Dimensions of Uncertainty in Evidential Grid Maps

Thomas Reineking and Joachim Clemens

Cognitive Neuroinformatics, University of Bremen, Germany
{reineking,clemens}@uni-bremen.de

Abstract. We show how a SLAM algorithm based on belief function theory can produce evidential occupancy grid maps that provide a mobile robot with additional information about its environment. While uncertainty in probabilistic grid maps is usually measured by entropy, we show that for evidential grid maps, uncertainty can be expressed in a three-dimensional space and we propose appropriate measures for quantifying uncertainty in these different dimensions. We analyze these measures in a practical mapping example containing typical sources of uncertainty for SLAM. As a result of the evidential representation, the robot is able to distinguish between different sources of uncertainty (e.g., a lack of measurements vs. conflicting measurements) which are indistinguishable in the probabilistic framework.

1 Introduction

In order to generate a spatial representation of an environment, a mobile robot needs to solve the problem of simultaneous localization and mapping (SLAM) [5]. Occupancy grid maps are a popular type of spatial representation for SLAM and discretize the environment using a grid structure where each grid cell is either occupied or empty (denoted by o and e) [6]. Usually, the state of a grid cell is modeled probabilistically with a single probability $P(o)$. The problem of this approach is that different states of belief are mapped to the same probability distribution. For example, a uniform probability distribution can represent a complete lack of measurements just as it can represent conflicting measurements. In this paper, we show that the uncertainty in occupancy grid maps has multiple dimensions that cannot be uniquely captured by a probabilistic representation. Instead, we propose to model grid maps using belief functions in order to make these different dimensions explicit.

Belief function theory [19,22] (also called Dempster-Shafer theory or evidence theory) can be viewed as a generalization of Bayesian probability theory. For belief functions, probability mass is not just assigned to the singletons of a domain (here o and e) but to all subsets of the domain (including $\{o, e\}$ and $\emptyset$). While a probability function can only capture the ratio between $P(o)$ and $P(e)$, a belief function can additionally make a lack of evidence explicit by the mass assigned to the disjunction $\{o, e\}$ and it can make conflicting evidence explicit by the mass assigned to $\emptyset$. As a result, belief functions are able to represent different

C. Freksa et al. (Eds.): Spatial Cognition 2014, LNAI 8684, pp. 283–298, 2014.

types of uncertainty and therefore enable the robot to distinguish belief states that are indistinguishable using a probabilistic model.

There are a number of works in which belief functions are used to model the uncertainty in occupancy grid maps [14,17]. The additional parameters provided by evidential grid maps have been used to solve problems like assessing the quality of maps [1] and detecting moving objects [13]. However, all of these works have in common that they do not consider the joint estimation problem underlying SLAM and only consider the mapping part (by assuming perfect localization information). In this case, the localization error, which is the major source of uncertainty for SLAM, is entirely ignored. In contrast, in [16] we described how SLAM can be modeled in the belief function framework. The resulting algorithm produces evidential grid maps that reflect the full uncertainty associated with SLAM, including localization uncertainty. The major contribution in this paper is an analysis of the different types of uncertainty in the evidential grid maps produced by the algorithm, both on a theoretical as well as on an empirical level.

The remainder of this paper is structured as follows. In Sect. 2, the belief function formalism is briefly introduced. The evidential SLAM approach based on this formalism is presented in Sect. 3 along with evidential sensor models for laser scanners. The different dimensions of uncertainty are analyzed theoretically in Sect. 4 and corresponding measures are proposed. Practical examples of the different types of uncertainty in generated grid maps are presented in Sect. 5. The paper concludes with a discussion in Sect. 6.

2 Belief Function Theory

The term "belief function" is a general term which can refer to several equivalent representations. The most fundamental belief representation is a *mass function*. A mass function m is a mapping $m : \mathcal{P}(\Theta) \to [0, 1]$ with

$$\sum_{X \subseteq \Theta} m(X) = 1, \tag{1}$$

where $\mathcal{P}(\Theta)$ is the power set of the (usually finite) domain Θ. The value $m(X)$ is the amount of belief strictly committed to set X. A mass assignment to a set X represents complete ignorance about the belief distribution over subsets of X. If a mass function satisfies $m(\emptyset) = 0$, it is called *normalized*. If it also satisfies $\sum_{x \in \Theta} m(x) = 1$, it is called *Bayesian* (it then simply represents a probability function).

A *plausibility function pl* expresses how much belief mass potentially supports a set X. It is therefore sometimes interpreted as an upper probability bound. The plausibility $pl(X)$ of $X \subseteq \Theta$ can be directly derived from the corresponding mass function m and is defined as

$$pl(X) = \sum_{Y \subseteq \Theta, Y \cap X \neq \emptyset} m(Y) \tag{2}$$

with $pl(\emptyset) = 0$.

Like probability functions, belief functions can be conditional. This is written as $m[Y](X)$, which means "the mass of set X given set Y" (note that the order is different from the usual conditional probability notation $P(X|Y)$).

For inference, each piece of evidence (e.g., a sensor measurement) is represented by a separate belief function and these belief functions are then successively combined in order to fuse the underlying evidence. The combination is usually performed using the *conjunctive rule* (also referred to as *Dempster's rule* if mass assignments to $\emptyset$ are not desired). Let e_1 and e_2 denote distinct pieces of evidence and let $m[e_1]$ and $m[e_2]$ denote the mass functions defined over the same domain Θ which are induced by the respective pieces of evidence. The mass function $m[e_1, e_2]$ resulting from the combination using the conjunctive rule $\textcircled{\cap}$ is defined as

$$m[e_1, e_2](X) = \sum_{Y \cap Z = X} m[e_1](Y)\, m[e_2](Z), \quad \forall X \subseteq \Theta. \tag{3}$$

3 Evidential FastSLAM

The evidential SLAM algorithm proposed in [16] uses a Rao-Blackwellized particle filter [3] to approximate the joint belief distribution of the map and the robot's path. It therefore constitutes a generalization of the well-known FastSLAM algorithm [12]. Let $x_{0:t} = x_0, \ldots, x_t$ denote the sequence of robot poses over time and let Y denote the map. The aim for evidential SLAM is to compute the joint belief distribution $m[z_{0:t}, u_{1:t}](x_{0:t}, Y)$ where $z_{0:t}$ is the sequence of measurements recorded over time and $u_{1:t}$ is the sequence of robot controls describing pose changes. In order to make computing this joint distribution feasible, it is assumed that the marginal distribution over the path is a probability density function. In this case, the joint belief distribution can be factorized into a probabilistic path component and a conditional evidential map component:

$$m[z_{0:t}, u_{1:t}](x_{0:t}, Y) = p(x_{0:t}|z_{0:t}, u_{1:t})\, m[x_{0:t}, z_{0:t}](Y). \tag{4}$$

This factorization corresponds to a generalized version of the product rule for probabilities, a proof of which is provided in [15]. The next two subsections describe how the path distribution $p(x_{0:t}|z_{0:t}, u_{1:t})$ and the conditional map distribution $m[x_{0:t}, z_{0:t}](Y)$ can be computed. In Sect. 3.3, the resulting particle filter algorithm is presented and, in Sect. 3.4, evidential sensor models for laser scanners are presented.

3.1 Localization

Because the path distribution is modeled probabilistically, computing it is similar to classical Markov localization [27] where pose changes are modeled in a prediction step and measurements are incorporated in a correction step. In the prediction step, the motion model $p(x_t|x_{t-1}, u_t)$ is applied to the prior distribution $p(x_{0:t-1}|z_{0:t-1}, u_{1:t-1})$ from time $t-1$ in order to compute the proposal distribution $p(x_{0:t}|z_{0:t-1}, u_{1:t})$ at time t, which is given by

$$p(x_{0:t}|z_{0:t-1}, u_{1:t}) = p(x_t|x_{t-1}, u_t)\, p(x_{0:t-1}|z_{0:t-1}, u_{1:t-1}). \tag{5}$$

In the correction step, the posterior $p(x_{0:t}|z_{0:t}, u_{1:t})$ is computed from the proposal distribution and the measurement likelihood using the generalized Bayesian theorem [23]. Though the prior and posterior for the path are assumed to be probability distributions, the likelihood can be a general belief function [15], in this case the plausibility $pl[x_{0:t}, z_{0:t-1}](z_t)$:

$$p(x_{0:t}|z_{0:t}, u_{1:t}) \propto pl[x_{0:t}, z_{0:t-1}](z_t)\, p(x_{0:t}|z_{0:t-1}, u_{1:t}). \tag{6}$$

The likelihood $pl[x_{0:t}, z_{0:t-1}](z_t)$ depends on the entire history of measurements and states. In order to avoid this problem, we condition the likehood on the estimated map Y with

$$pl[x_{0:t}, z_{0:t-1}](z_t) = \sum_{Y \subseteq \Theta_Y^M} pl[x_t, Y](z_t)\, m[x_{0:t-1}, z_{0:t-1}](Y), \tag{7}$$

where Θ_Y^M denotes the map space (defined below). Here, $pl[x_t, Y](z_t)$ represents the *forward sensor model* and $m[x_{0:t-1}, z_{0:t-1}](Y)$ represents the map belief at time $t-1$. The sum over the power set of the map space in (7) may appear intractable but, as shown in [16], by making appropriate independence assumptions, the likelihood can be efficiently computed with time complexity proportional to the number of grid cells.

3.2 Grid Mapping

The power of the original FastSLAM algorithm lies in the fact that cells in the map become approximately independent if conditioned on the robot's path $x_{0:t}$. The same principle is exploited in the evidential FastSLAM algorithm. Let Y_i denote the evidential variable representing the i-th grid cell with $Y_i \subseteq \Theta_Y = \{o, e\}$ and $1 \leq i \leq M$ where M denotes the total number of grid cells. The entire map Y is then a subset of the product space with $Y \subseteq \Theta_Y^M$. The joint distribution over all cells can be factorized into marginal cell distributions:

$$m[x_{0:t}, z_{0:t}](Y) = \prod_{i=1}^{M} m[x_{0:t}, z_{0:t}](Y_i). \tag{8}$$

In this case, each cell can be updated independently over time. This is done by combining the prior cell belief $m[x_{0:t-1}, z_{0:t-1}](Y_i)$ at time $t-1$ with the cell belief $m[x_t, z_t](Y_i)$ induced by the current measurement z_t using the conjunctive rule:

$$m[x_{0:t}, z_{0:t}](Y_i = \cdot) = m[x_{0:t-1}, z_{0:t-1}](Y_i = \cdot) \textcircled{\cap} m[x_t, z_t](Y_i = \cdot). \tag{9}$$

Here, $m[x_t, z_t](Y_i)$ represents the *inverse sensor model*. The initial belief $m(Y_i)$ at time $t = 0$ is assumed to be vacuous with $m(Y_i = \Theta_Y) = 1$, $1 \leq i \leq M$ (unless there is prior knowledge indicating otherwise). This expresses the total lack of evidence at the beginning and differs from a probabilistic model where the initial distribution is usually assumed to be uniform.

3.3 Algorithm

The algorithm for approximating the joint distribution in (4) is very similar to the original FastSLAM algorithm. It is based on a Rao-Blackwellized particle filter where each particle represents a complete path and a corresponding map belief function. This means, at time t, the k-th particle is a tuple $(x_{0:t}^{[k]}, m[x_{0:t}^{[k]}, z_{0:t}](Y))$. Like in most probabilistic particle filters, measurements are incorporated using importance sampling. The set of particles is updated recursively over time by repeatedly applying the following four steps in order to obtain the particle set at time t from the particle set at time $t-1$.

1. Prediction step: Sample a new pose $x_t^{[k]}$ from the motion model $p(x_t|x_{t-1}^{[k]}, u_t)$ for each particle in order to incorporate control u_t and update the robot's path $x_{0:t}^{[k]}$ as defined by (5).
2. Correction step: Compute importance weight $w_t^{[k]} = pl[x_{0:t}^{[k]}, z_{0:t-1}](z_t)$ for each particle using the forward sensor model $pl[x_t^{[k]}, Y](z_t)$ and the current map belief $m[x_{0:t-1}^{[k]}, z_{0:t-1}](Y)$ as defined by (7).
3. Map update: Update the current map belief $m[x_{0:t-1}^{[k]}, z_{0:t-1}](Y)$ of each particle using the inverse sensor model $m[x_t^{[k]}, z_t](Y)$ and the conjunctive rule of combination.
4. Resampling: Resample particles with probability proportional to the importance weights. This results in a particle set representing the joint path/map distribution $m[z_{0:t}, u_{1:t}](x_{0:t}, Y)$ reflecting all measurements and controls up to time t.[1]

The time complexity of the algorithm is $O(KM)$ (K denotes the number of particles) assuming the complexity of the sensor models is $O(M)$. The overall complexity is thus identical to that of classical FastSLAM aside from a constant overhead caused by the fact that, for each cell, three parameters need to be updated instead of one.

3.4 Sensor Models

For the SLAM algorithm described above, an evidential forward sensor model $pl[x_t, Y](z_t)$ and a corresponding inverse sensor model $m[x_t, z_t](Y)$ need to be specified. In this paper, we focus on laser scanners where $z_t = (z_{t;1}, \ldots, z_{t;N})^T$ is a vector of range measurements for different angles. The models considered here are adaptations of the evidential sonar models presented in [16] (the main difference is that a laser beam is much narrower than the measurement cone of a sonar sensor). Because the resulting equations are quite complex, they are omitted here and the interested reader is encouraged to refer to the original paper.

[1] The resampling step is only performed if the importance weights diverge too much, otherwise the importance weights are maintained over time in order to minimize sampling errors [4].

Examples of the forward and inverse sensor models for a single laser beam are shown in Fig. 1. The forward sensor model computes the plausibility of a single range measurement $z_{t;j}$ given the current state x_t and the evidential map Y. The forward model captures both noisiness and complete randomness of measurements. In order to make computing it tractable, the simplifying assumption is made that a laser beam always hits the closest occupied cell located along the beam's path. For cells representing the disjunctive state Θ_Y, the measurement plausibility is maximal but the plausibility of measuring more distant obstacles is not reduced because Θ_Y also includes the possibility that the cell is empty. The plausibility $pl[x_t, Y](z_t)$ of the entire vector z_t is simply defined as the product of the individual range measurement plausibilities $pl[x_t, Y](z_{t;j})$.

The inverse sensor model provides the robot with a belief distribution representing a "local map" obtained from a single range measurement $z_{t;j}$. A number of evidential inverse models have been proposed in the literature over time, e.g., [29,7]. All these works ignore the forward model and directly specify the inverse model, usually in a heuristic manner. In contrast, the inverse sensor model here is directly derived from the forward sensor model using the generalized Bayesian theorem. It thus reflects all the parameters of the forward sensor model. The inverse model for the entire measurement vector z_t results from combining the belief functions corresponding to the individual range measurements $z_{t;j}$ using Dempster's rule. The time complexity of evaluating both the forward and the inverse sensor model is proportional to the number of cells that intersect with the laser beams, and evaluating the models is quite fast in practice.

4 Dimensions of Uncertainty

In this section, we analyze the different dimensions of uncertainty in evidential grid maps on a theoretical level. In particular, we show how the different dimensions of uncertainty can be measured and we compare them to the Bayesian notion of uncertainty. To compare evidential and probabilistic grid maps, the question of how one representation relates to the other first needs to be addressed. In belief function theory, there is a simple operation called the "pignistic transformation" for projecting a belief function onto a probability function [24]. Usually, the pignistic transformation is applied in the context of decision making based on belief functions but, here, it is used to analyze the Bayesian uncertainty of an evidential grid map. The pignistic transformation for a cell is defined as

$$P(o) = \frac{m(o) + \frac{1}{2}m(\Theta_Y)}{1 - m(\emptyset)}. \tag{10}$$

The mass assigned to Θ_Y is split up and evenly distributed over o and e. In addition, normalization is performed to remove the mass on $\emptyset$.

In case of probabilistic grid maps, uncertainty can be measured using Shannon entropy [20], which is defined as

$$H(P) = -\sum_{x \in \Theta} P(x) \log_2 P(x). \tag{11}$$

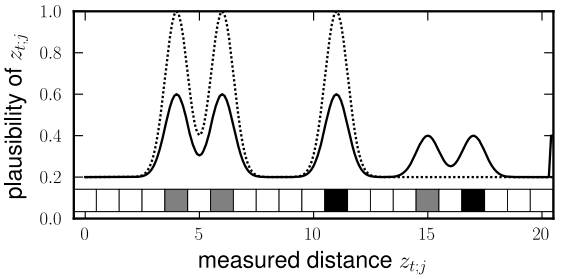

(a) Forward sensor model

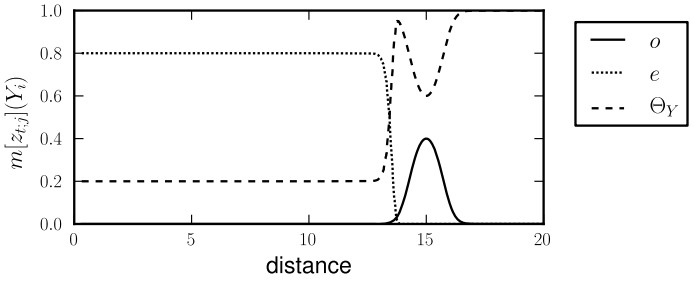

(b) Inverse sensor model

Fig. 1. Evidential sensor models for a laser scanner. The squares at the bottom in (a) represent the grid cells of a 1D map: white means e, black means o, and gray means Θ_Y. The dotted line in (a) represents the forward sensor model for a given map Y (without uncertainty) while the solid line represents the forward sensor model conditioned on an uncertain map (in the latter case, mass 0.5 is assigned to e for each of the black and gray cells). Fig. (b) shows the inverse sensor model for a measurement of $z_t = 15$. Note that the noise parameter in these examples is higher compared to the laser scanner considered in Sect. 5 in order to obtain a better visualization. (Figures adopted from [16].)

This measure of uncertainty yields 1 for a uniform cell distribution and it yields 0 if a cell is known with certainty to be either occupied or empty. In contrast, belief functions contain multiple dimensions of uncertainty and thus require multiple measures of uncertainty [10]. The two main dimensions are referred to as *non-specificity* and *conflict*. Non-specificity results from the fact that mass can be assigned to arbitrary non-singleton subsets of the domain where set cardinality represents a form of uncertainty (here, $|\Theta_Y| = 2$). Conflict corresponds more to the classical notion of entropy and reflects the uncertainty resulting from mass assigned to mutually exclusive states (o and e in this case). Because we are considering unnormalized mass functions with $m(\emptyset) \geq 0$, we further distinguish between what we call *internal conflict* and *external conflict*. While internal conflict represents the aforementioned entropy-like uncertainty associated with the masses assigned to o and e, external conflict represents how strongly the evidence underlying a belief function contradicts each other, which is represented

by the mass assigned to $\emptyset$. These three dimensions of uncertainty are described and analyzed in more detail in the following.

4.1 Non-Specificity (NS)

Non-specificity results from a lack of evidence where neither o nor e is supported. It is thus represented by the mass assigned to the disjunction Θ_Y. This intuition also has a well-justified theoretical basis in the form of the Hartley measure [8]. The Hartley measure states how much uncertainty results from the fact that the true state is contained in a set of possible states. It is simply defined as $\log_2 |X|$ where $|X|$ denotes the set cardinality. For belief functions, the generalized Hartley measure computes the amount of non-specificity associated with a mass function [9], and it is defined by

$$NS(m) = \sum_{X \subseteq \Theta, X \neq \emptyset} m(X) \log_2 |X|. \tag{12}$$

For the case of occupancy grid maps, it simply reduces to

$$NS(m) = m(\Theta_Y). \tag{13}$$

This definition thus confirms the intuition that the mass assigned to Θ_Y reflects the amount of non-specificity associated with the distribution of a grid cell.

4.2 Internal Conflict (IC)

The internal conflict essentially describes the relation between $m(o)$ and $m(e)$. Like Shannon entropy, it reaches its maximum if $m(o)$ and $m(e)$ are equal and it becomes 0 in case either $m(o)$ or $m(e)$ is 0. We refer to this dimension as "internal" because, other than the external conflict which results from combining multiple belief functions, internal conflict is a property of a single belief function. In order to quantify it, we use the measure of dissonance proposed in [28].[2] It is defined as

$$Dis(m) = -\sum_{X \subseteq \Theta} m(X) \log_2 pl(X). \tag{14}$$

Because the measure of dissonance can only handle normalized belief functions, we first perform normalization before computing the measure. The result is then scaled with $1 - m(\emptyset)$, which also ensures that the measure does not exhibit a discontinuity at $m(\emptyset) = 1$:

$$IC(m) = -\left(1 - m(\emptyset)\right)\left(m'(o) \log_2 pl'(o) + m'(e) \log_2 pl'(e)\right), \tag{15}$$

$$m'(X) = \begin{cases} \frac{m(X)}{1-m(\emptyset)} & \text{if } X \neq \emptyset \\ 0 & \text{if } X = \emptyset \end{cases}. \tag{16}$$

[2] Despite a long history of research, none of the measures proposed over time satisfy all the properties required of a general measure of internal conflict [9]. However, the limitation of the dissonance measure (violation of subadditivity for joint spaces) can be ignored here because we do not consider the internal conflict over joint spaces.

Here, m' and pl' denote the normalized mass and plausibility functions. Note that the internal conflict measure reduces to classical Shannon entropy in case of a Bayesian belief function.

4.3 External Conflict (EC)

The external conflict measures the conflict resulting from combining different measurements over time. The more these measurements contradict each other regarding the state of a cell, the higher the external conflict. The external conflict is represented by the mass assigned to $\emptyset$. Such an assignment to $\emptyset$ can also be interpreted as an open-world assumption where the true state is not captured by the set of possible states [21] (e.g., a cell is neither completely occupied nor completely empty but something in between).

Analogously to non-specificity, one possibility to measure external conflict would simply be to consider the mass assigned to $\emptyset$. However, this value quickly approaches 1 if sufficiently many measurements are combined where each measurement induces a small amount of external conflict. Instead, it is more useful to consider a logarithmic measure which makes it possible to differentiate between small differences of mass values assigned to $\emptyset$. Here, we use the *weight of conflict* proposed in [19], which is defined as

$$Con(m_1, \ldots, m_n) = -\log_2 \left(1 - \sum_{X_1 \cap \ldots \cap X_n = \emptyset} m_1(X_1) \cdot \ldots \cdot m_n(X_n) \right). \quad (17)$$

This measure states how strongly multiple belief functions $m_1, \ldots, m_n$ contradict each other if combined conjunctively. Because the conflict associated with combinations is represented by $m(\emptyset)$, the weight of conflict can be used as a measure of external conflict:

$$EC(m) = -\log_2(1 - m(\emptyset)). \quad (18)$$

4.4 Comparison

In Fig. 2, each of the proposed measures is plotted in relation to the underlying mass function. While non-specificity and internal conflict are classical dimensions of uncertainty for belief functions, the role of external conflict is a somewhat different one. A mathematical difference is obvious because NS and IC are bounded by 1 while the EC measure goes to infinity if $m(\emptyset)$ goes to 1. But the difference is also a conceptual one. Typically, belief functions representing a single piece of evidence (e.g., a single measurement) do not exhibit external conflict (i.e., $m(\emptyset) = 0$) and, only as a result of combination with other belief functions, external conflict is created. Therefore, external conflict does not describe a property of a single piece of evidence but rather describes a property of the combination process. This is also reflected by the fact that external conflict

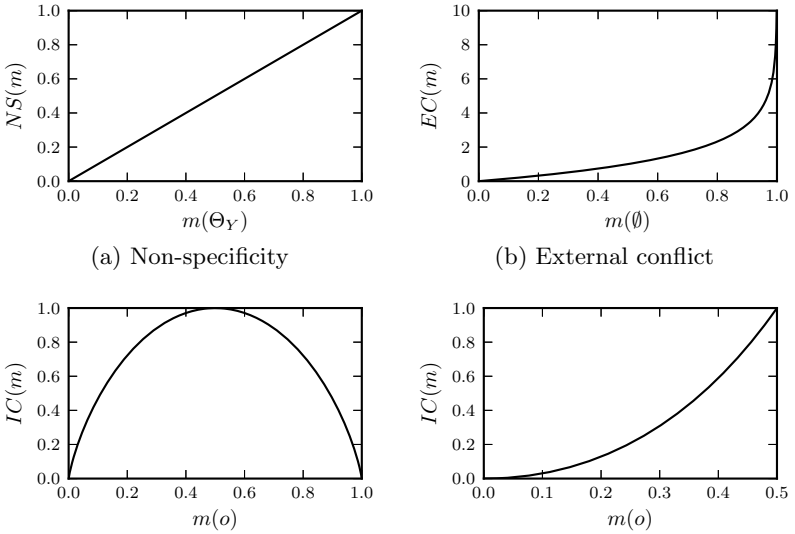

(a) Non-specificity

(b) External conflict

(c) Internal conflict where $m(e) = 1 - m(o)$

(d) Internal conflict where $m(e) = m(o)$ and $m(\Theta_Y) = 1 - m(o) - m(e)$

Fig. 2. Measures for different dimensions of uncertainty. The internal conflict shown in (c) and (d) is actually a function of three parameters and can therefore not be visualized in its entirety. Instead, (c) shows the IC for a Bayesian belief function, in which case it is identical to entropy. In (d), the IC is shown for the case where $m(e) = m(o)$ and where the remaining mass is assigned to Θ_Y.

grows monotonically with additional combinations. In this regard, external conflict represents a kind of "meta uncertainty" indicating whether the underlying evidence is in agreement or not.

A comparison of the evidential dimensions of uncertainty with the Bayesian notion of entropy is shown in Fig. 3. Here, several prototypical belief states are plotted in the 2D uncertainty space defined by non-specificity and internal conflict, which are then projected to the 1D uncertainty space corresponding to entropy. External conflict is ignored in these examples because entropy is invariant with respect to this dimension. It can be seen that a state of maximum non-specificity ($m(\Theta_Y) = 1$) has the same entropy like a state of maximum internal conflict ($m(o) = m(e) = 0.5$) when mapped to a probability function using (10). Thus, very different evidential belief states are mapped to the same entropy values, showing the ambiguity of the Bayesian representation.

5 Grid Mapping Example

This section describes a practical mapping example where different types of uncertainty are represented by an evidential grid map. The example is based on a run of a simulated mobile robot equipped with a laser scanner exploring a virtual

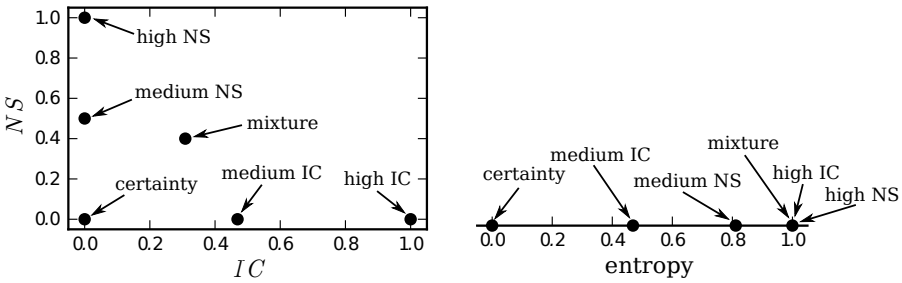

(a) Evidential uncertainty (b) Bayesian uncertainty

Fig. 3. Comparison of 2D evidential uncertainty (without EC) and 1D Bayesian uncertainty. The 2D points from (a) are projected to 1D points in (b).

2D environment.[3] The evidential SLAM algorithm presented in Sect. 3 was used to generate an evidential grid map by selecting the particle with the highest cumulative importance weight (using $K = 100$). This map, along with the ground truth and a probabilistic grid map obtained from the pignistic transformation, is shown in Fig. 4. The corresponding uncertainty measures are shown in Fig. 5.

What is directly noticeable is that the evidential uncertainty measures are sensitive to different types of uncertainty because they are maximal in different areas of the map. *NS* is high for areas where no or only few measurements have been recorded. *IC* and *EC* are high in the vicinity of obstacles due to measurement noise, although *EC* also indicates effects like localization errors and changes in the environment. In contrast, Shannon entropy appears to represent an "aggregate" measure of uncertainty where the different maxima of the evidential measures result in high entropy. In the following, the different causes of uncertainty in this example are analyzed in more detail.

Partial Exploration: The *NS* measure (i.e., the mass on Θ_Y) clearly shows which parts of the environment the robot has not fully explored yet. This is visible both for the room the robot did not enter as well as for the areas on the outside. For the probabilistic representation, this belief state is represented by a uniform distribution where entropy is maximal. However, entropy is also high in other parts of the map, making the uncertainty resulting from not having explored an area indistinguishable from the uncertainty resulting from other causes.

Measurement Noise: Measurement noise in the vicinity of obstacles results from both the noisiness of the sensor and from discretization errors caused by the grid-based representation. It is reflected by high values for the *IC* measure as well as for the *EC* measure. The difference between the two measures is

[3] For the simulation, the "Gridmap Navigation Simulator" was used, which is part of the Mobile Robot Programming Toolkit available at `www.mrpt.org`.

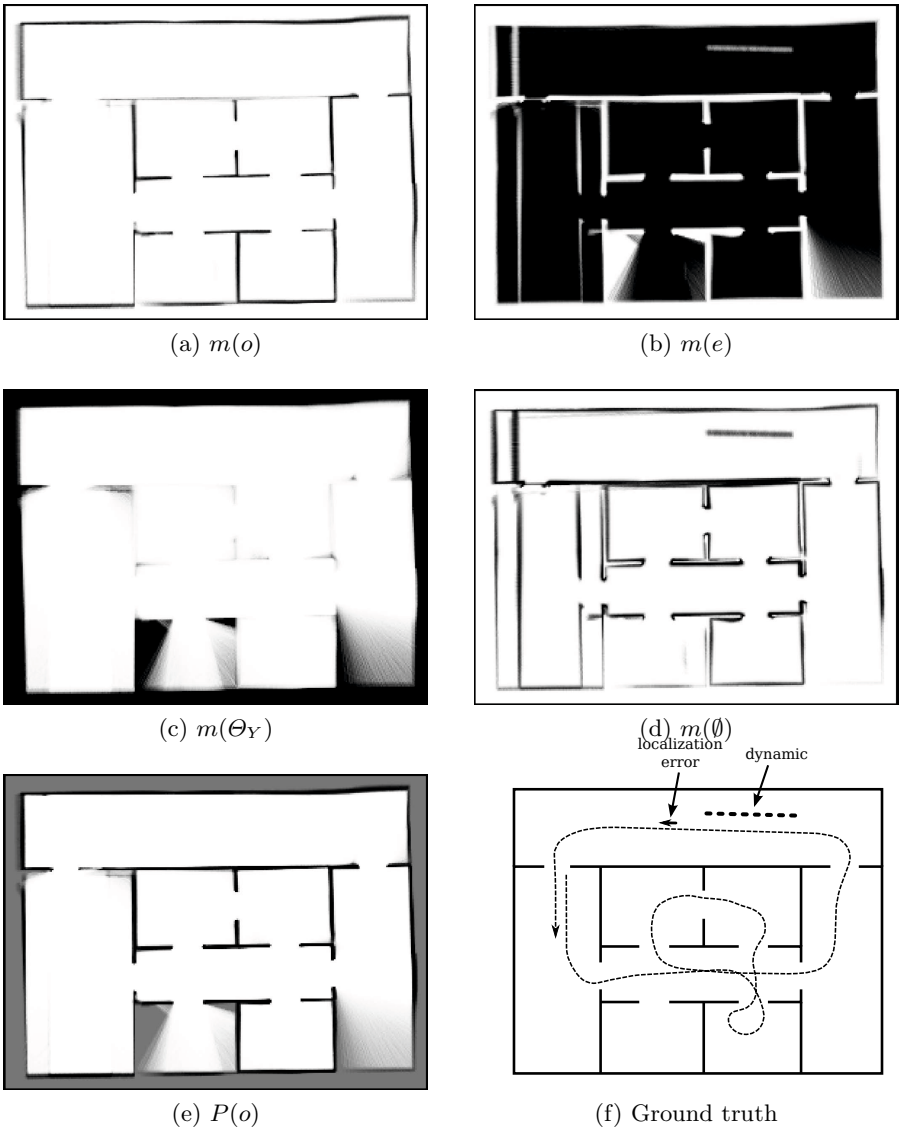

(a) $m(o)$ (b) $m(e)$

(c) $m(\Theta_Y)$ (d) $m(\emptyset)$

(e) $P(o)$ (f) Ground truth

Fig. 4. Grid map generated by the evidential FastSLAM algorithm. Plots (a) to (d) show the different mass components of the evidential grid map (black represents a mass value of one and white represents zero). The corresponding probabilistic grid map is shown in (e) while the ground truth, including the robot's path, is shown in (f). During the run, a localization error and dynamics in the environment were simulated at the indicated locations (the former by displacing the robot and the latter by simulating the presence of a moving object blocking some of the laser beams).

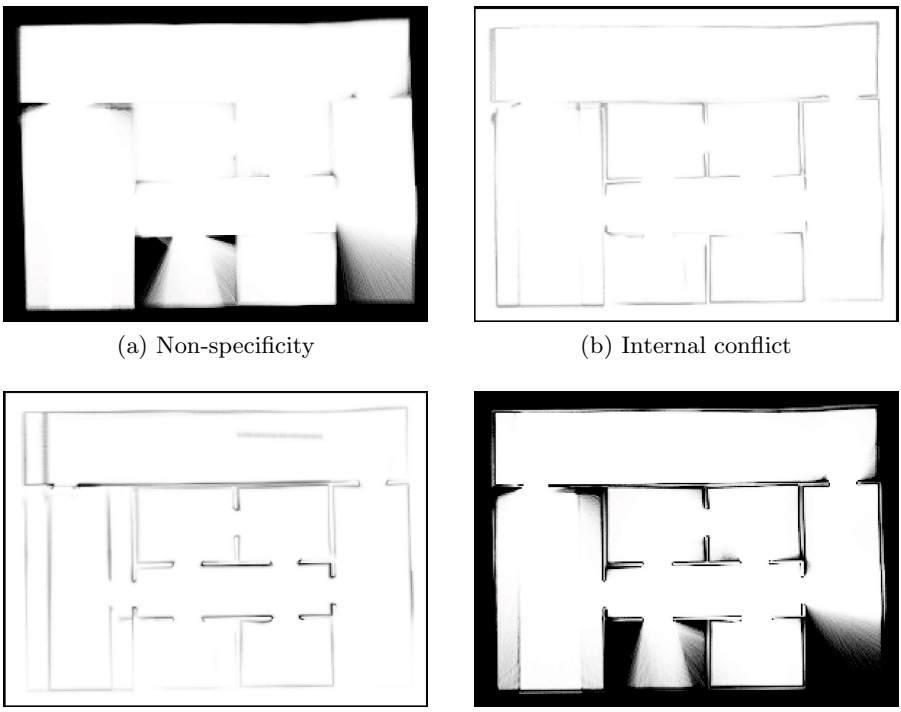

(a) Non-specificity (b) Internal conflict

(c) External conflict (d) Bayesian entropy

Fig. 5. Uncertainty measures corresponding to the map shown in Fig. 4. Note that the measures are scaled for improved visibility where a darker color indicates a higher value (the maximum in (a) and (d) is at 1, the maximum of (b) is at 0.34, and the maximum of (c) is at 7.58).

that IC responds most strongly if the underlying evidence is "balanced", e.g., if an equal number of measurements favors particular states (this also applies to entropy). In contrast, for a high EC value, it is sufficient that there is at least some contradictory evidence, even if the majority of evidence favors one state. In this regard, the EC measure can be interpreted as an "or-like" operation regarding the contradictory nature of the underlying evidence. In addition, both measures are sensitive to the number of measurements because EC grows with additional measurements if these contradict previous measurements at least by a small degree while the IC measure is scaled by the amount of mass assigned to $\emptyset$ (see (15)) and therefore tends to become smaller with more measurements. An example of this effect are the high EC values at the walls in the center hallway and at the lower wall of the top hallway caused by a large number of measurements where IC is relatively small. In contrast, in the lower right part of the map where few measurements were recorded, EC values are low and IC is high.

Localization Error: Localization errors usually lead to increased EC because measurements automatically contradict each other in this case. Like measurement noise, localization errors are most visible in the vicinity of obstacles. In order to better visualize this effect, we created an artificial localization error in the top hallway by displacing the robot by 2m (see (f) in Fig. 4). This error is clearly visible in Fig. 5 (c) where the two wall structures on the very left actually correspond to a single wall. In contrast, the entropy shown in Fig. 5 (d) gives almost no indication of this localization error (the same applies to (b) and (d) vs. (e) in Fig. 4).

Dynamic Environment: Another typical problem for SLAM is that the environment is usually assumed to be static while, in reality, environments often contain dynamic elements. We simulated the effect of an object moving alongside the robot by setting the laser measurements to a constant value (with additional noise) for a small range of angles in the top hallway (see (f) in Fig. 4). For example, this could correspond to a person walking next to the robot. While the robot can still reliably localize itself in this case, the measurements corresponding to the moving object clearly contradict the other measurements (current and past ones), which leads to EC (visible in Fig. 5 (c) as the gray line in the middle of the top hallway). By comparison, entropy fails to capture this effect because the majority of measurements do not indicate the presence of an obstacle (see also (b) and (d) vs. (e) in Fig. 4).

6 Discussion

In this paper, we have shown how a SLAM algorithm based on belief functions can produce evidential grid maps that provide a mobile robot with additional information about its environment. Compared to probabilistic grid maps, evidential grid maps contain multiple dimensions of uncertainty and we have proposed measures for quantifying uncertainty in these different dimensions. In a practical mapping example, we have illustrated how different causes of uncertainty can be distinguished in the multiple-dimensional uncertainty space corresponding to belief functions whereas they are indistinguishable in case of a probabilistic representation.

One interesting direction for future work would be the problem of fusing map information from different sources [11]. For probabilistic grid maps, map fusion is usually based on an absolute independence assumption [26], which can be quite problematic. In contrast, belief function theory could be of great value here due to the extensive work on different combination rules [18]. Another direction for future work that has not been considered in the context of evidential mapping is that of active information gathering. The idea of minimizing expected entropy [2,30] has been successfully applied in the past in order to make a robot drive to locations that reduce uncertainty about the environment [25], resulting in optimal exploration strategies. Using a multi-dimensional uncertainty representation, it becomes possible for the robot to discern different types of uncertainty

during exploration. For example, high *NS* indicates that an area has not been explored and that moving to this area is likely to reduce overall uncertainty. High *IC* is likely to represent measurement uncertainty where the use of additional sensors may be necessary to further reduce uncertainty because previous measurements have been inconclusive. In contrast, the *EC* measure is able to capture "meta uncertainty" indicating that something is wrong (e.g., failed localization or a change in the environment). Overall, we belief that evidential representations are a valuable tool for modeling spatial uncertainty and that being able to distinguish between different types of uncertainty can provide a mobile robot with important information about its environment.

Acknowledgements. This work was supported by DFG (SFB/TR 8 Spatial Cognition, project "A5-[ActionSpace]") and DLR (projects "Enceladus Explorer" and "KaNaRiA").

References

1. Carlson, J., Murphy, R., Christopher, S., Casper, J.: Conflict metric as a measure of sensing quality. In: Proceedings of the 2005 IEEE International Conference on Robotics and Automation, pp. 2032–2039 (2005)
2. Cassandra, A.R., Kaelbling, L.P., Kurien, J.A.: Acting under uncertainty: Discrete bayesian models for mobile-robot navigation. In: Proceedings of the 1996 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 1996, vol. 2, pp. 963–972. IEEE (1996)
3. Doucet, A., De Freitas, N., Murphy, K., Russell, S.: Rao-Blackwellised particle filtering for dynamic Bayesian networks. In: Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence, pp. 176–183 (2000)
4. Doucet, A., Godsill, S., Andrieu, C.: On sequential Monte Carlo sampling methods for Bayesian filtering. Statistics and Computing 10(3), 197–208 (2000)
5. Durrant-Whyte, H., Bailey, T.: Simultaneous localization and mapping: part i. IEEE Robotics & Automation Magazine 13(2), 99–110 (2006)
6. Elfes, A.: Using occupancy grids for mobile robot perception and navigation. Computer 22(6), 46–57 (1989)
7. Gambino, F., Ulivi, G., Vendittelli, M.: The transferable belief model in ultrasonic map building. In: Proceedings of the Sixth IEEE International Conference on Fuzzy Systems, vol. 1, pp. 601–608. IEEE (1997)
8. Hartley, R.V.L.: Transmission of information. The Bell Labs Technical Journal 7(3) (1928)
9. Klir, G.J.: Uncertainty and information: foundations of generalized information theory. Wiley (2005)
10. Klir, G.J., Smith, R.M.: On measuring uncertainty and uncertainty-based information: Recent developments. Annals of Mathematics and Artificial Intelligence 32(1-4), 5–33 (2001)
11. Kurdej, M., Moras, J., Cherfaoui, V., Bonnifait, P.: Controlling remanence in evidential grids using geodata for dynamic scene perception. International Journal of Approximate Reasoning 55(1, part 3), 355–375 (2014)

12. Montemerlo, M., Thrun, S., Koller, D., Wegbreit, B.: FastSLAM: A factored solution to the simultaneous localization and mapping problem. In: Proceedings of the National Conference on Artificial Intelligence, pp. 593–598 (2002)
13. Moras, J., Cherfaoui, V., Bonnifait, P.: Credibilist occupancy grids for vehicle perception in dynamic environments. In: 2011 IEEE International Conference on Robotics and Automation (ICRA), pp. 84–89 (May 2011)
14. Pagac, D., Nebot, E.M., Durrant-Whyte, H.: An evidential approach to map-building for autonomous vehicles. IEEE Transactions on Robotics and Automation 14(4), 623–629 (1998)
15. Reineking, T.: Belief Functions: Theory and Algorithms. Ph.D. thesis, University of Bremen (February 2014), http://nbn-resolving.de/urn:nbn:de:gbv:46-00103727-16
16. Reineking, T., Clemens, J.: Evidential FastSLAM for grid mapping. In: 16th International Conference on Information Fusion (FUSION), pp. 789–796 (July 2013)
17. Ribo, M., Pinz, A.: A comparison of three uncertainty calculi for building sonar-based occupancy grids. Robotics and Autonomous Systems 35(3-4), 201–209 (2001)
18. Sentz, K., Ferson, S.: Combination of evidence in Dempster-Shafer theory. Tech. rep., Sandia National Laboratories (2002)
19. Shafer, G.: A Mathematical Theory of Evidence. Princeton University Press, Princeton (1976)
20. Shannon, C.E.: A mathematical theory of communication. The Bell System Technical Journal 27, 379–423 (1948)
21. Smets, P.: Belief functions. In: Smets, P., Mamdani, E.H., Dubois, D., Prade, H. (eds.) Non Standard Logics for Automated Reasoning, pp. 253–286. Academic Press, London (1988)
22. Smets, P., Kennes, R.: The transferable belief model. Artificial Intelligence 66, 191–234 (1994)
23. Smets, P.: Belief functions: The disjunctive rule of combination and the generalized Bayesian theorem. International Journal of Approximate Reasoning 9, 1–35 (1993)
24. Smets, P.: Decision making in the TBM: the necessity of the pignistic transformation. International Journal of Approximate Reasoning 38, 133–147 (2005)
25. Stachniss, C., Grisetti, G., Burgard, W.: Information gain-based exploration using rao-blackwellized particle filters. In: Robotics: Science and Systems, vol. 2 (2005)
26. Thrun, S., Burgard, W., Fox, D.: Probabilistic robotics. MIT Press, Cambridge (2005)
27. Thrun, S., Fox, D., Burgard, W., Dellaert, F.: Robust Monte Carlo localization for mobile robots. Artificial Intelligence 128(1-2), 99–141 (2001)
28. Yager, R.R.: Entropy and specificity in a mathematical theory of evidence. International Journal of General System 9(4), 249–260 (1983)
29. Yang, T., Aitken, V.: Evidential mapping for mobile robots with range sensors. IEEE Transactions on Instrumentation and Measurement 55(4), 1422–1429 (2006)
30. Zetzsche, C., Wolter, J., Schill, K.: Sensorimotor representation and knowledge-based reasoning for spatial exploration and localisation. Cognitive Processing 9, 283–297 (2008)