

Text S2: Forward and Backward Inference in Spatial Cognition

Will D. Penny, Peter Zeidman and Neil Burgess

Forward and Backward Inference: General Formulation

In this section, for brevity we drop the dependence on \mathbf{u}_n and E in our notation, but this is nevertheless implied. Probabilistic path integration can be implemented by computing

$$p(\mathbf{x}_n|\mathbf{Y}_{n-1}) = \int p(\mathbf{x}_n|\mathbf{x}_{n-1})p(\mathbf{x}_{n-1}|\mathbf{Y}_{n-1})d\mathbf{x}_{n-1} \quad (1)$$

To optimally combine path integration with sensory input we compute

$$p(\mathbf{x}_n|\mathbf{Y}_n) = \frac{p(\mathbf{y}_n|\mathbf{x}_n)p(\mathbf{x}_n|\mathbf{Y}_{n-1})}{p(\mathbf{y}_n|\mathbf{Y}_{n-1})} \quad (2)$$

which is Bayes rule with prior $p(\mathbf{x}_n|\mathbf{Y}_{n-1})$ from probabilistic path integration, likelihood $p(\mathbf{y}_n|\mathbf{x}_n)$ from sensory input, and the normalisation term is given by

$$p(\mathbf{y}_n|\mathbf{Y}_{n-1}) = \int p(\mathbf{y}_n|\mathbf{x}_n)p(\mathbf{x}_n|\mathbf{Y}_{n-1})d\mathbf{x}_n \quad (3)$$

which is also known as the predictive density. Equations 1 and 2 together update $p(\mathbf{x}_{n-1}|\mathbf{Y}_{n-1})$ to $p(\mathbf{x}_n|\mathbf{Y}_n)$. This procedure is referred to as forward inference and is also known as filtering [38]. The goals of forward inference are to (i) produce the ‘filtering density’ $p(\mathbf{x}_n|\mathbf{Y}_n)$ and (ii) compute the likelihood $p(\mathbf{Y}_n)$. The likelihood of a data sequence can be computed using the product of predictive densities

$$p(\mathbf{Y}_n) = p(\mathbf{y}_1) \prod_{n=2}^N p(\mathbf{y}_n|\mathbf{Y}_{n-1}) \quad (4)$$

If forward inference runs up until time index N then it is possible to run a backward inference procedure which produces the posterior density $p(\mathbf{x}_n|\mathbf{Y}_N)$. This is the probability of the state \mathbf{x}_n given observations up to time index N . This includes observations after time index n and so corresponds to a retrospective update. This can be computed recursively using

$$p(\mathbf{x}_n|\mathbf{Y}_N) = \int p(\mathbf{x}_n|\mathbf{x}_{n+1}, \mathbf{Y}_n)p(\mathbf{x}_{n+1}|\mathbf{Y}_N)d\mathbf{x}_{n+1} \quad (5)$$

The first term in the integral can be thought of as a reverse flow term and is computed using Bayes rule

$$p(\mathbf{x}_n|\mathbf{x}_{n+1}, \mathbf{Y}_n) = \frac{p(\mathbf{x}_{n+1}|\mathbf{x}_n, \mathbf{Y}_n)p(\mathbf{x}_n|\mathbf{Y}_n)}{\int p(\mathbf{x}_{n+1}|\mathbf{x}_n, \mathbf{Y}_n)p(\mathbf{x}_n|\mathbf{Y}_n)d\mathbf{x}_n} \quad (6)$$

Equations 5 and 6 implement an online algorithm for backward inference in the sense that there is no need to store the time series of observations, \mathbf{Y}_n . In signal processing these updates are known as the gamma recursions [38] and it is these which we envisage could be implemented in the brain.

There is an alternative procedure for backward inference known as the beta recursions. We include this for completeness below, as there is an interesting connection to optimal control theory which we will also refer to in the discussion. However, the beta recursions require storage of the observations and so are not an online algorithm.

Alpha-Beta Decomposition

The distribution of hidden states \mathbf{x}_n given an observed trajectory \mathbf{Y}_1^N can be expressed using Bayes rule

$$\begin{aligned}
 p(\mathbf{x}_n | \mathbf{Y}_1^N) &= \frac{p(\mathbf{Y}_1^N | \mathbf{x}_n) p(\mathbf{x}_n)}{p(\mathbf{Y}_1^N)} \\
 &= \frac{p(\mathbf{Y}_1^{n-1} | \mathbf{x}_n) p(\mathbf{Y}_n^N | \mathbf{x}_n) p(\mathbf{x}_n)}{p(\mathbf{Y}_1^N)} \\
 &= \frac{p(\mathbf{Y}_1^{n-1}, \mathbf{x}_n) p(\mathbf{Y}_n^N | \mathbf{x}_n)}{p(\mathbf{Y}_1^N)} \\
 &= \frac{\alpha(\mathbf{x}_n) \beta(\mathbf{x}_n)}{p(\mathbf{Y}_1^N)}
 \end{aligned} \tag{7}$$

where $\alpha(\mathbf{x}_n) = p(\mathbf{Y}_1^{n-1}, \mathbf{x}_n)$ and $\beta(\mathbf{x}_n) = p(\mathbf{Y}_n^N | \mathbf{x}_n)$. In offline applications, where an agent may store the time series of observations, the alpha and beta recursions can therefore be run in parallel. The alpha-beta algorithm is also therefore known as the parallel update [118]. The alpha and beta updates can then be combined to estimate the posterior probability of the hidden states given both past and future observations, as described above.

Beta recursions

Following on from the previous section we have

$$\beta(\mathbf{x}_n) = p(\mathbf{Y}_n^N | \mathbf{x}_n) \tag{8}$$

This can be computed recursively using

$$\beta(\mathbf{x}_n) = p(\mathbf{y}_n | \mathbf{x}_n) \int p(\mathbf{x}_{n+1} | \mathbf{x}_n) \beta(\mathbf{x}_{n+1}) d\mathbf{x}_{n+1} \tag{9}$$

The beta recursions operate on future data and are the time-reversed equivalent of the alpha recursions. Under weak assumptions this is equivalent to Bellman's equation for recursively computing the optimal value of a state [119].