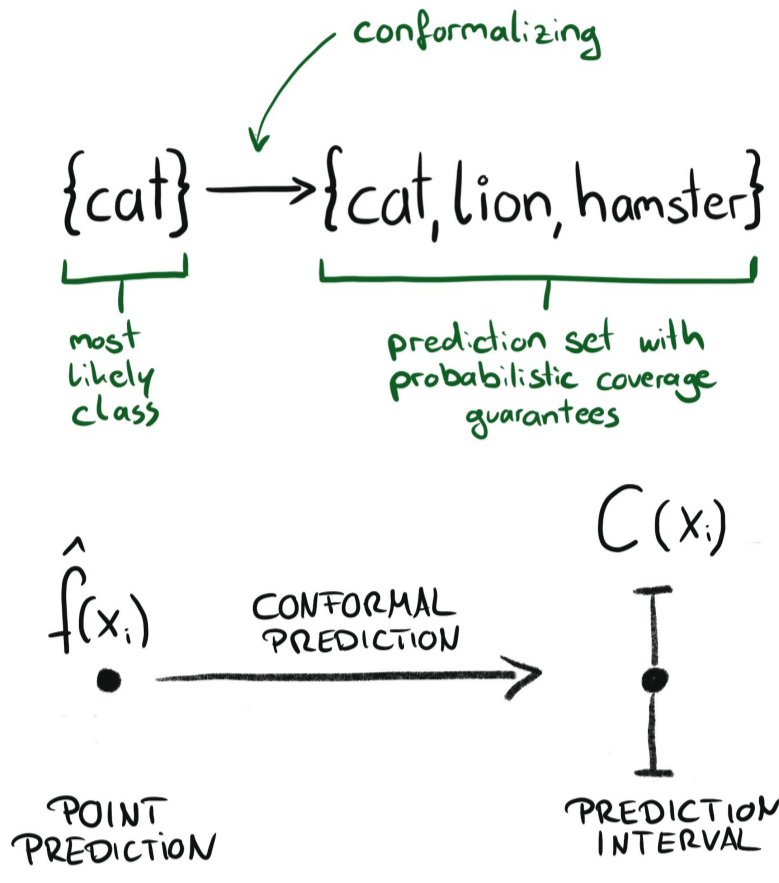


1. Introduction to Conformal Prediction

Conformal prediction (CP) [1] builds prediction regions that achieve the desired marginal coverage $1 - \alpha$ under nothing more than **exchangeability** - no other distributional assumptions.



A non-conformity score $S(x, y)$ is evaluated on an independent calibration set D_{cal} . The prediction set is

$$C(x) = \{y: S(x, y) \leq \hat{q}_{1-\alpha}\}. \quad (1)$$

where $\hat{q}_{1-\alpha}$ is the empirical $(1-\alpha)$ conformal quantile of the calibration scores.

2. Different Types of Coverage Guarantee

• Marginal Coverage Guarantee

$$P_{(X,Y)}(Y_{\{n+1\}} \in C(X_{\{n+1\}})) \geq 1 - \alpha. \quad (2)$$

• PAC-style Coverage Guarantee

$$P_{(X,Y)} \left[P(Y_{\{n+1\}} \in C(X_{\{n+1\}})) \geq 1 - \alpha \right] \geq 1 - \epsilon. \quad (3)$$

$P_{(X,Y)}$: The probability taken over the data distribution of a random test sample (e.g., X, Y).

$1 - \alpha$: The target **coverage level** (e.g., 0.90 or 0.95) indicating the proportion of times the true value falls inside the predicted set.

$1 - \delta$: The **confidence parameter** (or probability of success) indicating how often the coverage guarantee will actually hold.

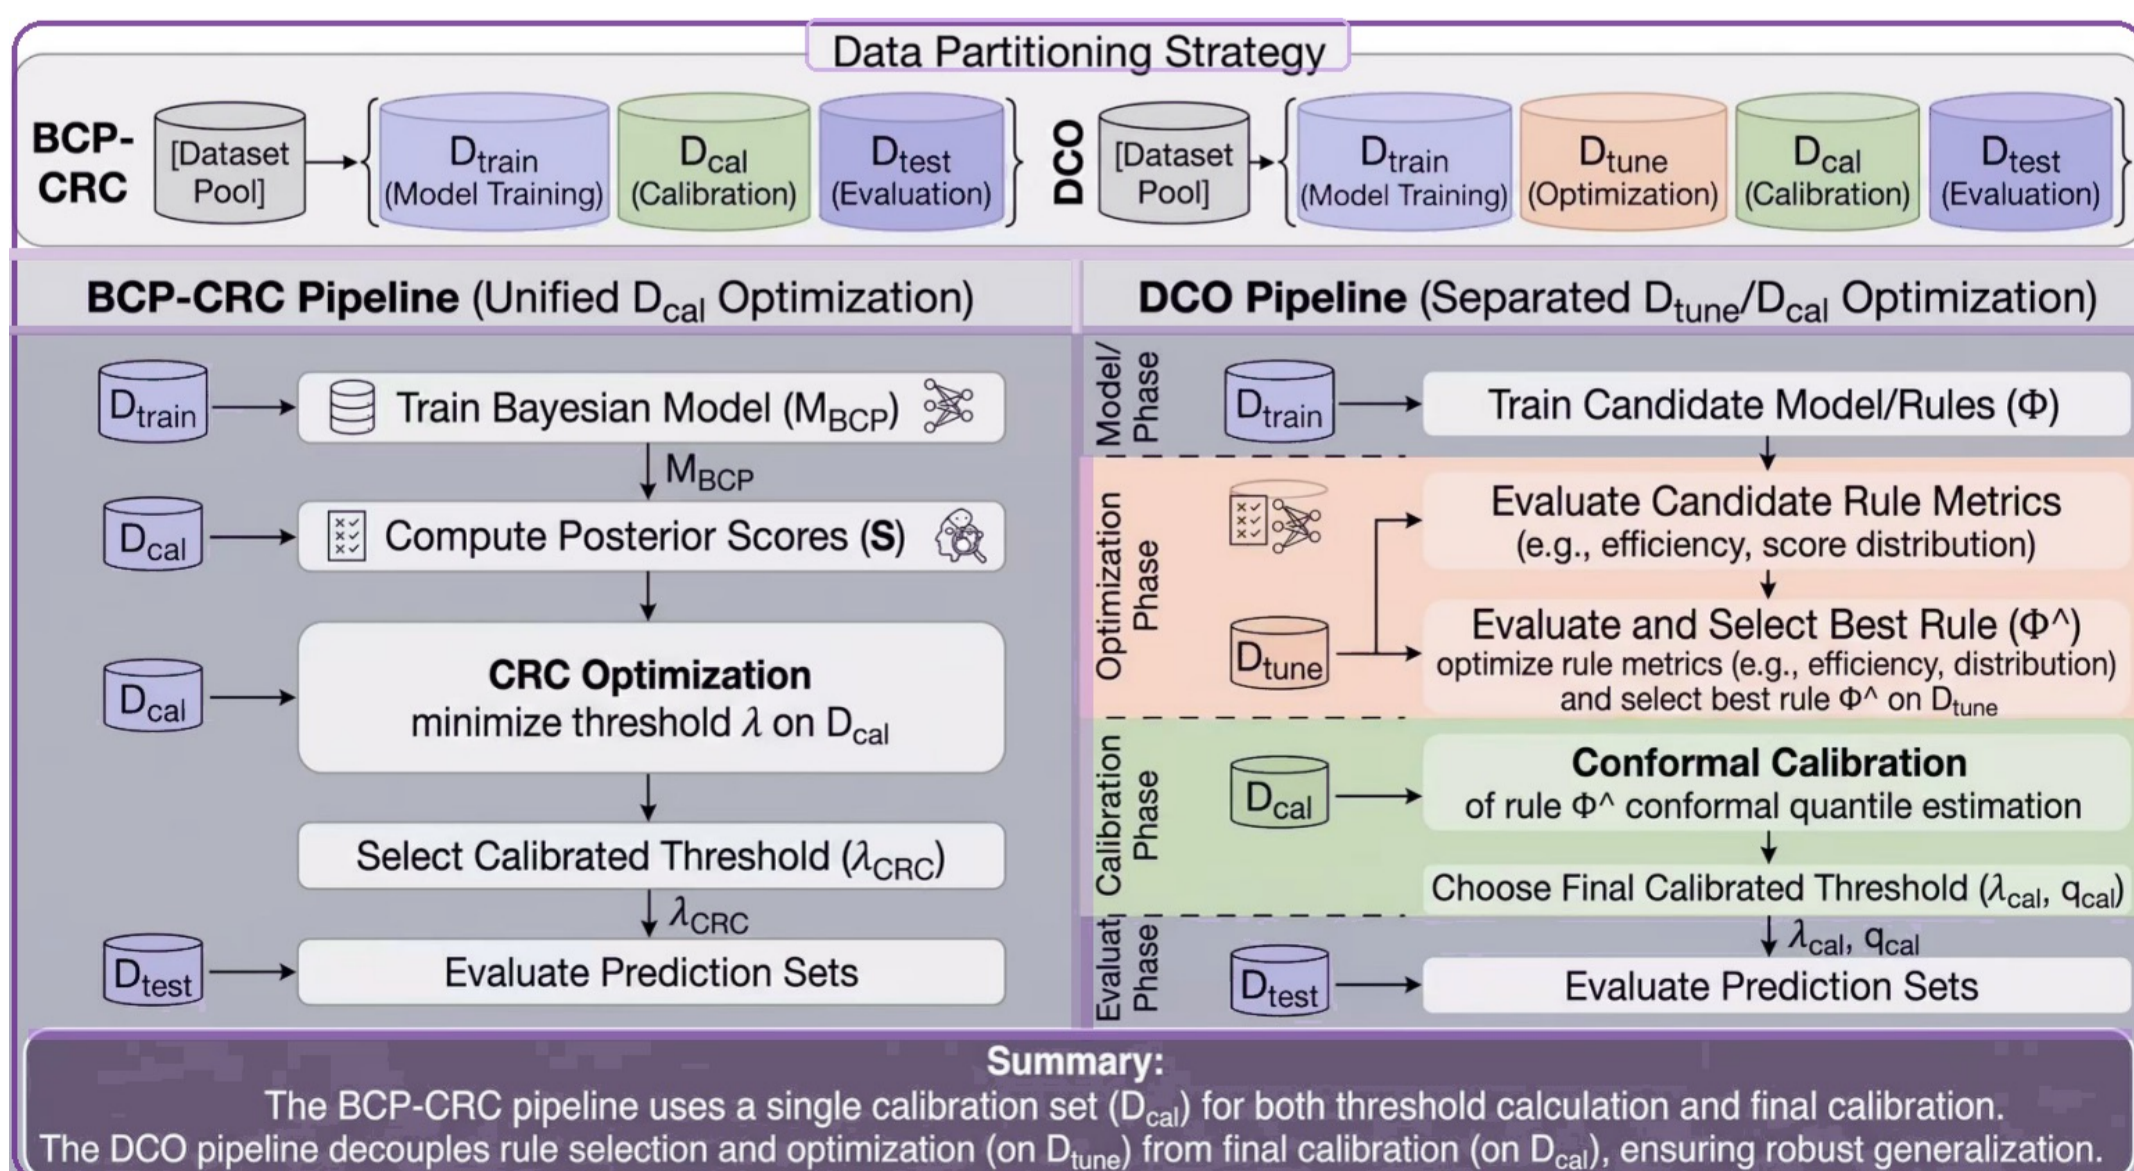
3. Optimising Conformal Prediction

Non-conformity Scores: design sharper non-conformity scores.

Bayesian integration: use posterior predictive uncertainty in the score.

Threshold search: minimise set size under a risk / mis-coverage constraint.

DCO principle: tune structure on D_{tune} , calibrate only on untouched D_{cal} .



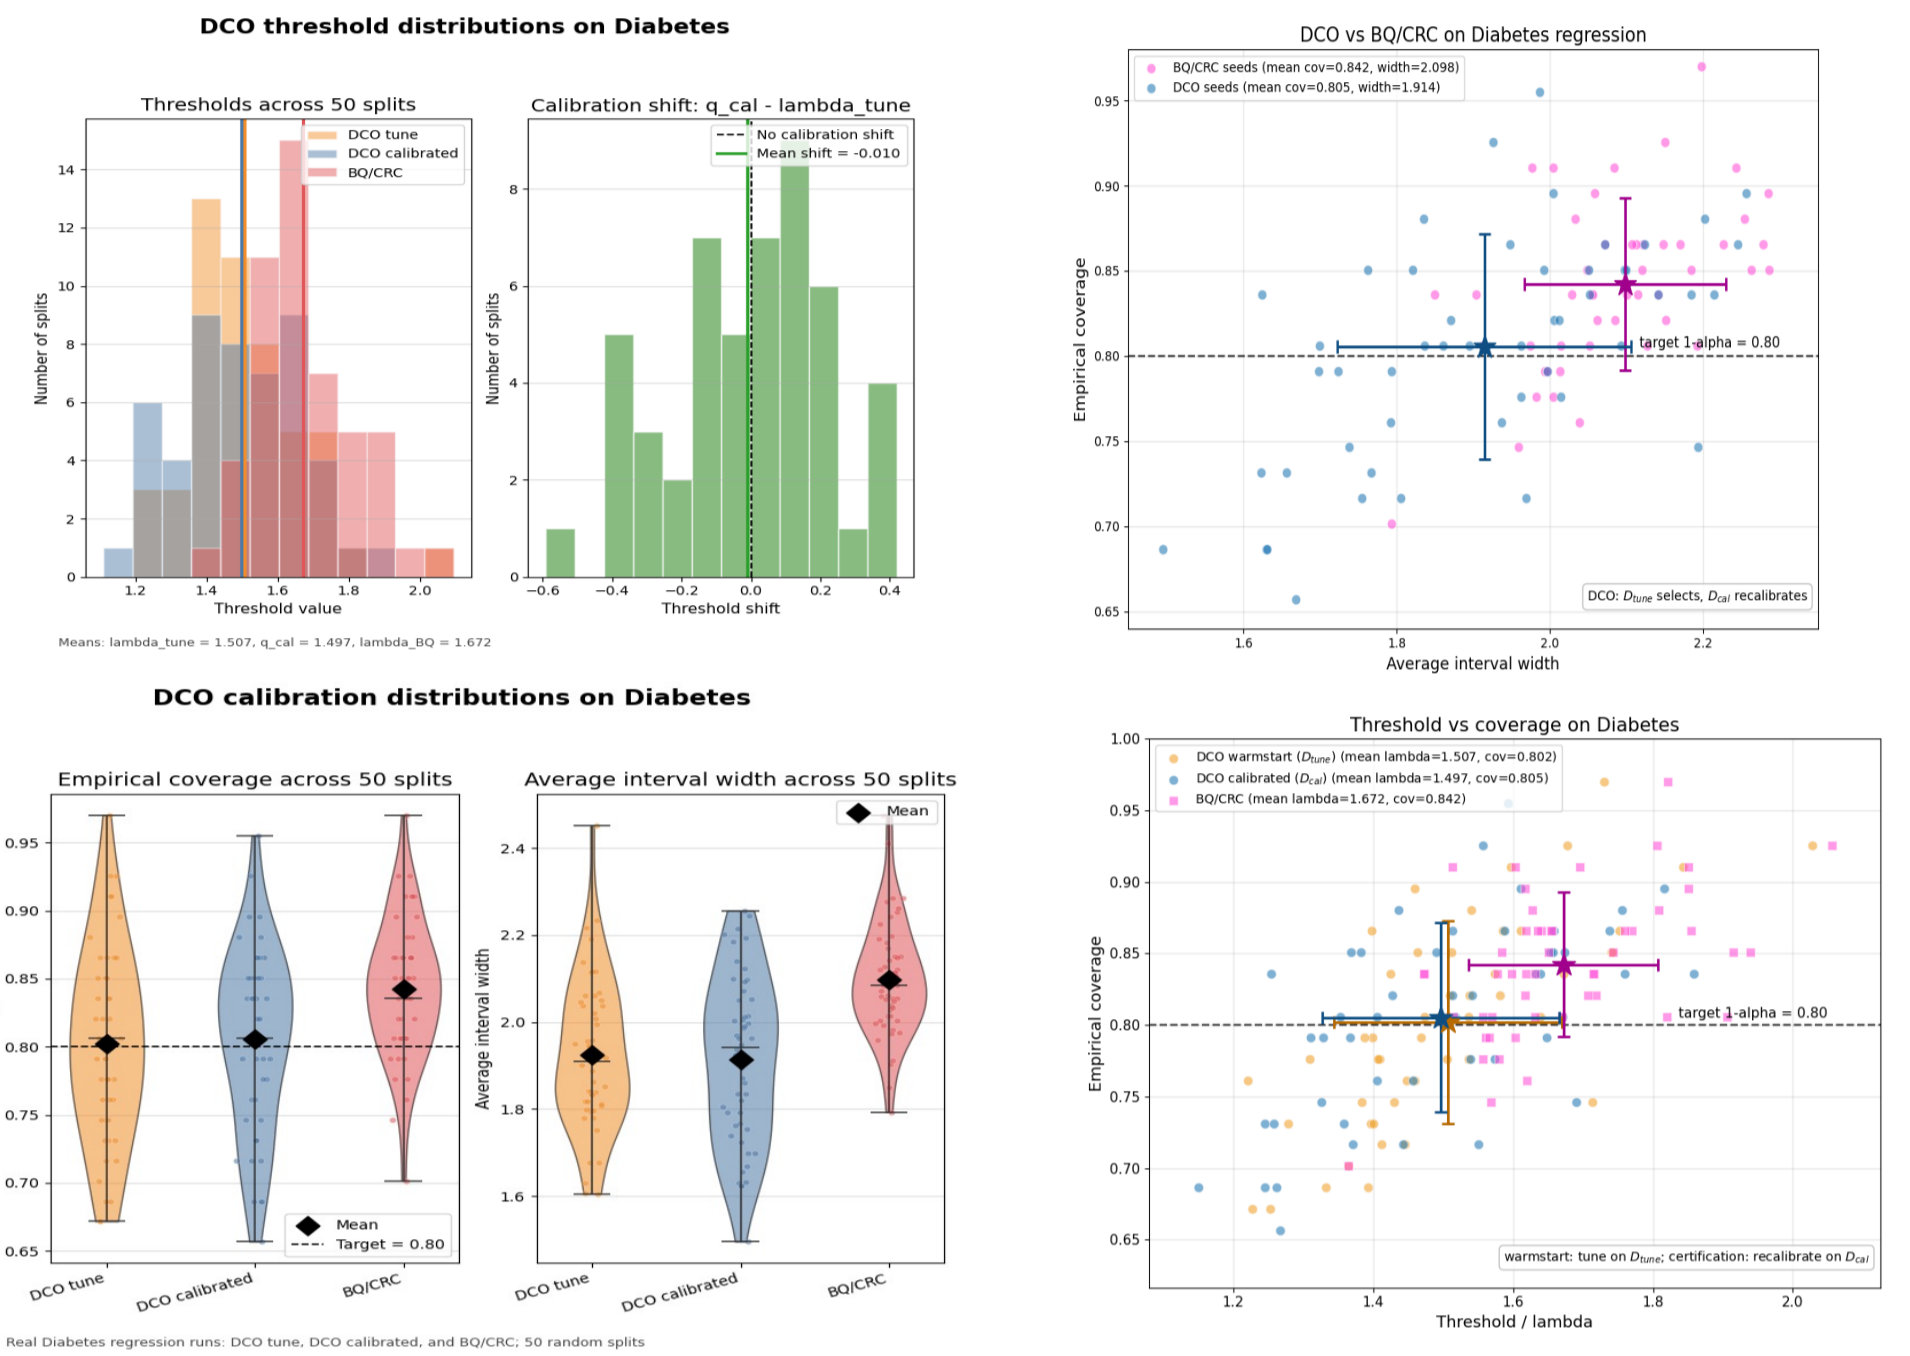
4. Regression: DCO separates tuning from calibration

- **Dataset:** Diabetes, $n = 442, d = 10$; all predictors and the response are standardized.
- **Sparse prior:**

$$\theta_j \sim \text{Laplace}(0, b), \quad \theta_0 \sim \mathcal{N}(0, 10), \quad \tau \sim \mathcal{N}^+(0, c)$$

- **Bayesian model:**

$$f_\theta(y | x) = \mathcal{N}(y | \theta^T x + \theta_0, \tau^2)$$

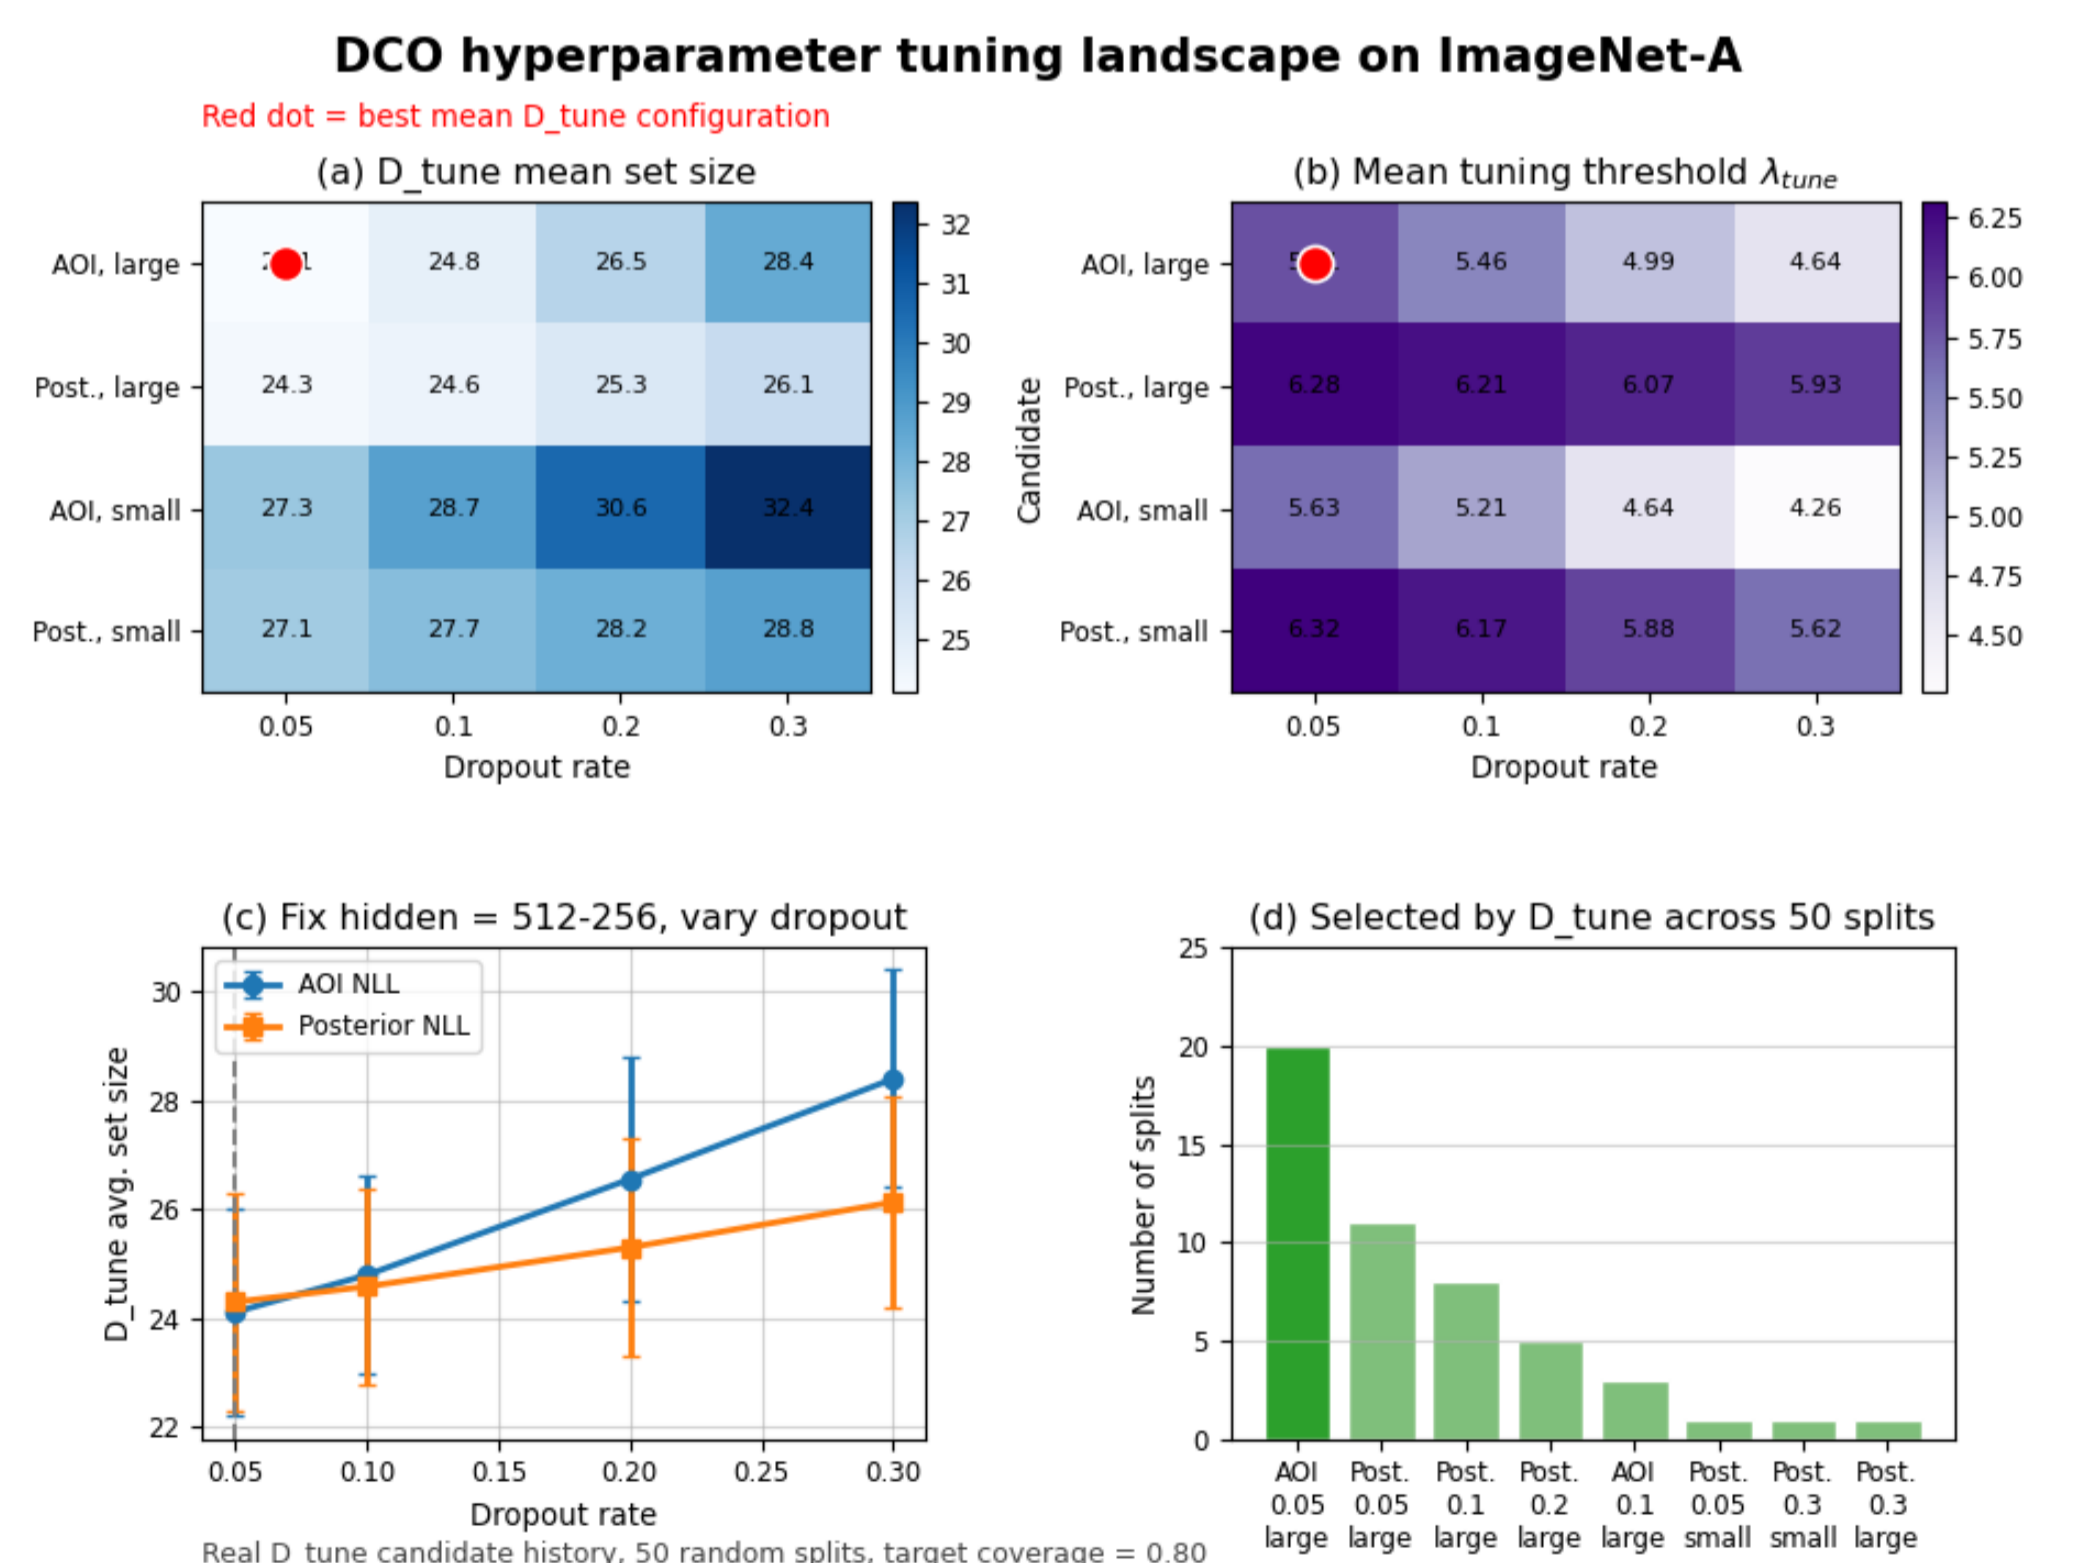


Method	Coverage (%)	Width	Certified	Notes
BQ/CRC	84.24 +/- 5.06	2.098 +/- 0.131	Yes	Conservative risk-bound calibration
Split CP	81.16 +/- 5.70	1.920 +/- 0.159	Yes	Residual split conformal baseline
CQR	80.66 +/- 5.95	1.990 +/- 0.159	Yes	Quantile-regression conformal baseline
DCO (ours)	80.54 +/- 6.62	1.914 +/- 0.192	Yes	Tune selects; calibration certifies
Direct tuning	80.21 +/- 7.08	1.924 +/- 0.173	No	Uses tuning threshold without recalibration

Diabetes regression, 50 random splits, target coverage = 80%

5. ImageNet-A Classification Example

- **Dataset:** ImageNet-A
- **Prediction set:** $C(x) = \{y: S_\phi(x, y) \leq q_{cal}\}$



Method	Coverage	Avg. Set Size
BQ/CRC	0.812 ± 0.011	26.52 ± 1.34
DCO-Warmstart	0.801 ± 0.011	25.26 ± 1.60
DirectTune	0.791 ± 0.015	23.83 ± 1.80

ImageNet-A classification, 50 random splits, target coverage = 0.80

