Inferring in Circles: Active Inference in Continuous State Space Using Hierarchical Gaussian Filtering of Sufficient Statistics

Peter Thestrup Waade[1,2](✉), Nace Mikus[1,3], and Christoph Mathys[1,4,5]

[1] Interacting Minds Centre (IMC), Aarhus University, Aarhus, Denmark
ptw@cas.au.dk
[2] Embodied Cognition Group (ECG), Aarhus University, Aarhus, Denmark
[3] Department of Cognition, Emotion, and Methods in Psychology, University of Vienna, Vienna, Austria
[4] Scuola Internazionale Superiore di Studi Avanzati (SISSA), Trieste, Italy
[5] Translational Neuromodeling Unit (TNU), Institute for Biomedical Engineering, University of Zurich and ETH Zurich, Zurich, Switzerland

Abstract. We create a continuous state space active inference agent based on the hierarchical Gaussian filter. It uses the HGF to track the sufficient statistics of noisy observations of a moving target that is performing a Gaussian random walk with drift and varying volatility. On the basis of this filtering, the agent predicts the target's position, and minimizes surprisal by staying close to it. Our simulated agent represents the first full implementation of this approach. It demonstrates the feasibility of supplementing active inference with HGF-filtering of the sufficient statistics of observations, which is particularly useful in noisy and volatile continuous state space environments.

Keywords: Active inference · Continuous state space · Sufficient statistics filtering · Precision-weighted prediction errors · Hierarchical gaussian filter

1 Introduction

Active inference [7] is a formal framework for programming and modelling agents that navigate their environment such that they sample evidence for being within a desired set of states. This is done by minimizing the surprisal of sensory observations relative to a generative model of the environment, in which preferences for states are encoded as prior expectations. Actions are then chosen that are expected to lead to less surprisal in the future. Evaluating surprisal exactly is usually computationally intractable. In practical implementations of active inference, a variational free energy approximation is therefore often used.

Active inference furnishes a modeling framework which unites action and perception under a shared optimization imperative. Models inherently include a

M. Kamp et al. (Eds.): ECML PKDD 2021 Workshops, CCIS 1524, pp. 810–818, 2021.
https://doi.org/10.1007/978-3-030-93736-2_57

balance between epistemic and pragmatic behavior [4], can be related to neurobiological theories such as predictive processing [6], and can be motivated from first principles in physics and information theory [2,8].

Recently, most active inference agents have been implemented as partially observable Markov Decision Processes (POMDP's) [18]. Here agents are limited to making discrete actions and observations in a discrete state space. By contrast, we here aim to (re-)extend active inference models to the continuous domain. We demonstrate a principled approach where an active inference agent filters the sufficient statistics of its observations with a hierarchical Gaussian filter [12,13], allowing it to perform goal-directed actions in a noisy and volatile continuous state space-environment.

2 Filtering Sufficient Statistics with Hierarchical Gaussian Filters

For agents inferring continuous states obscured by observational, informational, and environmental uncertainty, a fundamental challenge is to filter these various sources of noise from their observations. One principled way of solving this problem, which we use here and which is consistent with active inference in general, is to invert a generative model of what causes sensory observations. The hierarchical Gaussian filter's update equations implement such an inversion, where the generative model consists of a hierarchical cascade of random walks [13]. The update equations in the HGF are a more efficient alternative to variational Laplace, as detailed in [12]. Given a time series of observations, this allows for teasing apart observation noise, (potentially changing) volatility and (possibly state-dependent) regularities like drifts and biases. HGFs also provide precision-weighted predictions for future states, and can be used to infer a full predictive posterior probability distribution over observations in the future. This can be done by constructing the predictive distribution such that it reflects the uncertainty implied in the HGF's updates when filtering the sufficient statistics of the observations [14].

The decisive point here is that in a Gaussian model for a continuous univariate hidden state (i.e., Gaussian prior and Gaussian likelihood), the prior and posterior predictive distributions are *Gaussian-predictive* distributions, corresponding to a reparameterization of the generalized Student's-t distribution. This means that in addition to location and scale parameters, the predictive distribution also has a degree-of-freedom parameter which determines the fatness of its tails. An appropriate filter, such as an HGF, allows for inferring all three of these parameters. We put this to use here in order for our active inference agent to make the most appropriate predictions possible, i.e. predictions which minimize surprisal by optimizing all three of their aspects: point prediction (mean), uncertainty (variance), and fatness of tails (degrees of freedom). In the next section, we demonstrate how this can be accomplished in a simple active inference context.

3 Active Inference in Continuous State Space

We here provide a framework for a proof-of-principle simulation with a simple HGF-based active inference agent. The agent's objective is to stay close to a target which moves in continuous space with varying volatility. We will here first describe the *generative process* that forms the environment, and how it interfaces with the agent's *observations* o and *actions* a. Then we will describe how the agent makes actions as active inference based on inferences from the HGF.

The generative process consists of a total of three hidden states. The first is the target's position x_1 which moves in a Gaussian random walk with a constant drift ρ:

$$x_{1,t} \sim \mathcal{N}(x_{1,t-1} + \rho, x_{2,t}) \tag{1}$$

The second is the volatility of the random walk x_2, which changes in a pre-specified pattern between low and high levels of volatility. The last is then the agent's own position x_{agent}. In this example simulation, the position is fully determined by the agent's action a_{agent}, implemented here as being sampled from a delta distribution:

$$x_{\text{agent},t} \sim \delta(a_{\text{agent},t}) \tag{2}$$

The target's position x_1 is observed noisily, with observations o_1 normally distributed around the true position with standard deviation σ:

$$o_{1,t} \sim \mathcal{N}(x_{1,t}, \sigma) \tag{3}$$

The agent also observes its own position perfectly:

$$o_{\text{agent},t} \sim \delta(x_{\text{agent},t}) \tag{4}$$

To make inferences and predictions about the position and volatility of the target, the agent uses a standard HGF with a single volatility parent and a drift on the position. On each timestep t, this gives the agent a Gaussian belief about the target's position on the next timestep with mean $\hat{\mu}_{x_1,t}$ and precision $\hat{\pi}_{x_1,t}$. This lets it generate a full predictive posterior probability distribution for the observation on the next timestep. This distribution is a t-distribution with $\nu_t + 1$ degrees of freedom, with location $\hat{\mu}_{x_1,t}$ and precision $\hat{\pi}_{x_1,t}$:

$$p_{PP}(o_{1,t+1}|o_{1,1:t}) = t\left(o_{1,t}|\hat{\mu}_{x_1,t}, \hat{\pi}_{x_1,t}, \nu_t + 1\right) \tag{5}$$

where

$$\nu_t = \frac{\hat{\pi}_{x_1,t}}{\pi_\epsilon}, \tag{6}$$

and π_ϵ is the agent's prior belief about the input precision. In addition, the agent is equipped with a static prior distribution encoding its expectations (i.e.

preferences) for observations. The *goal prior*, as it will be referred to, is here a probability distribution over differences between the observed position of the target o_1 and the observation of the agent's own position o_{agent}. Specifically, it is a Gaussian distribution, with the mean μ_{GP} (usually at 0) encoding the preferred position relative to the target and the precision π_{GP} specifying the strength of this preference:

$$p_{GP}(o_1 - o_{\text{agent}}) = \mathcal{N}(o_1 - o_{\text{agent}}; \mu_{GP}, \pi_{GP}) \tag{7}$$

On each trial, the agent's surprisal is calculated as the negative log probability of its sensory input relative to the goal prior:

$$\Im(o_1 - o_{\text{agent}}) = -\ln\, p_{GP}(o_1 - o_{\text{agent}}) \tag{8}$$

In order to choose its action, the agent creates an expected surprisal distribution over possible control states a. First the expectation of the predictive posterior over observations of the target is assumed as the observation of the target. This gives a time-varying distribution over the agent's preferences for observations of its own position, given that the target is observed at its expectation. In the agent's model of the environment, Eqs. 2 and 4 are recapitulated, so we can substitute the expected observation o_{agent} with the agent's control states a_{agent}:

$$p_{GP,t}(a_{\text{agent}}) = p_{GP}(o_{\text{agent}} | o_1 = E(p_{PP}(o_1, t))) \tag{9}$$

The right side of this equation is the goal prior over observations of the agent, given that the target is observed at its expectation $E(p_{PP}(o_1, t))$. The left side of the equation $p_{GP,t}(a_{\text{agent}})$ then becomes what might be called a 'goal posterior', a distribution over preferences for actions. In order to incorporate the full uncertainty of the agent's predictions, however, this preference distribution is convolved with the full predictive posterior. Taking the negative log of the resulting probability distribution then yields the expected surprisal associated with each possible move, after including the full uncertainty:

$$\Im_{\text{expected},t}(a_{\text{agent}}) = -\ln\, p_{PP}(o_1) * p_{GP,t}(a_{\text{agent}}) \tag{10}$$

The agent then selects deterministically the action with the lowest associated expected surprisal.

$$a_{\text{agent},t} = \underset{a}{\operatorname{argmin}} \quad \Im_{\text{expected},t}(a_{\text{agent}})) \tag{11}$$

4 Example Simulation

We here show results from an example simulation with the environment and the HGF-based active inference implementation described in the previous section. Figure 1 shows a schematic of the inference, prediction and decision process of the active inference agent. A GIF demonstrating the agent moving to follow the noisy and volatile observations of the target can be found on this link: https://osf.io/x5v39/

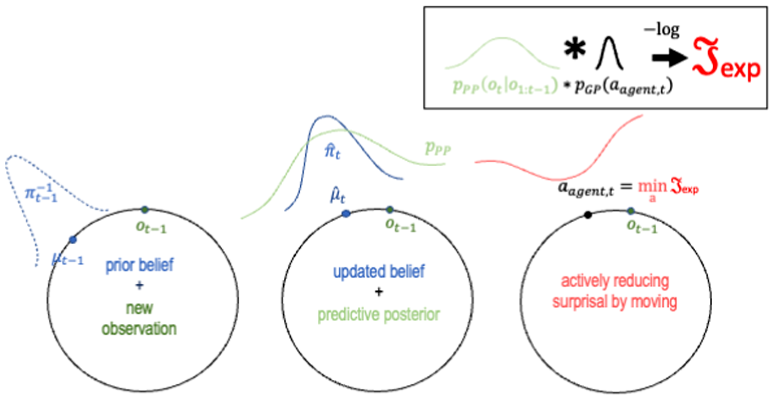

Fig. 1. Graphical sketch of the agent's action process, visualized on the circle. The agent starts with a Gaussian prior belief about the target's position with mean μ_{t-1} and precision π_{t-1}, and makes a new observation o_t. From that a new belief is computed with the HGF (also taking into consideration the drift ρ), and a predictive posterior t-distribution p_{PP} can be calculated. Finally, the predictive posterior is convolved with the 'goal posterior' $p_{GP}(a_{agent,t})$ i.e. the goal prior over agent positions given that the target is observed at its expectation (see Eq. 9). The negative log of the resulting probability distribution is the expected surprisal associated with the agent moving to different positions, of which the lowest is selected. If the static goal prior $p_{GP}(o_1 - o_{agent})$ is symmetric and centered around 0, μ_t, the mean of p_{PP} and the agent's action $a_{agent,t}$ coincide.

Figure 2 shows the inference on the target position, the prediction of future observations and the subsequent movement of the agent in an example simulation. Here the agent's goal prior is centered around 0, meaning that it consistently moves to the mean of its predictive posterior. Figure 3 shows the volatility of the environment and the agents inferred volatility in the same simulation. Note that this is a stochastic process, so even though the generating volatility is high it is not guaranteed that the target will move more. This means that the optimal inference on the volatility is not always the same as the volatility that generated the data (as in this case), although it should be when averaged over many simulations (this is also potentially true for any other hidden aspect of the environment). Figure 4 shows the agent's surprisal at its observation, relative to the goal prior. As expected, surprisal and predictive uncertainty is generally higher in those periods where the actual volatility is higher. In general, the agent performs its task well, demonstrating the feasibility of using the HGF for performing active inference. Code to replicate and modify this simulation can be found on https://github.com/ilabcode/hgf-active-inference

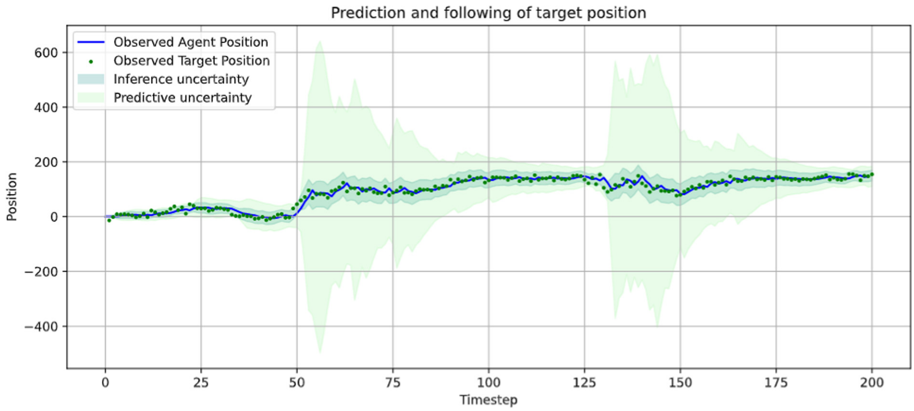

Fig. 2. Relative positions of the agent and the target. Inner shaded area is the inverse precision of the inference of the target position. Outer shaded area is the 68% confidence interval of the predictive posterior

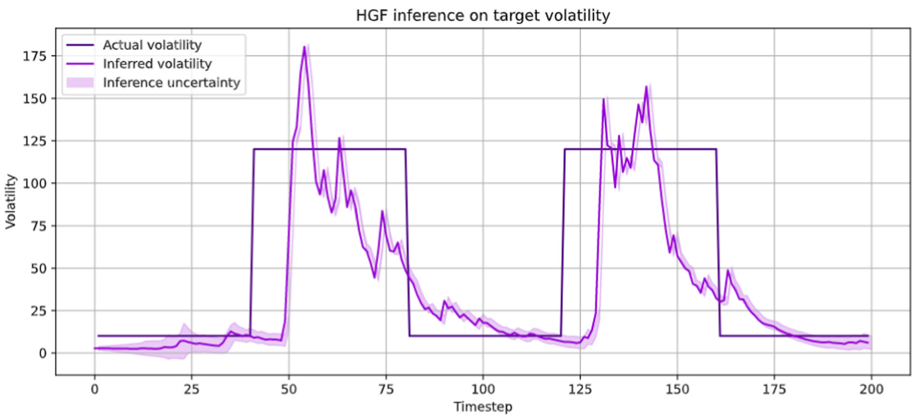

Fig. 3. The agent's HGF-based inference of the volatility (standard deviation) of the target's Gaussian random walk. Shaded area is the inverse precision of the inference.

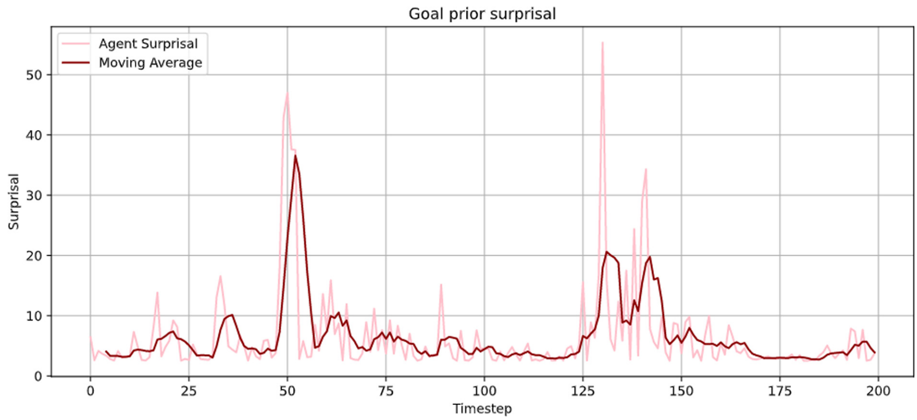

Fig. 4. Agent surprisal at making its observation, as calculated from the goal prior.

5 Discussion

We have here provided a proof-of-principle for active inference models in continuous state space that use hierarchical Gaussian filters to infer and predict the environment. This provides a way of constructing active inference agents that efficiently navigate volatile and noisy continuous state space-environments. The method is consistent with the theoretical framework of active inference, it is elegant, interpretable, and computationally efficient, and shows promise as a method for (re-)extending active inference models into the continuous domain.

There are multiple ways in which this method can be extended in order to employ the active inference framework still more effectively. Most importantly, there is no epistemic component in the current task, which can be included to fully utilize the advantages of active inference. Secondly, we have not here demonstrated the full flexibility of HGF-based active inference. In our simulation, we used a Gaussian goal prior centred around zero, meaning that the agent in practice always moves to the mean of its predictive distribution. Other distributions can be used for the goal prior, however, for example to make the agent prefer observing itself at a certain distance from the target, or be more sensitive to erring in one direction than the other. It is also possible to use more complex instances of the HGF as generative model, allowing for tracking an arbitrary number of possibly inter- or action-dependent hidden states. Since our approach provides a parametric predictive posterior, the expected surprisal can be evaluated directly, leading to the same result as variational methods would converge on. However, when the generative model is more complex, as for example when the agent must plan several steps ahead, or when its actions also influence the movement of the target, this could become less feasible. If so, approximate variational methods might be required.

It would also be valuable to make a more detailed comparison of the HGF-based active inference framework to older continuous state space approaches (for example the saccade models in [3] or the birdsong models in [5]). There are also newer mixed continuous-discrete models which combine discrete policy-level POMDP's with continuous movement and sensation models [9,10,15]. The predictions of the discrete model are then used as prior constraints on the continuous model, which in turn provides evidence for the discrete hypotheses entertained in the former. The main difference between these approaches and HGF-based active inference is that the HGF uses single-step update equations instead of the iterative variational Laplace approach, and that it is generically applicable across contexts. The HGF-based active inference framework can also be contrasted with recent attempts at scaling POMDP methods to complex and continuous domains by amortizing the specification of the generative model with deep learning techniques [1,11,19]. Here, an advantage of HGF-based active inference is that it, beyond specification of hyperparameters as is also the case in deep learning approaches, does not need to be trained, and that it is fully transparent and interpretable.

Finally, it still remains to equip the HGF-based active inference method with parameter and model structure learning capabilities, so that it can select the HGF-architecture that best explains observations. This is especially important when mapping the HGF to neuronal message passing [20]. Note that this can be combined with recent approaches where the hyperparameters of the HGF are learned online [16,17]. It also remains to apply it more complex environments, and to fit it to experimentally observed behaviour. This is feasible because it has been shown that the HGF is generic and adaptable, can be fit to experimental data, and can be mapped onto neuronal message-passing.

References

1. Çatal, O., Wauthier, S., De Boom, C., Verbelen, T., Dhoedt, B.: Learning generative state space models for active inference. Front. Comput. Neurosci. **14**, 103 (2020)
2. Friston, K.: A free energy principle for a particular physics. arXiv preprint arXiv:1906.10184 (2019)
3. Friston, K., Adams, R., Perrinet, L., Breakspear, M.: Perceptions as hypotheses: saccades as experiments. Front. Psychol. **3**, 151 (2012)
4. Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O'Doherty, J., Pezzulo, G.: Active inference and learning. Neurosci. Biobehav. Rev. **68**, 862–879 (2016). https://doi.org/10.1016/j.neubiorev.2016.06.022
5. Friston, K., Frith, C.: A duet for one. Conscious. Cogn. **36**, 390–405 (2015)
6. Friston, K., Kiebel, S.: Predictive coding under the free-energy principle. Philos. Trans. R. Soc. B Biol. Sci. **364**(1521), 1211–1221 (2009)
7. Friston, K.J., Daunizeau, J., Kiebel, S.J.: Reinforcement Learning or Active Inference? PLoS ONE **4**(7), e6421 (2009). https://doi.org/10.1371/journal.pone.0006421
8. Friston, K.J., Daunizeau, J., Kilner, J., Kiebel, S.J.: Action and behavior: a free-energy formulation. Biol. Cybern. **102**(3), 227–260 (2010). https://doi.org/10.1007/s00422-010-0364-z

9. Friston, K.J., Parr, T., de Vries, B.: The graphical brain: belief propagation and active inference. Netw. Neurosci. **1**(4), 381–414 (2017)
10. Friston, K.J., Sajid, N., Quiroga-Martinez, D.R., Parr, T., Price, C.J., Holmes, E.: Active listening. Hear. Res. **399**, 107998 (2021)
11. Van de Maele, T., Verbelen, T., Çatal, O., De Boom, C., Dhoedt, B.: Active vision for robot manipulators using the free energy principle. Front. Neurorobotics **15**, 14 (2021)
12. Mathys, C., Daunizeau, J., Friston, K.J., Stephan, K.E.: A Bayesian foundation for individual learning under uncertainty. Front. Hum. Neurosci. **5**, 39 (2011). https://doi.org/10.3389/fnhum.2011.00039
13. Mathys, C., et al.: Uncertainty in perception and the hierarchical gaussian filter. Front. Hum. Neurosci. **8**, 825 (2014). https://doi.org/10.3389/fnhum.2014.00825
14. Mathys, C., Weber, L.: Hierarchical gaussian filtering of sufficient statistic time series for active inference. In: IWAI 2020. CCIS, vol. 1326, pp. 52–58. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64919-7_7
15. Parr, T., Friston, K.J.: The discrete and continuous brain: from decisions to movement-and back again. Neural Comput. **30**(9), 2319–2347 (2018)
16. Şenöz, İ., De Vries, B.: Online variational message passing in the hierarchical gaussian filter. In: 2018 IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP), pp. 1–6. IEEE (2018)
17. Şenöz, İ., de Vries, B.: Online message passing-based inference in the hierarchical gaussian filter. In: 2020 IEEE International Symposium on Information Theory (ISIT), pp. 2676–2681. IEEE (2020)
18. Smith, R., Friston, K., Whyte, C.: A step-by-step tutorial on active inference and its application to empirical data. PsyArXiv (2021)
19. Tschantz, A., Baltieri, M., Seth, A.K., Buckley, C.L.: Scaling active inference. In: 2020 International Joint Conference on Neural Networks (IJCNN), pp. 1–8. IEEE (2020)
20. Weber, Lilian, A.E.: Perception as Hierarchical Bayesian Inference - Toward Non-Invasive Readouts of Exteroceptive and Interoceptive Processing. Doctoral thesis, ETH Zurich (2020)