

We thank the reviewers for their time and thoughtful feedback. We were glad to see that the reviewers had uniformly positive things to say about the work. Below, we clarify two specific questions.

1. Computing KL divergences. R3 and R4 raised concerns about computing the KL in the absence of a tractable probability density of X_t (induced by the choice of random utility U).

R4: “As far as I understood p_θ distribution is implicit, so I am puzzled about how KL-part of ELBO for spanning tree is calculated in Table 1.”

We computed the KL with respect to the random utility U . There are at least 3 possible KL terms for a relaxed ELBO: the KL from a prior over \mathcal{X} to the distribution of X , the KL from a prior over $\text{conv}(\mathcal{X})$ to the distribution of X_t , or the KL from a prior over \mathbb{R}^n to the distribution of U , see Section C.3 of [52, Maddison et al.] for a discussion of this topic. In our case, since we do not know of an explicit tractable density of X_t or X , we compute the KL with respect to U . This means that for the graph layout experiments, in both training and testing, we computed the KL between $U \sim \text{Gumbel}(\theta)$ and the Gumbel with location 0 prior (see section D.4.2).

R3: “How justifiable is this substitute?”

The KL divergence with respect to U is an *upper-bound* to the KL divergence with respect to X_t , due to a data processing inequality; therefore, the ELBO that we are optimizing is a *lower-bound* to the relaxed variational objective. Whether or not this is a good choice is an empirical question, but it seems to not be an issue for the graph layout experiments. Note, that *when optimizing the relaxed objective*, using a KL divergence with respect to X does not result in a lower-bound to the relaxed variational objective, as it is not necessarily an ELBO for the continuous relaxed model (see again Section C.3 of [52]).

2. Differentiability of the argmax.

R3: “Do the authors refer to a more general scheme that is being invalidated by the jump discontinuities (but not the non-differentiability)?”

That’s a great question. We agree that we should make this clearer. Indeed, we refer to Proposition 2.3 in Chapter 7.2 in [8, Asmussen & Glynn]. This gives a slightly more general condition for the exchange of expectation and differentiation than Leibniz rule. It does not require the existence of partial derivatives everywhere, but instead poses a Lipschitz condition. The jump discontinuities in SMTs violate this condition, see also Remark 2.6 in [8]. In contrast, the Euclidean projection for SSTs does not. As a result, it admits a reparameterization gradient, even though its partial derivatives do not exist everywhere.

If accepted, we plan to use the additional page allowance in the following way. First, we would add text addressing the clarifications requested by the reviewers. Second, we’ve since explored additional score function estimator baselines. We would add these to the appendix of the graph layout experiments. The baselines improved with our efforts, but the conclusions are unchanged (SSTs outperforming the other gradient estimator baselines). Finally, we would use the additional page allowance to add Fig. 1 (below) to aid in intuition.




Figure 1: Structured discrete objects can be represented by binary arrays. In these graphical representations, color indicates 1 and no color indicates 0. For example, “Spanning tree” is the adjacency matrix of an undirected spanning tree over 6 nodes; “Arborescence” is the adjacency matrix of a directed spanning tree rooted at node 3.