# Bounding Estimation Performance in Inverse Problems via Conformal Prediction


**Phil Schniter**

(Joint work with Jeffrey Wen and Rizwan Ahmad)

The Ohio State University



Supported in part by the NIH under grant R01-EB029957


Asilomar Conference on Signals, Systems, and Computers
October 2025

# Inverse problems

Inverse problems:

- Unknown quantity $\boldsymbol{x}_0 \in \mathbb{R}^d$
- Masked/noisy/distorted measurements $\boldsymbol{y} = \boldsymbol{h}(\boldsymbol{x}_0)$
  - Communications examples: wireless device localization, CSI estimation, RF tomography
  - Imaging examples: MRI, CT, inpainting, deblurring, super-resolution, phase retrieval
- Estimate $\widehat{\boldsymbol{x}} = \boldsymbol{r}(\boldsymbol{y})$

Challenge:

- Typically ill-posed: many hypotheses of $\boldsymbol{x}_0$ form good explanations of $\boldsymbol{y}$

Subject of this talk:

- We want to quantify the uncertainty of the estimate $\widehat{\boldsymbol{x}}$
- In particular, we want rigorous probabilistic bounds on the accuracy of $\widehat{\boldsymbol{x}}$

# Example: Wireless transmitter localization

- We'd like to estimate the 2D Tx location $\boldsymbol{x}_0$ from a pilot signal and Rx array measurements $\boldsymbol{y}$
- Ill-posed due to noise, multipath fading, array ambiguities, non-ideal antenna elements, etc
- The posterior $p(\boldsymbol{x}_0|\boldsymbol{y})$ is shown on the right:
- For a user-chosen error-rate $\alpha \in (0,1)$ and arbitrary estimator $\widehat{\boldsymbol{x}} = \boldsymbol{r}(\boldsymbol{y})$, can we construct an upper bound $\beta(\boldsymbol{y})$ such that

$$\Pr\left\{\|\widehat{\boldsymbol{x}} - \boldsymbol{x}_0\| \leq \beta(\boldsymbol{y})\right\} \geq 1 - \alpha \ ?$$

Beyond Point Estimates: Likelihood-Based Full-Posterior Wireless Localization

Haozhe Lei*[1], Hao Guo[1,2], Tommy Svensson[2], and Sundeep Rangan[1]

arXiv:2509.25719

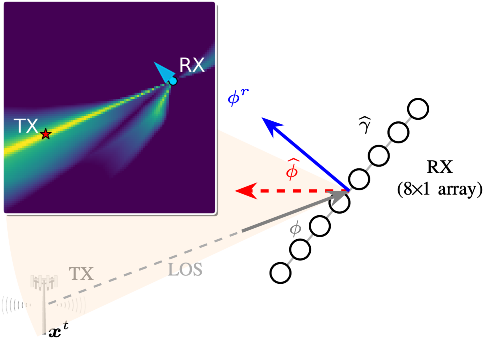

# Example: Accelerated MRI

- $\boldsymbol{x}_0$ is the true image
- $\boldsymbol{y}$ are $\frac{1}{R}$ sub-Nyquist measurements
- $\widehat{\boldsymbol{x}} = \boldsymbol{r}(\boldsymbol{y})$ is the reconstructed image

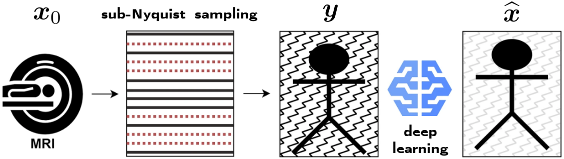


Problem: modern recovery methods can hallucinate, i.e., generate clean but inaccurate $\widehat{\boldsymbol{x}}$. For example:[1]

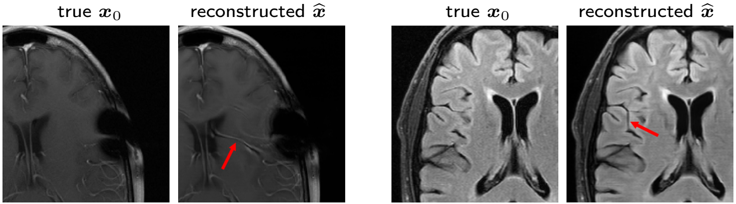


[1] Muckley et al.'21

# Probabilistic bounds on estimation accuracy

- We'd like to know the accuracy of $\widehat{\boldsymbol{x}}_0$ relative to the true $\boldsymbol{x}_0$
  - From now on, subscript "$0$" indicates "test" quantity vs calibration quantity

- To quantify accuracy, we'll use an arbitrary metric $\boxed{z_0 = m(\widehat{\boldsymbol{x}}_0, \boldsymbol{x}_0)}$, such as
  - $m(\widehat{\boldsymbol{x}}_0, \boldsymbol{x}_0) = -\|\widehat{\boldsymbol{x}}_0 - \boldsymbol{x}_0\|_p$ for wireless localization
  - $m(\widehat{\boldsymbol{x}}_0, \boldsymbol{x}_0) =$ PSNR or SSIM or LPIPS or DISTS for images

- Is it possible to guarantee the accuracy of $\widehat{\boldsymbol{x}}_0$, i.e., construct a lower bound $\boxed{\beta_0(\boldsymbol{y}_0)}$ such that

$$\Pr\{Z_0 \geq \beta_0(\boldsymbol{Y}_0)\} \geq 1 - \alpha$$

for some chosen error rate $\alpha$? (Here, capital letters denote random variables)

# What if we had a perfect posterior sampler?

- Suppose we had a perfect posterior sampler generating $n_{\text{post}}$ independent samples $\{\widetilde{\boldsymbol{x}}_0^{(j)}\}_{j=1}^{n_{\text{post}}}$

$$\{\widetilde{\boldsymbol{x}}_0^{(1)}, \dots, \widetilde{\boldsymbol{x}}_0^{(n_{\text{post}})}\} \sim p_{\boldsymbol{X}_0 \,|\, \boldsymbol{Y}_0}(\cdot \,|\, \boldsymbol{y}_0)$$

- Define the corresponding accuracy samples $\widetilde{z}_0^{(j)} \triangleq m(\widehat{\boldsymbol{x}}_0, \widetilde{\boldsymbol{x}}_0^{(j)})$:

$$\{\widetilde{z}_0^{(1)}, \dots, \widetilde{z}_0^{(n_{\text{post}})}\} \sim p_{Z_0 \,|\, \boldsymbol{Y}_0}(\cdot \,|\, \boldsymbol{y}_0)$$

- We can construct a lower bound $\beta_0$ that obeys $\Pr\{Z_0 \geq \beta_0 \,|\, \boldsymbol{Y}_0 = \boldsymbol{y}_0\} = 1 - \alpha$ using an empirical quantile using asymptotically many samples:

$$\beta_0 = \lim_{n_{\text{post}} \to \infty} \widehat{\beta}_{0,n_{\text{post}}} \quad \text{with} \quad \widehat{\beta}_{0,n_{\text{post}}} \triangleq \text{EmpQuant}(\alpha, \{\widetilde{z}_0^{(j)}\}_{j=1}^{n_{\text{post}}})$$

- Okay, but can we make this practical?

# A useful tool: split conformal prediction

- Given an estimate $\widehat{z}_0$ of the true $z_0 \in \mathbb{R}$, split conformal prediction[1] can construct a set $\mathcal{C}_\lambda(\widehat{z}_0)$ that contains $z_0$ with high probability. Here, $|\mathcal{C}_\lambda(\cdot)|$ grows with $\lambda \in \mathbb{R}$

- Given a user-chosen error rate $\alpha \in (0,1)$, it computes a $\widehat{\lambda}_\alpha \in \mathbb{R}$ using calibration data $d_{\textsf{cal}} = \{(z_i, \widehat{z}_i)\}_{i=1}^{n_{\textsf{cal}}}$ of size $n_{\textsf{cal}}$

- The prediction set guarantees marginal coverage

$$\Pr\left\{Z_0 \in \mathcal{C}_{\widehat{\lambda}_\alpha(D_{\textsf{cal}})}(\widehat{Z}_0)\right\} \geq 1 - \alpha$$

when the test & calibration pairs $\{(Z_0, \widehat{Z}_0), (Z_1, \widehat{Z}_1), \ldots, (Z_{n_{\textsf{cal}}}, \widehat{Z}_{n_{\textsf{cal}}})\}$ are statistically exchageable

[1] Vovk,Gammerman,Shafer'05, [2] Lei,G'Sell,Rinaldo,Tibshirani,Wasserman'18

# Approximate posterior sampling & conformal prediction

Recall we want to lower-bound the accuracy $z_0 = m(\widehat{\boldsymbol{x}}_0, \boldsymbol{x}_0)$ of estimate $\widehat{\boldsymbol{x}}_0 = \boldsymbol{r}(\boldsymbol{y}_0)$ of unknown $\boldsymbol{x}_0$

- Assume we have calibration samples $\{(\boldsymbol{x}_i, \boldsymbol{y}_i)\}_{i=1}^{n_{\mathsf{cal}}}$ in addition to the test measurments $\boldsymbol{y}_0$
- Generate approximate posterior samples $\{\widetilde{\boldsymbol{x}}_i^{(j)}\}_{j=1}^{n_{\mathsf{post}}}$ from $\boldsymbol{y}_i$ for each $i = 0, \dots, n_{\mathsf{cal}}$.
- Compute the accuracy samples $\widetilde{z}_i^{(j)} = m(\widehat{\boldsymbol{x}}_i, \widetilde{\boldsymbol{x}}_i^{(j)})$ for all $i = 0, \dots, n_{\mathsf{cal}}$ and $j = 1, \dots, n_{\mathsf{post}}$
- Construct approximate bounds $\widehat{\beta}_i = \operatorname{EmpQuant}\big(\alpha, \{\widetilde{z}_i^{(j)}\}_{j=1}^{c}\big)$ for $i = 0, \dots, n_{\mathsf{cal}}$
- Construct bound-violation scores $s_i = \widehat{\beta}_i - z_i$ for $i = 1, \dots, n_{\mathsf{cal}}$ (positive when bound is violated)
- Compute the bound calibration term

$$\widehat{\lambda}_\alpha(d_{\mathsf{cal}}) = \operatorname{EmpQuant}\Big(\tfrac{\lceil (1-\alpha)(n_{\mathsf{cal}}+1)\rceil}{n_{\mathsf{cal}}}, \{s_i\}_{i=1}^{n_{\mathsf{cal}}}\Big), \quad \text{where } d_{\mathsf{cal}} \triangleq \{s_i\}_{i=1}^{n_{\mathsf{cal}}}$$

- Finally, form the lower-bound as $\beta_0(\boldsymbol{y}_0, d_{\mathsf{cal}}) = \widehat{\beta}_0(\boldsymbol{y}_0) - \widehat{\lambda}_\alpha(d_{\mathsf{cal}})$

If $\{S_0, S_1, \dots, S_n\}$ are statistically exchangeable, then we have the marginal coverage guarantee

$$\Pr\big\{Z_0 \geq \beta_0(\boldsymbol{Y}_0, D_{\mathsf{cal}})\big\} \geq 1 - \alpha$$

# Illustration of lower-bound procedure (for accelerated MRI)

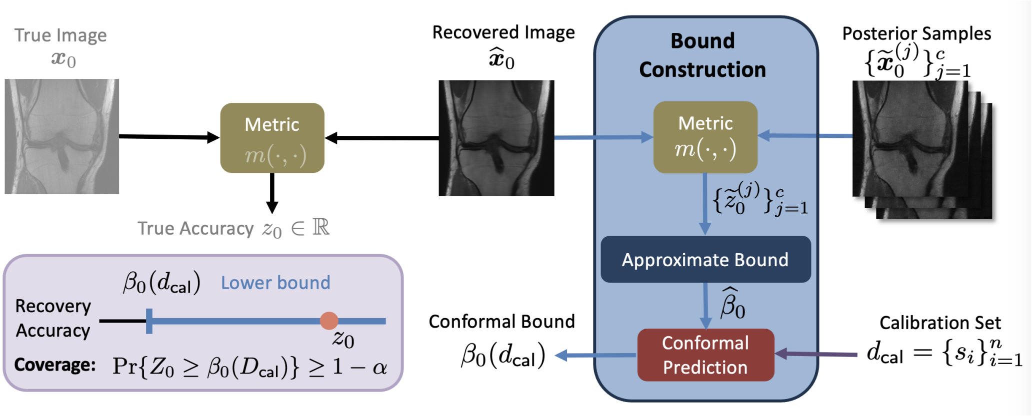


- This bounding methodology[1] works with any estimation problem, estimator, accuracy metric, and approximate posterior sampler

[1] Wen, Ahmad, Schniter'25a

# Example: Bounding recovery accuracy in MRI

- Scatter plots of $(z_0, \beta_0)$ from fastMRI knee recovery @ acceleration $R = 8$ using a conditional normalizing flow[1]:

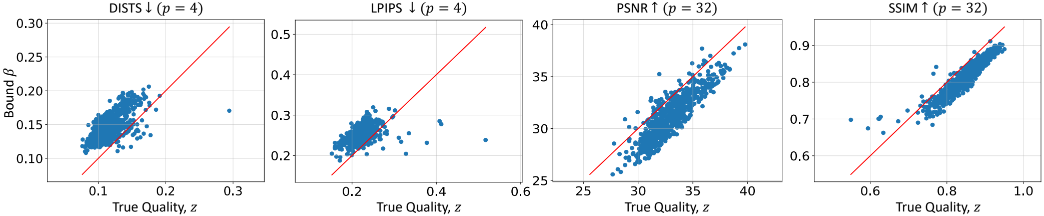


The red line indicates where the bound would be exact

- Validation of marginal coverage using $10\,000$ Monte-Carlo trials (each with a random $70\%$ test / $30\%$ calibration split):

| target coverage $1-\alpha$ | average empirical coverage |
|---|---|
| $0.95$ | $0.9504 \pm 0.0001$ |

[1] Wen, Ahmad, S'23

# Leveraging the accuracy-bound for multi-round measurements

- Often, the estimate $\widehat{\boldsymbol{x}}_0$ can be improved by gathering more measurements $\boldsymbol{y}$
  - wireless Tx localization: gather more pilot sequences
  - accelerated MRI: gather more k-space samples

- But often there's a cost to collecting more measurements

- Proposed idea: Collect measurements until the accuracy lower-bound surpasses a threshold

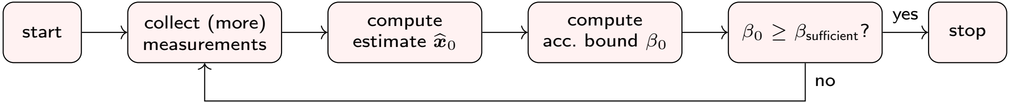


- Applied to MRI with $1-\alpha = 0.95$:

| method | avg acceleration | empirical coverage |
|---|---|---|
| single-round | $2.000 \pm 0.0000$ | $0.9505 \pm 0.0001$ |
| multi-round | $5.422 \pm 0.0001$ | $0.9461 \pm 0.0001$ |

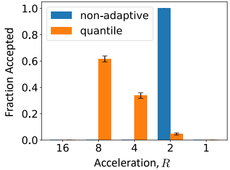

# Back to wireless transmitter localization

- Given the nature of the posterior, we may want to separately bound the angle and the radius
- This gives multiple accuracies to bound:

$$\boldsymbol{z}_0 = \begin{bmatrix} -\big|\angle(\widehat{\boldsymbol{x}}_0) - \angle(\boldsymbol{x}_0)\big| \\ -\big|\|\widehat{\boldsymbol{x}}_0\| - \|\boldsymbol{x}_0\|\big| \end{bmatrix} \in \mathbb{R}^2$$

and requires extending our conformal bounding method to multiple "targets"

Beyond Point Estimates: Likelihood-Based Full-Posterior Wireless Localization

Haozhe Lei[*1], Hao Guo[1,2], Tommy Svensson[2], and Sundeep Rangan[1]

arXiv:2509.25719

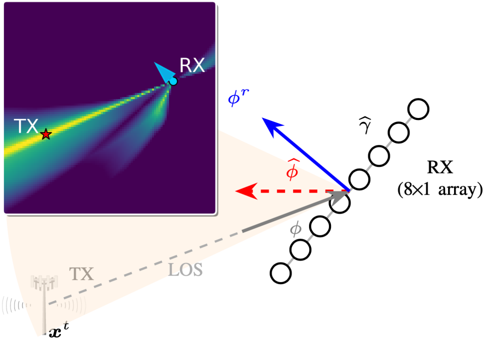

# Multi-target conformal prediction

- Say we have $K > 1$ scalar accuracies $\{z_{0,k}\}_{k=1}^K$ and we want to lower-bound each of them
- Our goal is to compute bounds $\{\beta_{0,k}\}_{k=1}^K$ that guarantee joint marginal coverage:

$$\Pr\left\{\bigcap_{k=1}^K Z_{0,k} \geq \beta_{0,k}(\boldsymbol{Y}_0, D_{\mathsf{cal}})\right\} \geq 1-\alpha,$$

  which means that, with probability at least $1-\alpha$, all bounds are *simultaneously* valid
- Several existing methods[1] accomplish this task, but they struggle with uniformity across targets: bounds for some targets are too loose while those for other targets are unnecessarily tight

[1] Messoudi et al.'20, Sampson&Chan'24, Sun&Yu'24, Park&Cho'25

# Proposed multi-target conformal prediction

- Since looser bounds correspond to larger *single*-target coverages $\Pr\{Z_{0,k} \geq \beta_{0,k}\}$, we propose to solve
$$\underset{(\beta_{0,1},\ldots,\beta_{0,K})}{\arg\min} \max_k \Pr\{Z_{0,k} \geq \beta_{0,k}\} \text{ s.t. } \Pr\left\{\cap_{k=1}^K Z_{0,k} \geq \beta_{0,k}\right\} \geq 1-\alpha$$
i.e., minimize the maximum single-target coverage subject to the joint-coverage constraint

- Forming the $k$th bound as $\beta_{0,k} = \widehat{\beta}_{0,k}(\boldsymbol{y}_0) - \lambda_k$ and bound-violation score as $s_{0,k} = \widehat{\beta}_{0,k}(\boldsymbol{y}_0) - z_{0,k}$, we have
$$\Pr\{Z_{0,k} \geq \beta_{0,k}\} = \Pr\{S_{0,k} \leq \lambda_k\} = F_{S_{0,k}}(\lambda_k) \text{ with CDF } F_{S_{0,k}}(\cdot)$$
$$\& \ \Pr\left\{\cap_{k=1}^K Z_{0,k} \geq \beta_{0,k}\right\} = \Pr\left\{\cap_{k=1}^K S_{0,k} \leq \lambda_k\right\} = \Pr\left\{\cap_{k=1}^K F_{S_{0,k}}(S_{0,k}) \leq F_{S_{0,k}}(\lambda_k)\right\}$$

- Thus the minimax design problem can be rephrased as
$$\min_{(\lambda_1,\ldots,\lambda_K)} \max_k F_{S_{0,k}}(\lambda_k) \text{ s.t. } \Pr\left\{\cap_{k=1}^K F_{S_{0,k}}(S_{0,k}) \leq F_{S_{0,k}}(\lambda_k)\right\} \geq 1-\alpha$$
$$\Leftrightarrow \min_{(\zeta_1,\ldots,\zeta_K)} \max_k \zeta_k \text{ s.t. } \Pr\left\{\cap_{k=1}^K F_{S_{0,k}}(S_{0,k}) \leq \zeta_k\right\} \geq 1-\alpha \quad \text{via } \zeta_k \triangleq F_{S_{0,k}}(\lambda_k)$$
$$\Leftrightarrow \min_{\zeta} \zeta \text{ s.t. } \Pr\left\{\cap_{k=1}^K F_{S_{0,k}}(S_{0,k}) \leq \zeta\right\} \geq 1-\alpha \quad \text{via } \zeta \triangleq \max_k \zeta_k$$

# Proposed multi-target conformal prediction (cont.)

- If we knew the joint statistics of the bound violation scores $\boldsymbol{S}_0 \triangleq [S_{0,1}, \ldots, S_{0,K}]$, we could solve for
$$\zeta_* = \arg\min_{\zeta} \ \zeta \ \text{ s.t. } \Pr\left\{ \cap_{k=1}^{K} F_{S_{0,k}}(S_{0,k}) \leq \zeta \right\} \geq 1-\alpha$$
and then form the $k$th bound as $\beta_{0,k} = \widehat{\beta}_{0,k}(\boldsymbol{y}_0) - \lambda_k$ with $\lambda_k = F_{S_{0,k}}^{-1}(\zeta_*)$

- But we don't know the statistics of $\boldsymbol{S}_0$. So we propose to compute an empirical CDF $\widehat{F}_{S_k}(\cdot)$ using a tuning set $d_{\mathsf{tune},k} = \{s_{i,k}\}_{i=n_{\mathsf{cal}}+1}^{n_{\mathsf{cal}}+n_{\mathsf{tune}}}$ of size $n_{\mathsf{tune}}$

- Next we compute transformed calibration scores $u_{i,k} \triangleq \widehat{F}_{S_k}(s_{i,k})$ from $d_{\mathsf{cal},k} = \{s_{i,k}\}_{i=1}^{n_{\mathsf{cal}}}$, and then
$$\widehat{\zeta}(d_{\mathsf{cal}}) = \mathrm{EmpQuant}\left(\tfrac{\lceil (1-\alpha)(n_{\mathsf{cal}}+1)\rceil}{n_{\mathsf{cal}}}; \{u_i\}_{i=1}^{n_{\mathsf{cal}}}\right) \text{ with } u_i \triangleq \max_k u_{i,k},$$
and finally set the $k$th bound as $\beta_{0,k} = \widehat{\beta}_{0,k}(\boldsymbol{y}_0) - \widehat{\lambda}_k$ with $\widehat{\lambda}_k = \widehat{F}_{S_k}^{-1}(\widehat{\zeta}(d_{\mathsf{cal}}))$

# Proposed multi-target conformal prediction (cont.)

We prove[1] the following about our proposed scheme:

- For any chosen error-rate $\alpha \in (0,1)$, it guarantees joint marginal coverage, i.e.,

$$\Pr\left\{ \bigcap_{k=1}^{K} Z_{0,k} \geq \beta_{0,k}(\boldsymbol{Y}_0, D_{\text{cal}}) \right\} \geq 1-\alpha,$$

  as long as $\{\boldsymbol{S}_0, \boldsymbol{S}_1, \dots, \boldsymbol{S}_{n_{\text{cal}}}\}$ are statistically exchangeable

- It is asymptotically minimax in that

$$\widehat{\zeta}(d_{\text{cal}}) \xrightarrow{a.s.} \zeta_* \quad \text{as} \quad n_{\text{cal}}, n_{\text{tune}} \to \infty$$

[1] Wen,Ahmad,Schniter'25b

# Bounding multiple accuracy metrics in accelerated MRI

- Say we want to guarantee the performance of accelerated MRI simultaneously in PSNR, SSIM, LPIPS, and DISTS metrics (i.e., $K = 4$ targets)
- Among existing methods that guarantee joint marginal coverage, ours (green) provides tighter bounds across $k$:

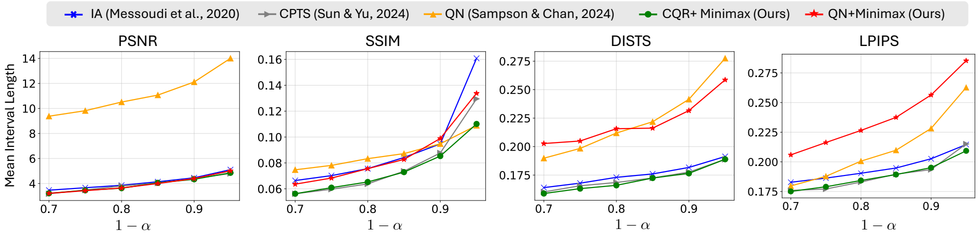

# Conclusion

- Due to the ill-posedness of many inverse problems, there's a need to bound estimator accuracy

- By combining approximate posterior sampling with conformal prediction, we proposed accuracy lower-bounds $\beta(\cdot)$ with probabilistic guarantees of the form

$$\Pr\left\{m(\widehat{\boldsymbol{x}}, \boldsymbol{x}_0) \leq \beta(\boldsymbol{y}, d_{\textsf{cal}}, \alpha)\right\} \geq 1 - \alpha$$

  that allow arbitrary estimators $\boldsymbol{r}(\cdot)$, accuracy metrics $m(\cdot)$, and error-rates $\alpha \in (0,1)$, assuming exchangeabile test & calibration scores

- Although our prior work focused on MRI imaging, the techniques are directly applicable to communications problems like wireless device localization, CSI estimation, and RF tomography