

# Planning to avoid ambiguous states through Gaussian approximations to non-linear sensors in active inference agents

Wouter M. Kouw<sup>[0000-0002-0547-4817]</sup>

Bayesian Intelligent Autonomous Systems laboratory  
TU Eindhoven, Eindhoven, Netherlands  
w.m.kouw@tue.nl

**Abstract.** In nature, active inference agents must learn how observations of the world represent the state of the agent. In engineering, the physics behind sensors is often known reasonably accurately and measurement functions can be incorporated into generative models. When a measurement function is non-linear, the transformed variable is typically approximated with a Gaussian distribution to ensure tractable inference. We show that Gaussian approximations that are sensitive to the curvature of the measurement function, such as a second-order Taylor approximation, produce a state-dependent ambiguity term. This induces a preference over states, based on how accurately the state can be inferred from the observation. We demonstrate this preference with a robot navigation experiment where agents plan trajectories.

**Keywords:** Active inference · Free energy minimization · Bayesian filtering · Non-linear sensing · Control systems · Planning · Navigation

## 1 Introduction

In nature, intelligent agents build a model to infer the causes of their sensations [2]. In engineering, we are able to utilize knowledge of the relevant physics to structure such a model. In particular, we often know how sensors measure states of the world. For example, we know how radar measures relative velocity and distance [19]. Measurement functions that are non-linear transformations of state variables pose challenges to state estimation, which are often dealt with using Gaussian approximations of the transformed variables [8,14]. We show that for certain Gaussian approximations, an active inference agent will prefer to avoid states because it already knows that state estimation will be difficult.

Active inference agents are based on free energy functionals that rank policies on explorative and goal-directed behaviour [5,7,6,16]. The expected free energy functional can be understood through its decomposition into a cross-entropy term between states and observations given action ("ambiguity"), and a Kullback-Leibler divergence between the posterior predictive and a goal prior distribution ("risk") [7,18,4]. We show that Gaussian approximations of a non-linear observation function that are itself linear in the covariance matrix, e.g.,

first-order Taylor and the unscented transform [9], lead to ambiguity terms that are constant over states. This echoes an earlier finding that agents with a linear Gaussian state-space model exhibit a constant ambiguity term [10]. However, utilizing a second-order Taylor approximation induces a non-constant ambiguity term. Under this model, the agent will avoid states where the non-linear measurement function curves strongly. Our contributions are:

- Analysis of ambiguity in expected free energy functions under three different Gaussian approximations.
- An experiment where a robot must plan a trajectory and navigate to a goal prior distribution, testing the effect of the ambiguity term.

## 2 Problem statement

We want to plan a trajectory for a robot across a plane. The robot’s state at time  $k$  is its planar position and time derivatives,  $x_k \in \mathbb{R}^{D_x}$ . The robot does not sense position directly, but has to infer it from noisy measurements  $y_k \in \mathbb{R}^{D_y}$ , produced by a sensor through a non-linear mapping  $g : \mathbb{R}^{D_x} \rightarrow \mathbb{R}^{D_y}$  and measurement noise  $v_k \in \mathcal{R}^{D_y}$ . It accepts control inputs  $u_k \in \mathbb{R}^{D_u}$  and moves according to linear dynamics with a transition matrix  $A \in \mathbb{R}^{D_x \times D_x}$ , control matrix  $B \in \mathbb{R}^{D_x \times D_u}$  and process noise  $e_k \in \mathbb{R}^{D_x}$ . Overall, we consider robot systems described with discrete-time state-space models of the form:

$$x_k = Ax_{k-1} + Bu_k + e_k, \quad e_k \sim \mathcal{N}(0, Q), \quad (1)$$

$$y_k = g(x_k) + v_k, \quad v_k \sim \mathcal{N}(0, R), \quad (2)$$

where  $Q, R$  are noise covariance matrices.

The goal is to find a sequence of  $T$  controls  $\bar{u}_k = u_{k+1}, \dots, u_{k+T}$  that produces future states close to a desired state  $x_*$ . Agents must plan every time-step. The challenge is that errors in state estimation may cause drastic changes in the planned trajectory, which can lead an agent astray.

**Example** Consider a robot with position and velocity states that must move from position  $x_0 = (0, -1)$  to  $x_* = (0, 1)$ . Its state transition, control and process noise covariance matrices are given by:

$$A = \begin{bmatrix} 1 & 0 & \Delta t & 0 \\ 0 & 1 & 0 & \Delta t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, B = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ \Delta t & 0 \\ 0 & \Delta t \end{bmatrix}, Q = \begin{bmatrix} \sigma_1^2 \frac{\Delta t^3}{3} & 0 & \sigma_1^2 \frac{\Delta t^2}{2} & 0 \\ 0 & \sigma_2^2 \frac{\Delta t^3}{3} & 0 & \sigma_2^2 \frac{\Delta t^2}{2} \\ \sigma_1^2 \frac{\Delta t^2}{2} & 0 & \sigma_1^2 \Delta t & 0 \\ 0 & \sigma_2^2 \frac{\Delta t^2}{2} & 0 & \sigma_2^2 \Delta t \end{bmatrix}, \quad (3)$$

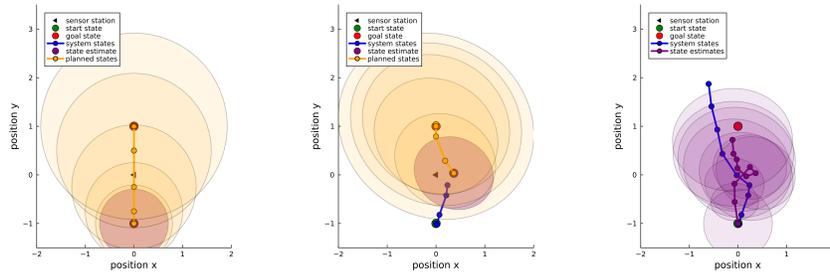
for  $\Delta t = 0.5$ ,  $\sigma_1 = \sigma_2 = 0.1$ . Measurements are produced by a sensor station at  $(0, 0)$  that reports relative angle  $\phi_k \in [-\pi, \pi]$  and relative distance  $d_k \in [0, \infty)$ . The mapping and measurement noise covariance matrix are:

$$g(x_k) = \begin{bmatrix} \phi_k \\ d_k \end{bmatrix} = \begin{bmatrix} \sqrt{x_{1k}^2 + x_{2k}^2} \\ \arctan(x_{1k}, x_{2k}) \end{bmatrix}, R = \begin{bmatrix} \rho_1^2 & 0 \\ 0 & \rho_2^2 \end{bmatrix}, \quad (4)$$

where  $\rho_1 = \rho_2 = 0.001$ . Suppose it uses an extended Kalman filter (first-order Taylor approximation) for state estimation and a finite-horizon model-predictive control objective of the form:

$$J_k(\bar{u}_k) = \sum_{t=k+1}^{k+T} ((A\hat{x}_{t-1} + Bu_t) - x_*)^T C ((A\hat{x}_{t-1} + Bu_t) - x_*) + \eta u_t^2, \quad (5)$$

where  $\hat{x}$  is the mean state,  $\hat{x}_t = A\hat{x}_{t-1} + Bu_t$ ,  $C$  is a cost matrix (ones for position, zeros for velocity) and  $\eta$  a regularization parameter. Minimizing this objective every time-step produces the control sequence  $\bar{u}_k^{\text{MPC}} = \arg \min J_k(\bar{u}_k)$ . Such an agent will first plan a trajectory moving directly forward, as described in Figure 1 (left). However, as it approaches the sensor station, its state estimate become progressively more inaccurate and it makes increasingly more drastic adjustments to the control plan (see  $k = 5$  in Figure 1 middle). Figure 1 (right) shows the executed trajectory over a trial of 10 steps, demonstrating that the agent lost track of the robot's state and did not successfully reach the target.



**Fig. 1.** (Left) Planned trajectory at  $k = 1$ , from start to goal directly over the sensor station. (Middle) Planned trajectory at  $k = 5$  showing a mismatch between true and estimated state resulting in a strong adjustment to the planned trajectory. (Right) Executed trajectory over a trial of 10 steps demonstrates the agent losing track of the robot when it approaches the sensor station.

### 3 Agent specification

#### 3.1 Probabilistic Model

The agent's model will have Gaussian prior distributions over states and controls,

$$p(x_0) = \mathcal{N}(x_0 | m_0, S_0), \quad \mathcal{N}(u_k | 0, \eta^{-1}I), \quad (6)$$

with mean  $m_0$ , covariance matrix  $S_0$ , precision  $\eta$  and identity matrix  $I$ . The agent's state transition will also be expressed as a Gaussian distribution:

$$p(x_k | x_{k-1}, u_k) = \mathcal{N}(x_k | Ax_{k-1} + Bu_k, Q). \quad (7)$$

Let the marginal state  $x_k$  be Gaussian distributed, i.e.,  $p(x_k) = \mathcal{N}(x_k | m_k, S_k)$ . We restrict our attention to approximations of the nonlinear sensor  $g(x_k)$  that produce Gaussian joint distributions over states and observations [14], i.e.,

$$p(y_k, x_k) \approx \mathcal{N}\left(\begin{bmatrix} x_k \\ y_k \end{bmatrix} \mid \begin{bmatrix} m_k \\ \mu_k \end{bmatrix}, \begin{bmatrix} S_k & \Gamma_k \\ \Gamma_k^\top & \Sigma_k \end{bmatrix}\right). \quad (8)$$

From the joint, we obtain a conditional distribution of observations given states:

$$p(y_k | x_k) \approx \mathcal{N}(y_k | \mu_k + \Gamma_k^\top S_k^{-1}(x_k - m_k), \Sigma_k - \Gamma_k S_k^{-1} \Gamma_k^\top). \quad (9)$$

This distribution is linear in  $x_k$ , and will allow for exact Bayesian filtering. But note that the parameters  $\mu_k$ ,  $\Gamma_k$  and  $\Sigma_k$  may be nonlinear functions of  $x_k$ , depending on the type of Gaussian approximation (specifics treated in Section 4), and may thus capture more of the effect of  $g(x_k)$ .

### 3.2 Inferring states

We assume that, when inferring states, the agent has observed the system output  $y_k = \hat{y}_k$  and input  $u_k = \hat{u}_k$ . Let  $\mathcal{D}_k \triangleq \{\hat{y}_i, \hat{u}_i\}_{i=1}^k$  refer to data observed thus far. Given the known executed control, state estimation follows the general Bayesian filtering equations [14]. Firstly, the prior predictive distribution is given by:

$$p(x_k | \hat{u}_k, \mathcal{D}_{k-1}) = \int p(x_k | x_{k-1}, \hat{u}_k) p(x_{k-1} | \mathcal{D}_{k-1}) dx_{k-1} = \mathcal{N}(x_k | \bar{m}_k, \bar{S}_k). \quad (10)$$

with  $\bar{m}_k \triangleq A m_{k-1} + B \hat{u}_k$  and  $\bar{S}_k \triangleq A S_{k-1} A^\top + Q$ . This prediction is corrected by the observation through Bayes' rule [14],

$$p(x_k | \mathcal{D}_k) = \frac{p(\hat{y}_k | x_k)}{p(\hat{y}_k | \mathcal{D}_{k-1})} p(x_k | \hat{u}_k, \mathcal{D}_{k-1}) = \mathcal{N}(x_k | m_k, S_k), \quad (11)$$

with  $m_k = \bar{m}_k + \Gamma_k \Sigma_k^{-1}(\hat{y}_k - \mu_k)$  and  $S_k = \bar{S}_k - \Gamma_k \Sigma_k^{-1} \Gamma_k$ .

### 3.3 Inferring controls

We will discuss the inference procedure first for a single step into the future, and then generalize to a finite horizon of length  $T$ . Predictions for the future state and observation are made by unrolling the generative model to  $t = k + 1$ :

$$p(y_t, x_t, u_t | \mathcal{D}_k) = p(y_t | x_t) p(x_t | u_t; \mathcal{D}_k) p(u_t). \quad (12)$$

We will use an expected free energy functional to infer a posterior distribution over the control  $u_t$  [13]:

$$\mathcal{F}_k[q] = \int q(y_t | x_t) \int q(x_t, u_t) \ln \frac{q(x_t, u_t)}{p(y_t, x_t, u_t | \mathcal{D}_k)} d(u_t, x_t) dy_t. \quad (13)$$

The variational model is specified to be:

$$q(y_t | x_t) \triangleq p(y_t | x_t), \quad q(x_t, u_t) \triangleq p(x_t | u_t; \mathcal{D}_k)q(u_t). \quad (14)$$

Constraining  $q(y_t | x_t)$  to the Gaussian approximation defined in Eq. 9 allows us to study deterministic approximations in an expected free energy minimization context. Given this variational model, Eq. 13 may be re-arranged to:

$$\mathcal{F}_k[q] = \int p(y_t | x_t) \int p(x_t | u_t; \mathcal{D}_k)q(u_t) \ln \frac{p(x_t | u_t; \mathcal{D}_k)q(u_t)}{p(y_t, x_t, u_t; \mathcal{D}_k)} d(x_t, u_t) dy_t \quad (15)$$

$$= \int q(u_t) \left( \int p(y_t, x_t | u_t; \mathcal{D}_k) \ln \frac{p(x_t | u_t; \mathcal{D}_k)q(u_t)}{p(y_t, x_t | u_t; \mathcal{D}_k)p(u_t)} d(y_t, x_t) \right) du_t \quad (16)$$

$$= \int q(u_t) \left( \ln \frac{q(u_t)}{p(u_t)} + \underbrace{\int p(y_t, x_t | u_t; \mathcal{D}_k) \ln \frac{p(x_t | u_t; \mathcal{D}_k)}{p(y_t, x_t | u_t; \mathcal{D}_k)} d(y_t, x_t)}_{\mathcal{J}_k(u_t)} \right) du_t. \quad (17)$$

We refer to  $\mathcal{J}_k(u_t)$  as the expected free energy *function* as it depends on the value of  $u_t$  not on its distribution. Under  $\mathcal{J}_k(u_t) = \ln(1/\exp(-\mathcal{J}_k(u_t)))$ , the expected free energy functional can be concisely expressed as:

$$\mathcal{F}_k[q] = \int q(u_t) \ln \frac{q(u_t)}{p(u_t) \exp(-\mathcal{J}_k(u_t))} du_t. \quad (18)$$

The above is a Kullback-Leibler divergence, which is minimal when

$$q^*(u_t) \propto p(u_t) \exp(-\mathcal{J}_k(u_t)). \quad (19)$$

The proportionality is due to the implicit constraint<sup>1</sup> that  $q^*(u_t)$  should integrate to 1. To work out the expectation in Eq. 17, we first decompose the joint over states and observations into

$$p(y_t, x_t | u_t; \mathcal{D}_k) = p(x_t | y_t, u_t; \mathcal{D}_k)p(y_t), \quad (20)$$

and then intervene on the marginal distribution over  $y_t$  with a distribution reflecting desired future observations (a.k.a. goal prior) [11]:

$$p(y_t) \rightarrow p(y_t | y_*) = \mathcal{N}(y_t | \mu_*, \Sigma_*). \quad (21)$$

The next step involves applying Bayes' rule in the inverse direction:

$$\frac{1}{p(x_t | y_t, u_t; \mathcal{D}_k)} = \frac{p(y_t | u_t; \mathcal{D}_k)}{p(y_t | x_t)p(x_t | u_t; \mathcal{D}_k)}, \quad (22)$$

<sup>1</sup> A more rigorous treatment would define a Lagrangian with normalization and marginalization constraints [17]. However, such a treatment is inconsequential when resorting to MAP estimation, as will be pursued later in the paper.

where the marginal prediction for the future observation is:

$$p(y_t | u_t; \mathcal{D}_k) = \int p(y_t | x_t) p(x_t | u_t; \mathcal{D}_k) dx_t \quad (23)$$

$$= \int \mathcal{N}\left(\begin{bmatrix} x_t \\ y_t \end{bmatrix} \mid \begin{bmatrix} \bar{m}_t \\ \mu_t \end{bmatrix}, \begin{bmatrix} \bar{S}_t & \Gamma_t \\ \Gamma_t^\top & \Sigma_t \end{bmatrix}\right) dx_t = \mathcal{N}(y_t | \mu_t, \Sigma_t). \quad (24)$$

Note that  $\mu_t$  and  $\Sigma_t$  depend on  $u_t$  through  $\bar{m}_t$ . Plugging Eqs. 21 and 22 into Eq. 20 yields:

$$\mathcal{J}_k(u_t) = \int p(y_t, x_t | u_t; \mathcal{D}_k) \ln \frac{p(x_t | u_t; \mathcal{D}_k)}{p(y_t | y_*)} \frac{p(y_t | u_t; \mathcal{D}_k)}{p(y_t | x_t) p(x_t | u_t; \mathcal{D}_k)} d(y_t, x_t) \quad (25)$$

$$= \underbrace{\int p(y_t, x_t | u_t; \mathcal{D}_k) \left[ -\ln \frac{p(y_t, x_t | u_t; \mathcal{D}_k)}{p(x_t | u_t; \mathcal{D}_k)} \right] d(y_t, x_t)}_{\text{ambiguity}} + \underbrace{\int \left[ \int p(y_t, x_t | u_t; \mathcal{D}_k) dx_t \right] \ln \frac{p(y_t | u_t; \mathcal{D}_k)}{p(y_t | y_*)} dy_t}_{\text{risk}}. \quad (26)$$

"Risk" refers to the Kullback-Leibler (KL) divergence between predicted and desired future observations. The inner integral in the risk term leads to a Gaussian distribution (Eq. 24) and the KL divergence between Gaussians is [3]:

$$\mathbb{E}_{p(y_t | u_t; \mathcal{D}_k)} \left[ \ln \frac{p(y_t | u_t; \mathcal{D}_k)}{p(y_t | y_*)} \right] = \frac{1}{2} \left( \ln \frac{|\Sigma_*|}{|\Sigma_t|} - D_y + \text{tr}(\Sigma_*^{-1}(\Sigma_t + \Psi_*)) \right). \quad (27)$$

where  $\Psi_* \triangleq (\mu_* - \mu_t)(\mu_* - \mu_t)^\top$ . "Ambiguity" refers to the conditional entropy of the future observations given the future states.

**Lemma 1.** *Ambiguity, as defined in Eq. 26, for a generative model described in Eq. 12 and a variational distribution described in Eq. 14, is:*

$$\mathbb{E}_{p(y_t, x_t | u_t; \mathcal{D}_k)} \left[ -\ln \frac{p(y_t, x_t | u_t; \mathcal{D}_k)}{p(x_t | u_t; \mathcal{D}_k)} \right] = \frac{D_y}{2} \ln(2\pi e) + \frac{1}{2} \ln |\Sigma_t - \Gamma_t^\top \bar{S}_t^{-1} \Gamma_t|. \quad (28)$$

The proof is in Appendix A. Note that the first term does not depend on the state  $x_t$ . Plugging Eqs. 27 and 28 into the expected free energy function (Eq. 26) produces:

$$\mathcal{J}_k(u_t) = \frac{1}{2} \left( \ln \frac{|\Sigma_*|}{|\Sigma_t|} + D_y \ln(2\pi) + \text{tr}(\Sigma_*^{-1}(\Sigma_t + \Psi_*)) + \ln |\Sigma_t - \Gamma_t^\top \bar{S}_t^{-1} \Gamma_t| \right). \quad (29)$$

Note that  $\Gamma_t$ ,  $\Sigma_t$  and  $\Psi_*$  depend on  $u_t$ . The above steps can be generalized to a longer time horizon  $t = k + 1, \dots, k + T$ . Because the prior is independent over time,  $p(\bar{u}) = \prod_{t=1}^T p(u_t)$  (see Eq. 6), the function  $\mathcal{J}_k(\bar{u})$  factorizes to a sum of recursive expected free energy functions  $\sum_{t=1}^T \mathcal{J}_k(u_t)$ .

We are interested in the most probable value under the approximate control posterior, i.e., the MAP estimate:

$$\hat{u} = \arg \max_{\bar{u} \in \mathcal{U}} q^*(\bar{u}) = \arg \min_{\bar{u} \in \mathcal{U}} \sum_{t=k+1}^{k+T} \mathcal{J}_t(u_t) - \ln p(u_t), \quad (30)$$

where  $\mathcal{U} \subset \mathbb{R}^T$  is the space of affordable controls over  $T$  steps. Constraints such as motor force limits can be imposed during optimization.

## 4 Gaussian approximations

We discuss the three most popular Gaussian approximations to non-linear transformations of Gaussian random variables: the first and second-order Taylor series approximations (used in extended Kalman filters) and the unscented transform (used in the unscented Kalman filter) [9,8][14, Ch. 5].

The first-order Taylor series approximation effectively linearizes the non-linear observation function  $g(x_t)$ . Since ambiguity is known to be constant over states under a linear observation function [10], it is no surprise that the first-order Taylor also leads to an ambiguity term that is constant over states.

**Theorem 1.** *Let  $G_x(\bar{m}_t)$  be the Jacobian of  $g$  with respect to  $x_t$ , evaluated at  $\bar{m}_t$ . Under a first-order Taylor approximation, the parameters  $\Sigma_t, \Gamma_t$  are:*

$$\Sigma_t = G_x(\bar{m}_t) \bar{S}_t G_x(\bar{m}_t)^\top + R, \quad \Gamma_t = \bar{S}_t G_x(\bar{m}_t)^\top. \quad (31)$$

With these parameters, the ambiguity term does not depend on the state  $x_t$ :

$$\mathbb{E}_{p(y_t, x_t | u_t; \mathcal{D}_k)} \left[ -\ln \frac{p(y_t, x_t | u_t; \mathcal{D}_k)}{p(x_t | u_t; \mathcal{D}_k)} \right] = \frac{D_y}{2} \ln(2\pi e) + \frac{1}{2} \ln |R|. \quad (32)$$

The proof is in Appendix B. Perhaps surprisingly, under the second-order Taylor approximation, the ambiguity term varies as a function of the state  $x_t$ .

**Theorem 2.** *Let  $G_{xx}^{(i)}(\bar{m}_t)$  be the Hessian of the  $i$ -th element of the non-linear observation function evaluated at  $\bar{m}_t$  and  $e_i$  be a canonical basis vector. The parameters  $\Sigma_t, \Gamma_t$  computed through a second-order Taylor approximation are:*

$$\begin{aligned} \Sigma_t &= G_x(\bar{m}_t) \bar{S}_t G_x(\bar{m}_t)^\top + \frac{1}{2} \sum_{i=1}^{D_y} \sum_{j=1}^{D_y} e_i e_j^\top \text{tr}(G_{xx}^{(i)}(\bar{m}_t) \bar{S}_t G_{xx}^{(j)}(\bar{m}_t) \bar{S}_t) + R \\ \Gamma_t &= \bar{S}_t G_x(\bar{m}_t)^\top. \end{aligned} \quad (33)$$

With these parameters, the ambiguity term depends on  $x_t$  through:

$$\begin{aligned} \mathbb{E}_{p(y_t, x_t | u_t; \mathcal{D}_k)} \left[ -\ln \frac{p(y_t, x_t | u_t; \mathcal{D}_k)}{p(x_t | u_t; \mathcal{D}_k)} \right] &= \\ \frac{D_y}{2} \ln(2\pi e) + \frac{1}{2} \ln \left| \frac{1}{2} \sum_{i=1}^{D_y} \sum_{j=1}^{D_y} e_i e_j^\top \text{tr}(G_{xx}^{(i)}(\bar{m}_t) \bar{S}_t G_{xx}^{(j)}(\bar{m}_t) \bar{S}_t) + R \right|. \end{aligned} \quad (34)$$

The proof is in Appendix C.

Interestingly, the ambiguity is also constant for the unscented transform.

**Theorem 3.** Define  $2D_x + 1$  sigma points as:

$$\chi_0 \triangleq \bar{m}_t, \quad \chi_i \triangleq \bar{m}_t + \sqrt{D_x + \lambda} [\sqrt{\bar{S}_t}]_i, \quad \chi_{D_x+i} \triangleq \bar{m}_t - \sqrt{D_x + \lambda} [\sqrt{\bar{S}_t}]_i, \quad (35)$$

where  $i = 1, \dots, D_x$ ,  $[\cdot]_i$  denotes the  $i$ -th column of a matrix, and  $\sqrt{S}$  denotes the matrix square root such that  $\sqrt{S}\sqrt{S} = S$ . The parameter  $\lambda \triangleq \alpha^2(D_x + \kappa) - D_x$  depends on free parameters  $\alpha$  and  $\kappa$ . Define  $2D_x + 1$  weights as:

$$w_0 \triangleq \frac{\lambda}{D_x + \lambda} + (1 - \alpha^2 + \beta), \quad w_i \triangleq \frac{1}{D_x + \lambda}, \quad (36)$$

for  $i = 1, \dots, 2D_x$  and  $\beta$  as an additional free parameter. Under these sigma points and weights, the parameters  $\mu_t$ ,  $\Sigma_t$ ,  $\Gamma_t$  are [14, Eq. 5.89]:

$$\begin{aligned} \mu_t &= \frac{\lambda}{D_x + \lambda} g(\chi_0) + \sum_{i=1}^{2D_x} \frac{1}{2(D_x + \lambda)} g(\chi_i), \quad \Gamma_t = \sum_{i=0}^{2D_x} w_i (\chi_i - \bar{m}_t)(g(\chi_i) - \mu_t)^\top, \\ \Sigma_t &= \sum_{i=0}^{2D_x} w_i (g(\chi_i) - \mu_t)(g(\chi_i) - \mu_t)^\top + R. \end{aligned} \quad (37)$$

Then, the ambiguity is independent of the state:

$$\mathbb{E}_{p(y_t, x_t | u_t; \mathcal{D}_k)} \left[ -\ln \frac{p(y_t, x_t | u_t; \mathcal{D}_k)}{p(x_t | u_t; \mathcal{D}_k)} \right] = \frac{D_y}{2} \ln(2\pi e) + \frac{1}{2} \ln |R|. \quad (38)$$

The proof can be found in the Appendix D. This result is conjectured to hold for other Gaussian approximations that are linear in their estimate of the covariance matrix, for example the Gauss-Hermite approximation [14, Ch. 6].

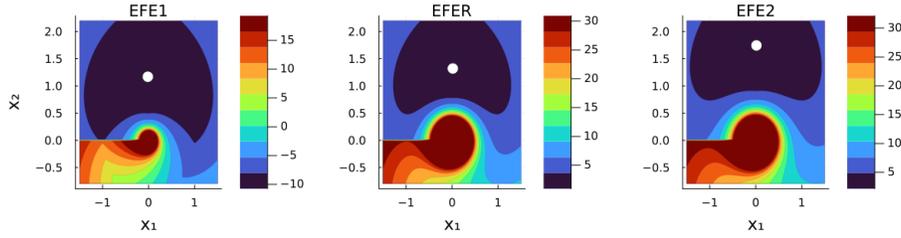
## 5 Experiments

Our experiment is as described in Section 2, with the nonlinear observation function  $g(\cdot)$  measuring relative angle and distance to a base station. Examples of sensors include Hall effect and ultrasound sensors. The robot starts at  $x_0 = [0 \ -1 \ 0 \ 0]$  and must reach  $x_* = [0 \ 1 \ 0 \ 0]$ . The agent's state prior distribution's parameters were  $m_0 = [0 \ -1 \ 0 \ 0]$  and  $S_0 = 0.5I$ . Its control prior precision was set to a tiny value,  $\eta = 1.0 \cdot 10^{-8}$ , so as to best study the effects of ambiguity and risk. It was given a goal prior of  $m_* = g(x_*)$  and  $S_* = 0.5I$ .

We will compare three agents<sup>2</sup>: firstly, an agent that uses the first-order Taylor approximation, referred to as EFE1. Secondly, an agent with a second-order Taylor approximation, referred to as EFE2. Thirdly, an agent with a second-order Taylor approximation but with only the risk term included, referred to as EFER. The difference between EFER and EFE2 reflects the effect of the ambiguity term, while the difference between EFE1 and EFE2 reflects the effect of

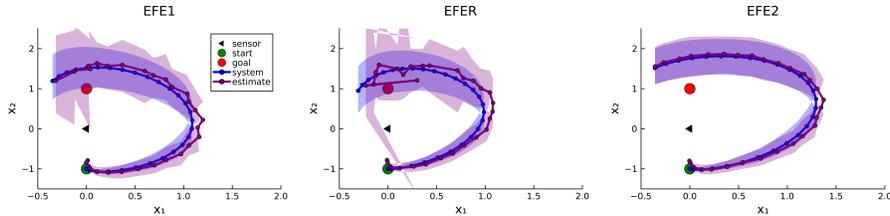
<sup>2</sup> Details and code at: <https://github.com/biaslab/IWAI2024-ambiguity>

the second-order Gaussian approximation. Figure 2 plots the value of the control objective function at every position in state-space, under a state covariance matrix of  $S_t = I$ . States close to the sensor station are red and will lead to high values under the control objective. Note that the area around the sensor station increases from EFE1 to EFER due to the curvature of the relative distance sensor. The white markers are the approximate minimizers for this choice of  $S_t$  matrix. Comparing EFER and EFE2, we can see that ambiguity increases the cost of being close to the sensor station.



**Fig. 2.** Value under three EFE functions over a plane: EFE1 is risk and ambiguity under a first-order Taylor approximation, EFER is risk only under a second-order Taylor approximation and EFE2 is both risk and ambiguity under a second-order Taylor approximation. White markers indicate minimizers. Note that each EFE function induces a different preference over states.

We ran 100 Monte Carlo experiments. Figure 3 plots the average trajectory of  $T = 30$  steps taken by the EFE1, EFER and EFE2 agents. Ribbons indicate the standard error of the mean at every time-point. Note that all agents avoid the sensor station, with EFE2 taking the widest curve (EFE1 and EFER turn at  $x_1 = 1.0$  while EFE2 turns at  $x_1 = 1.5$ ). EFE1 and EFER lose track of the robot in a number of experiments (like the model predictive controller in Sec. 2), leading to a more volatile average trajectory. EFE2 has the smoothest average trajectory, indicating that the ambiguity term helps planning.



**Fig. 3.** Trajectories of agents under three EFE functions, averaged over 100 Monte Carlo samples (ribbon is standard deviation of the mean). The robot starts at the green marker and must reach the red goal marker. All agents avoid the sensor station, with EFE2 taking the widest curve and having the smoothest average trajectory.

## 6 Discussion

One could argue that our analysis is more about model selection than inference, as each Gaussian approximation essentially constitutes a different generative model. In that sense, the experiments only indicate that richer approximations of nonlinear functions lead to better performance, which is not surprising. However, the result is more subtle than that since the unscented transform is richer than the first-order Taylor (produces a more accurate mean estimate [8]) but apparently still leads to constant ambiguity. No, the approximation must be sensitive to how the covariance matrix of the joint distribution over states and observations changes as a function of  $g$ 's curvature. It would be interesting to extend this work with parameter estimation, such as inferring the process noise covariance matrix using a Wishart distribution [15], or the state transition matrix with a Matrix-Normal distribution [1,12].

## 7 Conclusion

We examined active inference agents with linear Gaussian distributed dynamics and a non-linear measurement function. We found that the first-order Taylor series and unscented transform approximations to the non-linearly transformed states lead to expected free energy functions with ambiguity terms that are constant over states. A second-order Taylor approximation leads to a state-dependent ambiguity term, inducing a preference over states.

**Acknowledgments.** The author gratefully acknowledges financial support from the Eindhoven Artificial Intelligence Systems Institute (EAISI) at TU Eindhoven.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## A Appendix: proof of Lemma 1

*Proof.* The cross-entropy is split into two entropies that simplify according to:

$$\begin{aligned} & \mathbb{E}_{p(y_t, x_t | u_t; \mathcal{D}_k)} \left[ -\ln \frac{p(y_t, x_t | u_t; \mathcal{D}_k)}{p(x_t | u_t; \mathcal{D}_k)} \right] \\ &= - \int \mathcal{N} \left( \begin{bmatrix} x_t \\ y_t \end{bmatrix} \mid \begin{bmatrix} \bar{m}_t \\ \mu_t \end{bmatrix}, \begin{bmatrix} \bar{S}_t & \Gamma_t \\ \Gamma_t^\top & \Sigma_t \end{bmatrix} \right) \ln \mathcal{N} \left( \begin{bmatrix} x_t \\ y_t \end{bmatrix} \mid \begin{bmatrix} \bar{m}_t \\ \mu_t \end{bmatrix}, \begin{bmatrix} \bar{S}_t & \Gamma_t \\ \Gamma_t^\top & \Sigma_t \end{bmatrix} \right) d(y_t, x_t) \\ & \quad + \int \mathcal{N}(x_t | \bar{m}_t, \bar{S}_t) \ln \mathcal{N}(x_t | \bar{m}_t, \bar{S}_t) dx_t \end{aligned} \quad (39)$$

$$= \frac{D_x + D_y}{2} \ln(2\pi e) + \frac{1}{2} \ln \left| \begin{bmatrix} \bar{S}_t & \Gamma_t \\ \Gamma_t^\top & \Sigma_t \end{bmatrix} \right| - \left( \frac{D_x}{2} \ln(2\pi e) + \frac{1}{2} \ln |\bar{S}_t| \right) \quad (40)$$

$$= \frac{D_y}{2} \ln(2\pi e) + \frac{1}{2} \ln (|\bar{S}_t| \cdot |\Sigma_t - \Gamma_t^\top \bar{S}_t^{-1} \Gamma_t|) - \frac{1}{2} \ln |\bar{S}_t| \quad (41)$$

$$= \frac{D_y}{2} \ln(2\pi e) + \frac{1}{2} \ln |\Sigma_t - \Gamma_t^\top \bar{S}_t^{-1} \Gamma_t|. \quad (42)$$

## B Appendix: proof of Theorem 1

*Proof.* Plugging  $\Sigma_t, \Gamma_t$  from (31) into the result from Lemma 1, yields:

$$\frac{D_y}{2} \ln(2\pi e) + \frac{1}{2} \ln |\Sigma_t - \Gamma_t^\top \bar{S}_t^{-1} \Gamma| \quad (43)$$

$$\begin{aligned} &= \frac{D_y}{2} \ln(2\pi e) + \frac{1}{2} \ln |G_x(\bar{m}_t) \bar{S}_t G_x(\bar{m}_t)^\top + R - G_x(\bar{m}_t) \bar{S}_t^\top \bar{S}_t^{-1} \bar{S}_t G_x(\bar{m}_t)^\top| \\ &= \frac{D_y}{2} \ln(2\pi e) + \frac{1}{2} \ln |R|. \end{aligned} \quad (44)$$

The cancellation is due to  $\bar{S}_t$  being symmetric, i.e.,  $\bar{S}_t^\top \bar{S}_t^{-1} = \bar{S}_t \bar{S}_t^{-1} = I$ .

## C Appendix: proof of Theorem 2

*Proof.* Plugging  $\Sigma_t, \Gamma_t$  from (33) into the result from Lemma 1, yields:

$$\begin{aligned} \frac{D_y}{2} \ln(2\pi e) + \frac{1}{2} \ln |\Sigma_t - \Gamma_t^\top \bar{S}_t^{-1} \Gamma| &= \frac{D_y}{2} \ln(2\pi e) + \frac{1}{2} \ln |G_x(\bar{m}_t) \bar{S}_t G_x(\bar{m}_t)^\top| \quad (45) \\ &+ \frac{1}{2} \sum_{i=1}^{D_y} \sum_{j=1}^{D_y} e_i e_j^\top \text{tr} (G_{xx}^{(i)}(\bar{m}_t) \bar{S}_t G_{xx}^{(j)}(\bar{m}_t) \bar{S}_t) + R - G_x(\bar{m}_t) \bar{S}_t^\top \bar{S}_t^{-1} \bar{S}_t G_x(\bar{m}_t)^\top \\ &= \frac{D_y}{2} \ln(2\pi e) + \frac{1}{2} \ln \left| \frac{1}{2} \sum_{i=1}^{D_y} \sum_{j=1}^{D_y} e_i e_j^\top \text{tr} (G_{xx}^{(i)}(\bar{m}_t) \bar{S}_t G_{xx}^{(j)}(\bar{m}_t) \bar{S}_t) + R \right|. \end{aligned} \quad (46)$$

The covariance matrix  $\bar{S}_t$  is symmetric, i.e.,  $\bar{S}_t^\top \bar{S}_t^{-1} = \bar{S}_t \bar{S}_t^{-1} = I$ . Note that the Hessian  $G_{xx}^{(i)}(\bar{m}_t)$  depends on the inferred mean of the predicted state  $\bar{m}_t$ , meaning that ambiguity is not constant over state-space.

## D Appendix: proof of Theorem 3

*Proof.* Plugging  $\mu_t, \Sigma_t, \Gamma_t$  from (37) into the log-determinant term from the result in Lemma 1, gives:

$$\begin{aligned} \frac{1}{2} \ln |\Sigma_t - \Gamma_t^\top \bar{S}_t^{-1} \Gamma| &= \frac{1}{2} \ln \left| \sum_{i'=0}^{2D_x} w_{i'} (g(\chi_{i'}) - \mu_t)(g(\chi_{i'}) - \mu_t)^\top + R \right. \\ &\left. - \left( \sum_{i=0}^{2D_x} w_i (\chi_i - \bar{m}_t)(g(\chi_i) - \mu_t)^\top \right)^\top \bar{S}_t^{-1} \left( \sum_{j=0}^{2D_x} w_j (\chi_j - \bar{m}_t)(g(\chi_j) - \mu_t)^\top \right) \right|. \end{aligned} \quad (47)$$

The second term can be re-arranged to:

$$\begin{aligned} &\left( \sum_{i=0}^{2D_x} w_i (\chi_i - \bar{m}_t)(g(\chi_i) - \mu_t)^\top \right)^\top \bar{S}_t^{-1} \left( \sum_{j=0}^{2D_x} w_j (\chi_j - \bar{m}_t)(g(\chi_j) - \mu_t)^\top \right) \\ &= \sum_{i=0}^{2D_x} \sum_{j=0}^{2D_x} w_i (g(\chi_i) - \mu_t)(\chi_i - \bar{m}_t)^\top \bar{S}_t^{-1} w_j (\chi_j - \bar{m}_t)(g(\chi_j) - \mu_t)^\top. \end{aligned} \quad (48)$$

Note that for  $j = 0$ ,  $(\chi_j - \bar{m}_t) = (\bar{m}_t - \bar{m}_t) = 0$ . Let  $D_\lambda = D_x + \lambda$ . For  $j \geq 1$ :

$$(\chi_i - \bar{m}_t)^\top \bar{S}_t^{-1} w_j (\chi_j - \bar{m}_t) \quad (49)$$

$$\begin{aligned} &= (\bar{m}_t + (-1)^i \sqrt{D_\lambda} [\sqrt{\bar{S}_t}]_i - \bar{m}_t)^\top \bar{S}_t^{-1} \frac{1}{D_\lambda} (\bar{m}_t + (-1)^j \sqrt{D_\lambda} [\sqrt{\bar{S}_t}]_j - \bar{m}_t) \\ &= \frac{1}{D_\lambda} (\sqrt{D_\lambda})^2 (-1)^{i+j} [\sqrt{\bar{S}_t}]_i^\top \bar{S}_t^{-1} [\sqrt{\bar{S}_t}]_j \end{aligned} \quad (50)$$

$$= (-1)^{i+j} [\sqrt{\bar{S}_t}]_i^\top \bar{S}_t^{-1} [\sqrt{\bar{S}_t}]_j. \quad (51)$$

Furthermore, note that column selection  $[\cdot]_i$  is equivalent to right-multiplication with a canonical basis vector  $e_i$ ;

$$[\sqrt{\bar{S}_t}]_i^\top \bar{S}_t^{-1} [\sqrt{\bar{S}_t}]_j = (\sqrt{\bar{S}_t} e_i)^\top \bar{S}_t^{-1} (\sqrt{\bar{S}_t} e_j) = e_i^\top \sqrt{\bar{S}_t}^\top \bar{S}_t^{-1} \sqrt{\bar{S}_t} e_j. \quad (52)$$

Since  $\bar{S}_t$  is a normal matrix, the eigendecomposition  $\bar{S}_t = V \Omega V^{-1}$  generates an orthonormal eigenvector matrix  $V$ , implying  $V^{-1} = V^\top$ , and a diagonal matrix of eigenvalues  $\Omega$ . This means that  $\sqrt{\bar{S}_t} = V \Omega^{1/2} V^{-1}$ , and that:

$$\sqrt{\bar{S}_t}^\top \bar{S}_t^{-1} \sqrt{\bar{S}_t} = (V \Omega^{1/2} V^{-1})^\top V \Omega^{-1} V^{-1} (V \Omega^{1/2} V^{-1}) \quad (53)$$

$$= V \Omega^{1/2} V^{-1} V \Omega^{-1} V^{-1} V \Omega^{1/2} V^{-1} \quad (54)$$

$$= V V^{-1} = I. \quad (55)$$

Therefore,  $e_i^\top I e_j$  will be 1 for all  $i = j$  and 0 for  $i \neq j$ . We can thus identify two cases in the double sum in (48), one of which is always 0:

$$\sum_{i=0}^{2D_x} \sum_{j=0}^{2D_x} w_i (g(\chi_i) - \mu_t) (\chi_i - \bar{m}_t)^\top \bar{S}_t^{-1} w_j (\chi_j - \bar{m}_t) (g(\chi_j) - \mu_t)^\top \quad (56)$$

$$\begin{aligned} &= \sum_{i=0}^{2D_x} \sum_{j=i}^{2D_x} w_i (g(\chi_i) - \mu_t) (-1)^{(i+j)} 1 (g(\chi_j) - \mu_t)^\top \\ &\quad + \sum_{i=0}^{2D_x} \sum_{j \neq i}^{2D_x} w_i (g(\chi_i) - \mu_t) (-1)^{(i+j)} 0 (g(\chi_j) - \mu_t)^\top \end{aligned} \quad (57)$$

$$= \sum_{i=0}^{2D_x} w_i (g(\chi_i) - \mu_t) (g(\chi_i) - \mu_t)^\top, \quad (58)$$

where the  $(-1)^{(i+j)}$  drops out because for  $i = j$ ,  $i + j$  will always be even. One may now recognize that (47) has two terms that cancel each other:

$$\begin{aligned} &\frac{1}{2} \ln \left| \sum_{i'=0}^{2D_x} w_{i'} (g(\chi_{i'}) - \mu_t) (g(\chi_{i'}) - \mu_t)^\top + R - \sum_{i=0}^{2D_x} w_i (g(\chi_i) - \mu_t) (g(\chi_i) - \mu_t)^\top \right| \\ &= \frac{1}{2} \ln |R|. \end{aligned} \quad (59)$$

Using this result and Lemma 1, we have proven Theorem 3.

## References

1. Barber, D., Chiappa, S.: Unified inference for variational Bayesian linear Gaussian state-space models. *Advances in Neural Information Processing Systems* **19** (2006)
2. Conant, R.C., Ross Ashby, W.: Every good regulator of a system must be a model of that system. *International Journal of Systems Science* **1**(2), 89–97 (1970)
3. Cover, T.M.: *Elements of information theory*. John Wiley & Sons (1999)
4. Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V., Friston, K.: Active inference on discrete state-spaces: A synthesis. *Journal of Mathematical Psychology* **99**, 102447 (2020)
5. Friston, K.: The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience* **11**(2), 127–138 (2010)
6. Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., Pezzulo, G.: Active inference: a process theory. *Neural Computation* **29**(1), 1–49 (2017)
7. Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., Pezzulo, G.: Active inference and epistemic value. *Cognitive Neuroscience* **6**(4), 187–214 (2015)
8. Gustafsson, F., Hendeby, G.: Some relations between extended and unscented Kalman filters. *IEEE Transactions on Signal Processing* **60**(2), 545–555 (2011)
9. Julier, S.J., Uhlmann, J.K.: Unscented filtering and nonlinear estimation. *Proceedings of the IEEE* **92**(3), 401–422 (2004)
10. Koudahl, M.T., Kouw, W.M., de Vries, B.: On epistemics in expected free energy for linear Gaussian state space models. *Entropy* **23**(12), 1565 (2021)
11. van de Laar, T., Koudahl, M., van Erp, B., de Vries, B.: Active inference and epistemic value in graphical models. *Frontiers in Robotics and AI* **9**, 794464 (2022)
12. Luttinen, J.: Fast variational Bayesian linear state-space model. In: *European Conference on Machine Learning*, pp. 305–320. Springer (2013)
13. Millidge, B., Tschantz, A., Buckley, C.L.: Whence the expected free energy? *Neural Computation* **33**(2), 447–482 (2021)
14. Särkkä, S.: *Bayesian filtering and smoothing*, vol. 3. Cambridge University Press (2013)
15. Sarkka, S., Nummenmaa, A.: Recursive noise adaptive Kalman filtering by variational Bayesian approximations. *IEEE Transactions on Automatic Control* **54**(3), 596–600 (2009)
16. Schwartenbeck, P., Passetker, J., Hauser, T.U., FitzGerald, T.H., Kronbichler, M., Friston, K.J.: Computational mechanisms of curiosity and goal-directed exploration. *eLife* **8**, e41703 (2019)
17. Şenöz, İ., van de Laar, T., Bagaev, D., de Vries, B.: Variational message passing and local constraint manipulation in factor graphs. *Entropy* **23**(7), 807 (2021)
18. Tschantz, A., Seth, A.K., Buckley, C.L.: Learning action-oriented models through active inference. *PLOS Computational Biology* **16**(4), e1007805 (2020)
19. Zamiri-Jafarian, Y., Plataniotis, K.N.: A Bayesian surprise approach in designing cognitive radar for autonomous driving. *Entropy* **24**(5), 672 (2022)