

Hybrid Inference

Frank Dellaert & Varun Agrawal

January 2023

1 Hybrid Conditionals

Here we develop a hybrid conditional density, on continuous variables (typically a measurement x), given a mix of continuous variables y and discrete variables m . We start by reviewing a Gaussian conditional density and its invariants (relationship between density, error, and normalization constant), and then work out what needs to happen for a hybrid version.

GaussianConditional

A *GaussianConditional* is a properly normalized, multivariate Gaussian conditional density:

$$P(x|y) = \frac{1}{\sqrt{|2\pi\Sigma|}} \exp \left\{ -\frac{1}{2} \|Rx + Sy - d\|_{\Sigma}^2 \right\}$$

where R is square and upper-triangular. For every *GaussianConditional*, we have the following **invariant**,

$$\log P(x|y) = K_{gc} - E_{gc}(x, y), \quad (1)$$

with the **log-normalization constant** K_{gc} equal to

$$K_{gc} = \log \frac{1}{\sqrt{|2\pi\Sigma|}} \quad (2)$$

and the **error** $E_{gc}(x, y)$ equal to the negative log-density, up to a constant:

$$E_{gc}(x, y) = \frac{1}{2} \|Rx + Sy - d\|_{\Sigma}^2. \quad (3)$$

GaussianMixture

A *GaussianMixture* (maybe to be renamed to *GaussianMixtureComponent*) just indexes into a number of *GaussianConditional* instances, that are each properly normalized:

$$P(x|y, m) = P_m(x|y).$$

We store one *GaussianConditional* $P_m(x|y)$ for every possible assignment m to a set of discrete variables. As *GaussianMixture* is a *Conditional*, it needs to satisfy the a similar invariant to (1):

$$\log P(x|y, m) = K_{gm} - E_{gm}(x, y, m). \quad (4)$$

If we take the log of $P(x|y, m)$ we get

$$\log P(x|y, m) = \log P_m(x|y) = K_{gcm} - E_{gcm}(x, y). \quad (5)$$

Equating (4) and (5) we see that this can be achieved by defining the error $E_{gm}(x, y, m)$ as

$$E_{gm}(x, y, m) = E_{gcm}(x, y) + K_{gm} - K_{gcm} \quad (6)$$

where choose $K_{gm} = \max K_{gcm}$, as then the error will always be positive.

2 Hybrid Factors

In GTSAM, we typically condition on known measurements, and factors encode the resulting negative log-likelihood of the unknown variables y given the measurements x . We review how a Gaussian conditional density is converted into a Gaussian factor, and then develop a hybrid version satisfying the correct invariants as well.

JacobianFactor

A *JacobianFactor* typically results from a *GaussianConditional* by having known values $\bar{x}$ for the “measurement” x :

$$L(y) \propto P(\bar{x}|y) \quad (7)$$

In GTSAM factors represent the negative log-likelihood $E_{jf}(y)$ and hence we have

$$E_{jf}(y) = -\log L(y) = C - \log P(\bar{x}|y),$$

with C the log of the proportionality constant in (7). Substituting in $\log P(\bar{x}|y)$ from the invariant (1) we obtain

$$E_{jf}(y) = C - K_{gc} + E_{gc}(\bar{x}, y).$$

The *likelihood* function in *GaussianConditional* chooses $C = K_{gc}$, and the *JacobianFactor* does not store any constant; it just implements:

$$E_{jf}(y) = E_{gc}(\bar{x}, y) = \frac{1}{2} \|R\bar{x} + Sy - d\|_{\Sigma}^2 = \frac{1}{2} \|Ay - b\|_{\Sigma}^2$$

with $A = S$ and $b = d - R\bar{x}$.

GaussianMixtureFactor

Analogously, a *GaussianMixtureFactor* typically results from a *GaussianMixture* by having known values $\bar{x}$ for the “measurement” x :

$$L(y, m) \propto P(\bar{x}|y, m).$$

We will similarly implement the negative log-likelihood $E_{mf}(y, m)$:

$$E_{mf}(y, m) = -\log L(y, m) = C - \log P(\bar{x}|y, m).$$

Since we know the log-density from the invariant (4), we obtain

$$\log P(\bar{x}|y, m) = K_{gm} - E_{gm}(\bar{x}, y, m),$$

and hence

$$E_{mf}(y, m) = C + E_{gm}(\bar{x}, y, m) - K_{gm}.$$

Substituting in (6) we finally have an expression where K_{gm} canceled out, but we have a dependence on the individual component constants K_{gcm} :

$$E_{mf}(y, m) = C + E_{gcm}(\bar{x}, y) - K_{gcm}.$$

Unfortunately, we can no longer choose C independently from m to make the constant disappear. There are two possibilities:

1. Implement likelihood to yield both a hybrid factor *and* a discrete factor.
2. Hide the constant inside the collection of *JacobianFactor* instances, which is the possibility we implement.

In either case, we implement the mixture factor $E_{mf}(y, m)$ as a set of *JacobianFactor* instances $E_{mf}(y, m)$, indexed by the discrete assignment m :

$$E_{mf}(y, m) = E_{jfm}(y) = \frac{1}{2} \|A_m y - b_m\|_{\Sigma_{mfm}}^2.$$

In GTSAM, we define A_m and b_m strategically to make the *JacobianFactor* compute the constant, as well:

$$\frac{1}{2} \|A_m y - b_m\|_{\Sigma_{mfm}}^2 = C + E_{gcm}(\bar{x}, y) - K_{gcm}.$$

Substituting in the definition (3) for $E_{gcm}(\bar{x}, y)$ we need

$$\frac{1}{2} \|A_m y - b_m\|_{\Sigma_{mfm}}^2 = C + \frac{1}{2} \|R_m \bar{x} + S_m y - d_m\|_{\Sigma_m}^2 - K_{gcm}$$

which can be achieved by setting

$$A_m = \begin{bmatrix} S_m \\ 0 \end{bmatrix}, \quad b_m = \begin{bmatrix} d_m - R_m \bar{x} \\ c_m \end{bmatrix}, \quad \Sigma_{mfm} = \begin{bmatrix} \Sigma_m & \\ & 1 \end{bmatrix}$$

and setting the mode-dependent scalar c_m such that $c_m^2 = C - K_{gcm}$. This can be achieved by $C = \max K_{gcm} = K_{gm}$ and $c_m = \sqrt{2(C - K_{gcm})}$. Note that in case that all constants K_{gcm} are equal, we can just use $C = K_{gm}$ and

$$A_m = S_m, \quad b_m = d_m - R_m \bar{x}, \quad \Sigma_{mfm} = \Sigma_m$$

as before.

In summary, we have

$$E_{mf}(y, m) = \frac{1}{2} \|A_m y - b_m\|_{\Sigma_{mfm}}^2 = E_{gcm}(\bar{x}, y) + K_{gm} - K_{gcm}. \quad (8)$$

which is identical to the GaussianMixture error (6).