From Neurons to Robots: Towards Efficient Biologically Inspired Filtering and SLAM

Niko Sünderhauf and Peter Protzel

Department of Electrical Engineering and Information Technology
Chemnitz University of Technology, 09126 Chemnitz, Germany
{niko.suenderhauf,peter.protzel}@etit.tu-chemnitz.de

Abstract. We discuss recently published models of neural information processing under uncertainty and a SLAM system that was inspired by the neural structures underlying mammalian spatial navigation. We summarize the derivation of a novel filter scheme that captures the important ideas of the biologically inspired SLAM approach, but implements them on a higher level of abstraction. This leads to a new and more efficient approach to biologically inspired filtering which we successfully applied to real world urban SLAM challenge of 66 km length.

1 Introduction

Environmental perception, information processing under uncertainty and navigation in unknown environments are two of the core problems of ongoing research in mobile autonomous systems. We want to direct the readers attention towards a class of relatively new, biologically motivated approaches to filtering, navigation and SLAM and the change of paradigms they express. After reviewing some of this interesting work, we are going summarize our recently published derivation of a novel filter scheme that was inspired by these new approaches and show that the new filter can reproduce the results of the biologically more accurate, but computationally more expensive approaches.

2 Related Work

2.1 A Neural Implementation of the Kalman Filter

Very recently, Wilson and Finkel [17] presented how a Kalman filter can be approximated by a neural attractor network structure. The time-dependent activations $u_{i(t)}$ of the neurons in this network are governed by the internal network dynamics and by the external input. This input $I_{(t)}$ corresponds to the noisy sensor data and is incorporated into the network in an additive way: $u_{(t)} = D(u_{(t-1)}) + I_{(t)}$. Wilson and Finkel were able to show that the function D expressing the network dynamics can be adjusted so that the resulting model equations map directly onto the Kalman filter equations. However, the authors pointed out that this is only the case when the prediction errors are small, i.e. prior and evidence distribution coincide. In case of large prediction errors, "the output of the network diverges from that of the Kalman filter, but in a way that is both interesting and useful." Wilson and Finkel concluded that this behaviour makes the neural-based estimator more robust to changepoints and outliers in the sensor data.

R. Dillmann et al. (Eds.): KI 2010, LNAI 6359, pp. 341–348, 2010.

2.2 Probabilistic Population Codes

Another group of authors developed a model of how uncertainty is manged and processed in the human brain. Their model is called *Probabilistic Population Codes* [6] [2]. Based on experimental results that humans perform "near-optimal Bayesian inference", the authors concluded that neural structures in the human brain are able to represent probability distributions and combine them using a "close approximation of Bayes' law"[6]. According to the proposed model, if there are two population of neurons that respond to the same stimuli and represent a certain piece of information (e.g. the position of a target according to auditory and visual cues), both representations can be fused by *adding* the weighted activities of the two individual neuron populations, i.e. performing a linear combination of the population activity vectors. Under the assumption that the neurons respond to the stimulus in a Poisson-like fashion, the linear combination of the two individual population activities can be shown to be Bayes-optimal [6].

Although not directly stated by [6] and [2], this optimality only holds when the hypotheses represented in the two populations approximately coincide. In later work [8] and [3], the same authors addressed this problem and formulated the "causal inference model" that distinguishes between cases of small and large disagreement between the two populations. Depending on the amount of conflict, the model can either force a fusion of the data or keep the two independent hypotheses.

2.3 RatSLAM – A Biologically Inspired SLAM System

RatSLAM [9] is a biologically motivated approach to appearance-based SLAM (*Simultaneous Localization and Mapping*, see [4] and [1] for an introduction and overview) that was successfully applied to a very large (66 km) and demanding urban scenario [10]. This approach differs substantially from the established probabilistic SLAM algorithms [16] in that it makes no use of Bayesian calculus or probabilities. Instead, it borrows important key ideas from biological systems and mimics the behaviour of different types of brain cells that have been found to be involved in spatial navigation tasks in rodents, primates and humans. We shortly reviewed these cells and provided references to the relevant literature in [15].

Fig.1(a) sketches the RatSLAM system. The sole input is provided by the vision system. It provides coarse *self motion cues* by means of visual odometry and is furthermore able to recognize known places the robot already visited. Both the calculation of visual odometry and the place recognition are performed using very simplistic algorithms.

The core of the system is a 3-dimensional *continuous attractor network*, called the *pose cell network* (PCN) that was inspired by the head direction and place cells in rodent brains. Due to the attractor dynamics, mutual excitation and inhibition, self-preserving packets of local activity form in the network of pose cells. These local packets compete, trying to annihilate one another until a stable state is reached. The pose cell network is used to maintain an estimate of the system's current pose in (x, y, θ)-space. Fig. 1(b) illustrates the network and shows how the activity can wrap around the network borders.

Each cell in the PCN can receive additional stimuli from the *local view cells* which inject energy into the pose cell network. These local view cells are driven by the vision system's place recognition. Thus like their biological counterparts, the pose cells

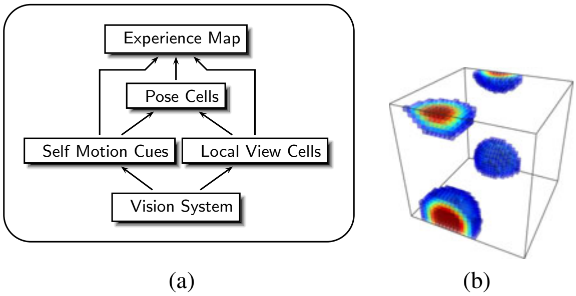

Fig. 1. (a) The general structure of RatSLAM. The vision system provides the input for odometry information and place recognition. The pose cells are inspired by the rodent head direction and place cells, while the experience map is responsible for loop closing and maintaining a topologically consistent map. **(b)** A screenshot of our C++ implementation of the pose cell network visualized using OpenGL. The network contains a single packet of activity. Notice how the activity wraps around the network borders.

are driven by both self-motion and visual cues. The main task of the PCN is to filter the many spurious false activations that are caused by the erroneous vision based place recognition system that generates a lot of false positive matches between different places. Without proper filtering, many wrong loop closures would be created.

Finally, on the top of the system we find the *experience map* which is responsible for managing a topologically and (to some extend) metrically consistent global map of the environment. It is a graph structure and consists of single *experiences*, each bound to a particular position in the state space and connected to previous and successive other experiences.

3 Deriving a Novel Filter Scheme from RatSLAM

The pose cell network (PCN) in RatSLAM is a three-dimensional attractor network, in general comparable to the one-dimensional network used by Wilson and Finkel in their work [17] mentioned above. While Wilson and Finkel designed their network in a way that it approximates a Kalman filter in case of small prediction errors, the authors of RatSLAM made no such attempt. Despite that, we discussed how the functionality of the pose cell network (PCN) can be compared to and interpreted from a Bayesian viewpoint in [13] and [14]. We pointed out the strong resemblance of the PCN to the histogram filter, a discrete version of the general Bayes filter.

Furthermore, we concluded that both approaches, PCN and histogram filters share the prediction step that predicts $\overline{\mathrm{bel}}(x_t)$ given the prior $\mathrm{bel}(x_{t-1})$ and the state transition model and control inputs. The main difference between both approaches is an *additive update* step performed by the PCN:

$$\mathrm{bel}(x_t) = \eta \left(\alpha p(z_t|x_t) + \overline{\mathrm{bel}}(x_t) \right) \tag{1}$$

where the posterior $\mathrm{bel}(x_t)$ is the weighted sum of the prediction $\overline{\mathrm{bel}}(x_t)$ and the evidence distribution $p(z_t|x_t)$ which is expressed by the local view cell's additive injection of energy into the pose cell network. The equation above contrasts the *multiplicative update* that is performed by Bayes filters, where $\mathrm{bel}(x_t) = \eta \cdot p(z_t|x_t) \cdot \overline{\mathrm{bel}}(x_t)$.

It is apparent that the *additive* incorporation of external data has as well been described in the work concerning probabilistic population codes and the work of Wilson and Finkel we reviewed above. Models of biologically motivated information processing seem to arrive naturally at this additive solution as neurons are generally understood to perform linear combinations, i.e. weighted *summation* of their inputs. Similar to [17], we found this additive solution to perform more optimal in cases of large prediction errors, i.e. when the prediction and evidence distributions largely disagree. Fig. 2(a) compares the additive and the Bayesian multiplicative solutions under a large prediction errror that is not covered by the involved covariances. This situation might occur in the event of a loop closure after the odometry accumulated a lot of under-estimated errors due to a wrong estimation or a degration of the measurement noise (over-confidence). Other possible causes are outliers or an erroneous data association.

Algorithm 3.1. Causal Update Filter (evidences, priors)

```
for all p_i in priors do
    for all e_j in evidences do
        if CONSENT(p_i, e_j) then
            new_posterior ← BAYESIAN_UPDATE(p_i, e_j)
            INCORPORATE(new_posterior, all_posteriors)
            p_i.used ← True ; e_j.used ← True
        end if
    end for
    if p_i.used = False then
        INCORPORATE(p_i, all_posteriors)
    end if
end for

for all e_j in evidences do
    if e_j.used = False then
        INCORPORATE(e_j, all_posteriors)
    end if
end for

RESCALE_AND_PRUNE( all_posteriors )
best_posterior ← FIND_POSTERIOR_PEAK(all_posteriors)
return best_posterior, all_posteriors
```

Seeking a more efficient implementation of the additive incorporation scheme without using neural attractor networks, we formulated the novel *Causal Update filter* (CUF) in [13] and [14]. Its core idea is to use the multiplicative update where the two distributions (prior and evidence) agree, and the additive update when they disagree. The filter's response therefore is a mixture of the additive and the Bayesian multiplicative update and resembles the causal inference model of [7]: The filter can either fuse two hypotheses or keep both, depending on how much they coincide. Fig. 2(b) illustrates this idea. To achieve this behaviour, and for performance reasons, the prior, evidence and posterior distributions are modelled as multimodal Gaussian distributions. The algorithm (see Algorithm 3.1 for pseudo-code) determines the pairwise consent between the involved Gaussians using a Mahalanobis-like measure: $d_{(N_1,N_2)} = \sqrt{(\mu_1 - \mu_2)^T(\Sigma_1 + \Sigma_2)^{-1}(\mu_1 - \mu_2)}$. A threshold on this consent measure decides whether the Gaussians are incorporated into the resulting posterior in a multiplicative or additive way. Each Gaussian is assigned a weight, that is not changed when the Gaussian is incorporated additively. When two Gaussians are fused because they are in consent, the resulting Gaussian is assigned the sum of the weights of the two parent Gaussians. The weights are used in a pruning and rescale step at the end of the algorithm where posterior Gaussians whose weights are too small are removed from the joint posterior. This way, we achieve a kind of voting scheme where posterior Gaussians that represent consent prior and evidence hypotheses are considered more important.

Using this novel Causal Update filter we were able to completely replace the PCN with a much more efficient approach, while maintaining its desired robustness.

4 Results

In order to prove the CUF to be functionally equal to the pose cell network it was derived from, we adapted our RatSLAM implementation and replaced the PCN by the CUF. Fig. 3 illustrates the general algorithmic layout, details can be found in [14]. We found the SLAM system with the CUF at its heart to be able to perform equally well on the 66 km dataset presented by Milford et al. in [10]. Fig. 2(c) shows the results of the experiment. Compared to the map given in [10] one can identify some flaws (e.g. in the

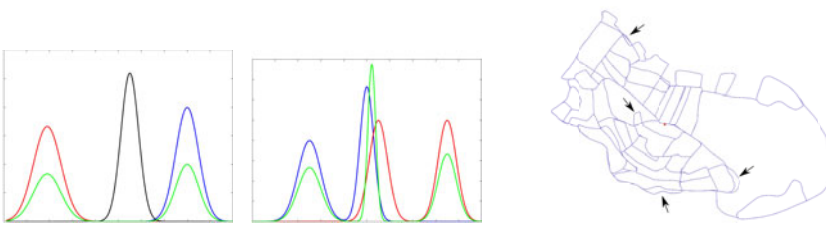

(a) (b) (c)

Fig. 2. **(a)** Bayesian multiplicative and additive update of dissenting Gaussian prior (blue) and evidence (red) distributions along with the respective Bayesian posterior (black) and additive posterior (green). Notice how the additive update bears more intuitive results by splitting the probability mass and keeping two independent hypotheses. **(b)** 1-dimensional example of the CUF. The prior (blue) and evidence (red) distribution are multimodal, each consisting of two peaks. The resulting posterior (green) consists of three Gaussians. Notice how the two hypotheses in consent have been incorporated multiplicatively and the remaining ones have been incorporated additively. **(c)** Final map of St. Lucia (path length 66 km, dimensions 1.8 by 3 km). The Causal Update filter was able to replace the pose cell network and managed to filter the erroneous data from the simplistic scene descriptors. The map has a few flaws (marked by the arrows) where the filter missed a loop closure, but covers the topologic layout and the general metric proportions of the environment.

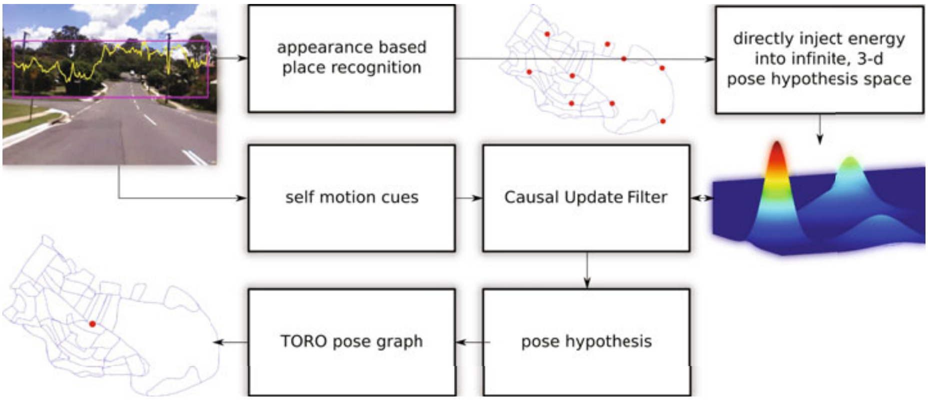

Fig. 3. The general structure of our modified, biologically inspired SLAM algorithm, using the Causal Update filter. The vision system provides the input for odometry information and place recognition. Notice that the place recognition is highly erroneous and produces a lot of false positives. The Causal Update filter at the center is derived from RatSLAM's pose cell network, but implements its core ideas on a higher level of abstraction. The TORO [5] pose graph is responsible for maintaining a topologically consistent map.

very top left or at the bottom) but in general the road network was successfully captured and the resulting map is topologically correct. Although it is of course not metrically correct, it can be considered semi-metric as the general proportions of the environment are represented adequately. A video showing the working SLAM system is available at our website `www.tu-chemnitz.de/etit/proaut`.

Compared to the original pose cell network, the CUF can be calculated much more efficiently: While incorporating odometry information into the PCN and calculating the network dynamics took 60 ms, the CUF can be calculated in well under 0.1 ms as it involves only a few operations. This massive speed-up comes at no costs or disadvantages, as the CUF is able functionally replace the pose cell network and to reproduce its results.

Furthermore, we reproduced the simple tracking example presented in [17]. In their experiment the authors tracked an object in a one-dimensional environment, using noisy position measurements as sensor input to the neural network structure that was tuned to approximate a Kalman filter. The results of this experiment using the CUF instead of the neural network are shown in Fig. 4. The first 75 timesteps resemble the small prediction error case in the experiment of Wilson and Finkel [17]. In this case, the CUF response is identical to a standard Bayes filter (a Kalman filter in this case, as all involved processes are linear). At time 75 the tracked object is kidnapped and the CUF's response diverges from that of a Kalman filter. We can see how a second, growing hypothesis is introduced at the new position and how the old hypothesis is still maintained, but constantly weakens due to the lack of sensory backup until it finally vanishes around time 120. This behaviour is exactly what was observed by [17] in the event of a so called changepoint. The CUF is able to reproduce these desirable results but can be calculated more efficiently, as no network dynamics are involved.

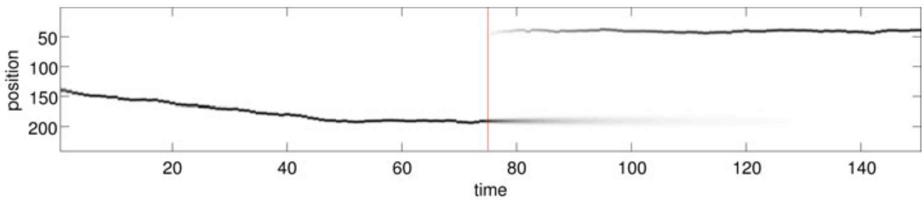

Fig. 4. The CUF in a 1-dimensional tracking experiment similar to those performed by [17]. Gray values indicate the weights of the CUF posteriors following the ground truth. The red line marks the kidnapping of the tracked object. See the text for further explanation.

5 Conclusions and Future Work

As we saw in the review of related work, the additive incorporation of information is common in biologically motivated models of filtering and information processing. Although this additive principle is inherent in neural systems, as neurons perform linear combinations, i.e. weighted summations of their inputs, it contrasts the multiplicative update scheme of Bayesian filter approaches.

Inspired by this fundamental difference between the world of Bayesian calculus and the biologically motivated approaches, we derived the Causal Update filter from the pose cell network of RatSLAM. While the pose cell network captured the additive incorporation principle by sticking closely to the neural nature of its biological archetypes (head direction and pose cells in the mammalian brain), our CUF is a higher abstraction of this principle, leading to an increased efficiency while maintaining the desirable robustness. Despite these abstractions, the CUF performed equally well in the same demanding, 66 km long urban SLAM scenario as the original RatSLAM algorithm [10]. Furthermore, we showed that the CUF is as well able to reproduce the results of a neural network model that was designed to approximate a Kalman filter [17].

In future work we will have to explore the possibilities of the novel CUF filter scheme and apply it to different filter problems where Bayesian techniques are the state of the art today. Seeing if the CUF can be applied to other problems and how well it performs there compared to the established Bayesian filters, remains an open but exciting question.

Equally important will be to establish an analysis and further understanding of the exact nature of the proposed filter. As the additive update can not be directly derived from the laws of Bayesian calculus, the filter seems to be non-Bayesian and closer to alternative formulations of probability that distinguish between *uncertainty* and *ignorance* like for instance Dempster-Shafer [11], the possibility theory of Zadeh [18] or the transferable belief model [12].

Acknowledgements

We would like to thank Michael Milford for providing the St. Lucia video footage that was presented in his paper [10]. Further material on RatSLAM is available at the RatSLAM website (`http://ratslam.itee.uq.edu.au`).

References

1. Bailey, T., Durrant-Whyte, H.: Simultaneous Localisation and Mapping (SLAM): Part II State of the Art. IEEE Robotics and Automation Magazine 13(3), 108–117 (2006)
2. Beck, J.M., Ma, W.J., Latham, P.E., Pouget, A.: Probabilistic population codes and the exponential family of distributions. Progress in Brain Research 165, 509–519 (2007)
3. Beierholm, U., Körding, K.P., Shams, L., Ma, W.J.: Comparing bayesian models for multisensory cue combination without mandatory integration. In: NIPS (2007)
4. Durrant-Whyte, H., Bailey, T.: Simultaneous Localisation and Mapping (SLAM): Part I The Essential Algorithms. IEEE Robotics and Automation Magazine 13(2), 99–110 (2006)
5. Grisetti, G., Stachniss, C., Burgard, W.: Non-linear constraint network optimization for efficient map learning. IEEE Transactions on Intelligent Transportation Systems 10(3), 428–439 (2009)
6. Ma, W.J., Beck, J.M., Latham, P.E., Pouget, A.: Bayesian inference with probabilistic population codes. Nature Neuroscience 9, 1432–1438 (2006)
7. Ma, W.J., Beck, J.M., Pouget, A.: Spiking networks for bayesian inference and choice. Current Opinion in Neurobiology 18(2), 217–222 (2008), Cognitive Neuroscience
8. Ma, W.J., Pouget, A.: Linking neurons to behavior in multisensory perception: A computational review. Brain Research 1242, 4–12 (2008)
9. Milford, M.J.: Robot Navigation from Nature. Springer, Heidelberg (March 2008)
10. Milford, M.J., Wyeth, G.F.: Mapping a Suburb with a Single Camera using a Biologically Inspired SLAM System. IEEE Transactions on Robotics 24(5) (October 2008)
11. Shafer, G.: A Mathematical Theory of Evidence. Princeton University Press, Princeton (1976)
12. Smets, P.: The combination of evidence in the transferable belief model. IEEE Pattern Analysis and Machine Intelligence 12, 447–458 (1990)
13. Sünderhauf, N., Neubert, P., Protzel, P.: The Causal Update Filter – A Novel Biologically Inspired Filter Paradigm for Appearance Based SLAM. In: Proc. of the IEEE International Conference on Intelligent Robots and Systems, IROS (2010)
14. Sünderhauf, N., Protzel, P.: Beyond RatSLAM: Improvements to a Biologically Inspired SLAM System. In: Proceedings of the IEEE International Conference on Emerging Technologies and Factory Automation, Bilbao (September 2010)
15. Sünderhauf, N., Protzel, P.: Learning from Nature: Biologically Inspired Robot Navigation and SLAM - A Review. In: Künstliche Intelligenz, German Journal on Artificial Intelligence, Special Issue on SLAM. Springer, Heidelberg (2010)
16. Thrun, Burgard, Fox: Probabilistic Robotics. The MIT Press, Cambridge (2005)
17. Wilson, R., Finkel, L.: A Neural Implementation of the Kalman Filter. In: Bengio, Y., Schuurmans, D., Lafferty, J., Williams, C.K.I., Culotta, A. (eds.) Advances in Neural Information Processing Systems, vol. 22, pp. 2062–2070 (2009)
18. Zadeh, L.: Fuzzy sets as the basis for a theory of possibility. Fuzzy Sets and Systems 1, 3–28 (1978)