Task-Driven Uncertainty Quantification in Inverse Problems via Conformal Prediction

Jeffrey Wen, Rizwan Ahmad, and Phil Schniter

THE OHIO STATE UNIVERSITY

Supported by NIH grant R01-EB029957

2025 BASP Frontiers Conference
Villars-sur-Ollon, Switzerland

Imaging inverse problems

- Measurements $\boldsymbol{y} = \mathcal{A}(\boldsymbol{x}_0)$ of unknown image $\boldsymbol{x}_0$
  - $\mathcal{A}(\cdot)$ masks, distorts, and/or corrupts $\boldsymbol{x}_0$ with noise.
  - Examples: denoising, deblurring, inpainting, super-resolution, phase retrieval, computed tomography (CT), magnetic resonance imaging (MRI), etc.
- Typical goals
  - Recover image $\boldsymbol{x}_0$
  - Extract quantitative information from $\boldsymbol{x}_0$ (e.g., probability of a pathology)
- Challenges
  - Ill-posed: Many hypotheses of $\boldsymbol{x}_0$ can explain $\boldsymbol{y}$
  - Hallucinations: Inaccurate $\widehat{\boldsymbol{x}}$ can still look "good," leading to false sense of trust

Uncertainty quantification (UQ)

- We'd like to quantify the uncertainty or error in the image recovery $\widehat{\boldsymbol{x}}$
- Especially important in safety-critical applications (e.g., medical imaging)
- Most UQ methods produce pixel-wise uncertainty maps like

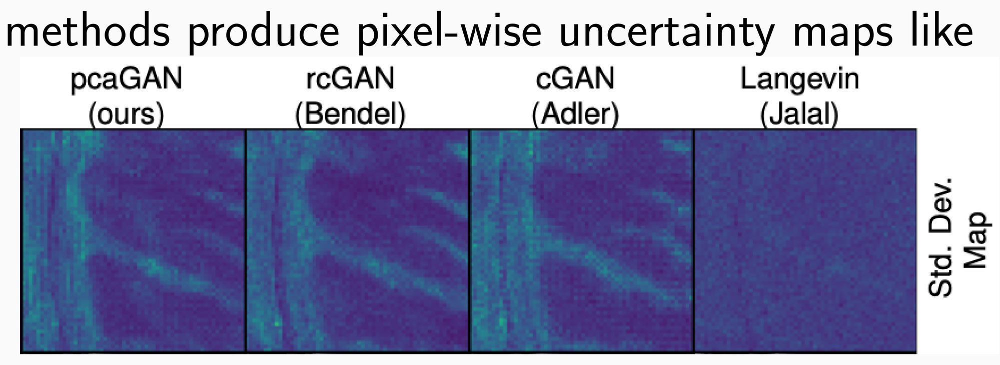

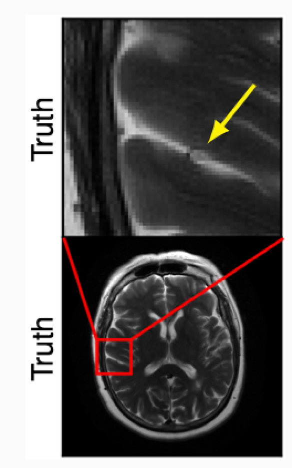

- But how useful are they?
  - Do they say when the recovery is accurate enough for the task at hand?

Probabilistic bounds on recovered-image accuracy

- Say we have an image-recovery method $\boldsymbol{r}(\cdot)$ that produces $\widehat{\boldsymbol{x}}_0 = \boldsymbol{r}(\boldsymbol{y}_0)$
- We'd like to know the accuracy' of $\widehat{\boldsymbol{x}}_0$ relative to the true $\boldsymbol{x}_0$
- To measure accuracy, we'll use an arbitrary image-quality metric $\boxed{z_0 = m(\widehat{\boldsymbol{x}}_0, \boldsymbol{x}_0)}$ like
  - PSNR, SSIM …higher preferred (we'll focus on this below)
  - LPIPS [1], DISTS [2] …lower preferred
- Can we guarantee the accuracy of $\widehat{\boldsymbol{x}}_0$? Can we construct a bound $\beta_0(\boldsymbol{y}_0)$ such that

$$\Pr\{Z_0 \geq \beta_0(\boldsymbol{Y}_0)\} \geq 1-\alpha \quad \text{for some chosen error rate } \alpha?$$

An impractical bound

- Say we have a perfect posterior sampler generating c i.i.d image samples $\{\widetilde{\boldsymbol{x}}_0^{(j)}\}_{j=1}^c$
- The corresponding image-accuracy samples are $\widetilde{z}_0^{(j)} \triangleq m(\widehat{\boldsymbol{x}}_0, \widetilde{\boldsymbol{x}}_0^{(j)}) \overset{iid}{\sim} p_{Z_0|\boldsymbol{Y}_0}(\cdot\,|\,\boldsymbol{y}_0)$
- An accuracy lower bound β_0 that obeys

$$\Pr\{Z_0 \geq \beta_0 \,|\, \boldsymbol{Y}_0 = \boldsymbol{y}_0\} = 1-\alpha$$

can be constructed using an infinite number of perfect posterior samples:

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

A conformal image-accuracy bound

- In practice, we have only a finite number c of imperfect posterior samples
- We propose to design a valid lower bound using conformal prediction [3]
  - Assume we have n calibration samples $\{(\boldsymbol{x}_i, \boldsymbol{y}_i)\}_{i=1}^n$ in addition to the test measurements $\boldsymbol{y}_0$
  - Construct approximate c-sample bounds: $\widehat{\beta}_i = \mathrm{EmpQuant}(\alpha, \{\widetilde{z}_i^{(j)}\}_{j=1}^c)$ for $i = 0, \dots, n$ and "bound-violation scores" $s_i \triangleq \widehat{\beta}_i - z_i$ for $i = 1, \dots, n$
  - Using the set $d_{\text{cal}} \triangleq \{s_i\}_{i=1}^n$ of calibration scores, compute a bound correction term

$$\widehat{\lambda}(d_{\text{cal}}) = \mathrm{EmpQuant}\left(\tfrac{\lceil(1-\alpha)(n+1)\rceil}{n}, \{s_i\}_{i=1}^n\right)$$

  - Form the final "test" lower bound as $\beta_0(\boldsymbol{y}_0, d_{\text{cal}}) = \widehat{\beta}_0 - \widehat{\lambda}(d_{\text{cal}})$
- This conformal bound obeys the marginal coverage guarantee

$$1-\alpha \leq \Pr\{Z_0 \geq \beta_0(\boldsymbol{Y}_0, D_{\text{cal}})\} \leq 1-\alpha+\tfrac{1}{n+1}$$

assuming that $\{S_0, S_1, \dots, S_n\}$ are statistically exchangeable

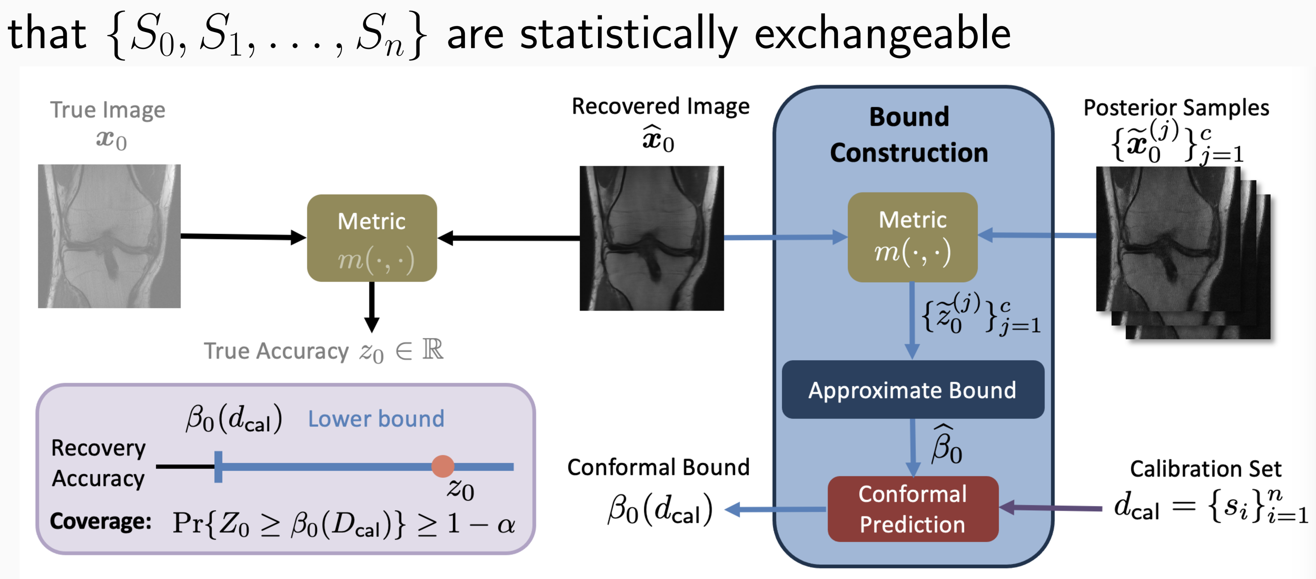

Example: Bounding recovery accuracy in MRI

- Scatter plots of (z_0, β_0) from fastMRI knee recovery @ acceleration $R = 8$ using a conditional normalizing flow:

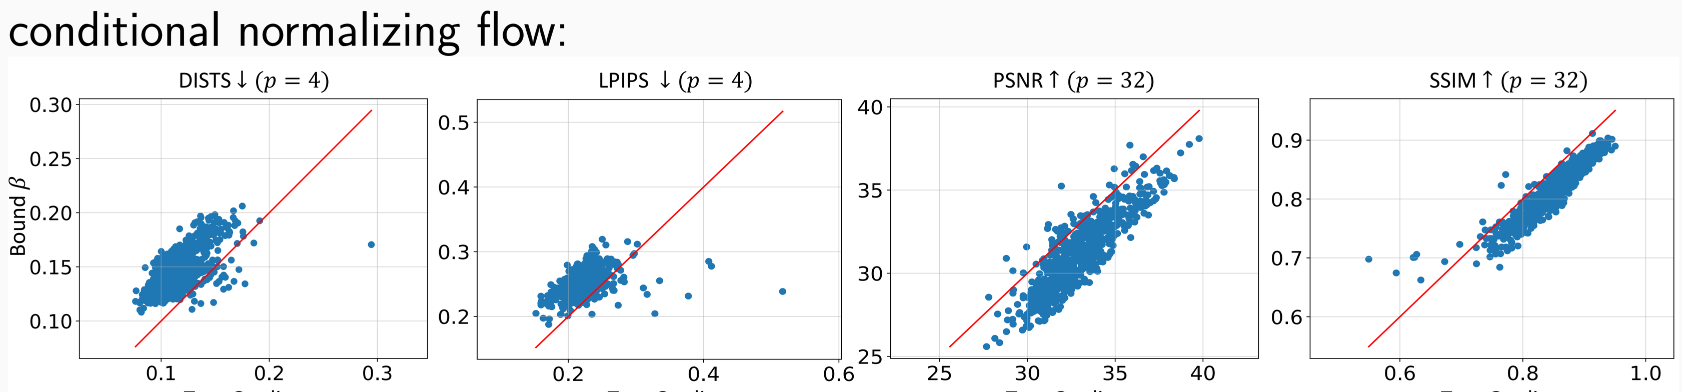

The red line indicates where the bound would be exact

- Marginal coverage validation via 10 000 Monte-Carlo trials (random 70% test / 30% calibration split):

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

Application: Multi-round MRI acquisition

- Consider acquiring over multiple rounds (i.e., $R \in \{16, 8, 4, 2, 1\}$), stopping as soon as the conformal upper bound on DISTS [4] is good enough (i.e., $\leq \beta_{\max}$)

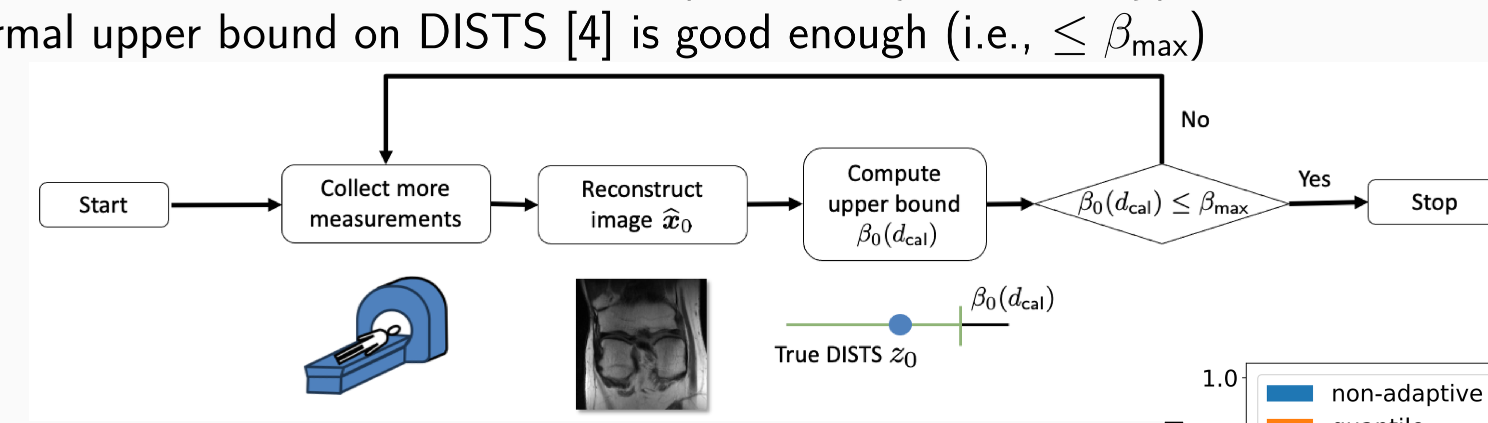

- Multi-round acquisition achieves much higher average acceleration than single-round acquisition:

method	avg acceleration	empirical coverage
single-round	2.000 ± 0.0000	0.9505 ± 0.0001
multi-round	5.422 ± 0.0001	0.9461 ± 0.0001

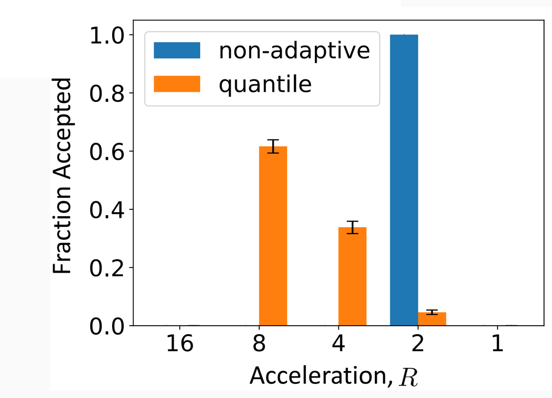

Quantitative / Task-based imaging

- Again consider measurements $\boldsymbol{y}_0 = \mathcal{A}(\boldsymbol{x}_0)$ and some recovery $\widehat{\boldsymbol{x}}_0 = r(\boldsymbol{y}_0)$
- But now say that our goal is to extract quantitative information about $\boldsymbol{x}_0$
- Example: Does the MRI knee image $\boldsymbol{x}_0$ indicate a meniscus tear?
  - Say we've trained & calibrated a soft-output binary classifier $f(\cdot)$ on clean images
  - Naively applying $f(\cdot)$ to imperfect recoveries $\widehat{\boldsymbol{x}}_0$ would give unreliable results
  - Instead, we want to estimate $f(\boldsymbol{x}_0)$ from $\boldsymbol{y}_0$ (without knowing $\boldsymbol{x}_0$)
- More generally, one may wish to estimate a generic $\boxed{z_0 = f(\boldsymbol{x}_0) \in \mathbb{R}}$ given $\boldsymbol{y}_0$
- Can one construct guaranteed upper and lower bounds on z_0?

Conformal prediction of the true task output [5]

- Our approach is similar to before, in that we combine posterior image sampling with conformal prediction
- But instead of a one-sided bound, we construct a prediction interval $C_\lambda = [\beta_l, \beta_u]$ that is guaranteed to contain the true task output z_0 with probability $\geq 1-\alpha$

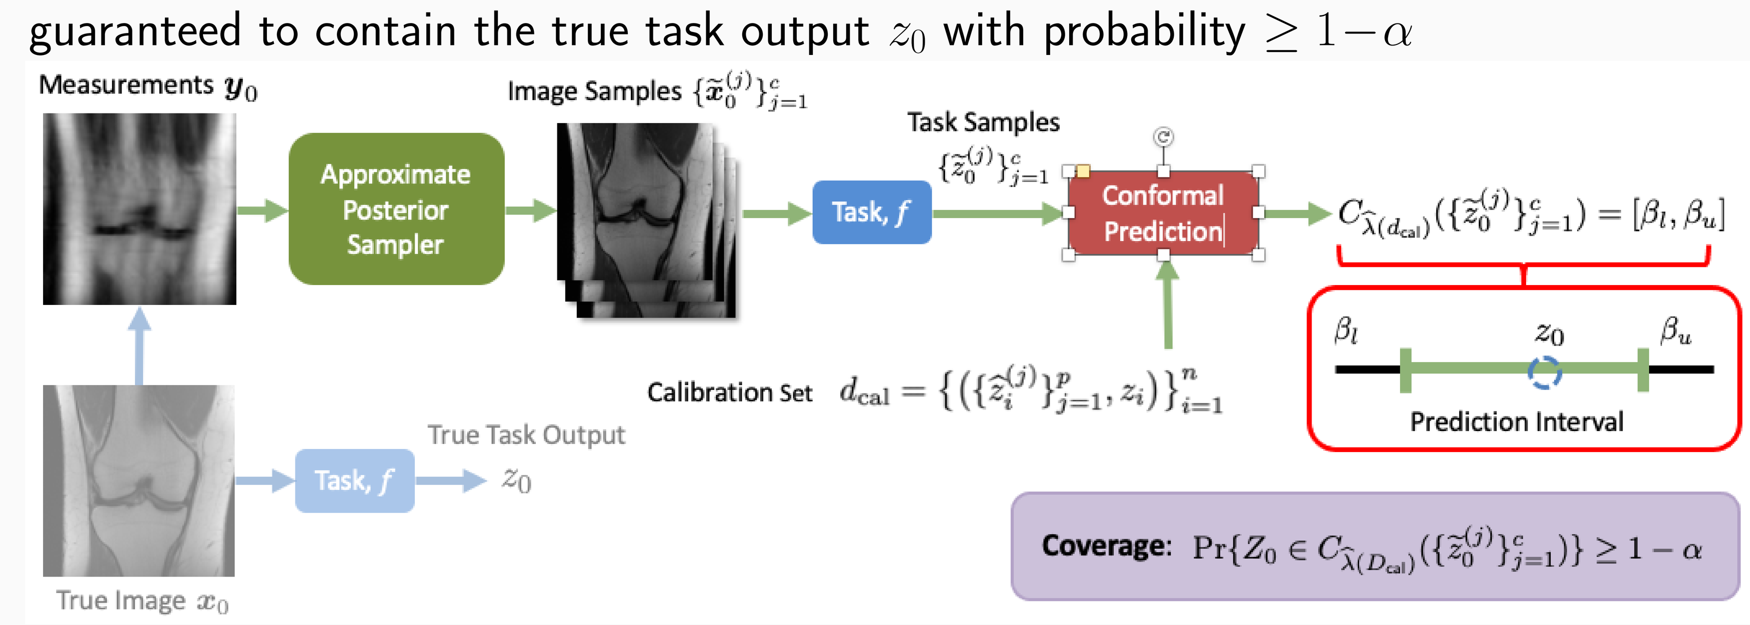

Example: Predicting meniscus tears in knee MRI

- We trained and calibrated a ResNet50 to output meniscus-tear probability $z_0 = f(\boldsymbol{x}_0)$ using clean images from fastMRI+
- From R-accelerated measurements $\boldsymbol{y}_0$, we compute a prediction interval C_λ that contains the true z_0 with 99% probability
- The conformal bound uses c posterior samples from a conditional normalizing flow

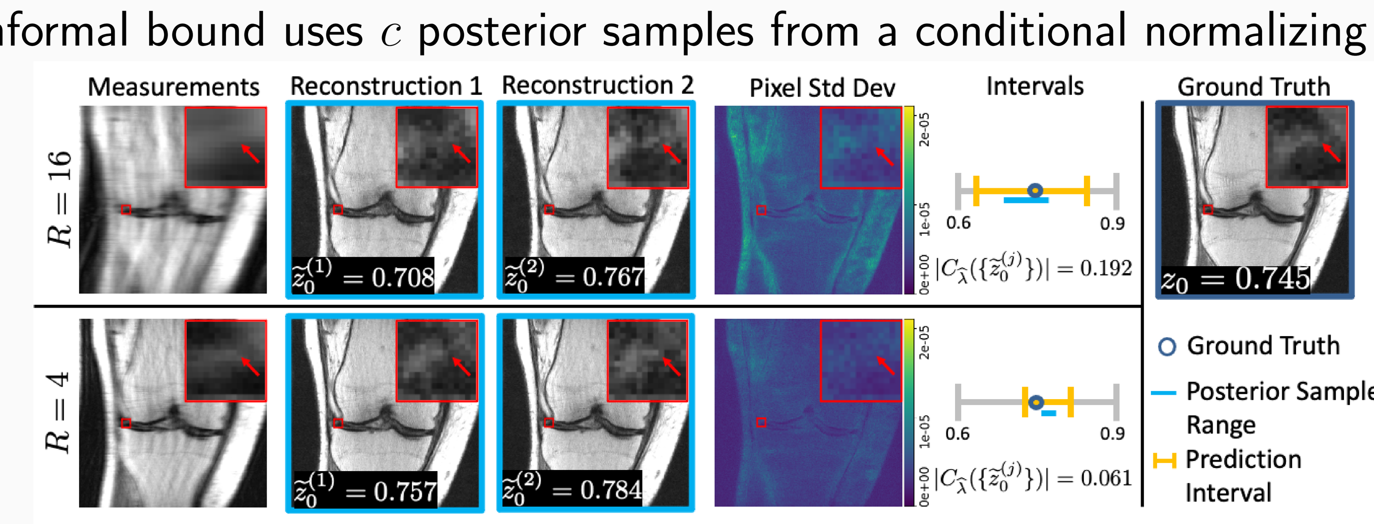

Application: Multi-round MRI acquisition

- Consider acquiring over multiple rounds (i.e., $R \in \{16, 8, 4, 2, 1\}$), stopping as soon as the task uncertainty is small enough ($|C_\lambda| \leq \tau$ for $\tau = 0.1$)

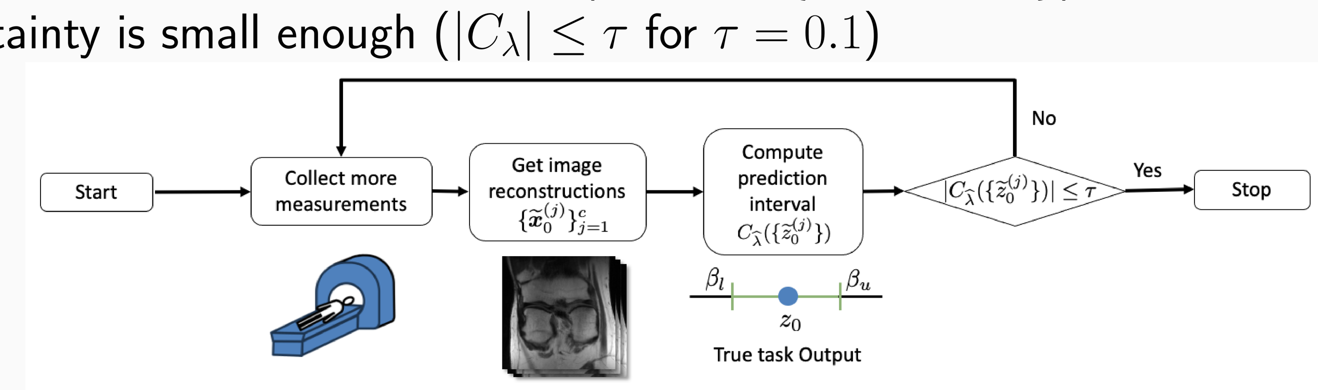

- The adaptive LWR and CQR schemes achieve much higher average acceleration rates than the non-adaptive AR scheme:

Method	Average Acceleration	Empirical Coverage	Average Center Error @ $R = 2$
AR	2.000	0.991 ± 0.008	0.032 ± 0.017
LWR	5.157	0.992 ± 0.005	0.020 ± 0.002
CQR	6.762	0.987 ± 0.008	0.044 ± 0.009

Center Error $\triangleq |z_0 - \frac{\beta_l + \beta_u}{2}|$

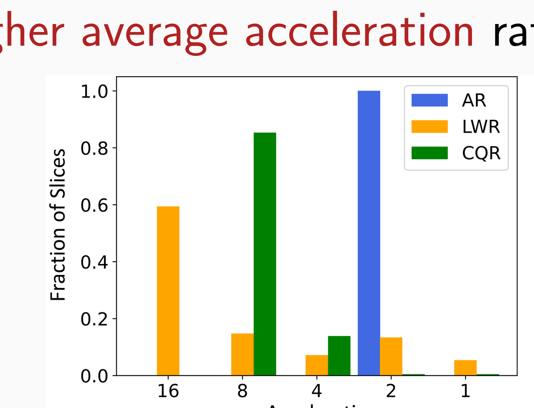

References

[1] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proc. IEEE Conf. Comp. Vision Pattern Recog.*, pp. 586–595, 2018.

[2] K. Ding, K. Ma, S. Wang, and E. P. Simoncelli, "Image quality assessment: Unifying structure and texture similarity," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 5, pp. 2567–2581, 2020.

[3] J. Lei, M. G'Sell, A. Rinaldo, R. J. Tibshirani, and L. Wasserman, "Distribution-free predictive inference for regression," *J. Am. Statist. Assoc.*, 2018.

[4] S. Kastryulin, J. Zakirov, N. Pezzotti, and D. V. Dylov, "Image quality assessment for magnetic resonance imaging," *IEEE Access*, vol. 11, pp. 14154–14168, 2023.

[5] J. Wen, R. Ahmad, and P. Schniter, "Task-driven uncertainty quantification in inverse problems via conformal prediction," in *Proc. Europ. Conf. Comp. Vision*, 2024.