Weak Lensing Mass Mapping with Uncertainty Quantification

Hubert Leterme, Postdoc
Affiliated to GREYC CNRS-Ensicaen (Caen, France)
Co-supervised at CosmoStat, CEA DAp
Joint ARGOS-TITAN-TOSCA workshop
2nd February 2024

Context and objectives

- **Gravitational lensing**: light rays emitted by distant objects (e.g., galaxies) are deflected by inhomogeneous matter density in the foreground.
- Dark matter cannot be detected by direct observations, but mass mapping can be performed from gravitational lensing observations.
- **Weak lensing regime**: two types of deformation:
  - **convergence**: isotropic dilation of the source;
  - **shear**: anisotropic stretching of the image.

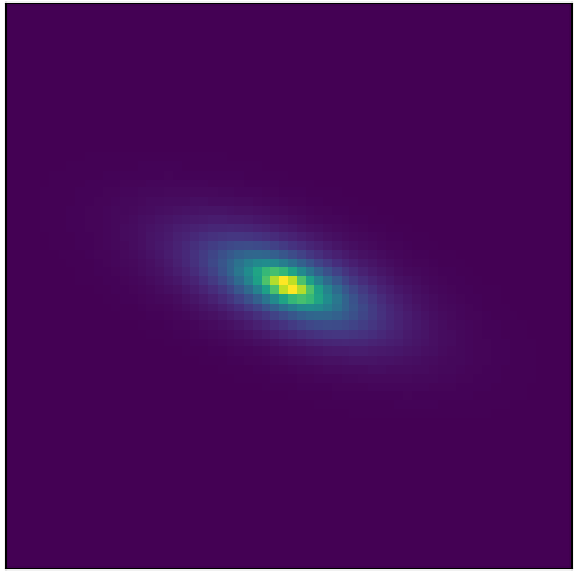

Source galaxy, unlensed

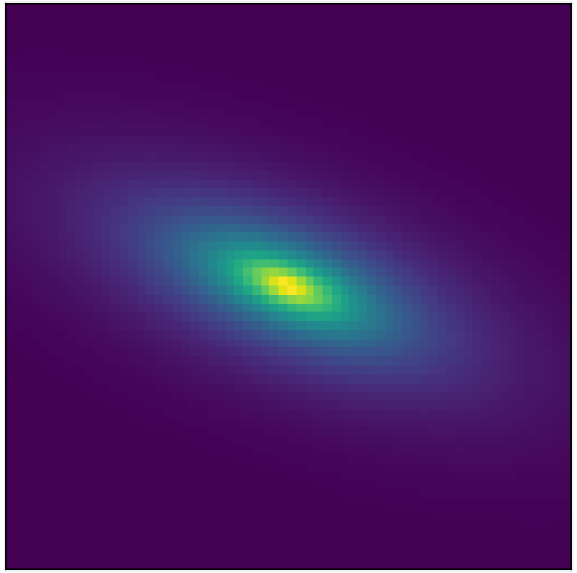

Convergence only
$\kappa = 1$

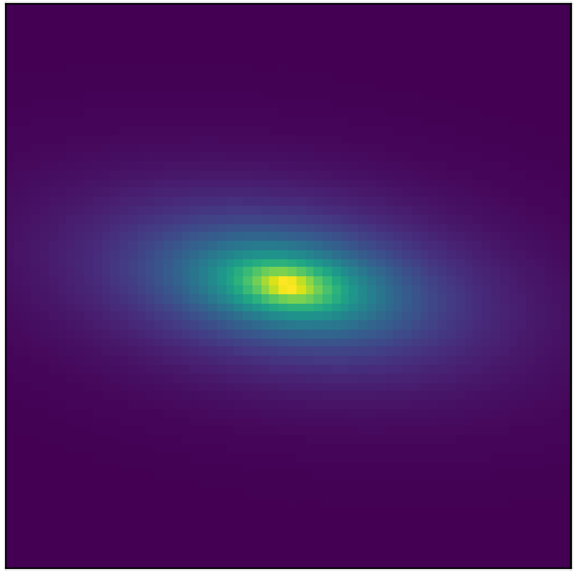

Convergence + shear
$\kappa = 1$ and $\gamma = (0.1 - 0.3\, i)$

Context and objectives

- Convergence map $\boldsymbol{\kappa} \in \mathbb{R}^K$ at a given redshift z: proportional to the projected mass along the line of sight ⇒ **variable of interest**.
- However, $\boldsymbol{\kappa}$ cannot be directly measured.
- Relationship between shear and convergence maps: $\boldsymbol{\gamma} = \mathbf{A}\boldsymbol{\kappa}$, with
  - $\boldsymbol{\gamma} \in \mathbb{C}^K$ true shear map (unknown);
  - $\mathbf{A} \in \mathbb{R}^{K \times K}$ Kaiser-Squires filter (known).

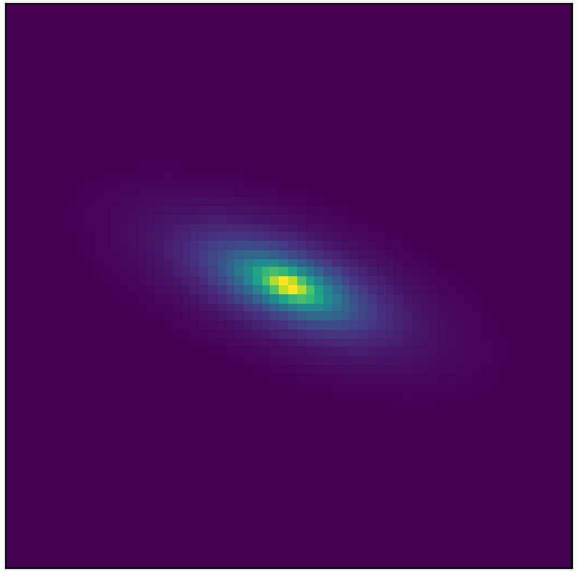

Source galaxy, unlensed

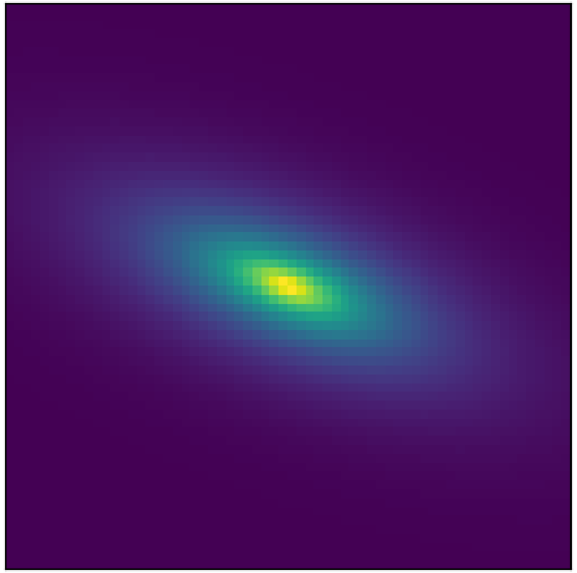

Convergence only
$\kappa = 1$

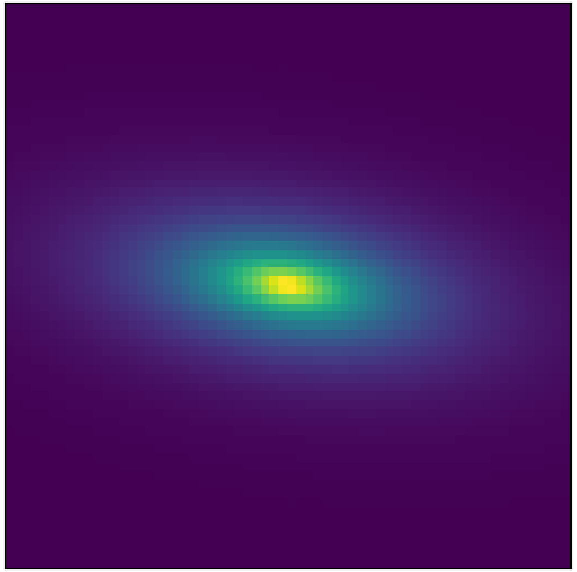

Convergence + shear
$\kappa = 1$ and $\gamma = (0.1 - 0.3\, i)$

Context and objectives

Example with the κTNG simulated dataset[1]

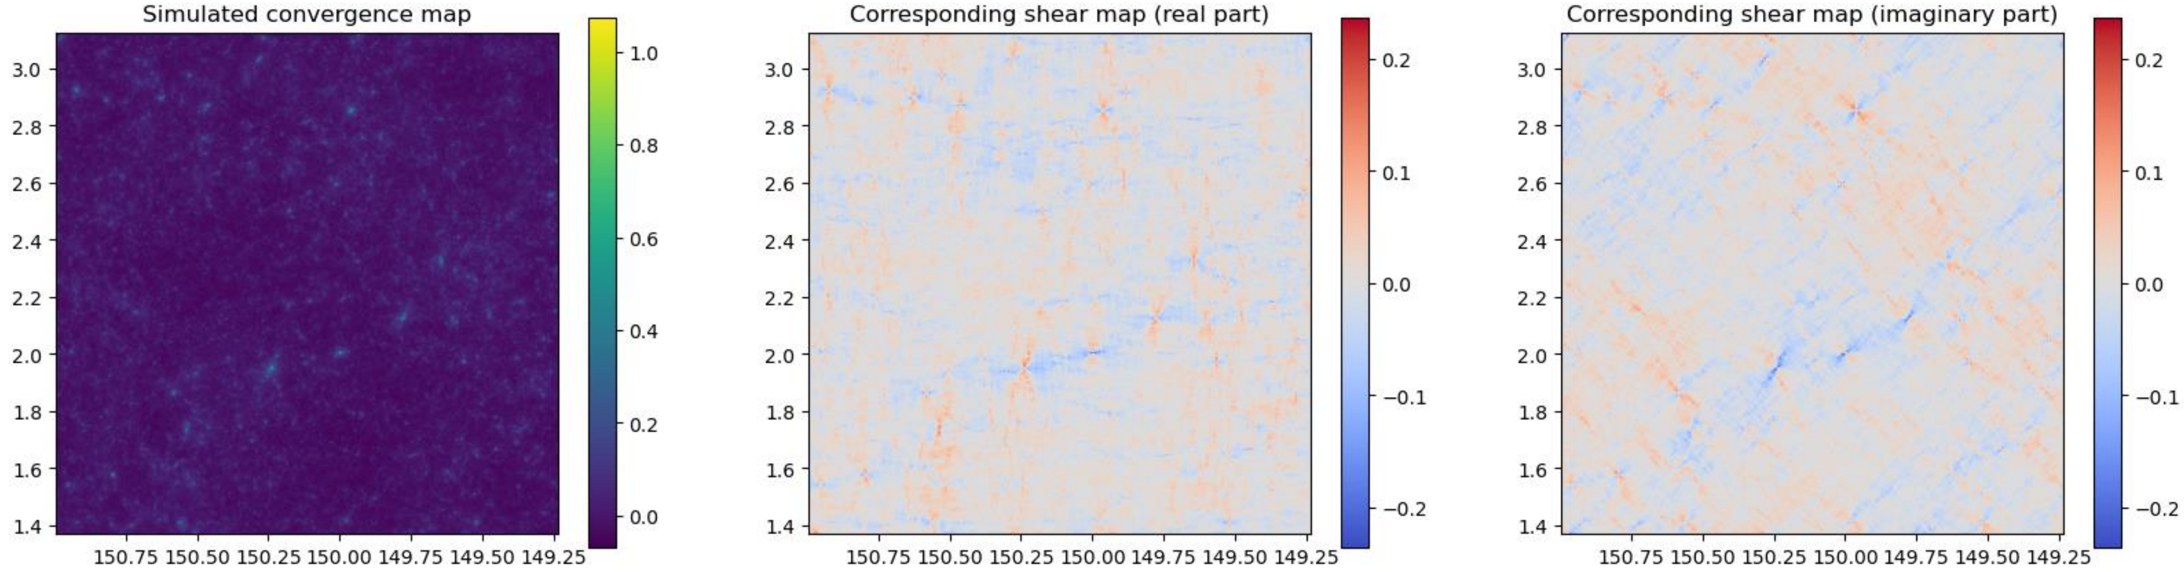

- As for $\boldsymbol{\kappa}$, the true shear map $\boldsymbol{\gamma}$ cannot be directly measured.
- Unbiased estimator of $\boldsymbol{\gamma}$, denoted by $\widehat{\boldsymbol{\gamma}}$, obtained by measuring galaxy ellipticities.
- Relation between $\widehat{\boldsymbol{\gamma}}$ (observable) and $\boldsymbol{\kappa}$ (quantity of interest):

$$\widehat{\boldsymbol{\gamma}} = \mathbf{A}\boldsymbol{\kappa} + \boldsymbol{n}.$$

- Level of noise: depends on the number N_k of observed galaxies at a given pixel k.

[1] http://columbialensing.org/

Context and objectives

Number of measured galaxies per pixel + mask (from the COSMOS shape catalog[1]):

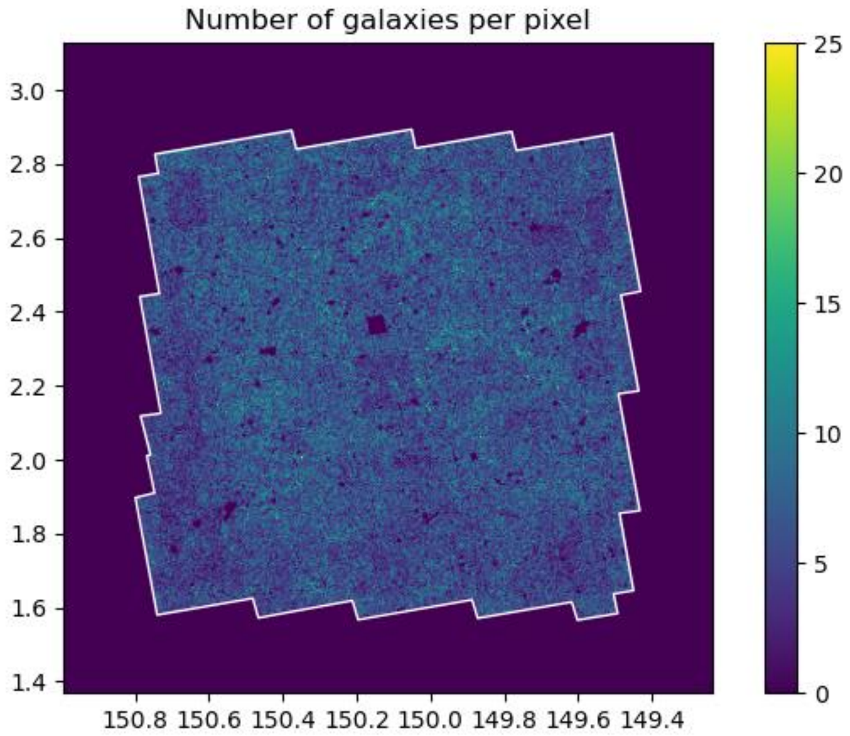

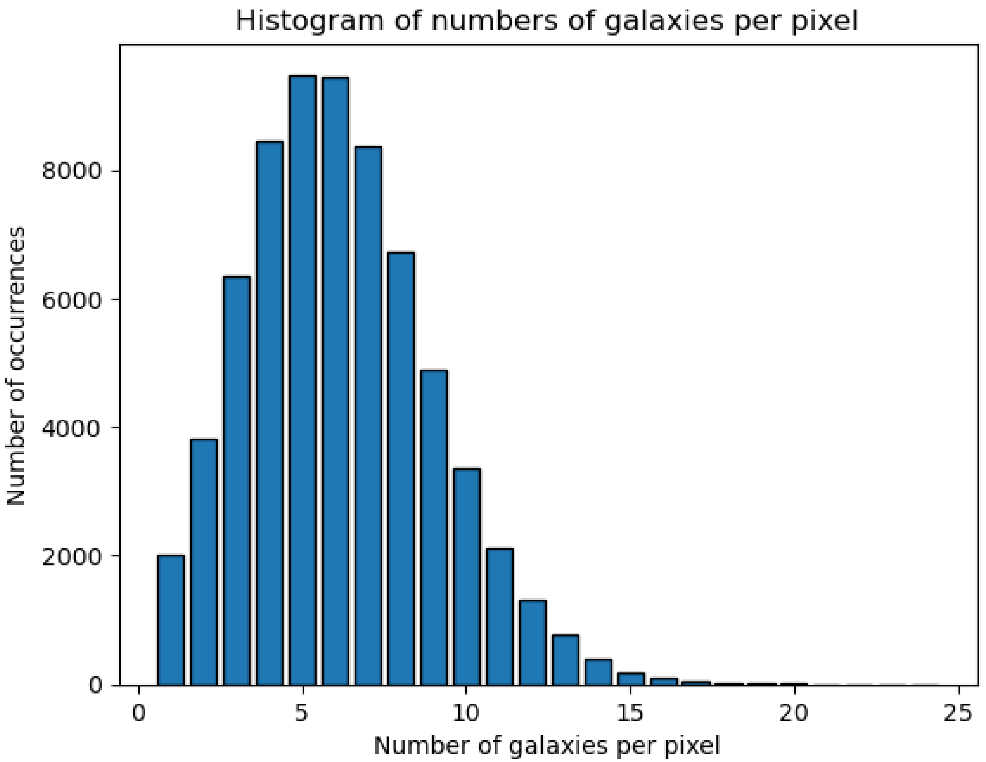

- At pixel k, $\sigma_k \propto \frac{1}{\sqrt{N_k}}$ (N_k number of measured galaxies).
- Masked data: noise set to an arbitrary large value.

[1] https://astro.uni-bonn.de/en/m/schrabba/research

Context and objectives

Noisy shear maps (ellipticities)

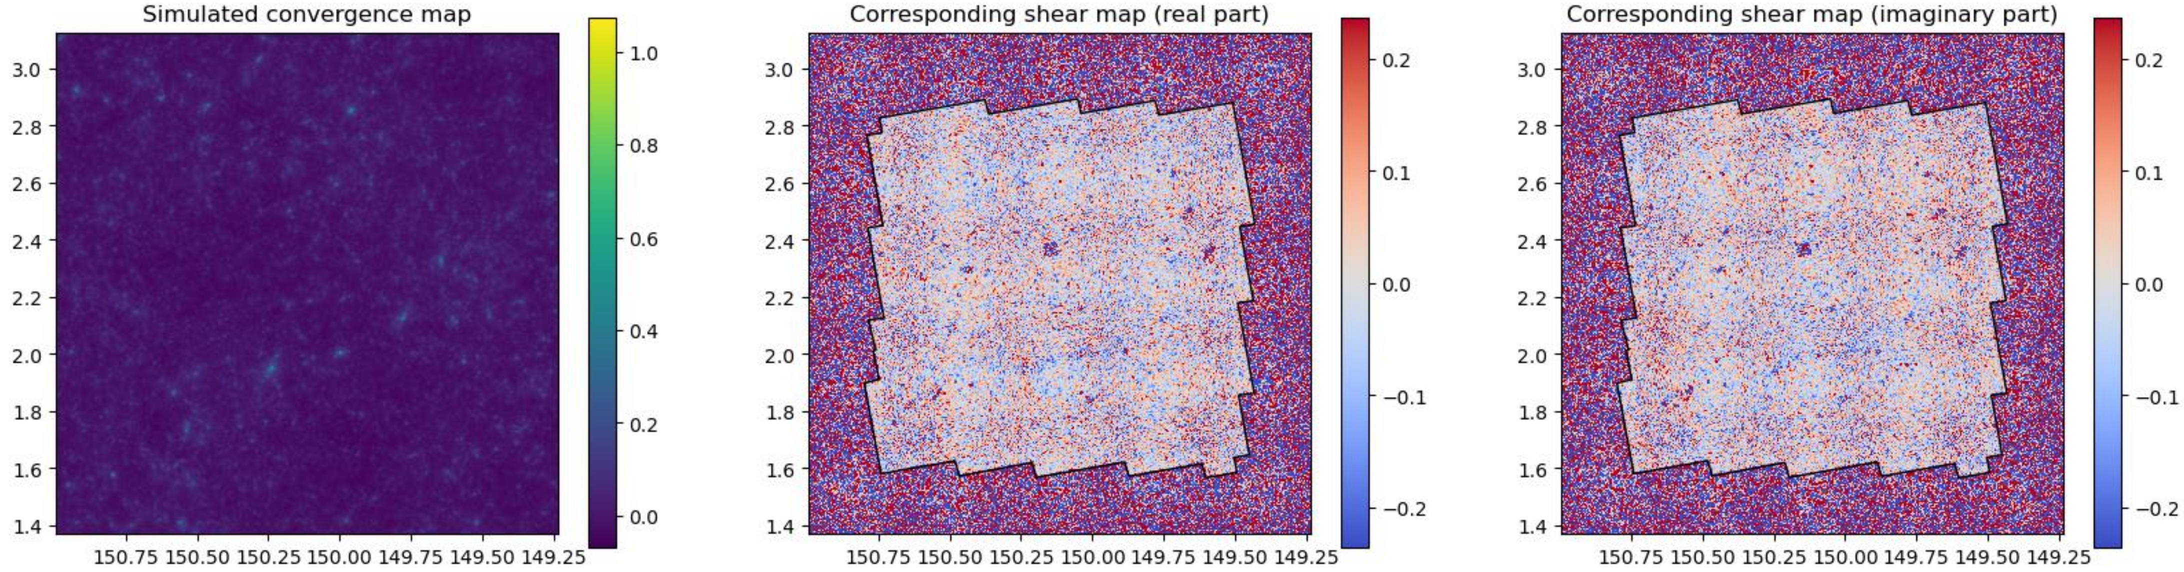

Objective: given $\hat{\boldsymbol{\gamma}}$, estimate $\hat{\boldsymbol{\kappa}}^-$ and $\hat{\boldsymbol{\kappa}}^+$ such that

$$\mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^-[k], \hat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha$$

for a given confidence level $\alpha \in \,]0,1[$, for any pixel k.

Requirement: we seek a “fast” mass mapping method (i.e., non-iterative).

Proposed approach

1. Compute heuristic bounds $\hat{\boldsymbol{\kappa}}^{-}$ and $\hat{\boldsymbol{\kappa}}^{+}$ (Kaiser-Quires, MSE solution, Wiener solution, plug-and-play algorithms, Glimpse, MCAlens…).
2. Post-processing: adjust bounds $\hat{\boldsymbol{\kappa}}^{-}$ and $\hat{\boldsymbol{\kappa}}^{+}$ using a **calibration set**.

→ Distribution-free UQ, does not assume any prior distribution.

→ Works for any heuristic prediction method, including deep learning.

A simple example: the Kaiser-Squires solution

- Exact solution if $\hat{\boldsymbol{\gamma}} = \boldsymbol{\gamma}$ (no noise, no mask): $\hat{\boldsymbol{\kappa}} = \mathbf{A}^{\dagger}\hat{\boldsymbol{\gamma}}$.
- In practice, the KS filter is followed by Gaussian smoothing: $\hat{\boldsymbol{\kappa}} = \mathbf{S}\mathbf{A}^{\dagger}\hat{\boldsymbol{\gamma}}$.

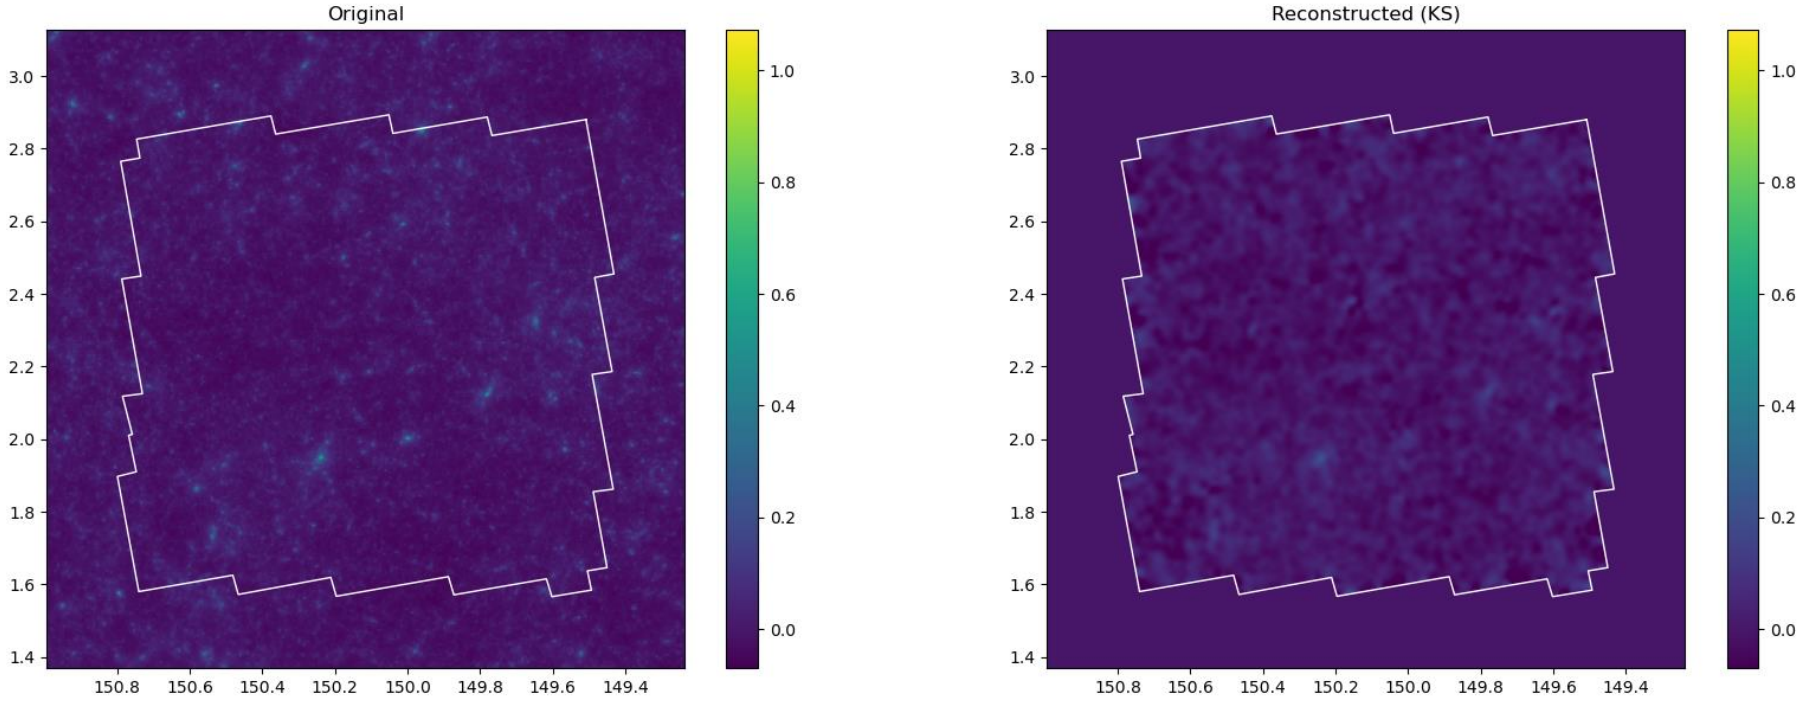

Kaiser-Squires bound estimation

Estimation of $\widehat{\boldsymbol{\kappa}}^-$ and $\widehat{\boldsymbol{\kappa}}^+$: we have $\widehat{\boldsymbol{\kappa}} \sim \boldsymbol{N}(\boldsymbol{S\kappa}, \mathbf{S}\mathbf{A}^\dagger\boldsymbol{\Sigma}\mathbf{A}^{\dagger *}\boldsymbol{S}^*)$.

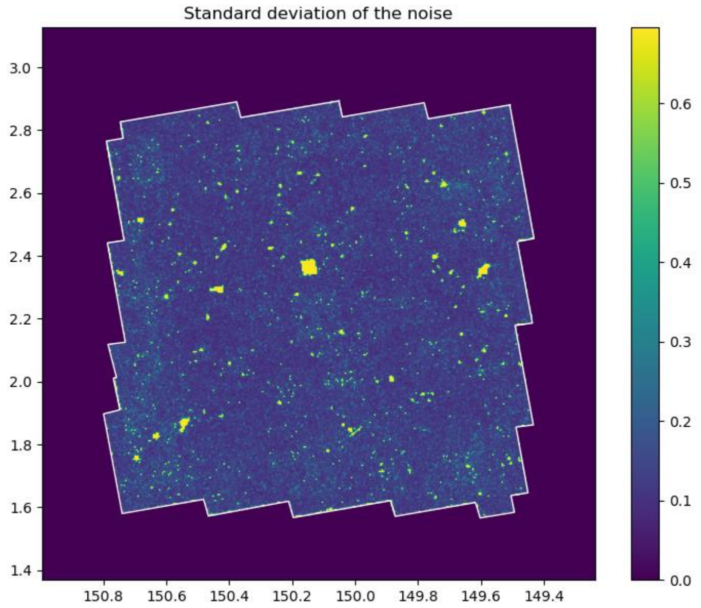

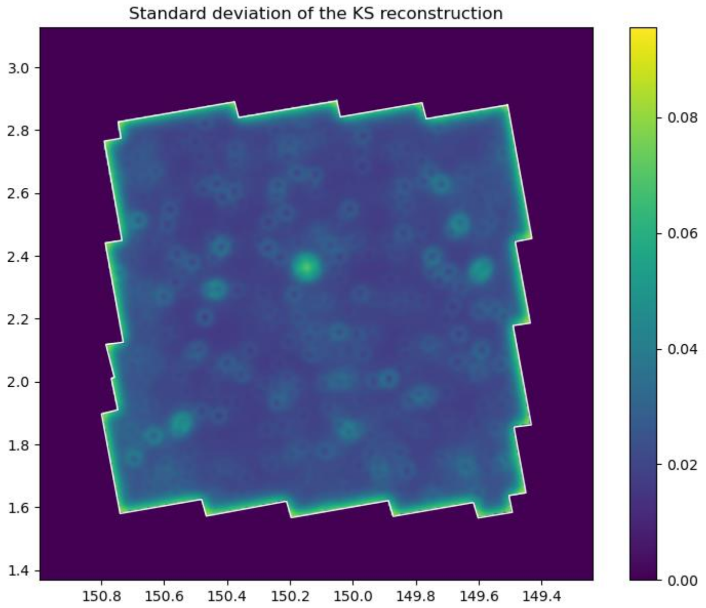

- $\boldsymbol{r}[k] \coloneqq \Phi_k^{-1}(1 - {}^{\alpha}/_{2})$ with Φ_k Gaussian CDF at pixel k;
- $\widehat{\boldsymbol{\kappa}}^- \coloneqq \widehat{\boldsymbol{\kappa}} - \boldsymbol{r}$ and $\widehat{\boldsymbol{\kappa}}^+ \coloneqq \widehat{\boldsymbol{\kappa}} + \boldsymbol{r}$;

$$\rightarrow \mathrm{P}\left\{\boldsymbol{S\kappa}[k] \notin \left[\widehat{\boldsymbol{\kappa}}^-[k], \widehat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha \text{ (approximately).}$$

- $\widehat{\boldsymbol{\kappa}}$ unbiased estimator of $\boldsymbol{S\kappa}$, but biased estimator of $\boldsymbol{\kappa}$. So are the bounds $\widehat{\boldsymbol{\kappa}}^-$ and $\widehat{\boldsymbol{\kappa}}^+$

$$\rightarrow \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\widehat{\boldsymbol{\kappa}}^-[k], \widehat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha??$$

Kaiser-Squires bound estimation

Random variable centered in $\boldsymbol{S\kappa}$

Estimation of $\widehat{\boldsymbol{\kappa}}^-$ and $\widehat{\boldsymbol{\kappa}}^+$: we have $\widehat{\boldsymbol{\kappa}} \sim \boldsymbol{N}(\boldsymbol{S\kappa}, \mathbf{S}\mathbf{A}^{\dagger}\boldsymbol{\Sigma}\mathbf{A}^{\dagger *}\boldsymbol{S}^*)$.

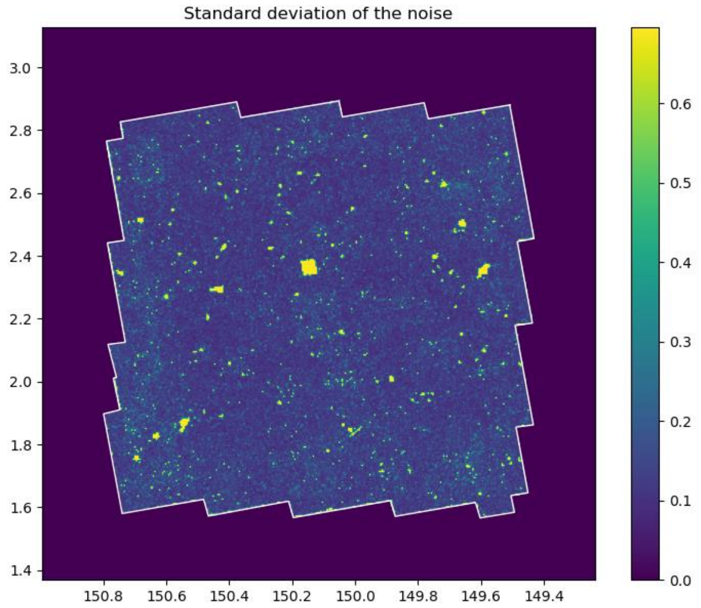

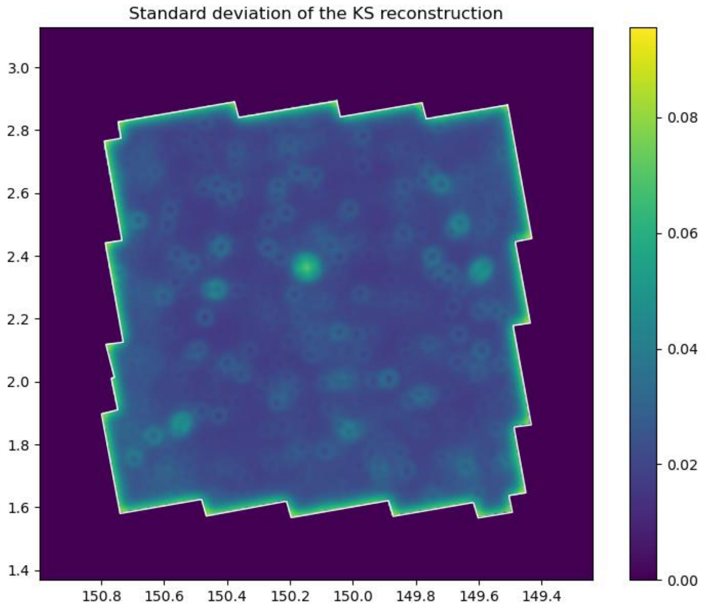

- $\boldsymbol{r}[k] \coloneqq \Phi_k^{-1}(1 - {}^{\alpha}/_{2})$ with Φ_k Gaussian CDF at pixel k;
- $\widehat{\boldsymbol{\kappa}}^- \coloneqq \widehat{\boldsymbol{\kappa}} - \boldsymbol{r}$ and $\widehat{\boldsymbol{\kappa}}^+ \coloneqq \widehat{\boldsymbol{\kappa}} + \boldsymbol{r}$;

$$\rightarrow \mathrm{P}\left\{\boldsymbol{S\kappa}[k] \notin \left[\widehat{\boldsymbol{\kappa}}^-[k], \widehat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha \text{ (approximately).}$$

- $\widehat{\boldsymbol{\kappa}}$ unbiased estimator of $\boldsymbol{S\kappa}$, but biased estimator of $\boldsymbol{\kappa}$. So are the bounds $\widehat{\boldsymbol{\kappa}}^-$ and $\widehat{\boldsymbol{\kappa}}^+$

$$\rightarrow \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\widehat{\boldsymbol{\kappa}}^-[k], \widehat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha??$$

Kaiser-Squires bound estimation

Random variable centered in $\boldsymbol{S\kappa}$

Estimation of $\widehat{\boldsymbol{\kappa}}^-$ and $\widehat{\boldsymbol{\kappa}}^+$: we have $\widehat{\boldsymbol{\kappa}} \sim \boldsymbol{N}(\boldsymbol{S\kappa}, \mathbf{SA}^\dagger \boldsymbol{\Sigma} \mathbf{A}^{\dagger *} \boldsymbol{S}^*)$.

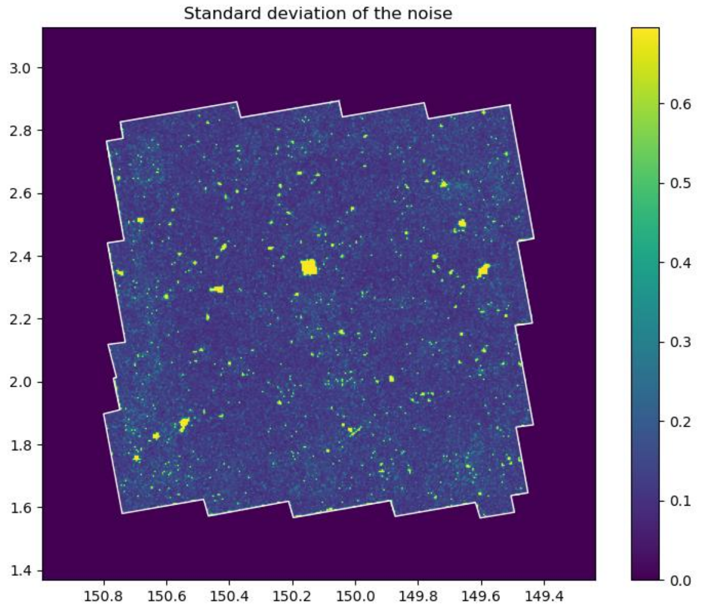

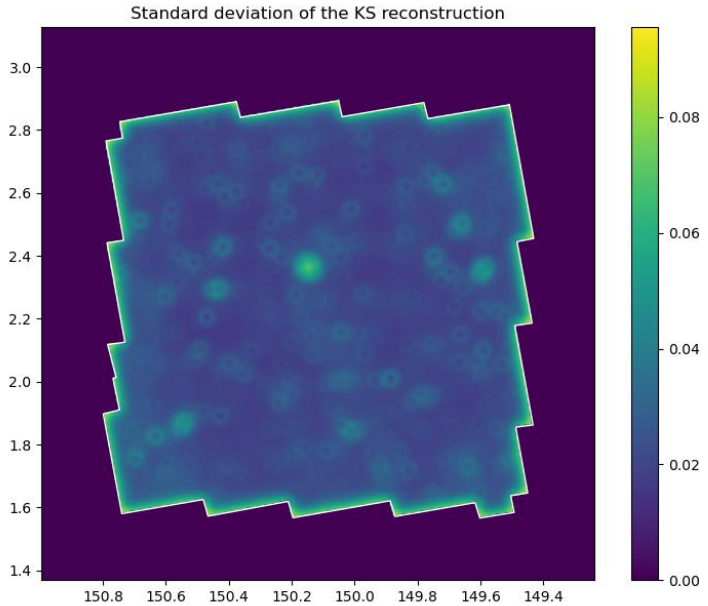

- $\boldsymbol{r}[k] \coloneqq \Phi_k^{-1}(1 - {}^{\alpha}/_{2})$ with Φ_k Gaussian CDF at pixel k;
- $\widehat{\boldsymbol{\kappa}}^- \coloneqq \widehat{\boldsymbol{\kappa}} - \boldsymbol{r}$ and $\widehat{\boldsymbol{\kappa}}^+ \coloneqq \widehat{\boldsymbol{\kappa}} + \boldsymbol{r}$;

$\rightarrow \mathrm{P}\left\{\boldsymbol{S\kappa}[k] \notin \left[\widehat{\boldsymbol{\kappa}}^-[k], \widehat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha$ (approximately).

- $\widehat{\boldsymbol{\kappa}}$ unbiased estimator of $\boldsymbol{S\kappa}$, but biased estimator of $\boldsymbol{\kappa}$. So are the bounds $\widehat{\boldsymbol{\kappa}}^-$ and $\widehat{\boldsymbol{\kappa}}^+$

$\rightarrow \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\widehat{\boldsymbol{\kappa}}^-[k], \widehat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha$?? Calibration needed!

Kaiser-Squires bound estimation

Target: $\alpha \approx 4{,}6\%$ (2σ-confidence)

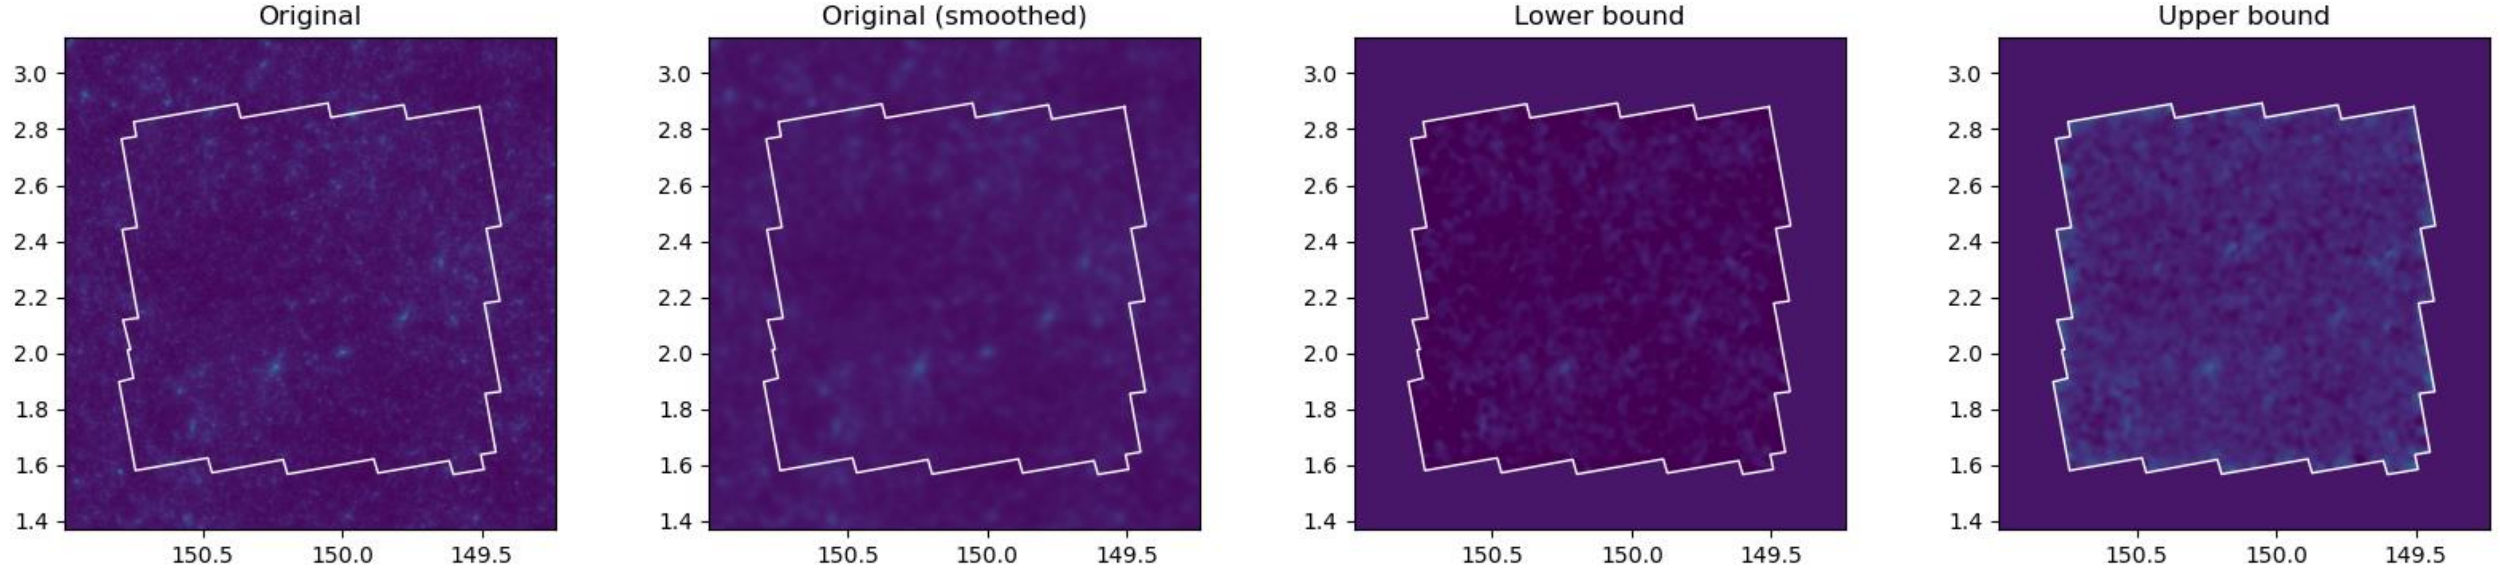

$\boldsymbol{\kappa}$ $\quad$ $\boldsymbol{S\kappa}$ $\quad$ $\boldsymbol{\hat{\kappa}^{-}}$ $\quad$ $\boldsymbol{\hat{\kappa}^{+}}$

Kaiser-Squires bound estimation

Target: $\alpha \approx 4{,}6\%$ (2σ-confidence)

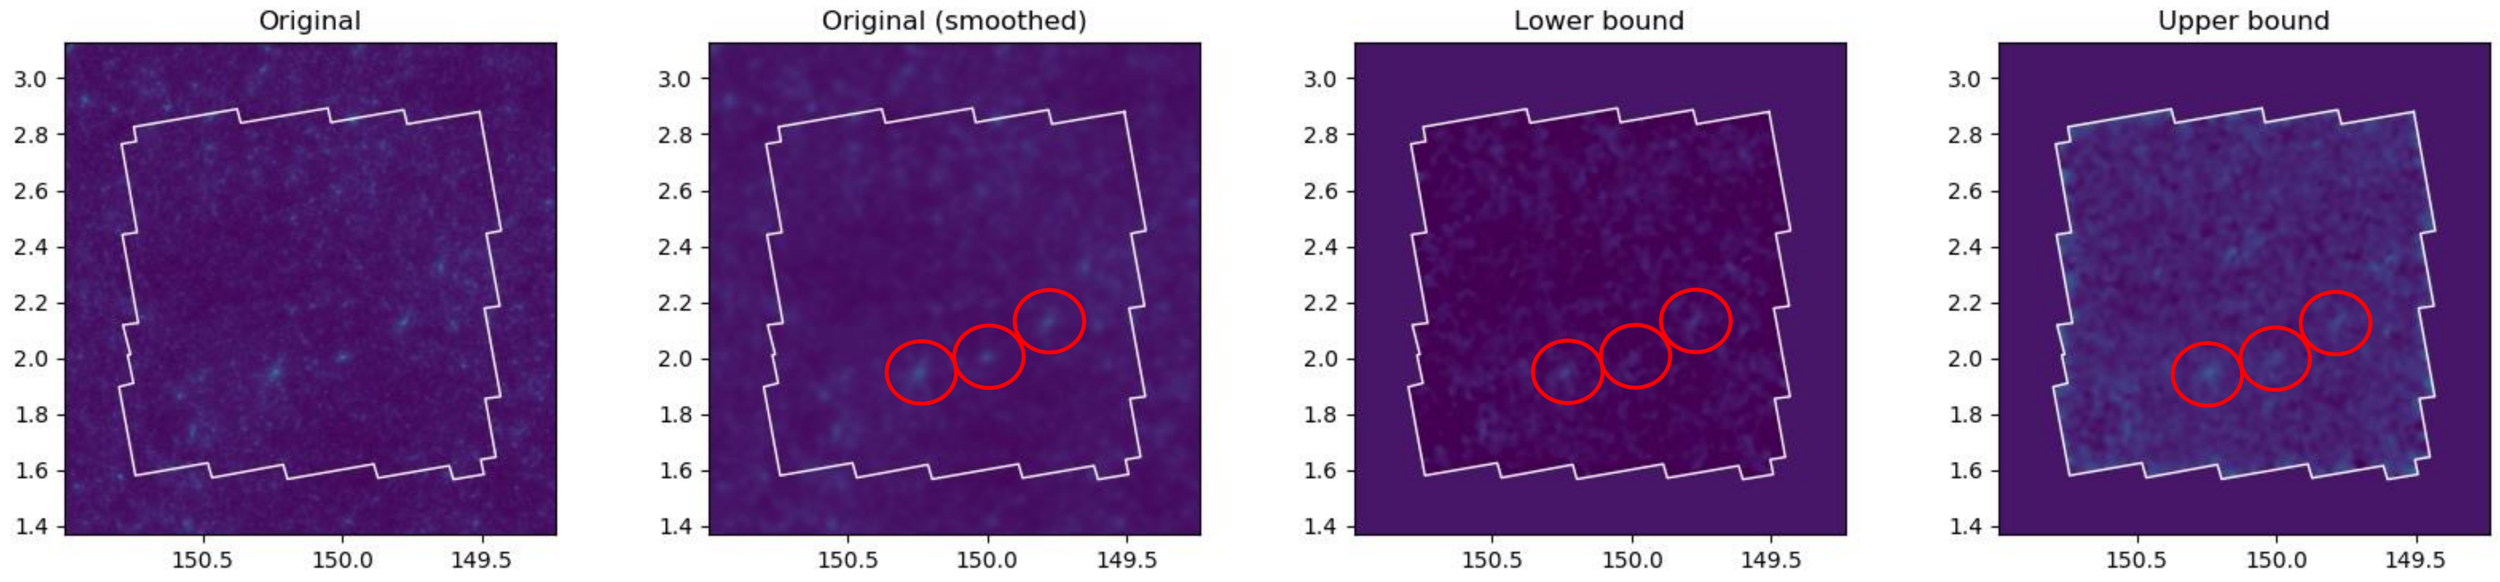

$\boldsymbol{\kappa}$ $\qquad$ $\boldsymbol{S\kappa}$ $\qquad$ $\widehat{\boldsymbol{\kappa}}^{-}$ $\qquad$ $\widehat{\boldsymbol{\kappa}}^{+}$

Correct identification of mass overdensities

Kaiser-Squires bound estimation

Target: $\alpha \approx 4{,}6\%$ (2σ-confidence)

High level of uncertainty due to masked data

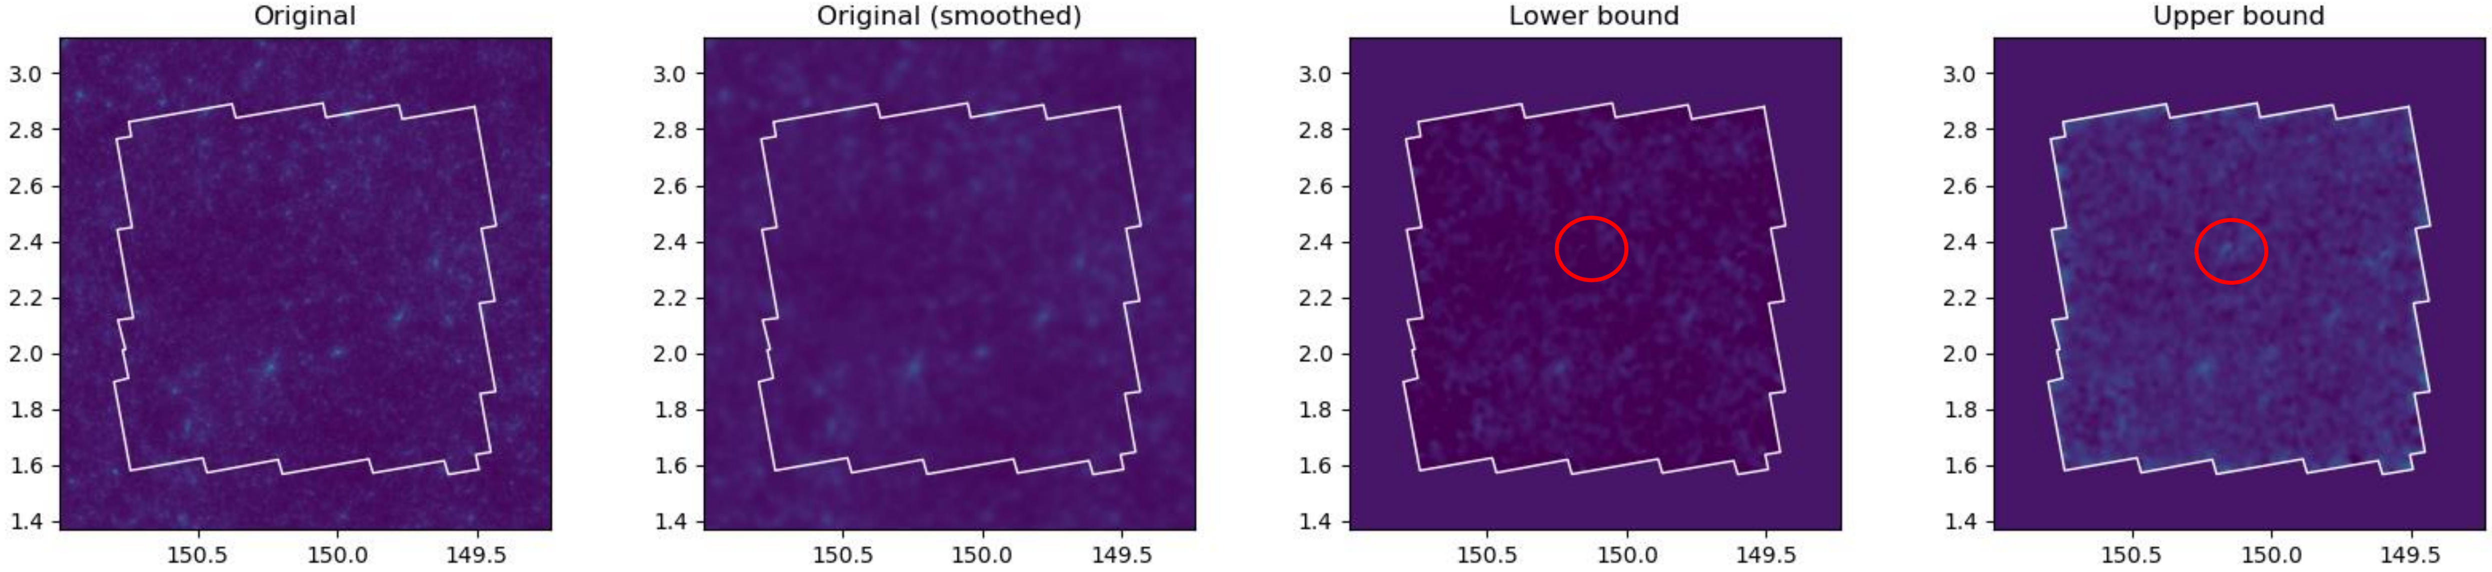

κ $\quad$ $S\kappa$ $\quad$ $\hat{\kappa}^-$ $\quad$ $\hat{\kappa}^+$

Kaiser-Squires bound estimation

$\alpha \approx 4{,}6\%$ (2σ-confidence)

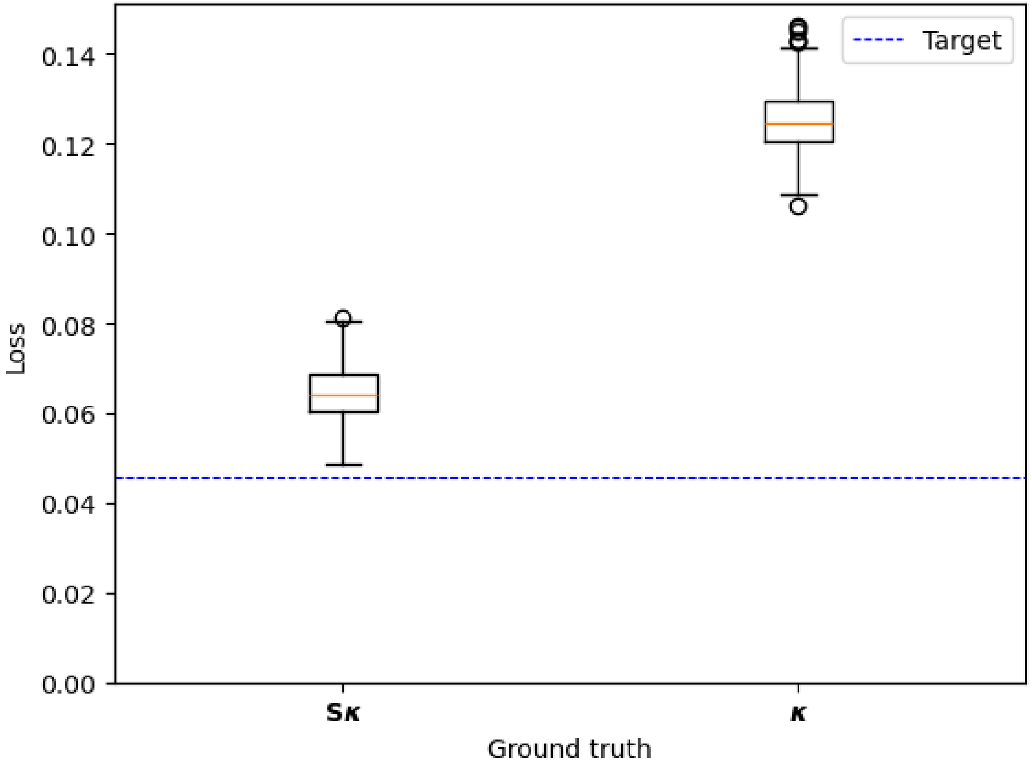

Kaiser-Squires bound estimation

$\alpha \approx 4{,}6\%$ (2σ-confidence)

$\mathrm{P}\{\boldsymbol{S\kappa}[k] \notin [\widehat{\boldsymbol{\kappa}}^{-}[k], \widehat{\boldsymbol{\kappa}}^{+}[k]]\}$

(empirical)

(unbiased estimator)

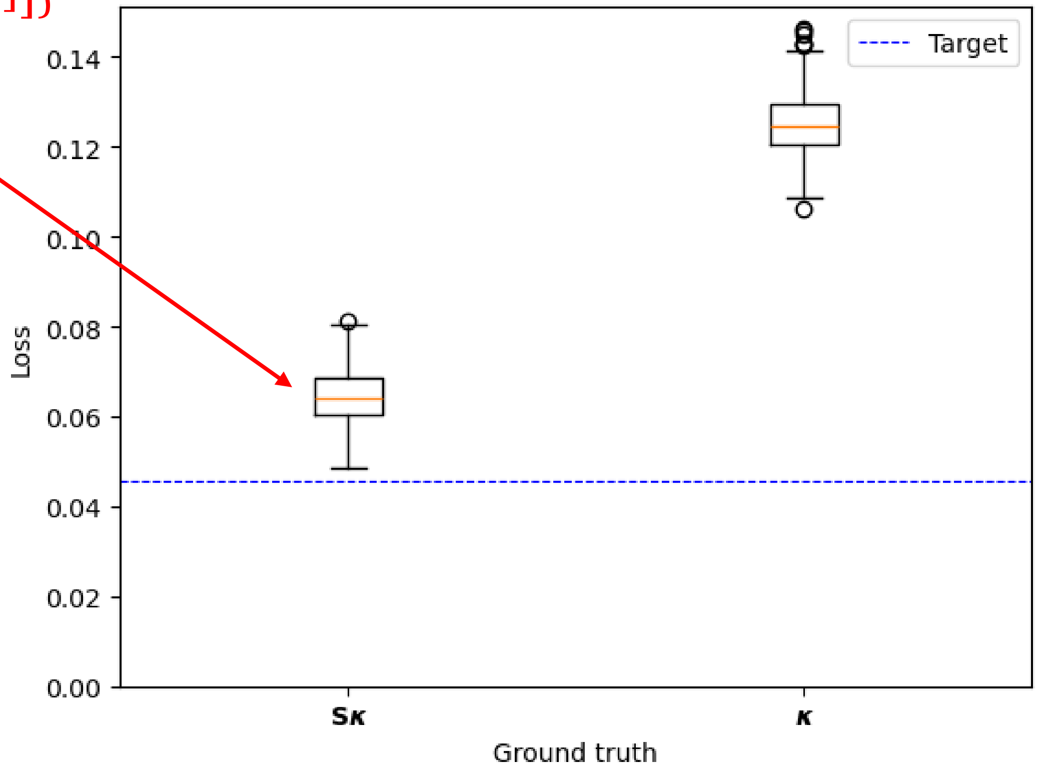

Kaiser-Squires bound estimation

$\alpha \approx 4{,}6\%$ (2σ-confidence)

$\mathrm{P}\{\boldsymbol{S\kappa}[k] \notin [\widehat{\boldsymbol{\kappa}}^{-}[k], \widehat{\boldsymbol{\kappa}}^{+}[k]]\}$
(empirical)
(unbiased estimator)

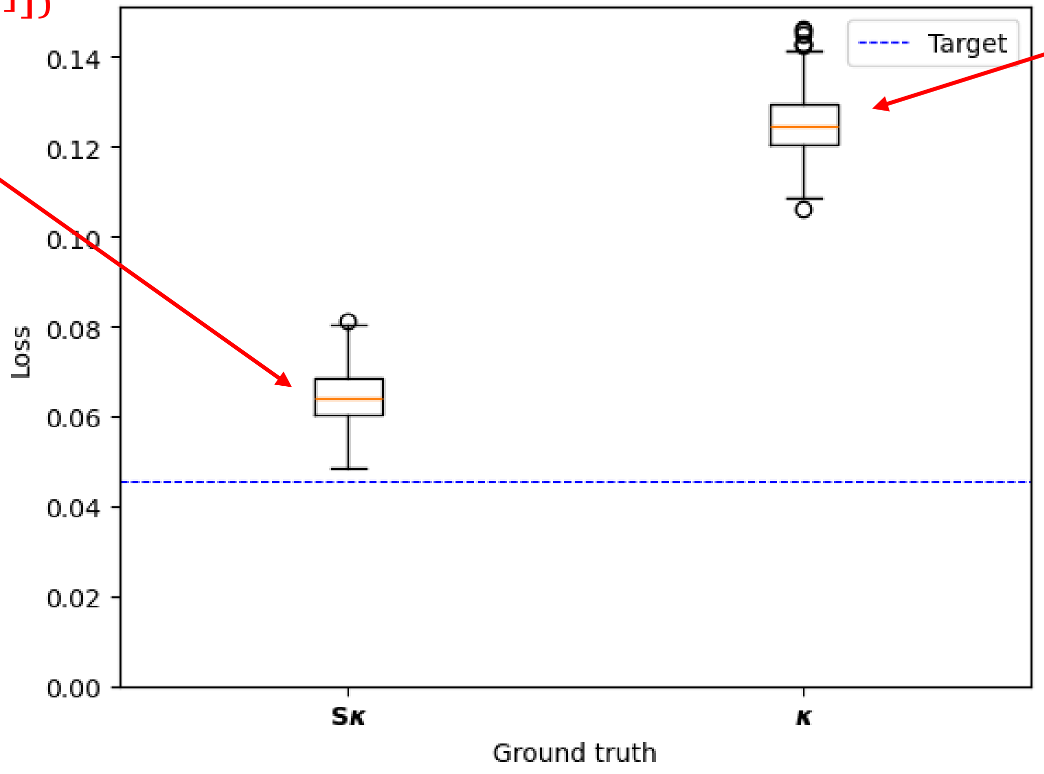

$\mathrm{P}\{\boldsymbol{\kappa}[k] \notin [\widehat{\boldsymbol{\kappa}}^{-}[k], \widehat{\boldsymbol{\kappa}}^{+}[k]]\}$
(empirical)
(biased estimator)

Kaiser-Squires bound estimation

$\alpha \approx 4{,}6\%$ (2σ-confidence)

$\mathrm{P}\{\boldsymbol{S\kappa}[k] \notin [\widehat{\boldsymbol{\kappa}}^-[k], \widehat{\boldsymbol{\kappa}}^+[k]]\}$
(empirical)
(unbiased estimator)

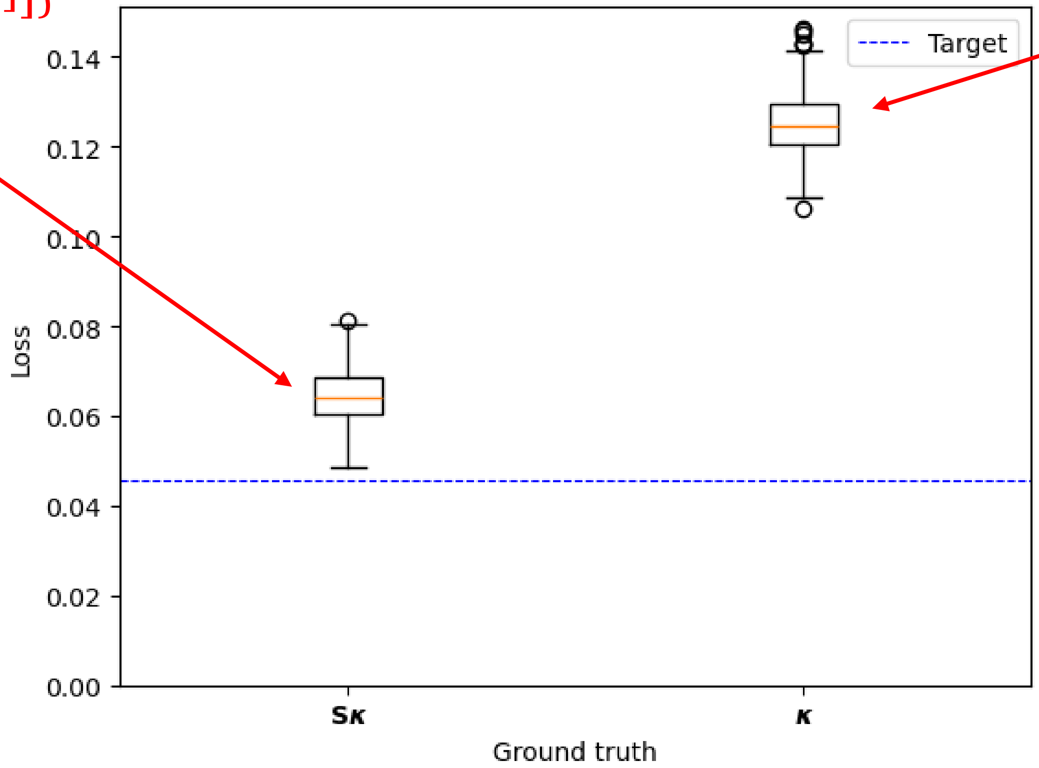

$\mathrm{P}\{\boldsymbol{\kappa}[k] \notin [\widehat{\boldsymbol{\kappa}}^-[k], \widehat{\boldsymbol{\kappa}}^+[k]]\}$
(empirical)
(biased estimator)

→ **Undercoverage**

Calibration procedure

Objective (reminder): given $\hat{\boldsymbol{\gamma}}$, estimate $\hat{\boldsymbol{\kappa}}^-$ and $\hat{\boldsymbol{\kappa}}^+$ such that

$$\mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^-[k], \hat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha.$$

Conformalized quantile regression (CQR):[1] performed on a **calibration set** $(\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n$. For each pixel k:

1. compute a prediction score on each calibration example:

$$\boldsymbol{e}_i[k] \coloneqq \max\{\hat{\boldsymbol{\kappa}}_i^-[k] - \boldsymbol{\kappa}_i[k], \boldsymbol{\kappa}_i[k] - \hat{\boldsymbol{\kappa}}_i^+[k]\};$$

2. get the $(1-\alpha)$-quantile of $(\boldsymbol{e}_i[k])_{i=1}^n$, denoted by $\boldsymbol{q}_{(1-\alpha)}[k]$;
3. adjust the bounds: set $\hat{\boldsymbol{\kappa}}^- \leftarrow \hat{\boldsymbol{\kappa}}^- - \boldsymbol{q}_{(1-\alpha)}$ and $\hat{\boldsymbol{\kappa}}^+ \leftarrow \hat{\boldsymbol{\kappa}}^+ + \boldsymbol{q}_{(1-\alpha)}$.

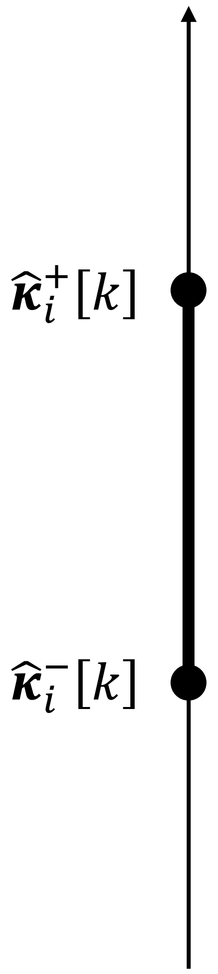

[1] Y. Romano, E. Patterson, and E. Candès, "Conformalized Quantile Regression," in *NeurIPS*, 2019.

Calibration procedure

Objective (reminder): given $\hat{\boldsymbol{\gamma}}$, estimate $\hat{\boldsymbol{\kappa}}^-$ and $\hat{\boldsymbol{\kappa}}^+$ such that

$$\mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^-[k], \hat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha.$$

Ground truth inside prediction bounds → $\boldsymbol{e}_i[k] < 0$

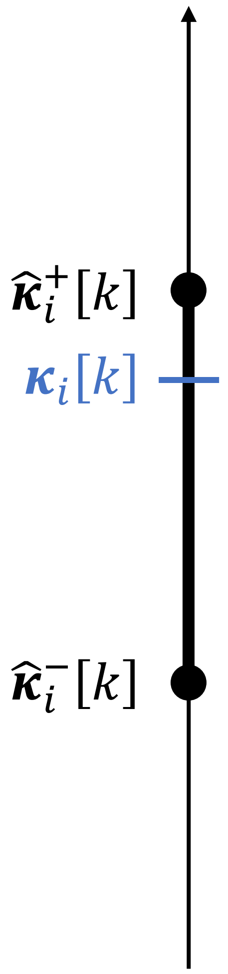

Conformalized quantile regression (CQR):[1] performed on a **calibration set** $(\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n$. For each pixel k:

1. compute a prediction score on each calibration example:

$$\boldsymbol{e}_i[k] \coloneqq \max\{\hat{\boldsymbol{\kappa}}_i^-[k] - \boldsymbol{\kappa}_i[k], \boldsymbol{\kappa}_i[k] - \hat{\boldsymbol{\kappa}}_i^+[k]\};$$

2. get the $(1-\alpha)$-quantile of $(\boldsymbol{e}_i[k])_{i=1}^n$, denoted by $\boldsymbol{q}_{(1-\alpha)}[k]$;
3. adjust the bounds: set $\hat{\boldsymbol{\kappa}}^- \leftarrow \hat{\boldsymbol{\kappa}}^- - \boldsymbol{q}_{(1-\alpha)}$ and $\hat{\boldsymbol{\kappa}}^+ \leftarrow \hat{\boldsymbol{\kappa}}^+ + \boldsymbol{q}_{(1-\alpha)}$.

[1] Y. Romano, E. Patterson, and E. Candès, "Conformalized Quantile Regression," in *NeurIPS*, 2019.

Calibration procedure

Ground truth outside prediction bounds → $\boldsymbol{e}_i[k] > 0$

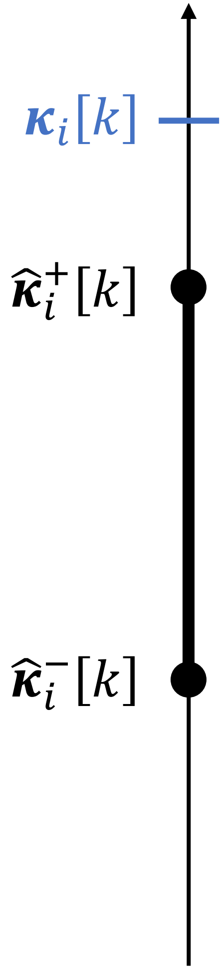

Objective (reminder): given $\widehat{\boldsymbol{\gamma}}$, estimate $\widehat{\boldsymbol{\kappa}}^-$ and $\widehat{\boldsymbol{\kappa}}^+$ such that

$$\mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\widehat{\boldsymbol{\kappa}}^-[k], \widehat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha.$$

Conformalized quantile regression (CQR):[1] performed on a **calibration set** $(\widehat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n$. For each pixel k:

1. compute a prediction score on each calibration example:

$$\boldsymbol{e}_i[k] \coloneqq \max\{\widehat{\boldsymbol{\kappa}}_i^-[k] - \boldsymbol{\kappa}_i[k], \boldsymbol{\kappa}_i[k] - \widehat{\boldsymbol{\kappa}}_i^+[k]\};$$

2. get the $(1-\alpha)$-quantile of $(\boldsymbol{e}_i[k])_{i=1}^n$, denoted by $\boldsymbol{q}_{(1-\alpha)}[k]$;
3. adjust the bounds: set $\widehat{\boldsymbol{\kappa}}^- \leftarrow \widehat{\boldsymbol{\kappa}}^- - \boldsymbol{q}_{(1-\alpha)}$ and $\widehat{\boldsymbol{\kappa}}^+ \leftarrow \widehat{\boldsymbol{\kappa}}^+ + \boldsymbol{q}_{(1-\alpha)}$.

[1] Y. Romano, E. Patterson, and E. Candès, “Conformalized Quantile Regression,” in *NeurIPS*, 2019.

Calibration procedure

Objective (reminder): given $\hat{\boldsymbol{\gamma}}$, estimate $\hat{\boldsymbol{\kappa}}^-$ and $\hat{\boldsymbol{\kappa}}^+$ such that

$$\mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^-[k], \hat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha.$$

Conformalized quantile regression (CQR):[1] performed on a **calibration set** $(\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n$. For each pixel k:

1. compute a prediction score on each calibration example:

$$\boldsymbol{e}_i[k] \coloneqq \max\{\hat{\boldsymbol{\kappa}}_i^-[k] - \boldsymbol{\kappa}_i[k], \boldsymbol{\kappa}_i[k] - \hat{\boldsymbol{\kappa}}_i^+[k]\};$$

2. get the $(1-\alpha)$-quantile of $(\boldsymbol{e}_i[k])_{i=1}^n$, denoted by $\boldsymbol{q}_{(1-\alpha)}[k]$;
3. adjust the bounds: set $\hat{\boldsymbol{\kappa}}^- \leftarrow \hat{\boldsymbol{\kappa}}^- - \boldsymbol{q}_{(1-\alpha)}$ and $\hat{\boldsymbol{\kappa}}^+ \leftarrow \hat{\boldsymbol{\kappa}}^+ + \boldsymbol{q}_{(1-\alpha)}$.

Well-calibrated → $\boldsymbol{q}_{(1-\alpha)}[k] = 0$
Undercoverage → $\boldsymbol{q}_{(1-\alpha)}[k] > 0$
Overcoverage → $\boldsymbol{q}_{(1-\alpha)}[k] < 0$

[1] Y. Romano, E. Patterson, and E. Candès, "Conformalized Quantile Regression," in *NeurIPS*, 2019.

Calibration procedure

Objective (reminder): given $\widehat{\boldsymbol{\gamma}}$, estimate $\widehat{\boldsymbol{\kappa}}^-$ and $\widehat{\boldsymbol{\kappa}}^+$ such that

$$\mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\widehat{\boldsymbol{\kappa}}^-[k], \widehat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha.$$

Conformalized quantile regression (CQR):[1] performed on a **calibration set** $(\widehat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n$. For each pixel k:

1. compute a prediction score on each calibration example:

$$\boldsymbol{e}_i[k] \coloneqq \max\{\widehat{\boldsymbol{\kappa}}_i^-[k] - \boldsymbol{\kappa}_i[k], \boldsymbol{\kappa}_i[k] - \widehat{\boldsymbol{\kappa}}_i^+[k]\};$$

2. get the $(1-\alpha)$-quantile of $(\boldsymbol{e}_i[k])_{i=1}^n$, denoted by $\boldsymbol{q}_{(1-\alpha)}[k]$;
3. adjust the bounds: set $\widehat{\boldsymbol{\kappa}}^- \leftarrow \widehat{\boldsymbol{\kappa}}^- - \boldsymbol{q}_{(1-\alpha)}$ and $\widehat{\boldsymbol{\kappa}}^+ \leftarrow \widehat{\boldsymbol{\kappa}}^+ + \boldsymbol{q}_{(1-\alpha)}$.

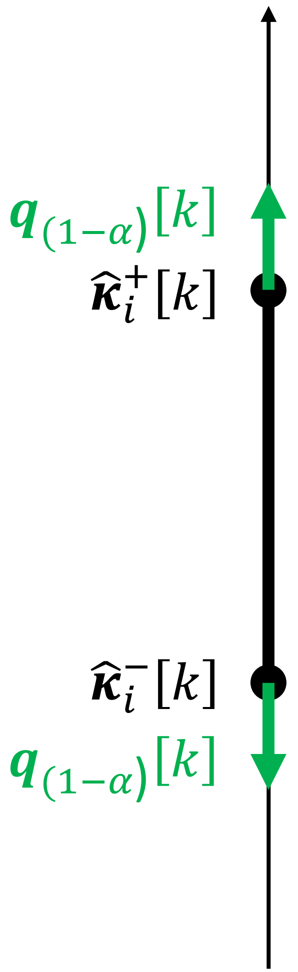

[1] Y. Romano, E. Patterson, and E. Candès, "Conformalized Quantile Regression," in *NeurIPS*, 2019.

Calibration procedure

Objective (reminder): given $\widehat{\boldsymbol{\gamma}}$, estimate $\widehat{\boldsymbol{\kappa}}^-$ and $\widehat{\boldsymbol{\kappa}}^+$ such that

$$\mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\widehat{\boldsymbol{\kappa}}^-[k], \widehat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha.$$

Conformalized quantile regression (CQR):[1] performed on a **calibration set** $(\widehat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n$. For each pixel k:

1. compute a prediction score on each calibration example:

$$\boldsymbol{e}_i[k] \coloneqq \max\{\widehat{\boldsymbol{\kappa}}_i^-[k] - \boldsymbol{\kappa}_i[k], \boldsymbol{\kappa}_i[k] - \widehat{\boldsymbol{\kappa}}_i^+[k]\};$$

2. get the $(1-\alpha)$-quantile of $(\boldsymbol{e}_i[k])_{i=1}^n$, denoted by $\boldsymbol{q}_{(1-\alpha)}[k]$;
3. adjust the bounds: set $\widehat{\boldsymbol{\kappa}}^- \leftarrow \widehat{\boldsymbol{\kappa}}^- - \boldsymbol{q}_{(1-\alpha)}$ and $\widehat{\boldsymbol{\kappa}}^+ \leftarrow \widehat{\boldsymbol{\kappa}}^+ + \boldsymbol{q}_{(1-\alpha)}$.

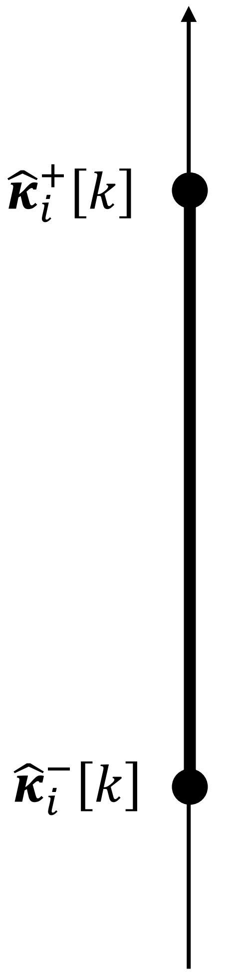

[1] Y. Romano, E. Patterson, and E. Candès, "Conformalized Quantile Regression," in *NeurIPS*, 2019.

Calibration procedure

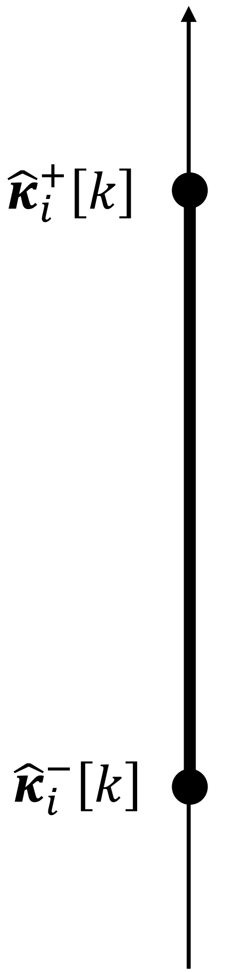

Objective (reminder): given $\hat{\boldsymbol{\gamma}}$, estimate $\hat{\boldsymbol{\kappa}}^-$ and $\hat{\boldsymbol{\kappa}}^+$ such that

$$\mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^-[k], \hat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha.$$

Conformalized quantile regression (CQR):[1] performed on a **calibration set** $(\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n$. For each pixel k:

1. compute a prediction score on each calibration example:
$$\boldsymbol{e}_i[k] \coloneqq \max\{\hat{\boldsymbol{\kappa}}_i^-[k] - \boldsymbol{\kappa}_i[k], \boldsymbol{\kappa}_i[k] - \hat{\boldsymbol{\kappa}}_i^+[k]\};$$
2. get the $(1-\alpha)$-quantile of $(\boldsymbol{e}_i[k])_{i=1}^n$, denoted by $\boldsymbol{q}_{(1-\alpha)}[k]$;
3. adjust the bounds: set $\hat{\boldsymbol{\kappa}}^- \leftarrow \hat{\boldsymbol{\kappa}}^- - \boldsymbol{q}_{(1-\alpha)}$ and $\hat{\boldsymbol{\kappa}}^+ \leftarrow \hat{\boldsymbol{\kappa}}^+ + \boldsymbol{q}_{(1-\alpha)}$.

THEOREM:[1] $\alpha - {}^{1}/_{n+1} \leq \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^-[k], \hat{\boldsymbol{\kappa}}^+[k]\right] \mid (\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n\right\} \leq \alpha$ for any pixel k.

[1] Y. Romano, E. Patterson, and E. Candès, "Conformalized Quantile Regression," in *NeurIPS*, 2019.

Calibration procedure

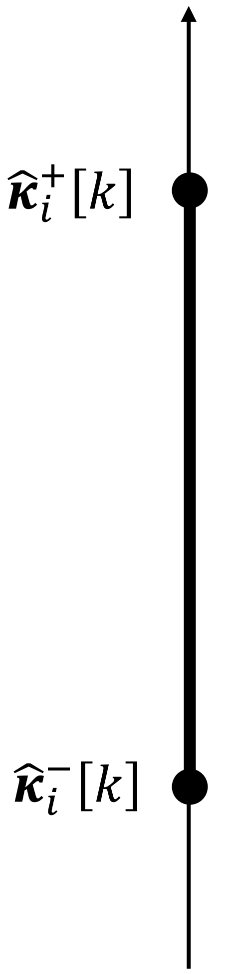

Objective (reminder): given $\hat{\boldsymbol{\gamma}}$, estimate $\hat{\boldsymbol{\kappa}}^-$ and $\hat{\boldsymbol{\kappa}}^+$ such that

$$\mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^-[k], \hat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha.$$

Conformalized quantile regression (CQR):[1] performed on a **calibration set** $(\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n$. For each pixel k:

1. compute a prediction score on each calibration example:

$$\boldsymbol{e}_i[k] \coloneqq \max\{\hat{\boldsymbol{\kappa}}_i^-[k] - \boldsymbol{\kappa}_i[k], \boldsymbol{\kappa}_i[k] - \hat{\boldsymbol{\kappa}}_i^+[k]\};$$

2. get the $(1-\alpha)$-quantile of $(\boldsymbol{e}_i[k])_{i=1}^n$, denoted by $\boldsymbol{q}_{(1-\alpha)}[k]$;
3. adjust the bounds: set $\hat{\boldsymbol{\kappa}}^- \leftarrow \hat{\boldsymbol{\kappa}}^- - \boldsymbol{q}_{(1-\alpha)}$ and $\hat{\boldsymbol{\kappa}}^+ \leftarrow \hat{\boldsymbol{\kappa}}^+ + \boldsymbol{q}_{(1-\alpha)}$.

THEOREM:[1] $\alpha - {}^1/_{n+1} \leq \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^-[k], \hat{\boldsymbol{\kappa}}^+[k]\right] \mid (\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n\right\} \leq \alpha$ for any pixel k.

Conditionally to a specific calibration set

[1] Y. Romano, E. Patterson, and E. Candès, "Conformalized Quantile Regression," in *NeurIPS*, 2019.

Calibration procedure

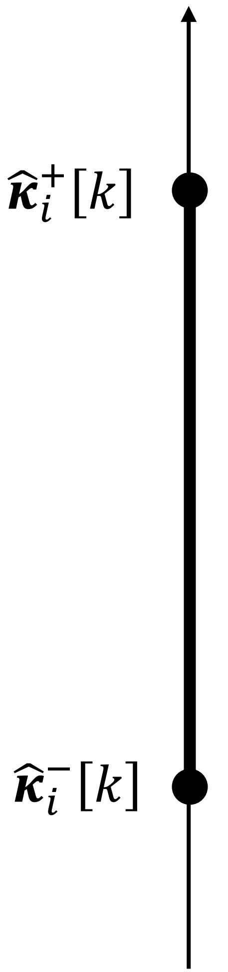

Objective (reminder): given $\hat{\boldsymbol{\gamma}}$, estimate $\hat{\boldsymbol{\kappa}}^-$ and $\hat{\boldsymbol{\kappa}}^+$ such that

$$\mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^-[k], \hat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha.$$

Conformalized quantile regression (CQR):[1] performed on a **calibration set** $(\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n$. For each pixel k:

1. compute a prediction score on each calibration example:

$$\boldsymbol{e}_i[k] \coloneqq \max\{\hat{\boldsymbol{\kappa}}_i^-[k] - \boldsymbol{\kappa}_i[k], \boldsymbol{\kappa}_i[k] - \hat{\boldsymbol{\kappa}}_i^+[k]\};$$

2. get the $(1-\alpha)$-quantile of $(\boldsymbol{e}_i[k])_{i=1}^n$, denoted by $\boldsymbol{q}_{(1-\alpha)}[k]$;
3. adjust the bounds: set $\hat{\boldsymbol{\kappa}}^- \leftarrow \hat{\boldsymbol{\kappa}}^- - \boldsymbol{q}_{(1-\alpha)}$ and $\hat{\boldsymbol{\kappa}}^+ \leftarrow \hat{\boldsymbol{\kappa}}^+ + \boldsymbol{q}_{(1-\alpha)}$.

THEOREM:[1] $\alpha - {}^1/_{n+1} \leq \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^-[k], \hat{\boldsymbol{\kappa}}^+[k]\right] \mid (\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n\right\} \leq \alpha$ for any pixel k.

Upper bound: coverage guarantee

[1] Y. Romano, E. Patterson, and E. Candès, "Conformalized Quantile Regression," in *NeurIPS*, 2019.

Calibration procedure

Objective (reminder): given $\widehat{\boldsymbol{\gamma}}$, estimate $\widehat{\boldsymbol{\kappa}}^-$ and $\widehat{\boldsymbol{\kappa}}^+$ such that

$$\mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\widehat{\boldsymbol{\kappa}}^-[k], \widehat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha.$$

Conformalized quantile regression (CQR):[1] performed on a **calibration set** $(\widehat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n$. For each pixel k:

1. compute a prediction score on each calibration example:

$$\boldsymbol{e}_i[k] \coloneqq \max\{\widehat{\boldsymbol{\kappa}}_i^-[k] - \boldsymbol{\kappa}_i[k], \boldsymbol{\kappa}_i[k] - \widehat{\boldsymbol{\kappa}}_i^+[k]\};$$

2. get the $(1-\alpha)$-quantile of $(\boldsymbol{e}_i[k])_{i=1}^n$, denoted by $\boldsymbol{q}_{(1-\alpha)}[k]$;
3. adjust the bounds: set $\widehat{\boldsymbol{\kappa}}^- \leftarrow \widehat{\boldsymbol{\kappa}}^- - \boldsymbol{q}_{(1-\alpha)}$ and $\widehat{\boldsymbol{\kappa}}^+ \leftarrow \widehat{\boldsymbol{\kappa}}^+ + \boldsymbol{q}_{(1-\alpha)}$.

THEOREM:[1] $\alpha - {}^1/_{n+1} \leq \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\widehat{\boldsymbol{\kappa}}^-[k], \widehat{\boldsymbol{\kappa}}^+[k]\right] \,|\, (\widehat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n\right\} \leq \alpha$ for any pixel k.

Lower bound: prevents overconservative prediction bounds

[1] Y. Romano, E. Patterson, and E. Candès, "Conformalized Quantile Regression," in *NeurIPS*, 2019.

Calibration procedure

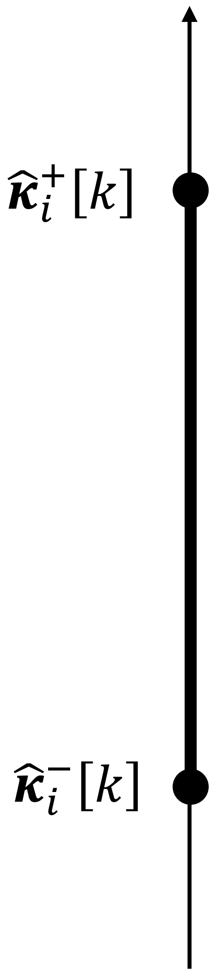

Objective (reminder): given $\hat{\boldsymbol{\gamma}}$, estimate $\hat{\boldsymbol{\kappa}}^-$ and $\hat{\boldsymbol{\kappa}}^+$ such that

$$\mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^-[k], \hat{\boldsymbol{\kappa}}^+[k]\right]\right\} \leq \alpha.$$

Conformalized quantile regression (CQR):[1] performed on a **calibration set** $(\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n$. For each pixel k:

1. compute a prediction score on each calibration example:

$$\boldsymbol{e}_i[k] \coloneqq \max\{\hat{\boldsymbol{\kappa}}_i^-[k] - \boldsymbol{\kappa}_i[k], \boldsymbol{\kappa}_i[k] - \hat{\boldsymbol{\kappa}}_i^+[k]\};$$

2. get the $(1-\alpha)$-quantile of $(\boldsymbol{e}_i[k])_{i=1}^n$, denoted by $\boldsymbol{q}_{(1-\alpha)}[k]$;
3. adjust the bounds: set $\hat{\boldsymbol{\kappa}}^- \leftarrow \hat{\boldsymbol{\kappa}}^- - \boldsymbol{q}_{(1-\alpha)}$ and $\hat{\boldsymbol{\kappa}}^+ \leftarrow \hat{\boldsymbol{\kappa}}^+ + \boldsymbol{q}_{(1-\alpha)}$.

THEOREM:[1] $\alpha - {}^1/_{n+1} \leq \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^-[k], \hat{\boldsymbol{\kappa}}^+[k]\right] \,|\, (\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^n\right\} \leq \alpha$ for any pixel k.

Works for any blackbox quantile predictor!

[1] Y. Romano, E. Patterson, and E. Candès, "Conformalized Quantile Regression," in *NeurIPS*, 2019.

CQR calibration

Target: $\alpha \approx 4{,}6\%$ (2σ-confidence)

$n = 100$ (size of the calibration set)

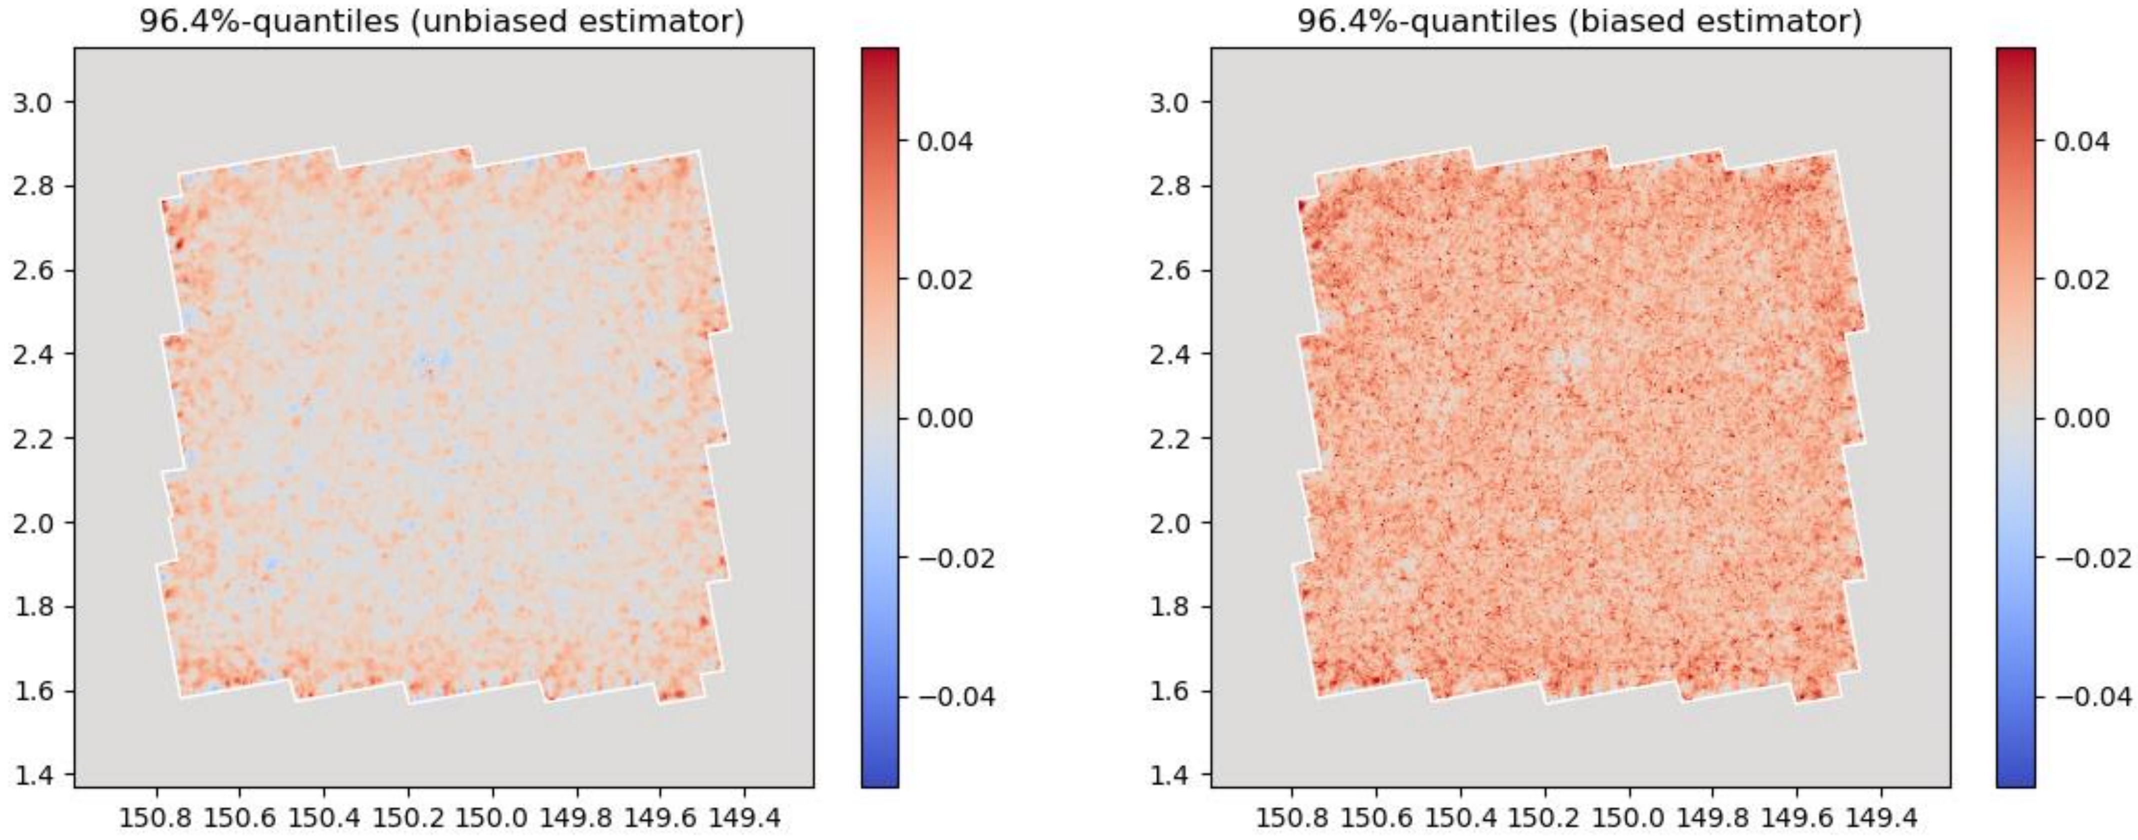

CQR calibration

Target: $\alpha \approx 4{,}6\%$ (2σ-confidence)

$n = 100$ (size of the calibration set)

Undercoverage near the edges

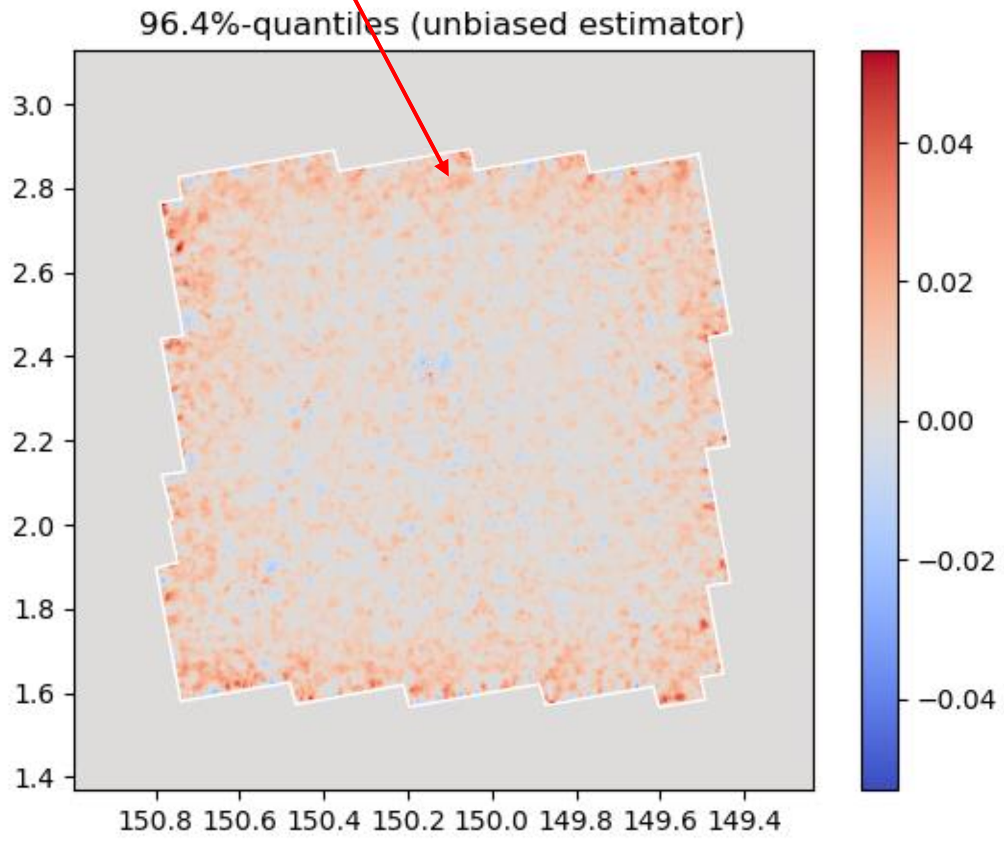

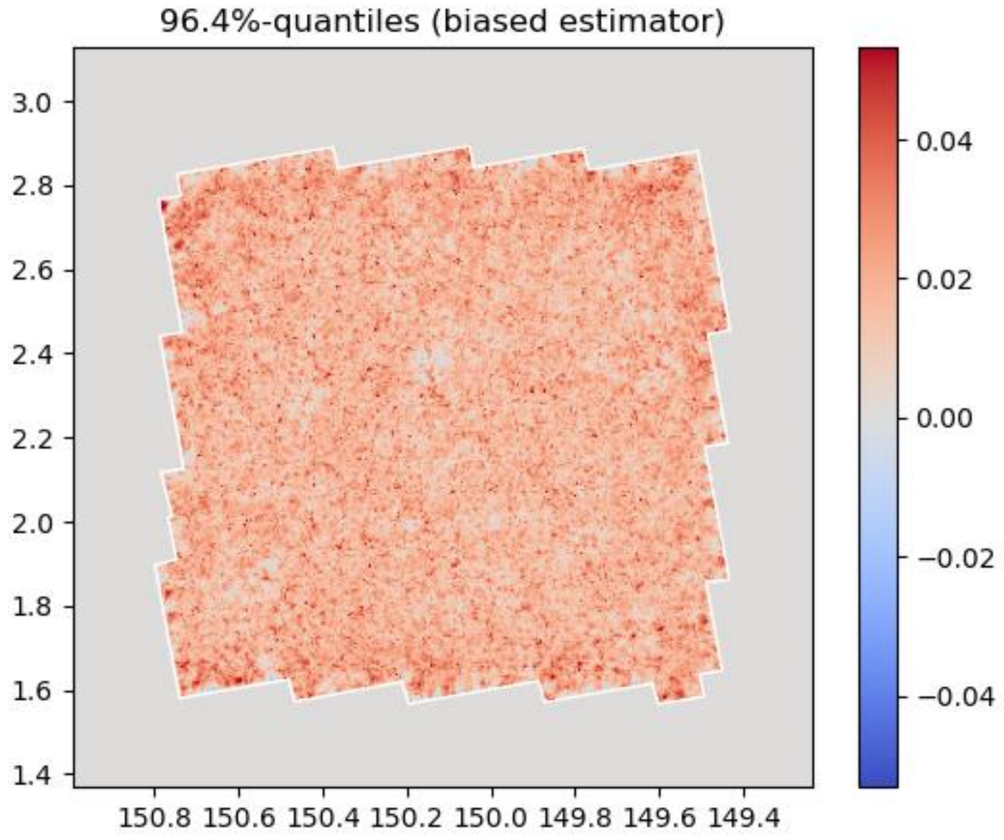

CQR calibration

Target: $\alpha \approx 4{,}6\%$ (2σ-confidence)

$n = 100$ (size of the calibration set)

Overcoverage near the masked regions

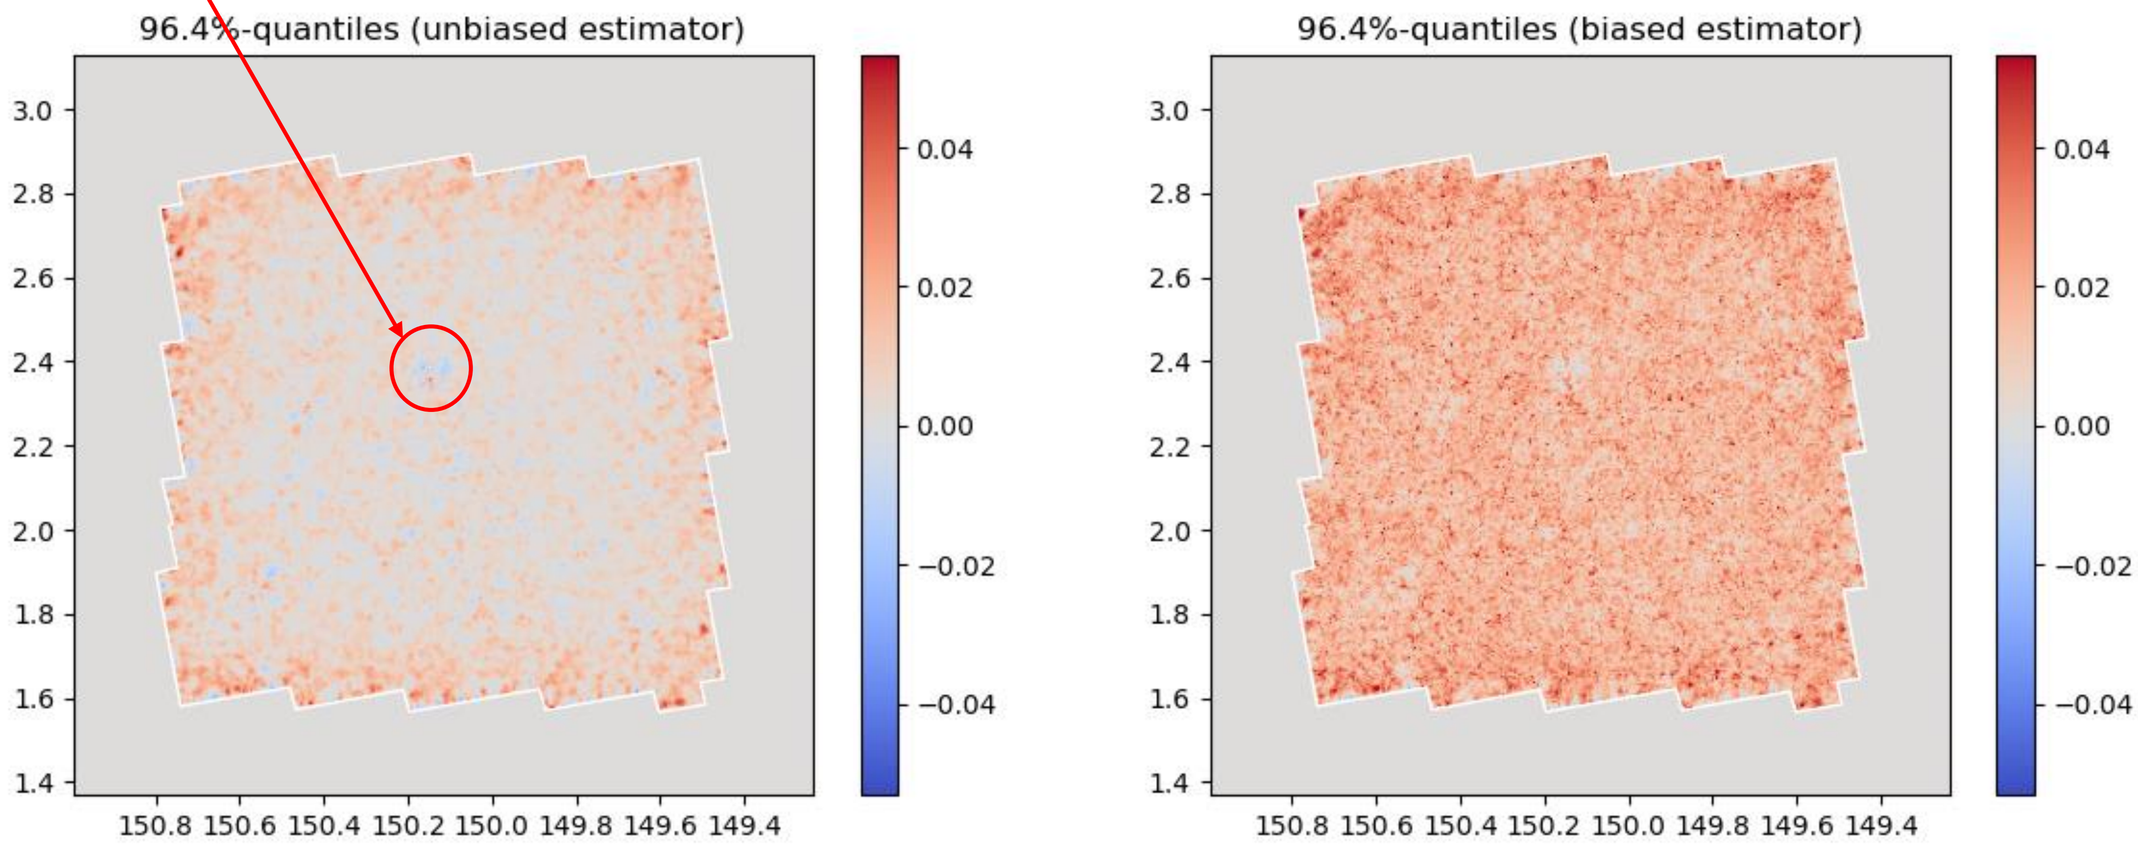

CQR calibration – losses

Target: $\alpha \approx 4{,}6\%$ (2σ-confidence)

$n = 100$ (size of the calibration set)

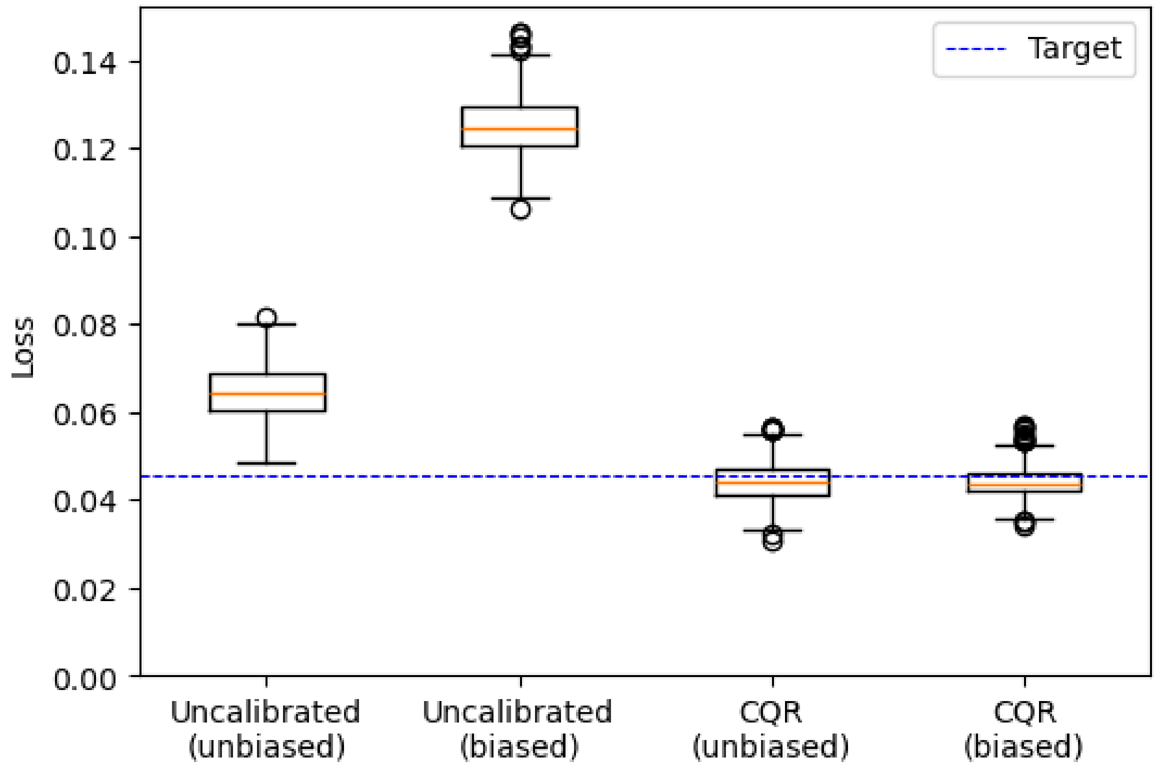

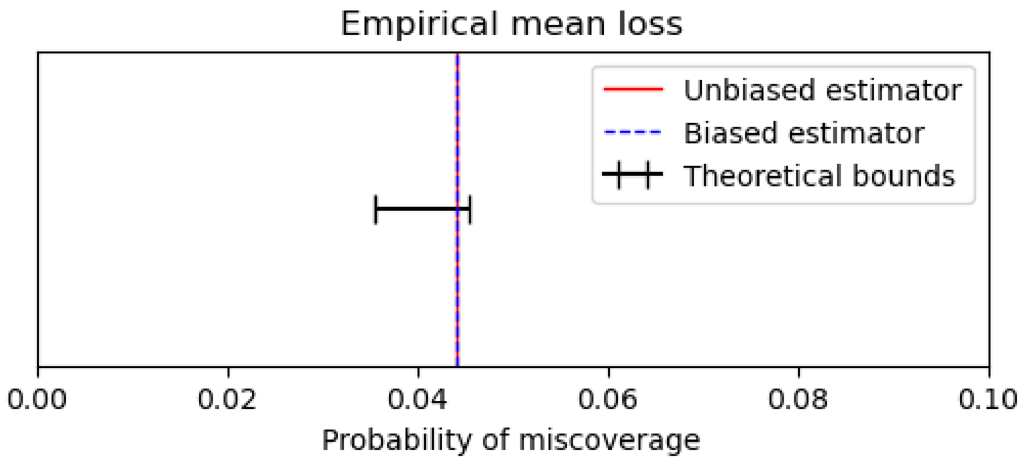

$$\alpha - {}^{1}\!/_{n+1} \leq \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^{-}[k], \hat{\boldsymbol{\kappa}}^{+}[k]\right] \mid (\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^{n}\right\} \leq \alpha$$

CQR calibration – losses

Target: $\alpha \approx 4{,}6\%$ (2σ-confidence)

$n = 100$ (size of the calibration set)

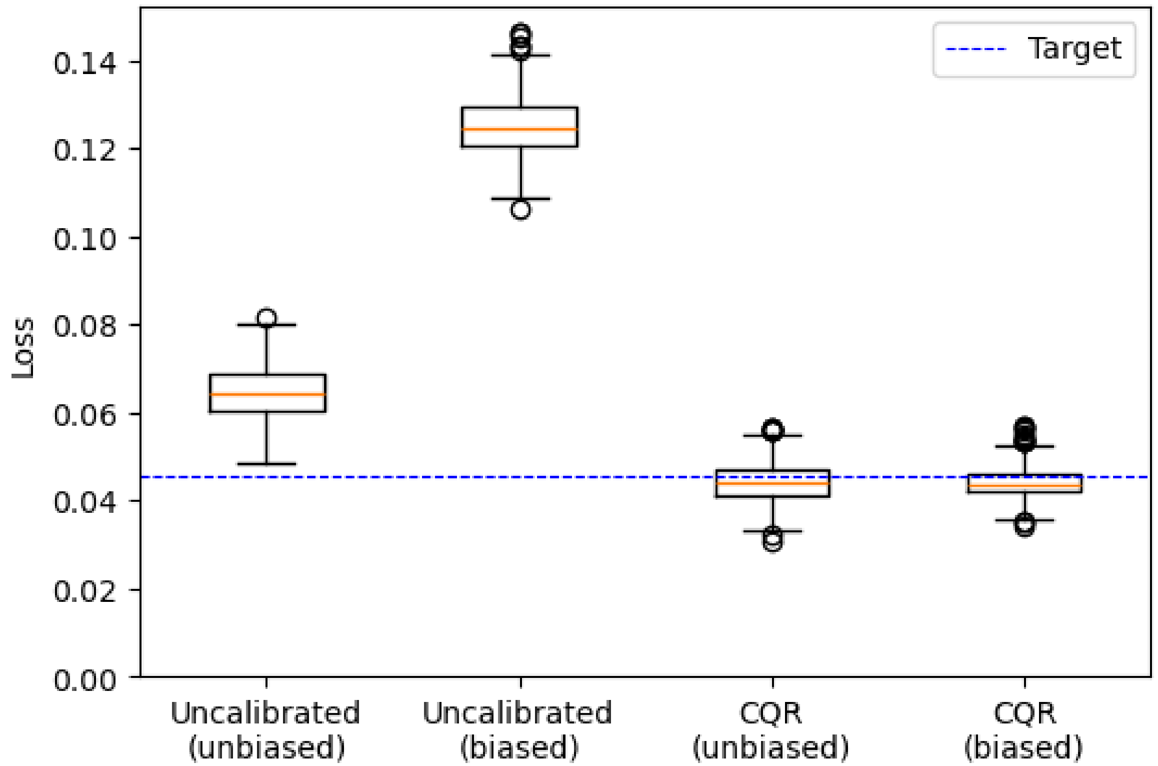

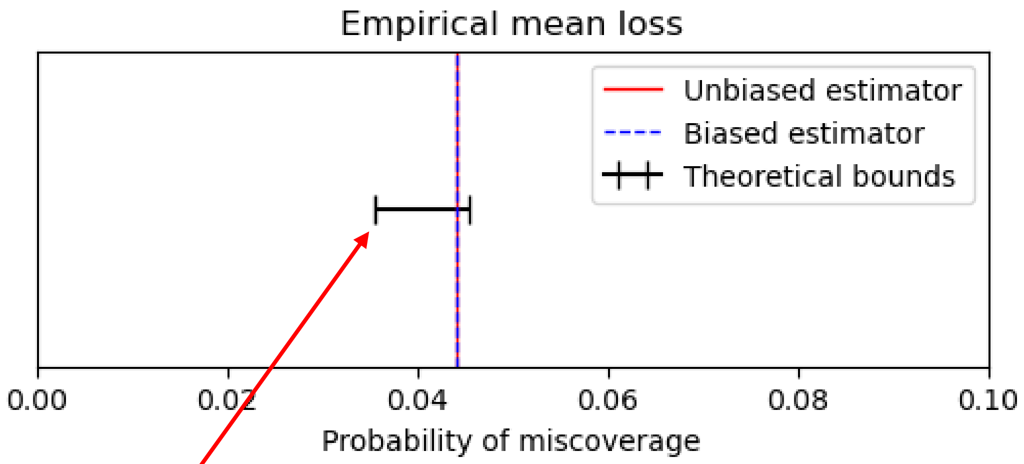

$$\alpha - {}^{1}/_{n+1} \leq \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^{-}[k], \hat{\boldsymbol{\kappa}}^{+}[k]\right] \mid (\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^{n}\right\} \leq \alpha$$

CQR calibration – losses

Target: $\alpha \approx 4{,}6\%$ (2σ-confidence)

$n = 100$ (size of the calibration set)

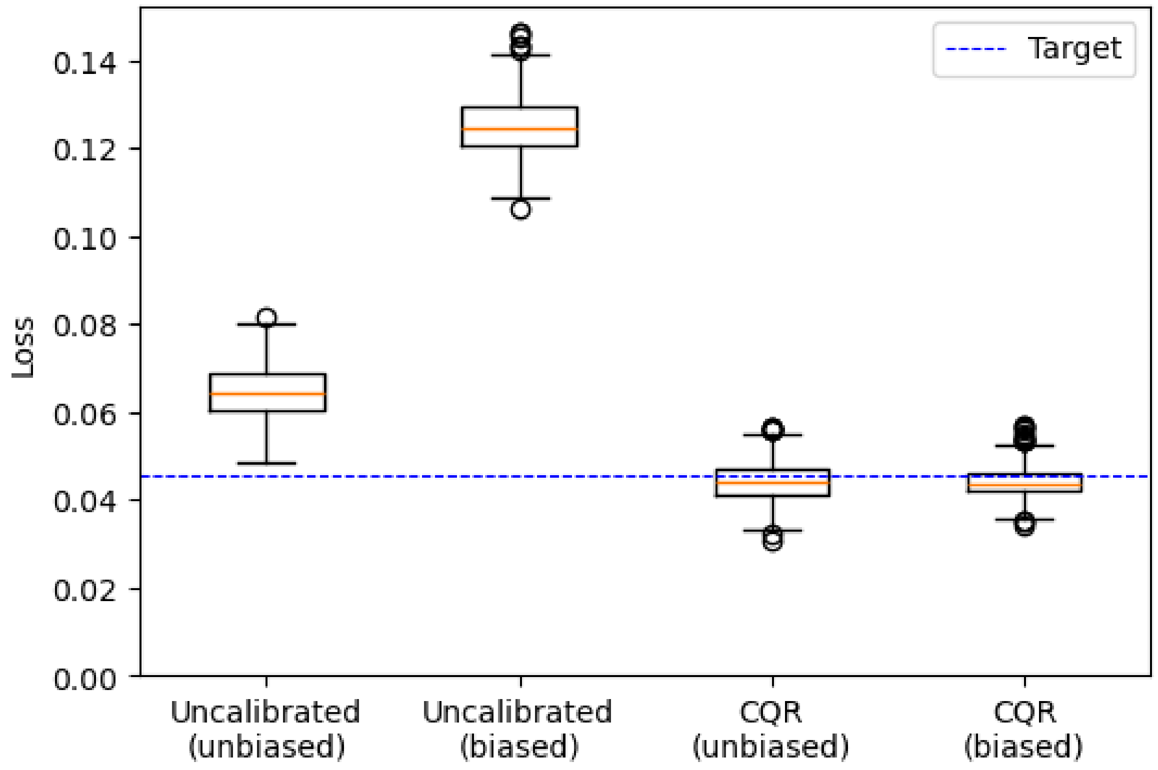

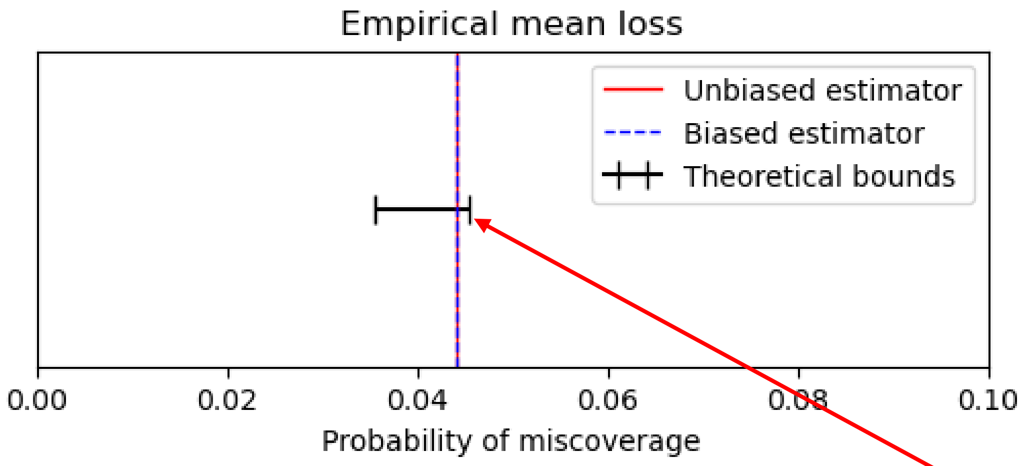

$$\alpha - {}^{1}/_{n+1} \leq \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^{-}[k], \hat{\boldsymbol{\kappa}}^{+}[k]\right] \mid (\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^{n}\right\} \leq \alpha$$

CQR calibration – losses

Target: $\alpha \approx 4{,}6\%$ (2σ-confidence)

$n = 100$ (size of the calibration set)

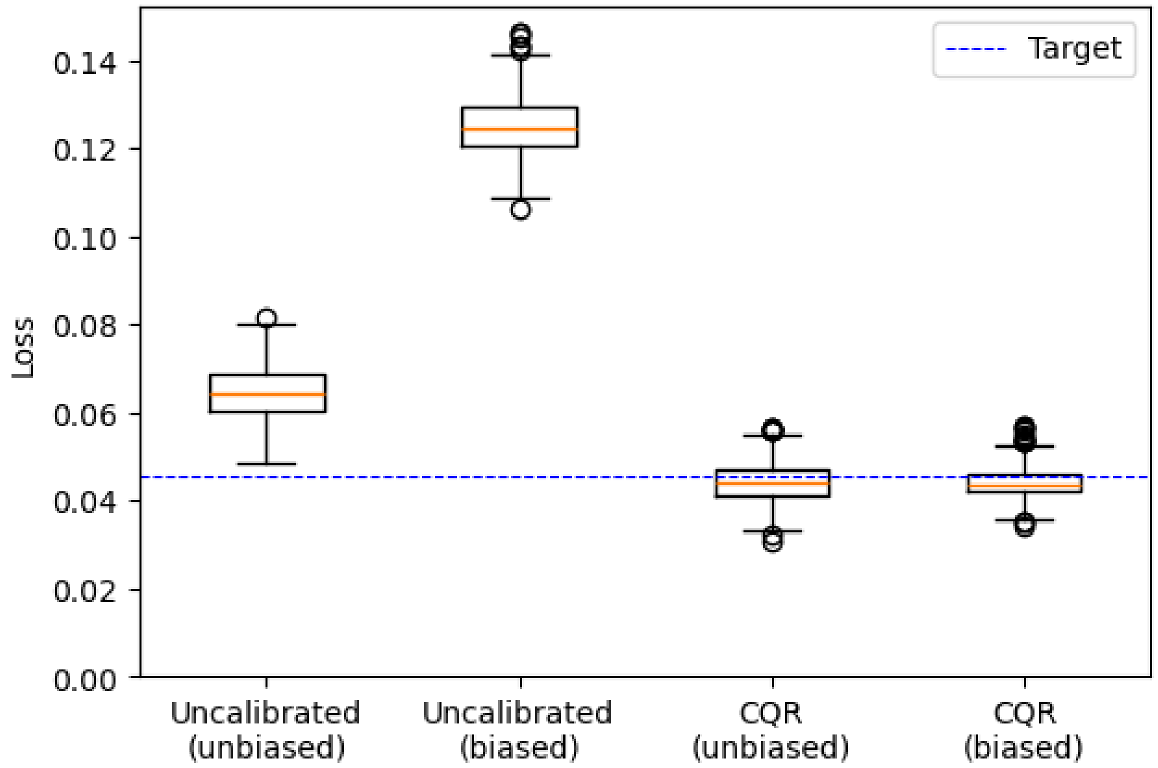

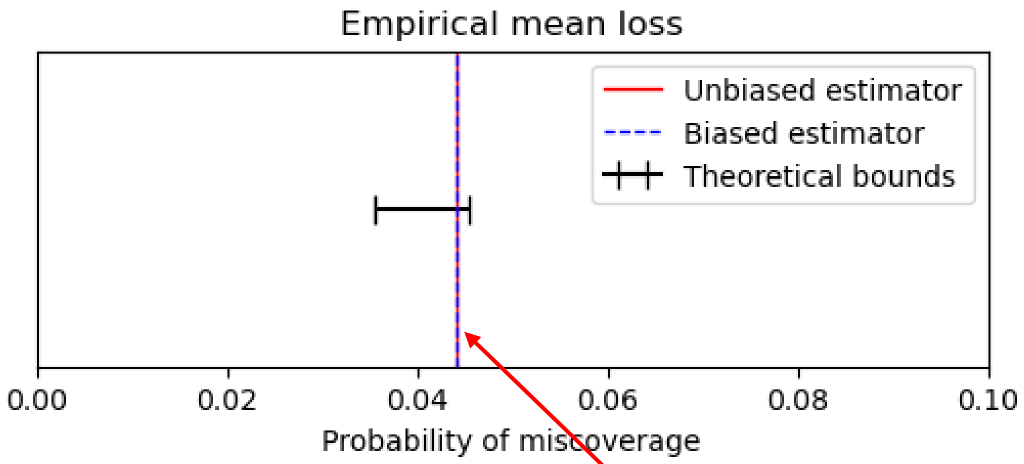

$$\alpha - {}^{1}/_{n+1} \leq \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\widehat{\boldsymbol{\kappa}}^{-}[k], \widehat{\boldsymbol{\kappa}}^{+}[k]\right] \mid (\widehat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^{n}\right\} \leq \alpha$$

CQR calibration – losses

Target: $\alpha \approx 4{,}6\%$ (2σ-confidence)

$n = 100$ (size of the calibration set)

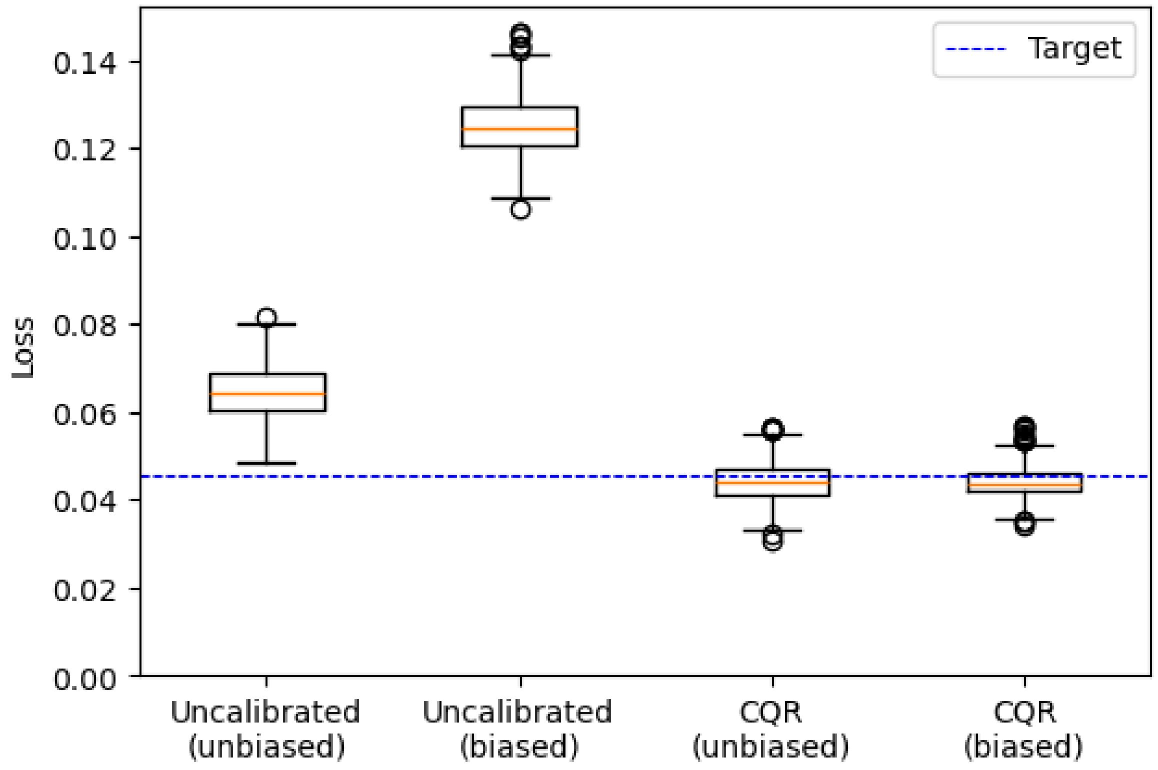

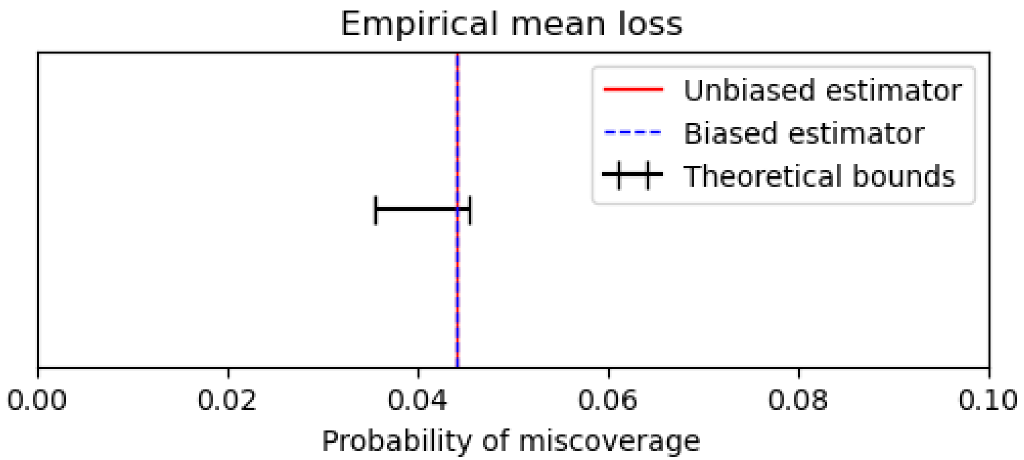

$$\alpha - {}^{1}/_{n+1} \leq \mathrm{P}\{\boldsymbol{\kappa}[k] \notin [\widehat{\boldsymbol{\kappa}}^{-}[k], \widehat{\boldsymbol{\kappa}}^{+}[k]] \mid (\widehat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^{n}\} \leq \alpha$$

→ These experiments are in line with the theoretical results.

CQR calibration – length of prediction sets

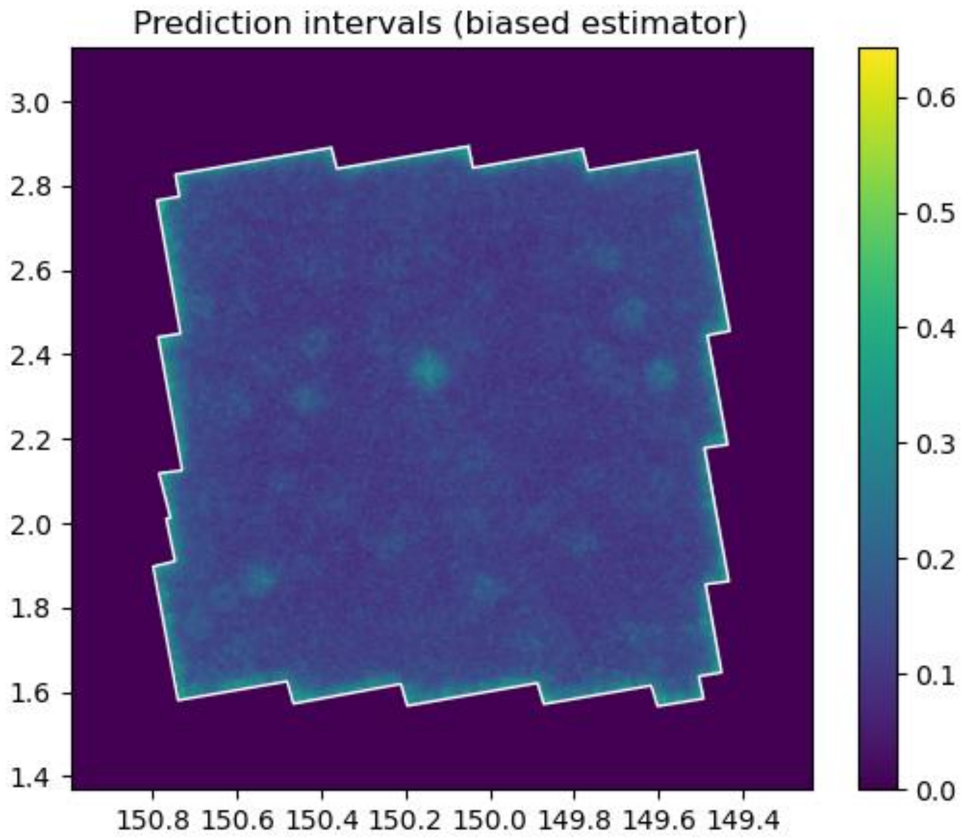

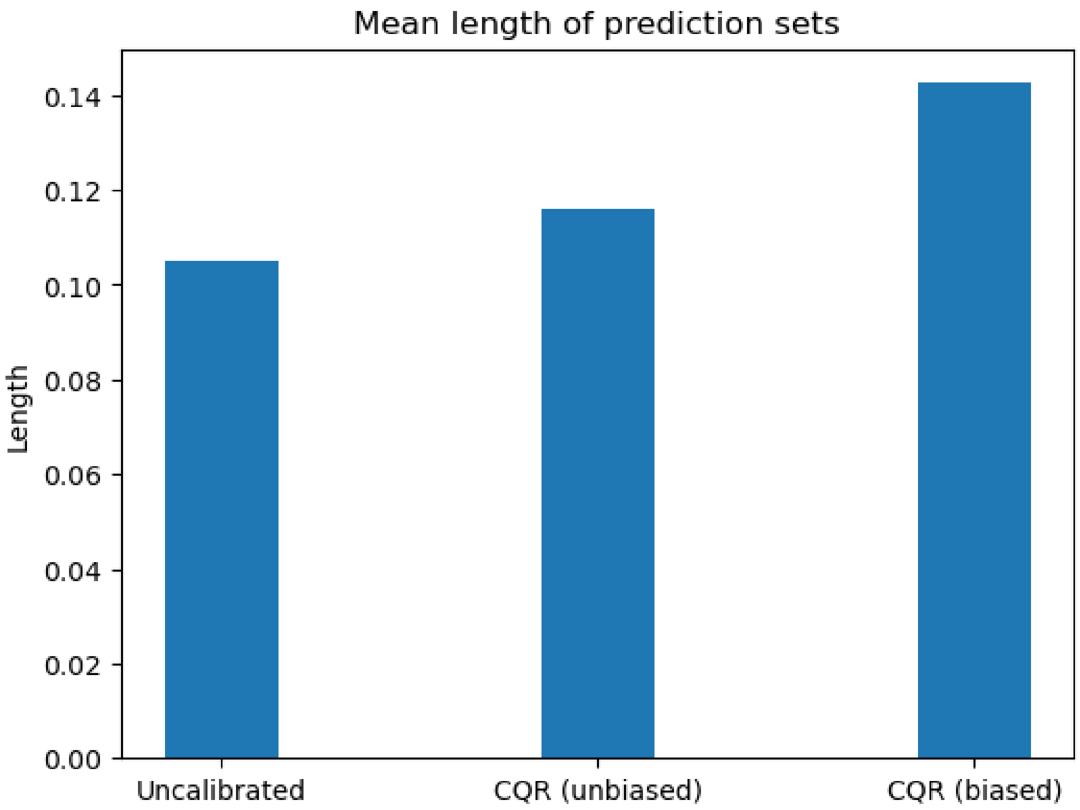

CQR calibration – error rate per pixel

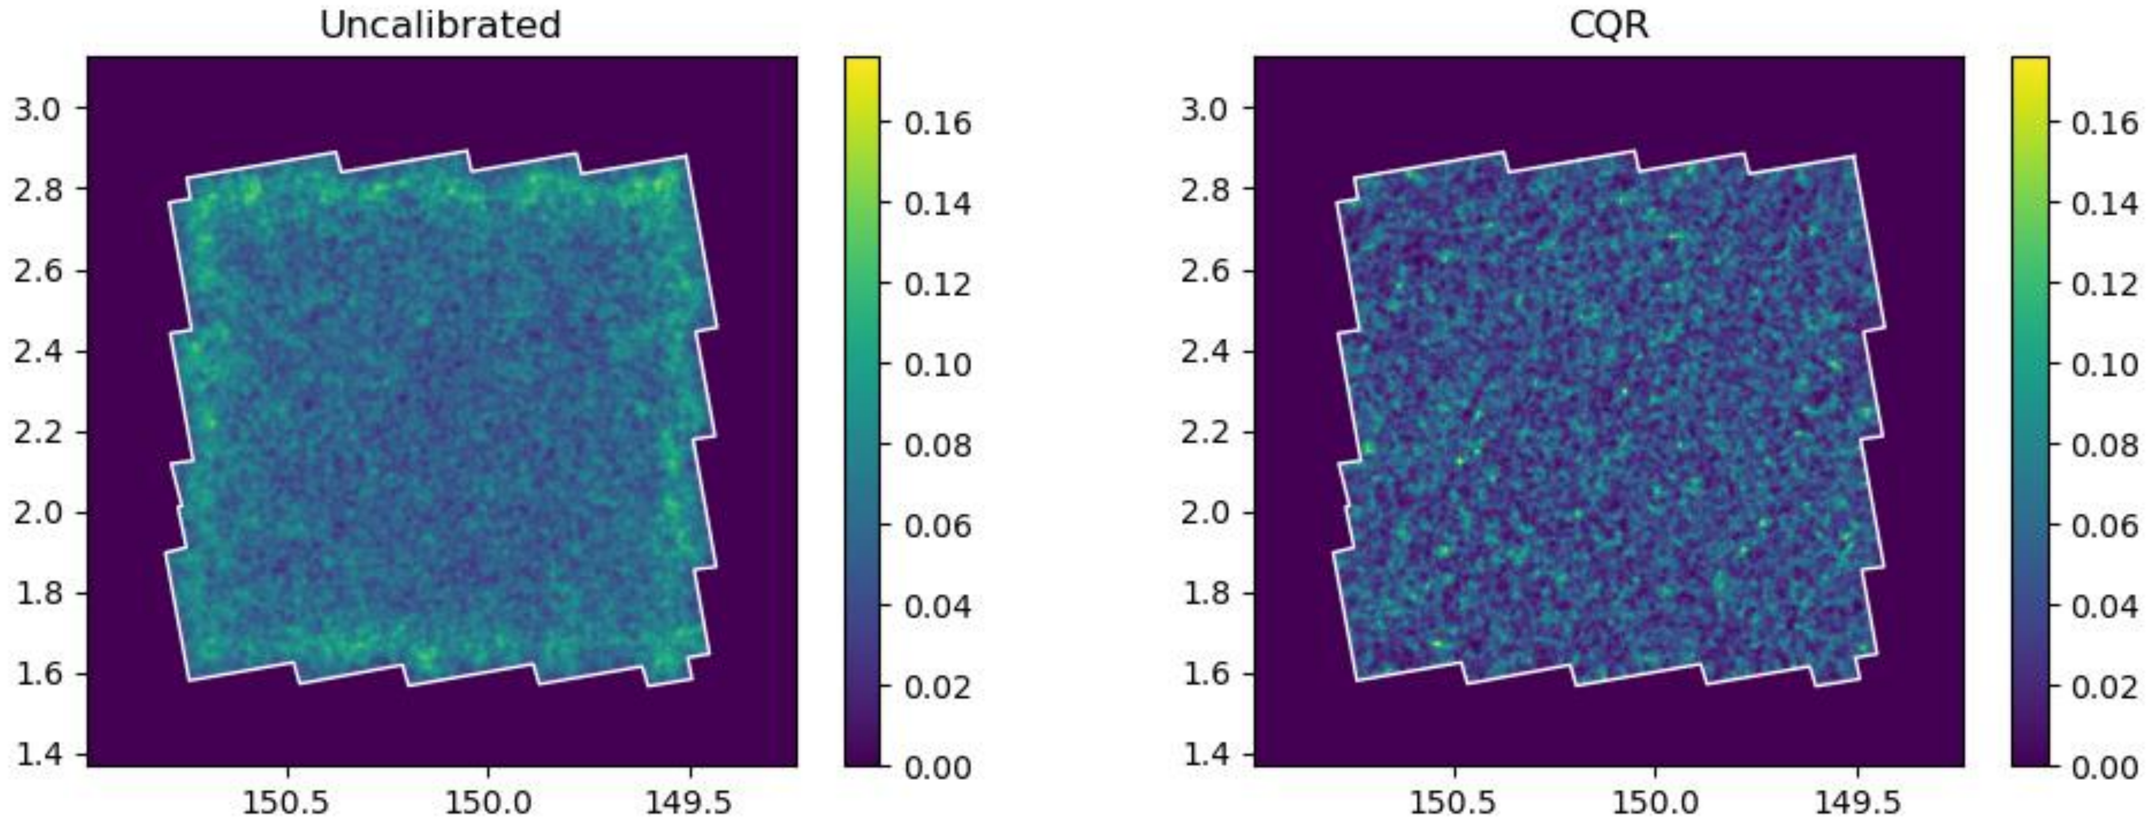

Ground truth = $\boldsymbol{S\kappa}$ (unbiased estimator)

CQR calibration – error rate per pixel

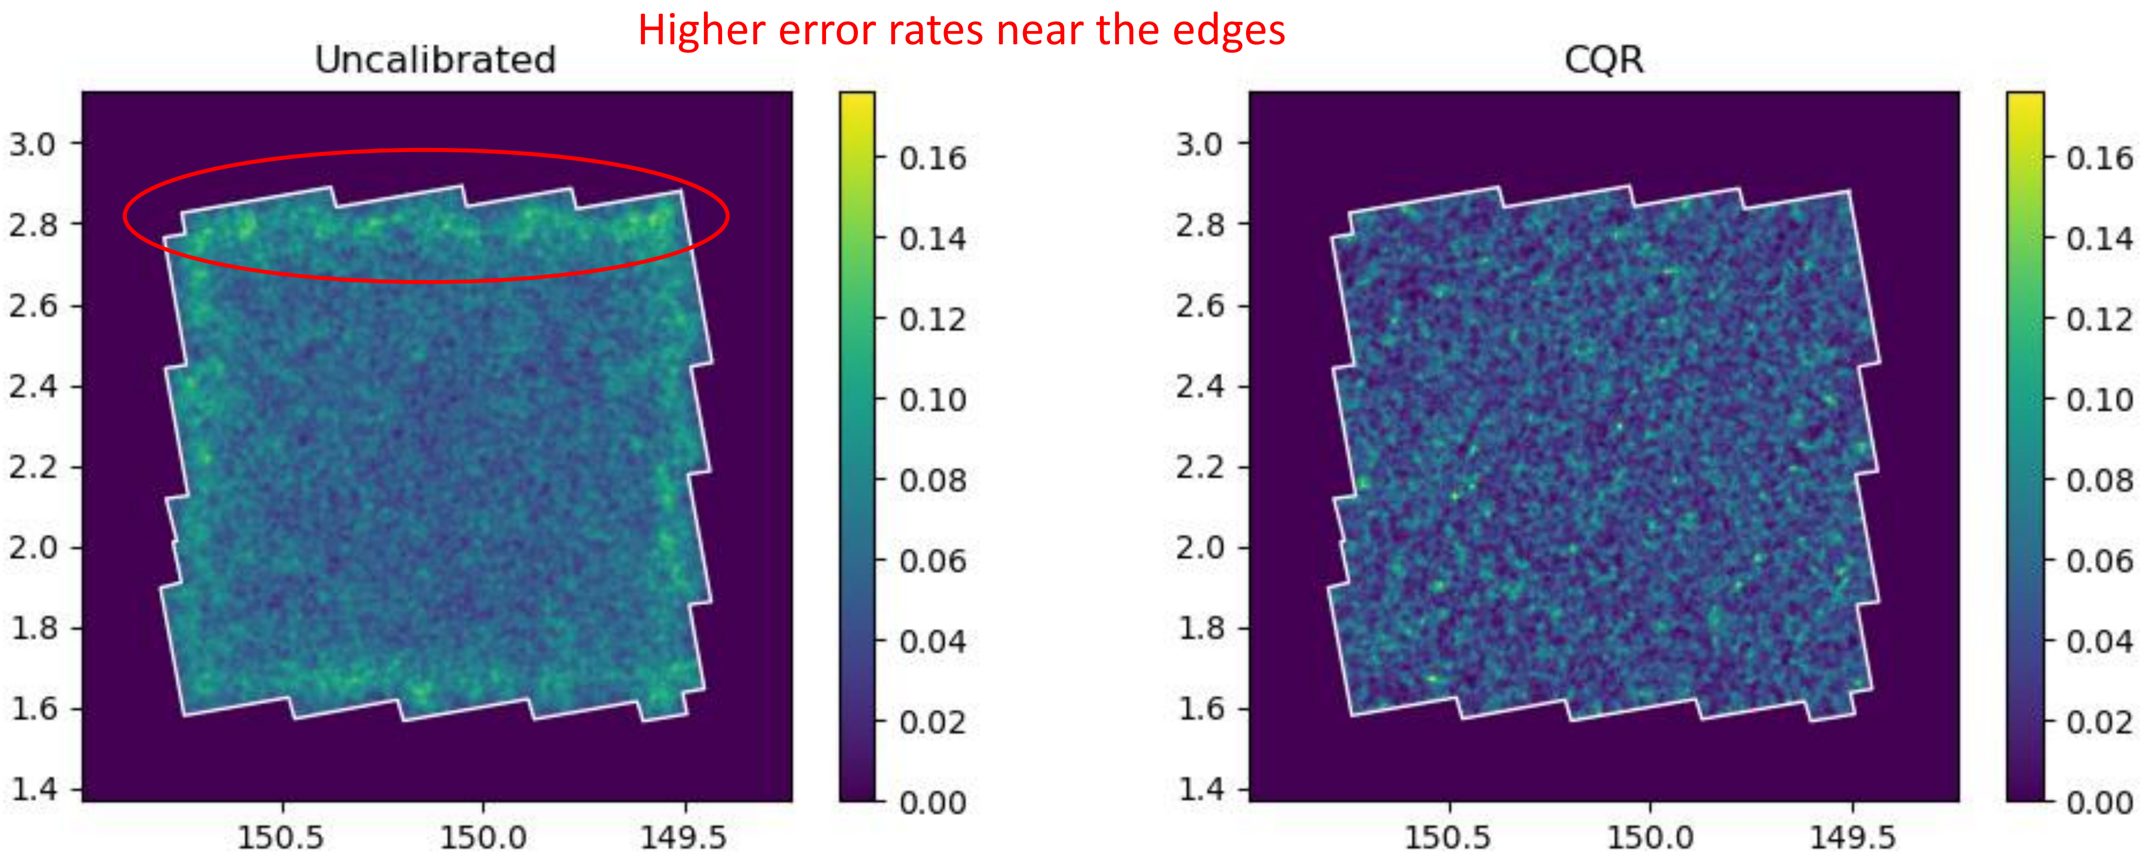

Ground truth = $\boldsymbol{S\kappa}$ (unbiased estimator)

CQR calibration – error rate per pixel

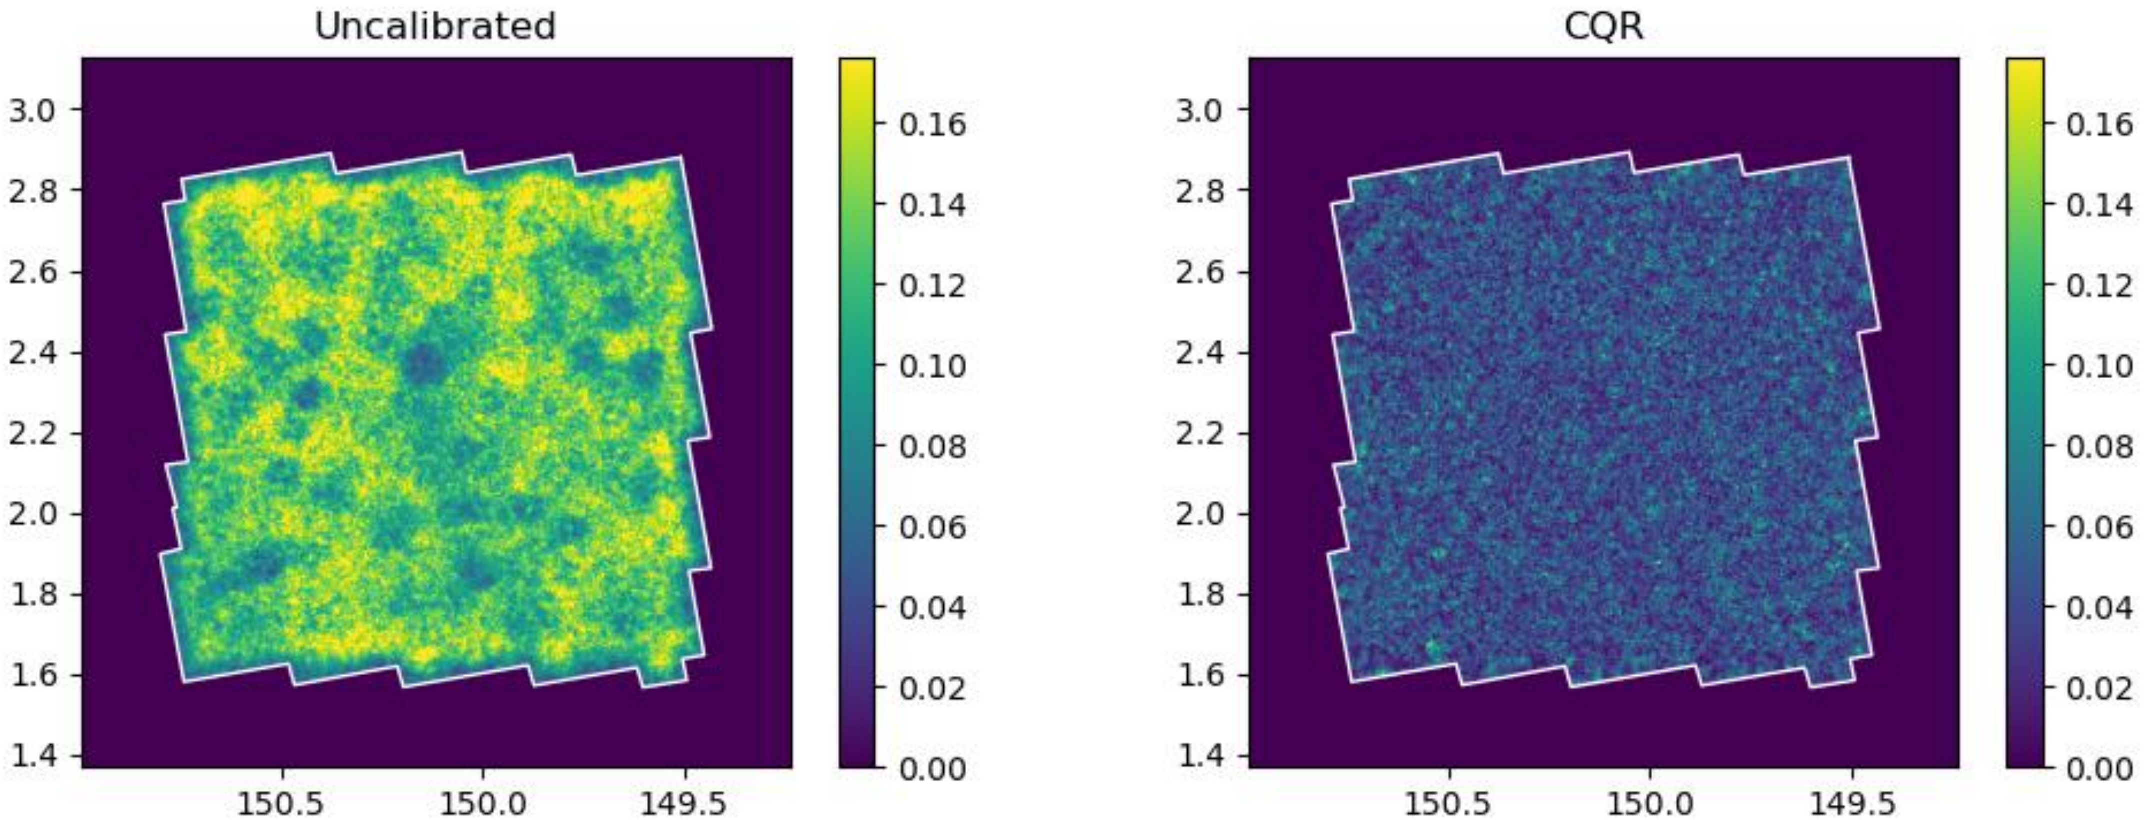

Ground truth = $\boldsymbol{\kappa}$ (biased estimator)

CQR calibration – error rate per pixel

Lower error rates around
masked regions (overcoverage)

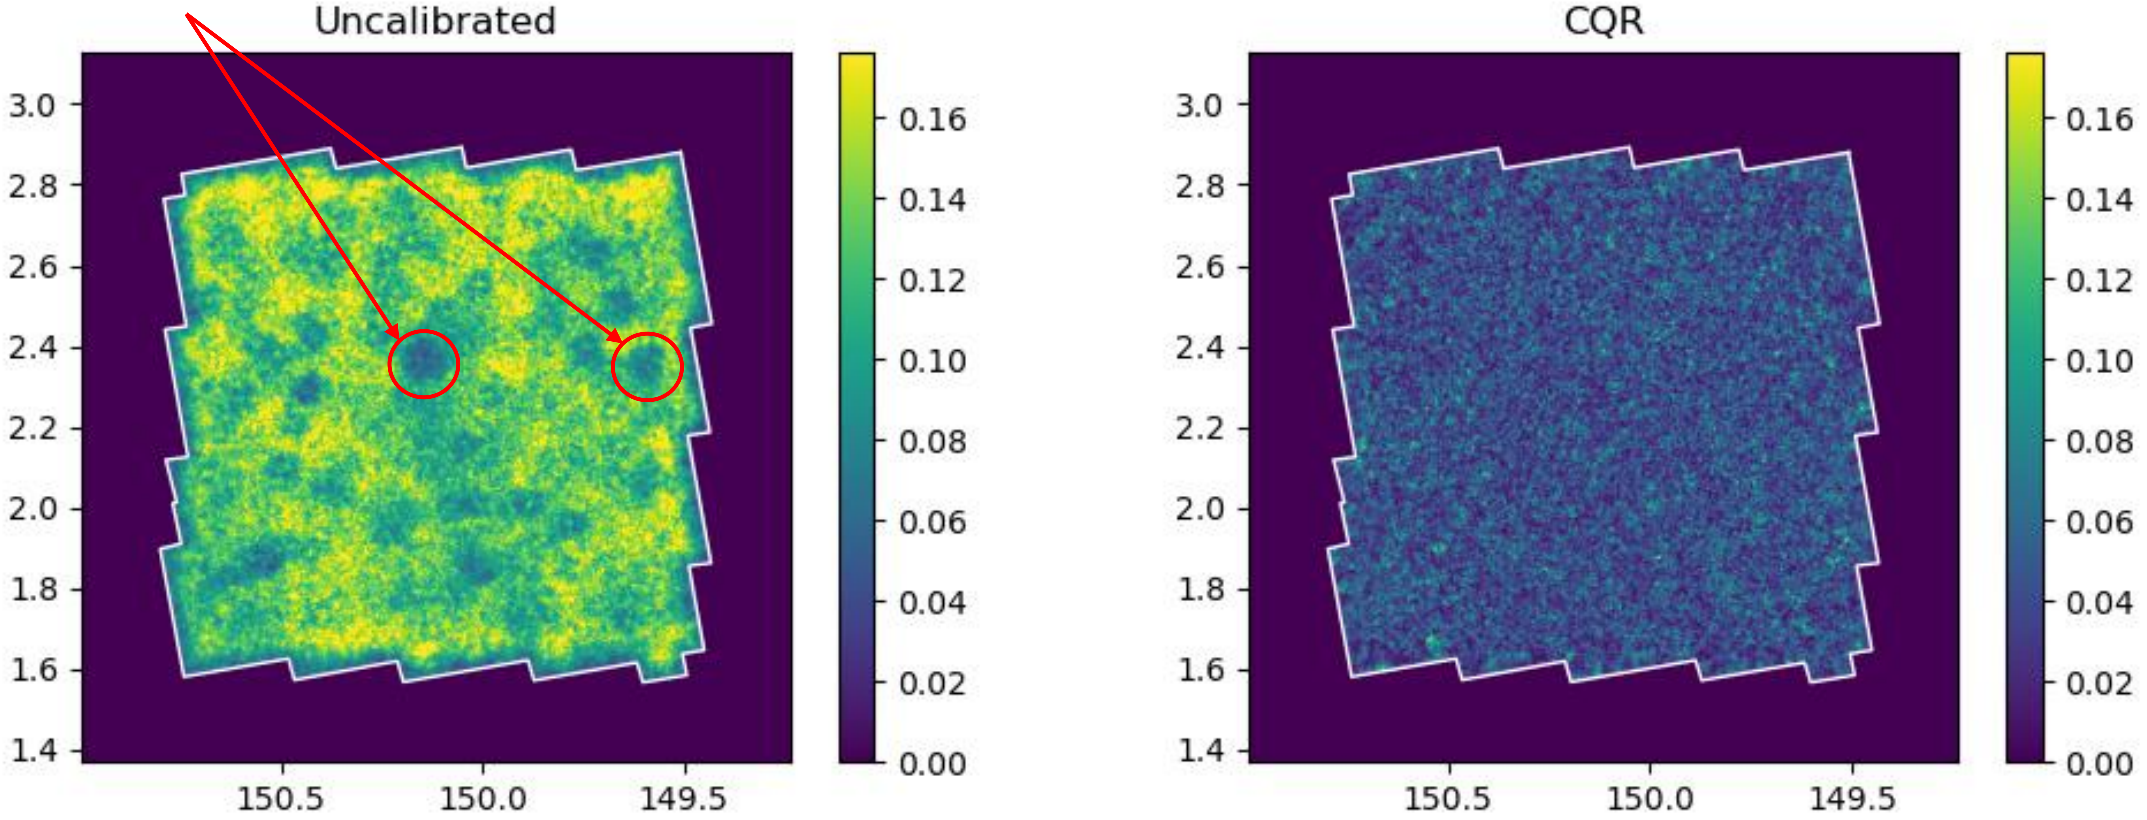

Ground truth = $\boldsymbol{\kappa}$ (biased estimator)

CQR calibration – still biased??

- Back to the theoretical guarantee:

$$\alpha - {}^{1}/_{n+1} \le \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^{-}[k], \hat{\boldsymbol{\kappa}}^{+}[k]\right] \mid (\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^{n}\right\} \le \alpha.$$

- Conditioned on the calibration set, but NOT on the neighboring pixels. E.g.,

$$\alpha - {}^{1}/_{n+1} \le \mathrm{P}\left\{\boldsymbol{\kappa}[k] \notin \left[\hat{\boldsymbol{\kappa}}^{-}[k], \hat{\boldsymbol{\kappa}}^{+}[k]\right] \mid \mathbf{S}\boldsymbol{\kappa}[k], (\hat{\boldsymbol{\gamma}}_i, \boldsymbol{\kappa}_i)_{i=1}^{n}\right\} \le \alpha.$$

→ CQR could fail to provide valid coverage guarantees in the regions of interest!

CQR calibration – still biased??

Ratio of pixels falling outside the predicted bounds: $\boldsymbol{\kappa}[k] \notin [\hat{\boldsymbol{\kappa}}^{-}[k], \hat{\boldsymbol{\kappa}}^{+}[k]]$, conditioned on the value of $\mathbf{S}\boldsymbol{\kappa}[k]$.

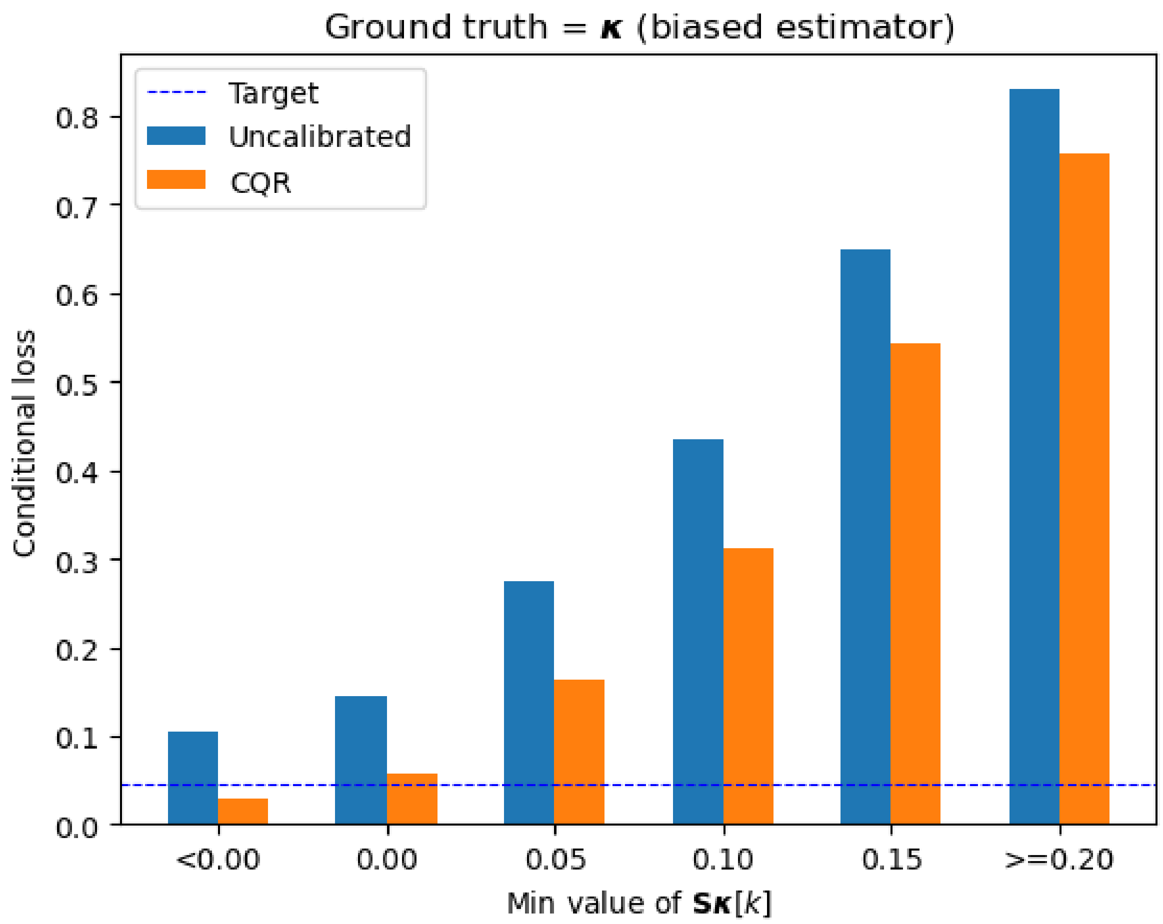

CQR calibration – still biased??

Ratio of pixels falling outside the predicted bounds: $\boldsymbol{\kappa}[k] \notin [\widehat{\boldsymbol{\kappa}}^{-}[k], \widehat{\boldsymbol{\kappa}}^{+}[k]]$, conditioned on the value of $\mathbf{S}\boldsymbol{\kappa}[k]$.

Way above target for large values of convergence!

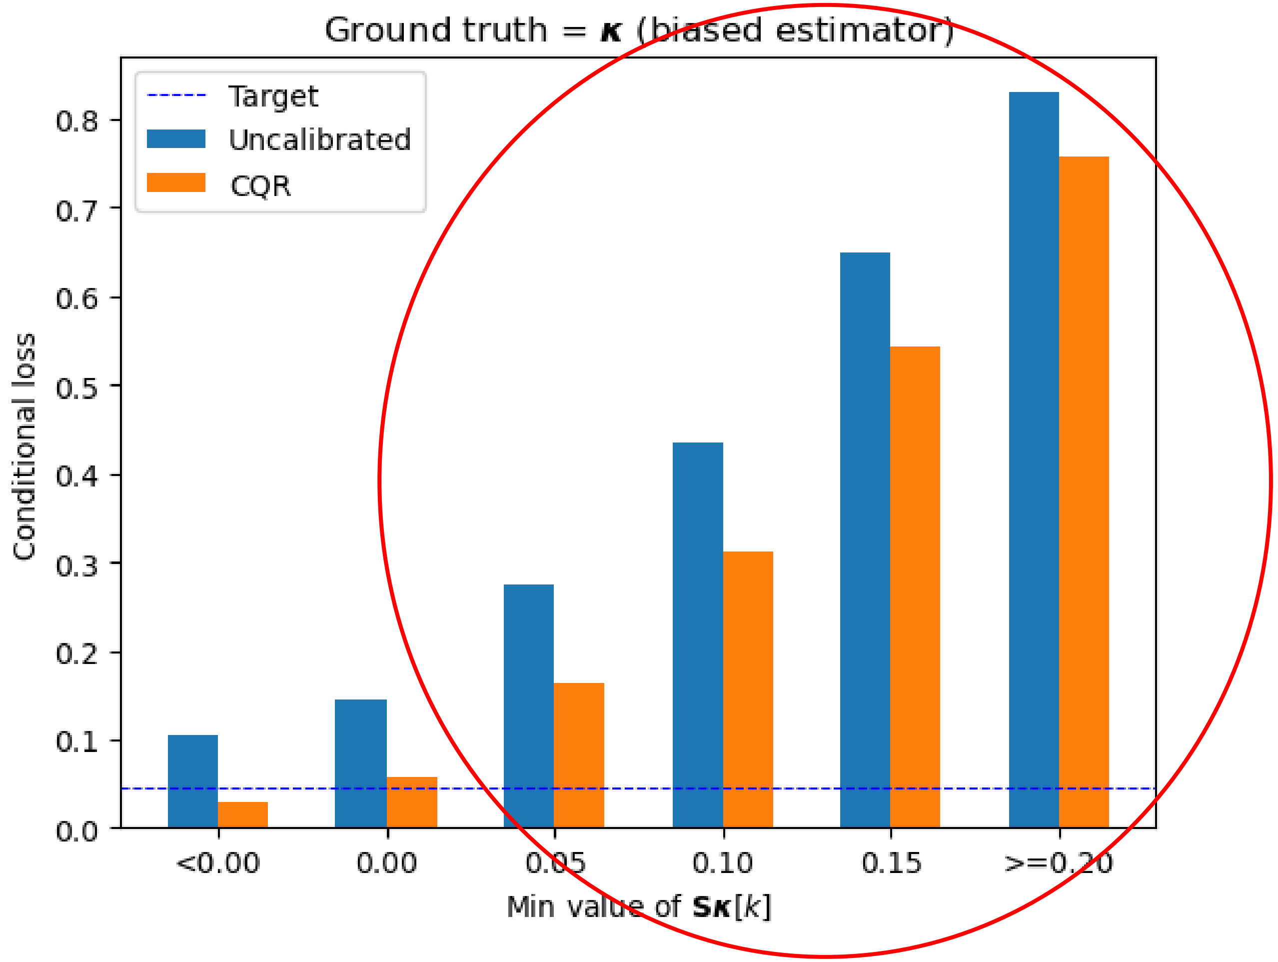

Next steps

- Run similar experiments on more advanced reconstruction methods.
- In particular, apply plug-and-play methods to this framework?
- Adapt the results to simulated galaxy shape catalogs predicting results from SKA (T-RECS).[1]
- Can we get coverage guarantees conditionally to the neighboring pixels? E.g., around peak values.
- Detection of extensive objects: existing work based on prediction masks.[2]
- Other avenue: exploit correlation between pixels?[3]

[1] A. Bonaldi et al., "The Tiered Radio Extragalactic Continuum Simulation (T-RECS)," Monthly Notices of the Royal Astronomical Society, vol. 482, no. 1, pp. 2–19, Jan. 2019.

[2] G. Kutiel, R. Cohen, M. Elad, and D. Freedman, "What's Behind the Mask: Estimating Uncertainty in Image-to-Image Problems." arXiv, Nov. 28, 2022.

[3] O. Belhasin, Y. Romano, D. Freedman, E. Rivlin, and M. Elad, "Principal Uncertainty Quantification with Spatial Correlation for Image Restoration Problems." arXiv, May 17, 2023.