Integrating Multiple Sources of Knowledge for the Intelligent Detection of Anomalous Sensory Data in a Mobile Robot

Manuel Castellano-Quero, Juan-Antonio Fernández-Madrigal, and Alfonso J. García-Cerezo

Systems Engineering and Automation Department.
University of Malaga, Campus de Teatinos, 29071, Malaga (Spain)
{mcastellano,jafernandez,ajgarcia}@uma.es

Abstract. For service robots to expand in everyday scenarios they must be able to identify and manage abnormal situations intelligently. In this paper we work at a basic sensor level, by dealing with raw data produced by diverse devices subjected to some negative circumstances such as adverse environmental conditions or difficult to perceive objects. We have implemented a probabilistic Bayesian inference process for deducing whether the sensors are working nominally or not, which abnormal situation occurs, and even to correct their data. Our inference system works by integrating in a rigorous and homogeneous mathematical framework multiple sources and modalities of knowledge: human expert, external information systems, application-specific and temporal. The results on a real service robot navigating in a structured mixed indoor-outdoor environment demonstrate good detection capabilities and set a promising basis for improving robustness and safety in many common service tasks.

Keywords: Robot sensors, Bayesian inference, Sensory data diagnosis.

1 Introduction

Service robotics is in current expansion [1]. Blossomed almost two decades ago mainly due to the methodological advances in probabilistic management of information [2], now standard robots are capable of dealing with relevant amounts of uncertainty in the real world intrinsically and efficiently, particularly at the lowest control levels. However, dealing with uncertainty (noise) is not enough: they should also be able to identify and deal intelligently with —make decisions about— abnormal situations in order to improve their expected performance; other than computational efficiency issues, there are no conceptual reasons that prevent them to achieve that even at the most basic levels of operation.

In this paper we focus on the particular problem of using sensory data as safely and robustly as possible beyond uncertainty, i.e., when those data are heavily modified by unexpected circumstances: adverse environmental conditions, the special nature of some perceived elements in the world, or even breakdowns in the sensor hardware. For obtaining such capability in today robots, they have to resort to multiple sources of knowledge besides the ones intrinsic to their own design and operation, i.e., regarding the mentioned circumstances, the normal behaviour of their sensor devices, and their dynamics. These sources and the modalities of knowledge they provide are really diverse: they can be human experts, external information systems, previous knowledge about the environment where the robots work, etc. In spite of that diversity, all of them should be integrated and used as rigorously and as homogeneously as possible to optimise tractability, robustness and safety.

Concerning our particular problem, there exist a number of tools that could be employed for detecting abnormal sensory data: neural networks [3], fuzzy inference [4], Bayesian inference [5], ad-hoc or heuristic approaches [6], etc. However, only Bayesian inference can provide a homogeneous and mathematically rigorous foundation (based on probability and statistics [7]) for fusing knowledge coming from a number of different sources; neural networks offer no explicit explanations about their deductions nor a rigorous basis for managing uncertainty, fuzzy logic is mostly suitable for expert knowledge, and other approaches are not based on a sound mathematical foundation, thus compromising their guarantees. Bayesian networks can reason not only in one direction (from data to conclusions), but can infer knowledge about any element in the network, i.e., we can deduce either which anomaly is present or what would be the actual information perceived by a sensor if some anomaly occurs, without re-building the network. In addition, Bayesian inference can be hybridised with other paradigms [8][9] and be integrated naturally with probabilistic methods used by modern robots. The main drawback of Bayesian inference with respect to other approaches is its high computational cost. In this paper we do not cope with that problem (our current and future work is being focused and planned to advance in that direction), but some approximation inference algorithms exist that limit the amount of computation during inference [10][11], thus we claim that the Bayesian approach is a highly promising one for our goals.

Bayesian inference, although widely used in robotic estimation (for localization and mapping [2]), has been scarcely and only sporadically used for sensory diagnosis (e.g., [12]). A general Bayesian inference system has been reported in [13], but that is a preliminary work focused on the interactive construction of the network rather than in its capabilities, and has been used only in particular, static robot situations, being, in consequence, only limitedly tested. In this paper we take a step further by augmenting the sources of knowledge that can be integrated and exploited by the network, particularly through the addition of sequential filters to its output data, concretely a kind of infinite impulse response [14] for probability estimation and a robust moving window median filter for value estimation [15] —a much more computationally efficient approach than using a dynamic Bayesian network, which would increase the cost of inference exponentially. This allows us to include temporal information, a source of knowledge of great relevance in a deduction process that has necessarily to deal with the dynamics of the robot-environment interaction. We have also tested our system along a complete navigation route of the mobile robot in a mixed indoor-outdoor scenario where the robot encounters some sensory issues, not only in particular, static situations, thus providing a richer study of the anomaly detection and sensory data recovering during the robot operation in a real application.

Our current implementation fuses geographical knowledge (the location of the robot), meteorological knowledge (about the weather in the season at that location), expert knowledge (of a human, concerning the possible anomalies that could occur in the sensory devices), environmental knowledge (some characteristics of the scenario and its elements), and temporal knowledge (about how the last sensory and estimation data affects the present). In principle, the number and diversity of knowledge sources are only constrained by the complexity of the resulting Bayesian network.

In our experiments the robot has used both very basic sensors (wheel encoders, bumpers, etc.) and more complex ones (laser rangefinders and RGB-D cameras). It has dealt with abnormal sensory data produced by the presence of dark-surface objects (difficult to perceive by range sensors that emit infrared radiation), thin objects (easy to be missed between consecutive beams), and adverse environmental conditions (excessive light when the robot goes outside, that affect visual devices). Thanks to the sequential filter, the dynamics of the detection process has been coped with adequately in all cases. All in all, our results show that with all this information, an operational robot can infer successfully whether its sensors are working nominally or some anomaly is likely to be present, and even correct the data coming from a doubtful sensor device with the ones from another sensor or with commonsense information.

The rest of the paper is as follows. Section 2 gives an overview of our inference system based on Bayesian networks. Section 3 details its most relevant elements for a particular case of mobile robot. Section 4 explains the inference capabilities of the system in several adverse circumstances encountered by the robot during its operation. Finally, section 5 summarizes the main conclusions of this work and sets some future lines of research.

2 Overview of the Bayesian Inference System

A Bayesian network defined on a set of variables $\mathbf{V}$ is a pair (G,Θ) consisting of a direct acyclic graph, G, over $\mathbf{V}$, called the *network structure*, and a set of Conditional Probability Tables (CPTs), Θ, for each variable in $\mathbf{V}$, called the *network parametrization*. The graph structure captures the causal relationships between variables through directed arcs, which indicate dependence relationships, while CPTs define probability distributions over the variables. For more in-depth treatment of Bayesian networks and inference, please refer to [7][16].

We are interested in inferring new knowledge from existing one, i.e., in deducing a probability distribution over a set of query nodes of the graph, $\mathbf{Q}$, given some evidence nodes $\mathbf{E}$, i.e., $P(\mathbf{Q}|\mathbf{E})$. This can be done by applying basic probability theory (the chain rule, the Bayes theorem, etc.) repeatedly, although that can be prohibitive even for small scale problems, so a Bayesian inference algorithm should aim to a reduced computational complexity.

There exist many inference methods for Bayesian networks, both exact, such as the well known junction tree algorithm [16], and approximate, such as loopy belief propagation and likelihood weighting [7]. In general, the former provide correct answers using more computational resources (junction tree is $O(n \cdot \exp(w))$ for a graph with n nodes and a treewidth of w, for instance), while the latter can be more efficient at the expense of less accurate answers —their quality depends on the problem and on the allowed number of iterations, since they are often any-time algorithms. In this paper we have used the junction tree algorithm since it has provided the best trade-off between quality of the results and computational cost for the size of our networks.

Our architecture for doing inference in sensory systems uses a basic component, a so-called *Bayesian sensor*, modelled through a Bayesian network, which represents not only a real sensor but also additional information that enables intelligent diagnosis and sensory enhancement, i.e., it is the component in charge of integrating the diverse sources of knowledge concerning the sensory data production. Figure 1 depicts the structure of this generic Bayesian sensor, formed by three different subnetworks (rectangles) and a multiplexer node (ellipse) that we explain below. Multiple components of this kind can be interconnected in our complete system; nodes of interest can also be enhanced by adding temporal filters to them (not shown in that figure). These subnetworks and filters will be dealt with in more detail, and deployed into their elements, in section 3.

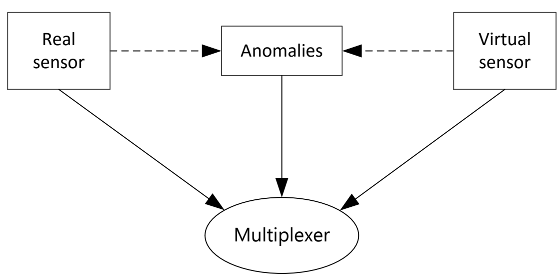

Fig. 1. Generic sensor based on a Bayesian network. Dashed lines are optional. Interconnections among these Bayesian sensors are done through their multiplexers.

The *real sensor* subnetwork represents an existing sensor on board the mobile robot. This subnetwork will contain one node whose values represent the measurements of the sensor.

The *virtual sensor* subnetwork receives information from other Bayesian sensors (directly or through some calculations) in order to emulate the behaviour of its real sensor when it is faulty and to deduce its data, i.e., for recovery.

The *anomalies* subnetwork indicates whether there are faults or abnormal situations in the associated sensor. This can be deduced, for example, using information from other sensors, or integrating external knowledge (wheather, location, expert, etc.).

Finally, the *multiplexer* node selects the inferred sensor measurement. If there is a high probability of abnormal behaviour, the virtual sensor will have more influence than the real one in this final result. In our current implementation, the multiplexer node is a discrete variable with the same values as the ones of the leaf nodes of the real and virtual sensor subnetworks.

Although the component depicted in Figure 1 is enough for integrating a diversity of knowledge coming from different sources, we also include sequential filters in some nodes in order to take into account temporal knowledge (dynamics). These filters work on the posterior distributions obtained during the robot operation, which may represent either anomalies or sensory data encoded by multiplexers.

3 Instantiation for a particular mobile robot

In this work we have used a *Turtlebot* robot [17] to implement the architecture defined above (see figure 5). It is a robot with a *Kobuki* mobile platform and a bunch of sensors: three bumpers, two magnetic encoders, three cliff detectors, a gyroscope, two wheel drop sensors, a *Kinect* RGB-D device and a *Hokuyo* 2D laser rangefinder. In our experiments we use the laser rangefinder, the RGB-D camera, the bumpers and the encoders.

For the sake of space, in this section we focus on one of the components of the whole network, the one corresponding to the laser rangefinder, since it contains enough complexity in the integration of several sources of knowledge and represents well the decisions made during the implementation of the other sensory devices in the entire system. Also notice that, from a software implementation point of view, some of the further described elements can be coded only once, since they affect several parts of the network the same way; we have considered these re-factoring issues appropriately in our software.

The real sensor subnetwork for the laser rangefinder in this robot contains as many nodes as elements we wish to represent from the vector of measurements (beams). Each node is a discrete random variable with a suitable discretization of the measured distances. The multiplexer node is replicated for each laser beam and it represents the final probability distributions we want to get. The structure is trivial from figure 1.

In the corresponding anomalies subnetwork we integrate environmental information (wheather, location) as well as data from other sensors to detect abnormal situations (see Figure 2). This is done in two complementary ways: through the connectivity (dependences) among nodes in the network shown in the figure and by filling their CPTs with suitable commonsense and expert knowledge regarding all that information. We omit here the content of the CPTs for their very large size; in short, the knowledge for filling them has been translated into imperative programming, and thus appropriate routines have been coded with it. For instance, bad lighting in indoors can produce wrong measurements, as well as the presence of rain outdoors and also extremely high temperatures. Environmental anomalies in this device are considered to affect all laser beams equally.

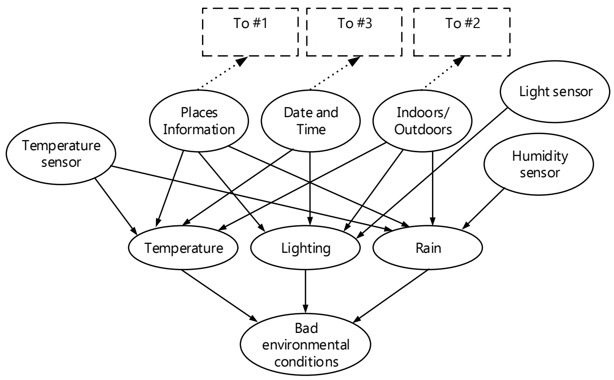

Fig. 2. Abnormal environmental conditions detection in the anomalies subnetwork of the laser rangefinder Bayesian network. The CPTs associated to edges in the graph code the integration of the different sources of knowledge.

There also exists systematic errors that can affect a laser rangefinder, such as the detection of too thin, transparent or black objects. For this part of its anomalies subnetwork we use an alternative distance sensor (the depth information provided by the *Kinect* camera), a RGB image and some nodes of the environmental part (see figure 3 (a)). The actual values of the laser rangefinder are compared with the ones from the alternative distance sensor; if there is a significant difference, then this anomaly is detected —we are here integrating expert knowledge about the sensors. We can also distinguish the absorbed radiation anomaly (black objects) by combining the information from the place where the robot is and the percentage of black pixels in the RGB image. The undetected object anomaly (too thin or transparent objects in the path of the laser beams) takes into account the sense of the difference between the laser and the alternative sensors, and also information known about the working place; for instance, if we are known to be in a place with transparent objects, e.g, windows, and one or more alternative sensors indicate shorter distances than the laser, the probability of anomalies for this reason will be high. There is also a particular node (reading error failure) that increases its probability when there is an important difference between measured distances while the mentioned anomalies are false.

All the described anomalies are summarized into a leaf node which finally indicates if a certain beam of the sensor has an abnormal state. This is depicted in figure 3 (b). This kind of leaf nodes are natural targets for integrating temporal knowledge, i.e., for considering dynamics. For that purpose, we have attached a sequential filter suitable for continuous data to them, concretely an exponentially

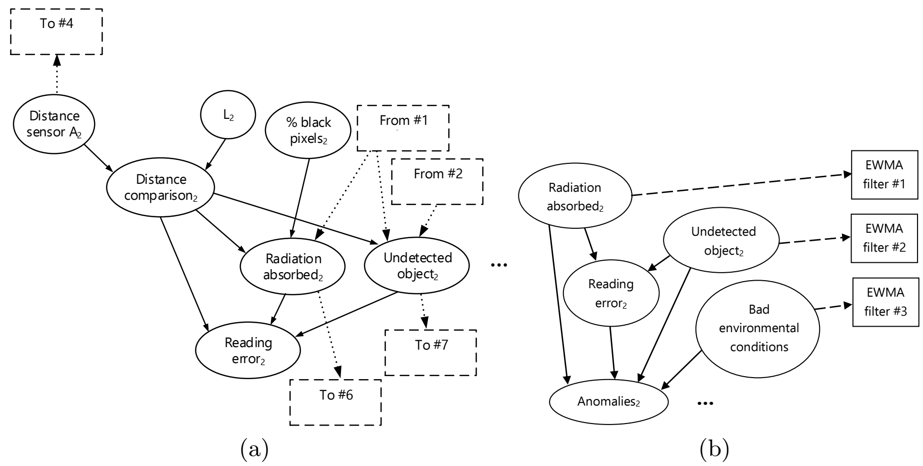

Fig. 3. (a) Systematic errors detection for the anomalies subnetwork corresponding to beam #2 of the laser rangefinder. (b) Anomalies integration for beam #2 of the laser rangefinder and their connection with the associated EWMA sequential filters.

weighted moving average filter (EWMA); more precisely, one to each of the three main anomalies considered in this sensor. Thus, the value of their probabilities are affected by present and past inferred distributions.

The second main element of the laser rangefinder sensor network is the virtual subnetwork (deployed in figure 4). It integrates the value of the bumpers and the linear speed of the platform under a specific laser scan angle, in addition to some anomalies, using expert knowledge. For instance, assuming reactive navigation, when the linear speed is low the probability of finding a nearby obstacle should be high; also, the bumper node gets information about collisions, which, in case of occurring, should set high the probability of having a short distance to obstacles.

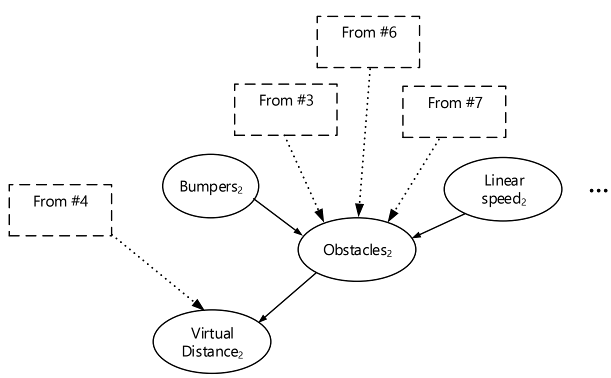

Fig. 4. Virtual subnetwork for the beam #2 of the laser rangefinder sensor.

The previously described anomaly subnetwork provides this one with knowledge about whether there is an undetected object (too thin, transparent or black). Furthermore, knowledge about the time of the day and date is important to estimate the amount of people or other kind of mobile obstacles present in the navigation scene. This is another example of the power of Bayesian networks to integrate diverse kinds of knowledge under a common, consistent formalism, namely graph connectivity and CPTs.

Finally, all the multiplexers in our entire Bayesian networks are also temporally filtered, but not with the EWMA filters attached to anomaly nodes, as before, since we are interested in inferred values of the nodes as outputs (discrete), and not in their probabilities (continuous). For that purpose we have used moving-window median filters, an also efficient and well-known solution in statistics for increasing the robustness of data against outliers [15].

4 Evaluation of the inference system

The Bayesian inference system described in the previous sections has been evaluated in a *Turtlebot* mobile robot in a real experiment described here. The robot has an on board netbook PC with an Intel Celeron N2840 at 2.16 GHz and 2 GB DDR3 that runs Ubuntu 14.04 with ROS [18]. Since we are not interested in dealing with computational limitations in this work, we have used another PC to remotely execute our Bayesian inference software, an Intel Core i3 3217U at 1.8 GHz and 6 GB DDR3 that runs Ubuntu 16.04 with ROS, where our software with the model described in section 3 has been implemented in MATLAB using the Bayes Net Toolbox (BNT) [19].

The Bayesian network has been endowed with 10 laser beams (10 multiplexer nodes in the network, corresponding to approximately 60 degrees of fov) with 10 possible distances for each one, ranging from 1 (no obstacle detected) to 10 (maximum distance). The resulting network has 148 nodes. In every experiment we have run a control loop in which sensory data are obtained, then the Bayesian architecture is evaluated with that evidence and, finally, velocity commands are sent to the robot if needed.

In the real experiment, the *Turtlebot* robot navigates along a route in the mixed indoor-outdoor scenario shown in figure 5. During the route, it has to deal with different sensory abnormalities (described in the previous section), that are represented by the posterior distributions associated with nodes of the anomalies subnetworks. These results are shown with their corresponding temporal filtering in figure 6, where we consider that there is an abnormal situation when its probability is reasonably high, e.g., equal or greater than 0.7.

Firstly, we analyse the results obtained for the undetected object anomaly (figure 6 (a). This issue arises whenever the sensors are faced to thin objects (chair or table legs, columns, cables, etc.) and also when they are not able to capture too distant objects. In this experiment, both situations take place. When the robot is halfway between points 1 and 2 (see Figure 5), it is pointing to the wall near point 4 —this happens during the experiment times of 50 and 60

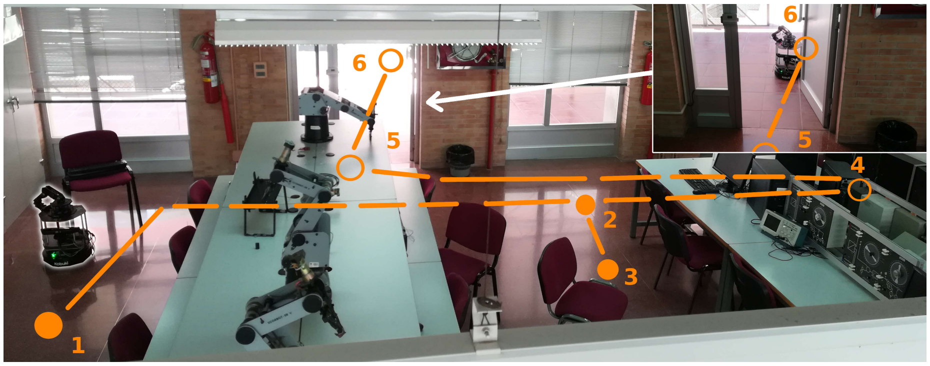

Fig. 5. Route followed during the experiment, with some points of interest. A video is available at http://150.214.109.139/robot2019.mp4. The robot is highlighted close to the point 1 (bottom-left); the outdoor portion of the route is zoomed in (top-right).

seconds, and the probability of anomaly shown in Figure 6 is nearly 0.9. That wall cannot be detected by the laser rangefinder since it is out of its range, but this is not the case for the depth camera, which is used to correct the final sensory value. After that, the robot moves between distinctive points 2 and 3 pointing towards the center chair (times between 70 and 90 seconds). The problem here is that the chair legs are too thin to be detected; thus the probability of undetected obstacle is again 0.9. Note that this deduction is possible, in particular, thanks to the integration of knowledge about the working environment of the robot —presence of chairs. When the robot is navigating from distinctive point 4 to 5 there are no nearby obstacles again, thus the problem of measurement range reappears.

Concerning the radiation absorbed anomaly (figure 6 (b)), that situation is provoked by the presence of dark surfaces, specially black ones, which absorb part of the infrared radiation emitted by the ranging sensors. In this particular experiment, the robot is in front of dark objects two times. One coincides with the first undetected object anomaly, as the robot is pointing to a few black-cases computer surfaces that lie around point 4. After navigating through the central corridor (points 2 and 3), the robot moves from point 2 to 4 (experiment times between 90 and 110 seconds) pointing again towards the dark objects. The inference about this problem is also possible due to the knowledge about the environment, as well as the fact that some measurements are lost.

A third abnormal situation appear in the outdoor part of the environment. Our robot goes outdoor between points 5 and 6, and there finds strong sunlight (experiment time between 160 and 220 seconds). Both ranging sensors are seriously affected due to the interferences produced by the infrared component of natural light, but that is correctly detected in the probability distribution shown in figure 6 (c). In this case the deduction is possible thanks to the integration of geographical and other external knowledge (location, season and weather).

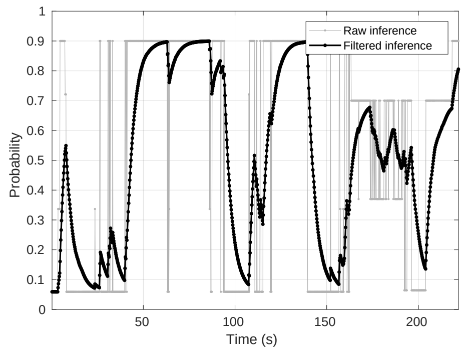

(a)

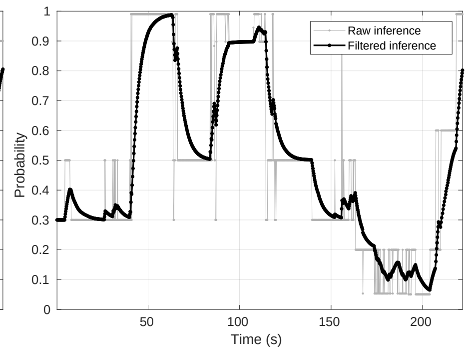

(b)

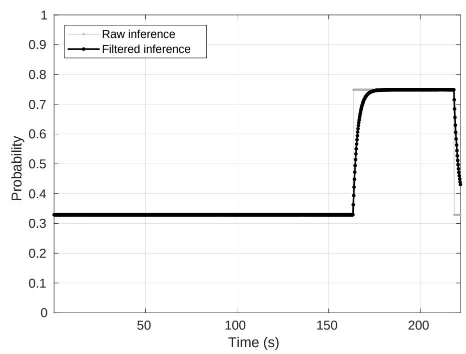

(c)

Fig. 6. Posterior distributions for anomalies in the laser rangefinder, both before temporal knowledge integration (gray) and after (black). (a) Undetected obstacle, beam #4. (b) Radiation absorbed, beam #4. (c) Environmental conditions (excess of light).

Our Bayesian inference system enables not only to detect abnormal situations as shown so far, but also to recover sensory data under these conditions. In order to illustrate this, the multiplexer nodes of the laser rangefinder have been used. These variables have probability distributions over the possible discretized distances to obstacles (from 1 to 10 in this case). We have selected the distance with the highest probability at each point along the experiment in order to analyse the recovery capabilities, considering the integration of temporal knowledge through the moving median filter, as explained before. We show the behaviour of a single beam, comparing the corresponding filtered multiplexer with the raw distances obtained by both the laser and depth sensors, in figure 7.

These results show that the proposed system is able to recover sensory data in quite adverse situations. As an example, consider the time interval between 120 and 140 seconds. As shown in figure 7 (b), the laser rangefinder detects no obstacle, while the filtered multiplexer indicates the presence of a distant object —it has been deduced that the depth camera operates under nominal conditions

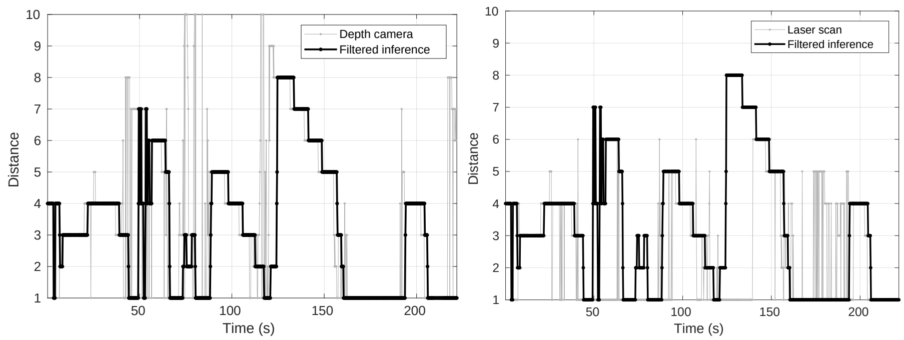

(a) (b)

Fig. 7. Most-likely distance to obstacles inferred by the multiplexer of beam #4 after temporal filtering (black). (a) and (b), in gray, raw evidence for that provided by the depth camera and the laser rangefinder, respectively. "1" means no detected obstacle; otherwise the number corresponds to an increasing distance range, from "2" to "10".

here. Furthermore, it is also inferred that the probability of undetected obstacle within this interval is high (figure 6 (a)). In this case, lost measurements have been corrected by the data coming from another sensor, integrating knowledge about the conditions of the experiment in this point and temporal information.

5 Conclusions and future work

In this work we have shown how Bayesian inference can be used to fuse multiple and heterogeneous sources of knowledge, external, expert and temporal, in a rigorous and consistent framework, so as to improve the robustness and power of robotic sensory systems. Our results show that the proposed inference system enables not only to infer faults in the sensors and their causes but also to recover sensory data even in those faulty situations.

In the future there are a number of issues to address. The computational cost of the inference method used here (junction tree) is not suitable for every robotic task, thus improvements are needed (e.g., parallelization, use approximate algorithms, or abstracting the network). Also, the Bayesian network should be created more autonomously and automatically. This should be done by a procedure that ensures the integration of human knowledge and at same time allows to discover the most likely structure of the network. Finally, we also plan to extend our inference system to different robotic platforms and applications.

Acknowledgements. This work has been supported by the Spanish government through the national grant FPU16/02243, by the University of Malaga through its local research program and the International Excellence Campus Andalucia Tech, and by the national research project DPI2015-65186-R.

References

1. International Federation of Robotics (IFR): Executive Summary World Robotics 2017. Service Robots, available at https://ifr.org/free-downloads/, consulted june 14th, 2019.
2. Thrun, S., Burgard, W., Fox, D.: Probabilistic Robotics. In: Intelligent Robotics and Autonomous Agents. The MIT Press (2005). ISBN: 0262201623.
3. Hu, M., Wang, Z., Yuan, Y., Qi, L.: On-Line Sensor Diagnosis of the Diesel Engine Cold Starting Based on RBFNN. In: IEEE Circuits and Systems International Conference on Testing and Diagnosis (2009).
4. Calderwood, S., McAreavey, K., Liu, W.: Context-dependent combination of sensor information in DempsterShafer theory for BDI. Knowledge and Information Systems. **51**(1), 259–285 (2017).
5. Mengshoel, O., Darwiche, A., Uckun, S.: Sensor Validation Using Bayesian Networks. In: 9th International Symposium on Artificial Intelligence, Robotics, and Automation in Space (2008).
6. Zhu, C., Wang, W. Q., Chen,H., So, H.C.: Impaired Sensor Diagnosis, Beamforming, and DOA Estimation With Difference Co-Array Processing. IEEE Sensors Journal. **15**(7), 3773–3780 (2015).
7. Darwiche, A.: Modeling and Reasoning with Bayesian Networks. Cambridge University Press (2009).
8. Osoba, O., Mitaim, S., Kosko, B.: Bayesian Inference With Adaptive Fuzzy Priors and Likelihoods. IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics). **41**(5), 1183–1197 (2011).
9. Giovanis, G., Papaioannou, I., Straub, D., Papadopoulos, V.: Bayesian updating with subset simulation using artificial neural networks. Computer Methods in Applied Mechanics and Engineering. **319**(1), 124–145 (2017).
10. Murphy, K., Weiss, Y., Jordan, M.: Loopy belief propagation for approximate inference: An empirical study. In: Proceedings of the Fifteenth Annual Conference on Uncertainty in Artificial Intelligence, UAI-99, San Francisco, CA, Morgan Kaufmann Publishers, 467–475 (1999).
11. Shachter, R.D., Peot, M.A: Simulation approaches to general probabilistic inference on belief networks. Uncertainty in Artificial Intelligence. **5**, 221–231 (1989).
12. Saha, B., Koshimoto, E., Quach, C.C., Hogge, E.F., Strom, T.H., Hill, B.L., Vazquez, S.L., Goebel, K.: Battery health management system for electric UAVs. In: 2011 Aerospace Conference, Big Sky, MT (2011).
13. Castellano-Quero, M., Fernández-Madrigal, J.A., García-Cerezo, A.: Interactive Construction of Bayesian Inference Networks for Robust Robot Sensorics. In: XIV Simposio CEA de Control Inteligente (2018).
14. Pham, H. (Ed.): Springer Handbook of Engineering Statistics. Springer Science and Business Media (2006).
15. Huber, P.J., Ronchetti, E.M.: Robust Statistics. Wiley (2009). ISBN: 978-0470129906.
16. Huang, C., Darwiche, A.: Inference in Belief Networks: A Procedural Guide. International Journal of Approximate Reasoning. **15**(3), 225–263 (1996).
17. Open Source Robotics Foundation, Turtlebot mobile robot official website, https://www.turtlebot.com, visited June 16, 2019.
18. Quigley, M., Conley, K., et al.: ROS: an open-source Robot Operating System. In: ICRA Workshop Open Source Software, 1–6 (2009).
19. Murphy, K.: The Bayes Net Toolbox for MATLAB. Computing Science and Statistics. **33**(2), 1024–1034 (2001).