008 \end{array}$
PSNR (†)	Im2Im-UQ Im2Im-Deep QUTCC	26.867 ± 2.923 29.711 ± 3.138 29.833 ± 3.156	33.135 ± 2.543 34.565 ± 2.393 34.948 ± 2.436	32.739 ± 3.535 33.557 ± 4.018 33.660 ± 4.143	36.163 ± 3.059 37.062 ± 2.968 37.498 ± 3.797	25.833 ± 0.734 34.038 ± 1.837 37.350 ± 1.936
LPIPS (↓)	Im2Im-UQ Im2Im-Deep QUTCC	0.343 ± 0.033 0.324 ± 0.043 0.323 ± 0.040	$\begin{array}{c} 0.153 \pm 0.025 \\ 0.125 \pm 0.015 \\ 0.121 \pm 0.015 \end{array}$	$\begin{array}{c} 0.420 \pm 0.092 \\ 0.414 \pm 0.103 \\ 0.408 \pm 0.102 \end{array}$	$\begin{array}{c} 0.294 \pm 0.071 \\ 0.299 \pm 0.071 \\ 0.284 \pm 0.072 \end{array}$	$\begin{array}{c} 0.360 \pm 0.033 \\ 0.297 \pm 0.029 \\ 0.312 \pm 0.026 \end{array}$

Table 2: Image reconstruction performance of Im2Im-UQ, Im2Im-Deep, and QUTCC: For each metric, the arrow indicates the direction of better performance.

A.3.2 QUANTILE CROSSING PERFORMANCE

In section 3.4 we describe the conformal calibration step, which is dependent on the quantile function being monotonic. To ensure the validity of the predicted quantiles, specifically to avoid the issue of quantile crossing, we quantified the number of quantile crossing occurrences between q = [0.1, 0.2, 0.3, ..., 0.9] in QUTCC (Tbl 3). Quantile crossing can undermine the interpretability of our uncertainty estimates, as

Task	Crossed Pixels	Total Pixels	Crossing Ratio
MRI	2.20e1	1.64e9	1.34e-8
QPI	5.00	2.62e8	$1.91e{-8}$
Gaussian	1.10e6	3.36e9	$3.29e{-4}$
Poisson	3.49e3	3.36e9	$1.04e{-6}$
Real Noise	3.38e4	1.05e9	$3.23e{-5}$

Table 3: Quantile crossing occurrences for QUTCC

it contradicts the notion that quantile functions should be non-decreasing/non-overlapping. The results indicate that across all imaging tasks, the ratio of quantile crossing occurrences is minimal.

A.4 ADDITIONAL UNCERTAINTY INTERVAL EVALUATIONS

A.4.1 SIZE-STRATIFIED RISK

We observed the size-stratified risk of all inverse tasks between Im2Im-Deep and QUTCC (Fig. 6). To calculate size-stratified risk, the prediction intervals are first binned into different sizes, ranging from smallest to largest. Then the risk is calculated across all the bins to ensure that the model's uncertainty estimates are well-calibrated across different levels of confidence. While both Im2Im-Deep and QUTCC have bins that exceed the α , generally, most bins fall under the chosen risk.

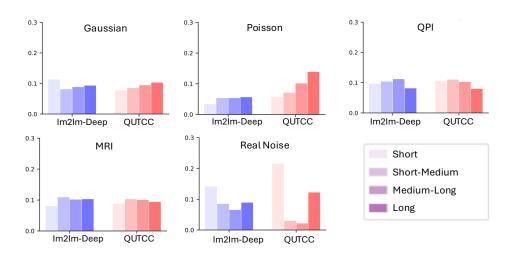


Figure 6: **Size-Stratified Risk of Im2Im-Deep vs. QUTCC**: We evaluate the size-stratified risk of Im2Im-Deep and QUTCC across all tasks. Overall, neither model exhibits a strong relationship between interval width and empirical risk, suggesting limited sensitivity to interval size. However, in the Gaussian and Poisson settings, both models display a mild trend toward improved calibration, or lower risk, for narrower prediction intervals.

A.4.2 ADDITIONAL VISUALIZATIONS

We also provide visualizations of the remaining imaging tasks not included in the main results (Fig. 7). For all sample tasks, both QUTCC and Im2Im-Deep effectively highlight regions with high reconstruction error. However, QUTCC demonstrates a more focused identification of areas with high uncertainty that align closely with the true error, indicating its greater precision in uncertainty estimation. While Im2Im-Deep is capable of identifying regions of error, it tends to assign elevated uncertainty across larger portions of the sample, making it challenging for downstream tasks to prioritize regions based on uncertainty interval sizes. This trend is consistent across all five imaging inverse problems.

A.4.3 BOUND ASYMMETRY

How does QUTCC achieve tighter confidence intervals while maintaining the same level of risk? As shown in Fig. 8, QUTCC produces asymmetric predictive intervals—its upper and lower bounds are adjusted independently based on localized uncertainty. In contrast, Im2Im-UQ applies a single global scaling factor λ uniformly to both bounds, which can be suboptimal in cases where only one side of the interval requires adjustment. This limitation often leads to unnecessarily widened intervals. In the red boxed region of Fig. 8, both QUTCC and Im2Im-UQ share a similar lower bound, yet QUTCC predicts a significantly tighter upper bound. Similarly, in the green boxed region, both methods align on the upper bound, but QUTCC yields a tighter lower bound. These examples highlight QUTCC's ability to adaptively adjust its interval predictions, leading to more precise interval estimates.

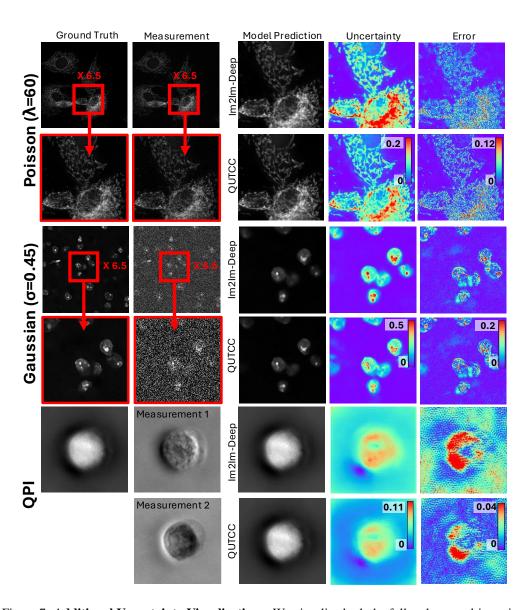


Figure 7: Additional Uncertainty Visualizations: We visualize both the full and zoomed-in regions of image reconstructions for QPI and denoising with Poisson, Gaussian and Real Noise. Consistent with observations presented in the results section, QUTCC produces more precise uncertainty estimates that closely align with localized regions of high reconstruction error. In contrast, Im2Im-Deep tends to highlight broader regions of uncertainty and lacks specificity, making it hard to distinguish areas of importance. This highlights QUTCC 's ability predict more informative uncertainty maps.

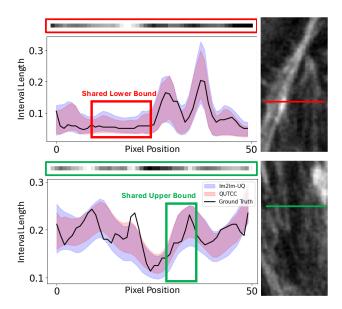


Figure 8: QUTCC produces narrower intervals through asymmetric bounds: We analyze the pixel-wise uncertainty bounds predicted by Im2Im-UQ and QUTCC and observe that QUTCC exhibits asymmetric behavior in its interval estimates. In Im2Im-UQ, both the upper and lower bounds are uniformly scaled by a global factor λ to satisfy coverage constraints, which limits flexibility in adapting to signal-specific uncertainty. In contrast, QUTCC learns to predict quantiles directly, enabling it to independently modulate upper and lower bounds based on signal characteristics. This results in more adaptive and efficient uncertainty intervals. For instance, in the red boxed region, QUTCC matches Im2Im-UQ's lower bound but produces a significantly tighter upper bound. Conversely, in the green boxed region, both models share an upper bound, yet QUTCC yields a tighter lower bound. Samples shown are Gaussian images with $\sigma=0.1$.

A.5 ADDITIONAL PDF RESULTS

In this section, we present additional results highlighting QUTCC's ability to estimate a conditional probability density function. In Fig. 9a, we show several of the predicted PDFs for image denoising with Gaussian noise with $\sigma=0.4$. Detailed views of the corresponding pixel-wise PDFs are presented in Fig. 9b for the low-intensity pixel and Fig. 9c for the high-intensity pixel. Each graph displays two PDFs: the blue PDF represents the quantile predictions from the uncalibrated model, while the green PDF represents the conformally calibrated quantiles that provide finite-sample statistical coverage guarantees. The conformal calibration procedure adjusts the quantile levels to ensure valid coverage properties. For instance, while the 25th and 75th percentiles theoretically provide 50% coverage, conformal calibration determines that the 20.7th and 69.2nd percentiles are required to achieve statistically guaranteed 50% coverage for this specific dataset and model. Additionally, Fig. 10 illustrates the evolution of the pixel-wise PDF for a fixed pixel coordinate under varying Gaussian noise levels $\sigma \in \{0.1, 0.3, 0.5\}$. As the noise standard deviation increases, the PDF widens. This widening corresponds to increased epistemic uncertainty in the model's predictions, as higher noise levels reduce the information content available for accurate pixel intensity estimation. This then directly widens prediction intervals to maintain the desired coverage guarantees.

A.6 LLM USAGE

LLMs were used to edit and correct grammar during writing, but were not involved in the research ideation process.

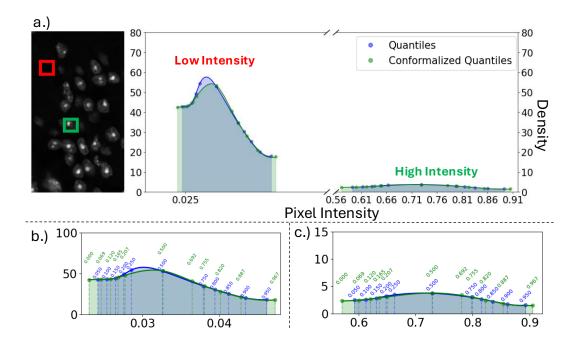


Figure 9: QUTCC predicts different pixel-wise PDFs based on different siganl intensity a) Comparison of pixel-wise PDFs for representative low-intensity and high-intensity pixels in a Gaussian measurement ($\sigma=0.4$). b) Detailed view of the low-intensity pixel PDF, exhibiting a narrow, high-density distribution concentrated around few intensity values, indicating low predictive uncertainty. c) Detailed view of the high-intensity pixel PDF, showing a broader, lower-density distribution with increased spread, reflecting higher predictive uncertainty in bright image regions. The blue lines show the uncalibrated model, while the green lines show the conformalized quantiles after calibration.

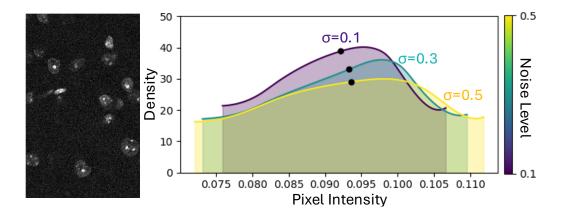


Figure 10: **PDF broadens as noise increases:** We observe the PDF of a single pixel under varying noise levels. At $\sigma=0.1$, the noise is low, and the PDF is compact. As the noise increases to $\sigma=0.3$ and $\sigma=0.5$, the PDF gradually broadens, while the mean prediction value remains relatively unchanged. This broadening occurs due to the increased uncertainty introduced by higher noise levels. QUTCC successfully predicts this increased uncertainty as the noise increases.