Localized Adaptive Risk Control

Matteo Zecchin and Osvaldo Simeone

Centre for Intelligent Information Processing Systems,

King's College London

King's Communications, Learning & Information Processing lab

Online Calibration

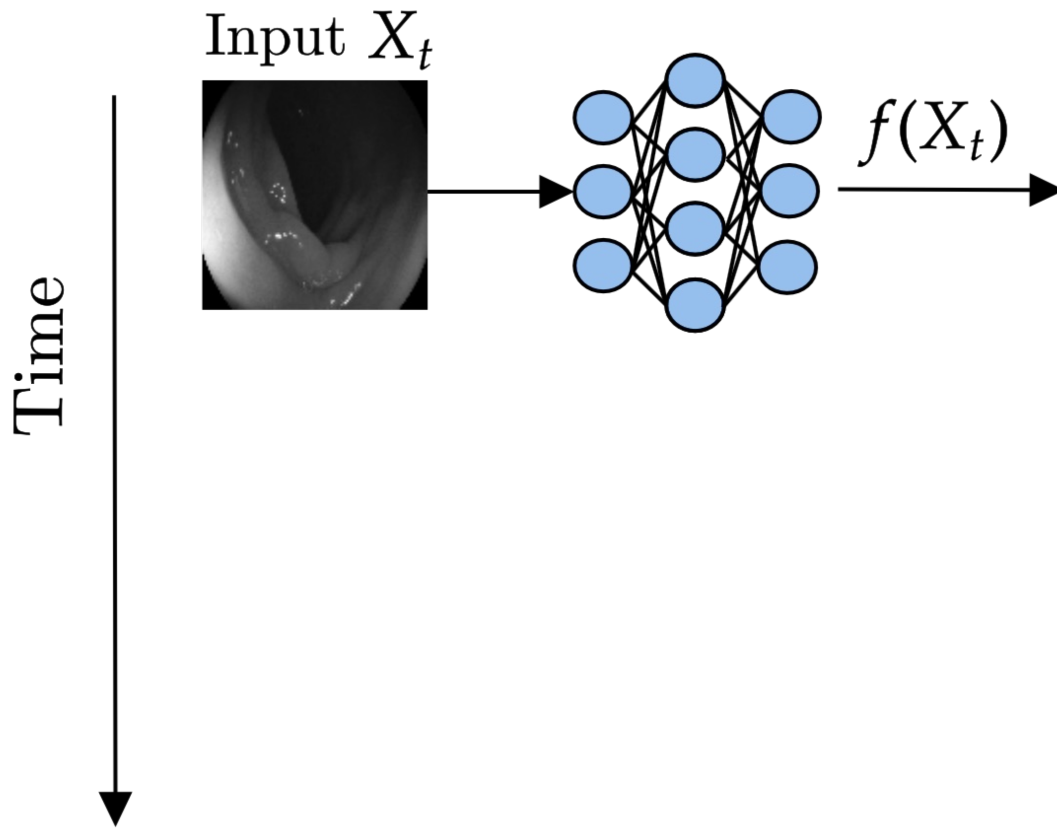

Online Calibration

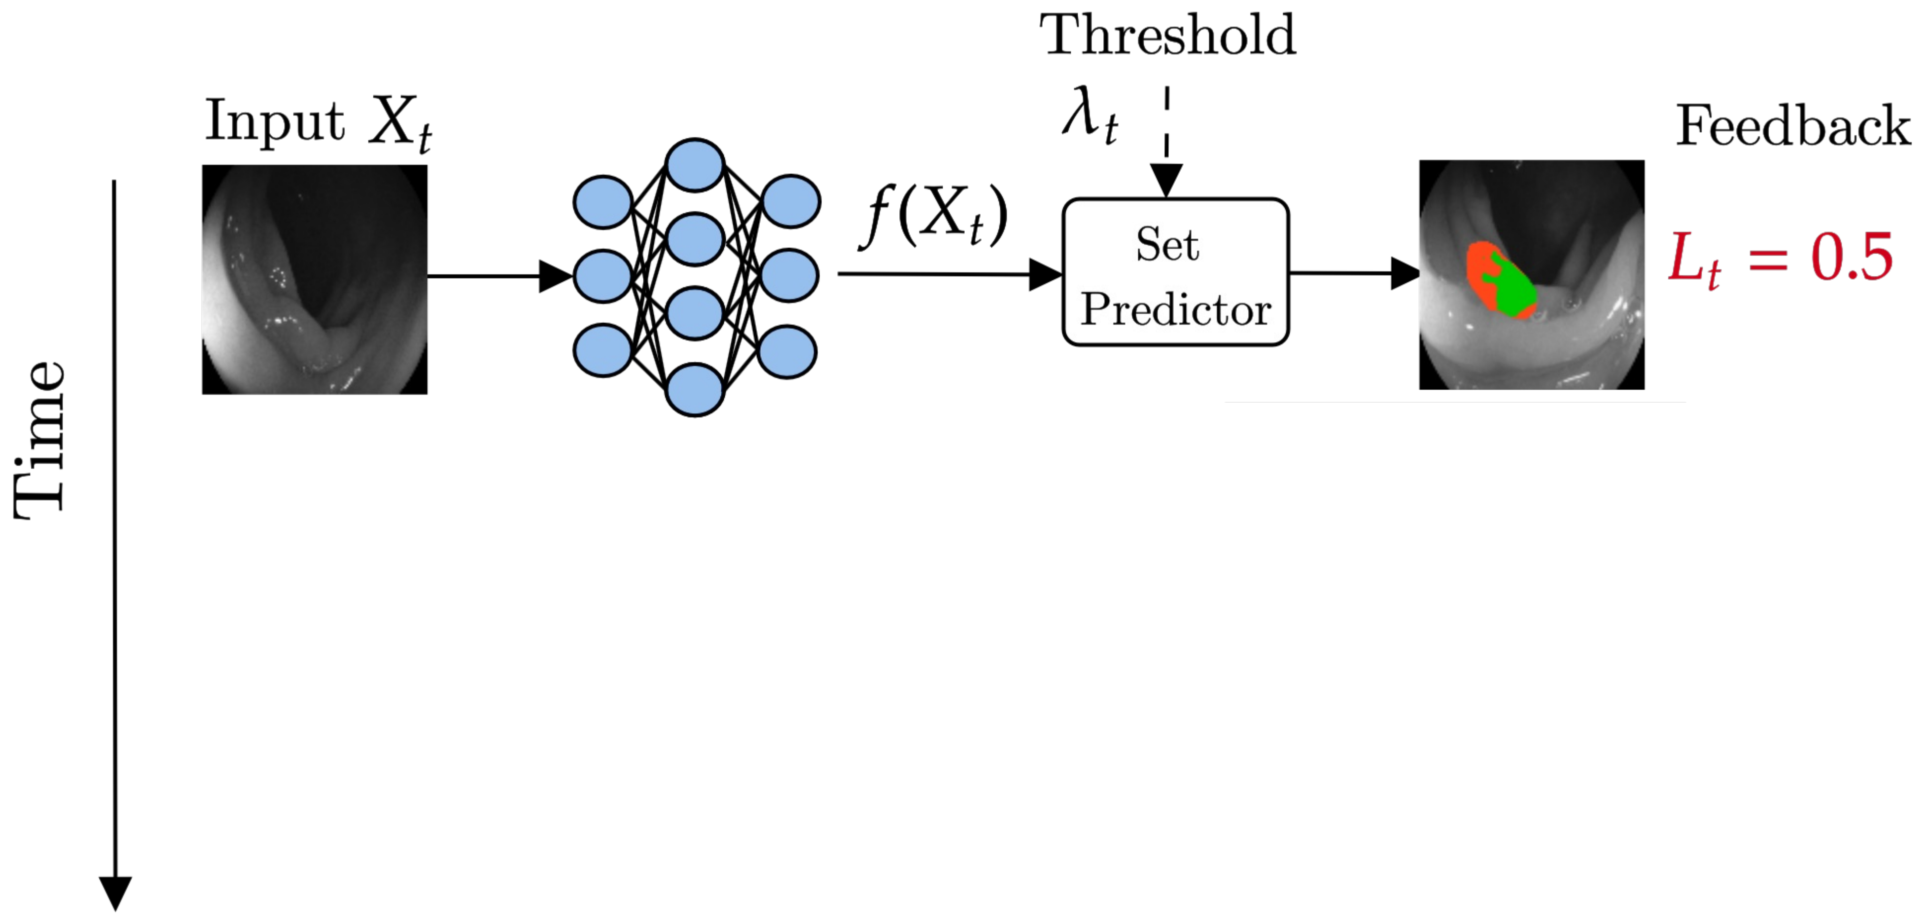

Online Calibration

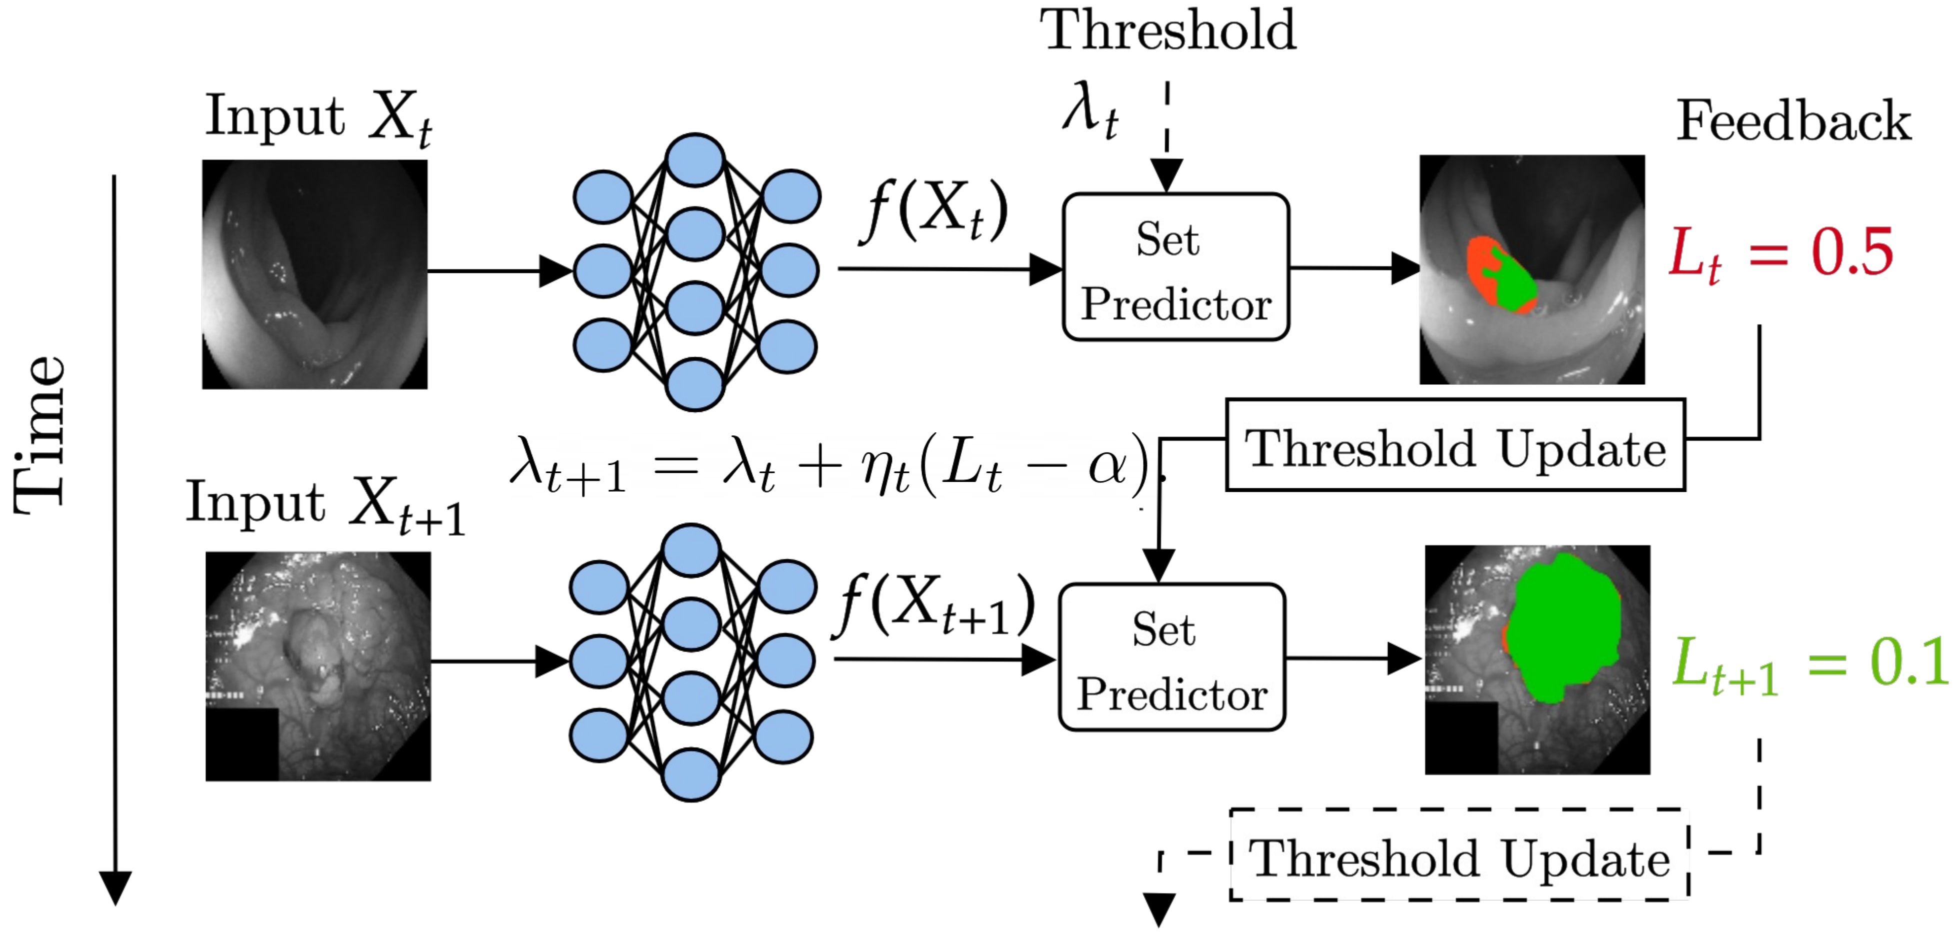

Online Calibration via Adaptive Risk Control

Worst-case deterministic guarantees (informal) [1]

For any data sequence $\{(X_t, Y_t)\}_{t\geq 1}$ and specified reliability level $\alpha \in [0, B]$, the long-term risk satisfies

$$\left| \frac{1}{T} \sum_{t=1}^{T} L_t - \alpha \right| \leq \frac{S_{\max} + \eta_1 B}{\sqrt{T}}$$

Asymptotic marginal guarantees (informal) [2]

For i.i.d. data sequences $(X_t, Y_t) \sim P_{XY}$ and loss $\mathcal{L}(C, y) = \mathbb{1}\{y \notin C\}$, the prediction sets satisfies asymptotic marginal coverage

$$\lim_{T\to\infty} \Pr\left[Y \notin C_T\right] \overset{p}{=} \alpha$$

[1] Feldman, Shai, et al. "Achieving Risk Control in Online Learning Settings." *Transactions on Machine Learning Research* (2024).
[2] Angelopoulos, Anastasios Nikolas, Rina Barber, and Stephen Bates. "Online conformal prediction with decaying step sizes." *ICML (2024)*.

Limitations of Marginal Risk Control

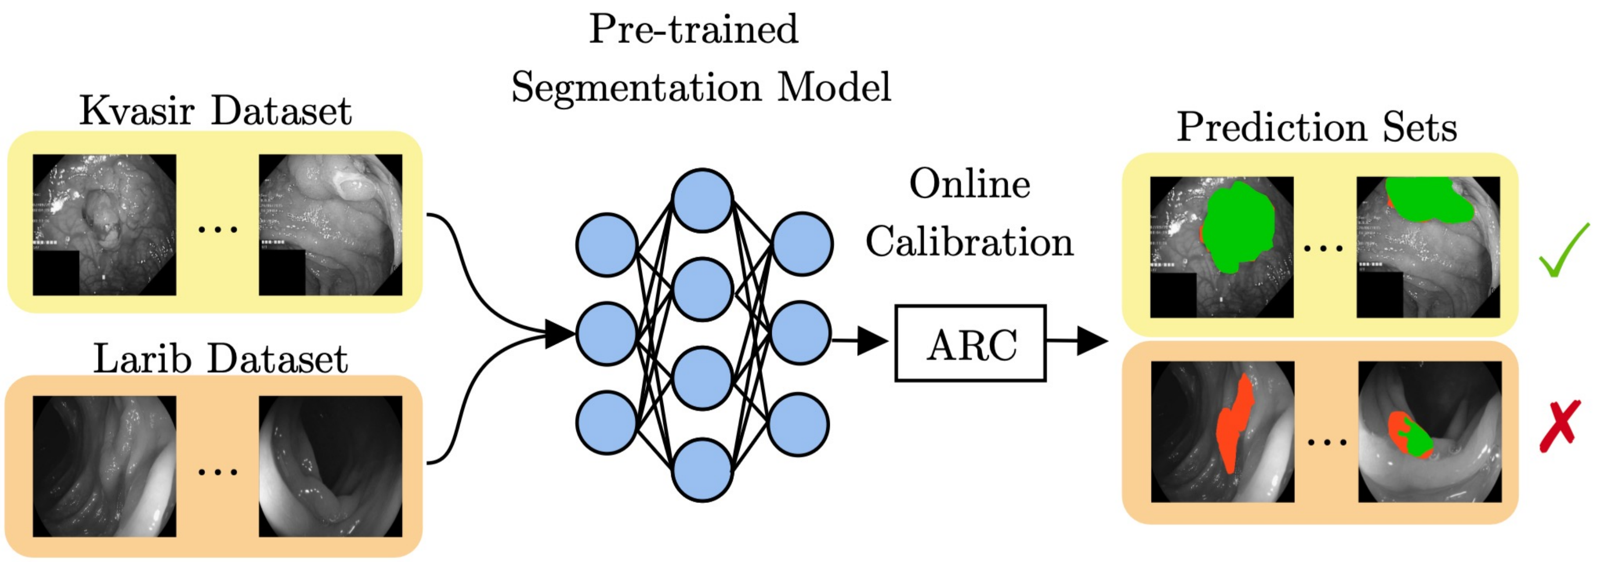

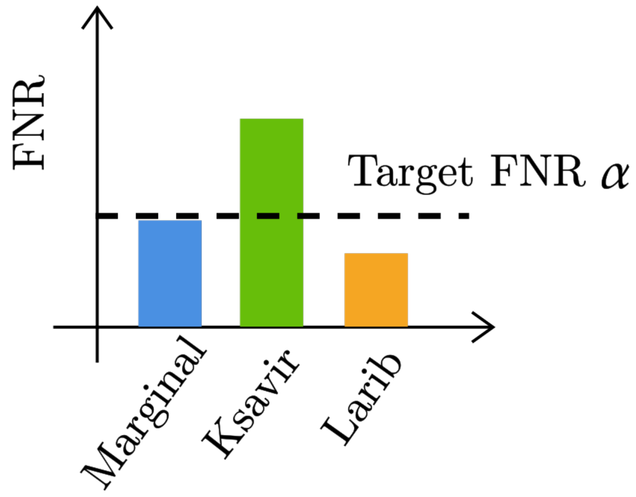

Uneven Risk Control

Contribution

We provide a novel online calibration algorithm that comes with worst-case deterministic guarantees as well as asymptotic localized risk control.

Localized Risk Control

Localized risk control[3] control the risk for a class of covariate shifts

$$\mathbb{E}_{X,Y,\mathcal{D}_{\text{cal}}}\left[\frac{w(X)}{\mathbb{E}_X[w(X)]}\mathcal{L}(C(X|\mathcal{D}_{\text{cal}}),Y)\right] \leq \alpha, \text{ for all } w(\cdot) \in \mathcal{W}.$$

The degree of localization of the functions in set $\mathcal{W}$ dictates the strength of the guarantee.

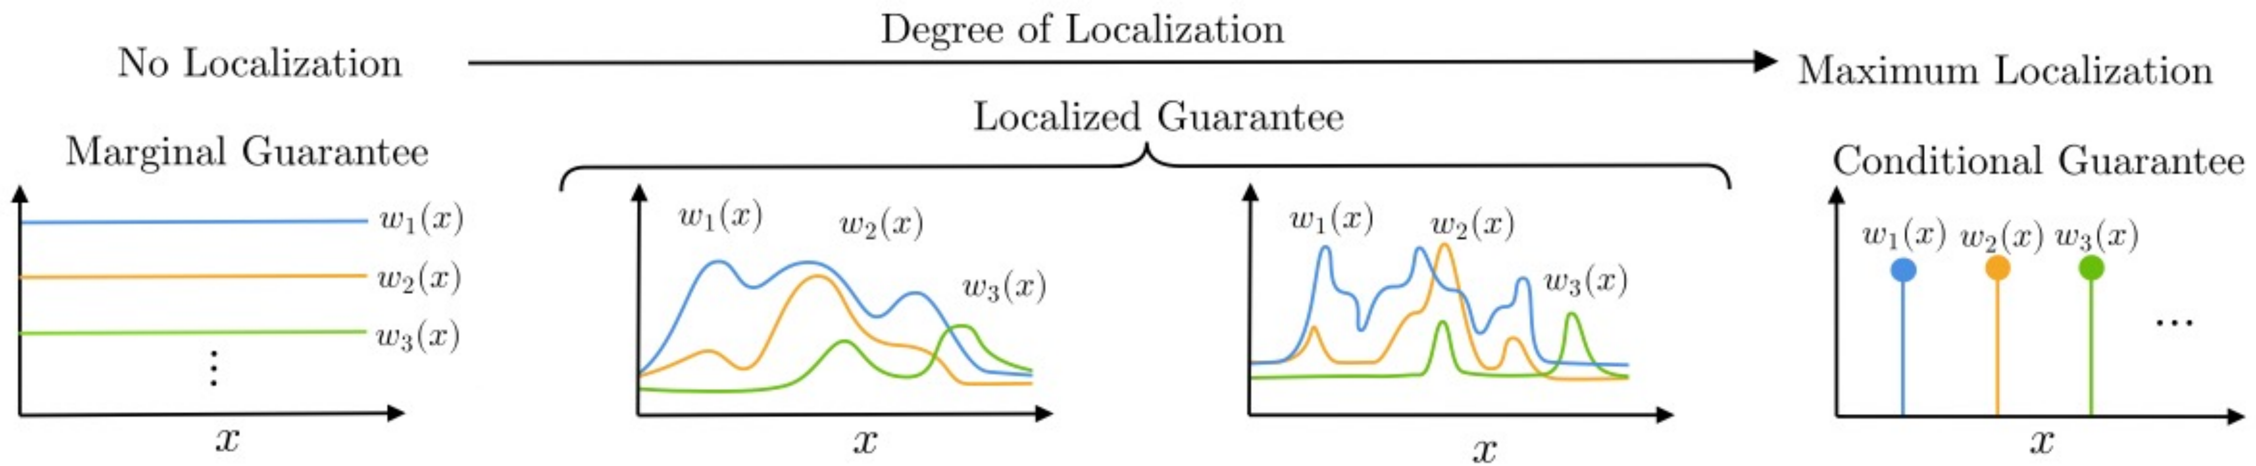

[3] Gibbs, Isaac, John J. Cherian, and Emmanuel J. Candès. "Conformal prediction with conditional guarantees." *arXiv:2305.12616* (2023).

Localized Adaptive Risk Control

Localized Adaptive Risk Control (L-ARC) generates prediction sets based on a threshold function $g_t(\cdot) \in \mathcal{G}$

$$C_t = C(X_t, g_t) := \{y \in \mathcal{Y} : s(X_t, y) \leq g_t(X_t)\}\,.$$

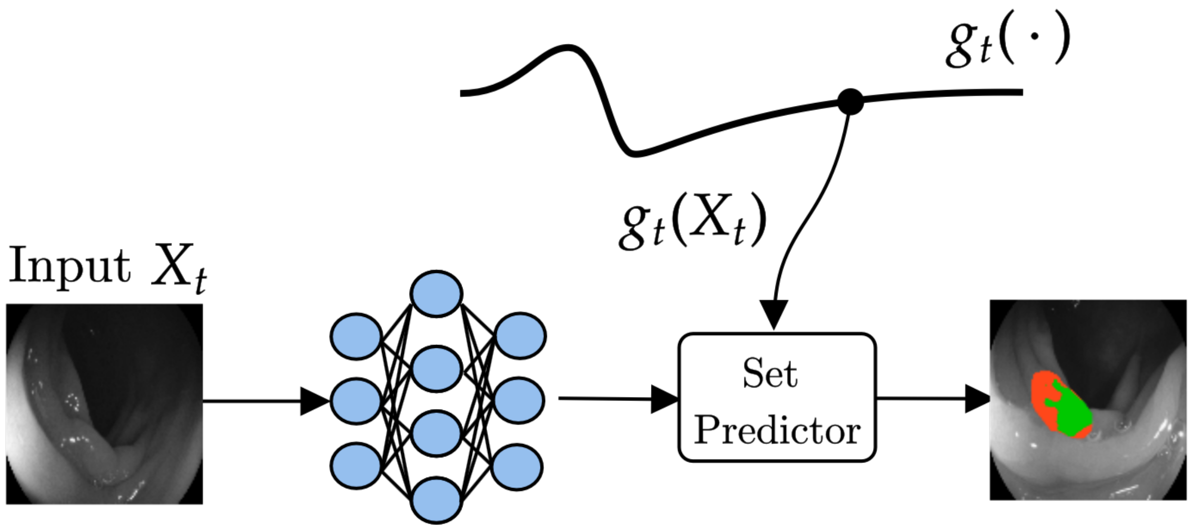

Localized Adaptive Risk Control

Localized Adaptive Risk Control (L-ARC) generates prediction sets based on a threshold function $g_t(\cdot) \in \mathcal{G}$

$$C_t = C(X_t, g_t) := \{y \in \mathcal{Y} : s(X_t, y) \leq g_t(X_t)\}.$$

The threshold $g_t(\cdot) = f_t(\cdot) + c_t$ consists of a function $f_t(\cdot)$ within an RKHS $\mathcal{H}$ and a constant c_t that are updated using an online kernel gradient descent rule.

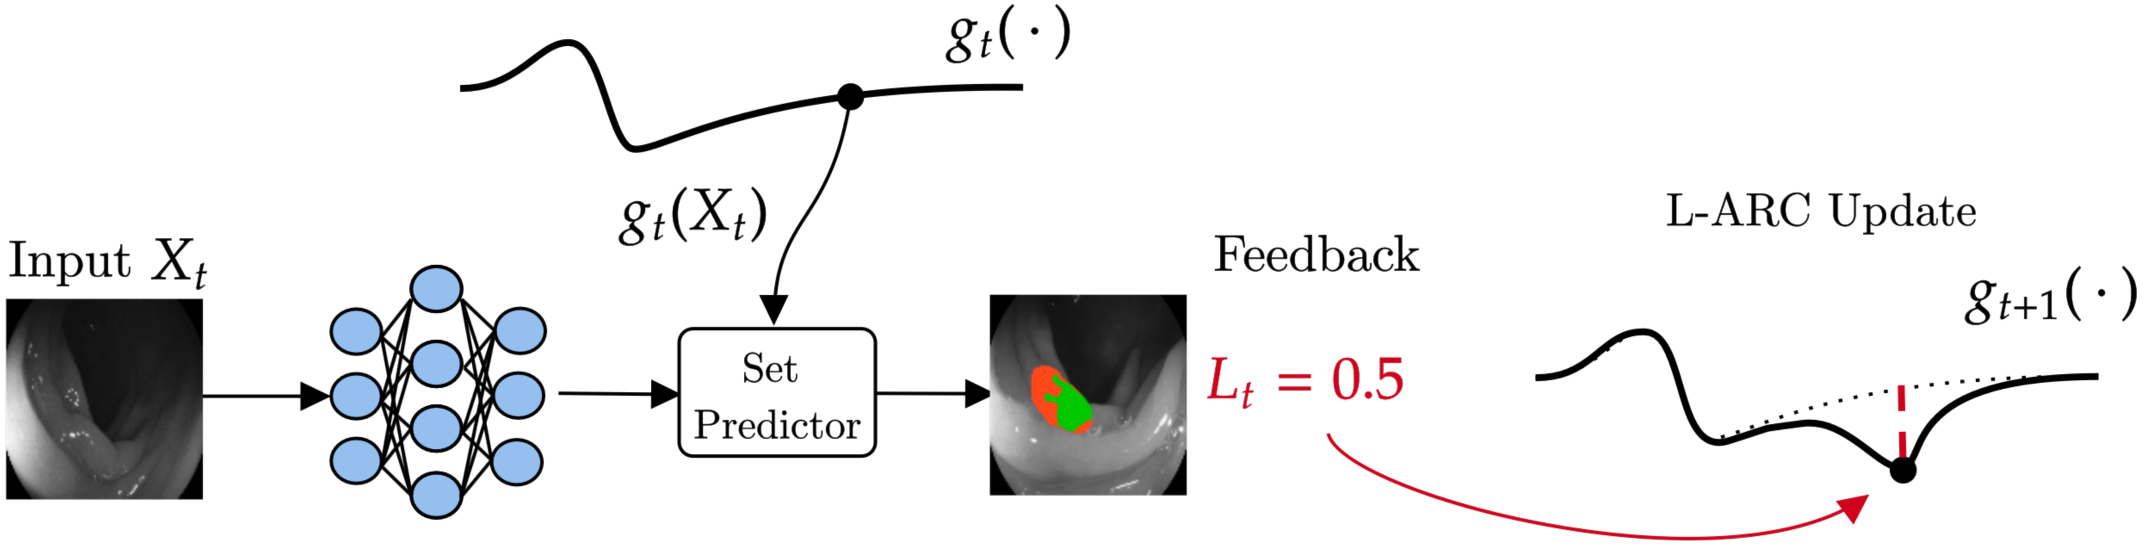

LARC Guarantees

Worst-case deterministic guarantees (informal)

For any sequence $\{(X_t, Y_t)\}_{t\geq 1}$ and reliability level $\alpha \in [0, B]$, the long-term risk satisfies

$$\left| \frac{1}{T} \sum_{t=1}^{T} \mathcal{L}(C(X_t, g_t), Y_t) - \alpha \right| \leq \frac{B(\mathcal{G})}{\sqrt{T}} + C(\mathcal{G}).$$

LARC Guarantees

Worst-case deterministic guarantees (informal)

For any sequence $\{(X_t, Y_t)\}_{t\geq 1}$ and reliability level $\alpha \in [0, B]$, the long-term risk satisfies

$$\left| \frac{1}{T} \sum_{t=1}^{T} \mathcal{L}(C(X_t, g_t), Y_t) - \alpha \right| \leq \frac{B(\mathcal{G})}{\sqrt{T}} + C(\mathcal{G}).$$

Asymptotic localized guarantees (informal)

For i.i.d. data sequences $(X_t, Y_t) \sim P_{XY}$ L-ARC prediction sets provide asymptotic *localized* risk control

$$\limsup_{T\to\infty} \mathbb{E}_{X,Y} \left[\frac{w(X)}{\mathbb{E}_X [w(X)]} \mathcal{L}(C(X, \bar{g}_T), Y) \right] \overset{p}{\leq} \alpha + A(\mathcal{G}, w).$$

LARC Guarantees

Worst-case deterministic guarantees (informal)

For any sequence $\{(X_t, Y_t)\}_{t\geq 1}$ and reliability level $\alpha \in [0, B]$, the long-term risk satisfies

$$\left| \frac{1}{T} \sum_{t=1}^{T} \mathcal{L}(C(X_t, g_t), Y_t) - \alpha \right| \leq \frac{B(\mathcal{G})}{\sqrt{T}} + C(\mathcal{G}).$$

Asymptotic localized guarantees (informal)

For i.i.d. data sequences $(X_t, Y_t) \sim P_{XY}$ L-ARC prediction sets provide asymptotic *localized* risk control

$$\limsup_{T\to\infty} \mathbb{E}_{X,Y} \left[\frac{w(X)}{\mathbb{E}_X[w(X)]} \mathcal{L}(C(X, \bar{g}_T), Y) \right] \stackrel{p}{\leq} \alpha + A(\mathcal{G}, w).$$

Proportional to the degree of localization of functions in $\mathcal{G}$ and zero for constant function, thereby recovering Adaptive Risk Control as a special case.

Example: Tumor Segmentation

Tumor segmentation with false negative ratio (FNR) guarantees.

Calibration data is from 5 different datasets.

In L-ARC, the length scale l of the kernel controls the localization of the threshold function.

Example: Tumor Segmentation

Tumor segmentation with false negative ratio (FNR) guarantees.

Calibration data is from 5 different datasets.

In L-ARC, the length scale l of the kernel controls the localization of the threshold function.

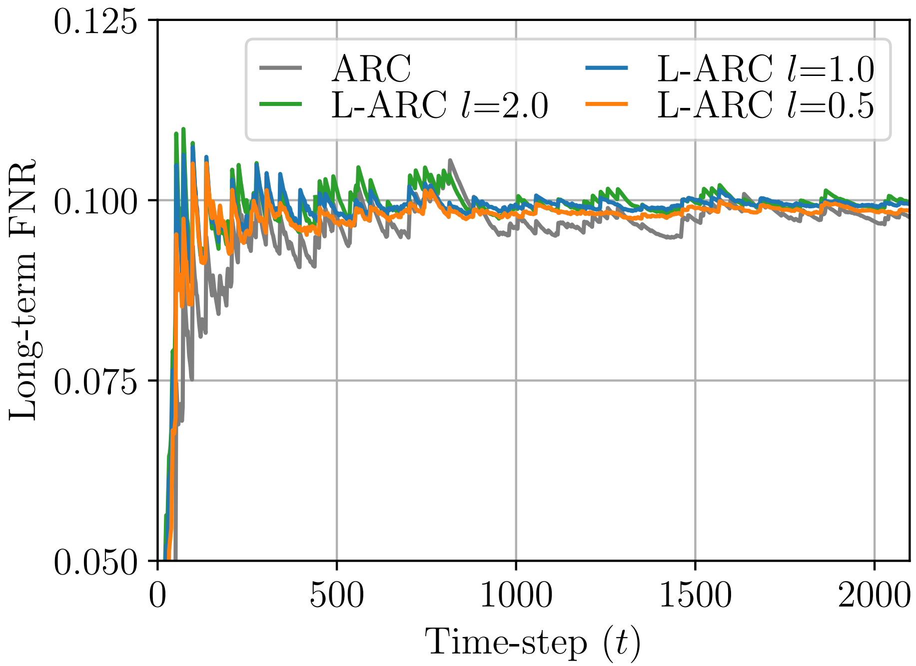

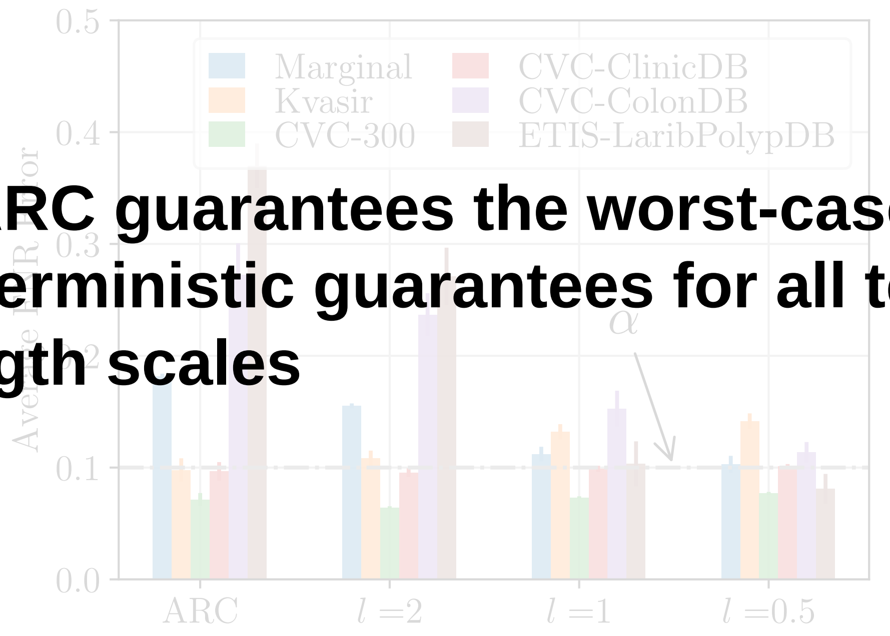

L-ARC guarantees the worst-case deterministic guarantees for all tested length scales

Example: Tumor Segmentation

Tumor segmentation with false negative ratio (FNR) guarantees.

Calibration data is from 5 different datasets.

In L-ARC, the length scale l of the kernel controls the localization of the threshold function.

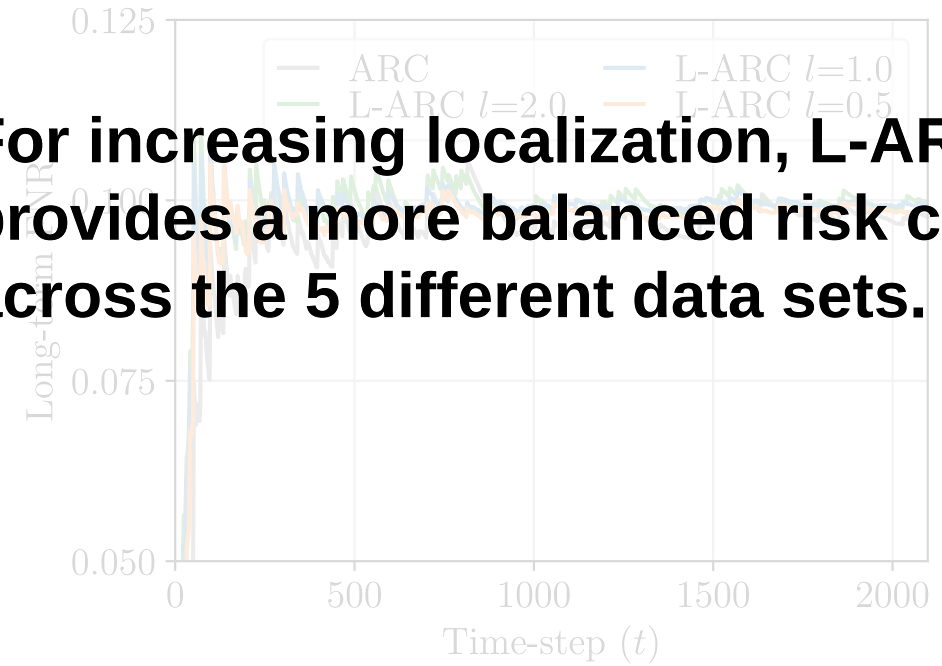

For increasing localization, L-ARC provides a more balanced risk control across the 5 different data sets.

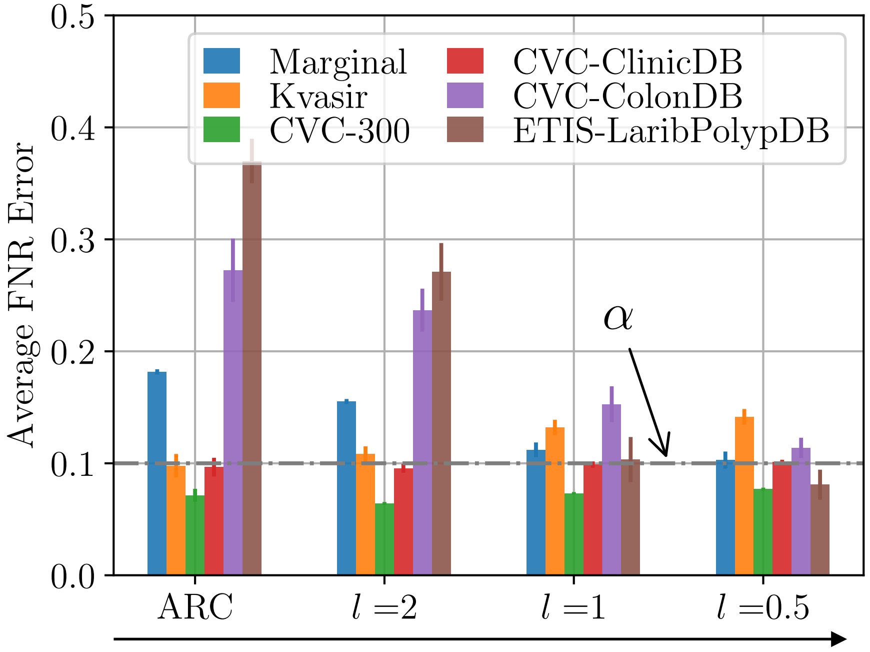

Localization

Conclusion

We proposed Localized Adaptive Risk Control (L-ARC), an online calibration scheme offering both **worst-case deterministic** and **asymptotic localized risk control**.

In a variety of experiments, L-ARC is shown to provide more balanced and fairer risk control as compared to ARC.

Thank you for your attention! For more details:

matteo.1.zecchin@kcl.ac.uk

CENTRIC

This work was supported by the CENTRIC European Project - H2020 Grant Agreement Number: 101096379