Inference Trees: Adaptive Inference with Exploration

Tom Rainforth[1] Yuan Zhou[1] Xiaoyu Lu[1] Yee Whye Teh[1] Frank Wood[1] Hongseok Yang[2] Jan-Willem van de Meent[3]
[1]University of Oxford; [2]KAIST, South Korea; [3]Northeastern University

Overview

- Inference trees (ITs) are a new adaptive Monte Carlo inference method building on ideas from Monte Carlo tree search (MCTS)
- Existing adaptive inference methods are implicitly based on pure exploitation, ITs explicitly aim to balance exploration and exploitation
- ITs use bandit strategies to adaptively sample from hierarchical partitions of the space, while simultaneously learning these partitions
- Particularly effective in combination with sequential Monte Carlo (SMC)

Take Home

- By explicitly allocating computational resources to exploration, we can construct more efficient and robust methods for adaptive inference
- Can capture long range dependencies and potentially improve beyond what can be achieved by proposal adaption alone in SMC

Paritioning the Target Space

- Aim: approximate posterior $p(x_{1:T}|Y),\ x_{1:T} \in \mathbb{X}_T$
- We partition $\mathbb{X}_T$ into separate regions $\{A_i\}_{i\in\mathcal{I}}$
- Run inference separately on each region, e.g. using SMC, then combine
- Allocate computational resources using MCTS
- Partitioning $\mathbb{X}_T$ directly is difficult $\Rightarrow$ partition in cumulative space $\mathfrak{Z}_T$

$$z_t = \eta_t(x_t; x_{1:t-1}) = \int_{-\infty}^{x_t} q_t(x'_t|x_{1:t-1})dx'_t \quad (1)$$

$$q_{i,t}(x_t|x_{1:t-1}, \{z_t \in Z_{i,t}\}) = \frac{q_t(x_t|x_{1:t-1})\mathbb{I}(z_t \in Z_{i,t})}{\int_0^1 \mathbb{I}(z_t \in Z_{i,t})dz_t} \quad (2)$$

- Using orthogonal partitioning of $\mathfrak{Z}_T$: $B_1, B_2, \dots, B_{\mathcal{I}}$
- This implies a non-orthogonal partitioning of $\mathbb{X}_T$: $A_1, A_2, \dots, A_{\mathcal{I}}$

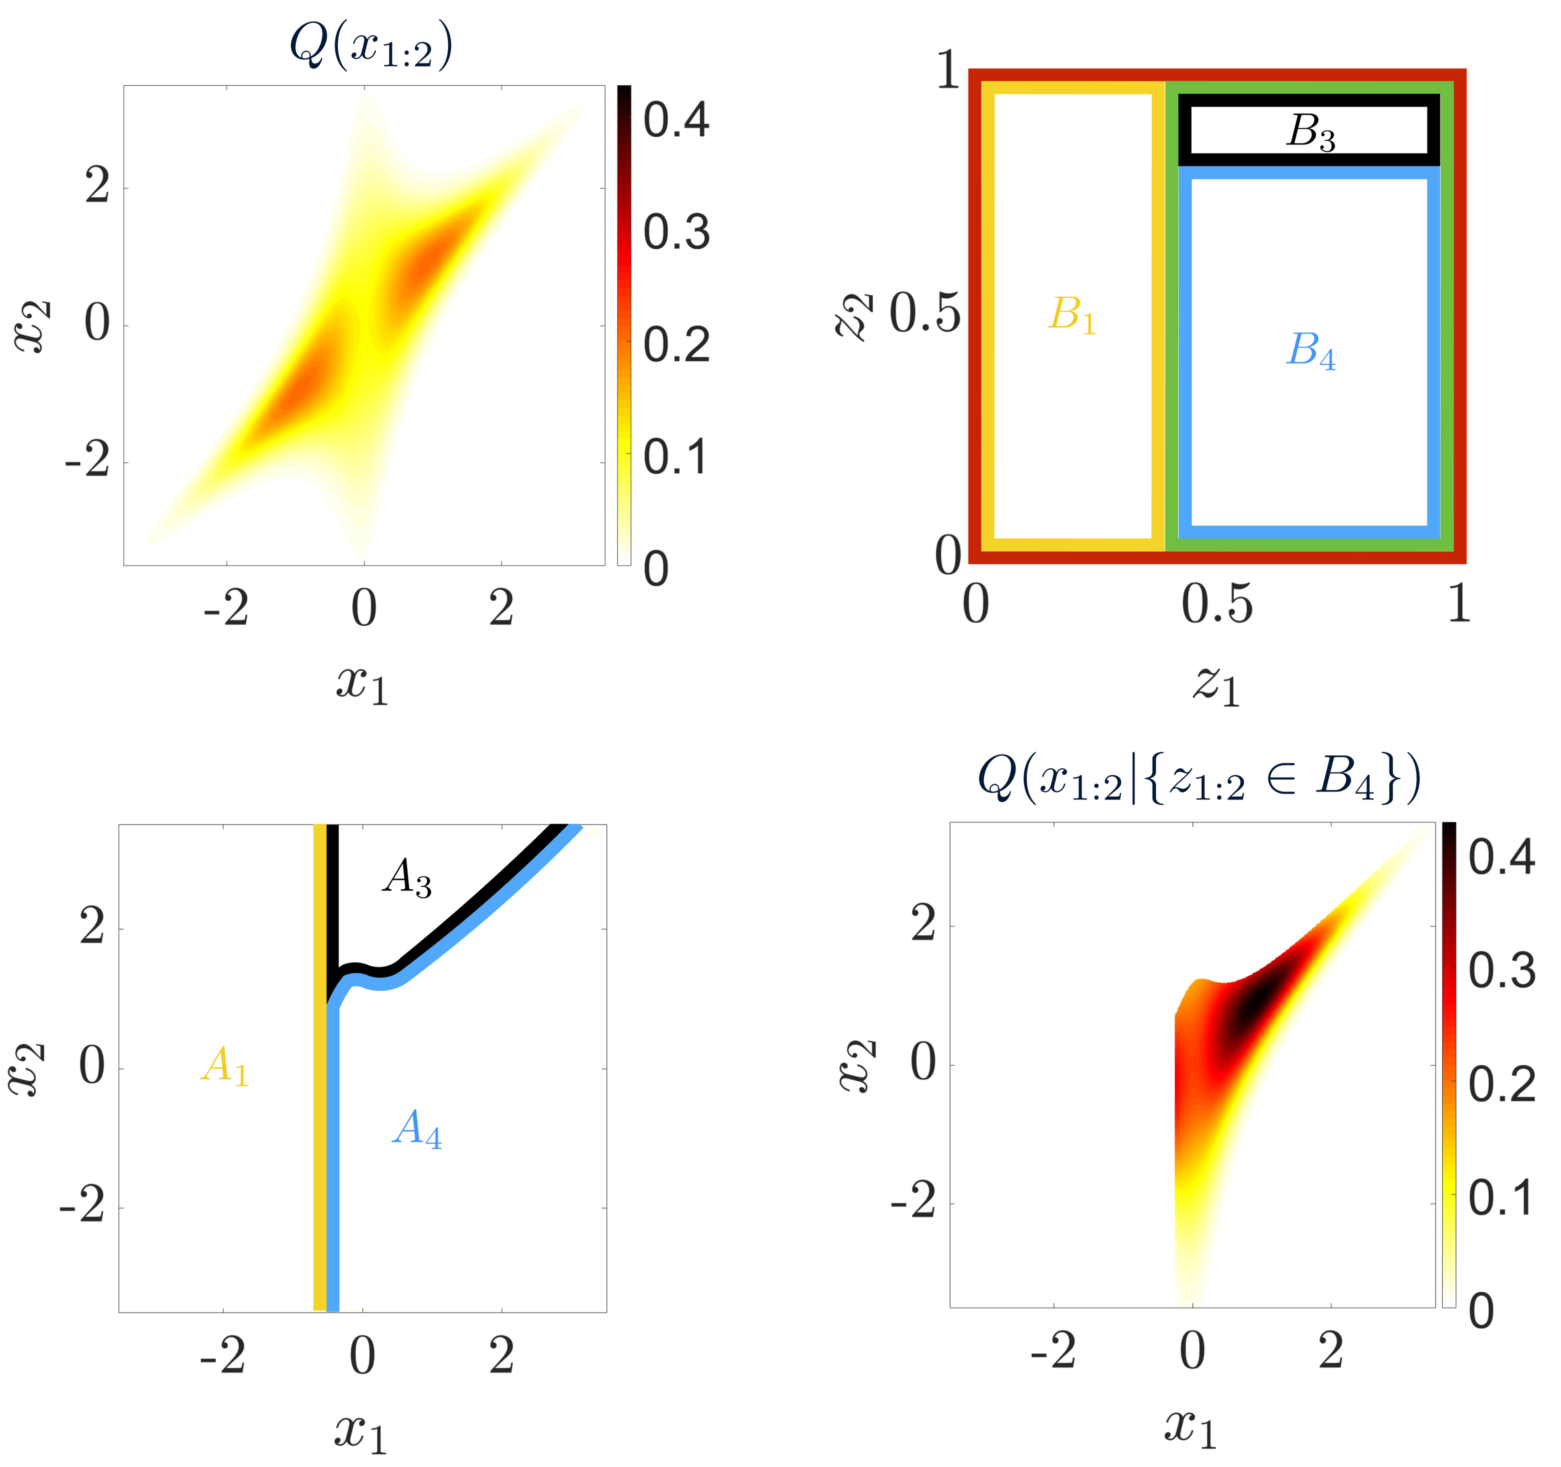

Combining Estimates

- Estimates from different partitions combined in an unweighted fashion

$$\mathbb{E}_{p(x_{1:T}|Y)}[f(x_{1:T})] = \frac{1}{p(Y)}\int p(x_{1:T},Y)f(x_{1:T})\,dx_{1:T}$$

$$= \frac{1}{p(Y)}\sum_{i\in\mathcal{I}} \mathbb{E}_{Q(x_{1:T}|\{z_{1:T}\in B_i\})}\left[\frac{p(x_{1:T},Y)f(x_{1:T})\prod_{t=1}^{T}\mathbb{I}(z_t \in Z_{i,t})}{Q(x_{1:T} \mid \{z_{1:T} \in B_i\})}\right]$$

$$\approx \frac{\sum_{i\in\mathcal{I}}\hat{\varrho}_i}{\sum_{i\in\mathcal{I}}\hat{\upsilon}_i} \quad \text{where} \quad \hat{\varrho}_i = \frac{1}{N_i}\sum_{n=1}^{N_i} w_n^i f\left(\hat{x}_{1:T,i}^n\right), \quad \hat{\upsilon}_i = \frac{1}{N_i}\sum_{n=1}^{N_i} w_n^i$$

- Node estimate is combination of local estimate and children

$$\hat{\mu}_j = \frac{c_j - (c_{l_j} + c_{r_j})}{c_j}\hat{\varrho}_j + \frac{c_{l_j} + c_{r_j}}{c_j}\left(\hat{\mu}_{l_j} + \hat{\mu}_{r_j}\right) \quad (3)$$

Tree Learning Algorithm

Three key components to training:

- **Traversal** – Choose a node to update using upper confidence bounding. This adaptively allocates resources to where they are needed.
- **Refinement** – Either run inference at chosen node or split it and run inference for children. The latter uses the existing samples at the node to choose a good split
- **Propagation** – Update the tree with information from new estimates.

Traversal

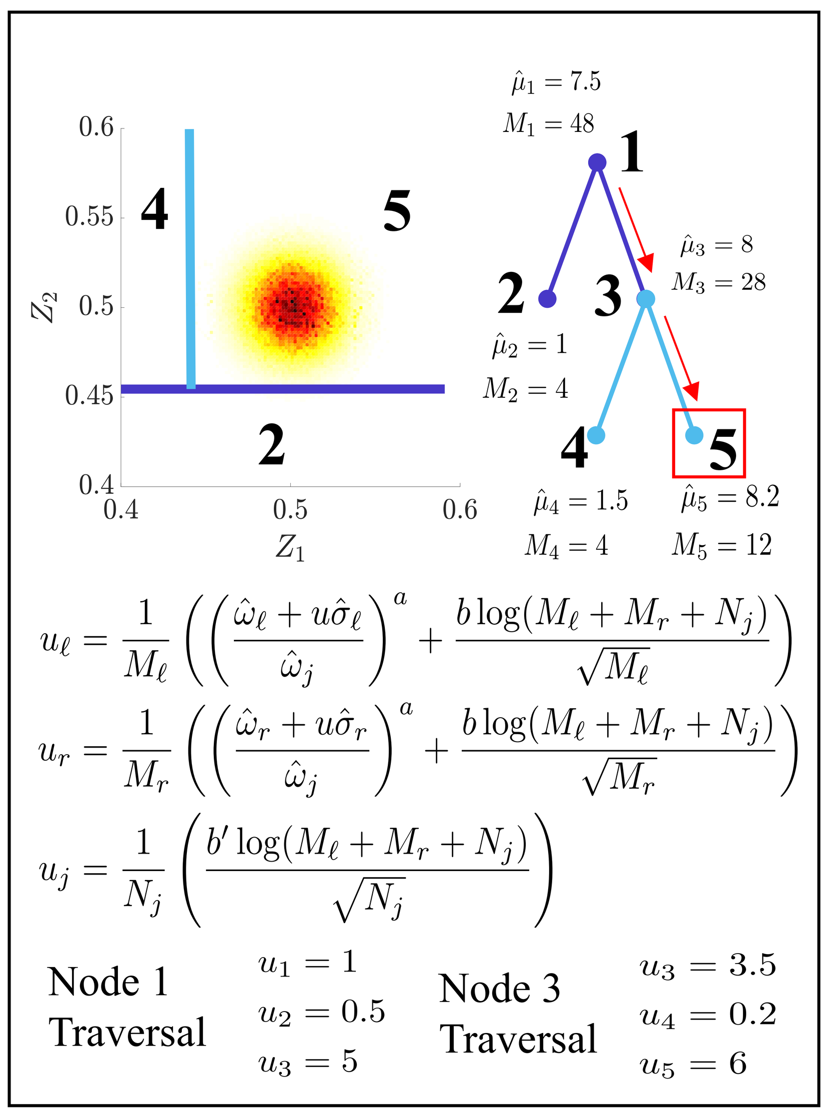

Refinement

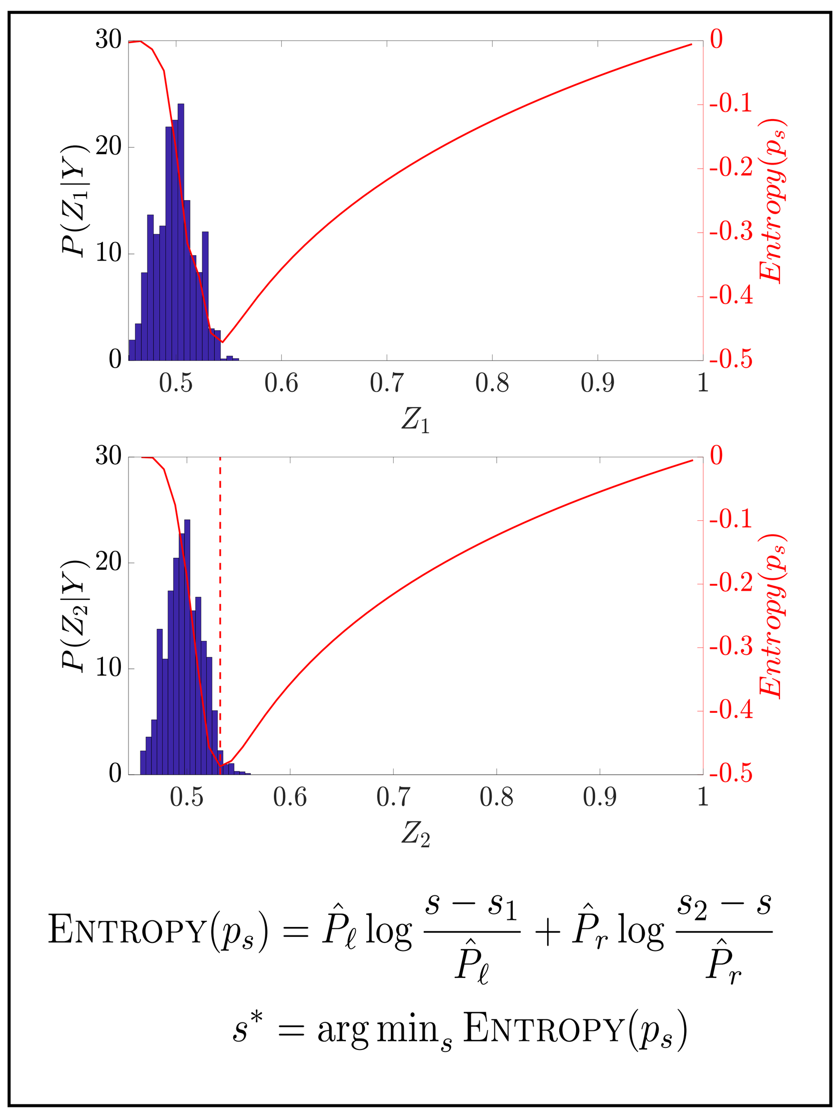

Propagation

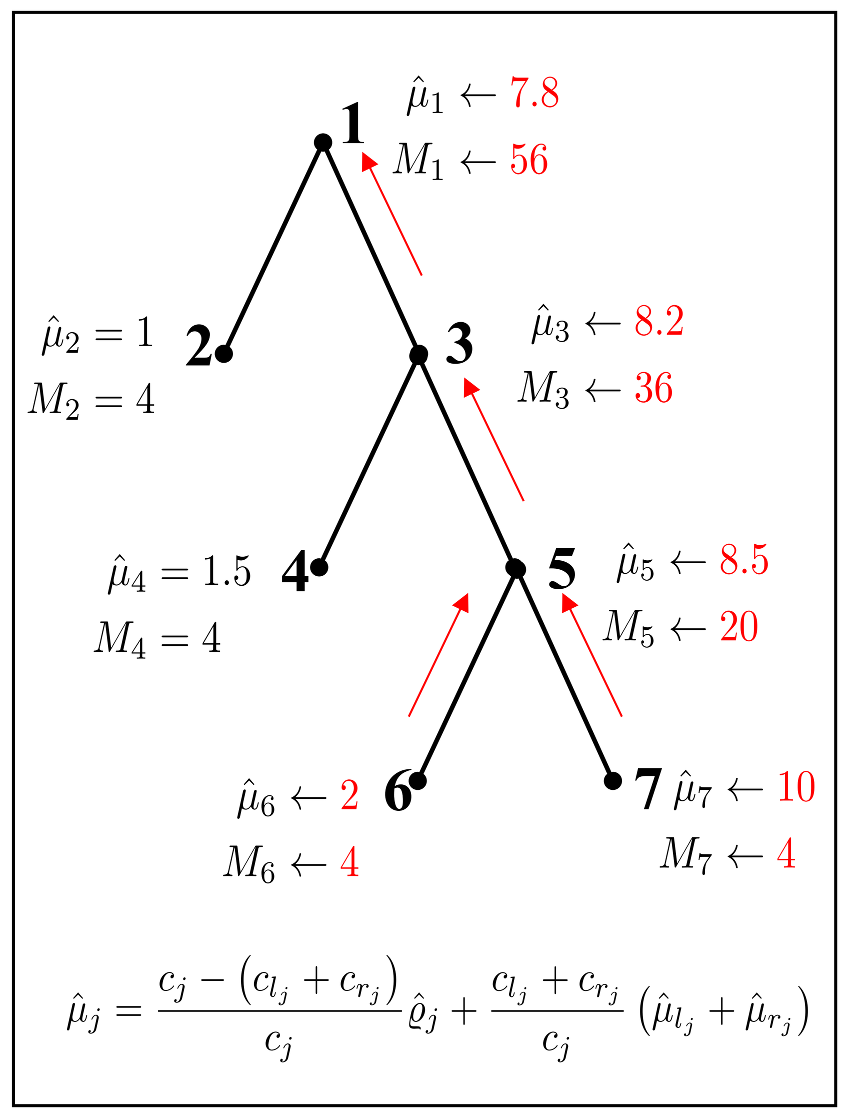

Consistent Regardless of Training Strategy

Consistent regardless of training strategy given some weak assumptions

- Assumption 1 – individual estimators converge
- Assumption 2 – nodes sampled from infinitely often if sibling is also
- Assumption 3 – depth of tree remains bounded

Theorem 1. *If Assumptions 1, 2, and 3 hold, then all estimates calculated using an IT converge weakly as $K \to \infty$.*

Experimental Results

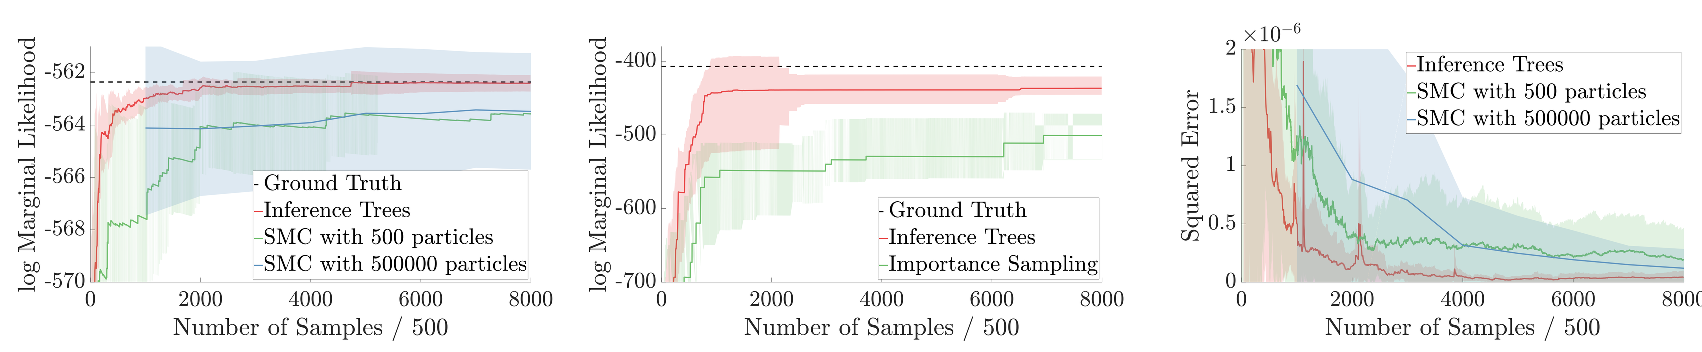

(a) Chaotic Dynamics (b) GMM (c) Network Tail Integral

References

[1] James Neufeld, Andras Gyorgy, Csaba Szepesvari, and Dale Schuurmans. Adaptive Monte Carlo via Bandit Allocation. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pages 1944–1952, 2014.

[2] Alexandra Carpentier, Remi Munos, and András Antos. Adaptive Strategy for Stratified Monte Carlo Sampling. *Journal of Machine Learning Research*, 16:2231–2271, 2015.

Acknowledgements

- BP, ESPRC, DARPA, MSIP

rainforth@stats.ox.ac.uk, yuan.zhou@cs.ox.ac.uk