Discrete approximation of a mixture distribution via restricted divergence

Christian Röver and Tim Friede*

Department of Medical Statistics,
University Medical Center Göttingen

February 17, 2017

Abstract

Mixture distributions arise in many application areas, for example as marginal distributions or convolutions of distributions. We present a method of constructing an easily tractable discrete mixture distribution as an approximation to a mixture distribution with a large to infinite number, discrete or continuous, of components. The proposed DIRECT (Divergence Restricting Conditional Tesselation) algorithm is set up such that a pre-specified precision, defined in terms of Kullback-Leibler divergence between true distribution and approximation, is guaranteed. Application of the algorithm is demonstrated in two examples.

Keywords: mixture distribution, discrete approximation, convolution, DIRECT.

Reference as:
C. Röver, T. Friede (2017). Discrete approximation of a mixture distribution via restricted divergence. *Journal of Computational and Graphical Statistics* **26**(1), 217–222.
doi: 10.1080/10618600.2016.1276840

*This project has received funding from the European Union's *Seventh Framework Programme for research, technological development and demonstration* under grant agreement number FP HEALTH 2013-602144 "*Innovative methodology for small populations research (InSPiRe)*".

1 Introduction

Mixture distributions with a large to infinite number of mixture components commonly occur in many fields of application (e.g., Seidel 2010). Common examples include e.g. marginal (posterior) distributions, convolutions of random variables, predictive distributions, distributions of test statistics, overdispersed sampling distributions, and many more.

If the mixture distribution's exact marginal density, distribution or quantile functions are not available in analytical form, then practical application of such mixtures is often very limited. Such mixtures may then often be approximated to a sufficient degree by a mixture of a lower, *finite* number of components. How exactly to select such a finite set of components however is not obvious. In the following we describe a general approach and an algorithm allowing to set up a finite mixture as an approximation to a mixture distribution with a large or infinite number of components in a completely automated way. The construction is based on the *Kullback-Leibler divergence* or *relative entropy* between distributions and as such aims at bounding the (expected) logarithmic ratio of exact and approximate probability densities.

2 Kullback-Leibler divergence

2.1 Definitions

The *Kullback-Leibler divergence* or *relative entropy* of two probability distributions with probability density functions p and q is defined as the *expected logarithmic ratio of densities* with respect to the former distribution (p),

$$\mathcal{D}_{\mathrm{KL}}\big(p(\theta)\big\|q(\theta)\big) \;=\; \int_{\Theta} \log\Big(\frac{p(\theta)}{q(\theta)}\Big)\, p(\theta)\,\mathrm{d}\theta \;=\; \mathrm{E}_{p(\theta)}\bigg[\log\Big(\frac{p(\theta)}{q(\theta)}\Big)\bigg] \tag{1}$$

(Cover & Thomas 1991, Ch. 2). In case of discrete probability distributions p and q, the integrals simplify to sums, but for simplicity we will stick to the integral notation in the following. The relative entropy is always positive, it is zero if the two distributions are identical ($p = q$), and larger otherwise. The divergence (in general) is *not* symmetric: $\mathcal{D}_{\mathrm{KL}}\big(p(\theta)\big\|q(\theta)\big) \neq \mathcal{D}_{\mathrm{KL}}\big(q(\theta)\big\|p(\theta)\big)$. The *symmetrized (KL-) divergence* is defined as

$$\mathcal{D}_{\mathrm{s}}\big(p(\theta)\big\|q(\theta)\big) \;=\; \mathcal{D}_{\mathrm{KL}}\big(p(\theta)\big\|q(\theta)\big) + \mathcal{D}_{\mathrm{KL}}\big(q(\theta)\big\|p(\theta)\big) \tag{2}$$

(Kullback & Leibler 1951). Unlike the *directed* divergence, $\mathcal{D}_{\mathrm{s}}$ is obviously symmetric. Note that, trivially but importantly,

$$\mathcal{D}_{\mathrm{s}}\big(p(\theta)\big\|q(\theta)\big) \;\geq\; \max\big\{\mathcal{D}_{\mathrm{KL}}\big(p(\theta)\big\|q(\theta)\big),\, \mathcal{D}_{\mathrm{KL}}\big(q(\theta)\big\|p(\theta)\big)\big\}, \tag{3}$$

i.e., the symmetrized divergence bounds both individual directed divergences. For simplicity, in the following we will mostly be focusing on *symmetrized* KL-divergences.

For example, the Kullback-Leibler divergence for two normal distributions with mean and variance parameters (μ_A, σ_A^2) and (μ_B, σ_B^2), respectively, is given by

$$\mathcal{D}_{\mathrm{KL}}\big(p(\theta|\mu_A,\sigma_A)\big\|p(\theta|\mu_B,\sigma_B)\big) \;=\; \tfrac{1}{2}\Big(\tfrac{(\mu_A-\mu_B)^2}{\sigma_B^2} + \tfrac{\sigma_A^2}{\sigma_B^2} + \log\big(\tfrac{\sigma_B^2}{\sigma_A^2}\big) - 1\Big) \tag{4}$$

(Kullback 1959, Ch. 9). The symmetrized divergence then results as

$$\mathcal{D}_{\mathrm{s}}\big(p(\theta|\mu_A,\sigma_A)\big\|p(\theta|\mu_B,\sigma_B)\big) \;=\; \tfrac{(\mu_A-\mu_B)^2}{\left(\frac{1}{2}(\sigma_A^{-2}+\sigma_B^{-2})\right)^{-1}} + \tfrac{(\sigma_A^2-\sigma_B^2)^2}{2\,\sigma_A^2\,\sigma_B^2}. \tag{5}$$

2.2 Motivation and interpretation

The Kullback-Leibler divergence is generally regarded as a measure of *discrepancy* between probability distributions. For example, when a simple parametric approximation to a more complicated distribution is sought, the approximation may reasonably be matched against the true distribution via minimization of the divergence (Bernardo & Smith 1994, O'Hagan 1994).

The divergence $\mathcal{D}_{\mathrm{KL}}$ relates to the logarithmic ratio of densities. The domain of main interest here is the limit of very similar p and q, i.e., almost equal numerator and denominator, when the density ratio is close to unity. In that case the logarithmic ratio approximately corresponds to the "relative difference" in densities: since $\log(x) \approx x-1$ for $x \approx 1$ (and hence $\log(a/b) \approx a/b - 1$ for $a \approx b$), a divergence of, say, 0.01 approximately corresponds to an (expected) 1% difference between numerator and denominator.

While there is no simple connection relating the divergence of two distributions to their moments, one can get an impression by considering the generic case of two normal distributions. For some fairly obvious parameter choices we get:

$$\sigma_B = \sigma_A,\;\; \mu_B = \mu_A + c\sigma_A \qquad \Rightarrow \qquad \mathcal{D}_{\mathrm{KL}}(p\|q) = \tfrac{1}{2}c^2, \tag{6}$$

$$\mathcal{D}_{\mathrm{s}}(p\|q) = c^2, \tag{7}$$

and

$$\mu_B = \mu_A,\ \ \sigma_B = (1+c)\sigma_A \qquad \Rightarrow \qquad \mathcal{D}_{\mathrm{KL}}(p\|q) = \tfrac{1}{2(1+c)^2} + \log(1+c) - \tfrac{1}{2} \approx c^2, \quad (8)$$

$$\mathcal{D}_{\mathrm{s}}(p\|q) = \tfrac{c^2(c+2)^2}{2(c+1)^2} \approx 2c^2, \quad (9)$$

where the latter approximations follow from Taylor expansion around $c = 0$.

From the above we can see that, for example, for equal variances, a difference in means by, say, 1% of a standard deviation corresponds to a symmetrized divergence $\mathcal{D}_{\mathrm{s}} = 0.01^2 = 0.0001$. For equal means on the other hand, standard deviations differing by 1% correspond to a symmetrized divergence of ≈ 0.0002.

3 Mixture distributions and discrete approximations

3.1 Definitions

Suppose a random variable Y follows a distribution with density $p(y|x)$ that depends on a parameter x. If that parameter is not fixed, but again is a random variable (X) with density $p(x)$, then the (marginal) distribution of Y is called a *mixture distribution*. The joint density of X and Y is given by $p(x, y) = p(y|x) \times p(x)$. What is commonly of interest is the marginal (unconditional) distribution of Y, whose density results by integration as $p(y) = \int p(x, y)\,\mathrm{d}x = \int p(y|x)\,p(x)\,\mathrm{d}x$. The (marginal) distribution of the underlying variable that is conditioned upon, $p(x)$, is called the *mixing distribution* (Seidel 2010) or *latent distribution* (Lindsay 1995).

Mixture distributions arise frequently in statistical problems, for example as marginal (posterior) distributions or as convolutions of random variables. In the following we will assume that X is one-dimensional, and that the domain of X is the real line, or a subset thereof (continuous or discrete).

3.2 Binning

In order to transition from continuous to discrete mixtures, we define a *binning* of the domain of X. Let $\{x_{(1)}, x_{(2)}, \ldots, x_{(k-1)}\} \subset \mathbb{R}$ be a set of bin margins with $x_{(1)} < x_{(2)} <$

$\cdots < x_{(k-1)}$. These define the (exhaustive and disjoint) set of k bins $\{\mathcal{X}_i\}_{i=1,\ldots,k}$ with

$$\mathcal{X}_i \;=\; \begin{cases} \{x : x \le x_{(1)}\} & \text{if } i = 1 \\ \{x : x_{(i-1)} < x \le x_{(i)}\} & \text{if } 1 < i < k \\ \{x : x_{(k-1)} < x\} & \text{if } i = k. \end{cases} \tag{10}$$

In addition, the set of k points $\{\tilde{x}_1, \ldots, \tilde{x}_k\}$ with $\tilde{x}_i \in \mathcal{X}_i$ defines a set of *reference points*, one for each bin. Each bin also has a probability π_i (with respect to $p(x)$) associated, which is given by

$$\pi_i \;=\; \mathrm{P}\big(x_{(i-1)} < x \le x_{(i)}\big) \;=\; \mathrm{P}\big(x \in \mathcal{X}_i\big). \tag{11}$$

3.3 The binned mixture

In addition to the probability density $p(x, y)$ given above, we define another probability distribution with density $q(x, y)$ that has the same marginal density (mixing distribution)

$$q(x) \;=\; p(x), \tag{12}$$

and whose conditional probability density is given by

$$q(y|x) \;=\; p(y|x{=}\tilde{x}_i) \quad \text{for } x \in \mathcal{X}_i. \tag{13}$$

So q is similar to p, but instead of conditioning on the "exact" x value as in the original definition above, this probability distribution conditions on the corresponding bin's reference value $\tilde{x}_i$, depending on which bin x belongs to. The joint distribution of X and Y again is defined through its joint density: $q(x, y) = q(x) \times q(y|x)$. The marginal density of Y again turns out as $q(y) = \int q(y|x)\, q(x)\, \mathrm{d}x$. Equivalently, the binning may be considered a discretization of the mixing distribution while keeping the conditional distribution the same. The discretized mixing distribution simply has the reference points $\{\tilde{x}_1, \ldots, \tilde{x}_k\}$ as its domain, while the associated bin probabilities $\{\pi_1, \ldots, \pi_k\}$ define the probability mass function. The reference points consequently act as "support points" for the discretized mixing distribution here; alternating between these points of view is sometimes helpful.

This "binned" approximation to the joint distribution of (X, Y) is useful, as the resulting marginal distribution of Y, $q(y)$, is a discrete sum of conditional densities (rather than an

integral), making numerical evaluation very easy. The marginal density simplifies to

$$q(y) \;=\; \sum_{i=1}^{k} \pi_i \, p(y|\tilde{x}_i). \tag{14}$$

Analogously, the cumulative distribution function (CDF) may also be expressed as a weighted sum of the component CDFs. Random number generation as well as computation of moments for finite mixtures is also straightforward (Lindsay 1995).

4 Constructing binned mixture approximations

4.1 Some preliminary results

For each bin i define the maximum symmetrized KL-divergence

$$d_i \;=\; \max_{x\in\mathcal{X}_i}\Big\{\mathcal{D}_{\mathrm{s}}\big(p(y|x)\big\|p(y|\tilde{x}_i)\big)\Big\} \;=\; \max_{x\in\mathcal{X}_i}\Big\{\mathcal{D}_{\mathrm{s}}\big(p(y|x)\big\|q(y|x)\big)\Big\}, \tag{15}$$

i.e., the maximum (symmetrized) divergence between distributions $p(y|x)$ corresponding to points within the ith bin and the corresponding ith reference point.

The *chain rule for relative entropy* states that

$$\mathcal{D}_{\mathrm{KL}}\big(p(x,y)\big\|q(x,y)\big) \;=\; \mathcal{D}_{\mathrm{KL}}\big(p(x)\big\|q(x)\big) + \mathrm{E}_{p(x)}\Big[\mathcal{D}_{\mathrm{KL}}\big(p(y|x)\big\|q(y|x)\big)\Big] \tag{16}$$

(Cover & Thomas 1991, Sec. 2.5). In other words, the divergence of two joint distributions is the sum of the divergence of the marginals and the expected divergence of the conditionals. Note that the expectation in (16) is also known as the *conditional relative entropy* (Cover & Thomas 1991, Sec. 2.5). For the *symmetrised* divergence immediately follows an analogous property:

$$\mathcal{D}_{\mathrm{s}}\big(p(x,y)\big\|q(x,y)\big) \;=\; \mathcal{D}_{\mathrm{s}}\big(p(x)\big\|q(x)\big) + \mathrm{E}_{p(x)}\Big[\mathcal{D}_{\mathrm{s}}\big(p(y|x)\big\|q(y|x)\big)\Big]. \tag{17}$$

In our case we have identical marginal distributions for X under both distributions, $p(x)=q(x)$, so that

$$\mathcal{D}_{\mathrm{KL}}\big(p(x)\big\|q(x)\big) \;=\; \mathcal{D}_{\mathrm{KL}}\big(q(x)\big\|p(x)\big) \;=\; 0 \tag{18}$$

and consequently

$$\mathcal{D}_{\mathrm{s}}\big(p(x,y)\big\|q(x,y)\big) \;=\; \mathrm{E}_{p(x)}\Big[\mathcal{D}_{\mathrm{s}}\big(p(y|x)\big\|q(y|x)\big)\Big]. \tag{19}$$

We are interested in the approximation of $p(x, y)$ through the simplified distribution $q(x, y)$, and in particular of $p(y)$ by $q(y)$. We know, again via the chain rule, that

$$\begin{aligned}
\mathcal{D}_{\mathrm{s}}\big(p(y)\big\|q(y)\big) &\overset{(17)}{=} \mathcal{D}_{\mathrm{s}}\big(p(x,y)\big\|q(x,y)\big) - \mathrm{E}_{p(x)}\Big[\mathcal{D}_{\mathrm{s}}\big(p(x|y)\big\|q(x|y)\big)\Big] && (20)\\
&\leq \mathcal{D}_{\mathrm{s}}\big(p(x,y)\big\|q(x,y)\big) && (21)\\
&\overset{(19)}{=} \mathrm{E}_{p(x)}\Big[\mathcal{D}_{\mathrm{s}}\big(p(y|x)\big\|q(y|x)\big)\Big] && (22)\\
&\leq \sum_i \pi_i d_i && (23)\\
&\leq \max_i d_i \;=:\; \delta. && (24)
\end{aligned}$$

So, by limiting the divergences of conditionals $p(y|x)$ and $q(y|x)$ within each single bin such that these remain $\leq \delta$ (24), we can now also bound the divergence of exact and approximate marginals $p(y)$ and $q(y)$ (20).

4.2 The proposed approach

Given the bin-wise divergences, we can now bound the divergence of exact and discretized marginals. The obvious question now is whether and how one can invert the argument and construct a grid approximation matching a pre-specified maximum divergence δ. For a given (reference) point x in the mixing distribution's domain, we can find a corresponding neighbourhood within which the divergence remains below δ. Once we have defined a single bin this way, we can also generate an exhaustive covering of the whole parameter space through such bins. We abbreviate this method as the DIRECT *(Divergence Restricting Conditional Tesselation) approach*, as it aims at a covering of the conditional's parameter space while bounding the divergence.

In some cases it is not possible to have a finite number of bins associated with finite bin-wise divergences. A "trick", if necessary, then is to simply ignore some fraction of parameter space (of the mixing distribution's domain) that is associated with a pre-set, arbitrarily small probability ϵ and do the binning on the remaining share of parameter space. Problems with unbounded divergences, or infinite numbers of necessary bins, commonly occur towards one or both of the parameter space's margins. Neglecting a certain fraction of parameter space that is associated with an (arbitrarily) small probability ϵ will usually

not pose a significant practical problem, as it will only add another bit to the error budget that needs to be considered in (almost) any numerical computation anyway.

4.3 The sequential DIRECT algorithm

We will in the following construct a binning so that the resulting discrete approximation of the exact marginal does not differ, in terms of symmetrized divergence, and with that of both directed divergences, from the exact ("continuous") marginal by more than a pre-specified amount. The number (k) of components and the placement of reference points will be determined automatically in the process. The idea is to sequentially divide the mixing distribution's domain into bins, while firstly ensuring that the divergences within bins are bounded, and secondly, if necessary, ignoring the mixing distribution's extreme left and/or right tails. In order to proceed, in the following we will assume that the divergence between any pair of points (x_1, x_2) in parameter space is Lipschitz continuous, at least within a range $[\tilde{x}_1, \tilde{x}_k]$ with $\mathrm{P}\big(X \notin [\tilde{x}_1, \tilde{x}_k]\big) \leq \epsilon$. This will ensure that the algorithm will work, although violations do not necessarily prevent a solution; even continuity is not strictly necessary. A possible implementation of the DIRECT approach is defined in Tab. 1.

Reference points $\tilde{x}_i$ and corresponding weights π_i now allow to define an approximation q as in (14). It is actually not necessary to also keep track of the exact bin margins $x_{(i)}$ once the bin weights π_i are determined. The maximum divergence of conditionals, and with that of the marginals, will now be $= \delta$, possibly up to a bit of probability ($\leq \epsilon$) beyond the first and/or last bins.

The essence here is to ensure condition (24) to be met. Possible boundary or singularity problems are circumvented by ignoring negligible bits of parameter space via specification of ϵ. Lipschitz continuity of the divergence will ensure that the relevant range may be covered using a finite number of bins. Note that the actual form of the latent (mixing) distribution is only used to determine the relevant range in parameter space, while the actual binning is otherwise independent. A number of variations of the DIRECT algorithm are conceivable; for example, it may or may not be sensible, or possible, to either have a reference point or a bin margin at the parameter space's boundary. Also, the relationship between x and $p(y|x)$ may not necessarily be monotonic, in which case it may be possible

Table 1: The sequential DIRECT algorithm (see Sec. 4.3).

1. Specify a maximum KL-divergence $\delta > 0$, some small probability $0 \leq \epsilon \ll 1$, and a starting reference point $\tilde{x}_1$. Sensible values for $\tilde{x}_1$ may for example be the minimum possible value, the $\frac{\epsilon}{2}$-quantile, or any value with $\mathrm{P}(X \leq \tilde{x}_1) < \epsilon$. Define $\epsilon_1 := \mathrm{P}(X \leq \tilde{x}_1) \geq 0$. Set $i = 1$.
2. Set $x^\star = \tilde{x}_1$. Obviously, $\mathcal{D}_\mathrm{s}\big(p(y|\tilde{x}_1)\big\|p(y|x^\star)\big) = 0$. Now increase $x^\star$ as far as possible while ensuring that $\mathcal{D}_\mathrm{s}\big(p(y|\tilde{x}_1)\big\|p(y|x^\star)\big) \leq \delta$. Use this point as the first bin margin: $x_{(1)} = x^\star$. Compute $\pi_1 = \mathrm{P}(x < x_{(1)})$. Set $i = i + 1$.
3. Increase $x^\star$ until $\mathcal{D}_\mathrm{s}\big(p(y|x_{(i-1)})\big\|p(y|x^\star)\big) = \delta$. Use this point as the next reference point: $\tilde{x}_i = x^\star$.
4. Increase $x^\star$ again until $\mathcal{D}_\mathrm{s}\big(p(y|\tilde{x}_i)\big\|p(y|x^\star)\big) = \delta$. Use this point as the next bin margin: $x_{(i)} = x^\star$.
5. Compute the bin weight $\pi_i = \mathrm{P}(x_{(i-1)} < X \leq x_{(i)})$.
6. If $\mathrm{P}(X > x_{(i)}) > (\epsilon - \epsilon_1)$, set $i = i + 1$ and proceed at step 3. Otherwise stop.

to devise more efficient non-sequential binning strategies.

5 Examples

5.1 Student-t distribution

A prominent example of a mixture distribution is the Student-t distribution. It arises as a continuous mixture of normal distributions with zero mean and scale $\sigma = \sqrt{\frac{\nu}{s}}$, where s is a draw from a χ^2 distribution with ν degrees of freedom (Johnson et al. 1994, Ch. 28). We can approximate the marginal Student-t distribution by a mixture of normal distributions, conditioning on a finite set of grid points in s, and compare against the true marginal which in this case we know to be a Student-t distribution.

Suppose we are interested in the case of $\nu = 5$ degrees of freedom. We set the tuning parameters to $\delta := 0.01$ and $\epsilon := 0.001$ and we use the χ^2_5 distribution's $\frac{\epsilon}{2}$-quantile as the

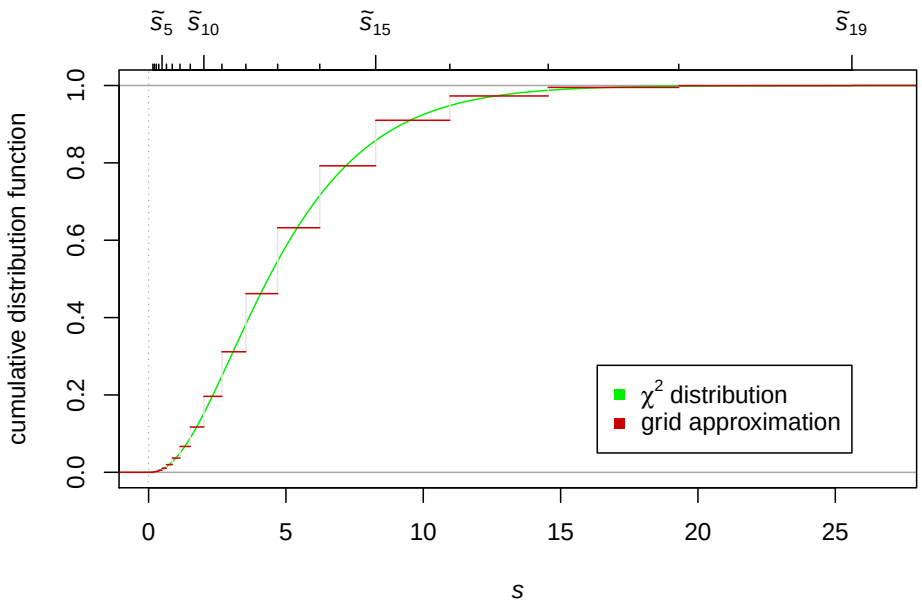

Figure 1: The underlying χ^2 mixing distribution (of the latent variable s) and the grid approximation that is effectively used instead in the Student-t example example (Sec. 5.1). The extra tick marks at the top indicate the 19 grid points used.

starting reference point ($\tilde{s}_1 := 0.158$). Applying the sequential DIRECT algorithm from Sec. 4.3 (utilizing expression (5)) results in a set of 19 reference points $\tilde{s}_i$. As a result from the implied differences in the corresponding conditionally normal distributions, the 19 reference points are very unequally spaced, with many points concentrated near zero and a coarser spacing at large values (see Fig. 1).

Fig. 2 illustrates the construction of the binning by showing the 13th bin and its two neighbouring bins with bin margins $s_{(i)}$ and reference points $\tilde{s}_i$. One can see that by construction within each bin the divergence relative to the corresponding reference point, $\mathcal{D}_{\mathrm{s}}\big(p(x|\sigma=\sqrt{\nu/s})\big\|p(x|\sigma=\sqrt{\nu/\tilde{s}_i})\big)$, remains below δ.

The 19-component normal mixture approximation is compared to the true marginal distribution in Fig. 3. The two densities are barely distinguishable, and their ratio is very close to unity; it only diverges towards the distributions' extreme tails. The numerically computed actual divergence in this case amounts to $\mathcal{D}_{\mathrm{s}}\big(p(x)\big\|q(x)\big) \approx 3.5\times10^{-5}$.

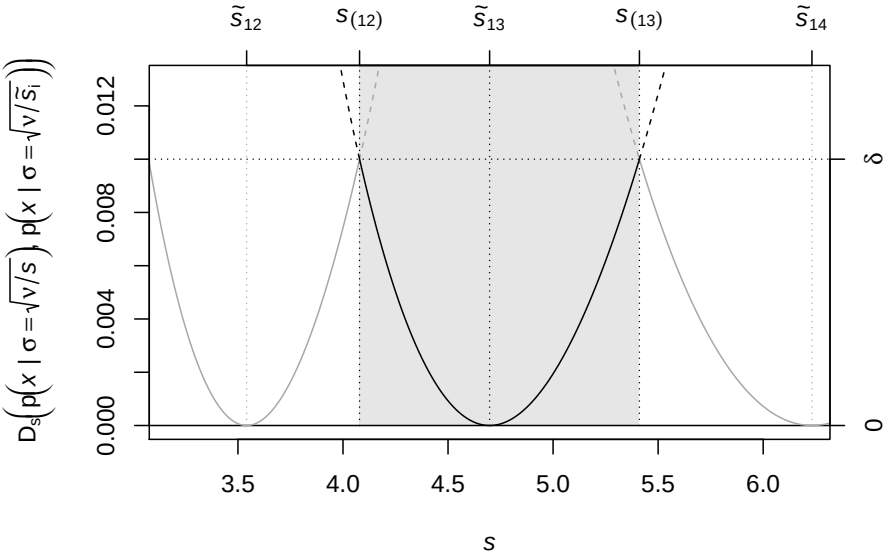

Figure 2: Illustration of how the binning is set up (Student-t example, Sec. 5.1). Bin margins $s_{(i)}$ and reference points $\tilde{s}_i$ are arranged such that within each bin the divergence relative to the corresponding reference point does not exceed the pre-set threshold δ.

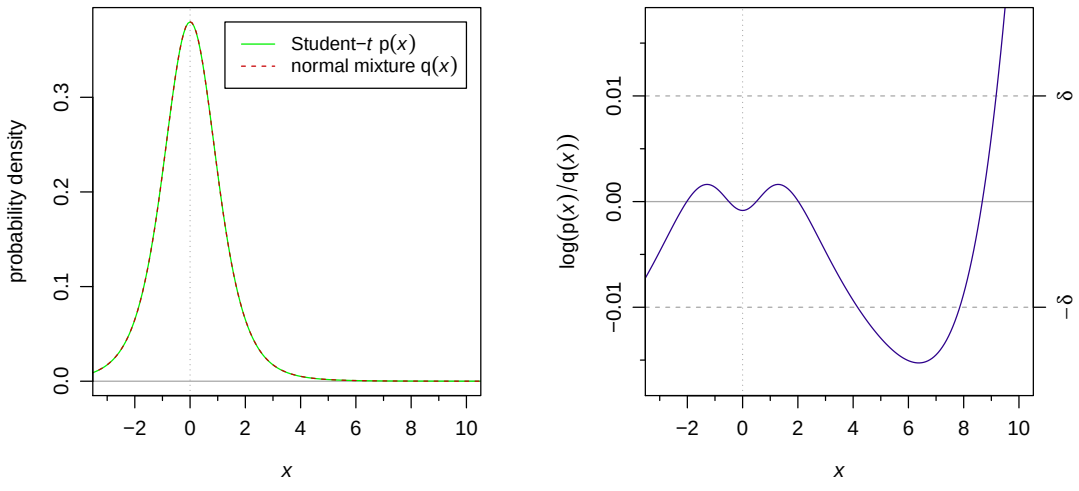

Figure 3: Comparison of the true mixture distribution, the Student-t distribution, to the grid approximation. The left panel shows the two probability density functions on top of each other; the two are essentially undiscernible at this scale. The right panel shows the logarithmic ratio of the densities as a function of x.

5.2 Convolution of two distributions

In the following we present the example of computing the convolution of two distributions. Suppose we have two random variables, X and Y, with densities $p_X(x)$ and $p_Y(y)$. We are interested in their sum $Z = X + Y$, and its density $p_Z(z)$. Here we take X and Y to follow skew-normal and logistic distributions, respectively, so that the solution is not trivial. We can turn the problem into that of a mixture distribution and subsequently apply the above algorithm by first considering the joint distribution of X and Z. Note that $\mathrm{P}(Z = z \,|\, X = x) \;=\; \mathrm{P}(Y = z - x)$, so the conditional distribution of $Z|X$ here is simply a "shifted" version of the (known) distribution P_Y. With that, we can rewrite the target density p_Z as a marginal density in terms of the (known) marginal $p_X(x)$ and the (known) conditional $p_Z(z|x) = p_Y(z - x)$:

$$p_Z(z) \;=\; \int p_Y(z - x)\, p_X(x)\, \mathrm{d}x. \tag{25}$$

This way it is obvious that convolution of two random variables may again be seen as a special case of a mixture distribution where the conditional $\mathrm{P}(Z|X)$ is mixed via the latent distribition $\mathrm{P}(X)$. Due to symmetry of the problem, the roles of X and Y may also be reversed.

In the following suppose that $p_X(x) > 0$ and $p_Y(y) > 0$ for all $x, y \in \mathbb{R}$, i.e., the domain of both X and Y is the whole real line. When applying the DIRECT algorithm to set up an approximation, it is important to note that the divergence $\mathcal{D}_\mathrm{s}\big(p_Z(z|x_1)\big\|p_Z(z|x_2)\big)$ required in steps 2–4 of the algorithm (Sec. 4.3) only depends on the (absolute) difference $|x_2 - x_1|$, since the conditional distributions $p_Z(z|\,\cdot\,)$ here only differ by a shift in location. This implies that the bin width $\tilde{x}_i - \tilde{x}_{i-1}$ is constant across all bins, and hence only needs to be determined once. This simplifies the grid construction to a few steps:

1. determine the bin half-width Δ_x such that $\mathcal{D}_\mathrm{s}\big(p_Y(y)\big\|p_Y(y - \Delta_x)\big) = \delta$.

2. determine minimum and maximum X values $\tilde{x}_1$ and $\tilde{x}_k$ e.g. as the $\frac{\epsilon}{2}$ and $1{-}\frac{\epsilon}{2}$ quantiles of p_X.

3. determine the remaining reference points $\tilde{x}_2$ to $\tilde{x}_{k-1}$ as well as their total number k by filling the interval with reference points that are at most $(\tilde{x}_i - \tilde{x}_{i-1}) \leq \Delta_x$ apart.

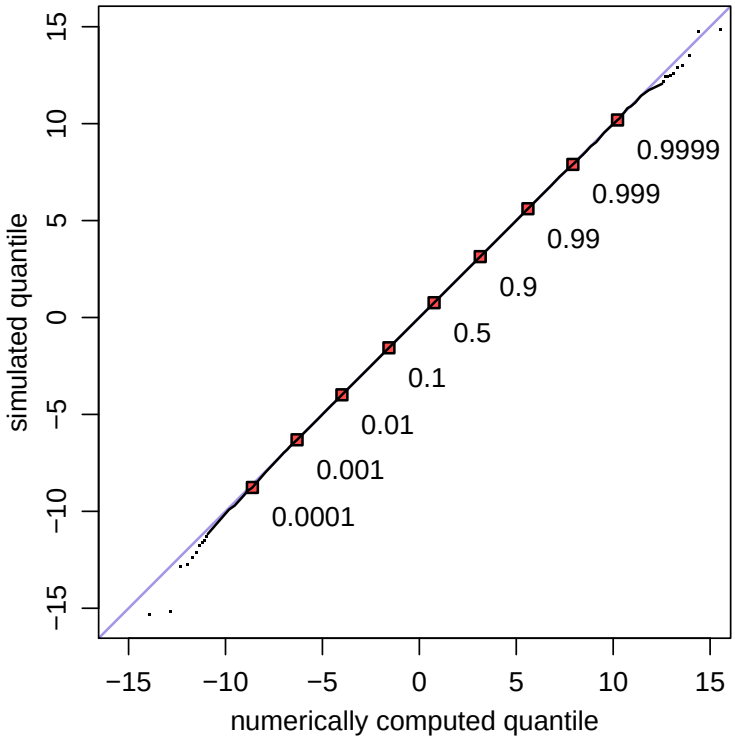

Figure 4: Q-Q-plot illustrating the accuracy of the convolution of a skew-normal and a logistic distribution by comparing quantiles computed numerically, using the DIRECT algorithm, and simulated quantiles.

A general implementation of the procedure in R is shown in the online supplement. Divergences here are computed numerically, without needing to have the corresponding formulas available in analytic form.

Consider the example of the sum of two random variables, one following a skew-normal distribution with shape parameter $\alpha = 4$ (Azzalini 2014, 2015), and one following a logistic distribution. Application of the DIRECT algorithm (using $\delta = 0.01$ and $\epsilon = 0.001$) results in a 13-component mixture of logistic distributions to approximate the convolution. Draws from the two summands' distributions may easily be simulated, so it is straightforward to also generate samples of their sum's distribution. Figure 4 illustrates the fit of the numerical approximation to 1 000 000 simulated samples via a quantile-quantile plot (Q-Q plot). Here the 10 smallest and largest samples are shown as individual dots, other quantiles are connected by a line, and selected quantiles are highlighted. Note that while

the design parameter ϵ was set to 0.001, the simulated and computed quantiles appear to match well even beyond tail probabilities of 0.001. The R code to reproduce these simulations is also provided in the online supplement. All computations here were carried out using R (R Core Team 2015).

6 Conclusions

The DIRECT approach introduced in this paper allows to generate finite mixtures as approximations to mixture distributions with a large or infinite number of mixture components. A formulation in terms of a finite mixture distribution then makes density function, cumulative distribution function, etc. easily accessible. The mismatch incurred by resorting to the approximation is efficiently controlled via two tuning parameters (δ and ϵ). The described algorithm allows for easy implementation in a completely automated fashion, as is also demonstrated in the examples. The setup relies on the computation of (symmetrized) divergences of (conditional) distributions; ideally these are available analytically, but numerical computation is also not a problem.

Variations of the DIRECT algorithm are conceivable. The bound derived in Sec. 4 may be met in many different ways; the described one is only a simple, general solution. For example, it may be possible, and possibly more efficient, to aim at the condition in (23) rather than (24) in order to bound the divergence. While for simplicity we concentrated on symmetrized divergences here, it may also make sense to directly aim for directed Kullback-Leibler divergences instead.

A generalization to higher dimensions of the latent mixing distribution should in general also be possible. Since the problem of covering of higher-dimensional spaces is considerably trickier, it may eventually be easiest to resort to random coverings here (Messenger et al. 2009, Röver 2010).

The algorithm was originally developed and eventually applied in the context of the `bayesmeta` R package (Röver 2015). In this meta-analysis application, one is faced with the common problem of inferring two parameters (τ and μ) via their posterior probability distribution. From their joint distribution ($p(\mu, \tau)$) one of the marginals, $p(\tau)$, may be derived analytically, while the conditionals $p(\mu|\tau)$ are normal. Primary interest usually lies

in μ, and application of the DIRECT algorithm facilitates quick and accurate computation of the marginal $p(\mu)$ without having to use, for example, Markov chain Monte Carlo (MCMC) methods (Friede et al. 2016).

References

Azzalini, A. (2014), *The skew-normal and related families*, Institute of Mathematical Statistics Monographs, Cambridge University Press, Cambridge.

Azzalini, A. (2015), 'The skew-normal probability distribution'. URL: `http://azzalini.stat.unipd.it/SN`.

Bernardo, J. M. & Smith, A. F. M. (1994), *Bayesian theory*, Wiley, Chichester, UK.

Cover, T. M. & Thomas, J. A. (1991), *Elements of information theory*, Wiley & Sons, New York.

Friede, T., Röver, C., Wandel, S. & Neuenschwander, B. (2016), 'Meta-analysis of few small studies in orphan diseases', *Research Synthesis Methods (in press)* . URL: `http://arxiv.org/abs/1601.06533`.

Johnson, N. L., Kotz, S. & Balakrishnan, N. (1994), *Continuous univariate distributions*, 2nd edn, Wiley, New York.

Kullback, S. (1959), *Information theory and statistics*, John Wiley & Sons, New York.

Kullback, S. & Leibler, R. A. (1951), 'On information and sufficiency', *The Annals of Mathematical Statistics* **22**(1), 79–86.

Lindsay, B. G. (1995), *Mixture models: theory, geometry and applications*, Vol. 5 of *NSF-CBMS Regional Conference Series in Probability and Statistics*, Institute of Mathematical Statistics, Hayward, CA, USA.

Messenger, C., Prix, R. & Papa, M. A. (2009), 'Random template banks and relaxed lattice coverings', *Physical Review D* **79**(10), 104017.

O'Hagan, A. (1994), *Bayesian inference*, Vol. 2B of *Kendall's advanced theory of statistics*, Wiley, NY, USA.

R Core Team (2015), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria. URL: `https://www.r-project.org/`.

Röver, C. (2010), 'Random template placement and prior information', *Journal of Physics: Conference Series* **228**(1), 012008.

Röver, C. (2015), 'bayesmeta: Bayesian random-effects meta analysis'. R package. URL: `http://cran.r-project.org/package=bayesmeta`.

Seidel, W. E. (2010), Mixture models, *in* M. Lovric, ed., 'International Encyclopedia of Statistical Science', Springer, Heidelberg, pp. 827–829.