Inference Trees: Adaptive Inference with Exploration

Tom Rainforth
Joint work with Yuan Zhou, Xiaoyu Lu, Yee Whye Teh,
Frank Wood, Hongseok Yang, and Jan Willem van de Meent
Preprint: arXiv:1806:09550

UNIVERSITY OF OXFORD

UNIVERSITY OF OXFORD

Overview

- Adaptive inference methods need to explore, not just exploit
- New class of inference algorithms — Inference Trees — based around Monte Carlo tree search
- Targeted exploration using density estimation of log sample weights

$$\pi(x) = \frac{\gamma(x)}{\omega}$$

Background: Importance Sampling

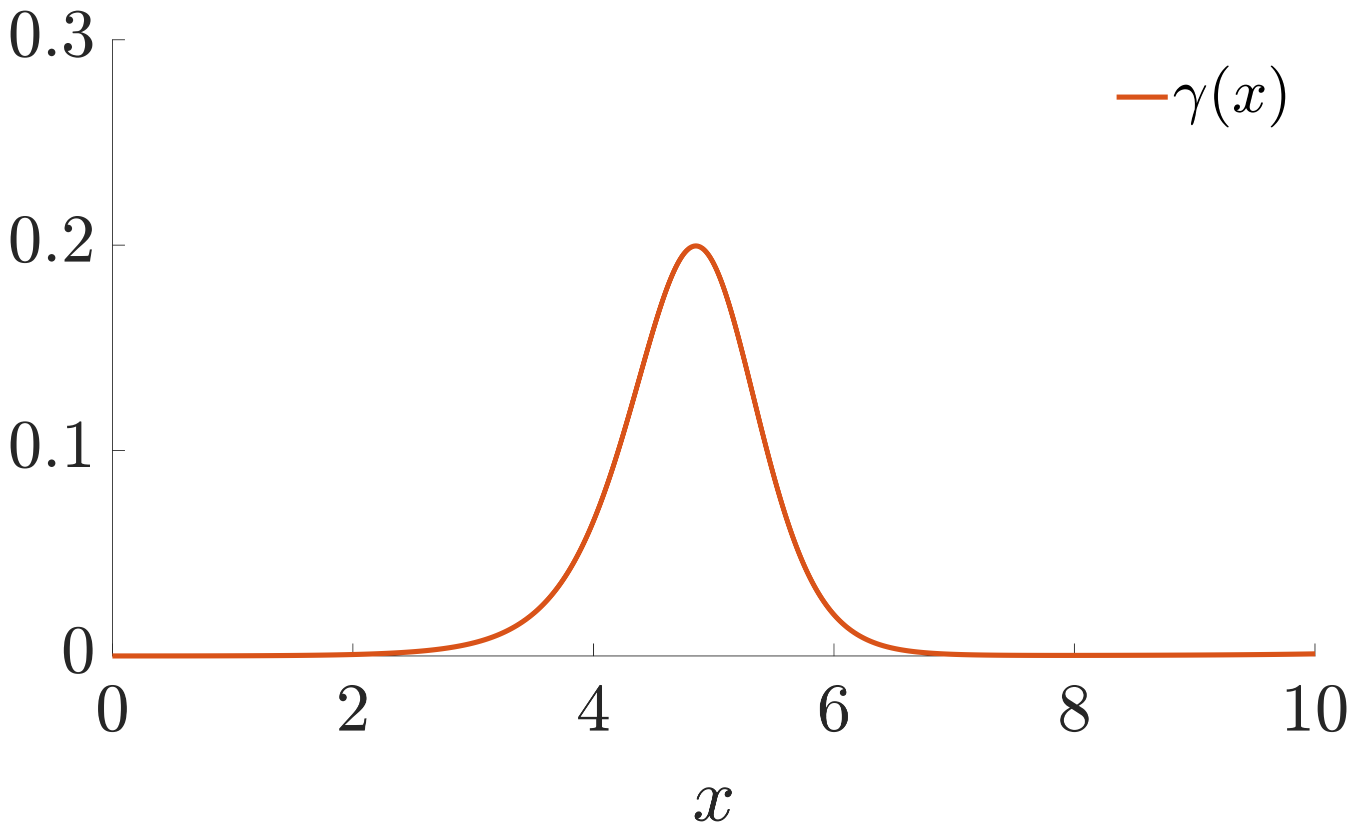

Background: Importance Sampling

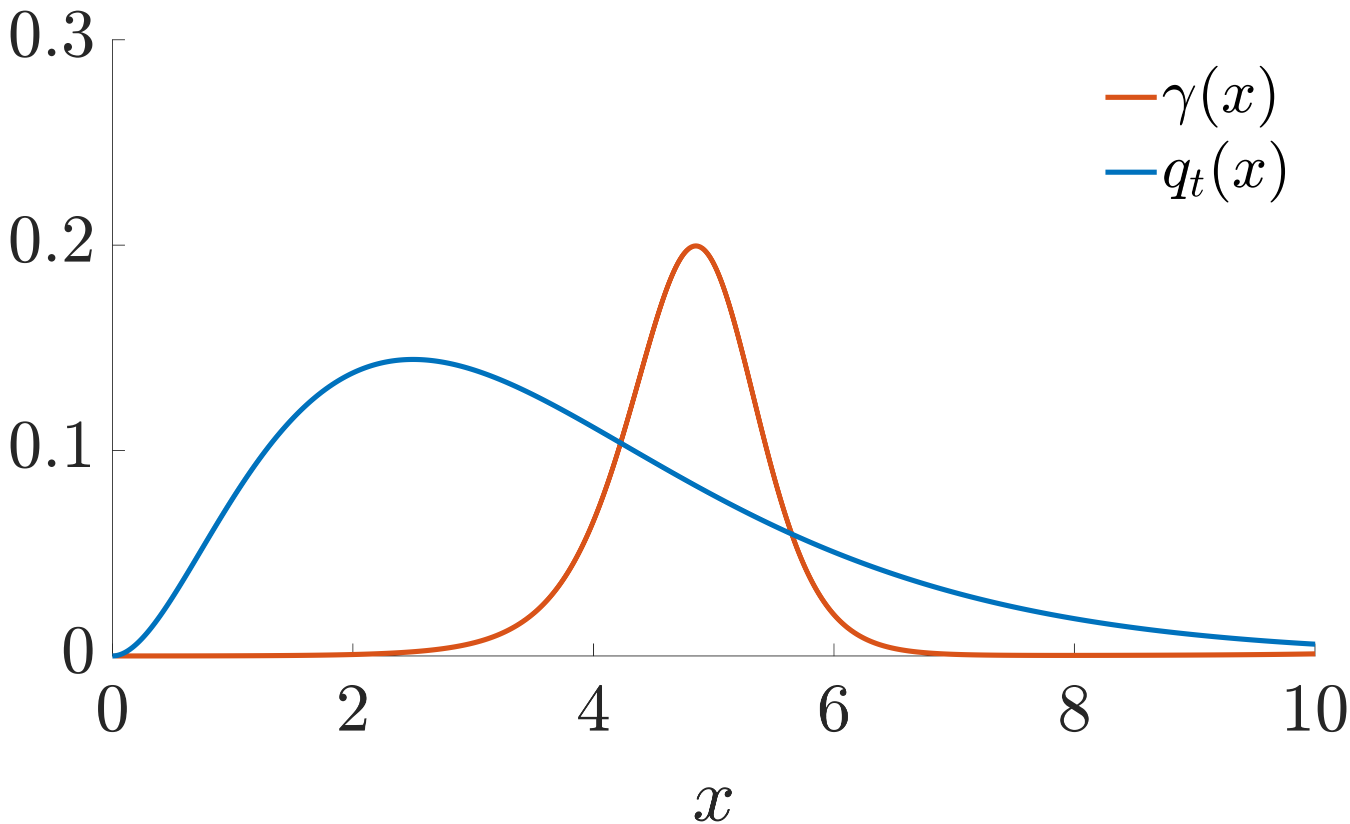

Background: Importance Sampling

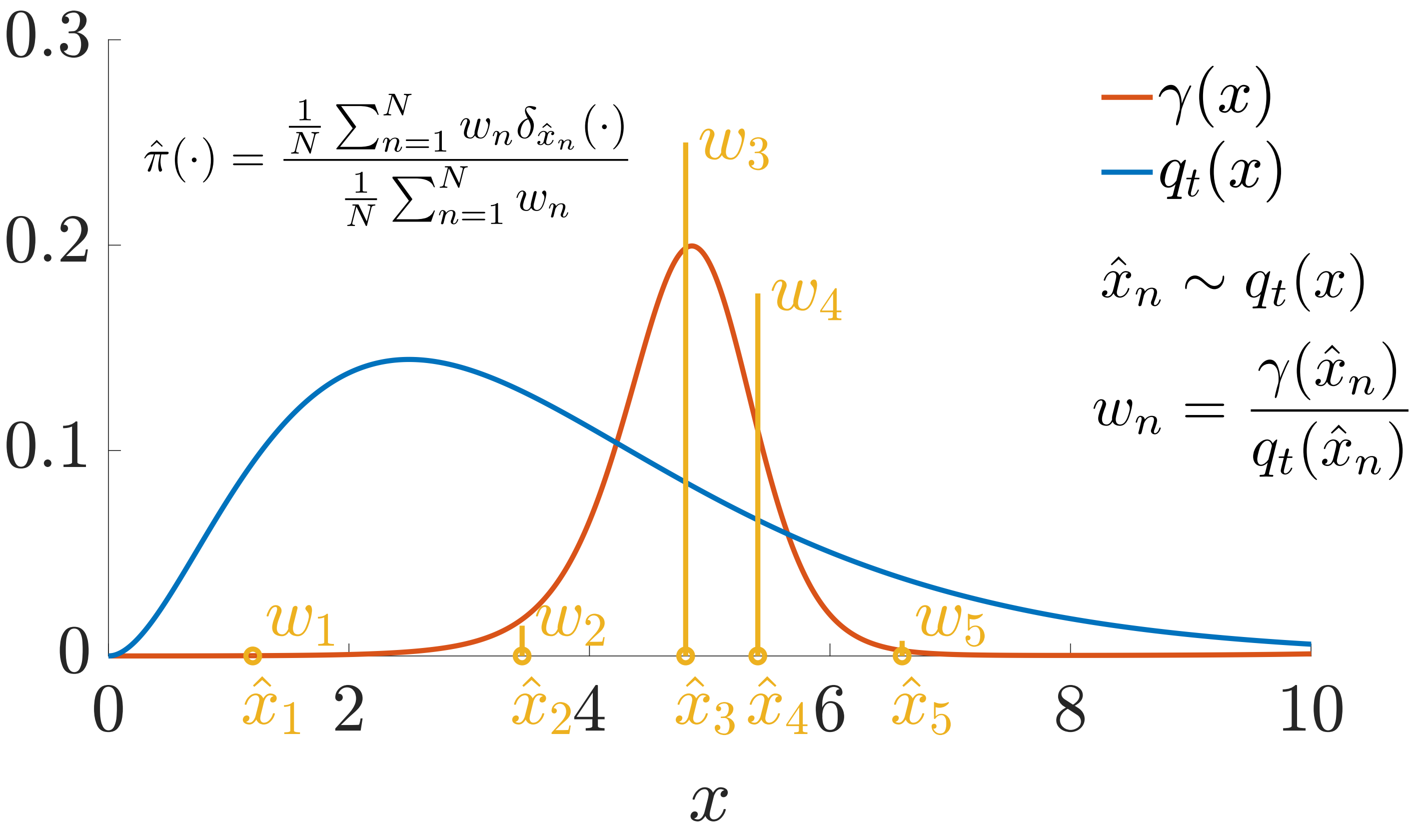

Background: Adaptive Monte Carlo

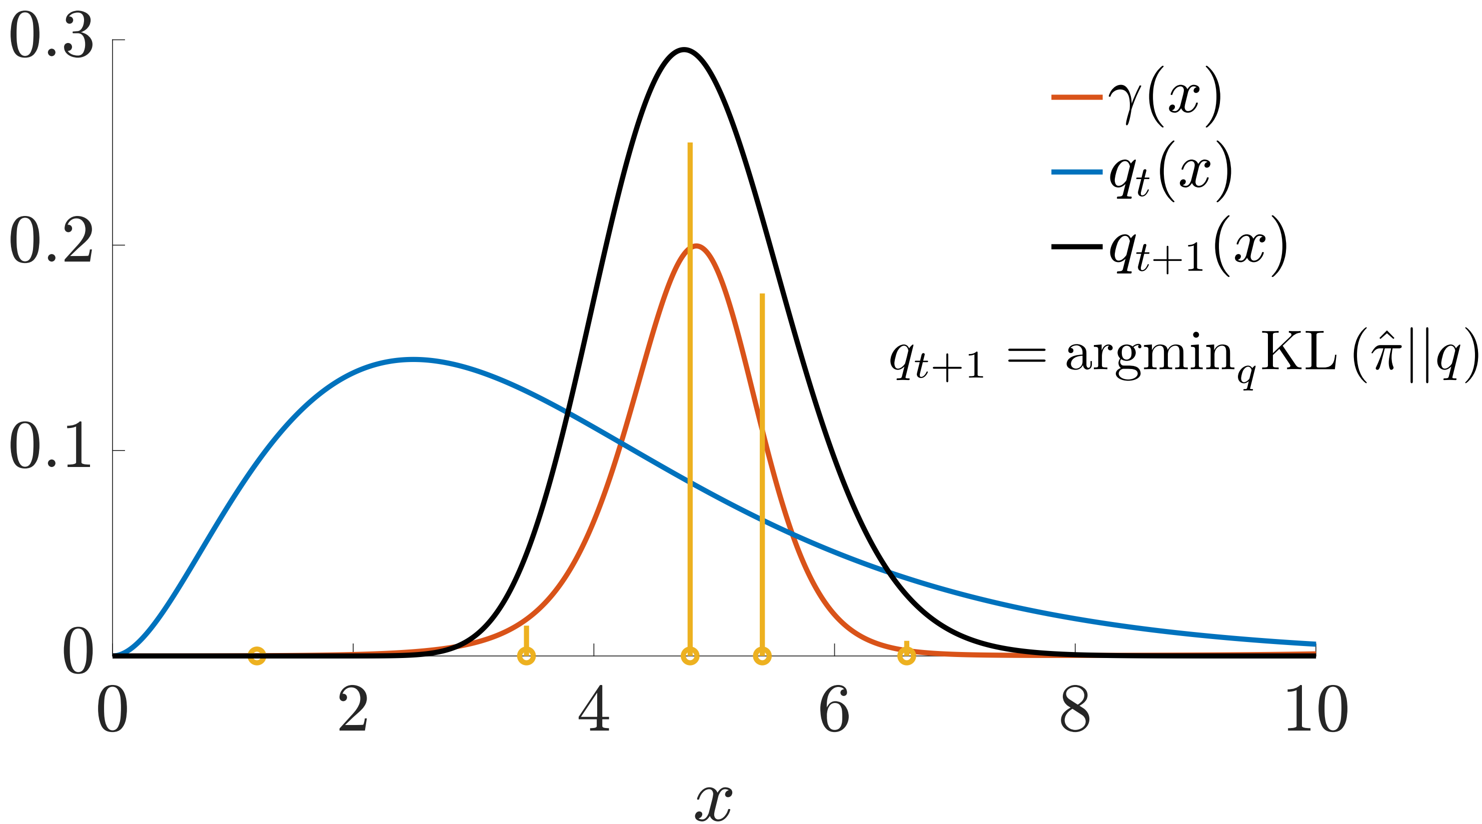

Background: Adaptive Monte Carlo

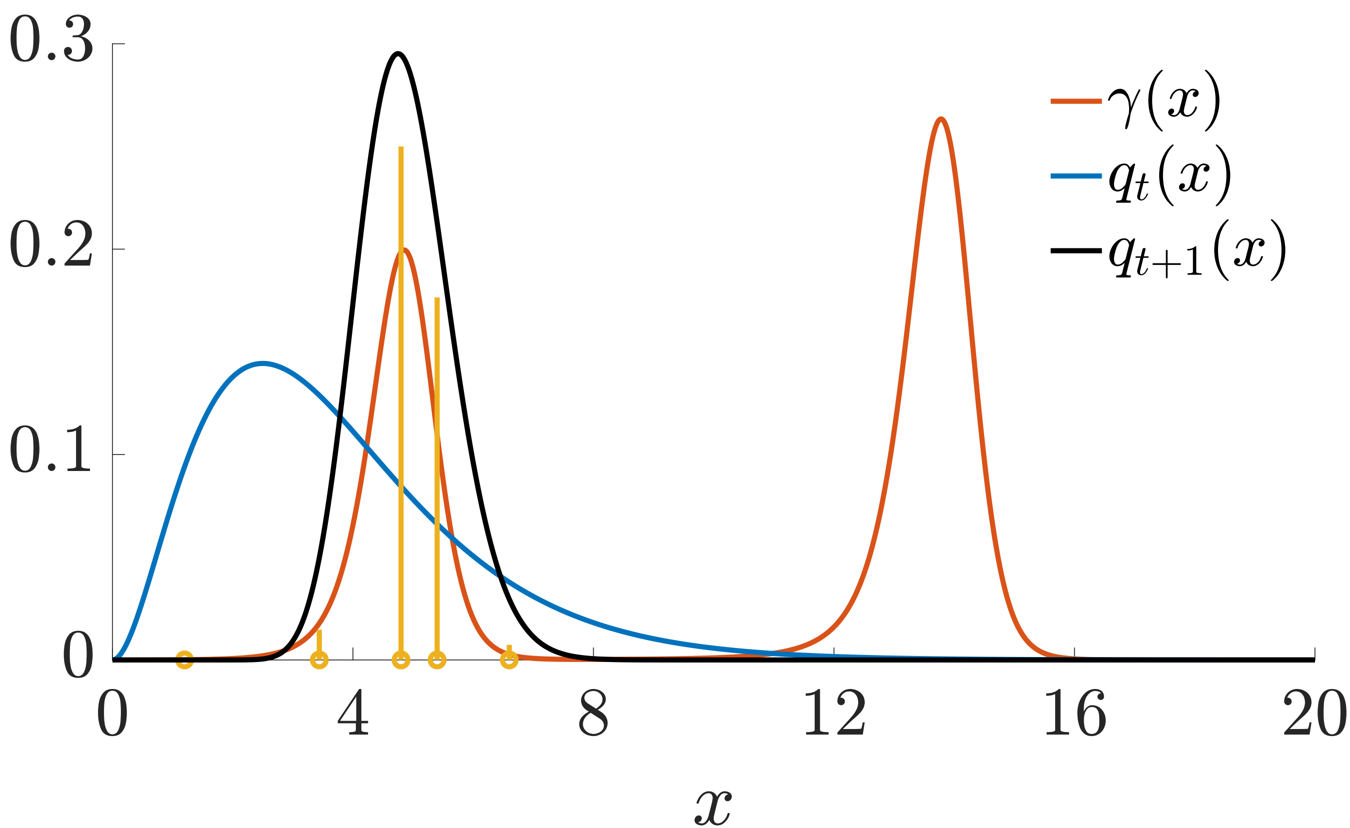

Exploitation vs Exploration

- Exploitation: sample in regions where we think the posterior mass is high
- Exploration: sample in regions where our uncertainty is high
- Utility from samples originates not only from direct contribution to estimator, but also information provided for future sampling

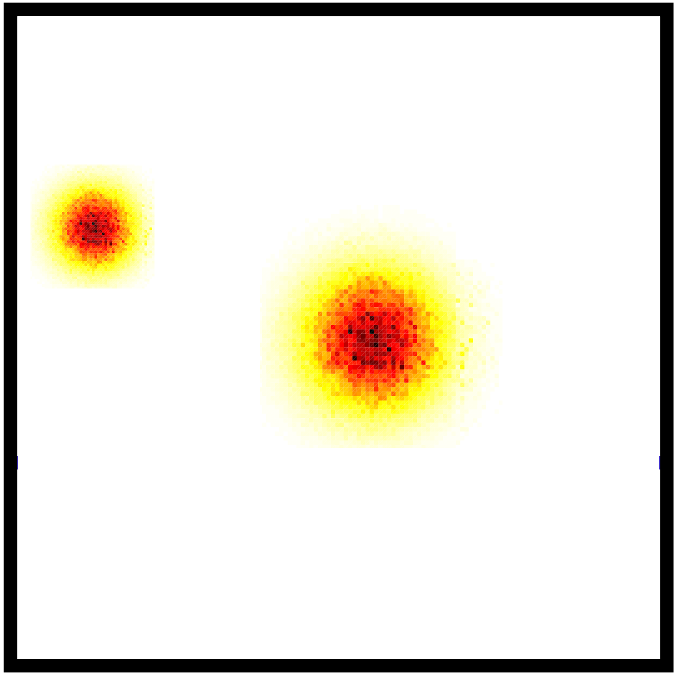

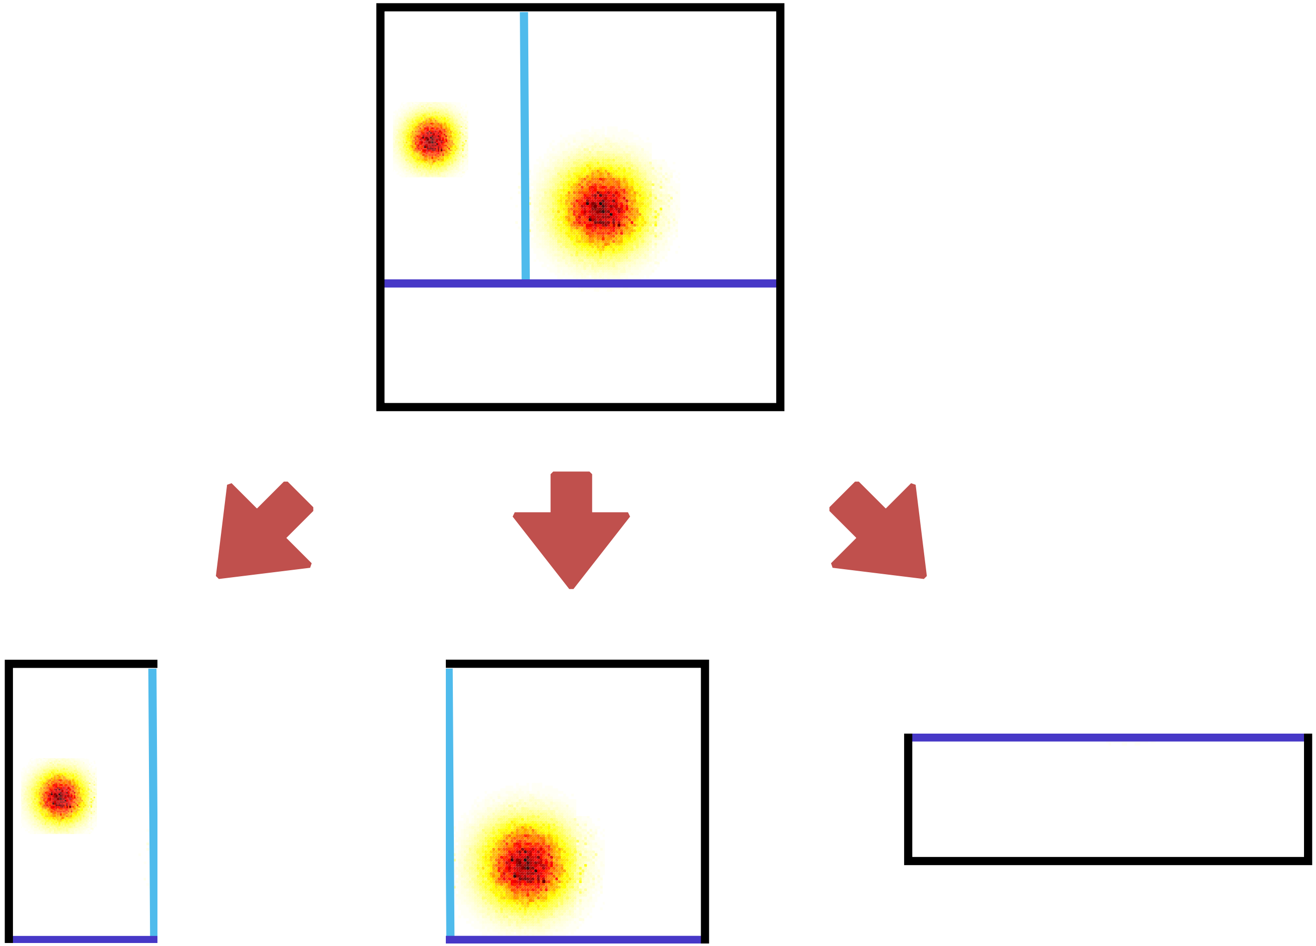

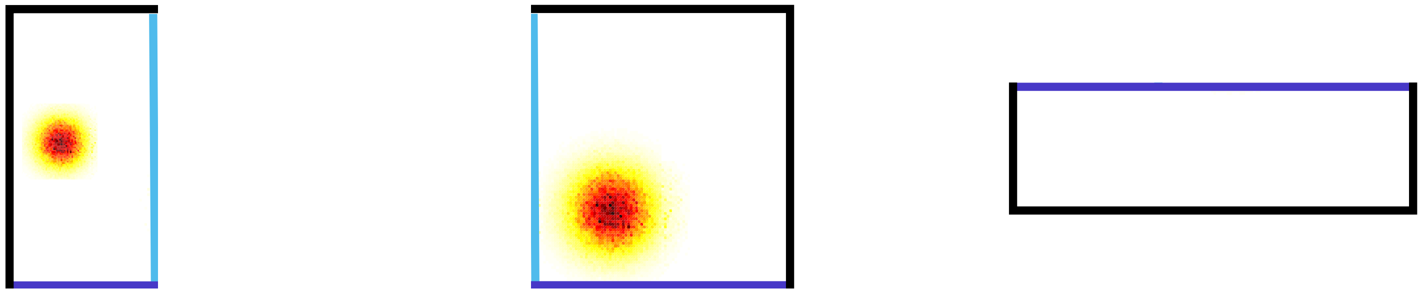

$$\hat{\gamma}_1(\cdot) \qquad \hat{\gamma}_2(\cdot) \qquad \hat{\gamma}_3(\cdot)$$

$$N_1 = 180 \qquad N_2 = 800 \qquad N_3 = 20$$

$$\hat{\gamma}(\cdot) = \hat{\gamma}_1(\cdot) + \hat{\gamma}_2(\cdot) + \hat{\gamma}_3(\cdot)$$

Why Split and Control?

- We no longer need to draw iid samples
  - More explicit control for resource allocation
  - Can gather additional information
- Easy to maintain consistency under adaptation

Inference Trees

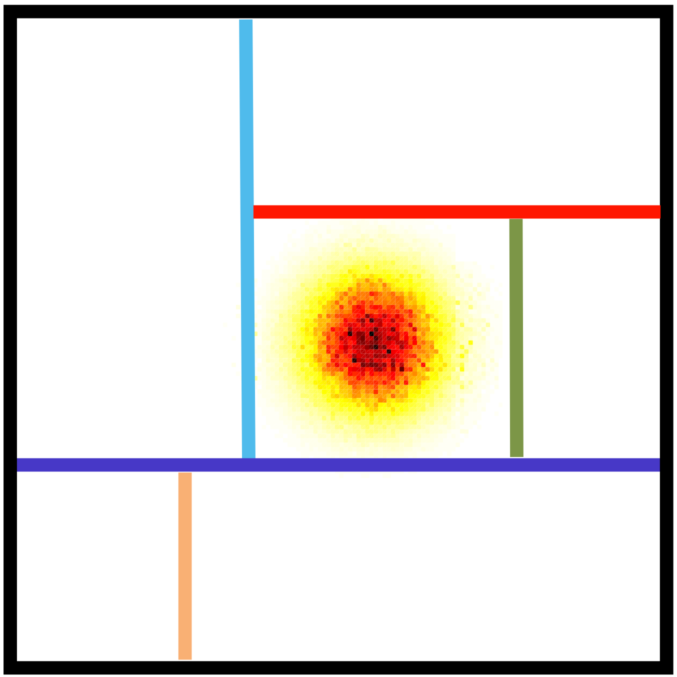

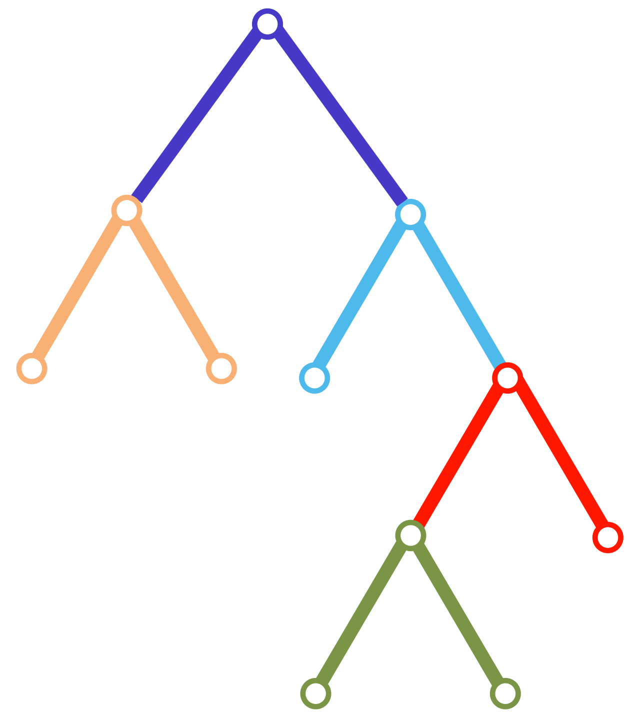

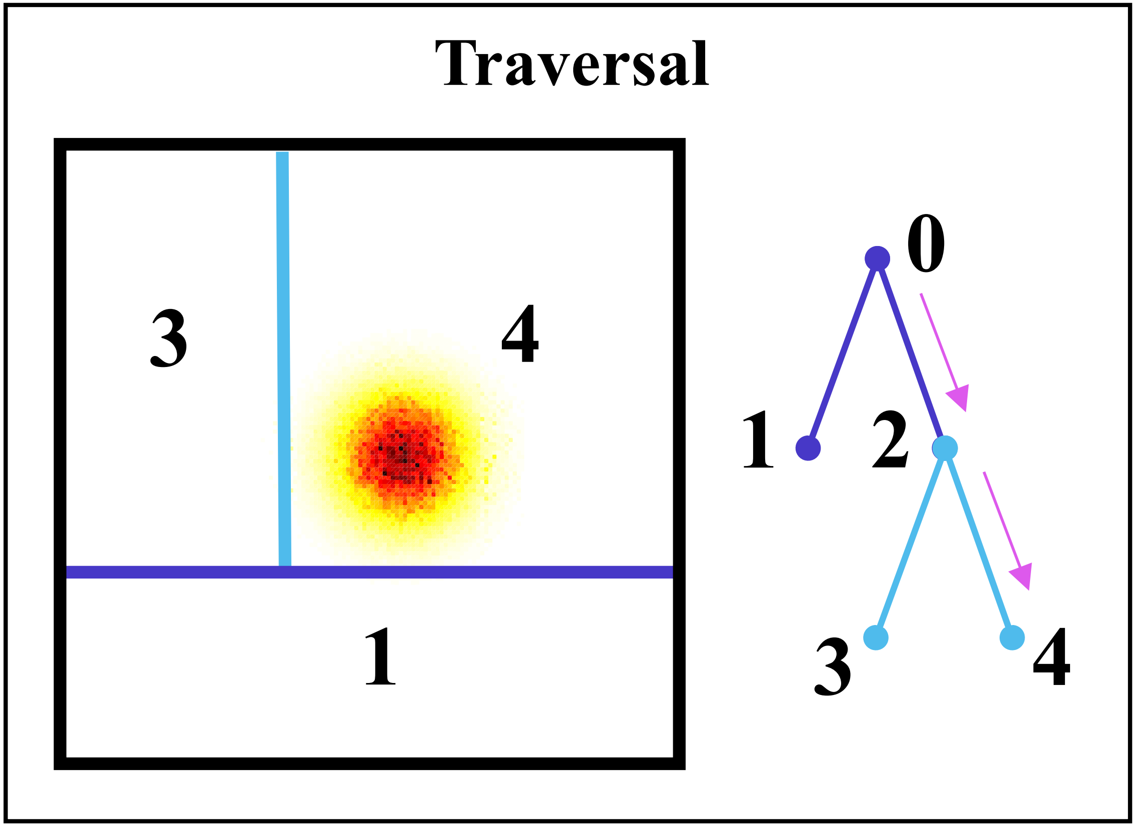
Traversal
3
4
1
0
1
2
3
4

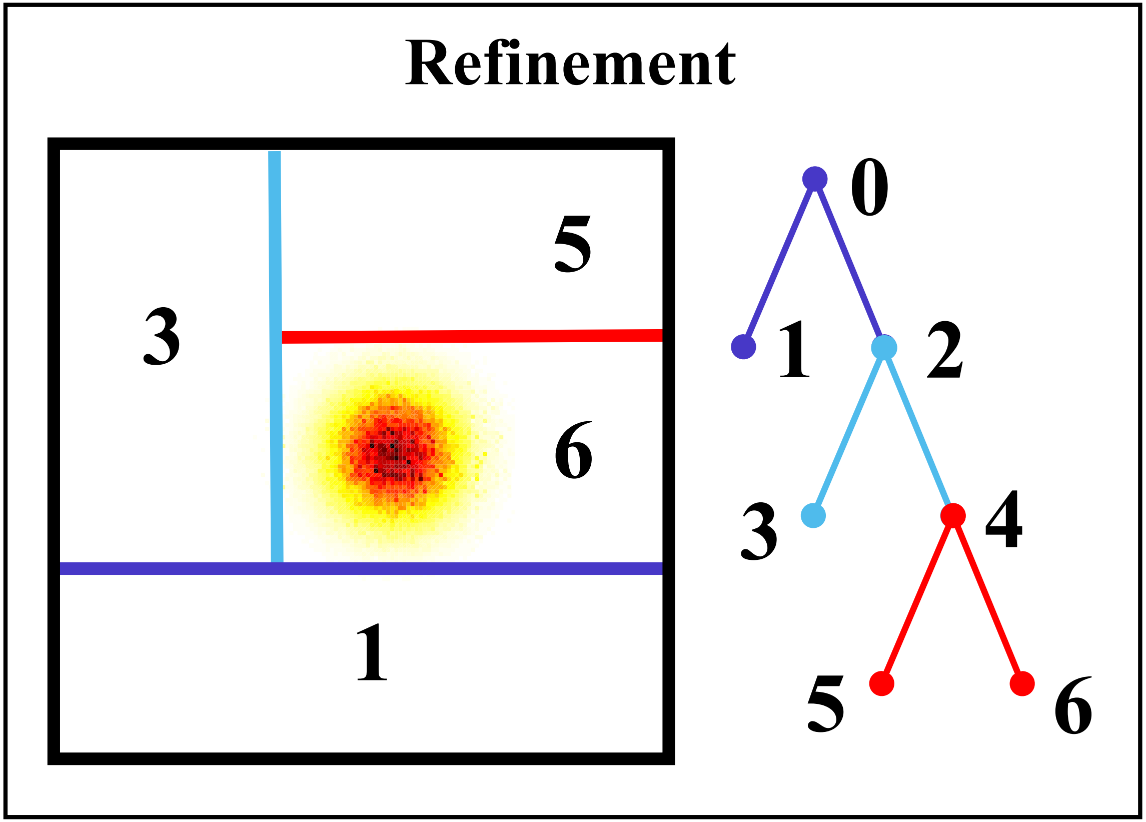
Refinement
5
3
6
1
0
1
2
3
4
5
6

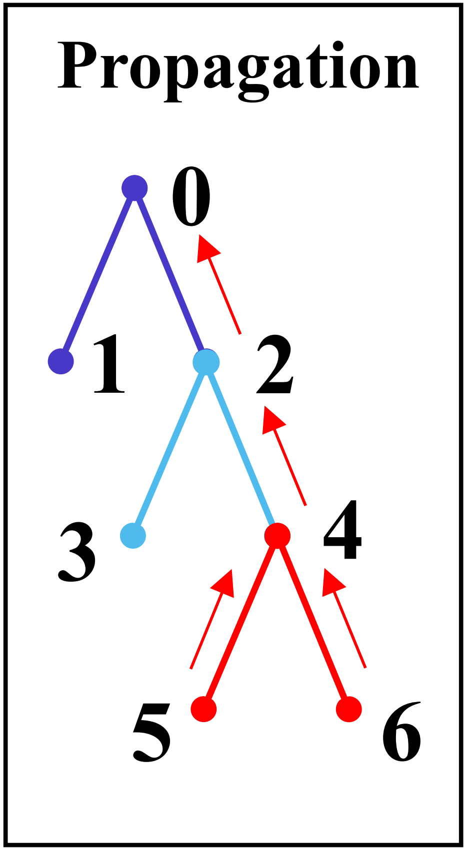
Propagation
0
1
2
3
4
5
6

Inference Tree Estimator / Propagation

$$\hat{\omega}_j = \frac{1}{N_j} \sum_{n=1}^{N_j} w_j^n$$

Inference Tree Estimator / Propagation

$$\hat{\omega}_j = c_j(\hat{\omega}_{\ell_j} + \hat{\omega}_{r_j}) + (1 - c_j)\frac{1}{N_j}\sum_{n=1}^{N_j} w_j^n$$

$$\hat{\gamma}_j(\cdot) = c_j(\hat{\gamma}_{\ell_j}(\cdot) + \hat{\gamma}_{r_j}(\cdot)) + (1 - c_j)\frac{1}{N_j}\sum_{n=1}^{N_j} w_j^n \delta_{\hat{x}_j^n}(\cdot)$$

$$\hat{\pi}_0(\cdot) = \frac{\hat{\gamma}_0}{\hat{\omega}_0}$$

Consistent Regardless of Traversal and Refinement Strategies:

Theorem 1.

$$\hat{\pi}_0(\cdot) \;\; \textit{converges weakly to } \pi(x)$$

UNIVERSITY OF OXFORD

What is the optimal allocation of samples?

One might expect we should sample in proportion to the marginal likelihood, but in fact we have

$$\tau_j = \sqrt{\omega_j^2 + (1+\kappa)\sigma_j^2}$$

Marginal likelihood Squared

Smoothness parameter

$\kappa \geq 0$

Variance of weights

UNIVERSITY OF OXFORD

Traversal

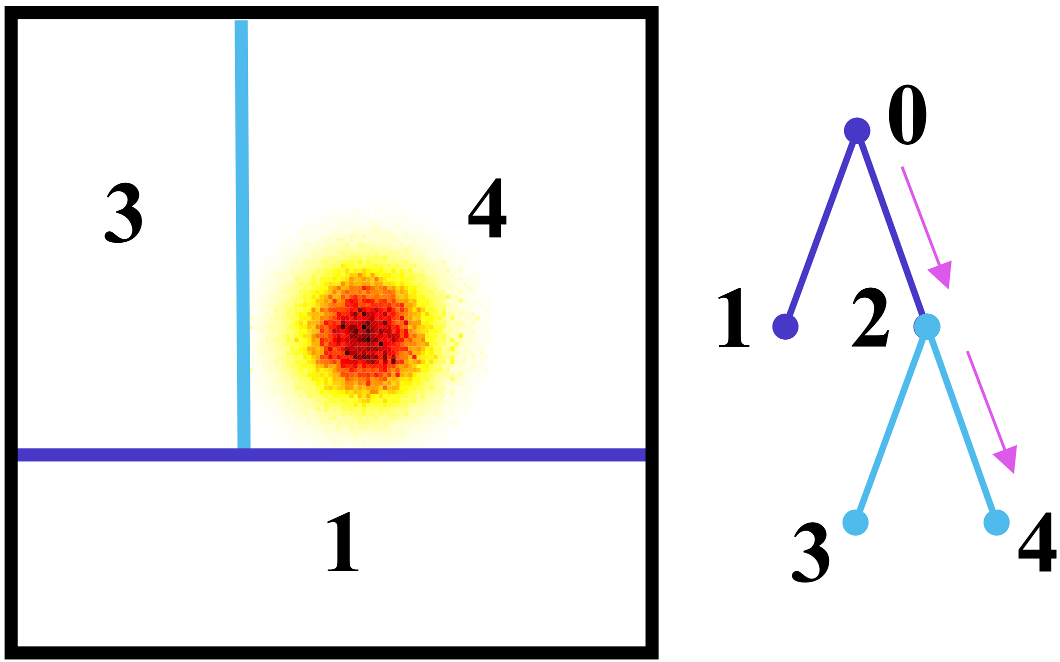

Multi Armed Bandits for Stratified Sampling

UCB: at each round, choose the arm j that maximises

$$u_j = \frac{1}{M_j}\left(\hat{r}_j + \frac{\beta \log \sum_i M_i}{\sqrt{M_j}}\right)$$

Estimated Reward

Optimism Boost

[Carpentier et al 2015]

Reward

$$\hat{r}_j = \hat{\tau}_j$$

UNIVERSITY OF OXFORD

Reward

$$\hat{r}_j = \underbrace{(1-\delta)\,\hat{\tau}_j}_{\text{Exploitation}} + \underbrace{\delta \hat{p}^s_j}_{\text{Targeted Exploration}}$$

UNIVERSITY OF OXFORD

Reward

$$\hat{r}_j = (1-\delta)\,\frac{\hat{\tau}_j}{\hat{\tau}_{\mathrm{pa}(j)}} + \delta\,\frac{\hat{p}^s_j}{\hat{p}^s_{\mathrm{pa}(j)}}$$

Exploitation

Targeted Exploration

How can we perform targeted exploration?

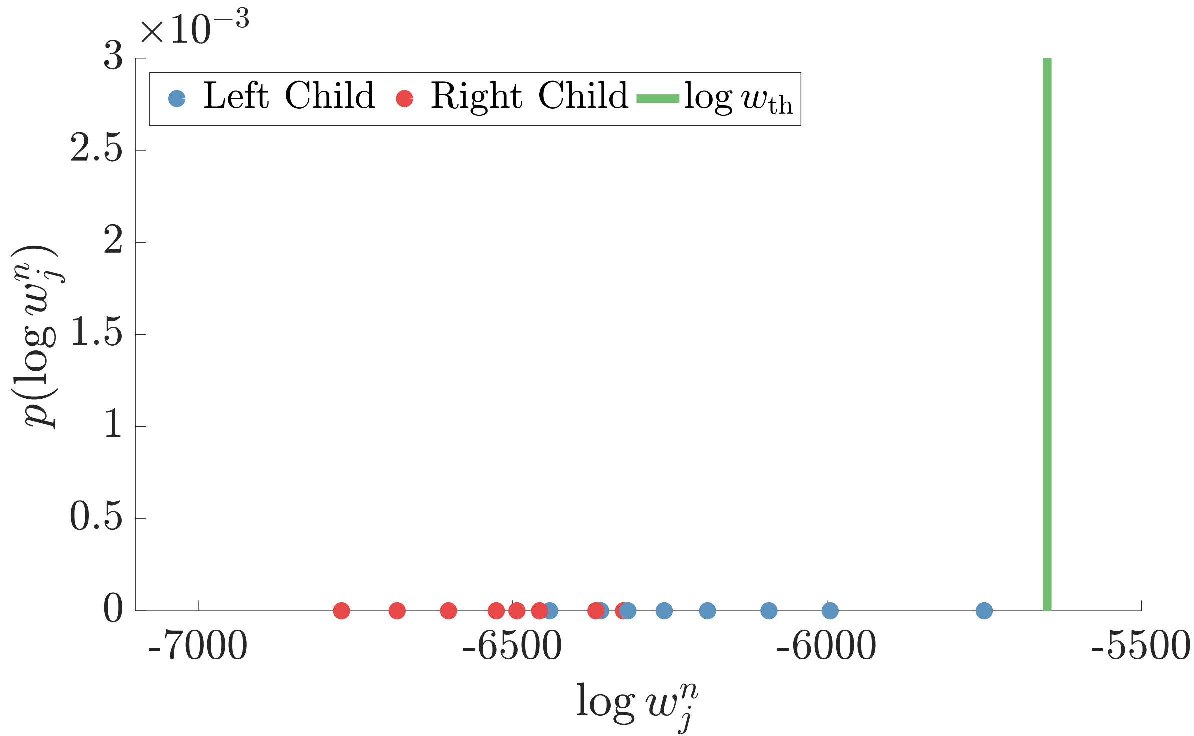

$$\text{Is } \omega_j > w_{\text{th}}?$$

Density estimation of the log weights

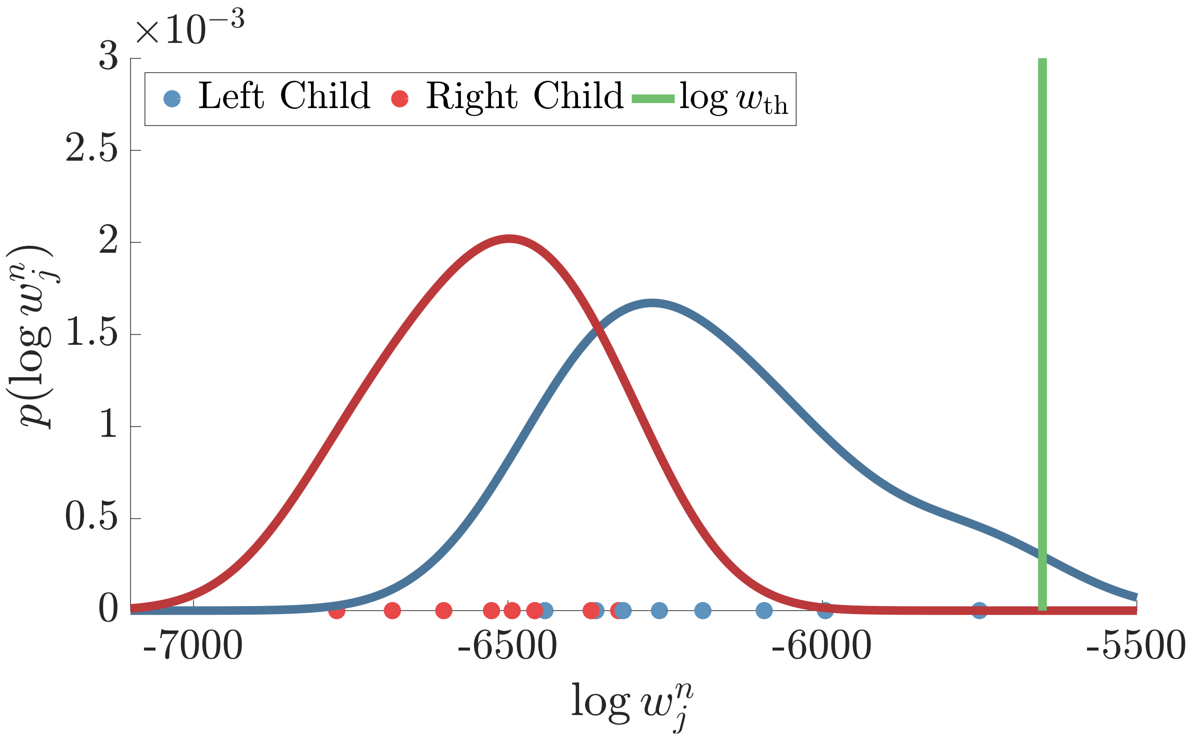

Is $\omega_j > w_{\text{th}}$?

Density estimation of the log weights

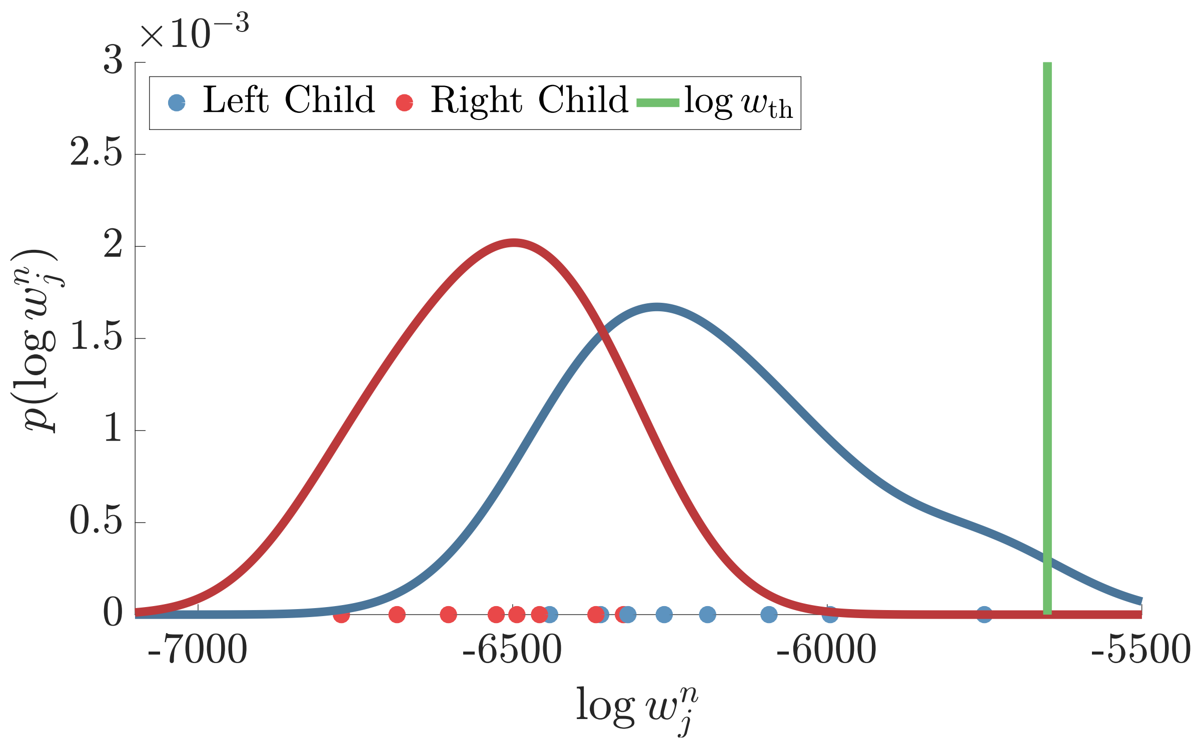

Will $\hat{\omega}_j(T)$ be larger than w_{th}?

Density estimation of the log weights

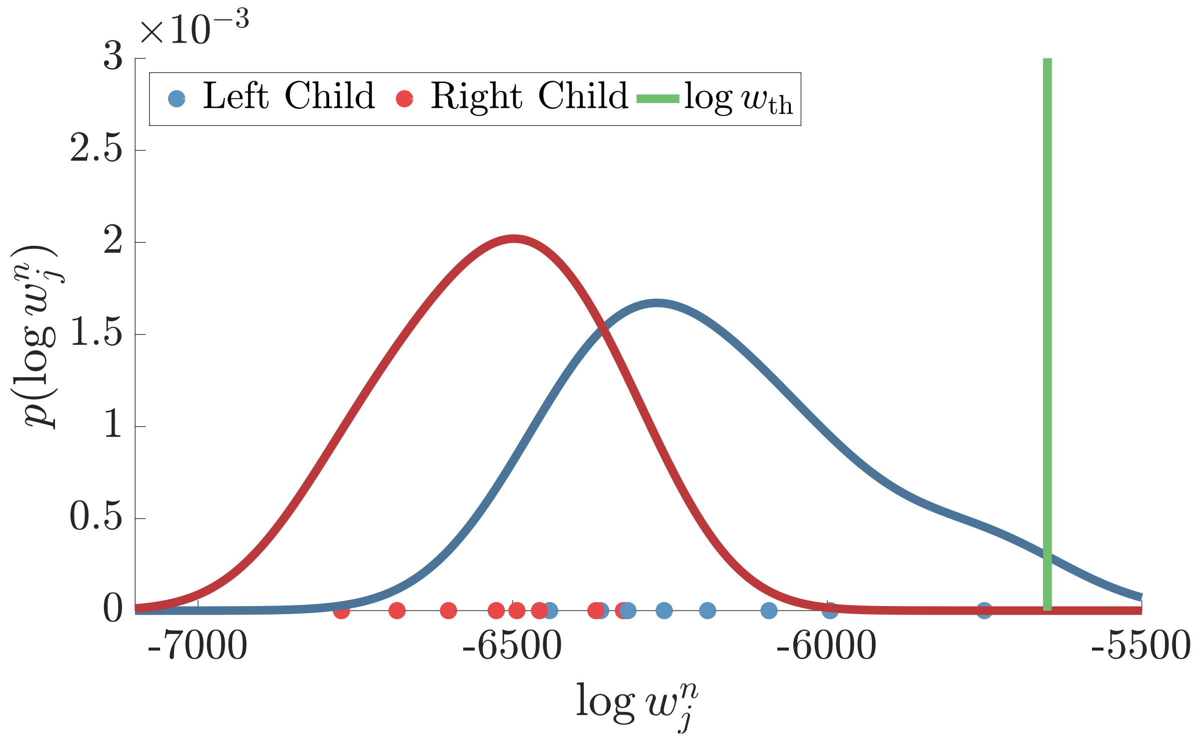

$$P\left(\hat{\omega}_j(T) > w_{\text{th}}\right) \approx P\left(\max(w_j^{1:T}) > w_{\text{th}}\right) \approx 1 - \Psi(\log w_{\text{th}})^T$$

Inference Tree Traversal Target

$$u_j = \frac{1}{M_j}\left((1-\delta)\frac{\hat{\tau}_j}{\hat{\tau}_{\text{pa}(j)}} + \delta\frac{\hat{p}^s_j}{\hat{p}^s_{\text{pa}(j)}} + \frac{\beta \log \sum_i M_i}{\sqrt{M_j}}\right)$$

Exploitation

Targeted Exploration

UNIVERSITY OF OXFORD

Split in Reparameterized Space

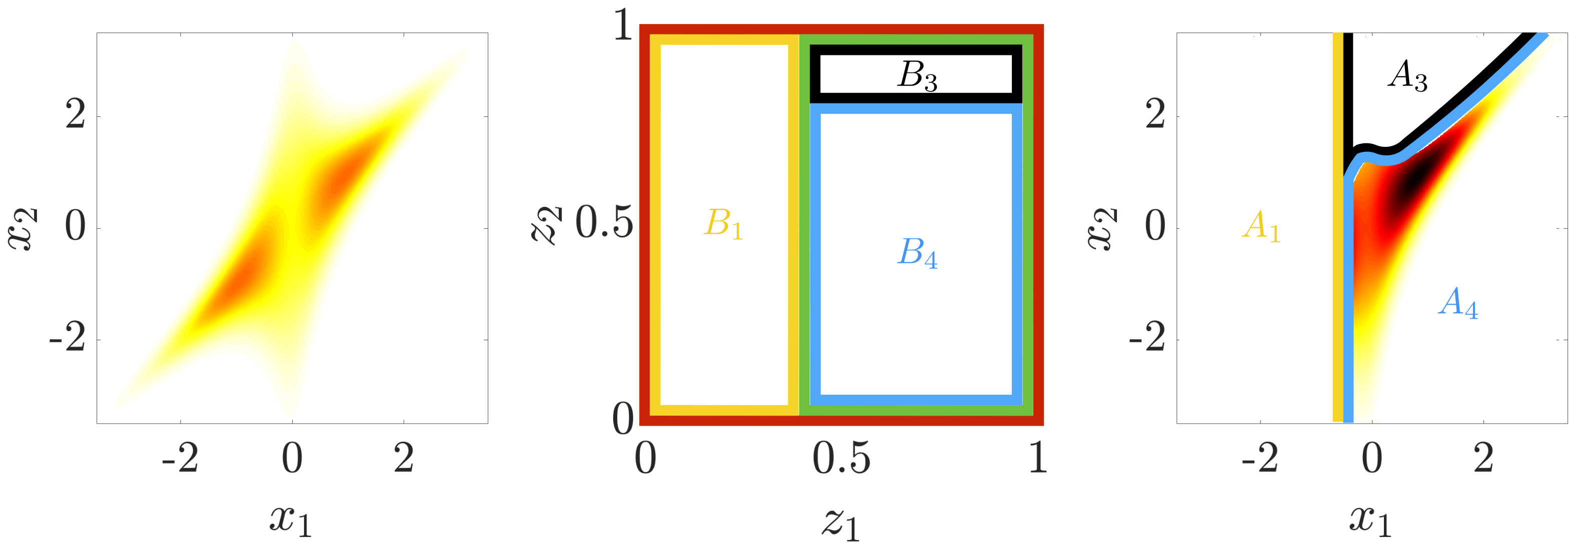

$$q(x|\{x \in A_j\}) = \frac{q(x)\mathbb{I}(x \in A_j)}{\|B_j\|}$$

Refinement

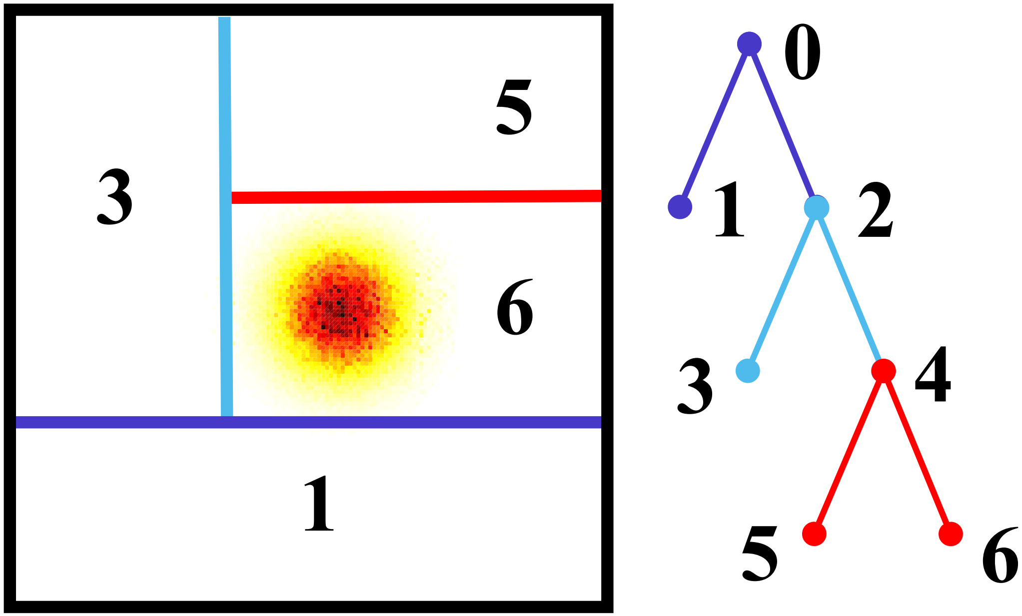

UNIVERSITY OF OXFORD

Split Criterion

- Prefer splits which lead to concentration of probability mass to few nodes

$$\textsc{Loss}(\text{split}) = \hat{\omega}_\ell \log \frac{\|B_\ell\|}{\hat{\omega}_\ell} + \hat{\omega}_r \log \frac{\|B_r\|}{\hat{\omega}_r}$$

- Sample a large number of candidate splits at random
- Choose candidate split which minimises loss criterion for existing samples

UNIVERSITY OF OXFORD

Inference Trees Recap

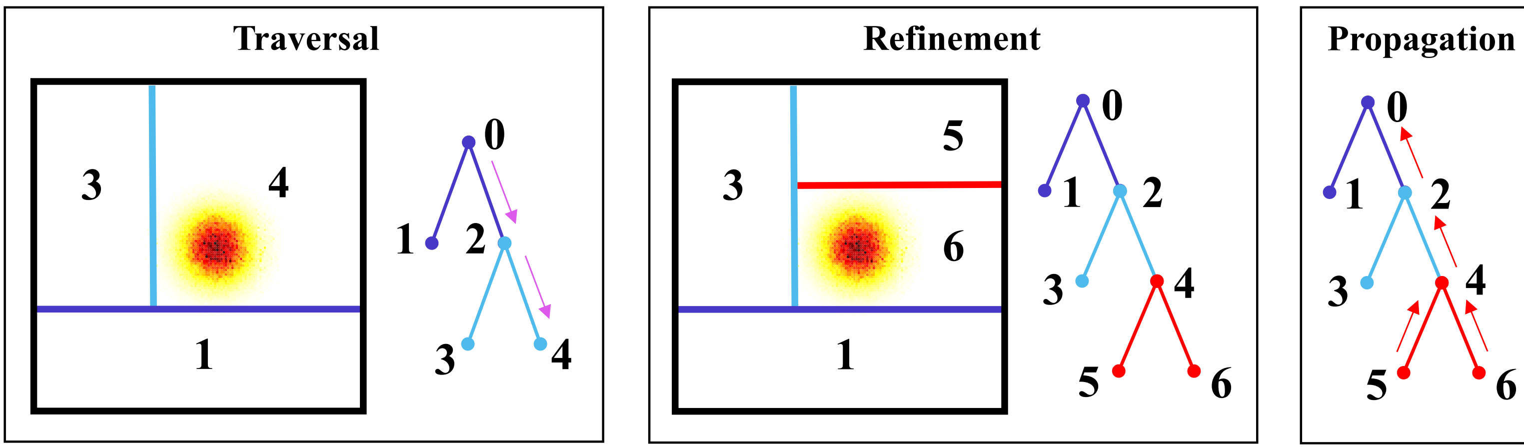

Experiments: Gaussian Mixture Model

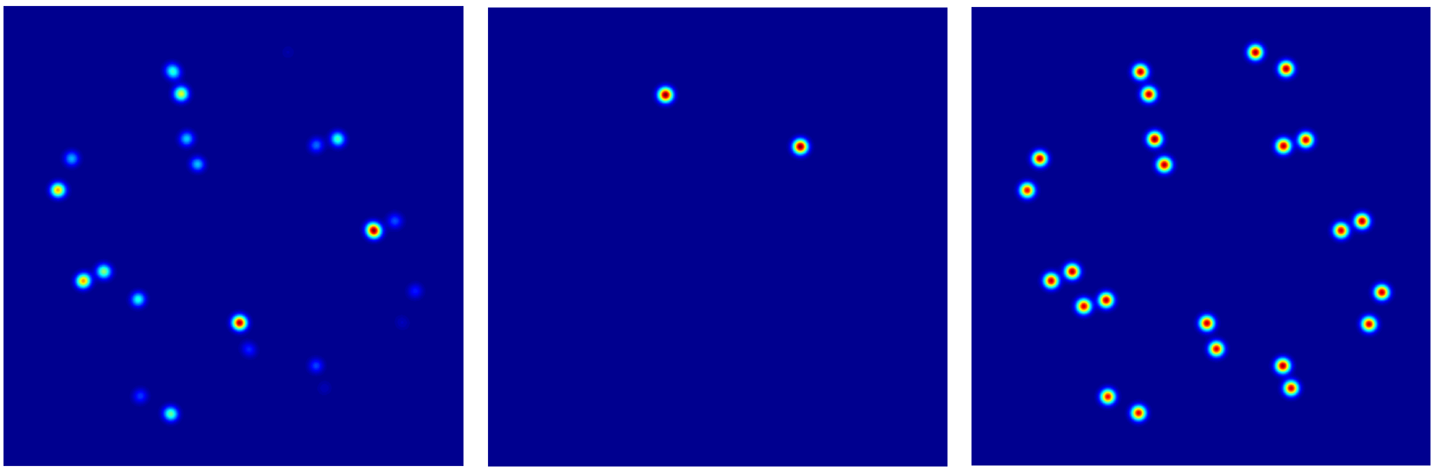

(a) PI-MAIS (b) Naïve IT (c) IT

Experiments: Gaussian Mixture Model

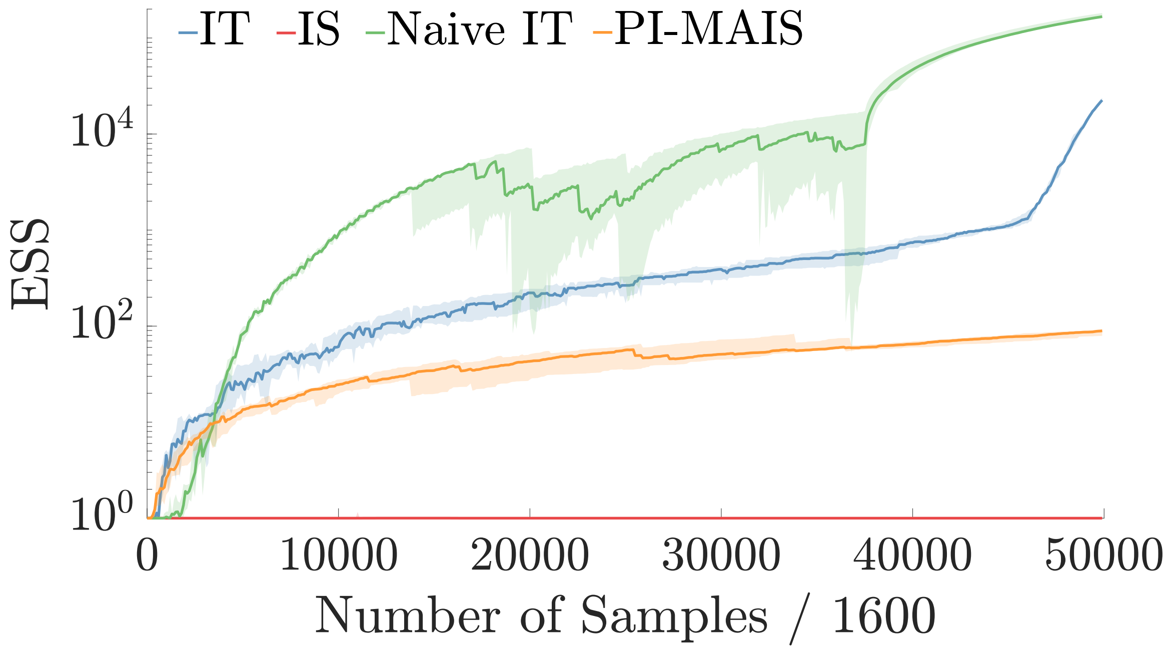

Experiments: Chaotic Dynamics Model

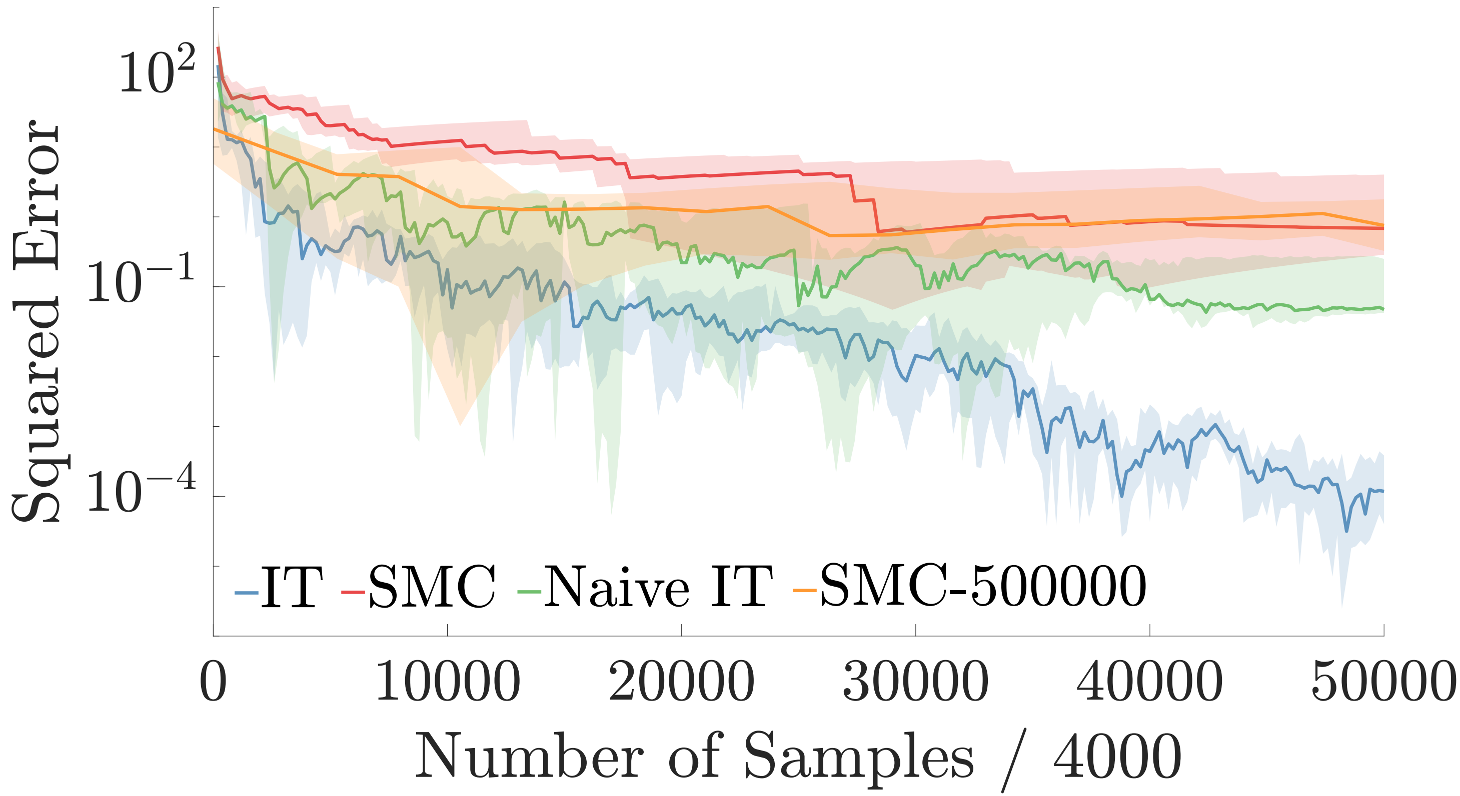

Experiments: Chaotic Dynamics Model

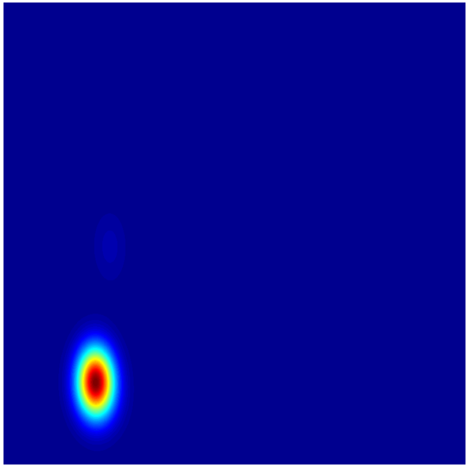

(d) PMMH

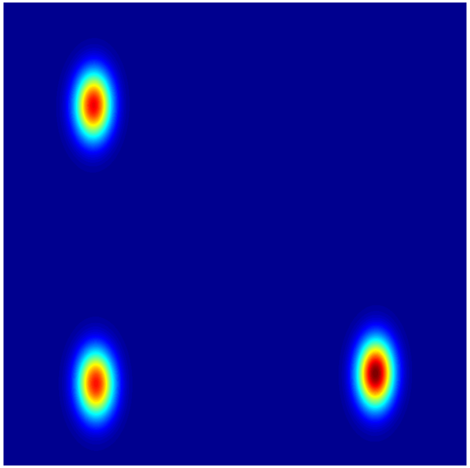

(e) Naïve IT

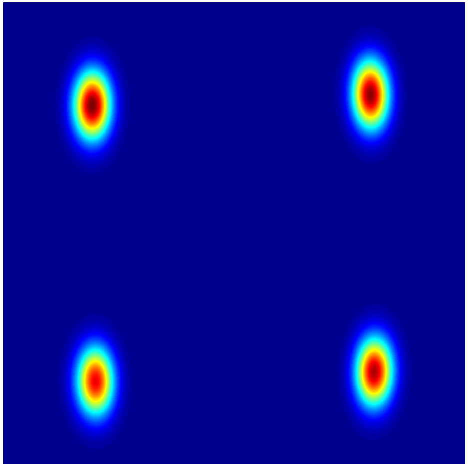

(f) IT

Recap

- Adaptive inference methods need to explore, not just exploit
- New class of inference algorithms — Inference Trees — that can outperform current adaptive importance sampling and SMC methods
- Targeted Exploration: new estimator for predicting presence of significant probability mass using poor quality samples

Thanks!

T. Rainforth, Y. Zhou, X. Lu, Y. W. Teh, F. Wood, H. Yang, and J.-W. van de Meent, “Inference Trees: Adaptive Inference with Exploration,” *arXiv preprint arXiv:1806.09550*, 2018.

rainforth@stats.ox.ac.uk

http://www.robots.ox.ac.uk/~twgr/