
The Hierarchical Adaptive Forgetting Variational Filter

Vincent Moens¹

Abstract

A common problem in Machine Learning and statistics consists in detecting whether the current sample in a stream of data belongs to the same distribution as previous ones, is an isolated outlier or inaugurates a new distribution of data. We present a hierarchical Bayesian algorithm that aims at learning a time-specific approximate posterior distribution of the parameters describing the distribution of the data observed. We derive the update equations of the variational parameters of the approximate posterior at each time step for models from the exponential family, and show that these updates find interesting correspondents in Reinforcement Learning (RL). In this perspective, our model can be seen as a hierarchical RL algorithm that learns a posterior distribution according to a certain stability confidence that is, in turn, learned according to its own stability confidence. Finally, we show some applications of our generic model, first in a RL context, next with an adaptive Bayesian Autoregressive model, and finally in the context of Stochastic Gradient Descent optimization.

1. Introduction

Learning in a changing environment is a difficult albeit ubiquitous task. One key issue for learning in such context is to discriminate between isolated, unexpected events and a prolonged contingency change. This discrimination is challenging with conventional techniques because they rely on prior assumptions about environment stability. When assuming fluctuating context, past experience will be forgotten immediately when an unexpected event occurs, but if that event was just noise, this erroneous forgetting might be very costly. In less variable contexts, model parameters will tend to change more gradually, thus sometimes missing

fluctuations when they happen faster than expected. Most models cover one of the two possibilities, and either gradually adapt their predictions to the new contingency or do it abruptly, but not both.

One classical solution to the problem of change detection is to compare the likelihood of the current observation given the previous posterior distribution with a default probability distribution (Kulhavy & Karny, 1984), representing an initial, naive state of the learner. Usually, the mixing coefficient (or forgetting factor) that is used to weight these two hypotheses is adapted to the current data in order to detect and account for the possible contingency change. This mixing coefficient can be implemented in a linear or exponential manner (Kulhavý & Kraus, 1996). We will focus here on the exponential case.

In the past decade, several Bayesian solutions to this problem based on the aforementioned strategy have been proposed (Smidl, 2004; Smidl & Gustafsson, 2012; Azizi & Quinn, 2015). However, they usually suffer from several drawbacks: many of them put a restrictive prior on the mixing coefficient, e.g. (Smidl, 2004; Masegosa et al., 2017), and cannot account for the fact that an unexpected event is unlikely to be caused by a contingency change if the environment has been stable for a long time.

We propose the Hierarchical Adaptive Forgetting Variational Filter (HAFVF). The core idea of the model is that the mixing coefficient can be learned as a latent variable with its own mixing coefficient. It is inspired by the observation that animals tend to decrease their flexibility (i.e. their capacity to adapt to a new contingency) when they are trained in a stable environment and that this flexibility is inversely correlated with the training length (Dickinson, 1985). We suggest that this strategy may be beneficial in many environments, where the stability of the system identified by a learner is a variable that can be learned as an independent variable with a certain confidence: in certain environments, contingency changes are inherently more less than in others. Although this assumption may not hold in every case, we show that it helps the algorithm to stabilize and discriminate contingency changes from accidents.

Accordingly, we frame our algorithm in a RL framework (Dearden & Russell, 1998). We explore how the forward learning algorithm can be extended to the forward-backward

¹COSY, Institute of Neuroscience, Université Catholique de Louvain, Brussels, Belgium. Correspondence to: Vincent Moens <vincent.moens@uclouvain.be>.

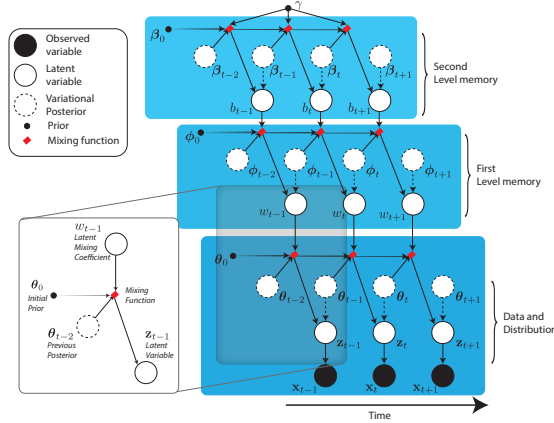


Figure 1. Directed Graph of the HAFVF. Our contribution with respect to previous works is to consider that the forgetting factor w has a mixture prior with forgetting factor b which acts on w as w acts on \mathbf{z} .

case. We show three applications of our model: first in the case of a simple RL task, next to fit an autoregressive model and finally for gradient learning in a Stochastic Gradient Descent (SGD) algorithm.

2. Hierarchical Model

Let $\mathbf{x} = \{x_1, x_2, \dots, x_T\}$ be a stream of data distributed according to a set of N distributions $\mathbf{p} = \{p_1(\mathbf{x}_{1:n_1} | \mathbf{z}_1), \dots, p_N(\mathbf{x}_{n_{N-1}+1:n_N} | \mathbf{z}_N)\}$, where the change trials $\mathbf{n} := \{n_1, n_2, \dots, n_N \leq T\} \in \mathbb{Z}^+$ are unknown and unpredictable. We make the following assumptions:

Assumption 1. Let $\{t_1, t_2\} \in T$ and $n_x < t_1 < t_2$, then $p(t_2 \in \mathbf{n}) \leq p(t_1 \in \mathbf{n})$.

Corollary 1. If $m(x_t, p_n)$ is a measure of the relative probability that x_t belongs to p_n wrt p_0 , and if $x_2 = x_1$, then $m(x_2, p_n) \geq m(x_1, p_n)$.

Assumption 1 and Corollary 1 state that the probability of seeing a contingency change decreases with time in a steady environment. This might seem counter-intuitive or even maladaptive in many situations, but it is a key assumption we use to discriminate artifacts from contingency change: after a long sequence, the amount of evidence needed to switch from the current belief to the naive belief is greater than after a short sequence. This assumption will lead us to build a model where, if the learner is very confident in his belief, it will take him more time to forget past observations, because he will need more evidence for a contingency change. Therefore, in this context, the learner aims not only to learn the distribution of the data at hand, but also a measure of the confidence in the volatility of the environment.

Assumption 2. In the set of probability distributions \mathbf{p} , all

elements have the same parametric form that belongs to the exponential family (EF) and have a conjugate prior that is also from the EF:

$$p_n \in \mathcal{P} \implies p_n(x_t | \mathbf{z}_n) = h(x_t) \exp \{ \mathbf{z}_n^T \mathbf{T}(x_t) - A(\mathbf{z}_n) \}$$

$$\text{and } p_n(\mathbf{z}_n) = \exp \{ \mathbf{T}(\mathbf{z}_n)^T \boldsymbol{\theta} - B(\boldsymbol{\theta}) \}.$$

We now focus on the problem of approximating the current posterior distribution of $\mathbf{z}_{n,x}$ given the current and past observations. For clarity, we will make the n subscripts implicit in the following. Let us first focus on the problem of estimating the posterior distribution of \mathbf{z} in the stationary case. After t steps, and given some prior distribution $p(\mathbf{z} | \boldsymbol{\theta}_0)$, the posterior distribution can be formulated recursively as:

$$p(\mathbf{z} | x_t, \mathbf{x}_{<t}) = \begin{cases} \frac{p(x_t | \mathbf{z}) p(\mathbf{z} | \mathbf{x}_{<t})}{p(x_t | \mathbf{x}_{<t})} & \text{for } t > 1 \\ \frac{p(x_t | \mathbf{z}) p(\mathbf{z} | \boldsymbol{\theta}_0)}{p(x_t)} & \text{otherwise.} \end{cases}$$

Given the restriction imposed by Assumption 2, this posterior probability distribution has a closed-form expression and can be estimated efficiently.

We enrich this basic model by first formulating the prior distribution of \mathbf{z} at t as a mixture of the previous posterior distribution and an arbitrary prior:

$$p_t(\mathbf{z} | \mathbf{x}_{<t}; \boldsymbol{\theta}_0, w) = \frac{p_{t-1}(\mathbf{z} | \mathbf{x}_{<t})^w p(\mathbf{z} | \boldsymbol{\theta}_0)^{1-w}}{Z(w, \mathbf{x}_{<t}, \boldsymbol{\theta}_0)}. \quad (1)$$

Following Assumption 2, the conjugate distribution $p_{t-1}(\mathbf{z} | \mathbf{x}_{<t})$ is also from the EF and reads

$$p_{t-1}(\mathbf{z} | \mathbf{x}_{<t}) := \exp \{ \mathbf{z}^T \boldsymbol{\theta}^\xi - \boldsymbol{\theta}^\eta A(\mathbf{z}) - B(\boldsymbol{\theta}) \}$$

where we have expanded $\mathbf{T}(\mathbf{z})$, where $\boldsymbol{\theta} = \{\boldsymbol{\theta}^\xi, \boldsymbol{\theta}^\eta\}$. $\boldsymbol{\theta}^\eta$ is the part of $\boldsymbol{\theta}$ that indicates the effective (prior) number of observations. If p_0 has the same form as $p_{t-1}(\mathbf{z} | \mathbf{x}_{<t})$, then the log-partition function Z can be computed efficiently (Mandt et al., 2014):

$$Z(w, \boldsymbol{\theta}_{t-1}, \boldsymbol{\theta}_0) = \exp \{ -wB(\boldsymbol{\theta}_{t-1}) - (1-w)B(\boldsymbol{\theta}_0) + B(w\boldsymbol{\theta}_{t-1} + (1-w)\boldsymbol{\theta}_0) \}.$$

Note that this result simplifies when combined with the numerator of Equation (1):

$$p(\mathbf{z} | \boldsymbol{\theta}_{t-1}, \boldsymbol{\theta}_0, w) = \exp \{ \mathbf{T}(\mathbf{z})^T \boldsymbol{\vartheta} - B(\boldsymbol{\vartheta}) \} \quad (2)$$

where $\boldsymbol{\vartheta} := w\boldsymbol{\theta}_{t-1} + (1-w)\boldsymbol{\theta}_0$. The latent variable $w \in [0; 1]$ weights the initial prior with the posterior at the previous trial. We incorporate this variable in the set of the latent variables, and, we put a mixture prior on w with a weight b : following this approach, the previous posterior probability of w conditions the current one (similarly to \mathbf{z}),

together with a prior that is blind to the stream of data up to now. Assuming that x , z and w each can be generated by changing distributions, the joint probability now reads:

$$\begin{aligned}
 p(x_t, \mathbf{z}, w, b | \mathbf{x}_{<t}; \boldsymbol{\theta}_0, \boldsymbol{\phi}_0, \boldsymbol{\beta}_0, \gamma) &:= p(x_t | \mathbf{z}) \times \\
 &\frac{p_{t-1}(\mathbf{z} | \mathbf{x}_{<t})^w p(\mathbf{z} | \boldsymbol{\theta}_0)^{1-w}}{Z(w, \mathbf{x}_{<t}, \boldsymbol{\theta}_0)} \times \\
 &\frac{p_{t-1}(w | \mathbf{x}_{<t})^b p(w | \boldsymbol{\phi}_0)^{1-b}}{Z(b, \mathbf{x}_{<t}, \boldsymbol{\phi}_0)} \times \\
 &\frac{p_{t-1}(b | \mathbf{x}_{<t})^\gamma p(b | \boldsymbol{\beta}_0)^{1-\gamma}}{Z(\gamma, \mathbf{x}_{<t}, \boldsymbol{\beta}_0)}
 \end{aligned} \quad (3)$$

where we have assumed that the posterior probability $p(\mathbf{z}, w, b | \mathbf{x})$ factorizes (Mean-Field assumption), and where $\{\boldsymbol{\theta}_0, \boldsymbol{\phi}_0, \boldsymbol{\beta}_0\}$ are the parameters of the naive, initial prior distributions over $\{\mathbf{z}, w, b\}$ respectively. The model presented in Equation (3) is not conjugate anymore, and the posterior probability does not generally have an analytical solution. We therefore introduce a variational posterior to approximate the posterior probability $p(\mathbf{z}, w, b | \mathbf{x})$. In short, Variational Inference (Jaakkola & Jordan, 2000) works by replacing the posterior by a proxy of an arbitrary form and finding the configuration of this approximate posterior that minimizes the Kullback-Liebler divergence between this distribution and the true posterior. This is virtually identical to maximizing the Expected Lower-Bound to the log-model evidence (ELBO).

For simplicity, we use a factorized variational posterior $q_t(\mathbf{z}, w, b) = q_t(\mathbf{z} | \boldsymbol{\theta}_t) q_t(w | \boldsymbol{\phi}_t) q_t(b | \boldsymbol{\beta}_t)$ where each factor has the same form as the prior distribution of its latent variable. Assuming that $q_{t-1}(\cdot) \approx p_{t-1}(\cdot)$ Equation (3) conveniently simplifies to:

$$\begin{aligned}
 p(x_t, \mathbf{z}, w, b | \mathbf{x}_{<t}; \boldsymbol{\theta}_0, \boldsymbol{\phi}_0, \boldsymbol{\beta}_0, \gamma) &\approx p(x_t | \mathbf{z}) \times \\
 &p(\mathbf{z} | w(\boldsymbol{\theta}_{t-1} - \boldsymbol{\theta}_0) + \boldsymbol{\theta}_0) \times \\
 &p(w | b(\boldsymbol{\phi}_{t-1} - \boldsymbol{\phi}_0) + \boldsymbol{\phi}_0) \times \\
 &p(b | \gamma(\boldsymbol{\beta}_{t-1} - \boldsymbol{\beta}_0) + \boldsymbol{\beta}_0).
 \end{aligned} \quad (4)$$

This model is shown in Figure 1. In what follows, we will restrict our analysis to the case where w and b are Beta distributed, meaning that the approximate posterior we will optimize for these two variables will also be a Beta distribution.

2.1. Update equations

Notation We first define the following notation: $d_\theta := \boldsymbol{\theta}_{t-1} - \boldsymbol{\theta}_0$ is the difference between the previous approximate posterior and the initial prior. We use $\boldsymbol{\vartheta} := w \boldsymbol{\theta}_{t-1} + (1-w) \boldsymbol{\theta}_0$ as the weighted prior parameters, and $\widehat{\boldsymbol{\vartheta}}$ as the expectation of $\boldsymbol{\vartheta}$ under $q(w)$. Similarly, $\boldsymbol{\varphi}$ and $\widehat{\boldsymbol{\varphi}}$ are the weighted prior over w and its expectation under q , respectively. Also, we will often abbreviate the summary

$$\text{statistics of } \mathbf{z} \text{ as } \mathbf{T}(\mathbf{z}) := \begin{bmatrix} \mathbf{z} \\ -A(\mathbf{z}) \end{bmatrix}.$$

We now focus on the problem of finding the approximate posterior configuration that maximizes the ELBO. Various techniques have been developed to solve this problem: whereas Stochastic Gradient Variational Bayes (Kingma & Welling, 2013) and Stochastic Variational Bayes (Hoffman et al., 2012) work well for large datasets, more traditional conjugate (Winn et al., 2005) or non-conjugate (Knowles & Minka, 2011) Variational Message Passing (VMP) algorithms are better suited for our problem. This technique indeed allows us to derive closed-form update equations that can be sequentially applied to each of the nodes of the factorized posterior distribution until a certain convergence criterion is met. We interpret these results in a Hierarchical Reinforcement Learning framework, where each level adapts its learning rate (LR) as a function of expected log-likelihood of the current observation given the past.

Fortunately, under the form of the approximate posterior we chose and using Conjugate VMP, the variational parameters of the posterior over the latent parameters \mathbf{z} have a simple form given the current value of $\boldsymbol{\phi}_t$ and $\boldsymbol{\beta}_t$. For a number of J observations observed at time t , we have:

$$\boldsymbol{\theta}_t^\xi = \widehat{\boldsymbol{\vartheta}}^\xi + \sum_{j=1}^J \mathbf{T}(x_{t,j}) \quad (5)$$

$$\boldsymbol{\theta}_t^\eta = \widehat{\boldsymbol{\vartheta}}^\eta + J \quad (6)$$

Equation (5) finds an interesting correspondent in the RL literature (Dearden & Russell, 1998). Consider the limit case where $\boldsymbol{\theta}_0 = 0$ and $J = 1$ (which is still analytically tractable following Equation (2)). As the expectation of a distribution of the EF has the general form $\mathbb{E}_{p(x|\mathbf{z})}[\mathbf{T}(x)] = dA(\mathbf{z})/d\mathbf{z}$, one can derive a similar posterior expectation of \mathbf{z} (Diaconis & Ylvisaker, 1979):

$$\mathbb{E}_{q(\mathbf{z}, w)}[\mathbf{z}] = \frac{\widehat{\boldsymbol{\vartheta}}^\xi + \mathbf{T}(x_t)}{\widehat{\boldsymbol{\vartheta}}^\eta + 1}$$

Now, replacing $\frac{1}{\widehat{\boldsymbol{\vartheta}}^\eta + 1}$ by α , the above expression becomes (Mathys, 2016)

$$\mathbb{E}_{q(\mathbf{z}, w)}[\mathbf{z}] = Q + \alpha(\mathbf{T}(x_t) - Q) \quad (7)$$

where $Q := \frac{\boldsymbol{\theta}_{t-1}^\xi}{\boldsymbol{\theta}_{t-1}^\eta}$ is the average \mathbf{z} at the time of the previous observation and α is the LR, whose value is inversely proportional to the effective memory $\boldsymbol{\theta}_{t-1}^\xi$ and to the current expected value of the forgetting factor $\mathbb{E}_{q(w)}[w]$ ¹. Equation (7) is a classical incremental update rule in single-stage

¹One can easily see that $\mathbb{E}_{q(w)}[w]$ dictates the memory of the learner. If $J = 1$ and assuming that $\mathbb{E}_{q(w)}[w]$ is stationary, we have: $\lim_{t \rightarrow \infty} \boldsymbol{\theta}_t^\eta = \boldsymbol{\theta}_0^\eta + \frac{1}{1 - \mathbb{E}_{q(w)}[w]}$

RL (Sutton & Barto, 1998), and our algorithm can be viewed as a special case of such algorithms where the LR is adapted online to the data at hand.

The update equations of ϕ_t is, however, not as simple to derive as θ_t , because $p(\mathbf{z} | \theta_{t-1}, \theta_0, w)$ is not conjugate to its Beta prior $p(\mathbf{z} | \phi_{t-1}, \phi_0, b)$. To solve this problem, we used a Non-Conjugate VMP approach (Knowles & Minka, 2011). Briefly, NCVMP minimizes an approximate KL divergence in order to find the value of the approximate posterior parameters that maximizes the ELBO. In order to use NCVMP, the first step is to derive the expected log-joint probability of the model, which we will need to differentiate wrt ϕ_t (or, in the case of the approximate posterior update rule for b, β_t). It quickly appears that part of this expression does not always have an analytical form for common exponential distributions: indeed, the expected value of $\mathbb{E}_{q(w)}[B(\boldsymbol{\vartheta})]$ is, in general, intractable and needs to be approximated. Expanding the Taylor series of this expression around $\hat{w} = \mathbb{E}_{q(w)}[w]$ up to the second order and taking the expectation, we have:

$$\mathbb{E}_{q(w)}[B(\boldsymbol{\vartheta})] \approx B(\hat{\boldsymbol{\vartheta}}) + \frac{1}{2} \mathbb{E}_{q(w)}[(w - \hat{w})^2] \nabla_{\hat{w}}^2 B(\hat{\boldsymbol{\vartheta}}) \quad (8)$$

Notice that the second term of the sum in Equation (8) can be expressed as $\frac{1}{2} \mathbb{V}\text{ar}_{q(w)}[w] d\boldsymbol{\theta}^T C(\mathbf{T}(\mathbf{z}) | \hat{\boldsymbol{\vartheta}}) d\boldsymbol{\theta}$, where $C(\mathbf{T}(\mathbf{z}) | \hat{\boldsymbol{\vartheta}})$ is the prior covariance of $\mathbf{T}(\mathbf{z})$. Hence, this penalty term becomes important when the product of the following factors increase: the distance between the previous posterior and initial prior $d\boldsymbol{\theta}$, the posterior variance of w and the prior covariance of $\mathbf{T}(\mathbf{z})$. This has the effect of favoring values of ϕ_t and w that have a low variance, especially when the two proposed distributions, q_{t-1} and p_0 , are very distant from each other.

We now derive the update equation for the approximate posterior of w . Let us first define

$$\begin{aligned} \delta_L &:= \frac{d}{d\hat{w}} \mathbb{E}_{q(\mathbf{z})} [\log p(\mathbf{z} | \hat{\boldsymbol{\vartheta}})] \\ \delta_C &:= -\frac{1}{2} \mathbb{V}\text{ar}_{q(w)}[w] \times \\ &\quad \frac{d}{d\hat{w}} d\boldsymbol{\theta}^T C(\mathbf{T}(\mathbf{z}) | \boldsymbol{\vartheta}) d\boldsymbol{\theta} \\ \delta_V &:= -\frac{1}{2} d\boldsymbol{\theta}^T C(\mathbf{T}(\mathbf{z}) | \hat{\boldsymbol{\vartheta}}) d\boldsymbol{\theta} \times \\ &\quad C(\log w | \phi)^{-1} \nabla_{\phi} \mathbb{V}\text{ar}_{q(w)}[w]. \end{aligned} \quad (9)$$

We obtain the following result:

Proposition 1. *Using Algorithm 1 of (Knowles & Minka, 2011), the update equation for ϕ_t has the form:*

$$\begin{aligned} \phi_t^\alpha &= \hat{\varphi}^\alpha + K(\phi_t^\alpha, \phi_t^\beta) \delta_L + K(\phi_t^\alpha, \phi_t^\beta) \delta_C + \delta_V^\alpha \\ \phi_t^\beta &= \hat{\varphi}^\beta - \underbrace{K(\phi_t^\beta, \phi_t^\alpha) \delta_L}_{u_1} - \underbrace{K(\phi_t^\beta, \phi_t^\alpha) \delta_C}_{u_2} + \underbrace{\delta_V^\beta}_{u_3} \end{aligned} \quad (10)$$

$$\text{where } K(x, y) := \frac{Mx + L(y)y}{(L(x)L(y) - M^2)(x + y)^2} > 0$$

$$L(x) := \psi_1(x) + M$$

$$M := -\psi_1(\phi_t^\alpha + \phi_t^\beta)$$

and $\psi_n(\cdot)$ is the n^{th} order polygamma function.

Proof. Follows directly from Algorithm 1 in (Knowles & Minka, 2011)². \square

The update equation in Equation (10) can be easily transposed for β_t .

In Proposition 1, we show that the update of ϕ_t can be decomposed in four terms: the first is the (weighted) prior φ , which acts as a reference for the update.

The second term, u_1 , depends upon δ_L , the derivative wrt \hat{w} of the expectation of the log probability $p(\mathbf{z} | \hat{\boldsymbol{\vartheta}})$ over \mathbf{z} , times a constant $K(\cdot, \cdot)$. δ_L has a simple form:

Lemma 1. *The derivative of the first order Taylor expansion of the expected log probability $\log p(\mathbf{z}) := \log p(\mathbf{z} | \hat{\boldsymbol{\vartheta}})$ around \hat{w} has the form*

$$\delta_L := \left(\left(\mathbb{E}_{q(\mathbf{z})}[\mathbf{T}(\mathbf{z})] - \mathbb{E}_{p(\mathbf{z})}[\mathbf{T}(\mathbf{z})] \right)^T d\boldsymbol{\theta} \right).$$

The proof is given in the supplementary materials. The expression of δ_L is easily understood as a measure of similarity between the current update of the variational posterior $q_t(\mathbf{z})$ and the previous posterior dependent prior $p(\mathbf{z})$. Note that a rather straightforward result of Equation (11) is that $\lim_{\theta_t^q \rightarrow \infty} \delta_L = 0$: as the posterior becomes stronger, the relative change that one can expect tends to zero, and the impact of δ_L on the update of ϕ can be expected to decrease. This is the behaviour we aimed at: a very strong posterior probability becomes more and more difficult to change as the training time increases.

Note also the opposite sign of the δ_L related increment in Equation (10) for ϕ_t^α and ϕ_t^β . This implies that if $\delta_L > 0$, then $u_1^\alpha > 0$, and the update of ϕ_t^α will tend to increase. The opposite is true for ϕ_t^β , showing that the posterior of w effectively deals with the similarity between the current observation and the previous ones.

The third and fourth term of Equation (10), u_2 and u_3 , are conditioned by the posterior variance of w and the prior variance of $\mathbf{T}(\mathbf{z})$. In brief, they push the update of ϕ in a direction that lowers the variance of both θ_t and ϕ . We will show in the next section a simple example of the relative contribution that each of these terms has in the update.

²The full development can be found in the supplementary materials.

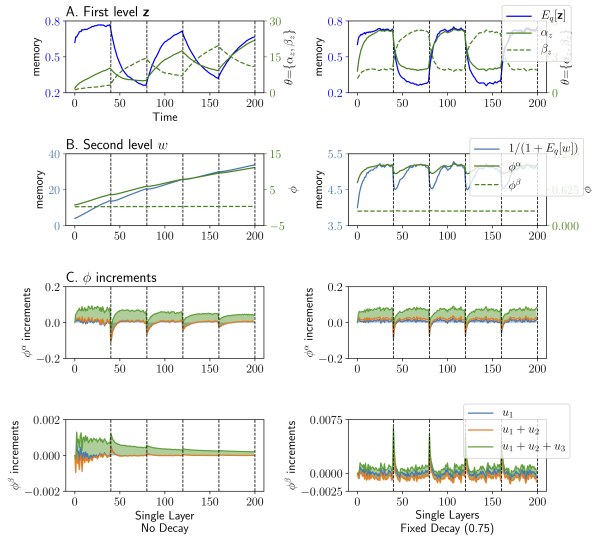


Figure 2. Binary learning with a single level of forgetting. Incremental (Left) and fixed-decay (Right) posterior learning of w . **A.** First level learning. The learner loses its capacity to forget as data are observed, because the expected effective memory (B.) tend to grow indefinitely when no decay was assumed. **C.** Trial-wise increment caused by u_1 , u_2 and u_3 . The effects of contingency changes decreased when no decay over w was considered.

An important consideration to make is that the value of ϕ must be > 0 , which implies that $u_1 + u_2 + u_3 > -\varphi$, a restriction that may be violated in practice, especially for low values of φ . In such cases, we reset the value of ϕ_t to some arbitrary value (typically $\phi \gg 0$) where the above inequality holds, and resume the update loop until convergence or until a certain amount of iterations is reached. Note that NCVMP is not guaranteed to converge, but, as suggested by (Knowles & Minka, 2011), exponential damping can improve the convergence of the algorithm.

2.2. Example: Binary distribution learning

In order to understand better the relative contribution of u_1 , u_2 and u_3 to the variational update scheme, we generated a sequence of 200 binary data distributed according to a binomial distribution, whose probability was switching between 0.8 and 0.2 every 40 trials. This distribution can be modelled as a hierarchy of beta distributions, where the first level is a Bernoulli distribution with a conjugate, Beta approximate posterior, and the one or two levels above are both Beta distributions measuring the stability of the level below. We simulated the learning process in three cases:

- A two-layer HAFVF model, where only the posterior over \mathbf{z} could be forgotten (incremental).

- A two-layer HAFVF model, where the posterior of w was being forgotten at a fixed rate (i.e. b fixed to 0.75).
- A three-layer HAFVF model, where the posterior of β was being forgotten at a rate of $\gamma = 0.999$.

In each of these examples, we used the following implementation: the beta prior of the first level was set to $\theta_0 = 1$. The value of ϕ_0 was set to $\{0.9, 0.1\}$, which showed to be a good compromise between informativeness and freedom to fit the data. If applicable, the top-level was set to $\beta = \{0.75, 0.25\}$.

In the first case, the fitting rapidly degenerated, as the memory grew at each trial. Figure 2, left column, gives a hint about the reason of this behaviour: each observation decreases the prior covariance $C(\mathbf{T}(\mathbf{z}) | \boldsymbol{\theta})$, which results in a positive increment for both ϕ_t^α and ϕ_t^β through u_3 . This can be viewed as a form of confirmation bias: because the posterior over w and \mathbf{z} are confident about the distribution of the data, they tend to reinforce each other and loose flexibility. Consequently, the impact of the contingency changes decreases as learning goes on. This might seem undesirable (and, in this pathological case, it is indeed the case), but in the case of datasets with outliers it can be very beneficial: a longer training in a stable environment will require a longer and/or stronger sequence of outliers to reset the parameters.

Adding a forgetting factor to the posterior of w can moderate the effect of overtraining. In the case of a fixed-forgetting for the posterior probability of w Figure 2, right column, the fitting is much more stable: the model is able to learn and forget the current distribution efficiently with a memory bounded at approximately 5 trials (i.e. $\mathbb{E}_{q(w)}[w] \approx 0.8$). This shows that adding a forgetting over the posterior of w effectively provides the flexibility we aim at: the contingency changes are efficiently detected, and the drop of δ_L (through u_1) triggers a resetting of the parameters in the following trials.

In the last case (Figure 3), the first level of the model acquires a higher memory than in the second example, due to the ability of the model to adapt the forgetting factor of w , which relaxes its bound. It is, however, more flexible than the first example. Long-term and flexible learning of the posterior probability of w therefore requires at least two more levels of forgetting (b and γ).

2.3. Forward-Backward algorithm

Let us consider the conjugate posterior of the distribution $p(x | \mathbf{z})$ from the EF $p(\mathbf{z})$ when the whole dataset has been observed. For a given $t \in 1 : T$, one can derive the posterior probability of \mathbf{z} given x_t as:

$$p(\mathbf{z} | x_t, \mathbf{x}_{-t}; \boldsymbol{\theta}_0) = \frac{p(x_t | \mathbf{z})p(\mathbf{z} | \mathbf{x}_{<t}, \mathbf{x}_{>t}; \boldsymbol{\theta}_0)}{p(x_t | \mathbf{x}_{<t}, \mathbf{x}_{>t})} \quad (11)$$

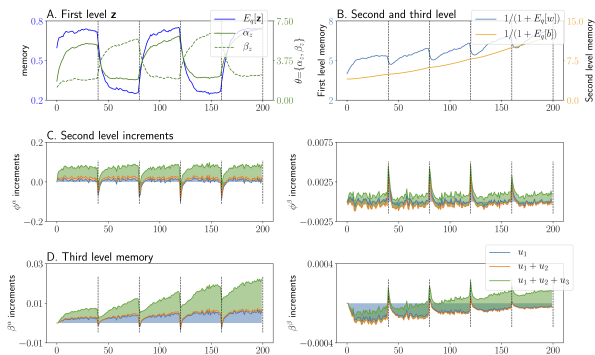


Figure 3. Binary learning with two levels of adaptive forgetting (w and b) and a third fixed level γ . **D.** is similar to **C.** for the third level updates $q(b)$.

Given Equation (2) and Equation (5), if $p(\mathbf{z} | \theta_0)$ is the conjugate prior of $p(x_t | \mathbf{x})$ and is from the EF, we can substitute the prior $p(\mathbf{z} | \mathbf{x}_{<t}, \mathbf{x}_{>t}; \theta_0)$ by $p(\mathbf{z} | \tilde{\mathbf{x}}_{<t}, \tilde{\mathbf{x}}_{>t}; \theta_0)$, where $\tilde{\mathbf{x}}_{<t}$ and $\tilde{\mathbf{x}}_{>t}$ are the effective samples retrieved from the forward and the backward application of the AFVF on the dataset, respectively. Formally, we have:

$$\begin{aligned} \theta_t^\xi &= f \theta_t^\xi + b \theta_t^\xi - \sum_{j=1}^J \mathbf{T}(x_{t,j}) - \theta_0^\xi & (12) \\ \theta_t^\eta &= f \theta_t^\eta + b \theta_t^\eta - J - \theta_0^\eta \end{aligned}$$

where the f and b superscripts index the forward and backward pass, respectively. In offline learning, this technique can increase the effective memory of the approximate posterior distribution just before and after the change trials.

3. Related work

Change detection is a broad field of research, where no optimal and general solution exists (KULHAVÝ & ZARROP, 1993). Consequently, assumptions about the structure of the system can lead to very different algorithms and results.

The Kalman Filter (KF) (Zarchan & Musoff, 2015) is a special case of Bayesian Filtering (BF) (Doucet et al., 2001) that has had a large success in the Signal Processing literature due to its sparsity and efficiency. It is, however, highly restrictive and its assumptions need to be relaxed in many instances. Specifically, we are here interested in cases where the evolution of the environment is unpredictable, in contrast with Hidden Markov Models and KF which explicitly give a structure to these changes (Bishop, 2006; Azizi & Quinn, 2015). One can discriminate two main approaches in order to deal with this problem: the first approach is to use a global approximation of BF such as Particle Filtering (PF)

(Smidl & Quinn, 2008; Smidl & Gustafsson, 2012; Özkan et al., 2013), which enjoys a bounded error but suffers from a lower accuracy than other local approximations. The second class of algorithms comprises the Stabilized Forgetting (SF) family of algorithms (KULHAVÝ & ZARROP, 1993; Azizi & Quinn, 2015; Laar et al., 2017), from which our model is a special case. SF suffers from an unbounded error, but it usually has a greater accuracy for a given amount of resources (Smidl & Quinn, 2008). Note that SF has been shown to be essential to reduce the divergence between the true posterior and its approximation in recursive Bayesian estimation (Kárný, 2014). As we apply the SF operator to estimate the posterior of \mathbf{z} and the mixture weight w (through the b weighted mixture prior), we ensure that the divergence is reduced for both of these latent variables.

Even though our model is described as a Stabilized Exponential Forgetting (Kulhavý & Kraus, 1996) algorithm and is well suited for signal processing, it can be generalized to models where there is no prediction of future states (e.g. smoothing of a signal, temporal difference learning etc.) Also, it overcomes other methods in several following ways:

First, it uses a Beta prior on the mixing coefficient. This is unusual (but not unique (Dedecius & Hofman, 2012)), as most of previous approaches used a truncated exponential prior (Smidl & Quinn, 2005; Masegosa et al., 2017) or a fixed, linear mixture prior that account for the stability of the process (Smidl & Gustafsson, 2012). In Linear SF, a Bernoulli prior with a Beta hyperprior has been proposed for the mixture weight (Laar et al., 2017). Our approach is designed to learn the posterior probability of the forgetting factor in a flexible manner. We show that this posterior probability depends upon its own (and possibly a mixture of) prior distribution and upon the prior covariance of the model parameters $C(\mathbf{T}(\mathbf{z}) | \hat{\vartheta})$. This makes the change detection more subtle than an all-or-none process, as one might observe with a Bernoulli distribution. It also enables us to accumulate evidence for a change of distribution across trials, which can help to discriminate outliers from real, prolonged contingency changes. This is, to our knowledge, an entirely novel feature in the adaptive forgetting literature.

The second important novelty of our model lies in its hierarchical learning of the environment stability. This is somehow similar to the Hierarchical Gaussian Filter (HGF) (Mathys, 2011; Mathys et al., 2014). The present model is, however, much more general, as the generic form we provide can be applied to several members of the EF. Also, although the KL divergence (error term) of our model is not bounded in the long run, it can be efficiently applied to a large subset of datasets and models, whereas the HGF often fails to fit processes that are highly stationary, with many datapoints and/or with abrupt contingency changes.

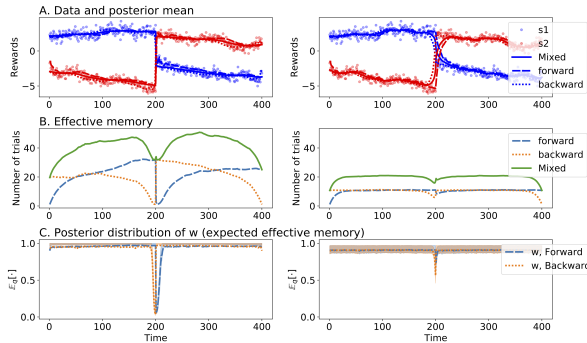


Figure 4. Experiment 1. Left column: weak prior over w . Right column: strong prior over w . Shaded areas represent the 95% posterior confidence interval. See text for more comments.

4. Experiments

The HAFVF was coded in the Julia language (Bezanson et al., 2017) using a Forward automatic differentiation algorithm (Revels et al., 2016) for the NCVMP for the RL and AR parts of this section, and using an analytical gradient for the SGD part.

4.1. Smoothing

We first look at the behaviour of the model in the simple case of estimating the current distribution of a random variable sampled from a moving distribution. We simulated two sequences of 2×200 datapoints where each pair of points was generated according to the same multivariate normal distribution with mean $\mu_1 = \{-2, +2\}$ and $\mu_2 = \{+2, -2\}$. We then added an independent random walk to these means.

We applied the Forward-Backward (FB) version of the HAFVF to these datasets. We used the same Normal Inverse Wishart prior for both of these results ($\mu_0 = 0$, $\kappa_0 = 0.1$, $\eta_0 = 3$, $\Lambda_0 = I$). The prior over w was manipulated to include a high confidence ($\phi_t^\alpha = 9$, $\phi_t^\beta = 1$) or a low confidence ($\phi_t^\alpha = 0.9$, $\phi_t^\beta = 0.1$) about the average value of w . Note that both of these priors had the same expected value. To avoid overfitting of early trials (which may happen using weak priors) while keeping the distribution flexible, we used a flat prior over b : $\beta_0 = 1$. The top level forgetting was ignored ($\gamma = 1$). Results are shown in Figure 4.

As the first setting had a weak prior over w , it had more freedom to adapt the posterior distribution to the current data. The effective memory trace (measured with the parameter κ_t) was greater when the environment was stable, and changed faster after the contingency change than when the prior was more confident, where the adaptation was slow and the effective memory did not increase much above the prior-defined threshold 10 (or 20 for the FB algorithm).

The behaviour of both models after the contingency change is informative about the effect that the prior had on the inference process: the weak-prior forgetting factor dropped immediately after an unexpected observation was made, which can be advantageous when sudden changes are expected, but maladaptive in the presence of outliers. The strong-prior model behaved in the opposite way, and handled the change more slowly than its weak-prior counterpart.

It is interesting to note that the posterior probability distribution of b (not shown in the figure) was also more flexible in the first model fit than in the second, because the observations in the level below were also more variable, due to less confident prior over w : this had the effect of increasing the gain in precision over w , which increased the strength of the posterior over b (through u_2 and u_3 in Equation (10)).

4.2. Autoregressive model

We fitted the HAFVF to a simulated a non-stationary sinusoidal signal of 400 datapoints issued from two separate systems with a low and high frequency. These signals were randomly generated as the sum of five sinusoidal waves, with the scope of observing whether the algorithm was able to adapt to the abrupt contingency change.

Because we also aimed at a more informative view on the performance of the algorithm in the presence of artifacts, we altered this signal by adding two impulses of 2 a.u. at $t = 100$ and $t = 300$.

We studied a single implementation of the model, with a relatively strong prior over w ($\phi = \{4.5, 0.5\}$) and a flat prior over b , ($\beta = \{1.0, 1.0\}$). The Forward-backward version of the algorithm was applied. We arbitrarily chose a forward and a backward order of 10 samples. Figure 5 shows the results of this experiment.

4.3. Stochastic Gradient Descent

SGD is a popular technique to find the minimum of (often computationally expensive) loss-functions over large datasets (Tran et al., 2015) or involving intractable integrals (Kingma & Welling, 2013) that can be sampled from. However, SGD can be unstable, especially with recurrent neural networks (Fabius & van Amersfoort, 2014) where an isolated, highly noisy sample in the sequence can lead to a degenerate sample of the gradient over the whole sequence. This effect is further magnified when the sample size is low.

We implemented a slightly modified version of the HAFVF in a SGD framework, intended to be similar to the Adam optimizer³ (Kingma & Ba, 2015). In short, we used two specific decays w_1 and w_2 for the posterior of the means and variances of the gradients, respectively, while ensuring

³More details can be found in the supplementary materials

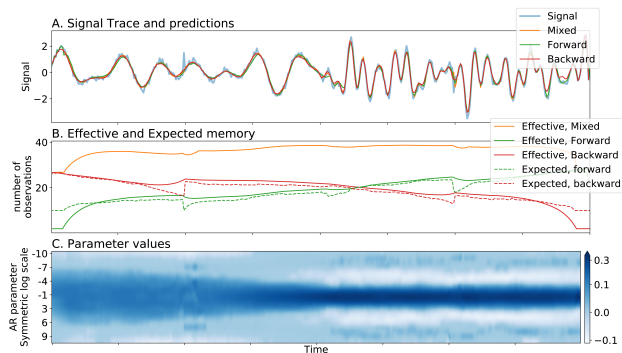


Figure 5. Experiment 2: Autoregressive model with a weak prior over w . **A.** Observations and simulated response of the models. The zoomed windows show the effect of the artifacts on the estimated mean value. **B.** Effective memory (the “effective number of observations” parameter of the posterior θ_t : namely κ_t) of the three parts of the algorithm (plain lines), and corresponding expected effective memory: $1/(1 + \mathbb{E}_{q(w)}[w])$ (dashed lines). Outliers had a limited impact on learning in both cases. **C.** Value of the AR mean weights through time. The model dealt adequately with the outliers (as the value of the parameters did not change substantially) and with the contingency change (as the values were adapted to the two different signals).

that $w_1 < w_2$. We modelled these posteriors as a set of Normal-Inverse-Gamma distributions. Each set of weights and biases of the multilayer perceptron was provided with its own hierarchical decay, to take advantage of the fact that some groups of partial derivatives might be more noisy than others. We used this algorithm with a strong prior over w_1 and w_2 ($\phi_1 = \{9, 1\}$ and $\phi_2 = \{9.5, 0.5\}$), to limit the effect of degenerated gradients on the approximate posteriors.

This algorithm was tested with a variational recurrent auto-encoding regression model inspired by (Moens & Zenon, 2018), where the output probability density was a first passage density of a Wiener process (Ratcliff, 1978). The simulated dataset was composed of 64 subjects performing a 500 trials long two alternative forced choice task (Britten et al., 1992), where choices and reaction times were the model was aiming to predict. At each step of the SGD process, 5 subjects were sampled, for a total of 2500 trials.

Figure 6 compares the results of the AdaFVF SGD optimizer with the Adam optimizer, executed with the default parameters ($\alpha = 0.001$, $\beta_1 = 0.9$, $\beta_2 = 0.999$). The AdaFVF showed to be less affected by degenerate samples than Adam, as can be seen from the ELBO trace and from the heat plots of the expected memories, for an estimated average negative ELBO of 1.08 for Adam and 0.85 for AdaFVF at the iteration 10000.

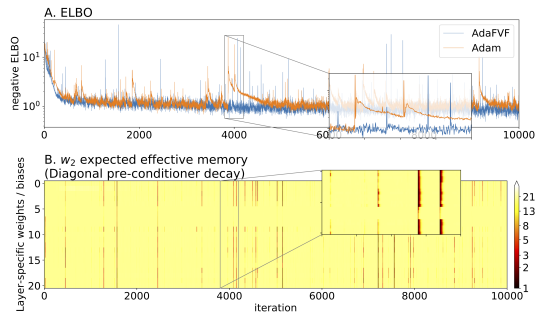


Figure 6. Experiment 3: SGD with the HAFVF. **A.** ELBO for the Adam optimizer and for the AdaFVF SGD. After an outlier was sampled, the AdaFVF forgot the gradient history, and reset its belief to the naive prior, thereby decreasing the relative contribution of this sample. **B.** Expected memory of the variance posteriors. The impact of outliers is highlighted by the zoomed windows.

5. Limitations, perspective and conclusion

Our algorithm has the following limitations: first, the exponential form we have given to the mixture distributions restricts the mixture prior to have a single mode. A linear form, similar to (Laar et al., 2017) could however also be implemented, at specific levels of the hierarchy of the whole model. It may also be difficult to choose adequate level-specific priors. A generic feature in adaptive forgetting is that the naive prior of the lower level $p_0(\mathbf{z})$ is usually crucial but hard to specify. For the two top levels, we propose as a rule of thumb to use a weak prior in situations where abrupt contingency changes are expected. Such configurations will usually provide a higher memory to the model but are, however, more affected by outliers than stronger priors. Performing updates using mini-batches of datapoints should mitigate this effect. If the environment is meta-stable (i.e. change points occur at a given frequency), derivation of an online type II maximum likelihood optimization algorithm of the prior parameters could improve the algorithm efficiency.

The HAFVF and variants could lead to many promising developments in RL related fields, where they might help to prevent unnecessary forgetting of past events during exploration, in signal processing and more distant fields such as deep learning, where they could be used to prevent the occurrence of catastrophic forgetting.

In conclusion, we present a new generic model aimed at coping with abrupt or slow signal changes and presence of artifacts. This model flexibly adapts its memory to the volatility of the environment, and reduces the risk of abruptly forgetting its learned belief when isolated, unexpected events occur. The HAFVF constitutes a promising tool for decay adaptation in RL, system identification and SGD.

Acknowledgements

We thank the reviewers for their careful reading and precious comments on the manuscript. We also thank O. Solopchuk and A. Z enon, who greatly contributed to the development and writing of the present manuscript. The present work was supported by grants from the ARC (Actions de Recherche Concert ees, Communaut  Francaise de Belgique).

References

- Azizi, S. and Quinn, A. A data-driven forgetting factor for stabilized forgetting in approximate Bayesian filtering. In *2015 26th Irish Signals and Systems Conference (ISSC)*, volume 11855, pp. 1–6. IEEE, jun 2015. ISBN 978-1-4673-6974-9. doi: 10.1109/ISSC.2015.7163747. URL <http://ieeexplore.ieee.org/document/7163747/>.
- Bezanson, J., Edelman, A., Karpinski, S., and Shah, V. B. Julia: A Fresh Approach to Numerical Computing. *SIAM Review*, 59(1):65–98, jan 2017. ISSN 0036-1445. doi: 10.1137/141000671. URL <http://arxiv.org/abs/1411.1607><http://epubs.siam.org/doi/10.1137/141000671>.
- Bishop, C. *Pattern Recognition And Machine Learning*. Springer-Verlag New York, 2006. ISBN 978-0-387-31073-2. URL <http://www.springer.com/us/book/9780387310732>.
- Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. a. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 12(12):4745–65, dec 1992. ISSN 0270-6474. doi: 10.1111.123.9899. URL <http://www.ncbi.nlm.nih.gov/pubmed/1464765>.
- Dearden, R. and Russell, S. Bayesian Q-learning. *AAAI*, 1998. URL <http://www.aaai.org/Papers/AAAI/1998/AAAI98-108.pdf>.
- Dedecius, K. and Hofman, R. Autoregressive model with partial forgetting within Rao-Blackwellized particle filter. *Communications in Statistics: Simulation and Computation*, 41(5):582–589, 2012. ISSN 03610918. doi: 10.1080/03610918.2011.598992.
- Diaconis, P. and Ylvisaker, D. Conjugate Priors for Exponential Families. *The Annals of Statistics*, 7(2):269–281, mar 1979. ISSN 0090-5364. doi: 10.1214/aos/1176344611. URL <http://projecteuclid.org/euclid.aos/1176344611>.
- Dickinson, A. Actions and Habits: The Development of Behavioural Autonomy. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 308(1135):67–78, feb 1985. ISSN 0962-8436. doi: 10.1098/rstb.1985.0010. URL <http://rstb.royalsocietypublishing.org/cgi/doi/10.1098/rstb.1985.0010>.
- Doucet, a., De Freitas, N., and Gordon, N. Sequential Monte Carlo Methods in Practice. *Springer New York*, pp. 178–195, 2001. ISSN 1530-888X. doi: 10.1198/tech.2003.s23. URL <http://www-sigproc.eng.cam.ac.uk/{~}ad2/book.html>.
- Fabius, O. and van Amersfoort, J. R. Variational Recurrent Auto-Encoders. (2013):1–5, dec 2014. URL <http://arxiv.org/abs/1412.6581>.
- Hoffman, M., Blei, D. M., Wang, C., and Paisley, J. Stochastic Variational Inference. (2), 2012. ISSN 1532-4435. doi: citeulike-article-id:10852147. URL <http://arxiv.org/abs/1206.7051>.
- Jaakkola, T. S. and Jordan, M. I. Bayesian parameter estimation via variational methods. *Statistics And Computing*, 10(1):25–37, 2000. ISSN 0960-3174. doi: 10.1023/A:1008932416310. URL papers2://publication/uuid/6B83D6F3-56DD-4219-A399-4FBC826D8223.
- K rn y, M. Approximate Bayesian recursive estimation. *Information Sciences*, 285(1):100–111, 2014. ISSN 00200255. doi: 10.1016/j.ins.2014.01.048.
- Kingma, D. P. and Ba, J. L. Adam: a Method for Stochastic Optimization. *International Conference on Learning Representations 2015*, pp. 1–15, 2015.
- Kingma, D. P. and Welling, M. Auto-Encoding Variational Bayes. dec 2013. URL <http://arxiv.org/abs/1312.6114><http://www.aanda.org/10.1051/0004-6361/201527329>.
- Knowles, D. and Minka, T. P. Non-conjugate variational message passing for multinomial and binary regression. *Nips*, pp. 1–9, 2011. URL <http://eprints.pascal-network.org/archive/00008459/>.
- Kulhavy, R. and Karny, M. Tracking of slowly varying parameters by directional forgetting. *Preprints 9th IFAC Congress*, 10:178–183, 1984. ISSN 07411146. doi: 10.1016/S1474-6670(17)61051-6. URL citeulike-article-id:10156625.
- Kulhav y, R. and Kraus, F. On Duality of Exponential and Linear Forgetting. *IFAC Proceedings Volumes*, 29(1):5340–5345, jun 1996. ISSN 14746670. doi: 10.1016/S1474-6670(17)58530-4. URL <http://linkinghub.elsevier.com/retrieve/pii/S1474667017585304>.

- KULHAVÝ, R. and ZARROP, M. B. On a general concept of forgetting. *International Journal of Control*, 58(4):905–924, oct 1993. ISSN 0020-7179. doi: 10.1080/00207179308923034. URL <http://www.tandfonline.com/doi/abs/10.1080/00207179308923034>.
- Laar, T. V. D., Cox, M., Diepen, A. V., and Vries, B. D. Variational Stabilized Linear Forgetting in State-Space Models. (Section ID):848–852, 2017.
- Mandt, S., McInerney, J., Abrol, F., Ranganath, R., and Blei, D. Variational Tempering. 41, 2014. URL <http://arxiv.org/abs/1411.1810>.
- Masegosa, A., Nielsen, T. D., Langseth, H., Ramos-Lopez, D., Salmeron, A., and Madsen, A. L. Bayesian Models of Data Streams with Hierarchical Power Priors. *International Conference on Machine Learning (ICM)*, 70:2334–2343, 2017. URL <http://proceedings.mlr.press/v70/masegosa17a.html>.
- Mathys, C. A Bayesian foundation for individual learning under uncertainty. *Frontiers in Human Neuroscience*, 5(May):1–20, 2011. ISSN 16625161. doi: 10.3389/fnhum.2011.00039. URL <http://journal.frontiersin.org/article/10.3389/fnhum.2011.00039/abstract>.
- Mathys, C. How could we get nosology from computation? *Computational Psychiatry: New Perspectives on Mental Illness*, 20:121–138, 2016. URL <https://mitpress.mit.edu/books/computational-psychiatry>.
- Mathys, C. D., Lomakina, E. I., Daunizeau, J., Iglesias, S., Brodersen, K. H., Friston, K. J., and Stephan, K. E. Uncertainty in perception and the Hierarchical Gaussian Filter. *Frontiers in Human Neuroscience*, 8(November):825, nov 2014. ISSN 1662-5161. doi: 10.3389/fnhum.2014.00825. URL <http://journal.frontiersin.org/article/10.3389/fnhum.2014.00825/abstract>.
- Moens, V. and Zenon, A. Recurrent Auto-Encoding Drift Diffusion Model. *bioRxiv*, 2018. doi: 10.1101/220517. URL <https://www.biorxiv.org/content/early/2018/05/13/220517>.
- Özkan, E., Šmídl, V., Saha, S., Lundquist, C., and Gustafsson, F. Marginalized adaptive particle filtering for nonlinear models with unknown time-varying noise parameters. *Automatica*, 49(6):1566–1575, jun 2013. ISSN 00051098. doi: 10.1016/j.automatica.2013.02.046. URL <http://linkinghub.elsevier.com/retrieve/pii/S0005109813001350>.
- Ratcliff, R. A theory of memory retrieval. *Psychological Review*, 85(2):59–108, 1978. doi: <http://dx.doi.org/10.1037/0033-295X.85.2.59>.
- Revels, J., Lubin, M., and Papamarkou, T. Forward-Mode Automatic Differentiation in Julia. (April):4, 2016. URL <http://arxiv.org/abs/1607.07892>.
- Smídl, V. *The Variational Bayes Approach in Signal Processing*. PhD thesis, 2004. URL <http://staff.utia.cz/smídl/Public/Thesis-final.pdf>.
- Smídl, V. and Gustafsson, F. Bayesian estimation of forgetting factor in adaptive filtering and change detection. In *2012 IEEE Statistical Signal Processing Workshop (SSP)*, number 1, pp. 197–200. IEEE, aug 2012. ISBN 978-1-4673-0182-4. doi: 10.1109/SSP.2012.6319658. URL <http://ieeexplore.ieee.org/document/6319658/>.
- Smídl, V. and Quinn, A. Mixture-based extension of the AR model and its recursive Bayesian identification. *IEEE Transactions on Signal Processing*, 53(9):3530–3542, sep 2005. ISSN 1053-587X. doi: 10.1109/TSP.2005.853103. URL <http://ieeexplore.ieee.org/document/1495888/>.
- Smídl, V. and Quinn, A. Variational Bayesian Filtering. *IEEE Transactions on Signal Processing*, 56(10):5020–5030, oct 2008. ISSN 1053-587X. doi: 10.1109/TSP.2008.928969. URL <http://ieeexplore.ieee.org/document/4585346/>.
- Sutton, R. S. and Barto, A. G. Introduction to Reinforcement Learning. *Learning*, 4:1–5, 1998. ISSN 10743529. doi: 10.1.1.32.7692. URL <http://dl.acm.org/citation.cfm?id=551283>.
- Tran, D., Toulis, P., and Airoldi, E. M. Stochastic gradient descent methods for estimation with large data sets. *VV*(October), sep 2015. URL <http://arxiv.org/abs/1509.06459>.
- Winn, J., Bishop, C., and Jaakkola, T. Variational Message Passing. *Journal of Machine Learning Research*, 6:661–694, 2005. ISSN 1532-4435. doi: 10.1007/s002130100880.
- Zarchan, P. and Musoff, H. (eds.). *Fundamentals of Kalman Filtering: A Practical Approach, Fourth Edition*. American Institute of Aeronautics and Astronautics, Inc., Reston, VA, jan 2015. ISBN 978-1-62410-276-9. doi: 10.2514/4.102776. URL <http://arc.aiaa.org/doi/book/10.2514/4.102776>.