Low-Variance Gradient Estimates for the Plackett-Luce Distribution

Artyom Gadetsky* Kirill Struminsky* Christopher Robinson
Novi Quadrianto Dmitry Vetrov

NATIONAL RESEARCH UNIVERSITY

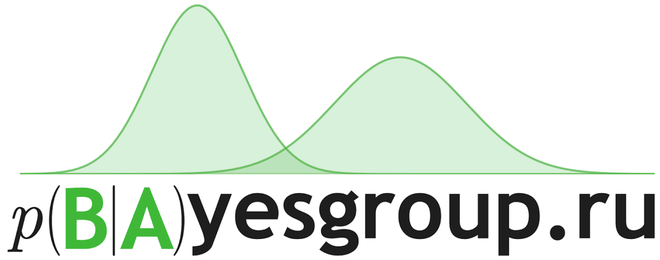

UNIVERSITY OF SUSSEX

The Plackett-Luce Distribution

- *Permutations* occur in:
  1. Information retrieval
  2. Combinatorial optimization (TSP, VRP, QAP, …)
  3. Causal inference
- The Plackett-Luce model defines *a distribution over permutations* $p(b \mid \theta)$
- Various tasks, including (1.-3.), require solving

$$\min_{\theta} \mathbb{E}_{p(b|\theta)} f(b)$$

The Gradient Estimates

- To minimize $\mathbb{E}_{p(b|\theta)}f(b)$ w.r.t θ we need reliable estimates of $\nabla_\theta \mathbb{E}_{p(b|\theta)}f(b)$
- *RELAX* (Grathwohl et al.) is a solution for the expectation w.r.t. *the categorical distribution*
- We extend RELAX to $b \sim p(b \mid \theta)$ from *the Plackett-Luce distribution*

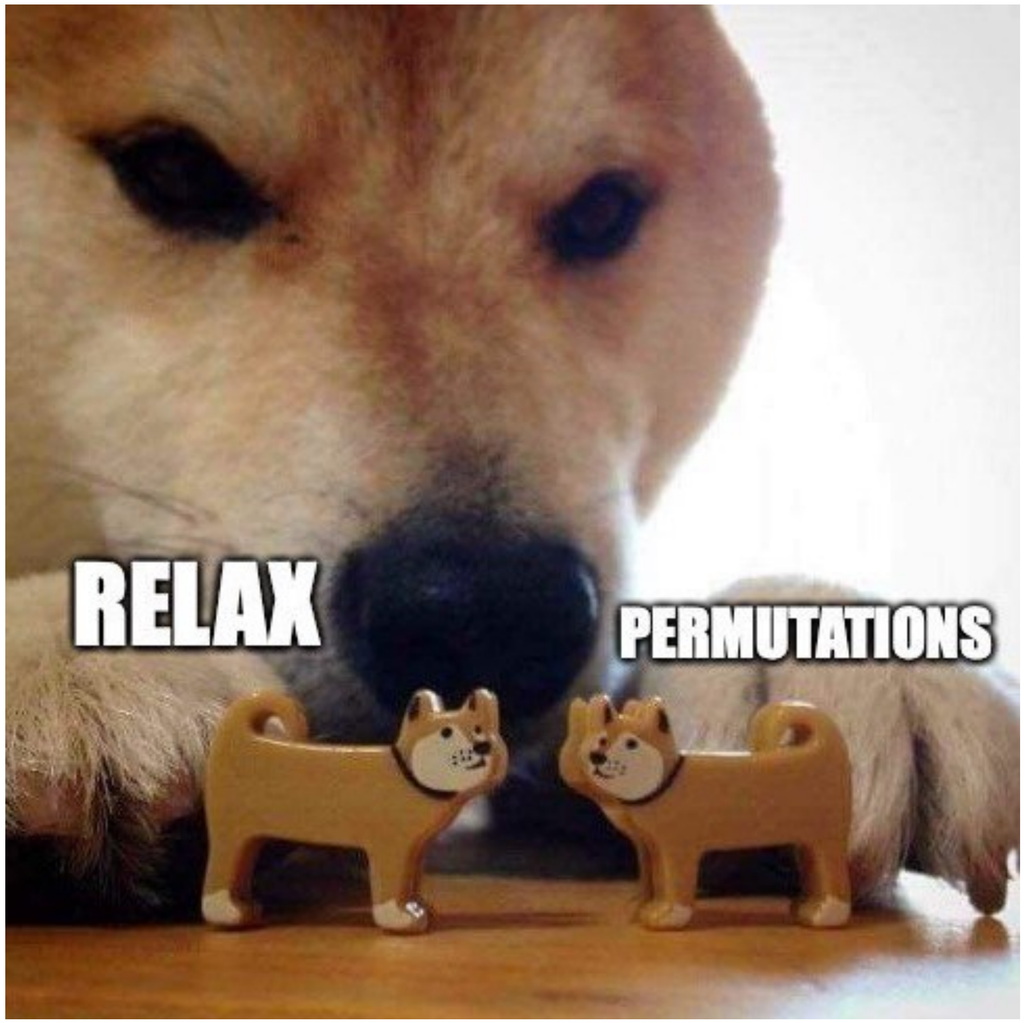

Grathwohl W. et al. Backpropagation through the Void: Optimizing control variates for black-box gradient estimation. – 2018.