

Learning as filtering: implications for spike-based plasticity

Jannes Jegminat^{*,a,b} and Jean-Pascal Pfister^{a,b}

^aDepartment of Physiology, University of Bern, 3012 Bern, Switzerland

^bInstitute of Neuroinformatics and Neuroscience Center Zurich, ETH and the University of Zurich, 8057 Zurich, Switzerland

This manuscript was compiled on August 10, 2020

Abstract

Most normative models in computational neuroscience describe the task of learning as the optimisation of a cost function with respect to a set of parameters. However, learning as optimisation fails to account for a time varying environment during the learning process; and the resulting point estimate in parameter space does not account for uncertainty. Here, we frame learning as filtering, i.e., a principled method for including time and parameter uncertainty. We derive the filtering-based learning rule for a spiking neuronal network - the Synaptic Filter - and show its computational and biological relevance. For the computational relevance, we show that filtering in combination with Bayesian regression improves performance compared to a gradient learning rule with optimal learning rate in terms of weight estimation. Furthermore, the filtering-based rule outperforms gradient-based rules in the presence of model mismatch, indicating a better generalisation performance. The dynamics of the mean of the Synaptic Filter is consistent with the spike-timing dependent plasticity (STDP) while the dynamics of the variance makes novel predictions regarding spike-timing dependent changes of EPSP variability. Moreover, the Synaptic Filter explains experimentally observed negative correlations between homo- and heterosynaptic plasticity.

Significance statement

The task of learning is commonly framed as parameter optimisation. Here, we adopt the framework of learning as filtering where the task is to continuously define the uncertainty about the parameters to be learned. We apply it to synaptic plasticity in a spiking neuronal network. Filtering includes a time varying environment and parameter uncertainty on the level of the learning task. We show that learning as filtering can qualitatively explain two biological experiments on synaptic plasticity that cannot be explained by learning as optimisation. Moreover, we make a new prediction and improve performance with respect to a gradient learning rule. Thus, learning as filtering is promising candidate for learning models.

1 Introduction

In computational neuroscience, most normative models frame learning as optimisation of a static cost function with respect to a set of parameters, such as synaptic efficacies [Stemmler and Koch, 1999, Seung, 2003, Lengyel et al., 2005, Booij and tat Nguyen, 2005, Gütig and Sompolinsky, 2006, Urbanczik and Senn, 2014, Xu et al., 2013, Urbanczik and Senn, 2014, Bohte et al., 2002, Bohte and Mozer, 2005] or a neuron’s excitability [Triesch, 2005a,b]. Different cost functions have been used to reproduce or predict experimental findings, such as a measure of sparseness and information preservation [Olshausen and Field, 1996], mutual information [Stemmler and Koch, 1999], the probability of timed postsynaptic spiking [Pfister et al., 2006, Brea et al., 2013], the mutual information of input and output spike trains [Chechik, 2003, Toyozumi et al., 2005], the network sensitivity [Bell and Parra, 2005] and free energy [Buckley et al., 2017].

However learning as optimisation has the drawback of not taking parameter uncertainty into account [MacKay, 1992]. When few training data are available compared to the number of model parameters, the parameter space is not sufficiently constrained, i.e., multiple parameter instantiations yield comparable model

*jannes.ini@uzh.ch

performance. Optimisation selects the best performing parameter, thereby ignoring the inherent parameter uncertainty present in a (probabilistic) model. This contributes to the problem of overfitting, i.e., the resulting performance on the training data does not generalise to the testing data [Gal and Ghahramani, 2016, Blundell et al., 2015]. Moreover, many decision making models require as input not only the most likely prediction but also prediction uncertainty [Bell, 1982]. To obtain accurate prediction uncertainty, the contribution of parameter uncertainty must be taken into account (e.g. [Henning et al., 2018]). Thus parameter uncertainty is computationally relevant for avoiding overfitting and the estimation of prediction uncertainty.

Learning as static optimisation is further limited because it lacks a principled way of accounting for a dynamic environment during learning. Often the data distribution is assumed to be static, i.e., independent of time. However, in many settings the environment and, thus, the data distribution, are dynamic. For example, the association between a location and the availability of food is not static when the source of food depletes over time. Dynamic environments pose the additional challenge of determining the speed of learning. A slow learner fails to adapt to quickly changing environmental statistics while an overly fast learner might disregard past data prematurely. The question of how to account for a dynamic environment during learning is closely related to continual learning, i.e., the task of sequentially learning from multiple data sets while maintaining (testing) performance on all previously observed ones [Kirkpatrick et al., 2017, Parisi et al., 2019]. Here, the dynamics of the environment translate into the sequential availability of data sets.

In neuroscience, time and uncertainty play an important role. Many studies have shown that a dynamic environment affects how animals learn [Kraemer and Golding, 1997, Zimmermann et al., 2018]. For instance, flies learn odour association and adapt their forgetting rate of old associations [Shuai et al., 2010, Brea et al., 2014]. Similar experiments have been conducted with rodents [Fassih et al., 2014, Akrami et al., 2011]. Uncertainty of rewards has been studied in prefrontal and cingulate cortex on the basis of reinforcement learning models [Rushworth and Behrens, 2008]; and several neuro-modulators have been identified that influence choices under uncertainty [Doya, 2008, Cocker et al., 2012]. The uncertainty related to a whisker stimulus is directly related to neuronal activity in rat barrel cortex [Stüttgen and Schwarz, 2008]. Uncertainty has also been linked to neuronal codes [Ma, 2010, Stüttgen and Schwarz, 2008] and many aspects of perception and decision making [Knill and Pouget, 2004]. In the context of plasticity, uncertainty of synaptic weights has been linked to spine turnover [Kappel et al., 2015].

Normative models of learning can benefit from going beyond the framework of static optimisation by including parameter uncertainty and time in the learning task. However, it remains unclear which framework could prove to be a fruitful alternative. Here, we propose learning as filtering. Filtering, developed by mathematicians in the 60's [Kushner, 1964, 1967], is a principled way to include time and uncertainty. It continuously computes the posterior distribution (also called the filtering distribution) of a latent variable from all the observations up to time t . We apply learning as filtering to synaptic plasticity, a field in which the need for new learning paradigms has become apparent [Brea and Gerstner, 2016].

In a continuous-time, spiking neuronal network, we derive the update rule for the synaptic weight distribution and call it the Synaptic Filter. The Synaptic Filter is computationally relevant because it outperforms a gradient rule with optimised learning rate parameter in a dynamic weight estimation task, confirming a previous result [Aitchison and Latham, 2015]. Going beyond performance measured in weight space, we study the predictive performance and find that the Synaptic Filter outperforms the optimised gradient rule as well, and that it is robust to the presence of model mismatch. From the biological perspective, the Synaptic Filter makes three experimental predictions. First, the mean synaptic weight change depends on the precise timing of pre- and postsynaptic spikes and is therefore reminiscent of spike-timing dependent plasticity (STDP), which yields long term potentiation of the synaptic strength (LTP) if the postsynaptic spikes follows the presynaptic spikes and long term depression (LTD) otherwise [Markram et al., 1997, Bi and Poo, 1998]. Normative models of STDP have provided a consistent view on the pre-post LTP lobe. Pre-before-post pairs induce LTP reinforcing causality. Therefore the time constant of LTP reflects the EPSP time constant. However, normative models do not provide a unifying view on the LTD window [Pfister et al., 2006]. Here, we provide a novel computational rationale for the LTD lobe, namely to compensate for a change in bias. Secondly, based on the hypothesis that EPSP variability encodes synaptic weight uncertainty [Aitchison and Latham, 2014, 2015] we formulate the novel prediction that EPSP variability also changes as a function of the precise timing of the spikes. Finally, weight changes induced by joint pre- and postsynaptic activity at one synapse can induce weight changes at synapses that did not receive presynaptic input, reminiscent of the phenomenon of heterosynaptic plasticity. Indeed, our learning rule can explain the negative correlation between homo- and heterosynaptic plasticity observed in experiments [Royer and Paré, 2003].

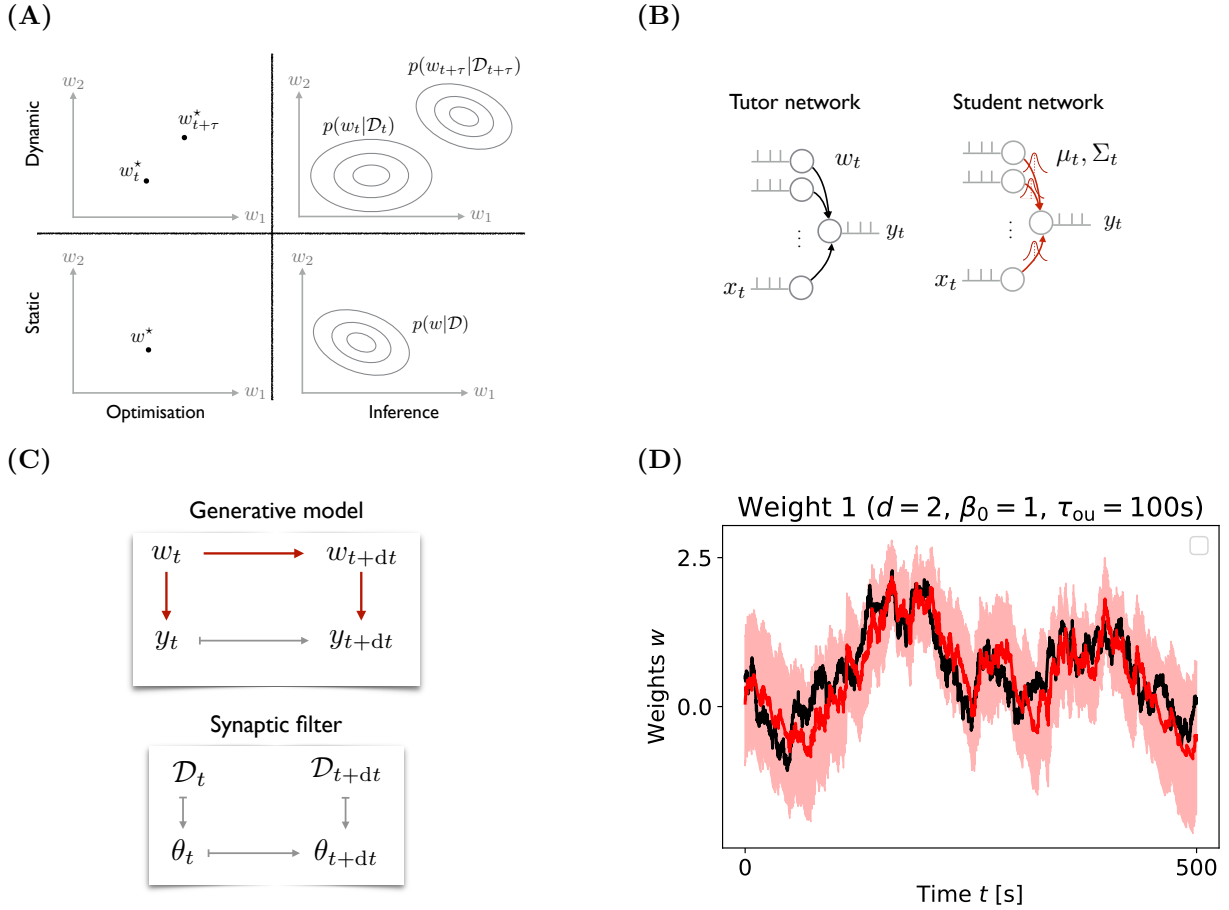


Figure 1: (A) Learning tasks can be static or dynamic, and deterministic or stochastic. (B) The generative model represents the assumption that the observed spike train y_t was generated from a tutor network with the same input x_t and hidden weights w_t . (C) Graphical model representation of the generative model (top), the Synaptic Filter (bottom) with deterministic dependencies shown in gray and probabilistic ones in red. (D) Time series of a ground truth weight (black) in the tutor network and weight distribution (red, shaded area = 2-SD) learned by the student network.

2 Results

2.1 The Synaptic Filter

The goal of learning is to find predictive functions from training data \mathcal{D} which map inputs x to outputs y typically based on a parametrised generative model. The generative model specifies how the output y of the predictive function is computed from the parameters w and the input x . Accounting for a dynamic environment on the level of the generative model, makes the data and the parameter w time dependent. Accounting for the fact that many parameter instantiations are compatible with the training data corresponds to computing predictions based on parameter uncertainty. Thus, a given static generative model for learning can be generalised by including the assumption of a dynamic environment or parameter uncertainty.

Including both, time and parameter uncertainty, yields the framework of learning as filtering. Figure 1 (A) illustrates how filtering generalises a static optimisation, which is the dominant learning framework. The learning task of static optimisation is to find a point estimate w^* in parameter space given the generative model and the data \mathcal{D} . By including parameter uncertainty, the learning task generalises to inferring the posterior distribution over parameters $p(w|\mathcal{D})$. In the limit of infinite data and for a convex parameter landscape, the parameter distribution collapses around a point, yielding similar results for parameter optimisation and inference, i.e., $p(w|\mathcal{D}) \approx \delta(w-w^*)$. However, in many problems equivalent minima exist and the amount of data is limited such that the posterior distribution is not well approximated by a point estimate. Another extension of static optimisation considers dynamically changing environmental statistics, i.e., the data distribution is time dependent. In this case, the learning task is to track the optimal parameter as a function of time w_t^* . In a filtering framework, which includes both, parameter uncertainty and time, the task is to compute the so-called filtering distribution over the weights as a function of time $p(w_t|\mathcal{D}_t)$ given all previous observations up to time t .

A generative model to study spike-based plasticity

To study learning as filtering in the context of spike-based plasticity, we consider a generative model with time dependent parameters $w_t \in \mathbb{R}^d$, the weights. In contrast to learning as static optimisation, we do not assume that the weights are fixed. Changes in the weights reflect changes in the statistics of the environment. Here, we assume that the weights follow an Ornstein-Uhlenbeck (OU) process with mean $\mu_{\text{ou}} = 0$, diagonal equilibrium covariance matrix $\Sigma_{\text{ou}} = \sigma_{\text{ou}}^2 \mathbb{1}$ with (non-zero elements) $\sigma_{\text{ou}}^2 = 1$ and time scale τ_{ou} . The limit of a large OU time constant, represents a static environment while the limit $\tau_{\text{ou}} \rightarrow 0$ represents an environment that changes too fast for meaningful learning.

At each moment in time, the weights relate the input spike trains and output spike of a single neuron via the observation probability $p(y_t|w_t, x_{0:t})$. Here, we assume that output spikes y_t are generated from an inhomogeneous Poisson neuron with instantaneous firing rate given by an exponential gain function $g(u_t) = g_0 \exp(\beta u_t)$ with base rate g_0 . The membrane potential u_t is a sum of presynaptic inputs weighted by w_t . The determinism parameter β controls how strongly the membrane potential affects the output spiking. For $\beta = 0$, the output spikes are independent of the membrane u_t and reflect only the base rate g_0 . With the additional assumption that the time scale of the membrane potential τ_m is much smaller than the one of the weights, i.e., $\tau_m \ll \tau_{\text{ou}}$, we ensure the Markovianity of the generative model (see Equation (7) in Materials and Methods). The generative model is represented as graphical model in Figure 1 (C). To ensure that performance metrics can be compared across various dimensions d , we let the determinism parameter scale with the dimension $\beta \propto \beta_0 d^{-1/2}$ such that the expected firing rate of the output neuron, and hence the rate of observations from the hidden weights, becomes independent of the dimension d (see Section 4.3.1 in Materials and Methods).

An assumed density filter solution: the Synaptic Filter

The generative model can be conceptualised as the belief that a tutor network with ground truth weights w_t , illustrated in Figure 1 (B), generates the observed output spikes from given inputs. Learning as filtering corresponds to a student network that continuously computes the distribution over the ground truth weights $p(w_t|\mathcal{D}_t)$ given the history of inputs and outputs $\mathcal{D}_t = \{(x, y)_\tau\}_{\tau=0}^t$.

Generally, the filtering distribution $p(w_t|\mathcal{D}_t)$ is intractable. Here, we obtain an approximated solution with an Assumed Density Filter (see Supplementary Information for the derivation), i.e., the exact filtering distribution $p(w_t|\mathcal{D}_t)$ is approximated with the parametric distribution $q_{\theta_t}(w_t)$. For the proposed generative model, we call the Gaussian Assumed Density Filter the *Synaptic Filter*. The distribution parameters $\theta_t := (\mu_t, \Sigma_t)$ denote the mean μ_t and covariance matrix Σ_t of the Gaussian. An Assumed Density Filtering reduces the problem to updating the distribution parameters θ_t based on observations:

$$\dot{\mu}_t = \beta \Sigma_t x_t^\epsilon (y_t - \gamma_t) + \tau_{\text{ou}}^{-1} (\mu_{\text{ou}} - \mu_t), \quad (1)$$

$$\dot{\Sigma}_t = -\beta^2 \gamma_t (\Sigma_t x_t^\epsilon) (\Sigma_t x_t^\epsilon)^\top + 2\tau_{\text{ou}}^{-1} (\Sigma_{\text{ou}} - \Sigma_t), \quad (2)$$

where γ_t is the expected firing rate, i.e., the expectation of the firing rate $g(u_t)$ computed based on the approximated filtering distribution $q_{\theta_t}(w_t)$ and x_t^ϵ denotes the presynaptic activation (see Materials and Methods). The first term in Equations (1) and (2) comes from the observations, which is why it scales with β . The update of the mean has the structure of 3-factor learning rule [Frémaux and Gerstner, 2016] with classical Hebbian factors, i.e., the pre-synaptic activation x_t^ϵ and the difference between observed and expected output $y_t - \gamma_t$, and the covariance Σ_t as third factor with a modulatory function, which has been linked to the computation of surprise [Liakoni et al., 2019]. The second term in Equations (1) and (2) comes from the hidden dynamics of the weights, which is why it is proportional to the inverse time constant τ_{ou}^{-1} . For $\beta = 0$ or $\tau_{\text{ou}} \rightarrow \infty$ the updates become independent of observations or the environmental dynamics respectively. In general, the covariance update in Assumed Density Filtering depends on the observations y_t . However, the combination of a Gaussian filtering distribution and an exponential gain function yields an update (Equation (2)) independent of the observations; an interesting similarity with the Kalman filter. Figure 1 (D) illustrate that the synaptic filter (red) successfully tracks the weights of the tutor network (black). In the Supplementary Information (Section 2), we show that the Synaptic Filter is a good approximation of the true filtering distribution.

In addition to the Synaptic Filter, we also derive and analyse the *Diagonalised Synaptic Filter*, an Assumed Density Filter based on a diagonal Gaussian (see Supplementary Information). The updates Equations (1) and (2) differ only in that the covariance matrix is diagonal and the off-diagonal updates are omitted.

2.2 MSE performance of the Synaptic Filter

The natural performance metric in filtering is the normalised Mean Square Error (MSE). It quantifies how closely the mean of the filtering distribution follows the weights of the tutor in Figure 1 (B). The MSE is defined by $\text{MSE} := d^{-1} \langle (w_t - \mu_t)^\top (w_t - \mu_t) \rangle_t$.

In the following, the MSE performance of the Synaptic Filter and the Diagonal Synaptic Filter are evaluated for a range of values for the determinism parameter β_0 with fixed dimension $d = 5$, and a range of dimensions d for a fixed value of the determinism $\beta_0 = 1$. Additionally, we compare the MSE of a gradient rule with optimised learning rate against the Synaptic Filter. The (expected) observation rate across dimensions is comparable because we chose the following scaling of $\beta \propto \beta_0 d^{-1}$. This cancels the linear scaling of the membrane potential variance with dimensionality of the model (Section 4.3.1 in Materials and Methods).

The MSE of both filtering models decreases as function of β_0 , as shown in fig. 2 (A). The Synaptic Filter (red line) performs slightly better than its diagonalised counterpart (black line). As β_0 increases the observations convey more information about the ground truth weights w_t , hence allowing for a more accurate estimation. At $\beta_0 = 0$, the MSE of all models is close to 1, a value that corresponds to the MSE of the prior.

For all four filtering models, the MSE increases as a function of the dimension d , as depicted in fig. 2 (B). The Diagonal Synaptic Filter has a slightly higher MSE. Increasing the dimension increases the difficulty of the filtering task because more weights compete for explaining the observation. For instance, at $d = 1$, observations can be attributed to $w_{t,0}$ uniquely but at $d = 2$ the information in the observations is distributed across both weights, $w_{t,0}$ and $w_{t,1}$. Thus, the scope of learning the weights via filtering is limited to the regime in which the observations convey enough information per dimension and per time.

As a benchmark for the Synaptic Filter, we use a gradient rule with a scalar learning rate¹ (see Section 4.1.3 in Material and Methods). As shown in Figure 2 (C), the MSE of the gradient rule (gray) is higher than the MSE (red) of the Synaptic Filter for a large range of learning rates. The MSE as a function of learning rate η exhibits a minimum when the combined effect of delayed learning and overshooting are minimal. Delayed learning occurs at low learning rates because the gradient does not converge before the ground truth weights change. At high learning rates, the update steps of the gradient rule are too large which leads to overshooting. In contrast, the Synaptic Filter optimally tunes the learning rate for each weight individually and at each moment based on the amount of information available in the data.

The Synaptic and the Diagonal Synaptic Filter have the highest MSE performance in low dimensions d and with high determinism β_0 . In this regime, observations contain the largest amount of information per weight. The existence of a limited high performance regime is a feature of the dynamic learning task, not a limitation of the filters. In particular, the Synaptic Filter outperforms the optimised gradient rule.

Despite the fact that the benefit of learning as filtering seems to be limited to a regime of low dimensionality and that biological neuron have up to 10^4 synaptic inputs, learning as filtering is a relevant framework for modelling synaptic plasticity. First, the sparsity of the neuronal code could limit the effective number of input dimensions at any point in time. Thus it is possible that learning in many brain systems takes place inside the learning regime. Secondly, the framework can be applied to high dimensions by aggregating the majority of synapses and inputs into a single variable, e.g., the bias $w_{t,0}$, and modelling only the remaining small group of synapses explicitly. It would be interesting to investigate whether multiple of these low-dimensional models yield good performance based on observations generated from a high-dimensional tutor network.

2.3 Predictive performance of the Synaptic Filter

The MSE is a function of the hidden variable, in our case the weights w_t , but it is not directly sensitive to whether the predictions of the network are useful, e.g., a low MSE does not necessitate poor predictions. The MSE is particularly limiting in presence of a model mismatch, when the student network in Figure 1 (B) makes false assumptions about how the tutor network maps inputs to outputs. If the model mismatch follows from a differences in weight dimensionality, the MSE cannot be defined. In the case of the nervous system, model mismatch is always present because sensory data is generated from physical processes which cannot be represented in detail by the brain.

This motivates the study of predictive performance which measures how well the student network predicts the emission of the next output spike y_t . The measure is equivalent with the Bayesian model evidence. Prediction performance can be evaluated even when the generative model is incorrect.

The Synaptic Filters make predictions on the basis of Bayesian regression (see Section 4.3 in Materials and Methods), a method that takes advantage of the filtering distribution to compute the posterior predictive distribution:

$$p(y_t | \mathcal{D}_t, x_t^\epsilon) = \int p(w_t | \mathcal{D}_t) p(y_t | w_t, x_t^\epsilon) dw_t. \quad (3)$$

Intuitively, the marginalisation over the weights w_t represents a weighted average over the probability of an output spike y_t . Thus, Bayesian regression takes advantage of parameter uncertainty encoded in the filtering distribution, including correlations between parameters when represented. To study the importance of including parameter uncertainty during prediction, we included a MAP prediction based on replacing the posterior in

¹This implicitly defines a Euclidean metric in weight space. A more general learning choice corresponds to a matrix-valued learning rate.




Figure 2: The Synaptic Filter has the best overall MSE and predictive performance. **(A)** The MSE of Synaptic Filter (red line) and the Diagonal Synaptic Filter (black line) decrease as the determinism β_0 increases. At $\beta_0 = 0$, the MSE corresponds to the equilibrium variance of the prior, $\sigma_{ou}^2 = 1$. The Diagonal Synaptic Filter performs slightly worse. **(B)** The MSE increases as a function of dimension. Again, the Diagonal Synaptic variants (black) perform worse. **(C)** The MSE of the Synaptic Filter (red) is lower than the MSE of a gradient learning rule (gray) for a range of learning rates η . The symbols indicates three combinations of determinism β_0 and dimension d . Consistently with **(A, B)**, the lowest MSE (black) is obtained at the lowest dimension and highest determinism, i.e. $\beta_0 = 1$ and $d = 5$. **(D)** The predictive performance is measured by the Bayes factor relative to a gradient learning (gray) rule with optimised learning rate. The Synaptic Filter (solid red line) and the Diagonal Synaptic Filter (solid black line) have superior performance and the Synaptic Filter performs best overall. Using Maximum a Posteriori (MAP, dashed lines), which does not include uncertainty during prediction, yields lower performance than using Bayesian Regression (BR, solid lines). **(E)** In the presence of model mismatch, $d_{tutor} = 5 \neq d$, the Synaptic Filter has the best overall performance. At high dimensions the optimal gradient rule performs equally well. The Diagonal Synaptic Filter and filters with MAP prediction overfit in the regime $d_{tutor} \ll d$. Dots and errors in **(A-E)** denote the mean and SEM from 100 simulations. The error bars take the gradient rule's SEM into account via the root mean square. See Section 4.3 in Materials and Methods for simulations details.

Equation (3) with a delta-function around its maximum. Based on the posterior predictive and the data \mathcal{D}_t , we compute the log evidence for four models \mathcal{M} , i.e., the Synaptic Filter, the Diagonal Synaptic Filter and their MAP-version (see Equation (12) in Materials and Methods), from time 0 to t : $\log p(\mathcal{D}_t | \mathcal{M})$. The model evidence is reported relative to a null model given by the baseline firing rate g_0 . The model evidence is a testing error in the sense that the prediction has not been informed by the current output y_t , i.e., the parameters of the weight distribution have not been updated. As a benchmark, we use a gradient rule with optimal learning

rate (Section 4.3 in Materials and Methods).

In the standard setting, i.e., without model mismatch, the Synaptic Filter and the Diagonal Synaptic Filter outperform the gradient rule, as shown in fig. 2 (D). The Synaptic Filter has the best performance, followed by the Diagonal Synaptic Filter and finally the gradient rule. The performance gain of the Synaptic Filter and Diagonal Synaptic Filter over the gradient method can be attributed to two factors. The first factor is their ability to estimate the value of the weight w_t of the tutor network better, consistent with the gain in MSE performance in fig. 2 (C). The second factor becomes evident by comparing the Bayesian regression prediction with the MAP prediction, i.e., including weight uncertainty via Bayesian regression improves performance.

Next, we wondered whether the introduction of a model mismatch would affect the model evidence of the Synaptic Filters compared to the benchmark. Specifically, we considered a tutor network with $d_{\text{tutor}} = 5$ and a student network with varying dimension $d \in (1, \dots, 15)$. Our rationale for this type of model mismatch was twofold. First, differences in model evidence can no longer be explained in terms of the MSE because the MSE is not defined; thus, the differences in model evidence can be more directly attributed to the use of Bayesian regression or MAP for prediction. Secondly, a central argument for using parameter uncertainty is that it helps to avoid overfitting. The risk of overfitting is high when the tutor has fewer dimensions than the student, i.e., $d_{\text{tutor}} < d$. In this case, the process that generates the data has less degrees of freedom than the model that aims at explaining the data.

Figure 2 (E) shows that the Synaptic Filters perform better than the optimised gradient rule when their dimension is smaller or equal to the tutor $d \leq d_{\text{tutor}}$. However, when the tutor has less dimensions than the student networks, $d > d_{\text{tutor}}$, the performance of the Diagonal Synaptic Filter drops while the Synaptic Filter continues to outperform the gradient method. This result can be explained by the fact that the Synaptic Filter includes correlations between weights while the Diagonal Synaptic Filter does not. In the case of the Synaptic Filter, weights become negatively correlated as multiple combinations of weights compete for explaining the data. This weight uncertainty cannot be reduced by additional observations because its fundamental cause is the model mismatch, i.e., extra dimensions of the student compared to the tutor. With Bayesian regression (BR) it is possible to include this uncertainty in the predictions. In contrast, not including weight correlations as in the case of the Synaptic Filter with Maximum a Posteriori (MAP) predictions leads to lower performance.

It might seem surprising that the gradient rule, which does not account for any parameter uncertainty, eventually outperforms all Synaptic Filters. The reason is that the optimisation of the loglikelihood with respect to the learning rate adds considerable power to the model. For instance, small learning rates effectively limit the risk of overfitting because weights do not converge quickly enough. Overall, the Synaptic Filter has higher predictive performance than the Diagonal Synaptic Filter and a gradient learning rule with optimal learning rate. This result was obtained in two conditions, with and without model mismatch. In particular, the model mismatch condition showed that the Synaptic Filter in combination with Bayesian regression is robust to overfitting because it accounts for the full posterior, including weight correlations.

2.4 The Synaptic Filter is consistent STDP

Spike-time dependent plasticity (STDP) refers to the property of a synapse to exhibit long-term potentiation (LTP) if the presynaptic spike comes before the postsynaptic spike and long-term depression (LTD) otherwise [Markram et al., 1997, Bi and Poo, 1998]. The results are usually depicted as STDP curve, i.e., the normalised change in synaptic weight as a function of the time interval between the pre and the postsynaptic spike.

Normative models of STDP have explained the LTP lobe, i.e., the weight change for pre- before postsynaptic spiking, in terms of causality reinforcement [Booij and tat Nguyen, 2005, Gütig and Sompolinsky, 2006, Urbanczik and Senn, 2014, Xu et al., 2013, Bohte and Mozer, 2005, Pfister et al., 2006]. The delay of the postsynaptic relative to the presynaptic spike represents the degree to which the occurrence of the postsynaptic spike can be attributed to the presynaptic activity trace and its decay resembles the shape of LTP lobe. In contrast, a post-before-pre spike pair has no causal relationship. Indeed, the computational rationale for the LTD lobe has remained a matter debate with proposals including the regulation of the postsynaptic firing rate and temporal locality [Pfister et al., 2006].

We wondered whether the Synaptic Filter could reproduce the STDP curve, in particular the LTD part. Our rationale was that postsynaptic spiking predominately affects the bias $w_{t,0}$, which represents the neuronal excitability. From this perspective, STDP is a secondary (differential) effect, i.e., changes in the synaptic weight account for the prediction error left unexplained by the adjustment of bias. Immediately after the occurrence of a postsynaptic spike, the expected firing rate is increased, leading to more LTD (Equation (1)). Thus, the time scale of the LTD lobe corresponds to the time scale of the transition probability of the prior $\tau_{\text{ou,bias}}$. In the simulations, we assumed $\tau_{\text{ou,bias}} = \tau_{\text{m}}$ but set all other transition time scales to values much larger than the duration of the experiment (Section 4.4 Materials and Methods).

To study the effect of the bias on the STDP curve, we applied a STDP protocol with one spike pair to three versions of the Synaptic Filter, i.e., a single synapse without bias and the Diagonal Synaptic Filter and Synaptic Filter with bias. To avoid effects from transients, the protocol was applied after the bias had reached its steady

state. In contrast to biological experiments with up to 60 spike pairs, we simulated a single spike pair to avoid complications from saturation effects and induction times.

In all three experiment, the resulting STDP curve shows an exponentially-shaped LTP lobe while the LTD lobe occurs only when the bias is included, as shown in Figure 3 (A). The LTP lobe resembles the exponential shape of the presynaptic activation because when the postsynaptic spike occurs, the weight update is proportional to the current amplitude of the presynaptic trace. The LTD lobe is present when the bias is included (black and red lines). Without the bias (gray line), the LTD part of the STDP curve is independent of the spike timing. Moreover, the bias lowers the amplitude of LTP. Both observations, LTD lobe and lower LTP, are caused by the modulation of the expected firing rate γ_t due to the bias $w_{t,0}$. The bias acts as low-pass filter of the postsynaptic spike train. Its value is maximal immediately after the occurrence of a postsynaptic spike at t_{post} and relaxes back to its equilibrium afterwards. The faster the pre follows the postsynaptic spike, the larger the value of the expected firing rate $\gamma_{t_{\text{pre}}}$ when the presynaptic spike occurs. Since LTD is proportional to γ_t , shorter intervals between pre and postsynaptic spikes lead to more LTD, in correspondence with biological STDP experiments. The dynamics of the variables of the Synaptic Filter are shown in detail the Supplementary Information.

The Synaptic Filter and Diagonal Synaptic Filter is consistent with the STDP. The appearance of the LTD lobe is contingent on inclusion of the bias. Indeed, since the underlying mechanism does not require the inclusion of uncertainty, a similar result could be obtained in a learning as optimisation framework as well, e.g., as the post-only term in the expansion of a generic update function [Gerstner and Kistler, 2002]. The fact that only a single synapse with bias was considered does not impair generality because additional unstimulated synapses do not affect the result. From an experimental perspective, the proposed mechanism can be tested by comparing the time scale of the negative lobe with the time scale of the excitability of a neuron, which is controlled by the bias.

2.5 The Synaptic Filter predicts spike-timing dependent changes of the variance

So far, we have assumed that the posterior mean μ_t can be measured experimentally as average EPSP amplitude. From here on, we assume the sampling hypothesis to be true, i.e., that the posterior variance σ_t^2 corresponds to the EPSP variance [Aitchison and Latham, 2015]. We wondered whether, in this case, the Synaptic Filter predicts spike-timing dependent changes in the EPSP variance.

Studying the same three conditions as in the previous section, we found that the variance decreases for all conditions and all pre-post timings, as shown in Figure 3 (B). In a Bayesian framework this is expected because inputs spikes, which represent informative data, decrease uncertainty. Interestingly, the reduction depends on the spike-timing. For a single synapse without bias (gray line), the effect is weak and confined to the causal pairings, i.e., $t_{\text{pre}} < t_{\text{post}}$. Including the bias (black and red lines) increases the amplitude of the variance change. Moreover, it adds a qualitatively new feature: a negative lobe in the regime $t_{\text{pre}} < t_{\text{post}}$. The underlying mechanism is the same as in the case of the LTD lobe of the mean (Figure 3 (A)). Changes in the synaptic weight and in the bias modulate the expected firing rate γ_t . In the models with bias, a postsynaptic spike increases the expected firing rate temporally and, hence, the potential for variance reduction when a presynaptic spiking occurs close in time. When both spikes coincide $t_{\text{pre}} = t_{\text{post}}$ the reduction is maximal because the temporally increased bias and the presynaptic activation increase the expected firing rate superlinearly. Thus, with the sampling hypothesis, the Synaptic Filter makes the novel prediction of spikes-timing dependent changes of the EPSP variance.

2.6 The Synaptic Filter explains heterosynaptic plasticity

From a hebbian perspective on plasticity, it is required that the presynaptic neuron’s activation takes part in postsynaptic neuron’s activation. Heterosynaptic plasticity contradicts a purely hebbian view on learning because plasticity occurs without presynaptic activation [Abraham and Goddard, 1983, Royer and Paré, 2003]. For example, LTP induction at hippocampal synapses leads to LTD at synapses that did not receive stimulation [Lynch et al., 1977]. It has been argued that the role of heterosynaptic plasticity is complementary to homosynaptic, hebbian plasticity, which can destabilize neuronal dynamics through run-away weights [Bailey et al., 2000, Chistiakova et al., 2015].

We wondered whether the Synaptic Filter is consistent with heterosynaptic plasticity. Our starting point was that heterosynaptic plasticity could be linked to the explaining-away effect in Bayesian reasoning. Explaining-away occurs when one infers the posterior over multiple causes for an observation. When additional observations provide evidence for only one of the causes, the competing causes are “explained away”. For example, hearing a triggered alarm (an observation) is best explained by a burglar. However, upon learning that an earthquake occurred when the alarm was set off, the posterior probability for the burglar becomes much lower.

In the spirit of explaining-away, we designed a preconditioning protocol to set up two synapses as competing causes for observations, i.e., the postsynaptic activity. The preconditioning protocol consists of synchronous




Figure 3: The Synaptic Filter exhibits STDP of the mean and variance. **(A)** The change of the mean of the filtering distribution as a function of the temporal difference between a pre-post spike pair. For a single weight (excluding the bias, gray line), the Synaptic Filter produces only the LTP lobe ($t_{\text{pre}} < t_{\text{post}}$) while the LTD lobe ($t_{\text{post}} < t_{\text{pre}}$) is independent of spike-timing. Inclusion of the bias, either with off-diagonal covariance elements (red line) or without (black line) yields the LTD lobe; and the magnitude of LTP decreases. **(B)** For the same protocol and learning rules, the variance σ^2 of the weight exhibits a spike-timing dependent decrease. When the bias (black and red lines) is included, the change in variance resembles a symmetrised LTD lobe, i.e., it scales as the inverse of $|t_{\text{pre}} - t_{\text{post}}|$. Without the bias, the amplitude of change is reduced and the dependence on spike-timing disappears for $t_{\text{post}} < t_{\text{pre}}$. See Section 4.4 in Materials and Methods for simulations details.

spikes trains at both inputs without postsynaptic spiking. In a second step, a STDP protocol was applied to the first synapse but plasticity was reported from both. The weight change at the first and second synapses is our prediction for homo- and heterosynaptic plasticity respectively. Both protocols are shown in Figure 4 **(A)**. We obtained the homo- and heterosynaptic STDP curves with and without preconditioning. The effect of preconditioning is illustrated in Figure 4 **(B)**: the equilibrium weight distribution, which is the initial condition of the STDP-step, becomes negatively correlated. We simulate a 3-dimensional Synaptic Filter with two synapses and bias. To test whether correlations between weights are important for heterosynaptic plasticity, we also include the Diagonal Synaptic Filter. The same time constants and initial conditions are assumed as in the STDP experiment (Section 4.4 Materials and Methods).

The homosynaptic STDP curve (black lines) appear robustly in all experiments, as shown in Figure 4 **(C, D)**, i.e. with and without preconditioning (dashed vs solid lines) and in both models, the Synaptic Filter and Diagonal Synaptic Filter. Preconditioning lowers the overall amplitude of the STDP curve, which was to be expected because presynaptic activity reduces the variance which acts as learning rate. Only a single experiment yields the heterosynaptic STDP curve: the Synaptic Filter with preconditioning (solid red line), shown in Figure 4 **(D)**. The heterosynaptic curve has the same shape as the homosynaptic STDP curve but with opposite sign. In contrast, the Diagonal Synaptic Filter exhibits no heterosynaptic STDP, i.e., the red curves in Figure 4 **(C)** are flat; and the Synaptic Filter without preconditioning has a flat heterosynaptic STDP curve (dashed red line, **(D)**) as well. These results make clear that the mechanism of heterosynaptic plasticity in Synaptic Filter requires weight correlations. The Diagonal Synaptic Filter cannot represent weight correlations, which is why it never exhibits heterosynaptic plasticity. The Synaptic Filter, in contrast, can represent weight correlations but they have to be induced by the preconditioning protocol (because we made the idealised assumption that the equilibrium distribution has strictly uncorrelated weights). The correlations are encoded in the covariance matrix Σ_t as off-diagonal elements. The covariance matrix acts as inverse metric of the parameter space. It scales the update of the mean weight via $\Sigma_t x_t^\epsilon$ [Surace et al., 2020]. Thus, activity at one of the weights can lead to plasticity at correlated weights. Therefore the Synaptic Filter exhibits heterosynaptic plasticity only in combination with the preconditioning protocol.

The observation that homo- and heterosynaptic STDP curves have opposite signs is explained by the negativity of the off-diagonal entries in the covariance matrix. From a mathematical perspective, the sign of the updates of the off-diagonal elements are negative when correlations are present. This follows from using non-negative inputs (see Section 4.2 Material and Methods). Intuitively, the negative correlation between two weights (with same-sign inputs) encodes how much they can explain-away each other.

Next, we wondered whether the negative correlation between homo- and heterosynaptic plasticity was consistent with experimental data. We quantified their relation by first plotting the values of the homo- and heterosynaptic STDP curves against each other and subsequently fitting a line, shown in Figure 4 **(E)**. The negative slope means that homosynaptic LTP is correlated with heterosynaptic LTD. The amplitude of heterosynaptic plasticity is around three quarters of the amplitude of homosynaptic plasticity. While the value of the slope depends on the number of spikes in the preconditioning protocol and other model parameters, the negativity of the slope is a robust feature of the Synaptic Filter caused by negativity of the weight correlations. A similar relation between homo- and heterosynaptic plasticity was reported in projections from basoateral amygdala (BLA) to intercalated cells (ITC) of the amygdala [Royer and Paré, 2003]. The authors used extra-

cellular low- and high-frequency stimulation to induce LTD and LTP respectively. They associated the induced weight change in the stimulated connection with homosynaptic plasticity and weight changes at other connections with heterosynaptic plasticity. Their main result aggregates plasticity results from multiple recorded ITC cells in a single figure, replotted in Figure 4 (F). The data confirms a robust negative correlation between homo- and heterosynaptic plasticity, consistent with the prediction of the Synaptic Filter. Points in Figure 4 (E, F) are comparable in the sense that they originate from the variability in the induction mechanism. In the case of the Synaptic Filter, only one parameter, the spike-timing differed between points. In a biological plasticity experiment, the number of sources of variability is much higher, including somatic properties, initial conditions of synapses, speed of signal transmission and effectiveness of extracellular stimulation. On the premise that the aggregate outcome of this variability modulates the effectivity of biological plasticity in a similar way as spike-timing in the stimulated experiment, Figure 4 (E, F) provides evidence for the Synaptic Filter. Moreover, the Synaptic Filter explains the inverse correlation between homo- and heterosynaptic plasticity in terms of the explaining-away effect. The inverse correlation has also been found in Hippocampus [Lynch et al., 1977] and more generally, see [Chistiakova and Volgushev, 2009] for a review.

3 Discussion

In this study, we showcased the framework of learning as filtering through the Synaptic Filter, an Assumed Density Filter for the weights in a spiking network. The main advantage of learning as filtering is that it accounts for the dynamics of the environment and weight uncertainty in a mathematically principled way. In a dynamic learning task, the Synaptic Filter outperforms a gradient rule with optimised learning rate in weight space. In combination with Bayesian regression, the Synaptic Filter also has better prediction performance. The representation of weight correlations proved particularly important to prevent overfitting in the presence of a model mismatch. The relevance of the Synaptic Filter to biological plasticity is threefold. It exhibits STDP of the mean weight, including the negative lobe; it predicts STDP of the weight variance; and based on weight correlations, it predicts heterosynaptic LTD and homosynaptic LTP consistently with experimental evidence. Thus, the Synaptic Filter combines computational benefits with biological insights into plasticity.

The framework learning as filtering can be used to derive additional learning rules. Here, we considered the simple case of a Gaussian weight distribution, an OU-process as prior and an exponential gain function in combination with spike-based observations. Alternative choices yield new learning rules. Parametrising the weight distribution through the binomial model of stochastic release represents an exciting possibility to study dynamics of EPSP variability and quantal parameter plasticity. The case of log-normally distributed weights has been studied in discrete time [Aitchison and Latham, 2015]. From a biological perspective, log-normal or binomial models have the advantage of obeying Dale’s law. However, this advantage comes at the cost of intractability, the requirement of additional approximation or more complicated update-rules. To avoid these complications, we chose Gaussian synapses in this work. Another option is to study gated plasticity through a hierarchical weight distribution in which an additional hidden variable infers whether a synapse should be plastic or not. Moreover, the framework can encompass different types of observations, for example continuous rates instead of spikes. Closely related to the observed variable is the choice of gain function. While the exponential offers simplicity, the sigmoidal and soft-max yield analytically tractable learning rules under additional approximations. Thus, learning as filtering offers a rich set of options to study learning and synaptic learning in particular.

The generalisation of the single neuron analysis to the case of a recurrent neuronal network is straightforward as long as the output neurons’ activities are conditionally independent of each other. In a recurrent neuronal network with only visible neurons, this condition is satisfied because past spiking affects the current output spikes only through the presynaptic kernels in the membrane potential. Formally, the joint probability distribution of output spikes conditioned on past spiking must factorise. Thus, the Synaptic Filter can be used to gain insights into learning dynamic weight distributions not only in single neuron but also in the more complex setting of a recurrent neuronal network model.

For the derivation of the learning rules, we make the assumption that the output neuron receives an external and neuron-specific supervision signal. Recent studies have addressed the question of how such a signal could be computed in biological networks [Sacramento et al., 2018] and in continuous time [Scellier and Bengio, 2017]. In our model, the supervision signal takes the form of spikes of the output neuron. This assumption does not exclude the possibility that biological neurons receive one type of spike as supervision signal to guide plasticity, and generate another type of spike to make predictions [Urbanczik and Senn, 2014, Schiess et al., 2016]. Experimental findings in cerebellum and cortex are compatible with this idea. Indeed, complex and simple spikes in Purkinjee cells, and bursts and normal firing in cortical pyramidal neurons play distinct roles for plasticity [Yang and Lisberger, 2014, Jacob et al., 2012]. Therefore, the assumption of a continuously provided, event based supervision signal does not impair the biological relevance of the Synaptic Filter.

The Synaptic Filter represents correlations between weights. From a biological perspective, this suggest

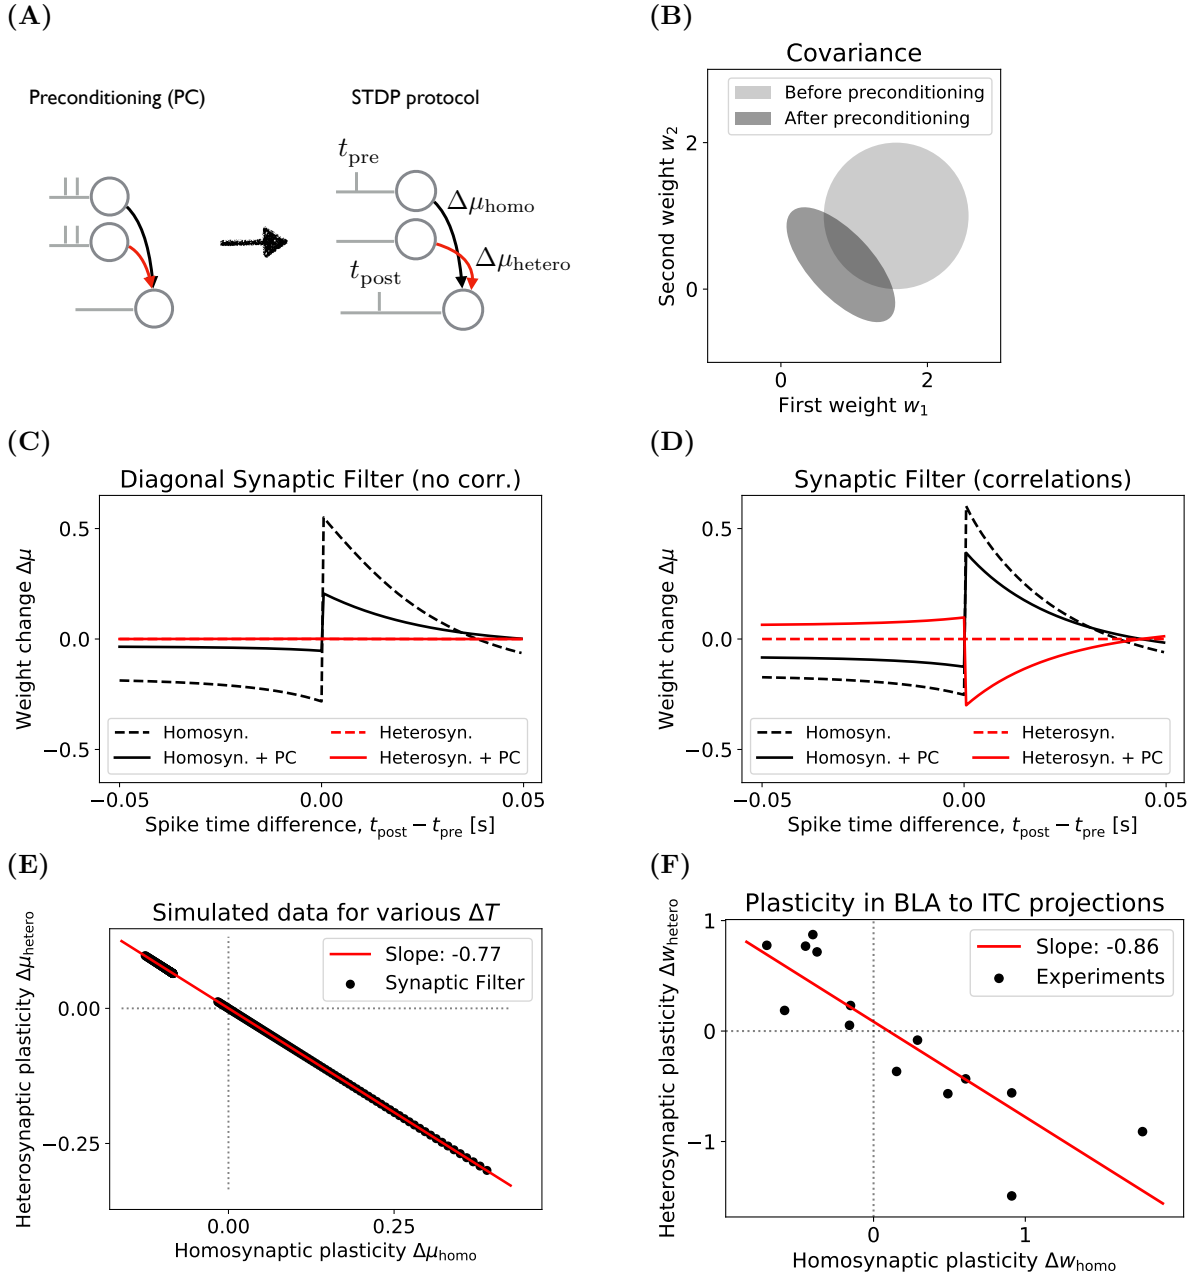


Figure 4: The Synaptic Filter explains experimentally observed anticorrelation of homo- and heterosynaptic plasticity. **(A)** Two inputs drive a neuron. During (optional) preconditioning (PC), two synchronous input spikes are delivered. Changes in the weights in the first and second weight in response to a STDP protocol are reported as homosynaptic (black) and heterosynaptic (red) plasticity respectively. See Section 4.4 in Materials and Methods for more details. **(B)** The effect of PC on the equilibrium weight distribution, visualised by contours, of the Synaptic Filter (light gray). After PC, the weight distribution (dark gray) has lower mean and weights become anticorrelated. **(C)** The Diagonal Synaptic Filter exhibits homosynaptic STDP (black lines) but no heterosynaptic plasticity (red lines). PC reduces plasticity (solid black line). **(D)** The Synaptic Filter exhibits homo- and heterosynaptic plasticity (solid black and red lines) after PC. Without PC (dashed lines), the Synaptic Filter behaves like its diagonal counterpart shown in **C**. **(E)** The Synaptic Filter predicts that homo- and heterosynaptic plasticity are anticorrelated. The x - and y -locations of the black points correspond to the solid black and red lines in **(C)**. **(F)** Anticorrelated homo- and heterosynaptic plasticity was found in synaptic projections from BLA to ITC neurons, Figure redrawn from [Royer and Paré, 2003].

two directions for future research. First, how can correlations between synaptic weights be implemented by biological synapses? On the premise, that EPSP samples represent weight uncertainty, non-linear summation at dendrites is a potential candidate. However, non-linear dendritic summation would be confined to spatially close synapses, e.g., ones located on the same dendritic branch. Thus, this mechanism cannot account for the full covariance matrix which scales as $\mathcal{O}(d^2)$ with the number of synapses d . One possibility to address this is to approximate the covariance matrix with a limited number off-diagonals. Another approximation aggregates the effect of spatially distant synapses on the membrane potential in a single aggregation variable, e.g., the bias. Secondly, the Synaptic Filter was derived under the assumption of positive presynaptic inputs while the sign of the weights could either be positive or negative. As a consequence, weight correlations are never positive. Based

on the negative weight correlations the Synaptic Filter could explain the negative correlations between homo- and heterosynaptic plasticity (shown in Figure 4). As an extension of this result, it would be interesting to generalise the Synaptic Filter to the case of positive and negative inputs, representing excitatory or inhibitory neurons. We hypothesise that a generalised Synaptic Filter would make the prediction of positively correlated homo- and heterosynaptic plasticity between synapses from inhibitory and excitatory neurons.

Because uncertainty controls the speed of learning the Synaptic Filter can in combination with the sampling-hypothesis link synaptic variability to synapse specific metaplasticity, which has been observed experimentally [Chistiakova and Volgushev, 2009]. The Synaptic Filter predicts that synaptic variability and learning speed reduce upon presynaptic stimulation but relax back to maximal value on the time scale of the OU-prior. Indeed, consistent with an OU time scale of hours, plasticity experiments have shown that LTP saturates temporally but recovers within hours [Abraham and Bear, 1996].

The presented work is closely related to the Know-Thy-Neighbour (KTN) theory [Pfister et al., 2009]. It assumes that synapses solve the problem of estimating the presynaptic membrane potential from the arriving spike train within the filtering framework. Consequently, it links the dynamics of short-term synaptic depression to the mean and variance updates of the filtering distribution. The Synaptic Filter and the KTN theory formalise plasticity via Assumed Density Filtering with an Ornstein-Uhlenbeck prior. At any point in time, the synapses encode a posterior distribution over the hidden variable given past observations, i.e., membrane potential and presynaptic spikes in the KTN model and ground truth weight with pre- and postsynaptic spikes in our model. In both models the classically defined synaptic efficacies, the averaged postsynaptic responses, corresponds to the mean of the posterior; and the variance plays the role of learning rate in the update equations. Our novel contributions are the extension of the KTN framework to multiple and potentially correlated hidden variables and a more complex emission process, i.e., the emissions are generated by the sum of the hidden variables weighted by the presynaptic trace. The KTN model is formally equivalent to the one-dimensional Synaptic Filter when the bias is the only hidden variable. On the level of the biological interpretation, KTN focuses on short-term plasticity while our work makes a connection to experiments in the context of long-term plasticity.

A previous study has addressed the computational role of synaptic uncertainty [Kappel et al., 2015]. The authors propose that spine turnover implements samples from the posterior distribution over synaptic weights via Langevin sampling. Their work differs from ours because they consider a static inference task (bottom right in Figure 1 (A)), not filtering. One consequence of the static nature of their task is that the online version of their learning rule includes a fixed data set size as external parameter. Compared to previous learning rules in the context of filtering [Aitchison and Latham, 2015], we make four additional contributions. First, we connect the learning task to the rich literature of filtering. In particular, this facilitates a simple, rigorous and continuous-time treatment. Secondly, we go beyond the assumption of a diagonal Gaussian Assumed Density Filtering by including weight correlations; and show that their importance for filtering performance. Thirdly, we show that the filtering distribution can be used in combination with Bayesian regression to improve predictive performance, a more relevant performance measure for learning than the MSE, which is computed in weight space. Finally, based on the spiking, continuous-time analysis, the Synaptic Filter recovers the phenomenon of spike-timing dependent plasticity, i.e., the mean synaptic increases if the postsynaptic spike follows the presynaptic spikes closely, and decreases if the spike order is reversed. Moreover, the Synaptic Filter predicts spike-timing dependent changes of the EPSP variability. Finally, it explains the negative correlation between homo- and heterosynaptic plasticity in terms of the Bayesian explaining-away effect.

Overall this article provides evidence that learning as filtering is a serious candidate for a computational principle underlying plasticity and provides testable predictions.

4 Materials and Methods

4.1 The generative model and learning rules

In this section, we define the filtering problem in terms of a generative model for plasticity. In addition, the update equations of the learning models are introduced, i.e., the Synaptic Filter, the Diagonal Synaptic Filter and the gradient learning rule.

4.1.1 The generative model

Learning as static optimisation has the goal of finding a parameter value that minimizes a cost function of given training data. Here, we consider a different framework for learning: learning as filtering. The goal of filtering is to continuously compute the probability distribution over dynamic, hidden variables based on continuously emitted observations. In contrast to optimisation, filtering includes time and parameter uncertainty in a principled manner on the level of the task.

A generative model specifies an observation process that relates the hidden parameters to observations and a transition probability that represents assumptions about the dynamics of the hidden parameters. To apply

the framework of learning as filtering to synaptic plasticity, we consider the following generative model.

The fundamental assumption is that d time dependent parameters, the synaptic weights $w_t \in \mathbb{R}^d$, govern how input spikes x_t give rise to output spikes y_t via the observation process $p(y_t|x_{0:t}, w_t)$ (explained in more detail below). The mapping between inputs and outputs is not static but changes due to the dynamics of w_t . For the transition probability, we assume that the hidden weights evolve according to an OrnsteinUhlenbeck (OU) process with time scale τ_{ou} , mean $\mu_{\text{ou}} = 0$ and diagonal covariance matrix $\Sigma_{\text{ou}} = \sigma_{\text{ou}}^2 \mathbb{1}$ (with non-zero elements) $\sigma_{\text{ou}}^2 = 1$:

$$dw_t = \tau_{\text{ou}}^{-1}(\mu_{\text{ou}} - w_t)dt + \sqrt{2\sigma_{\text{ou}}^2\tau_{\text{ou}}^{-1}}\mathbb{1}dV_t, \quad (4)$$

where V_t is a d -dimensional Wiener process. The process in Equation (4) can be represented as Gaussian transition probability: $p(w_t|w_{t-dt}) = \mathcal{N}(w_t - w_{t-dt} + \tau_{\text{ou}}^{-1}(\mu_{\text{ou}} - w_{t-dt})dt, 2\sigma_{\text{ou}}^2\tau_{\text{ou}}^{-1}dt)$.

For the observations, we assume a stochastic leaky-integrate-and-fire neuron, also called Spike-Response Model [Gerstner and Kistler, 2002]. The output spikes y_t are generated stochastically from a membrane potential u_t via an inhomogeneous Poisson point process. Thus, the output spikes represent a sum of delta distributions with spike times $\{t_j^f\}$: $y_t = \sum_f \delta(t - t_j^f)$. To connect the membrane potential to the firing rate of the Poisson process, we choose an exponential gain function:

$$y_t \sim \text{PoissonProcess}(g(u_t)) \quad (5)$$

$$g(u_t) = g_0 \exp(\beta u_t). \quad (6)$$

The determinism parameter β controls how strongly changes in the membrane affect the firing rate, and g_0 is the baseline firing rate. An exponential gain function represents a neuron close to the onset of firing but excludes saturation effects of the firing rate at high values of the membrane potential. The exponential gain function has been established as a phenomenological model of neocortical pyramidal neurons [Jolivet et al., 2006]. The membrane potential is a leaky integrator, with time constant $\tau_m = 25$, of the weighted sum of input spikes x_t . The leaky integration is represented by an exponential kernel $\epsilon_t = e^{-t/\tau_m}\Theta(t)$ where $\Theta(\cdot)$ represents the Heavyside function:

$$u_t = (w^\top x * \epsilon)_t \approx w_t^\top (x * \epsilon)_t, \quad (7)$$

The approximation in Equation (7) is valid in the regime $\tau_m \ll \tau_{\text{ou}}$, i.e., that the membrane dynamics is much faster than the dynamics of the ground truth weights assumed by the generative model. This assumption simplifies the generative model by casting it as a Markov process. Otherwise, the current observation y_t would depend on the entire history of hidden weights through the low-pass filtered membrane potential. The current observation does, however, depend on the history of input spikes via the convolutions $x_t^\epsilon := (x * \epsilon)_t$. This type of history dependence does not complicate the analysis because it can be straightforwardly taken into account in the spiking probability: $p(y_t|x_{0:t}, w_t)$ can be replaced by $p(y_t|x_t^\epsilon, w_t)$. Moreover, it has the biological interpretation of a presynaptic trace.

The notion of a bias can be included in the generative model by setting one of the inputs to one permanently. A bias stabilises the performance simulations and yields interesting biological insights. Thus, we adopt the convention that weight $w_{t,0}$ represents the bias and the input with index $i = 0$ is permanently set to $x_{t,0}^\epsilon = 1$. The generative model remains d -dimensional in terms of hidden variables but has only $d - 1$ presynaptic inputs.

4.1.2 The Synaptic Filter: update equations and prediction

The goal of learning as filtering is to compute the distribution over the hidden weights $p(w_t|\mathcal{D}_t)$ given all previously observed input and output spikes $\mathcal{D}_t := \{x_t, y_t\}_{\tau=0}^t$. On a formal level, the Markovian structure of the generative model, enables a recursive solution of the filtering problem:

$$p(w_t|\mathcal{D}_t) \propto p(y_t|x_t^\epsilon, w_t) \int p(w_t|w_{t-dt})p(w_{t-dt}|\mathcal{D}_{0:t-dt})dw_{t-dt}. \quad (8)$$

The Kushner-Stratonovich Equation [Kushner, 1967] gives a formal solution for all moments of the filtering distribution. However, for most generative models the solution is intractable because of the closure problem, i.e., the evolution of lower moments depends on higher moments of the filtering distribution. One way to address the closure problem is Assumed Density Filtering. The central idea is to replace the exact filtering distribution with a proposal distribution q parameterised by θ_t :

$$p(w_t|\mathcal{D}_t) \approx q_{\theta_t}(w_t). \quad (9)$$

When substituting the approximation Equation (9) into the right-hand side of Equation (8), the resulting posterior will generally not lay in the family of the proposal distribution q_θ . Therefore, one has to decide how to best approximate the result with a member of the proposal family [Sugiyama et al., 2012].

Here, we derive the *Synaptic Filter*, an approximate solution based on a Gaussian proposal density with mean $\mu_t \in \mathbb{R}^d$ and covariance matrix $\Sigma_t \in \mathbb{R}^{d \times d}$. The results for the *Diagonal Synaptic Filter* are identical with exception that off-diagonal elements of Σ_t are set to zero. The evolution of the distribution parameters $\theta_t = (\mu_t, \Sigma_t)$ can be computed from the Kushner-Stratonovich Equation. To remain in the Gaussian family, higher moments are omitted (see Supplementary Information). For the generative model specified above, the resulting update equations for μ_t and Σ_t are Equations (1) and (2).

The expected firing rate γ_t is a central quantity in the Synaptic Filter. It represents the predicted firing rate based on the filtering distribution:

$$\gamma_t := \int_w q_{\theta_t}(w_t)g(u_t)dw_t = g_0 \exp(\beta\mu_t^\top x_t^\epsilon + \frac{1}{2}\beta^2(x_t^\epsilon)^\top \Sigma_t x_t^\epsilon). \quad (10)$$

The expected firing rate depends not only on the mean of the filtering distribution but also on the covariance. As expected from a convex gain function, the covariance increases the expected firing rate compared to a scenario where only the mean is taken into account. The expected firing rate enters the computation of the error signal $y_t - \gamma_t$ in the update Equation (1) of the mean and controls the reduction of the covariance in Equation (2). In addition, the expected firing rate appears naturally when combining the Synaptic Filter with Bayesian regression. Bayesian regression is a method for making predictions in the presence of parameter uncertainty. The central idea is to marginalise over the parameters, as in Equation (3). For the Synaptic Filter, Equation (3) can be evaluated, by rewriting the probability of an output spike $p(y_t|x_t^\epsilon, w_t)$ as Bernoulli probability. The probability that a output spike occurs in the infinitesimal interval dt , indicated by $dN_t := y_t dt \in \{0, 1\}$ is:

$$\langle p(dN_t = 1|x_t^\epsilon, w_t) \rangle = \langle g(u_t)dt \rangle = \gamma_t dt, \quad (11)$$

where we used Equation (10) in the last step. From normalisation over both states of dN_t and converting the Bernoulli probability back to a point emission process, we obtain the posterior predictive distribution of the Synaptic Filter:

$$p(y_t|x_t^\epsilon, \theta_t) = \text{PoissonProcess}(\gamma_t). \quad (12)$$

4.1.3 Gradient learning rule

As a performance benchmark, we use the gradient learning rule. Assuming updates are proportional to the gradient of the log output probability yields:

$$\dot{w}_t^{\text{ML}} = \eta \beta x_t^\epsilon \beta (y_t - g_t^{\text{ML}}), \quad (13)$$

where η is the learning rate parameter and $g_t^{\text{ML}} = g_0 \exp(\beta(w_t^{\text{ML}})^\top x_t^\epsilon)$. We did not absorb β in the learning rate η to make the values of η more comparable to values of the variance in the Bayesian learning rule and to use the same scaling of β with the dimension as in the Synaptic Filters (see Section 4.3 in Materials and Methods).

4.2 The weight correlations of the Synaptic Filter are mostly negative

The weight correlations in the Synaptic Filter are represented by the off-diagonal elements of the covariance matrix Σ_t . In the 2-dimensional case and for positive inputs ($x_t^\epsilon \geq 0$) these elements $\Sigma_{t,i \neq j}$ are always negative. This follows from the following two observations. First, the change of weight correlations is negative when the initial weight distribution is diagonal, i.e., $\Sigma_0 = \sigma_0^2 \mathbb{1}$. Secondly, a negatively correlated weight distribution cannot evolve into a positively correlated weight distribution without assuming a diagonal form in between.

The change of the covariance is given by Equation (2). Omitting the temporal index for clarity and assuming $i \neq j$, the update of an off-diagonal element is:

$$\dot{\Sigma}_{ij} = -\gamma(\Sigma x^\epsilon)_i(\Sigma x^\epsilon)_j - 2\tau_{\text{ou}}^{-1}\Sigma_{ij}, \quad (14)$$

where we used that $(\Sigma_{\text{ou}})_{ij} = 0$. For the initial condition given by a diagonal covariance matrix, this expression simplifies to:

$$\dot{\Sigma}_{ij} = -\gamma \Sigma_{ii} x_i^\epsilon \Sigma_{jj} x_j^\epsilon. \quad (15)$$

Since all factors in this expression are positive but the overall sign is negative, an initially diagonal weight distribution can only evolve towards a negatively correlated distribution. Because in 2 dimensions a transition from a negatively correlated to a positively correlated weight distribution is not possible without a diagonal state in between, positive correlations cannot occur. Conditions under which this result holds for $d > 2$ are discussed in the Supplementary Information Section 6.

4.3 Computational performance: hyperparameters and simulation details

Here, we describe how the computational performance discussed in Sections 2.2 and 2.3 in Results was evaluated. In Section 4.3.1, the scaling of the membrane potential with dimensions is introduced. The following sections describe how the predictive performance and the model mismatch were implemented and the technical details of the simulations.

4.3.1 Scaling of the membrane potential with dimensions

The performance of the learning rules is reported as a function of the dimension d of the generative model. Varying d , influences the statistics of the firing rate $g(u_t)$ and, hence, the amount of information available for learning. In our model we scale the determinism β with the number of inputs such that the expected firing rate $\langle g(u_t) \rangle$ becomes (approximately) independent of the number of inputs d . In this section, the expected value is taken with respect to the statistics of the generative model, not the filtering distribution.

To compute $\langle g(u_t) \rangle$, we first approximate the membrane potential statistics as Gaussian. With $\langle w_t \rangle = \mu_{\text{ou}} = 0$, the mean and variance of the membrane are:

$$\langle u_t \rangle = \langle w_t^\top \rangle \langle x_t^\epsilon \rangle = 0 \quad (16)$$

$$\text{var}[u_t^2] = \sum_{i=0}^{d-1} \text{var}[w_{t,i} x_{t,i}^\epsilon] = \sigma_{\text{ou}}^2 (1 + (d-1) \langle (x_{t,i}^\epsilon)^2 \rangle) \approx d \frac{\sigma_{\text{ou}}^2 \tau_m \nu_0}{2} \quad (17)$$

where we used that for two independent random variables a, b with a being zero-mean: $\text{var}[ab] = \text{var}[a] \mathbb{E}[b^2]$ and dropped the special treatment of the bias in the final step. We also assumed that all presynaptic neurons fire at the same rate $\langle x_{t,i} \rangle = \nu_0$. Approximating the membrane statistics by their first and second moment, given by Equations (16) and (17), we used the Gaussian expectation of an exponential (see Supplementary Information) to obtain the expected firing rate in the generative model as a function of the dimension d :

$$\langle g(u_t) \rangle = g_0 \exp\left(\frac{1}{4} d \beta^2 \sigma_{\text{ou}}^2 \tau_m \nu_0\right). \quad (18)$$

To ensure a comparable observation rate across dimensions, we remove the dependence on d in Equation (18) by making the determinism β parameter a function of the dimension:

$$\beta \rightarrow \beta(d) = c \beta_0 d^{-1/2}. \quad (19)$$

The proportionality constant $c \beta_0$ is split into a factor of order $\beta_0 = \mathcal{O}(1)$ that is varied in the experiments and a constant c that ensures that the neuron's firing rate only rarely exceeds $g_{\text{max}} = 50\text{Hz}$. Specifically, the 5-sigma environment of the membrane statistics, $\sigma_u^{(5)} := \langle u_t \rangle \pm 5 \sqrt{\text{var}[u_t^2]}$ (given by Equations (16) and (17)), defines the condition for c :

$$g_{\text{max}} \stackrel{!}{=} g_0 \exp\left(\beta(d) \sigma_u^{(5)}\right) \Rightarrow c \equiv \frac{\log\left(\frac{g_{\text{max}}}{g_0}\right)}{5 \sqrt{\frac{1}{2} \sigma_{\text{ou}}^2 \tau_m \nu_0}} \quad (20)$$

In the simulations, we vary the proportionality constant β_0 , which we refer to as determinism parameter for simplicity.

4.3.2 Measuring predictive performance via the model evidence

In Section 2.3, the predictive performance of a model \mathcal{M} , e.g., the Synaptic Filter, is measured in terms of the evidence $p(\mathcal{M}|\mathcal{D}_t)$. The log evidence is directly related to the predictive distribution evaluated on the data \mathcal{D}_t . Assuming a flat model prior, we have:

$$\log p(\mathcal{D}_t|\mathcal{M}) = \int_0^t \left(\log \int p(y_\tau | x_\tau^\epsilon, w_\tau) p(w_\tau | \mathcal{D}, \mathcal{M}) dw_\tau \right) d\tau. \quad (21)$$

The model dependent parameter distribution $p(w_\tau | \mathcal{D}, \mathcal{M})$ corresponds to the filtering distribution for the Synaptic Filter and to a delta-distribution centered on the current estimator $w_\tau^{(\text{ML})}$ for the gradient rule (see Section 4.1.3). Thus, the argument of the log in Equation (21) is the predictive distribution evaluated on the data. In practise, we use the same discretisation as in Equation (11) to evaluate Equation (21) based on the size of the simulation time steps Δt . For instance, the log Bayes factor between the Synaptic Filter and a null-model \mathcal{M}_0 based on the baseline firing rate g_0 is:

$$\log \frac{p(\mathcal{D}_t|\mathcal{M}_{\text{SF}})}{p(\mathcal{D}_t|\mathcal{M}_0)} \approx \sum_{k=0}^{T/\Delta t} \log \left(\frac{\gamma_{t_k}}{g_0} \right)^{\Delta N_{t_k}} + \log \left(\frac{1 - \gamma_{t_k} \Delta t}{1 - g_0 \Delta t} \right)^{1 - \Delta N_{t_k}}, \quad (22)$$

where $t_k := k \Delta t$ and $\Delta N_{t_k} = 1$ if a spike occurred in the interval $[t_k, t_k + \Delta t]$ and zero otherwise.

4.3.3 Predictive performance of the optimised gradient rule

In Section 2.3, we use the log predictive distribution to evaluate learning rules. Because the gradient rules does not include parameter uncertainty, this metric is equivalent to the loglikelihood. For the optimisation of the loglikelihood with respect to the learning rate, we obtained the loglikelihood performance of the gradient rule for 11 log-spaced values of the learning rate in an interval $\eta \in [0.05, 2]$, which contains the optimal learning rate, and interpolated with a 3rd order polynomial (see Supplementary Information). Based on the fit, we selected the maximal loglikelihood value. The same procedure was used to compute the SEM of the loglikelihood.

The intuition behind the fact that $[0.05, 2]$ contains the optimal learning rate is that the learning rate replaces the variance in the update of the Diagonal Synaptic Filter, as eminent when comparing Equation (1) and Equation (13). The variance values of the Diagonal Synaptic Filter are generally below their equilibrium value $\sigma_t^2 \leq \sigma_{\text{ou}}^2 = 1$ but well above 0.05. Because the Diagonal Synaptic Filter adjusts the variance of each synapse optimally, it is expected that the optimal value of η has the same order of magnitude. As shown in the Supplementary Information, the optimal learning rate was indeed located in the interval $[0.05, 2]$.

4.3.4 Model mismatch

Bayesian regression takes advantage of the filtering distribution to compute the posterior predictive distribution. A central advantage of the posterior predictive is that it alleviates overfitting. In Section 2.3, we introduce a model mismatch between the model that generates the observations, i.e., the tutor network in Figure 1 (B), on the one hand, and, on the other hand, the generative model on which the learning rules are based. We study model mismatch in Section 2.3 as follows.

When the generative model includes excess dimensions compared to the tutor, $d > d_{\text{tutor}}$, we provide the student network with additional $d - d_{\text{tutor}}$ Poissonian input with firing rate ν_0 that have no effect on the generation of the observations. When the generative model has less dimensions compared to the tutor, $d < d_{\text{tutor}}$, tutor and generative model share the first d input neurons but the tutor included additional $d_{\text{tutor}} - d$ ones, which are used to generate observations. Both, generative model and tutor, share the same value of the determinism parameter $\beta = \beta(d_{\text{tutor}})$. In the simulations, we kept the dimension of the tutor at a constant value $d_{\text{tutor}} = 5$ and varied the dimension of the generative model.

4.3.5 Hyperparameters and simulated time

The membrane time constant and baseline firing rate are $\tau_m = 25\text{ms}$ and $g_0 = 1\text{Hz}$. The $d - 1$ input spikes at dimensions $0 < i < d$ were drawn from Poisson neurons with $\nu_0 = 40\text{Hz}$. Thus, the expected synaptic activation is $\langle x_t^e \rangle = \tau_m \nu_0 = 1$, i.e., on average each input neuron emits a spike per membrane time constant. The first dimension denotes the bias and does not receive spiking input.

For the MSE simulations discussed in Section 2.2, we used 100 simulations with an OU time scale of $\tau_{\text{ou}} = 100\text{s}$ and duration $T_{\text{sim}} = 10\tau_{\text{ou}}$. Shorter time scales yielded overall higher and less dynamic MSE curves for the hyperparameters studied.

For the evaluation of the predictive performance, reported in Section 2.3, a shorter time scale was used $\tau_{\text{ou}} = 5\text{s}$. Because the predictive distribution is a noisier metric than the MSE and our computational resources were limited, we reduced the τ_{ou} in order to be able to run more τ_{ou} -periods and, consequently, reduce the SEM. We chose $T_{\text{sim}} = 200\tau_{\text{ou}}$. As before, 100 simulations we used. A second rationale for reducing the time scale was that the optimisation of the learning rate η was computationally demanding.

4.3.6 Simulation details: time steps and error tolerance

For the Synaptic Filters and gradient rule in Section 2.3 we used a time step of $\Delta t = 0.5\text{ms}$ and for the particle filter $\Delta t = 1\text{ms}$. These time steps were a compromise between minimization of the frequency with which discretisation errors occurred and the requirement to average over sufficiently many τ_{ou} -period to improve the errors. When a discretisation in the firing rate occurred, i.e., $g(u_t)\Delta t > 1$, we corrected it by enforcing a value of 1. From all conditions, this problem occurred most frequently for $\beta_0 = 2$ (irrespectively of the dimension). However, even for $\beta_0 = 2$, only 10^{-4} of the time steps needed a correction, which we regarded within error tolerance. In the case of the particle filter, discretisation errors can lead to negative particle weights, which we corrected by setting them to 0 and renormalising all particle weights afterwards. Again, high values of β_0 caused the highest frequency of discretisation errors. For $\beta_0 = 2$ and $\beta_0 = 1.67$ negative particle weights occurred in 0.5% and 0.1% of the time steps.

4.3.7 Initial conditions

The initial value of the tutor's weights was $w_{t=0} = \mu_{\text{ou}}$. The distributional parameters were initialised as $\Sigma_{t=0} = \Sigma_{\text{ou}}$ and $\mu_{t=0} \sim \mathcal{N}(\mu_{\text{ou}}, \Sigma_{\text{ou}})$. For the particle filter, the initial positions of the particles, indexed with

l , were drawn from the prior as well: $v_{t=0}^{(l)} \sim \mathcal{N}(\mu_{\text{ou}}, \Sigma_{\text{ou}})$. After initialisation, a burn-in period of τ_{ou} was simulated.

4.4 Biological predictions: parameters and technical details

In this section, we specify the values of the hyperparameters, protocols, initial conditions and simulation parameters used in the simulated STDP experiments in Sections 2.4 to 2.6.

4.4.1 Hyperparameters

For the simulated experiments, the membrane time constant and baseline firing rate were set to their standard values, $\tau_{\text{m}} = 25\text{ms}$ and $g_0 = 1\text{Hz}$ and the determinism parameter was always set to $\beta \equiv 1$ (unlike in the performance simulations in Figure 2). Scaling β with dimensions as we did in the study of computational performance would have made it difficult to compare the weight change observed in the single synapse model ($d = 1$), the synapse with bias model ($d = 2$) and the heterosynaptic experiments ($d = 3$).

For the simulated experiments, we assume that the bias $w_{t,0}$ changes on the same time scale as the membrane potential, $\tau_{\text{ou,bias}} = \tau_{\text{m}} = 25\text{ms}$. The time scale of the weights and weight correlations is set to $\tau_{\text{ou}} = 10000\text{s}$. This implies that the correlations between the bias and other weights decay on the order of τ_{m} (see Supplementary Information).

With the exception of the bias, the equilibrium values of the transition probability are set to $\mu_{\text{ou}} = 0$ for the weights and $\sigma_{\text{ou}}^2 = 1$ for the diagonal covariance elements and off-diagonal covariance elements are set to zero: $\Sigma_{\text{ou},i \neq j} = 0$. The bias is treated differently because it represents the neuronal excitability, as explained in the main text. To prevent run-away dynamics of the bias in the absence of spikes, we chose $\mu_{\text{ou},0} = 1$. The prior variance $\sigma_{\text{ou},0}^2$ determines how strongly the bias changes in response to input and output spikes. For the STDP experiments in Figure 3, we chose $\sigma_{\text{ou}}^2 = 2$ and for the heterosynaptic experiments in Figure 4, we chose $\sigma_{\text{ou}}^2 = 1$.

4.4.2 Protocols and initial conditions

The STDP protocol consisted of pre- and postsynaptic spikes with 200 different values for the delay (shown on the x -axis of the STDP curve). The preconditioning protocol consists of a presynaptic spike pair with 5ms delay simultaneously applied to both presynaptic inputs. Prior to applying any protocol, the Synaptic Filter was simulated without any input or output spikes for $T_{\text{wait}} = 6\tau_{\text{m}}$ such that the mean and variance values of the bias converge to their equilibria. The same waiting time was simulated after the preconditioning protocol. The simulated STDP curve was computed based on the value of the weight directly before the application of the STDP protocol and the value of the weight after $2T_{\text{wait}}$. The initial conditions for the distribution parameters of the Synaptic Filter were chosen $\mu_{t=0,i} = 1$ with $i \in (0, \dots, d-1)$ for the weights and $\Sigma_{t=0} = \mathbb{1}\sigma_0^2$ with $\sigma_0^2 = 1$ for the covariance.

4.4.3 Technical details: simulations and fits

We solved the ODEs of the Synaptic Filters with the Euler method. The time step was $\Delta t = 0.1\text{ms}$ for the STDP experiments presented in Sections 2.4 and 2.5. Since the preconditioning protocol induces sharp decreases of the variance, a time step of $\Delta t = 0.01\text{ms}$ was used for the simulations in Section 2.6 to ensure that the variance remained positive.

The slopes reported in Figure 4 (**E**, **F**) correspond to a linear fit with a least-squares objective and two free parameters, slope and offset. The data shown in Figure 4 (**F**) were extracted manually.

5 Supplementary Information

The Supplementary Information addresses three main questions. Section 5.1 asks whether the sampling hypothesis is compatible with MSE performance. In Section 5.2, we analyse whether the Synaptic Filter is a faithful solution filtering problem. Thirdly, in Section 5.3, we answer how the update equations of the Synaptic Filter (Equations (1) and (2) in the Main Text, here Equations (58) and (59)) are derived.

Additional short sections provide additional explanations for the results in the Main Text. Section 5.4 details how the predictive performance was evaluated for the gradient rule and Section 5.5 shows how the variables of the Synaptic Filter evolve during the plasticity protocols. In the final Section, we discuss when weight correlations are negative in the Synaptic Filter with more than two dimensions.

5.1 The sampling hypothesis is compatible with MSE performance

The sampling hypothesis states that EPSP are samples from the filtering distribution [Aitchison and Latham, 2015]. We wondered whether the sampling hypothesis could be used to inspire a version of the Synaptic Filter that maintains good filtering performance while at the same time computing the expected firing rate through EPSP samples.

Based on two heuristic modifications of the update equations of the Synaptic Filter Equations (58) and (59), we derive the Sampling Synaptic Filter. The first modification is that the expected firing rate γ_t is replaced with the sampling based firing rate $\gamma_t^s := g(u_t^s)$. The sampling based membrane potential u_t^s is directly motivated by the sampling hypothesis, i.e., it is a low-pass filtered sum of spike-triggered samples from all synapses. The second modification aims at slowing down learning to the sampling time scale τ_s . This modification is motivated by the empirical observation that good filtering performance requires the sampling based firing rate γ_t^s to change on the same time scale as other spike-related terms in the update Equations (58) and (59). In the following, both modifications are discussed in more detail.

The sampling based firing rate $\gamma_t^s := g(u_t^s)$ is computed on the basis of a sampling based membrane potential u_t^s , which is modelled as leaky-integrator of spike-triggered EPSPs w_t^s :

$$u_t^s = ((w^s)^\top x * \epsilon)_t. \quad (23)$$

The sum $(w^s)^\top x$ takes a value different from zero only when a spike arrives at synapse $0 < i < d$ and a sample from the marginal filtering distribution is drawn:

$$w_{t,i}^s \sim \mathcal{N}(\mu_{t,i}, \Sigma_{t,ii}). \quad (24)$$

Since no two spikes occur in the same moment, the Sampling Synaptic Filter does not include correlations between weights. The bias is treated as a special case because it does not receive spikes and, thus, the idea of spike-triggered sampling cannot be applied. We make the assumption that Sampling Synaptic Filter includes the bias in the sampling based firing rate γ_t^s in the same way as the Synaptic Filter, i.e., we adopt the following convention for the first component of the sum:

$$(w_0^s x_0 * \epsilon)_t \equiv \mu_{t,0} + \frac{1}{2} \beta \sigma_{t,00}^2. \quad (25)$$

With Equation (25) and in the case of $d = 1$, the expected firing rate of the Synaptic Filter and sampling based firing rate in the Sampling Synaptic Filter are identical: $\gamma_t^s = \gamma_t$.

The second modification aims at slowing down learning in order to suppress fast feedback dynamics between weights and expected firing rate. Without this suppression, filtering performance deteriorates because the first modification prevents the fast feedback from working properly. Consider how the fast feedback works in the case of the Synaptic Filter. A large expected firing rate γ can reduce its own value indirectly by reducing the mean and variance values of all weights with active inputs through the update Equations (58) and (59). This reduction can be faster than the membrane time scale τ_m . In the Sampling Synaptic Filter, this fast feedback mechanism does not work anymore. The sampling based firing rate γ_t^s depends on samples of past filtering distributions and it cannot influence the value of these samples through the update Equations (58) and (59). The sampling based firing rate γ_t^s , responds to changes in the filtering distribution with a delay of τ_m . Thus, the weight updates can overshoot which leads to poor performance.

One way to address the broken fast feedback is by introducing a delay in the updates in Equations (58) and (59) such that the delayed response of γ_t^s to changes in the filtering distribution becomes irrelevant. Specifically, we consider spike related quantities as fast and replace them by their low-pass filtered version. With the exponential kernel α with *sampling time scale* τ_s , the update Equations (58) and (59) of the Sampling Synaptic Filter are defined as follows:

$$\dot{\mu}_t = \beta \Sigma_t ((\alpha * x^\epsilon y)_t - (\alpha * x^\epsilon \gamma^s)_t) - \tau_{\text{ou}}^{-1} (\mu_t - \mu_{\text{ou}}) \quad (26)$$

$$\dot{\Sigma}_t = -\beta^2 \Sigma_t (\alpha * x^\epsilon (x^\epsilon)^\top \gamma^s)_t \Sigma_t - 2\tau_{\text{ou}}^{-1} (\Sigma_t - \Sigma_{\text{ou}}). \quad (27)$$

In all simulations, we chose $\tau_s = 4\tau_m = 100\text{ms}$. Additionally, we test the Diagonal Sampling Synaptic Filter which is updated according to Equations (26) and (27) with the only difference that a diagonal covariance matrix is enforced. While both versions of the Sampling Synaptic Filter compute the sampling membrane based on the marginal filtering distribution, i.e., without considering correlations between samples, the Sampling Synaptic Filter includes the off-diagonal elements of the covariance matrix during learning while the Diagonal Sampling Synaptic Filter uses a diagonal covariance matrix.

Because the model of the membrane potential differs between the sampling filters and their deterministic counterparts, Equations (26) and (27) do not converge to Equations (58) and (59) in a non-trivial limiting case. Furthermore, the history of the weight samples w_t^s must be included in the data $\mathcal{D}_t^s := (w_{0:t-d}^s, \mathcal{D}_t)$. The key advantage of the Sampling Synaptic Filter over the Synaptic Filter is that it does not require the instantaneous

evaluation of the expected firing rate γ_t but instead uses a biologically more plausible expected firing rate γ_t^s inspired by the sampling hypothesis.

In the following, the MSE performance of the Synaptic Filter, the Sampling Synaptic Filter, the Diagonal Synaptic Filter and the Diagonal Sampling Synaptic Filter are evaluated for the same range of values for the determinism parameter β_0 and the dimension d as in the Main Text. The results for the Synaptic Filter and the Diagonal Synaptic Filter are replotted.

Taking into account the sampling hypothesis does not substantially impair MSE performance, as shown in Figure 5. The sampling filters behave similarly to their deterministic counterparts, including a small performance gain from using the covariance matrix (red lines) in the update equations. Still, the Synaptic Filter (red solid line) remains the best model overall.

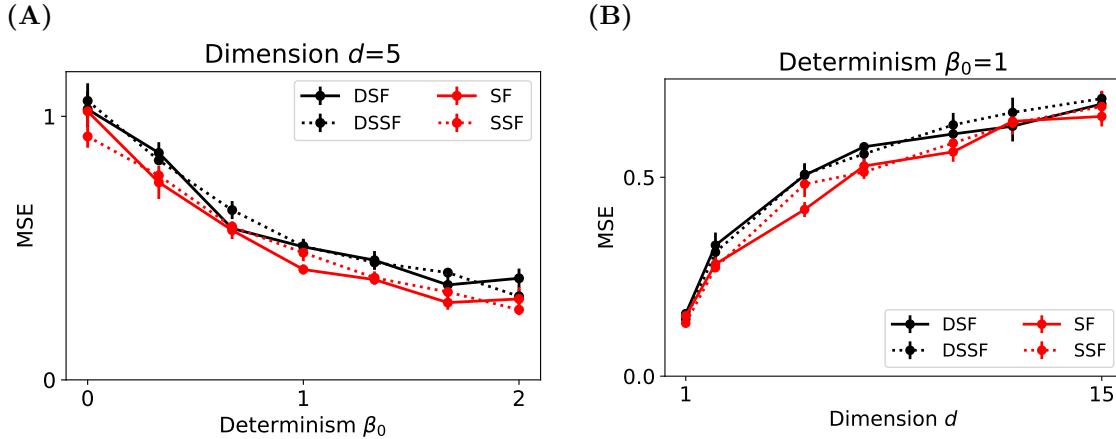


Figure 5: The Sampling Synaptic Filters have similar MSE performance to their deterministic counterparts for all values of the dimension d and determinism β_0 . **(A)** The MSE of Synaptic Filter (red solid), Sampling Synaptic Filter (red dashed) and their diagonal counterparts (black) decrease as the determinism β_0 increases. **(B)** For all filtering models, the MSE increases as a function of dimension. Again, the diagonal variants (black) perform worse. Dots and errors denote the mean and SEM from 100 simulations. The simulated time per run was $10\tau_{ou} = 1000s$.

5.2 The Synaptic Filter solves the filtering problem

5.2.1 The normalised moments of the Synaptic Filter correspond to the moments of the exact filtering distribution

The Synaptic Filter is an approximated solution to the filtering problem. Thus, we need to determine the level of accuracy of the approximation. Specifically, we check if the moments of the Synaptic Filter are consistent with the corresponding moments of the exact filtering distribution. We test a range of values for the determinism parameter β_0 with fixed dimension $d = 5$, and a range of dimensions d for a fixed value of the determinism $\beta_0 = 1$.

To compare the moments of the Synaptic Filter $q_{\theta_t}(w_t)$ with the moments of the exact filtering distribution $p(w_t|\mathcal{D}_t)$, we compute the normalised estimators $z^{(1)}$ and $z^{(2)}$ of the first and second moment respectively. If the moments of q converge to the first two moments of the exact distribution, the normalised estimators converge $z^{(1)} \rightarrow 0$ and $z^{(2)} \rightarrow 1$ (see Section 5.2.2). We use these conditions to measure the quality of the Synaptic Filters. Additionally, we obtain an approximation to the exact filtering distribution $p(w_t|\mathcal{D}_t)$ with a particle filter (see Section 5.2.3). The particle filter converges to the exact solution of the filtering problem in the limit of a large number of particles.

The results in Figure 6 **(A-D)** show that, for a range of dimensions d and values of the determinism β_0 , the normalised estimators of the Synaptic Filter (SF) confirm consistency between the approximated and the exact filtering distribution, i.e. $z_{SF}^{(1)} \approx 0$ and $z_{SF}^{(2)} \approx 1$. The Synaptic Filter (red solid line) and the Diagonal Synaptic Filter (black solid line) show similar performance at $d = 1$ (Figure 6 **(B, D)**) because in this case the covariance matrix is a scalar and, hence, both models are identical. However, at higher dimensions the Diagonal Synaptic Filter exhibits a small deviation from $z^{(1)} = 0$ and a strong positive deviation from $z^{(2)} = 1$. For the later, the deviation grows linearly with the determinism β_0 , as shown in Figure 6 **(C)**.

The reason is that due to its omission of correlations, the Diagonal Synaptic Filter underestimates the overall weight uncertainty, and hence overestimates the overall precision and $z^{(2)}$, which is proportional to the precision matrix, i.e., the inverse of the covariance matrix. Correlations arise from the likelihood term in the update Equation (59), which is proportional to β_0 . This explains the scaling of the deviation with β_0 . The superior performance of the Synaptic Filter shows that the off-diagonal elements of the covariance matrix are important to obtain a good approximation to the filtering distribution.

The estimators computed based on the particle filter (gray) are generally consistent with the exact distribution and with the Synaptic Filter, i.e., they satisfy $z_{\text{PF}}^{(1)} \rightarrow 0$ and $z_{\text{PF}}^{(2)} \rightarrow 1$. This was expected since particle filters are asymptotically exact in the limit of infinitely many particles. The systematic deviation $z_{\text{PF}}^{(2)} > 0$ in (Figure 6 (D)) arises because particle filters suffer from the curse of dimensionality, which leads to underestimates of the covariance in higher dimension.

The Sampling Synaptic Filter (red dashed line) and the Diagonal Sampling Synaptic Filter (black dashed line) estimate the second moment $z^{(2)}$ with similar accuracy as their counterparts without sampling (solid lines), as shown in Figure 6 (C, D). The Sampling Synaptic Filter performs well, i.e., $z_{\text{SSF}}^{(2)} \rightarrow 1$, while the Diagonal Sampling Synaptic Filter deviates strongly from $z^{(2)} = 1$. Both sampling filters perform well in terms of the first normalised moment, shown in Figure 6 (A, B). The fact that both sampling filters estimate the moments of the exact filtering distribution with comparable accuracy as their deterministic counterparts implies that the explicit inclusion of the sampling hypothesis in the updates does not impair filtering performance. For this conclusion to hold, we had to assume a sampling time scale τ_s that was many orders of magnitude smaller than the time scale of the weight evolution τ_{ou} (see Main Text).

The analysis of the first and second normalised moment estimators, $z^{(1)}$ and $z^{(2)}$, shows that the mean and covariance computed by the Synaptic Filter correspond closely to the mean and covariance of the exact filtering distribution. A particle filter solution to the filtering problem confirms this. The fact that the Diagonal Synaptic Filter and the Diagonal Sampling Synaptic Filter perform poorly shows that off-diagonal elements in the covariance matrix must be included to match the moments of the exact filtering distribution.

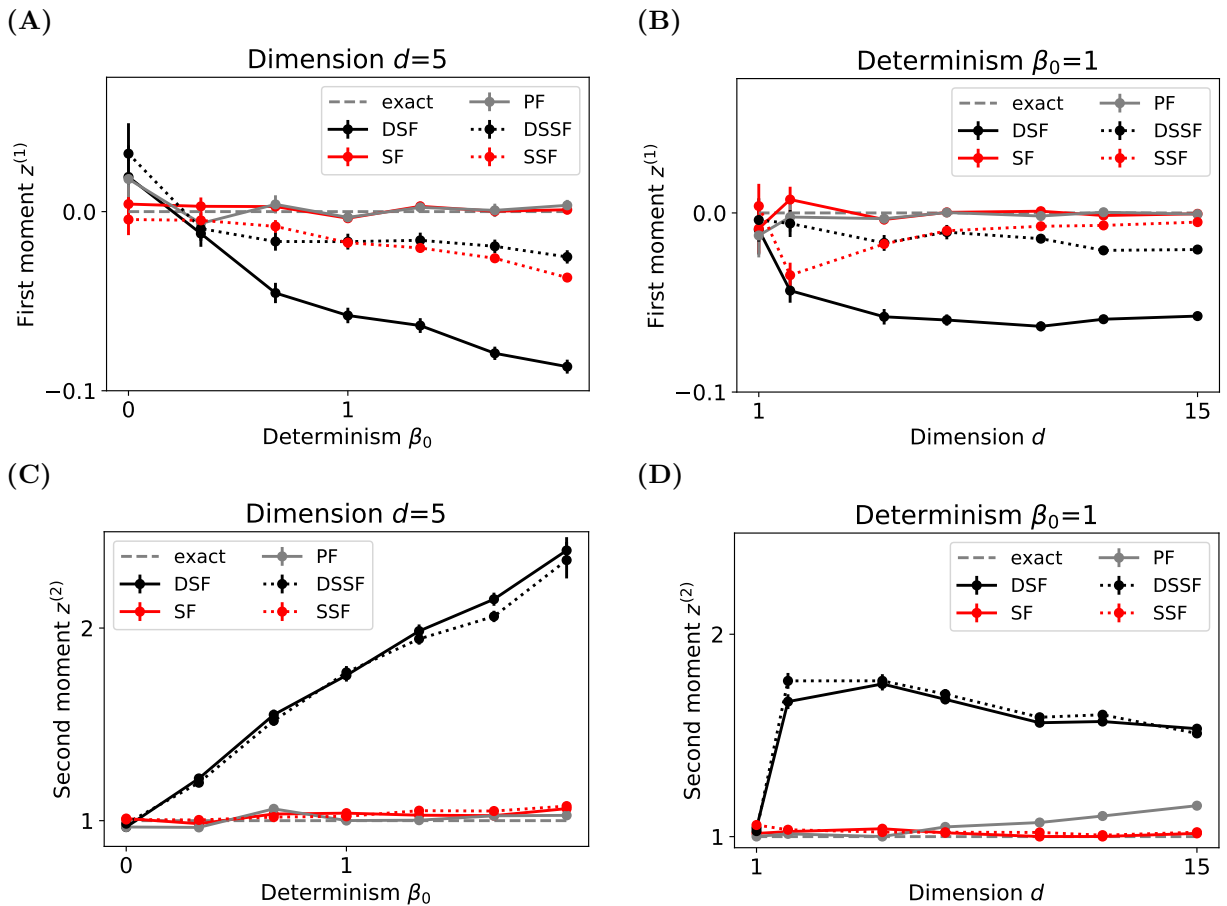


Figure 6: The first and second moments of the Synaptic Filter (SF, solid red line) match the corresponding moment of the exact filtering distribution, i.e. the normalised moments of the Synaptic Filter satisfy $z_{\text{SF}}^{(1)} \approx 0$ and $z_{\text{SF}}^{(2)} \approx 1$. (A) For $d = 5$ and a $0 \leq \beta_0 \leq 2$, the first normalised moment $z^{(1)}$ of the Synaptic Filter and the particle filter (PF, gray line) are close to 0 while the Sampling Synaptic Filter (SSF, solid black), Diagonal Synaptic Filter (DSF, dashed red line), and Diagonal Sampling Synaptic Filter (DSSF, dashed black line) deviate from 0. At $\beta_0 = 0$, all filtering distribution resemble the prior. As the determinism β_0 increases, the deviation increases as well. (B) For $\beta_0 = 1$ and $1 \leq d \leq 15$, the first normalised estimator $z^{(1)}$ of all models except for the Diagonal Synaptic Filter are close. The dimension d corresponds to the number of presynaptic inputs plus one (for the bias) so when $d = 1$, the Synaptic Filter and Sampling Synaptic Filter (red lines) are equivalent to their diagonalised (black). However, the diagonalised versions perform worse for $d > 2$. (C) The second normalised estimator $z^{(2)}$ of the Synaptic Filter, the Sampling Synaptic Filter and the particle filter are close to 1 while the diagonalised versions (black) overestimate $z^{(2)}$. The deviation is linear in β_0 . (D) As before, the $z^{(2)}$ is close to one for Synaptic Filter and Sampling Synaptic Filter and deviates for the diagonal models. The particle filter is consistent with $z^{(2)} = 1$ at low dimensions but deviates systematically for increasing d . Dots and error bars denote the mean and SEM from 100 simulations. The simulated time per run was $10\tau_{ou} = 1000s$.

5.2.2 Time averaged normalised moment estimators

The naive approach to quantify whether the moments of a proposal distribution, i.e., an approximation to the exact filtering distribution, match the moments of the exact filtering distribution, is computing the moments of both. However, this requires the potentially expensive computation of the moments of the exact filtering distribution. To avoid this problem, we use an implicit comparison via the normalised moments. In particular, we calculate the first two normalised moments $z^{(1)}$ and $z^{(2)}$, which form the basis of the analysis in Section 5.2.1.

Consider a proposal distribution $q_\theta(w_t)$ with mean μ_t and covariance Σ_t . The moments of the proposal can be compared to the respective moments μ_t^* and Σ_t^* of the exact filtering distribution $p(w_t|\mathcal{D}_t)$ by averaging over many realisations of the generative model, i.e., over the history of the hidden weight $w_{0:t}$ and data generated \mathcal{D}_t . To see this, we rewrite the distribution over the realisations of the generative model using the definition of conditional probabilities:

$$p(w_{0:t}, \mathcal{D}_t) = p(w_t, w_{0:t-dt} | \mathcal{D}_t) p(\mathcal{D}_t). \quad (28)$$

Now, we define the time dependent, normalised moments $z_t^{(1)}$ and $z_t^{(2)}$ as a function of the hidden weight w_t and the moments of the proposal:

$$z_t^{(1)} := d^{-1} \sum_{i=1}^d (\Sigma_t^{-\frac{1}{2}} (w_t - \mu_t))_i \quad (29)$$

$$z_t^{(2)} := d^{-1} (w_t - \mu_t)^\top \Sigma_t^{-1} (w_t - \mu_t). \quad (30)$$

Since Equations (29) and (30) do not depend on $w_{0:t-dt}$, the expectation with respect to realisations of the generative model (Equation (28)) is equivalent to the expectation with respect to the posterior and the data, i.e. $p(w_t|\mathcal{D}_t)p(\mathcal{D}_t)$:

$$\mathbb{E}[z_t^{(1)} | w_{0:t}, \mathcal{D}_t] = d^{-1} \sum_{i=1}^d \mathbb{E}[(\Sigma_t^{-\frac{1}{2}} (\mu_t^* - \mu_t))_i | \mathcal{D}_t] \quad (31)$$

$$\mathbb{E}[z_t^{(2)} | w_{0:t}, \mathcal{D}_t] = d^{-1} \text{Tr} (\mathbb{E}[\Sigma_t^{-1} (\mu_t^* - \mu_t) (\mu_t^* - \mu_t)^\top + \Sigma_t^{-1} \Sigma_t^* | \mathcal{D}_t]), \quad (32)$$

where we evaluated w_t using the mean and covariance of the exact posterior. To obtain Equation (32), we further used that the trace operator leaves a scalar quantity invariant and permits cyclic permutation. Note that the moments of the exact posterior and the proposal depend on the history of the data \mathcal{D}_t , which is why the expectation on the right-hand side of Equations (31) and (32) remains. However, in the special case that the proposal is equivalent to the exact posterior, we have $\mu_t^* = \mu_t$ and $\Sigma_t^{-1} \Sigma_t^* = \mathbb{1}$, and consequently:

$$\mathbb{E}[z_t^{(1)} | w_{0:t}, \mathcal{D}_t] = 0 \quad (33)$$

$$\mathbb{E}[z_t^{(2)} | w_{0:t}, \mathcal{D}_t] = 1. \quad (34)$$

Thus, Equations (31) and (32) can be compared to Equations (33) and (34) to test whether the moments of a proposal distribution match the moments of the exact posterior without explicitly calculating the latter.

The quantities shown in the Main Text are time averaged estimates of Equations (31) and (32) computed from 100 realisations of the generative model. The time average runs over 10 τ_{ou} time constants. The average over time is justified because the generative model is stationary, i.e., there is no explicit dependence on time. Indexing realisations by k , we define the normalised moments by:

$$z^{(1)} := \frac{1}{100} \sum_{k=1}^{100} \langle z_{t,k}^{(1)} \rangle_t \quad (35)$$

$$z^{(2)} := \frac{1}{100} \sum_{k=1}^{100} \langle z_{t,k}^{(2)} \rangle_t. \quad (36)$$

5.2.3 Particle filter

To compute an asymptotically correct approximation to the posterior distribution (without the approximation of the Assumed Density Filter) we use a particle filter with $L = 2^{13}$ particles. The update of the position $v^{(l)}$ and importance weight $a^{(l)}$ of the l^{th} particle are given by:

$$dv_t^{(l)} = \frac{\mu_{\text{ou}} - v_t^{(l)}}{\tau_{\text{ou}}} dt + \sqrt{2\sigma_{\text{ou}}^2 \tau_{\text{ou}}^{-1}} \mathbb{1} dV_t \quad (37)$$

$$da_t^{(l)} = a_t^{(l)} \left(\frac{g(\beta(v_t^{(l)})^\top x_t^\epsilon)}{\langle g_t \rangle} - 1 \right) (y_t - \langle g_t \rangle) dt, \quad (38)$$

where the average $\langle g_t \rangle = \sum_{l=1}^L a_t^{(l)} g(\beta(v_t^{(l)})^\top x_t^e)$ is computed based on the particle weights and their location. When the effective particle number $N_t = (\sum_{l=1}^L (a_t^{(l)})^2)^{-1} < \frac{3}{4}L$, we resample the particle location and set all weights to $1/L$. Additional details about particle filtering are found in the literature, e.g. [Doucet et al., 2000, Kutschireiter et al., 2020].

5.3 Derivation of the Synaptic Filter

The derivation of the Synaptic Filter extends the work of Pfister et al. [Pfister et al., 2009] to multidimensional hidden variables and a more complex observation process. The starting point of the derivation is the general framework of filtering with point observations and hidden diffusion dynamics. Then, the assumed density filter is introduced as a strategy to solve the filtering problem. The next two sections specify the generative model of the Synaptic Filter and show how it yields the update equations used in the Main Text.

5.3.1 Filtering with point process observations

Given the continuous time spiking observations $y_t = \sum_f \delta(t - t^{(f)})$ from hidden weights $w_t \in \mathbb{R}^d$, our goal is to derive the update equations for the parameters θ_t of the proposal distribution $q_{\theta_t}(w_t)$ of an assumed density filter. The generative model is specified in terms of a prior transition probability and point emission process. For the transition, we consider diffusion processes of the form:

$$dw_t = a(w_t)dt + b(w_t)dV_t, \quad (39)$$

where V_t is a d -dimensional Wiener process and $a \in \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $b \in \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ are deterministic functions. The observation $dN_t \in \{0, 1\}$ indicates whether a spike is present in the infinitesimal interval dt :

$$dN_t \sim \text{Poisson}(g_t(w_t)dt), \quad (40)$$

where the (deterministic) gain function $g_t(w_t)$ relates the hidden weights to observations.

This non-linear filtering problem has a general solution given by a Kushner-Stratonovic type of equation for point processes [Kushner, 1964, 1967, Brémaud, 1981]. Using the Laplacian $\mathcal{L}(\cdot) = a^\top \nabla(\cdot) + \frac{1}{2} \text{Tr}(bb^\top \nabla \nabla^\top(\cdot))$ all moments ϕ_t of the posterior obey the following formal solution:

$$d\langle \phi_t \rangle = \langle \mathcal{L}(\phi_t) \rangle dt + \frac{\text{cov}(\phi_t, g_t)}{\gamma_t} d\delta_t, \quad (41)$$

where the expectation $\langle \cdot \rangle$ is taken with respect to the filtering distribution and we introduced the error signal $d\delta_t = dN_t - \gamma_t dt$ based on the expected firing rate $\gamma_t := \langle g_t \rangle$.

Generally, Equation (41) is intractable due to the closure problem, i.e., the evolution of n -th moment depends on the $n + 1$ -th moment. For instance, for the choice $g_t \equiv w_t \in \mathbb{R}$, the evolution of the first moment ($\phi_t \equiv w_t$) depends on the variance $\text{cov}(w_t^2)$. However, when Equation (41) is used to compute the evolution of the variance ($\phi_t \equiv w_t^2$), third order terms appear: $\text{cov}(w_t^2, w_t)$.

5.3.2 Assumed density filter with Gaussian proposal

One strategy to apply Equation (41) is assumed density filtering (e.g. [Minka, 2001]). The central idea is to replace the exact filtering distribution $p(w_t|\mathcal{D}_t)$ with a more tractable proposal distribution $q_{\theta_t}(w_t)$ with parameters $\theta_t \in \mathcal{S} \subset \mathcal{R}^m$. The fact that the proposal distribution belongs to a parametric family, limits its degrees of freedom. Thus, (in the absence of degeneracy) the number of parameters m determines how many moments have to be computed in order to fully specify the evolution of the proposal distribution. The closure problem can be avoided.

The Synaptic Filter and the Diagonal Synaptic Filter are assumed density filters with a Gaussian proposal distribution:

$$q_{\theta_t}(w_t) := \mathcal{N}(w_t; \mu_t, \Sigma_t). \quad (42)$$

In the case of the Diagonal Synaptic Filter Σ_t contains only diagonal elements. The general notation Σ_t (with or without off-diagonal elements) allows for the simultaneous derivation of both filters. To relate the parameters of the proposal $\theta_t = (\mu_t, \Sigma_t)$, we follow a simple moment-matching strategy, i.e., we use Equation (41) to directly compute the updates for θ_t . In the case of Gaussian distributions, moment matching is optimal in the sense that it minimizes the Kullback-Leibler divergence: $\mathcal{D}_{KL}(p(w_t|\mathcal{D}_t)|q_{\theta_t}(w_t))$. However, the question of how to optimally project the exact filtering distribution $p(w_t|\mathcal{D}_t)$ onto the proposal is an active research area in information geometry [Sugiyama et al., 2012].

To compute the evolution of θ_t , we specify Equation (41) for the first two central moments μ_t and Σ_t ([Kutschireiter et al., 2020], Equations (101) and (102)):

$$d\mu_t = \langle a_t \rangle dt + \frac{\text{cov}(w_t, g_t)}{\gamma_t} d\delta_t \quad (43)$$

$$\begin{aligned} d\Sigma_t^2 &= (\text{cov}(a_t, w_t^\top) + \text{cov}(w_t, a_t^\top) + \langle b_t b_t^\top \rangle) dt \\ &+ \frac{\text{cov}(w_t w_t^\top, g_t) - \mu_t \text{cov}(w_t^\top, g_t) - \text{cov}(w_t, g_t) \mu_t^\top}{\gamma_t} d\delta_t \\ &- \frac{\text{cov}(w_t, g_t) \text{cov}(w_t^\top, g_t)}{\gamma_t^2} dN_t. \end{aligned} \quad (44)$$

Expectations $\langle \cdot \rangle$ are now evaluated with respect to the proposal density $q_{\theta_t}(w_t)$ rather than with respect to the exact filtering distribution $p(w_t | \mathcal{D}_t)$. For the Diagonal Synaptic Filter, one considers only the updates of the diagonal elements in Equation (44). In general, the updates of the diagonal elements can have complex dependencies on off-diagonal elements, e.g., if we had assumed a non-diagonal matrix b , the term bb^\top would have introduced such mixing of components. However, since b was assumed to be a diagonal matrix and w and a a vectors, no mixing occurs in Equation (44). This is why in our case the updates of diagonal covariance elements in Diagonal Synaptic Filter and the Synaptic Filter are identical.

5.3.3 OU-prior and exponential gain function

To evaluate the terms in the moment evolution Equations (43) and (44), we must make specific choices for the transition probability, often simply referred to as *prior* and the gain function g_t in the point emission process Equation (40).

In our work, the emission probability reflects the output spiking of a neuron with membrane potential:

$$u_t = w_t^\top (x * \epsilon)_t \quad (45)$$

where ϵ denotes the spike response kernel to given input spike trains $x_{t,i} = \sum_f \delta(t - t_i^{(f)})$. Since the zero-th weight $w_{t,0}$ acts as bias, we adopt the convention $(x_0 * \epsilon)_t = 1$. The assumption of an exponential gain function has analytical advantages and corresponds to a neuron close to the onset of a sigmoidal gain function:

$$g_t(w_t) \equiv g_0 \exp(\beta w_t^\top x_t^\epsilon) \equiv g(u_t). \quad (46)$$

The determinism parameter β can be absorbed in the units of x_t^ϵ and is omitted from the rest of the derivation. In addition, we drop the temporal index, since the right-hand side in Equations (43) and (44) is evaluated exclusively at time t . For brevity, we refer to the proposal distribution of the assumed density filter $q_{\theta_t}(w_t)$ as filtering distribution from here on.

For the transition probability of the hidden weights, we use an OU-process with equilibrium values μ_{ou} and $\Sigma_{\text{ou}} = \mathbb{1}\sigma_{\text{ou}}^2$, and relaxation time scale τ_{ou} :

$$a(w) \equiv -\tau_{\text{ou}}^{-1}(w - \mu_{\text{ou}}), \quad b(w) \equiv \sqrt{\frac{2\sigma_{\text{ou}}^2}{\tau_{\text{ou}}}} \mathbb{1}. \quad (47)$$

With Equations (46) and (47) and the update Equations (43) and (44) and the proposal density Equation (42), the assumed density filter is fully specified. The rest of derivation is dedicated towards the explicit computation of the terms in the update Equations (43) and (44).

5.3.4 Explicit computation of terms in the update equations

Prior

First, the terms related to the transition probability in Equations (43) and (44) are evaluated based on Equation (47):

$$\langle a(w) \rangle = a(\mu) \quad (48)$$

$$\langle bb^\top \rangle = 2\sigma_{\text{ou}}^2 \tau_{\text{ou}}^{-1} \mathbb{1} \quad (49)$$

$$\begin{aligned} \text{cov}(a, w^\top) &= \langle aw^\top \rangle - \langle a \rangle \langle w^\top \rangle \\ &= -\tau_{\text{ou}}^{-1} (\mu \mu^\top + \Sigma - \mu_{\text{ou}} \mu^\top) + \tau_{\text{ou}}^{-1} (\mu - \mu_{\text{ou}}) \mu^\top = -\tau_{\text{ou}}^{-1} \Sigma, \end{aligned} \quad (50)$$

For the stimulated STDP experiments in the Main Text, we consider a fast time scale τ_m for the bias and a slow time scale τ_{ou} for the remaining weights. The update of the covariance element Σ_{0i} , which represents the

correlations between the bias and the i^{th} weight, contains both time scales because of the first two terms on the left-hand side of Equation (44):

$$\text{cov}(a_i, w_0) + \text{cov}(w_i, a_0) = -(\tau_m^{-1} + \tau_{\text{ou}}^{-1})\Sigma_{0i}. \quad (51)$$

However, since the values of the time scales differ by nine orders of magnitude in the STDP simulations, the contribution of τ_{ou} can be safely ignored.

Observations

Next, we evaluate the terms in Equations (43) and (44) that depend on the gain function. We begin by showing that the membrane potential is Gaussian under the statistics of the filtering distribution. Then we evaluate the Gaussian expectation of the gain function. Finally, the expectation of the gain function is used to compute the covariance-terms in Equations (43) and (44).

The weighted sums of a Gaussian random variables yields another Gaussian random variable. With the input kernels x^ϵ given and $q_\theta(w) = \mathcal{N}(w; \mu, \Sigma)$, the membrane potential is therefore Gaussian with mean and covariance:

$$\mathbb{E}[u|\theta] = \mu^\top x^\epsilon := \bar{u} \quad (52)$$

$$\text{Var}[u|\theta] = (x^\epsilon)^\top \Sigma x^\epsilon := \sigma_u^2. \quad (53)$$

Thus, for any function $g(u)$, the expectation over the d -dimensional filtering distribution can be replaced by the 1-dimensional expectation over the membrane distribution. For the exponential gain function this yields:

$$\mathbb{E}[g|\theta] = g_0 \int e^u \mathcal{N}(u; \bar{u}, \sigma_u^2) du = g_0 \exp(\bar{u} + \frac{1}{2}\sigma_u^2) \underbrace{\int \mathcal{N}(u; \bar{u} + \sigma_u^2, \sigma_u^2) du}_{=1} := \gamma \quad (54)$$

where the second equality follows from completion of the square in the exponent of the Gaussian. The expected firing rate $\gamma := g_0 \exp(\bar{u} + \frac{1}{2}\sigma_u^2)$ plays a central role in the update equations and for making predictions with the filtering distribution in Bayesian regression.

To compute the covariance terms in Equations (43) and (44) we use the fact that the expectation over the filtering distribution commutes with derivatives with respect to x^ϵ :

$$\text{cov}[g, w] := \langle wg \rangle - \langle w \rangle \langle g \rangle = \langle \nabla_x g \rangle - \mu \gamma = \nabla_x \gamma - \mu \gamma = (\mu + \Sigma x^\epsilon) \gamma - \mu \gamma = \Sigma x^\epsilon \gamma \quad (55)$$

The result $\nabla_x \gamma = (\mu + \Sigma x^\epsilon) \gamma$ can be reused to compute the higher order covariance term in Equation (44):

$$\begin{aligned} \text{cov}[g, ww^\top] &:= \langle ww^\top g \rangle - \langle g \rangle \langle ww^\top \rangle \\ &= \nabla_x \nabla_x^\top \gamma - \gamma(\mu \mu^\top + \Sigma) \\ &= \nabla_x (\gamma(\mu + \Sigma x^\epsilon))^\top - \gamma(\mu \mu^\top + \Sigma) \\ &= \gamma(\mu + \Sigma x^\epsilon)(\mu + \Sigma x^\epsilon)^\top + \gamma \Sigma - \gamma(\mu \mu^\top + \Sigma) \\ &= \gamma \mu (\Sigma x^\epsilon)^\top + \gamma (\Sigma x^\epsilon) \mu^\top + \gamma (\Sigma x^\epsilon) (\Sigma x^\epsilon)^\top. \end{aligned} \quad (56)$$

Combining results to obtain update equations of the Synaptic Filter

With the results for the observation terms, i.e., Equations (55) and (56), the variance update Equation (44) simplifies considerably. Most terms related to gain function, i.e., terms proportional to $dt\delta$ and dN , cancel:

$$\begin{aligned} d\Sigma &\propto \frac{\text{cov}(ww^\top, g) - \mu \text{cov}(w^\top, g) - \text{cov}(w, g) \mu^\top}{\gamma} d\delta - \frac{\text{cov}(w, g) \text{cov}(w^\top, g)}{\gamma^2} dN \\ &= (\Sigma x^\epsilon) (\Sigma x^\epsilon)^\top d\delta - (\Sigma x^\epsilon) (\Sigma x^\epsilon)^\top dN = -\gamma (\Sigma x^\epsilon) (\Sigma x^\epsilon)^\top dt. \end{aligned} \quad (57)$$

We obtain the update equations for the Synaptic Filter and the Diagonal Synaptic Filter (for which Σ is a diagonal matrix), by substituting the results of prior and observations, i.e., Equations (48) to (50), (55) and (56), into the update Equations (43) and (44) for the moments. In addition, we use $y := dN/dt$ to switch to standard ODE-notation:

$$\dot{\mu} = \Sigma x^\epsilon (y - \gamma) + \tau_{\text{ou}}^{-1} (\mu_{\text{ou}} - \mu), \quad (58)$$

$$\dot{\Sigma} = -\gamma (\Sigma x^\epsilon) (\Sigma x^\epsilon)^\top + 2\tau_{\text{ou}}^{-1} (\Sigma_{\text{ou}} - \Sigma), \quad (59)$$

where $y := dN/dt$. In the Main Text, the determinism parameter β scales the input variable x^ϵ .

5.4 Evaluating the predictive of the gradient rule

To obtain the predictive performance of the gradient rule, we must maximise the loglikelihood as a function of the learning rate. This is challenging because the loglikelihood is noisy and due to limited computational resources, we could only evaluate few values of the learning rate.

Our strategy was to evaluate the loglikelihood at 11 log-spaced values for the learning rate in the interval $[0.05, 2]$. The choice of the interval is motivated in Main Text by noting that the learning rate replaces the variance in the Diagonal Synaptic Filter. Figure 7 shows that the interval contains the optimal learning rate indeed.

To maximise the loglikelihood with respect to the learning rate, we fit a 3rd order polynomial to a 7 point neighbourhood around the maximum and selected its maximum value of the loglikelihood. Compared to selected the maximal value of the simulated loglikelihoods, using the fit has two advantages. First, it reduces the risk of selecting a statistical outlier as the maximum loglikelihood, as shown in Figure 7 (A). Secondly, it compensates for the sparseness of our evaluation of the loglikelihood by interpolation, as seen in Figure 7 (B). The predictive performance benchmark shown in the Main Text corresponds to maximum value of the polynomial, obtained for each dimension separately.




Figure 7: Two typical examples of the polynomial fit to the loglikelihood of the gradient learning rule. (A) The simulation at $d = 3$ showcases the denoising effect of the fit. Compared to the outlier (central black dot), the maximal loglikelihood based on the fit is lower. (B) The example at $d = 8$ show that the optimal learning rate based on the fit shifts is smaller compared to the optimal learning rate obtained by selecting the simulation (black dot) with the highest loglikelihood value.

5.5 Variables of the Synaptic Filter during simulated biological experiments

The Synaptic Filter can explain the biological phenomena of STDP and the negative correlation between homo- and heterosynaptic plasticity. The goal of this section is to show how the interaction of the variables of the Synaptic Filter produces the aforementioned biological observations. In the following, we show the time series of the mean μ_t and covariance matrix Σ_t , alongside the simulated protocols.

The STDP curve shown in the Main Text arises from a series of simulations with varying spike-timings. Figure 8 (A) and (B) show Synaptic Filter variables for the cases $t_{\text{pre}} - t_{\text{post}} = \pm 10\text{ms}$, linked to the positive and negative STDP lobe respectively. The strength of potentiation in (A) is related to the amount of presynaptic activation present when the postsynaptic spike occurs. The reason is that the product $y_t x_t^\epsilon$ is proportional to the update of the mean in Equation (58). The strength of depression in (B) depends on amplitude of the mean of the bias when the presynaptic spike occurs. The mean bias increases the expected firing rate γ_t , which modulates the update via the term $-\gamma_t x_t^\epsilon$. Thus the timescale of the negative lobe is directly associated with the prior timescale of the bias $\tau_{\text{ou,bias}}$, which controls how quickly the bias returns to its equilibrium value. From this mechanistic perspective on the STDP protocol, it becomes clear that the negative lobe dependence on the presence of the bias. However, this result do not qualitatively depend on including the dynamics of the covariance matrix.

The simulations of the heterosynaptic plasticity protocol are similar to the STDP protocol; however, they include a preconditioning protocol and an additional synaptic weight, whose change in strength we label as heterosynaptic plasticity. Figure 9 (A) and (B) show Synaptic Filter variables for the cases $t_{\text{pre}} - t_{\text{post}} = \pm 10\text{ms}$. The preconditioning protocol consists of two presynaptic spikes at both inputs with minimal delay. This correlated input causes negative weight correlations via Equation (59). When the pre-before-post, shown in (A), and post-before-pre, shown in (B), protocols are applied, the negative weight correlation between the first and the second weight leads to opposing directions of plasticity. Mathematically, this is manifested in the prefactor $\Sigma_t x_t^\epsilon$ in the updates, Equation (58). When weight correlations are present, the covariance matrix

converts presynaptic activation in the first weight into a non-zero coefficient in the second weight, i.e., the matrix prefactor mixes presynaptic activation between inputs. Thus, the results for heterosynaptic plasticity depend on the presence of weight correlations.

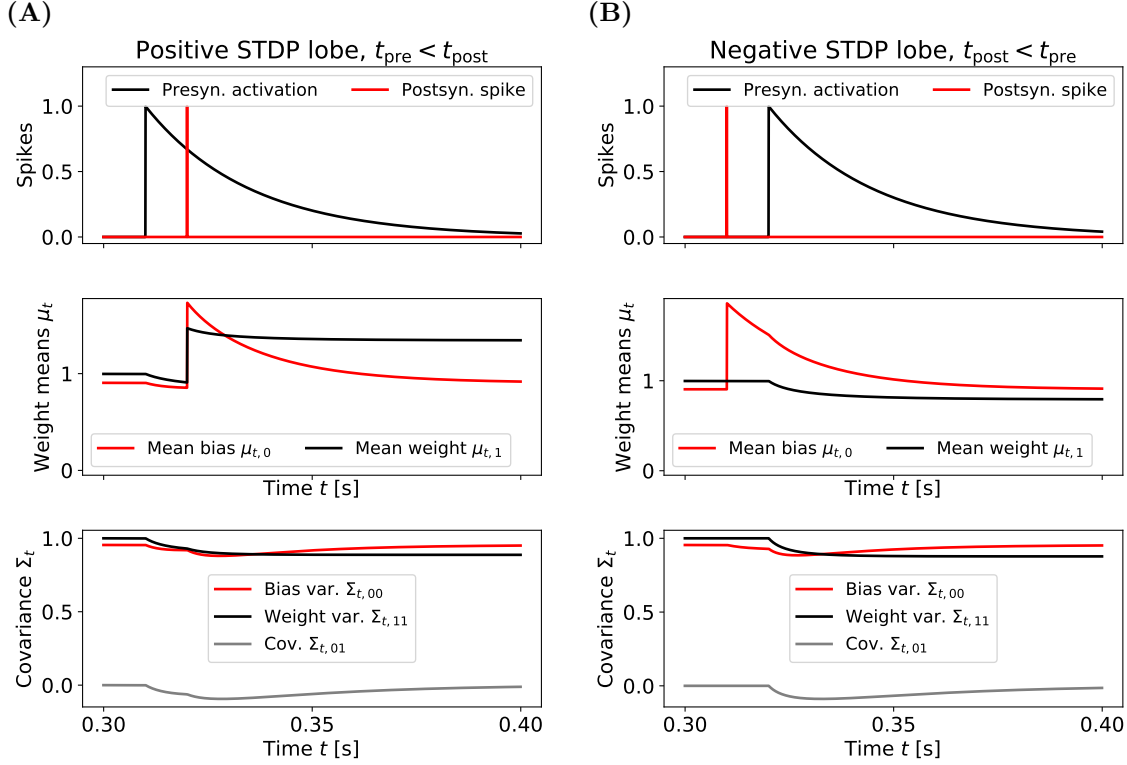


Figure 8: The dynamics of the variables of the Synaptic Filter during the STDP protocol. **(A, top)** shows a protocol for the positive lobe with a presynaptic spike (black) followed by a postsynaptic spike (red) with 10ms delay. The presynaptic activation paired with a postsynaptic spikes increases the mean value of the bias (red) and synaptic weight (black), shown in **(A, middle)**. The returns to its equilibrium value on a timescale of $\tau_{\text{ou},\text{bias}} = 25\text{ms}$. **(A, bottom)** shows that spiking activities reduces the elements of the covariance matrix. The variance of the bias returns quickly to its equilibrium value. **(B, top)** shows a protocol for the negative lobe. Importantly the depression of the mean weight (black) in **(B, middle)** is modulated by mean value of the bias (red). The smaller the delay of the presynaptic spike, the larger is the decrease in the mean weight.

5.6 Negativity of weight correlations

In the Materials and Methods in the Main Text, we showed that the weight correlations, i.e., the off-diagonal elements of the covariance matrix Σ , are always negative in two dimensions. In the following, we show that for constant input $x_i^\xi \equiv x_0 = \text{const}$ the off-diagonal elements are negative for any d .

To show this, we represent the covariance matrix in the orthogonal, normalised basis: $\hat{x}_i^\top \hat{x}_j = \delta_{ij}$ for all $i, j \in (0, \dots, d-1)$. The basis is chosen such that its first basis vector is parallel to the input vector: $\hat{x}_0^\top x_0 = \|x_0\|$. In this basis, the covariance matrix is:

$$\Sigma = \sum_{i=0}^{d-1} a_i \hat{x}_i \hat{x}_i^\top, \quad (60)$$

where $a_i \geq 0$ because the covariance is positive semi-definite. Using this representation in the covariance update Equation (59) and projecting the dynamics onto the basis yields update equations for the coefficients:

$$\dot{a}_i = \hat{x}_i^\top \dot{\Sigma} \hat{x}_i = -\delta_{i0} \gamma a_i^2 \|x_0\|^2 - 2\tau_{\text{ou}}^{-1} (a_i - \sigma_{\text{ou}}^2), \quad (61)$$

where we used $\Sigma_{\text{ou}} = \mathbb{1}\sigma_{\text{ou}}^2$ to obtain the second term.

The initial condition of the covariance matrix is $\Sigma = \mathbb{1}\sigma_{\text{ou}}^2$. Thus, the initial condition of the coefficients is $a_i = \sigma_{\text{ou}}^2$. It follows from Equation (61) that the updates of the coefficients are non-positive $\dot{a}_i \leq 0$. The values of a_i decrease until a fixed point is reached. This implies, because the elements of the first basis vector are non-negative, $\hat{x}_{0,k} \geq 0$ for all $k \in (0, \dots, d-1)$, that all elements of the covariance matrix, including the off-diagonals, decrease as well until the fixed points of the coefficients a_i are reached.

Empirically, we find that weight correlations are non-positive in the case of time-dependent inputs as well. The following analysis of a change from one static input to another supports this observation because it concludes that the first coefficient still has the fastest reduction rate after the change.




Figure 9: The dynamics of the variables of the Synaptic Filter during the heterosynaptic protocol with preconditioning. **(A, top)** shows the presynaptic activation (black, gray) of both synapses and the postsynaptic spikes (red) for a pre-before-post protocol. The initial waiting time of 100ms has been removed. **(A, middle)** shows the mean weights (black, gray) and the bias (red). The bias returns to its equilibrium value after the preconditioning protocol. At $t = 300$ ms, the pre-post protocol is applied to the first weight, leading to potentiation and depression in the first and second mean weight respectively. **(A, bottom)** shows elements of the covariance matrix. The preconditioning protocol reduces the variance values. In particular, the covariance between the weights drops below zero. A non-zero value of the negative covariance causes the heterosynaptic plasticity. **(B)** shows a post-before-pre protocol. In contrast to the previous case, the arrival of a presynaptic spike at $t = 300$ ms causes homosynaptic depression (black) and heterosynaptic potentiation (gray), as shown in **(B, middle)**.

The time-dependent input vector is x_t^e . Changes in the input vector correspond to a transformation of the basis vectors $\{\hat{x}_i\}$, used in Equation (60). Under such transformations, the first basis vector \hat{x}_1 remains in the first quadrant because its entries are always non-negative. Without loss of generality, we assume that such a transformation is characterised by a rotation $\phi \in [0, \pi/2]$ in the plane of the first two vectors of the basis and under the constraint that the first basis vector remains in the first quadrant. Thus, the first and second coefficient change as follows:

$$a'_0 = \cos(\phi)a_0 + \sin(\phi)a_1 \quad (62)$$

$$a'_1 = -\sin(\phi)a_0 + \cos(\phi)a_1. \quad (63)$$

From the analysis of the static case and the presence of an additional decay term in Equation (61) for the first coefficient, we conclude that $a_0 < a_1$ before the change of the input vector occurs. Thus Equation (62) implies that $a'_0 > a_0$, which leads to a faster decay of the first coefficient after the change of input. Equation (63) implies that $a'_1 < a_1$, resulting in a slower recovery to the equilibrium value σ_{on}^2 than without the change in input. However, the effect of changing the basis vectors is neglected here.

Intuitively, the transformation, which corresponds to changes in the input, shifts the direction of maximal reduction of the covariance away from the direction in which the covariance has been maximally reduced under the previous input. Thus, the coefficient a_0 remains the one with the highest rate of reduction compared to the other coefficients.

Turning this argument into a proof in the case of time-varying inputs would require the investigation of all edge cases of the coefficient dynamics and is beyond the scope of this work.

References

Martin Stemmler and Christof Koch. How voltage-dependent conductances can adapt to maximize the information encoded by neuronal firing rate. *Nature neuroscience*, 2(6):521–527, 1999.

- H Sebastian Seung. Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron*, 40(6):1063–1073, 2003.
- Máté Lengyel, Jeehyun Kwag, Ole Paulsen, and Peter Dayan. Matching storage and recall: hippocampal spike timing-dependent plasticity and phase response curves. *Nature neuroscience*, 8(12):1677–1683, 2005.
- Olaf Booi and Hieu tat Nguyen. A gradient descent rule for spiking neurons emitting multiple spikes. *Information Processing Letters*, 95(6):552–558, 2005.
- Robert Güttig and Haim Sompolinsky. The tempotron: a neuron that learns spike timing-based decisions. *Nature neuroscience*, 9(3):420–428, 2006.
- Robert Urbanczik and Walter Senn. Learning by the dendritic prediction of somatic spiking. *Neuron*, 81(3):521–528, 2014.
- Yan Xu, Xiaoqin Zeng, Lixin Han, and Jing Yang. A supervised multi-spike learning algorithm based on gradient descent for spiking neural networks. *Neural Networks*, 43:99–113, 2013.
- Sander M Bohte, Joost N Kok, and Han La Poutre. Error-backpropagation in temporally encoded networks of spiking neurons. *Neurocomputing*, 48(1-4):17–37, 2002.
- Sander M Bohte and Michael C Mozer. Reducing spike train variability: A computational theory of spike-timing dependent plasticity. In *Advances in neural information processing systems*, pages 201–208, 2005.
- Jochen Triesch. Synergies between intrinsic and synaptic plasticity in individual model neurons. In *Advances in neural information processing systems*, pages 1417–1424, 2005a.
- Jochen Triesch. A gradient rule for the plasticity of a neurons intrinsic excitability. In *International Conference on Artificial Neural Networks*, pages 65–70. Springer, 2005b.
- Bruno A Olshausen and David J Field. Code for Natural Images. *Nature*, 381(June):607–609, 1996.
- Jean-Pascal Pfister, Taro Toyozumi, David Barber, and Wulfram Gerstner. Optimal spike-timing-dependent plasticity for precise action potential firing in supervised learning. *Neural computation*, 18(6):1318–1348, 2006.
- Johanni Brea, Walter Senn, and Jean-Pascal Pfister. Matching recall and storage in sequence learning with spiking neural networks. *Journal of neuroscience*, 33(23):9565–9575, 2013.
- Gal Chechik. Spike-timing-dependent plasticity and relevant mutual information maximization. *Neural computation*, 15(7):1481–1510, 2003.
- Taro Toyozumi, Jean-Pascal Pfister, Kazuyuki Aihara, and Wulfram Gerstner. Generalized bienenstock-cooper-munro rule for spiking neurons that maximizes information transmission. *Proceedings of the National Academy of Sciences*, 102(14):5239–5244, 2005.
- Anthony J Bell and Lucas C Parra. Maximising sensitivity in a spiking network. In *Advances in neural information processing systems*, pages 121–128, 2005.
- Christopher L Buckley, Chang Sub Kim, Simon McGregor, and Anil K Seth. The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, 81:55–79, 2017.
- David JC MacKay. Bayesian interpolation. *Neural computation*, 4(3):415–447, 1992.
- Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. *arXiv preprint arXiv:1505.05424*, 2015.
- David E Bell. Regret in decision making under uncertainty. *Operations research*, 30(5):961–981, 1982.
- Christian Henning, Johannes von Oswald, João Sacramento, Simone Carlo Surace, Jean-Pascal Pfister, and Benjamin F Grewe. Approximating the predictive distribution via adversarially-trained hypernetworks. *NeurIPS workshop*, 2018.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.

- German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 2019.
- Philipp J Kraemer and Jonathan M Golding. Adaptive forgetting in animals. *Psychonomic Bulletin & Review*, 4(4):480–491, 1997.
- Jan Zimmermann, Paul W Glimcher, and Kenway Louie. Multiple timescales of normalized value coding underlie adaptive choice behavior. *Nature communications*, 9(1):3206, 2018.
- Yichun Shuai, Binyan Lu, Ying Hu, Lianzhang Wang, Kan Sun, and Yi Zhong. Forgetting is regulated through rac activity in drosophila. *Cell*, 140(4):579–589, 2010.
- Johanni Brea, Robert Urbanczik, and Walter Senn. A normative theory of forgetting: lessons from the fruit fly. *PLoS computational biology*, 10(6):e1003640, 2014.
- Arash Fassihi, Athena Akrami, Vahid Esmaeili, and Mathew E Diamond. Tactile perception and working memory in rats and humans. *Proceedings of the National Academy of Sciences*, 111(6):2331–2336, 2014.
- Athena Akrami, Arash Fassihi, Vahid Esmaeili, and Mathew E Diamond. Tactile working memory in rat and human: Prior competes with recent evidence. *PLoS One*, 6(5):e19551, 2011.
- Matthew FS Rushworth and Timothy EJ Behrens. Choice, uncertainty and value in prefrontal and cingulate cortex. *Nature neuroscience*, 11(4):389, 2008.
- Kenji Doya. Modulators of decision making. *Nature neuroscience*, 11(4):410–416, 2008.
- Paul J Cocker, Katherine Dinelle, Rick Kornelson, Vesna Sossi, and Catharine A Winstanley. Irrational choice under uncertainty correlates with lower striatal d2/3 receptor binding in rats. *Journal of Neuroscience*, 32(44):15450–15457, 2012.
- Maik C Stüttgen and Cornelius Schwarz. Psychophysical and neurometric detection performance under stimulus uncertainty. *Nature neuroscience*, 11(9):1091, 2008.
- Wei Ji Ma. Signal detection theory, uncertainty, and poisson-like population codes. *Vision research*, 50(22):2308–2319, 2010.
- David C Knill and Alexandre Pouget. The bayesian brain: the role of uncertainty in neural coding and computation. *TRENDS in Neurosciences*, 27(12):712–719, 2004.
- David Kappel, Stefan Habenschuss, Robert Legenstein, and Wolfgang Maass. Synaptic sampling: A bayesian approach to neural network plasticity and rewiring. In *Advances in Neural Information Processing Systems*, pages 370–378, 2015.
- Harold J Kushner. On the differential equations satisfied by conditional probability densities of markov processes, with applications. *Journal of the Society for Industrial and Applied Mathematics, Series A: Control*, 2(1):106–119, 1964.
- Harold J Kushner. Dynamical equations for optimal nonlinear filtering. *Journal of Differential Equations*, 3(2):179–190, 1967.
- Johanni Brea and Wulfram Gerstner. Does computational neuroscience need new synaptic learning paradigms? *Current opinion in behavioral sciences*, 11:61–66, 2016.
- Laurence Aitchison and Peter E Latham. Synaptic sampling: A connection between psp variability and uncertainty explains neurophysiological observations. *arXiv preprint arXiv:1505.04544*, 2015.
- Henry Markram, Joachim Lübke, Michael Frotscher, and Bert Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps. *Science*, 275(5297):213–215, 1997.
- Guo-qiang Bi and Mu-ming Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of neuroscience*, 18(24):10464–10472, 1998.
- Laurence Aitchison and Peter E Latham. Bayesian synaptic plasticity makes predictions about plasticity experiments in vivo. *arXiv preprint arXiv:1410.1029*, 2014.
- Sébastien Royer and Denis Paré. Conservation of total synaptic weight through balanced synaptic depression and potentiation. *Nature*, 422(6931):518–522, 2003.

- Nicolas Frémaux and Wulfram Gerstner. Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. *Frontiers in neural circuits*, 9:85, 2016.
- Vasiliki Liakoni, Alireza Modirshanechi, Wulfram Gerstner, and Johanni Brea. An approximate bayesian approach to surprise-based learning. *arXiv preprint arXiv:1907.02936*, 2019.
- Wulfram Gerstner and Werner M Kistler. *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press, 2002.
- WC Abraham and Graham V Goddard. Asymmetric relationships between homosynaptic long-term potentiation and heterosynaptic long-term depression. *Nature*, 305(5936):717–719, 1983.
- Gary S Lynch, Thomas Dunwiddie, and Valentin Gribkoff. Heterosynaptic depression: a postsynaptic correlate of long-term potentiation. *Nature*, 266(5604):737–739, 1977.
- Craig H Bailey, Maurizio Giustetto, Yan-You Huang, Robert D Hawkins, and Eric R Kandel. Is heterosynaptic modulation essential for stabilizing hebbian plasticity and memory. *Nature Reviews Neuroscience*, 1(1):11–20, 2000.
- Marina Chistiakova, Nicholas M Bannon, Jen-Yung Chen, Maxim Bazhenov, and Maxim Volgushev. Homeostatic role of heterosynaptic plasticity: models and experiments. *Frontiers in computational neuroscience*, 9: 89, 2015.
- Simone Carlo Surace, Jean-Pascal Pfister, Wulfram Gerstner, and Johanni Brea. On the choice of metric in gradient-based theories of brain function. *PLOS Computational Biology*, 16(4):e1007640, 2020.
- Marina Chistiakova and Maxim Volgushev. Heterosynaptic plasticity in the neocortex. *Experimental brain research*, 199(3-4):377, 2009.
- João Sacramento, Rui Ponte Costa, Yoshua Bengio, and Walter Senn. Dendritic cortical microcircuits approximate the backpropagation algorithm. In *Advances in Neural Information Processing Systems*, pages 8721–8732, 2018.
- Benjamin Scellier and Yoshua Bengio. Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. *Frontiers in computational neuroscience*, 11:24, 2017.
- Mathieu Schiess, Robert Urbanczik, and Walter Senn. Somato-dendritic synaptic plasticity and error-backpropagation in active dendrites. *PLoS computational biology*, 12(2), 2016.
- Yan Yang and Stephen G Lisberger. Purkinje-cell plasticity and cerebellar motor learning are graded by complex-spike duration. *Nature*, 510(7506):529–532, 2014.
- Vincent Jacob, Leopoldo Petreanu, Nick Wright, Karel Svoboda, and Kevin Fox. Regular spiking and intrinsic bursting pyramidal cells show orthogonal forms of experience-dependent plasticity in layer v of barrel cortex. *Neuron*, 73(2):391–404, 2012.
- Wickliffe C Abraham and Mark F Bear. Metaplasticity: the plasticity of synaptic plasticity. *Trends in neurosciences*, 19(4):126–130, 1996.
- Jean-Pascal Pfister, Peter Dayan, and Máté Lengyel. Know thy neighbour: A normative theory of synaptic depression. In *Advances in neural information processing systems*, pages 1464–1472, 2009.
- Renaud Jolivet, Alexander Rauch, Hans-Rudolf Lüscher, and Wulfram Gerstner. Predicting spike timing of neocortical pyramidal neurons by simple threshold models. *Journal of computational neuroscience*, 21(1): 35–49, 2006.
- Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. Density-ratio matching under the bregman divergence: a unified framework of density-ratio estimation. *Annals of the Institute of Statistical Mathematics*, 64 (5):1009–1044, 2012.
- Arnaud Doucet, Simon Godsill, and Christophe Andrieu. On sequential monte carlo sampling methods for bayesian filtering. *Statistics and computing*, 10(3):197–208, 2000.
- Anna Kutschireiter, Simone Carlo Surace, and Jean-Pascal Pfister. The hitchhikers guide to nonlinear filtering. *Journal of Mathematical Psychology*, 94:102307, 2020.
- Pierre Brémaud. *Point processes and queues: martingale dynamics*, volume 50. Springer, 1981.
- Thomas Peter Minka. *A family of algorithms for approximate Bayesian inference*. PhD thesis, Massachusetts Institute of Technology, 2001.