

Sussex Research

Low-variance black-box gradient estimates for the Plackett-Luce distribution

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

Publication date

18-06-2020

Licence

This work is made available under the All Rights Reserved licence and should only be used in accordance with that licence. For more information on the specific terms, consult the repository record for this item.

Document Version

Accepted version

Citation for this work (American Psychological Association 7th edition)

Gadetsky, A., Struminsky, K., Robinson, C., Quadrianto, N., & Vetrov, D. (2020). *Low-variance black-box gradient estimates for the Plackett-Luce distribution* (Version 1). University of Sussex. https://hdl.handle.net/10779/uos.23306855.v1

Published in

Thirty-Fourth AAAI Conference on Artificial Intelligence

Link to external publisher version

https://doi.org/10.1609/aaai.v34i06.6572

Copyright and reuse:

This work was downloaded from Sussex Research Open (SRO). This document is made available in line with publisher policy and may differ from the published version. Please cite the published version where possible. Copyright and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners unless otherwise stated. For more information on this work, SRO or to report an issue, you can contact the repository administrators at sro@sussex.ac.uk. Discover more of the University's research at https://sussex.figshare.com/

Conditional Reparametrization for the Plackett-Luce Distribution

We prove Proposition 1 in this section. We first discuss the properties of Gumbel distribution. Then we discuss the generative processes for the densities used for $p(z \mid b, \theta)$ in Eq. 20. Then we show that $p(b \mid z)p(z \mid \theta) = p(b \mid \theta)p(z \mid b, \theta)$ for the unconditional Gumbel density $p(z \mid \theta)$ and the Plackett-Luce distribution $p(b \mid \theta)$.

Density for the Gumbel distribution and the truncated Gumbel distribution

The density function of the Gumbel distribution with location parameter μ is

$$\phi_{\mu}(z) = \exp(-z + \mu) \exp(-\exp(-z + \mu))$$
 (31)

and the cumulative density function is

$$\Phi_{\mu} = \exp(-\exp(-z+\mu)). \tag{32}$$

Our derivation of the conditional distribution $p(b \mid z, \theta)$ relies on the additive property of the cumulative density function of the Gumbel distribution

$$\Phi_{\log(\exp\mu + \exp\nu)}(z) = \exp(-\exp(z)(\exp\mu + \exp\nu)) = \Phi_{\mu}(z)\Phi_{\nu}(z), \quad (33)$$

which we enfold in the following auxiliary claim.

Lemma 2. For permutation $b \in S_k$, score vector $\theta \in \mathbb{R}^k$ and i = 1, ..., k and the argument vector $z \in \mathbb{R}^k$ we have

$$\phi_{\theta_{b_i}}(z_{b_i})\Phi_{\log(\sum_{j=i+1}^k \exp \theta_{b_j})}(z_{b_i}) \tag{34}$$

$$= \frac{\exp \theta_{b_i}}{\sum_{j=i}^k \exp \theta_{b_j}} \phi_{\log(\sum_{j=i}^k \exp \theta_{b_j})}(z_{b_i}).$$
(35)

Proof. For brevity, we denote $\exp \theta_i$ as p_i . We then rewrite the density $\phi_{\log p_{b_i}}(z_{b_i})$ through the exponent $\exp(-z_{b_i} + \log p_{b_i})$ and c.d.f. $\Phi_{\log p_{b_i}}(z_{b_i})$ and apply the additive property in Eq. 38:

$$\phi_{\log p_{b_i}}(z_{b_i})\Phi_{\log(\sum_{j=i+1}^k p_{b_j})}(z_{b_i}) \tag{36}$$

$$= p_{b_i} \exp(-z_{b_i}) \Phi_{\log p_{b_i}}(z_{b_i}) \Phi_{\log(\sum_{j=i+1}^k p_{b_j})}(z_{b_i}) \quad (37)$$

= $p_k \exp(-z_k) \Phi_{b_i} \exp(-z_{b_i}) \quad (38)$

$$= p_{b_i} \exp(-z_{b_i}) \Phi_{\log(\sum_{j=i}^k p_{b_j})}(z_{b_i})$$
(38)
$$\sum_{j=1}^k p_{b_j}(z_{b_j}) = \sum_{j=1}^k p_{b_j}(z_{b_j}) = \sum_{j=1}^$$

$$= p_{b_i} \frac{\sum_{j=i}^{j=i} p_{b_j}}{\sum_{j=i}^{k} p_{b_j}} \exp(-z_{b_i}) \Phi_{\log(\sum_{j=i}^{k} p_{b_j})}(z_{b_i})$$
(39)

$$= \frac{p_{b_i}}{\sum_{j=1}^k p_{b_j}} \phi_{\log(\sum_{j=i}^k p_{b_j})}(z_{b_i}).$$
(40)

The last step collapses the exponent and the c.d.f. into the density function $\phi_{\log(\sum_{i=1}^{k} p_{b_i})}(z_{b_i})$.

Finally, to define the density of conditional distribution $p(b \mid z, \theta)$ we define the density of the truncated Gumbel distribution $\phi_{\mu}^{z_0}(z) \propto \phi_{\mu}(z)I[z \leq z_0]$:

$$\phi_{\mu}^{z_0}(z) = \frac{\phi_{\mu}(z)}{\Phi_{\mu}(z_0)}(z)I[z \le z_0], \tag{41}$$

where the superscript z_0 denotes the truncation parameter.

Reparametrization for the Gumbel distribution and the truncated Gumbel distribution

The reparametrization trick requires representing a draw from a distribution as a deterministic transformation of a fixed distribution sample and a distribution parameter. For a sample z from the Gumbel distribution $\mathcal{G}(\mu, 1)$ with location parameter μ the representation is

$$u = \mu - \log(-\log v), \ v \sim \text{uniform}[0, 1].$$
(42)

For the Gumbel distribution truncated at z_0 (Maddison, Tarlow, and Minka 2014) proposed an analogous representation

2

$$z = \mu - \log(-\log v + \exp(-z_0 + \mu))$$
$$= -\log\left(-\frac{\log v}{\exp \mu} + \exp(-z_0)\right)$$
(43)

$$v \sim \operatorname{uniform}[0, 1].$$
 (44)

In particular, the sampling schemes in Eq. 10 and Eq. 20 generate samples from the truncated Gumbel distribution.

The derivation of the conditional distribution

We now derive the conditional distribution and the sampling scheme proposed in Proposition 1.

The joint distribution of the permutation b and the Gumbel samples z is

$$p(b, z \mid \theta) = p(b \mid z)p(z \mid \theta)$$
(45)

$$= \phi_{\theta_{b_1}}(z_{b_1}) \prod_{i=2}^{n} \left(\phi_{\theta_{b_i}}(z_{b_i}) I[z_{b_{i-1}} \ge z_{b_i}] \right)$$
(46)

We first multiply and divide the joint density by the c.d.f. $\Phi_{\log(\sum_{i=2}^{k} \exp \theta_{b_i})}(z_{b_1})$ and apply Lemma 2

$$\frac{\Phi_{\log(\sum_{i=2}^{k} \exp \theta_{b_i})}(z_{b_1})}{\Phi_{\log(\sum_{i=2}^{k} \exp \theta_{b_i})}(z_{b_1})}\phi_{\theta_{b_1}}(z_{b_1})\prod_{i=2}^{k}\dots$$
(47)

$$= \frac{\exp \theta_{b_1}}{\sum_{i=1}^k \exp \theta_{b_i}} \frac{\phi_{\log(\sum_{i=1}^k \exp \theta_{b_i})}(z_{b_1})}{\Phi_{\log(\sum_{i=2}^k \exp \theta_{b_i})}(z_{b_1})} \prod_{i=2}^k \dots$$
(48)

Next, we apply Lemma 2 to combine the c.d.f. in the denominator $\Phi_{\log(\sum_{j=i}^{k} \exp \theta_{b_j})}(z_{b_{i-1}})$ and the term $\phi_{\theta_{b_i}}(z_{b_i})I[z_{b_{i-1}} \ge z_{b_i}]$ inside the product

$$\frac{\phi_{\theta_{b_i}}(z_{b_i})I[z_{b_{i-1}} \ge z_{b_i}]}{\Phi_{\log(\sum_{j=i}^k \exp \theta_{b_j})}(z_{b_{i-1}})}$$

$$\tag{49}$$

$$=\frac{\phi_{\theta_{b_i}}(z_{b_i})I[z_{b_{i-1}} \ge z_{b_i}]}{\Phi_{\log(\sum_{i=i}^k \exp \theta_{b_i})}(z_{b_{i-1}})} \frac{\Phi_{\log(\sum_{j=i+1}^k \exp \theta_{b_j})}(z_{b_i})}{\Phi_{\log(\sum_{j=i+1}^k \exp \theta_{b_j})}(z_{b_i})}$$
(50)

$$= \frac{\exp \theta_{b_i}}{\sum_{j=i}^k \exp \theta_{b_j}} \frac{\phi_{\log(\sum_{j=i}^k \exp \theta_{b_i})}^{s_{i-1}}(z_{b_i})}{\Phi_{\log(\sum_{j=i+1}^k \exp \theta_{b_j})}(z_{b_i})}$$
(51)

and obtain the truncated distribution $\phi_{\log(\sum_{j=i}^{k} \exp \theta_{b_i})}^{z_{b_{i-1}}}(z_{b_i})$ along with one factor of the Plackett-Luce probability $\frac{\exp \theta_{b_i}}{\sum_{j=i}^{k} \exp \theta_{b_j}}$. Also, after the transformation the summation index in the denominator c.d.f. changes from *i* to i + 1. This gives us an induction step that we apply sequentially for i = 2, ..., k - 1. For i = k the denominator c.d.f. $\Phi_{\log \exp \theta_k}(z_{b_{k-1}})$ and the product term $\phi_{\log \exp \theta_k}(z_{b_k})I[z_{k-1} \ge z_k]$ combine into the truncated Gumbel distribution with density $\phi_{\log \exp \theta_k}^{z_{b_{k-1}}}(z_{b_k})$. As a result, we rearrange $p(b, z \mid \theta)$ into the product of the truncated Gumbel distribution densities $\pi(z \mid b, \theta)$ and

As a result, we rearrange $p(b, z \mid \theta)$ into the product of the truncated Gumbel distribution densities $p(z \mid b, \theta)$ and the probability of the Plackett-Luce distribution $p(b \mid \theta)$:

$$\phi_{\log\sum_{j=1}^{k} \exp\theta_{b_j}}(z_{b_1}) \prod_{i=1}^{k} \frac{\exp\theta_{b_i}}{\sum_{j=i}^{k} \exp\theta_{b_j}} \prod_{i=2}^{k} \phi_{\log\sum_{j=i}^{k} \exp\theta_j}^{z_{b_{i-1}}}(z_{b_i}).$$
(52)

Finally, to obtain the claim of Proposition 1 we the replace the sum of exponents with a normalized sum $\Theta_i = \frac{\sum_{j=i}^{k} \exp \theta_{b_j}}{\sum_{j=1}^{k} \exp \theta_{b_j}}$ and apply the reparametrized sampling scheme defined in Eq. 43.