

BACHELOR THESIS  
ARTIFICIAL INTELLIGENCE

**Radboud University**



---

**Robot Localization and  
Navigation through Predictive  
Processing using LiDAR**

---

*Author:*  
Daniel Burghardt  
S1012240

*First supervisor:*  
dr. P. Lanillos  
Donders Institute for  
Brain, Cognition and  
Behaviour  
p.lanillos@donders.ru.nl

*Second supervisor:*  
dr. L. Ambrogioni  
Donders Institute for  
Brain, Cognition and  
Behaviour  
l.ambrogioni@donders.ru.nl



July 7, 2021

**Abstract.** Knowing the position of the robot in the world is crucial for navigation. Nowadays, Bayesian filters, such as Kalman and particle-based, are standard approaches in mobile robotics. Recently, end-to-end learning has allowed for scaling-up to high-dimensional inputs and improved generalization. However, there are still limitations to providing reliable laser navigation. Here we show a proof-of-concept of the predictive processing-inspired approach to perception applied for localization and navigation using laser sensors, without the need for odometry. We learn the generative model of the laser through self-supervised learning and perform both online state-estimation and navigation through stochastic gradient descent on the variational free-energy bound. We evaluated the algorithm on a mobile robot (TIAGo Base) with a laser sensor (SICK) in Gazebo. Results showed improved state-estimation performance when comparing to a state-of-the-art particle filter in the absence of odometry. Furthermore, conversely to standard Bayesian estimation approaches our method also enables the robot to navigate when providing the desired goal by inferring the actions that minimize the prediction error.

**Keywords:** Predictive Processing · Robot localization · Robot navigation · Laser sensor · LiDAR.

*Note* This Bachelor Thesis has been written in paper format and has been submitted to the 2021 International Workshop on Active Inference (IWA) in conjunction with ECML/PKDD 2021.

## 1 Introduction

Localization algorithms are part of our daily life and core for robotics. Recursive Bayesian estimation composes the current state-of-art in the field and has been essential for the development of localization, mapping, navigation and searching applications [18, 11]. Bayesian filters [5], e.g., Kalman and particle filters, are able to estimate the state of a system from noisy sensor observations formalized as a hidden Markov model. These approaches are useful also in the case of non-linear modeled systems and out-of-sequence measurements [3]. The particle filter (PF) is an approximate Bayesian method that tractably computes the posterior distribution of the state of any system given the observations. The state distribution is represented by individual particles, which are evaluated and weighted recursively. Particles with higher probability get bigger weights and are re-sampled into more particles in its neighborhood, whereas particles with smaller weights get fewer new samples close to them [13].

In recent years, novel approaches based on deep neural networks, have been proposed to improve localization using high-dimensional inputs. Regression solutions, for instance, compute the absolute position of the system from only visual information [9]. However, these methods have lower accuracy than previous approaches that exploit prior information, such as geometry [17]. In particular, LiDAR-based Navigation with representation learning (e.g., using autoencoders)

and reinforcement learning has shown downgraded performance in navigation tasks [8].

We describe how the predictive processing (PP) approach to perception [4, 12] can aid in localization and navigation tasks. Under PP, the agent, following the Free Energy Principle (FEP) [6], tries to minimize the error in the predicted observations by either performing corrective actions to match the expected internal state or by updating this internal state based on what it has experienced through the senses. In this work, we present a proof-of-concept based on the Pixel-Active Inference model [16], proposed for humanoid body perception and action, to perform laser-based localization and navigation without the need for odometry. Our approach combines the power of deep networks regression with variational Bayesian filtering to provide a better reliable state estimation than PFs in our proof-of-concept environment—See Fig. 1.

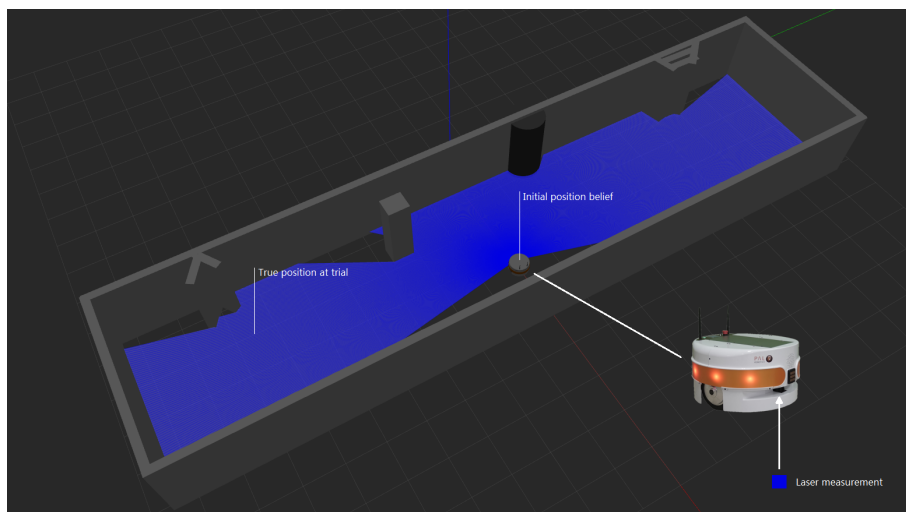


Fig. 1: Proof-of-concept environment designed for the experiments in Gazebo. In the localization experiment the robot true position is randomly set and the initial position belief is initialized to the center of the map. Laser range finder measurements are displayed as the blue shading.

## 1.1 Results

We evaluated our laser-based Active Inference algorithm against a particle filter [13] in the Gazebo simulator, using a commercial mobile robot with a laser rangefinder sensor (TIAGo Base, pmb-2) [2], interfaced with Robot Operating System (ROS) [1]. All experiments were conducted in a designed corridor-like map described in Fig. 1. Localization and navigation results are summarized in Fig. 2.




Fig. 2: Localization and navigation evaluation. Mean and standard deviation of the positioning absolute error of our model (FEP) compared against the PF algorithm. Furthermore, our model is able to perform navigation using the same Bayesian filtering framework. The green line shows the error the mean absolute error to the goal.

**Localization and estimation** Firstly, we evaluated the localization accuracy when initializing the robot to a random position in the space when a map was given. Thus, testing absolute positing with laser measurements in static situations. For the particle filter, the particle initial probabilities were randomly spread in the environment and reset at the beginning of every trial. In our algorithm, we initialized the initial belief in the center of the environment. We computed the ground truth positional error in 100 trials for 50 iterations. Our algorithm converged much more consistently to the true state over all trials, whereas the PF struggled to deliver consistent results, as shown in Fig. 2 by the rather large standard deviation in the blue shaded area. It is important to highlight that the PF is tuned for using the robot’s odometry. However, for the sake of fair comparison solely laser-sensor values were used as observations.

Secondly, we evaluated the localization performance when traversing the environment from one side to the other, by performing small teleports (to override odometry) to simulate robot movement while keeping the rotation angle constant. Results showed a more stable over time state estimation by our algorithm when compared to the PF, which seemed to suffer from the absence of odometry information.

**Navigation** For the assessment of the navigation algorithm’s performance, we ran an experiment consisting of 50 trials in which the robot had to navigate from a starting point to a goal position. The initial belief state was set to the robot’s initial true position, to evaluate the performance of navigation without the effects of localization in the first iterations. In every trial, both the initial position and the goal state of the robot were chosen randomly, with the constraint that they should be at least 12 meters (in Gazebo coordinates) apart from each other. The

task was considered complete when the robot got in a range of 0.8m from the target. The results are plotted in green in Fig. 2.

We observed that the robot initially quickly approximates the goal, with a big drop in the distance to the goal in the first couple of iterations. As it gets closer to the goal, the “velocity” of the robot (in the experiment described by the step sizes) decreases. This is a result of the diminishing gradient in every step of the algorithm, due to the stochastic gradient descent. Additionally, we computed the average number of iterations that it took the algorithm to get in the desired 0.8m range of the target. Over the 50 trials of similar travel distance ( $\sim 11m$  to  $\sim 13.5m$ ), the average number of iterations was 12.5. This number is naturally closely related to the optimal step size found.

## 2 Methods

**Robot** The TIAGo Base mobile robot uses the SICK TiM571 laser sensor, which has a 0.0m – 25m range and a 270° aperture angle. In all experiments we limited the robot’s movement to 2 degrees of freedom, i.e., moving forward, backward and sideways.

**Localization** We define the true state (position) of the robot at instant  $k$  as  $\mathbf{x}_k = (x, y) \in \mathbb{R}^2$  and the position belief of the robot as  $\tilde{\mathbf{x}}_k$ . We further define the observation  $\mathbf{o}_k$  as the laser measurements. Estimation is solved by computing the posterior distribution  $p(\mathbf{x}|\mathbf{o})$  by optimizing the Variational Free Energy (VFE). The algorithm is sketched in Fig. 3a. Under the mean-field and Laplace approximation this is equivalent to minimizing the error between the sensory input  $\mathbf{o}_k$  and the predicted sensory input  $\hat{\mathbf{o}}_k$ . While regression approaches (Fig. 3b) compute the pose directly from the visual input, stochastic neural filtering continuously refines the state through an error signal. We perform state estimation through perceptual inference, minimizing the VFE as follows:

$$\tilde{\mathbf{x}} = \underset{\tilde{\mathbf{x}}}{\operatorname{argmin}} F(\tilde{\mathbf{x}}, \mathbf{o}) \rightarrow \tilde{\mathbf{x}}_{k+1} = \tilde{\mathbf{x}}_k + \alpha \partial_{\tilde{\mathbf{x}}} g(\tilde{\mathbf{x}}_k) \Sigma_{\mathbf{o}}^{-1} (\mathbf{o}_k - g(\tilde{\mathbf{x}}_k)) \quad (1)$$

Where  $\alpha$  is the step size and  $\partial_{\tilde{\mathbf{x}}}$  denotes the derivative with respect to  $\tilde{\mathbf{x}}$ . This is computed iteratively using gradient descent on the prediction error—sensor measurement  $\mathbf{o}_k$  minus the predicted sensory input  $g(\tilde{\mathbf{x}}_k)$ —weighted by the variance  $\Sigma_{\mathbf{o}}$ . Both the predicted observations and the partial derivative of the error are computed by means of a deep neural network forward pass and its Jacobian [16], respectively.

**Predicting the observations** We compute the sensor likelihood  $p(\mathbf{o}_k|\tilde{\mathbf{x}}_k)$  using a transposed convolutional neural network (Fig. 3), which augments the dimensionality of the input from  $\mathbf{x}$  to the laser-sensor input size (e.g.,  $2 \rightarrow 622$ ). The input is firstly fed into two fully connected layers, and each transposed convolution layer is followed by a regular convolution layer, based on the work of [16]. At every layer, we used the ReLU activation function.

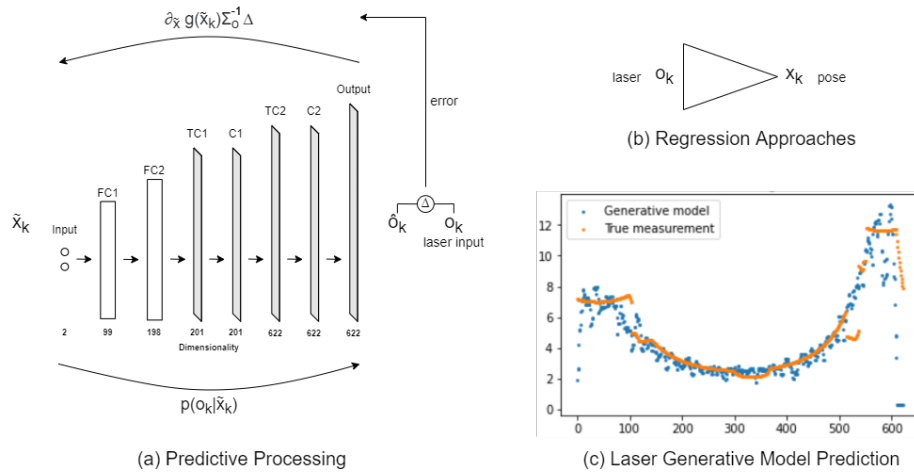


Fig. 3: (a) The Predictive Processing algorithm’s architecture. (b) Graphical representation of regression approaches to robot state estimation. (c) Generative model prediction vs. true laser measurement on a test sample.

The network was trained on 13000 normalized random samples collected in the Gazebo simulation. Each sample consists of the true position of the robot and the laser-sensor measurements at that location. The training was performed with 20 batches of 500 samples, using the L1 loss and Adam optimizer [10].

**Navigation.** Analogously, our algorithm infers the action in the same way that state estimation is computed, namely performing *active inference*. Actions also minimize the VFE in the predicted observations.

$$\mathbf{a} = \underset{\mathbf{a}}{\operatorname{argmin}} F(\tilde{\mathbf{x}}, \mathbf{o}(\mathbf{a})) \quad (2)$$

We define the goal as the preference or the intention of the agent  $\tilde{\mathbf{x}}_{goal}$  to arrive to a sensory state  $o_{goal}$  [15]. Estimation and control are computed as follows<sup>1</sup>:

$$\tilde{\mathbf{x}}_{k+1} = \tilde{\mathbf{x}}_k + \alpha [\partial_{\tilde{\mathbf{x}}} g(\tilde{\mathbf{x}}_k) \Sigma_{\mathbf{o}}^{-1} (\mathbf{o}_k - g(\tilde{\mathbf{x}}_k)) + \partial_{\tilde{\mathbf{x}}} g(\tilde{\mathbf{x}}_k) \Sigma_{\tilde{\mathbf{x}}}^{-1} \beta (\mathbf{o}_k - g(\tilde{\mathbf{x}}_{goal}))] \quad (3)$$

$$\mathbf{a}_{k+1} = \mathbf{a}_k + \gamma \partial_{\mathbf{a}} \tilde{\mathbf{x}} \partial_{\tilde{\mathbf{x}}} g(\tilde{\mathbf{x}}_k) \Sigma_{\mathbf{o}}^{-1} (\mathbf{o}_k - g(\tilde{\mathbf{x}}_k)) \quad (4)$$

where  $\beta$  weights the goal attractor and  $\gamma$  is the action step size. Note that each term computes the weighted prediction error mapped to the latent space.

The estimated state  $\tilde{\mathbf{x}}$ , now biased by the desired goal, generates a new predicted observation at every new iteration that is transformed into an action  $\mathbf{a}$ , which minimizes the VFE. Thus, performing a movement in the direction of the

<sup>1</sup> This update equation assumes that the Hessian of the goal dynamics is  $-1$  as proposed in [16].

goal. The pseudo-code described in Alg. 1 illustrates the process. The algorithm converges when the observation fits the predicted laser sensor measurements.

---

**Algorithm 1:** FEP localization and navigation algorithm

---

```

 $\tilde{\mathbf{x}} \leftarrow$  initial belief;
 $\mathbf{o}_{goal} \leftarrow g(\tilde{\mathbf{x}}_{goal});$  // Generate goal
while true do
     $\mathbf{o}_k \leftarrow$  Normalize(laser input);
     $\hat{\mathbf{o}}_k \leftarrow g(\tilde{\mathbf{x}});$  // Predicted observation
     $\tilde{\mathbf{x}} \leftarrow$  Eq. 3;
     $\mathbf{a} \leftarrow$  Eq. 4;
    PerformAction( $a$ );
end

```

---

### 3 Conclusions

This work shows a proof-of-concept on how predictive processing, i.e. active inference agents, can perform laser-based localization and navigation tasks. The results obtained in the localization experiment, where we compared against a state-of-the-art alternative (particle filter), show the potential of predictive stochastic neural filtering in robot localization, and estimation in general [7, 14]. Furthermore, the navigation experiment showcased how to compute actions as a dual filtering process. Nevertheless, the environment used in our experiments is rather simplistic compared to demonstrations of current sota algorithms. Therefore, we foresee further development and experimentation in terms of integration of odometry information, the introduction of extra degrees of freedom and connection to the robot’s non-linear dynamics.

### References

1. Ros documentation (Jun 2020), <http://wiki.ros.org/>
2. Tiago base (Oct 2020), <http://wiki.ros.org/Robots/TIAGo-base>
3. Besada-Portas, E., Lopez-Orozco, J.A., Lanillos, P., De la Cruz, J.M.: Localization of non-linearly modeled autonomous mobile robots using out-of-sequence measurements. *Sensors* **12**(3), 2487–2518 (2012)
4. Clark, A.: Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences* **36**(3), 181–204 (2013)
5. Fox, V., Hightower, J., Liao, L., Schulz, D., Borriello, G.: Bayesian filtering for location estimation. *IEEE Pervasive Computing* **2**(3), 24–33 (2003). <https://doi.org/10.1109/MPRV.2003.1228524>
6. Friston, K.: The free-energy principle: a unified brain theory? *Nature reviews neuroscience* **11**(2), 127–138 (2010)

7. Friston, K.J., Trujillo-Barreto, N., Daunizeau, J.: Dem: a variational treatment of dynamic systems. *Neuroimage* **41**(3), 849–885 (2008)
8. Gebauer, C., Bennewitz, M.: The pitfall of more powerful autoencoders in lidar-based navigation. arXiv preprint arXiv:2102.02127 (2021)
9. Kendall, A., Grimes, M., Cipolla, R.: Posenet: A convolutional network for real-time 6-dof camera relocalization. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2938–2946 (2015)
10. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2017)
11. Lanillos, P.: Minimum time search of moving targets in uncertain environments. Ph.D. thesis, PhD thesis (2013)
12. Lanillos, P., Cheng, G.: Adaptive robot body learning and estimation through predictive coding. In: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 4083–4090. IEEE (2018)
13. Liu, B., Cheng, S., Shi, Y.: Particle filter optimization: A brief introduction pp. 95–104 (2016)
14. Millidge, B., Tschantz, A., Seth, A., Buckley, C.: Neural kalman filtering. arXiv preprint arXiv:2102.10021 (2021)
15. Oliver, G., Lanillos, P., Cheng, G.: An empirical study of active inference on a humanoid robot. *IEEE Transactions on Cognitive and Developmental Systems* (2021)
16. Sancaktar, C., van Gerven, M.A.J., Lanillos, P.: End-to-end pixel-based deep active inference for body perception and action. *2020 Joint IEEE 10th International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)* (Oct 2020). <https://doi.org/10.1109/icdl-epirob48136.2020.9278105>, <http://dx.doi.org/10.1109/ICDL-EpiRob48136.2020.9278105>
17. Sattler, T., Zhou, Q., Pollefeys, M., Leal-Taixe, L.: Understanding the limitations of cnn-based absolute camera pose regression. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 3302–3312 (2019)
18. Thrun, S.: Simultaneous localization and mapping. In: *Robotics and cognitive approaches to spatial mapping*, pp. 13–41. Springer (2007)

## Appendix

### Moving state estimation

As discussed in the Localization and Estimation section, we have run two sets of experiments to evaluate the localization algorithm's performance. In the first setting the robot was motionless, whereas in the subsequent experiment it traversed the map from one side to the other. The results of one of the trials on this second experiment are seen in Fig. 4.

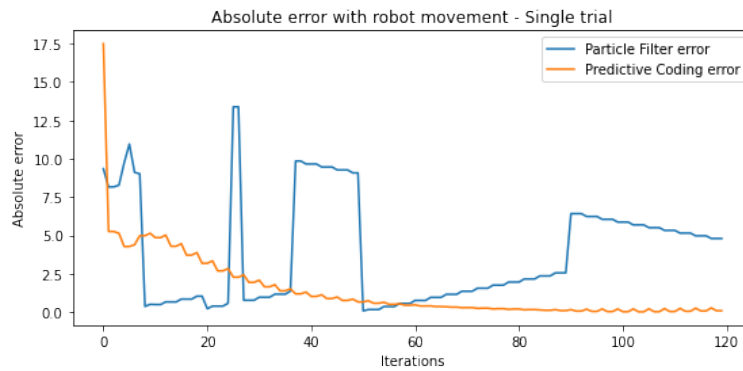


Fig. 4: Comparison to Particle Filter in moving setting

Performing the small teleports to simulate robot movement not only kept constant the robot's rotation angle but also isolated any possible incoming odometry information - potentially favouring performance of the PF. Movements were performed at every 3 iterations to let the algorithms correct the state estimations after the movements. This can be observed in the tooth saw artifacts in the figure.