

Towards Stochastic Fault-Tolerant Control Using Precision Learning and Active Inference

Baioumy, Mohamed; Pezzato, Corrado; Corbato, Carlos Hernández; Hawes, Nick; Ferrari, Riccardo

DOI

[10.1007/978-3-030-93736-2_48](https://doi.org/10.1007/978-3-030-93736-2_48)

Publication date

2022

Document Version

Final published version

Published in

Machine Learning and Principles and Practice of Knowledge Discovery in Databases

Citation (APA)

Baioumy, M., Pezzato, C., Corbato, C. H., Hawes, N., & Ferrari, R. (2022). Towards Stochastic Fault-Tolerant Control Using Precision Learning and Active Inference. In M. Kamp, M. Kamp, I. Koprinska, & E. A. (Eds.), *Machine Learning and Principles and Practice of Knowledge Discovery in Databases: Proceedings of the International Workshops of ECML PKDD 2021* (pp. 681-691). (Communications in Computer and Information Science; Vol. 1524 CCIS). Springer. https://doi.org/10.1007/978-3-030-93736-2_48

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Towards Stochastic Fault-Tolerant Control Using Precision Learning and Active Inference

Mohamed Baioumy¹(✉), Corrado Pezzato², Carlos Hernández Corbato², Nick Hawes¹, and Riccardo Ferrari³

¹ Oxford Robotics Institute, University of Oxford, Oxford, UK
{mohamed,nickh}@robots.ox.ac.uk

² Cognitive Robotics, Delft University of Technology, Delft, The Netherlands
{c.pezzato,c.h.corbato}@tudelft.nl

³ Delft Center for Systems and Control, Delft University of Technology, Delft, The Netherlands
{r.ferrari,m.wisse}@tudelft.nl

Abstract. This work presents a fault-tolerant control scheme for sensory faults in robotic manipulators based on active inference. In the majority of existing schemes a binary decision of whether a sensor is healthy (functional) or faulty is made based on measured data. The decision boundary is called a threshold and it is usually deterministic. Following a faulty decision, fault recovery is obtained by excluding the malfunctioning sensor. We propose a stochastic fault-tolerant scheme based on active inference and precision learning which does not require a priori threshold definitions to trigger fault recovery. Instead, the sensor precision, which represents its health status, is learned online in a model-free way allowing the system to gradually, and not abruptly exclude a failing unit. Experiments on a robotic manipulator show promising results and directions for future work are discussed.

1 Introduction

Safety is paramount for autonomous systems designed for operating in the real world. External dangers in the environment such as steep and slippery terrain encountered by planetary rovers [15] can compromise entire missions. In addition to external dangers, internal system components can also fail and possibly lead to dangerous outcomes if a proper fault-tolerant (FT) control scheme is not present. Building systems that are robust to the presence of faulty components, such as sensors and actuators, is addressed in the FT literature [8, 24, 32]. Generally speaking, FT control consists of *fault detection*, which provides a signal representing whether a system component is faulty; *fault isolation*, which identifies the exact faulty component, and *fault recovery*, which typically contains a switching or a re-tuning procedure of the running controllers to accommodate for the fault.

Several methods are available for fault detection, but model-based methods are among the most powerful and appealing, as they provide theoretical guarantees [8]. These methods rely on monitoring system outputs using mathematical models to generate ‘symptoms’ called *residual signals*. These signals are then compared to carefully designed detection *thresholds*: the sensor is ‘faulty’ if a threshold is exceeded or ‘healthy’ otherwise. To recover from a fault, the recovery actions are usually performed through controller reconfiguration, that entails adapting the controller parameters, or switching to another controller or to backup sensors and actuators [22]. When modelling external dangers or monitoring faulty systems, robust detection thresholds are essential. Robust thresholds used in existing work (such as [7] or [32]) are often *deterministic*, but this is sub-optimal. For instance, if the safety threshold for a rover on a slippery terrain slope is 15°C , this means that a slope of 14.9 is safe but 15.1 is unsafe. Additionally, a slope of 15.1°C and 40°C are ‘equally unsafe’.

In this paper we build upon two ideas in the literature. First, the usage of a stochastic fault tolerant formulation (e.g. [9,30]). This allows the agent to overcome the issues mentioned above. Additionally, we leverage an unbiased active inference controller (u-AIC) [3], evolved from previous active inference controllers (AIC) [1,6,25]. Active inference is a promising framework for FT control which has already been shown to facilitate fault-detection, isolation and recovery for robotic systems with sensory faults [3,26].

Besides fault tolerance, active inference showed promising performance in many control and state-estimation problems in robotics [16,17]. Particularly interesting are the works on robot arm control [23,25,29], which highlighted the adaptive properties of active inference. Active inference also shares similarities with the control as inference framework [18]. A more extensive analysis of active inference and its relation to control as inference can be found in [14,20].

The main contribution of this paper is a FT controller for robot manipulators with sensory faults based on unbiased active inference with a stochastic decision boundary. Unlike previous work [3], here we model the precision (inverse covariance) of each sensor in our system and determine the probability of the sensor being healthy to be proportional to its precision. Our approach allows for fault-tolerant behaviour without needing any threshold definition a priori, and without the need to design additional ad-hoc recovery mechanisms. Finally, this work can be used stand-alone or in conjunction with other methods for fault-detection and isolation in order to estimate the faults.

2 Problem Statement and Background

The FT scheme in this paper is derived for a class of systems, namely serial robot manipulators equipped with sensors for joint position and velocity, and end-effector location. In the following, the problem and the setup are described, and some background knowledge on u-AIC for torque control from [3] is presented.

Problem Setup. Consider a robotic manipulator with state $\mathbf{x}$ comprising of its joint positions and velocities $\mathbf{x} = [q \ \dot{q}]^\top$. The available sensors provide noisy joint position and velocities $\mathbf{y}_q, \mathbf{y}_{\dot{q}}$ readings. In addition, the end-effector Cartesian position $\mathbf{y}_v$ is available through a visual sensor. The system's output is represented by $\mathbf{y} = [\mathbf{y}_q, \mathbf{y}_{\dot{q}}, \mathbf{y}_v] \in \mathbb{R}^d$. The proprioceptive sensors and the visual sensor are affected by zero mean Gaussian noise $\boldsymbol{\eta} = [\boldsymbol{\eta}_q, \boldsymbol{\eta}_{\dot{q}}, \boldsymbol{\eta}_v]$. Additionally, the visual sensor is affected by barrel distortion. The system is controlled through an u-AIC [3] which steers the robot arm to a (changing) desired configuration in joint space $\boldsymbol{\mu}_d$, providing the control input $\mathbf{u} \in \mathbb{R}^m$ as torques to the joints.

Background: Unbiased Active Inference Controller. In this section we briefly describe the u-AIC as introduced in [3], to which an interested reader is referred for more details on the derivations of the following equations. The novel FT method presented in this paper in Sect. 3 builds upon the u-AIC, but instead of employing an ad-hoc hard update of the precision of a faulty sensor after fault detection, it relies on online precision learning during operations.

Let us consider $\mathbf{x} = [q \ \dot{q}]^\top$ and let us define a probabilistic model where actions are modelled explicitly:

$$p(\mathbf{x}, \mathbf{u}, \mathbf{y}_v, \mathbf{y}_q, \mathbf{y}_{\dot{q}}) = \underbrace{p(\mathbf{u}|\mathbf{x})}_{\text{control}} \underbrace{p(\mathbf{y}_v|\mathbf{x})p(\mathbf{y}_q|\mathbf{x})p(\mathbf{y}_{\dot{q}}|\mathbf{x})}_{\text{observation model}} \underbrace{p(\mathbf{x})}_{\text{prior}} \quad (1)$$

Note that with the u-AIC the information about the desired goal to be reached is encoded in the distribution $p(\mathbf{u}|\mathbf{x})$. In this paper, as in [1], we assume that an accurate dynamic model of the system is not available to keep the solution system agnostic and to highlight once again the adaptability of the controller.

The u-AIC aims at finding the posterior over states as well as the posterior over actions $p(\mathbf{x}, \mathbf{u}|\mathbf{y}_v, \mathbf{y}_q)$. The posteriors are approximated using a variational distribution $Q(\mathbf{x}, \mathbf{u})$. We can make use of the mean-field assumption ($Q(\mathbf{x}, \mathbf{u}) = Q(\mathbf{x})Q(\mathbf{u})$) and the Laplace approximation, and assume the posterior over the state $\mathbf{x}$ Gaussian with mean $\boldsymbol{\mu}_x$ [13]. Similarly for the actions, the posterior $\mathbf{u}$ is assumed Gaussian with mean $\boldsymbol{\mu}_u$. By defining the Kullback-Leibler divergence between the variational distribution and the true posterior, one can derive an expression for the so-called free-energy F as [3]:

$$F = -\ln p(\boldsymbol{\mu}_u, \boldsymbol{\mu}_x, \mathbf{y}_v, \mathbf{y}_q, \mathbf{y}_{\dot{q}}) + C \quad (2)$$

Considering Eq. (1) and assuming Gaussian distributions, F becomes:

$$F = \frac{1}{2}(\boldsymbol{\varepsilon}_{y_q}^\top \Sigma_{y_q}^{-1} \boldsymbol{\varepsilon}_{y_q} + \boldsymbol{\varepsilon}_{y_{\dot{q}}}^\top \Sigma_{y_{\dot{q}}}^{-1} \boldsymbol{\varepsilon}_{y_{\dot{q}}} + \boldsymbol{\varepsilon}_{y_v}^\top \Sigma_{y_v}^{-1} \boldsymbol{\varepsilon}_{y_v} + \boldsymbol{\varepsilon}_x^\top \Sigma_x^{-1} \boldsymbol{\varepsilon}_x + \boldsymbol{\varepsilon}_u^\top \Sigma_u^{-1} \boldsymbol{\varepsilon}_u + \ln |\Sigma_u \Sigma_{y_q} \Sigma_{y_{\dot{q}}} \Sigma_{y_v} \Sigma_x|) + C, \quad (3)$$

The terms $\boldsymbol{\varepsilon}_{y_q} = \mathbf{y}_q - \boldsymbol{\mu}$, $\boldsymbol{\varepsilon}_{y_{\dot{q}}} = \mathbf{y}_{\dot{q}} - \boldsymbol{\mu}'$, $\boldsymbol{\varepsilon}_{y_v} = \mathbf{y}_v - \mathbf{g}_v(\boldsymbol{\mu})$ are the sensory prediction errors respectively for position, velocity, and visual sensory inputs. The controller represents the states internally as $\boldsymbol{\mu}_x = [\boldsymbol{\mu}, \boldsymbol{\mu}']^\top$. The relation

between internal state and observation is expressed through the generative model of the sensory input $\mathbf{g} = [\mathbf{g}_q, \mathbf{g}_{\dot{q}}, \mathbf{g}_v]$. Position and velocity encoders directly measure the state, thus $\mathbf{g}_q$ and $\mathbf{g}_{\dot{q}}$ are linear (identity) mappings. To define $\mathbf{g}_v$, instead, we use a *Gaussian Process Regression* (GPR). This is particularly useful because we can model the noisy and distorted sensory input from the camera, and at the same time we can compute a closed form for the derivative of the process with respect to the beliefs $\boldsymbol{\mu}$, required for the state update laws. For details, see [3].

Additionally, ε_u is the prediction error on the control action while ε_x is the prediction error on the state. The latter is computed considering a prediction of the state $\hat{\mathbf{x}}$ at the current time-step such that $\varepsilon_x = (\boldsymbol{\mu}_x - \hat{\mathbf{x}})$. The prediction is a deterministic value $\hat{\mathbf{x}} = [\hat{q} \ \hat{\dot{q}}]^\top$ which can be computed in the same fashion as the prediction step of, for instance, a Kalman filter. The prediction is approximated propagating forward in time the current state belief using the following simplified discrete time model:

$$\hat{\mathbf{x}}_{k+1} = \begin{bmatrix} I & I\Delta t \\ 0 & I \end{bmatrix} \boldsymbol{\mu}_{x,k} \quad (4)$$

where I represents an unitary matrix of suitable size. This form assumes that the position of each joint is thus computed as the discrete time integral of the velocity, using a first-order Euler scheme. This approximation can be avoided if a better dynamic model of the system is available, and in that case predictions can be made using the model itself. Finally, by choosing the distribution $p(\mathbf{u}|\mathbf{x})$ to be Gaussian with mean $f^*(\boldsymbol{\mu}_x, \boldsymbol{\mu}_d)$, we can steer the systems toward the target $\boldsymbol{\mu}_d$ without biasing the state estimation. This results in $\varepsilon_u = (\boldsymbol{\mu}_u - f^*(\boldsymbol{\mu}_x, \boldsymbol{\mu}_d))$.

In the u-AIC state-estimation and control are achieved using gradient descent the free-energy. This leads to:

$$\dot{\boldsymbol{\mu}}_u = -\kappa_u \frac{\partial F}{\partial \boldsymbol{\mu}_u}, \quad \dot{\boldsymbol{\mu}}_x = -\kappa_\mu \frac{\partial F}{\partial \boldsymbol{\mu}_x}, \quad (5)$$

where κ_u and κ_μ are the gradient descent step sizes.

3 Precision Learning for Fault-Tolerant Control

In previous work [3], the u-AIC is used in combination with an established FT approach to achieve fault detection and recovery. In particular, the sensory prediction errors in the free-energy are used as residual signals for fault detection purposes. The statistical properties of the residuals are analysed offline and healthy boundaries are defined. At runtime, a healthy residual set is computed and if the current residual is outside the admissible set, the relative sensor is marked as faulty. When a fault is detected, the precision (or inverse covariance) of the sensor is abruptly set to zero, that is $P = \Sigma^{-1} = 0$, to exclude that sensor from the optimization of the free-energy. This idea is summarised in Fig. 1.

In this work, we propose a different approach to achieve fault recovery through online precision learning with u-AIC instead ad-hoc hard switches in the controller's parameters. Figure 2 shows the difference with respect to [3].




Fig. 1. Fault-tolerant pipeline from [3]. The term Σ_f^{-1} represents the precision of the detected faulty sensor.




Fig. 2. New fault-tolerant pipeline with precision learning, in contrast to previous work [3] from Fig. 1.

Learning Sensory Precision. For a sensor y , we can update an inverse precision matrix Σ_y^{-1} using gradient descent on F as done in [1, 2]:

$$\dot{\Sigma}_y^{-1} = -\kappa_\sigma \frac{\partial F}{\partial \Sigma_y^{-1}}. \quad (6)$$

However, we need to ensure that precision remains a positive number. Performing gradient descent does not inherently guarantee that.

First, consider a one-dimensional problem where state x and observation y are scalars. The observations is affected by zero-mean Gaussian noise with a variance of σ^2 (also a scalar). The scalar precision is defined as the inverse variance $\omega = 1/\sigma^2$. As explained, performing gradient descent on the free-energy with respect to ω may result in it being negative. A simple solution is to perform a reparameterization with a strictly positive function such as an exponential. I.e. we assume that $\omega = \exp \zeta$ and we perform gradient descent on ζ :

$$\dot{\zeta} = -\kappa_\zeta \frac{\partial F}{\partial \zeta} \quad (7)$$

where κ_ζ is the gradient step-size. Another way is to set a lower bound on the variance (as done in [5]). Both methods ensure the variance being positive.

Diagonal Precision Matrix. Guaranteeing a positive semi-definite matrix in an n -dimensional case is not as straightforward. However, in the context of a robotic manipulator, one can reasonably assume that the observation noise on

each sensor is independent [1,23,25]. This means that the covariance (and precision) matrices are diagonal.

$$P = \begin{bmatrix} \omega_1 & & & \\ & \omega_2 & & \\ & & \dots & \\ & & & \omega_n \end{bmatrix}$$

Given this assumption, every element on the diagonal is positive and can be updated in the same fashion as the scalar case (Eq. (7)).

Fault-Tolerant Control as Precision Learning. Consider the sum of the sensory prediction errors in the free-energy from Eq. (3):

$$F = \frac{1}{2}(\epsilon_{y_q}^\top \Sigma_{y_q}^{-1} \epsilon_{y_q} + \epsilon_{y_{\dot{q}}}^\top \Sigma_{y_{\dot{q}}}^{-1} \epsilon_{y_{\dot{q}}} + \epsilon_{y_v}^\top \Sigma_{y_v}^{-1} \epsilon_{y_v} + \dots) + C, \quad (8)$$

Intuitively, when a sensor is faulty, the related sensory prediction error will necessarily be higher since sensory readings and internal beliefs will drift away. After a fault, the estimated precision through our precision learning scheme will be much lower than the original $P = \Sigma^{-1}$. Thus its weight in the free-energy F , and so in the state-estimation and control equations as in Eq. (5) will naturally become lower than the other healthy sensors. Its weight essentially adjusts *proportionally to the degree of the sensor being faulty*. Note that this allows for automatic fault recovery but it does not provide explicit fault detection. In case the latter is needed for a potential user or an additional supervisory system, traditional techniques can be used as the one presented in [3] in conjunction with precision learning.

FT control for sensory faults can now be done using precision learning in several ways. The first way is to use it as a stand-alone and activate precision learning for all sensors during operation. In this case, no other methods are needed, no thresholds are designed and the recovery emerges naturally. As mentioned before, the drawback is that the users can not be ‘alerted’ for the presence of a fault (since there is no explicit fault-detection). The second way, which addresses this issue, is to use an established algorithm for fault detection (such as the one presented in [3]) and then, only after a fault is detected, allow precision update.

Interestingly, performing precision learning as presented in this section can make the state-estimation noisier since the agents only relies on the current observation (rather than a batch of last k observations) for the update and both the uncertainty of the state and precision are not quantified. An additional approach would then be to consider the last k observations for the update, but this is out of the scope of this work.

To summarise, the precision learning in this paper can either be activated at all times or *only after a fault is detected*. Activating the precision learning at all times with a small step-size for the gradient seems to work best.

4 Results

We apply the methods in Sect. 3 on a 2-DOF robotic manipulator. We test three scenarios: a) precision learning at all times, b) precision learning only when a fault is detected and c) a deterministic update as done in [3]. Note that the latter has access to a model and uses data to determine a threshold offline. This is not the case for the first two options where only model-free precision learning is performed. The results are summarized in the Table 1. In the simulations, the sensors are injected with zero-mean Gaussian noise. The standard deviation of the noise for encoders and velocity sensors is set to $\sigma_q = \sigma_{\dot{q}} = 0.001$, while the one for the camera is set to $\sigma_v = 0.01$. The camera is also affected by barrel distortion with coefficients $K_1 = -1.5e^{-3}$, $K_2 = 5e^{-6}$, $K_3 = 0$ (values are similar to work from [19,28]).

Table 1. Mean Squared Error (MSE) for different methods of fault-tolerant control. PL indicates precision learning

	Joints with encoder fault	Joints without encoder fault
No fault-tolerance	0.0036	0.0020
PL at all time	5.422 e−5	4.527 e−5
PL + fault-detection	6.097 e−5	4.134 e−5
Deterministic fault recovery	0.5946 e−5	0.3579 e−5

The agents starts in configuration $\mathbf{x}_0$, then moves to the targets $\mathbf{x}_1$ and $\mathbf{x}_2$. At $t = 8$ s a fault is injected. The encoder fault is such that the output related to the first joint freezes. For a discrete step k it holds then $\mathbf{y}_q(k) = [q_1(k_f), q_2(k)]^\top$ for $k \geq k_f$ and $k_f = 8$. The fault detection and recovery of such a fault, as well as the system’s response, are reported below in Fig. 3.

As seen in Fig. 3, the system is not able to reach the set-point after the occurrence of the fault if online precision update is not allowed. The robot arm reaches a different configuration to minimise the free-energy, which is built fusing the sensory information from the (faulty) encoders and the (healthy) camera. However, when the faulty encoder is adjusted using precision learning, the agent is able to reach the final configuration.

Figure 3 reports the results when precision learning is being done during the full operational time. Alternatively, one could only use precision learning when a fault is detected. This yields a response that is almost identical. The Mean Squared Error (MSE) between the belief and the true position ($\boldsymbol{\mu}_x - \mathbf{x}$) is computed on a sample of test runs and reported in the Table 1. The results are reported for both the joint whose encoder is faulty, and joints with healthy encoders. In both cases, hard update of the precision to zero has the lowest MSE; however, the approaches based on precision learning do not require any




Fig. 3. System's response in the case of encoder fault with and without precision learning applied at all times. The fault is injected at $t = 8$ s and indicated with a dot-dashed line.

previous information or a threshold definition thus it is simpler to implement. Yet, precision learning has a satisfactory performance while accommodating a sensory fault.

5 Improving Precision Learning: A Discussion

In this paper, we perform a simple modification to the unbiased active inference controller: adding precision learning for all sensors. We show that this results in stochastic fault-tolerance to sensory faults, i.e. the precision of a faulty sensor will decrease automatically making its relative weight in the control and estimation laws smaller. This eliminates the need to learn a threshold from data offline. Additionally, no ad-hoc recovery action is required. The controller automatically adjusts to the new precision.

In the experiments, we compared precision learning to a state-of-the-art method. Precision learning was an order of magnitude worse in performance but still satisfactory. Note that precision learning did *not* require any data or training offline to determine thresholds or recovery strategies. Finally, precision learning performs stochastic fault-detection rather than deterministic.

Most importantly, this approach based on precision learning can be improved in many ways. First, rather than computing a point-mass estimate, we can explicitly model the precision as a random variable and perform inference on it.

We can perform Bayesian inference by modelling the precision as a random variable and computing a posterior over it. In the one dimensional case we use a Gamma prior on the precision ω as

$$\Gamma(\omega; a, b) = \frac{b^a}{\Gamma(a)} \omega^{a-1} e^{-\omega b}.$$

Given that the observation model is Gaussian, this choice is beneficial since it is the conjugate prior [4, 21], where a and b are the parameters of the distribution and $\Gamma(a) = (a-1)!$ is a factorial function. For example, $\Gamma(5) = 4! = 24$. Now to compute the posterior, we multiply the prior with the Gaussian likelihood model of $p(y|\omega)$ and obtain the posterior which is also a Gamma distribution as shown below.

$$\begin{aligned} p(\omega) &= \Gamma(\omega; a, b) \propto \omega^{a-1} e^{-\omega b} \\ p(\omega|y) &\propto p(y|\omega)p(\omega) \propto \omega^{0.5+a-1} e^{-\omega(b + \frac{(y-C)^2}{2})} \\ p(\omega|y) &= \Gamma(\omega; a + \frac{1}{2}, b + \frac{(y-C)^2}{2}) \end{aligned}$$

The last equation shows a simple update rule to modify the belief over the precision for every data point. In the optimization for the state, the following quantities are used: expected precision $\mathbb{E}[\omega] = a/b$, $Mode[\omega] = (a-1)/b$ and $Var[\omega] = a/b^2$. In the n -dimensional case, the same procedure can be done but with a Wishart distribution rather than a Gamma.

Additionally, we could use a batch of k observation to learn the precision rather than just one observation. Many approaches for covariance/precision estimation have been successful in robotics e.g. [27, 31, 33, 34]. Additionally, many other approaches within the active inference literature can be used for effective precision learning such as dynamic expectation maximization (DEM) [11, 12]. These will be explored and compared in future work.

6 Conclusions

This paper presents a fault-tolerant controller based on active inference. We model the precision (inverse covariance) of each sensor in our system and determine the probability of the sensor being healthy to be proportional to its precision. Rather than reasoning about whether a sensor is faulty or not, we reason about the degree to which the sensor is faulty. We present gradient based approaches to approximate the precision matrices of the system. The results show that the precision learning is a promising approach for fault-tolerant control. It allows for robust behaviour without needing any threshold definition a priori, without designing additional ad-hoc recovery mechanisms, and can be used stand-alone or in conjunction with other methods. The results using precision learning was satisfactory but an order of magnitude away from the to state-of-the-art. However, precision learning was not trained on data offline and performs a stochastic update. Bayesian methods can be used to improve the performance of the approach. Additionally, in all cases regarding precision learning, the performance can be improved by considering the last k observations rather than just one. Future work will address this.

References

1. Baioumy, M., Duckworth, P., Lacerda, B., Hawes, N.: Active inference for integrated state-estimation, control, and learning. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2021)
2. Baioumy, M., Mattamala, M., Duckworth, P., Lacerda, B., Hawes, N.: Adaptive manipulator control using active inference with precision learning. In: UKRAS (2020)
3. Baioumy, M., Pezzato, C., Ferrari, R., Corbato, C.H., Hawes, N.: Fault-tolerant control of robotic systems with sensory faults using unbiased active inference. In: European Control Conference (ECC) (2021)
4. Bishop, C.M.: Pattern Recognition and Machine Learning. Springer, Heidelberg (2006)
5. Bogacz, R.: A tutorial on the free-energy framework for modelling perception and learning. *J. Math. Psychol.* **76**, 198–211 (2017)
6. Buckley, C.L., Kim, C.S., McGregor, S., Seth, A.K.: The free energy principle for action and perception: a mathematical review. *J. Math. Psychol.* **81**, 55–79 (2017)
7. Budd, M., Lacerda, B., Duckworth, P., West, A., Lennox, B., Hawes, N.: Markov decision processes with unknown state feature values for safe exploration using gaussian processes. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (2020)
8. Chen, J., Patton, R.J.: Robust Model-Based Fault Diagnosis for Dynamic Systems. Springer, Boston (1999). <https://doi.org/10.1007/978-1-4615-5149-2>
9. Fang, S., Blanke, M., Leira, B.J.: Mooring system diagnosis and structural reliability control for position moored vessels. *Control. Eng. Pract.* **36**, 12–26 (2015)
10. Fox, C.W., Roberts, S.J.: A tutorial on variational Bayesian inference. *Artif. Intell. Rev.* **38**(2), 85–95 (2012)
11. Friston, K., Stephan, K., Li, B., Daunizeau, J.: Generalised filtering. *Math. Prob. Eng.* **2010** (2010)
12. Friston, K.J., Trujillo-Barreto, N., Daunizeau, J.: DEM: a variational treatment of dynamic systems. *Neuroimage* **41**(3), 849–885 (2008)
13. Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., Penny, W.: Variational free energy and the Laplace approximation. *Neuroimage* **34**(1), 220–234 (2007)
14. Imohiosen, A., Watson, J., Peters, J.: Active inference or control as inference? A unifying view. In: IWAI 2020. CCIS, vol. 1326, pp. 12–19. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64919-7_2
15. Inotsume, H., Kubota, T., Wettergreen, D.: Robust path planning for slope traversing under uncertainty in slip prediction. *IEEE Robot. Autom. Lett.* **5**(2), 3390–3397 (2020)
16. Lanillos, P., Cheng, G.: Adaptive robot body learning and estimation through predictive coding. In: IROS (2018)
17. Lanillos, P., Cheng, G.: Active inference with function learning for robot body perception. In: International Workshop on Continual Unsupervised Sensorimotor Learning (ICDL-Epirob) (2018)
18. Levine, S.: Reinforcement learning and control as probabilistic inference: tutorial and review. arXiv preprint [arXiv:1805.00909](https://arxiv.org/abs/1805.00909) (2018)
19. Marshall, M., Lipkin, H.: Kalman filtering visual servoing control law. In: IEEE Proceedings of the International Conference on Mechatronics and Automation (2014)

20. Millidge, B., Tschantz, A., Seth, A.K., Buckley, C.L.: On the relationship between active inference and control as inference. In: IWAI 2020. CCIS, vol. 1326, pp. 3–11. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64919-7_1
21. Murphy, K.P.: Machine Learning: A Probabilistic Perspective. MIT Press (2012)
22. Narendra, K.S., Balakrishnan, J.: Adaptive control using multiple models. *IEEE Trans. Autom. Control* **42**, 171–187 (1997)
23. Oliver, G., Lanillos, P., Cheng, G.: An empirical study of active inference on a humanoid robot. *IEEE Trans. Cogn. Develop. Syst.* (2021). <https://doi.org/10.1109/TCDS.2021.3049907>
24. Paviglianiti, G., Pierri, F., Caccavale, F., Mattei, M.: Robust fault detection and isolation for proprioceptive sensors of robot manipulators. *Mechatronics* **20**(1), 162–170 (2010)
25. Pezzato, C., Ferrari, R., Corbato, C.H.: A novel adaptive controller for robot manipulators based on active inference. *IEEE Robot. Autom. Lett.* **5**(2), 2973–2980 (2020)
26. Pezzato, C., Baioumy, M., Corbato, C.H., Hawes, N., Wisse, M., Ferrari, R.: Active inference for fault tolerant control of robot manipulators with sensory faults. In: IWAI 2020. CCIS, vol. 1326, pp. 20–27. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64919-7_3
27. Pfeifer, T., Lange, S., Protzel, P.: Dynamic covariance estimation—a parameter free approach to robust sensor fusion. In: 2017 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI), pp. 359–365. IEEE (2017)
28. Piepmeyer, J., McMurray, G., Lipkin, H.: Uncalibrated dynamic visual servoing. *IEEE Trans. Robot. Autom.* **20**, 143–147 (2004)
29. Pio-Lopez, L., Nizard, A., Friston, K., Pezzulo, G.: Active inference and robot control: a case study. *J. R. Soc. Interface* **13**(122), 20160616 (2016)
30. Rostampour, V., Ferrari, R.M., Teixeira, A.M., Keviczky, T.: Privatized distributed anomaly detection for large-scale nonlinear uncertain systems. *IEEE Trans. Autom. Control* **66**, 5299–5313 (2020)
31. Shetty, A., Gao, G.X.: Covariance estimation for GPS-LiDAR sensor fusion for UAVs. In: Proceedings of the 30th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS+ 2017, pp. 2919–2923 (2017)
32. Van, M., Wu, D., Ge, S., Ren, H.: Fault diagnosis in image-based visual servoing with eye-in-hand configurations using Kalman filter. *IEEE Trans. Ind. Electron.* **12**(6), 1998–2007 (2016)
33. Vega-Brown, W., Bachrach, A., Bry, A., Kelly, J., Roy, N.: Cello: a fast algorithm for covariance estimation. In: 2013 IEEE International Conference on Robotics and Automation, pp. 3160–3167. IEEE (2013)
34. Vega-Brown, W., Roy, N.: CELLO-EM: adaptive sensor models without ground truth. In: 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 1907–1914. IEEE (2013)