18th International Conference on Information Fusion
Washington, DC - July 6-9, 2015

Context Driven Tracking using Particle Filters

Rik Claessens*†, Gregor Pavlin*‡ and Patrick de Oude*
*University of Liverpool, United Kingdom
†Thales Research & Technology Delft, the Netherlands
‡University of Amsterdam, the Netherlands
Email: {rik.claessens, gregor.pavlin, patrick.deoude}@d-cis.nl

***Abstract*—This paper introduces a novel approach to robust tracking that combines Particle Filters (PFs) and estimation of physical constraints using Bayesian Networks (BNs). Heterogeneous Context Data (CD) describing the environment in which tracked objects move, is fused with the help of BNs. The resulting uncertain constraints are incorporated into the filtering process through a modification of the importance weights. Causal probabilistic models representing relations between the tracked objects and their environment are used to derive an updating rule that allows theoretically sound incorporation of uncertain constraints into PF. The approach allows incorporation of new types of CD without requiring any adaptation of the PF algorithm itself. The experimental results confirm that the presented method significantly improves the tracking accuracy in a relevant class of problems characterized by partial sensor coverage and low updating frequencies.**

I. Introduction

Tracking of objects of interest (i.e. targets) is an important element of increasingly complex security applications. Example applications include wildlife preservation, search & rescue, patrol missions and surveillance in littoral and urban environments. However, accurate target tracking in these types of domains is challenging because the data is noisy and the observation frequencies are low compared to more traditional tracking applications using radars in maritime and aviation domains. In addition, reasoning about an object's whereabouts requires the use of heterogeneous data types stemming from disparate sources and, the knowledge about the dynamics of the targets is uncertain to a great extent.

The approach presented in this paper addresses these challenges by combining a basic PF tracking algorithm [1], [2], [3] and estimation of physical constraints based on the fusion of heterogeneous CD. Examples of relevant CD in case of vehicle tracking are road maps and intelligence about mobility, such as reports on flooding, obstacles, traffic jams, etc.

Physical constraints are estimated in a separate process, outside of the tracking process. Consequently, the presented approach allows incorporation of different types of uncertain CD without requiring any adaptation of the PF algorithm. A key challenge tackled by the presented paper is theoretically sound interfacing between the systems fusing CD and the tracking processes. The chosen interfacing approach depends to a great extent on the used knowledge representations and inference algorithms. Firstly, it is assumed that CD originates from inherently noisy sources. Therefore the estimated constraints are also uncertain[1]. Secondly, given the representations used for the observation and process models in a PF, it is reasonable to formulate the uncertain constraints with the help of probability distributions conditioned on CD. As non-trivial correlations may exist between the different types of relevant CD, the uncertain constraints are estimated with the help of BNs [4], [5].

CD is introduced to the filtering process of PF at the updating step, by multiplying the importance weights with the outcomes of BNs. The multiplication of importance weights is similar to the solution presented in [6], [7], although a different derivation approach is taken. However, as exposed in [7], naive use of such multiplication of importance weights can result in significant inaccuracies if the constraints are uncertain. By using causal probabilistic models, we *i*) show why naive approaches to updating importance weights with uncertain constraints fail and *ii*) derive an updating strategy that allows correct use of Context Data. Causal models provide a theoretically sound basis for the derivation of updating rules that are straightforward and result in correct estimation using uncertain constraints under realistic conditions.

Overall, the presented solution fuses heterogeneous CD and runtime observations of a target by seamlessly combining two related types of theoretically sound approaches to Bayesian modelling and inference. Experimental results show that the presented updating strategy in combination with modified importance weights results in improved accuracy of the tracking processes using uncertain constraints, estimated from noisy CD.

The presented results are relevant for many real world applications where the knowledge about the environment is often not perfect. This aspect is usually not addressed systematically in the related work on combining CD and PFs, such as [8], [9], [10], [11].

The structure of the paper is as follows: Section II introduces tracking enhanced with context data and discusses an theoretically sound updating strategy. The experimental validation of the proposed methods is reported in Section III, followed by the conclusions and future work in Section IV.

[1]In this paper we assume uncertain "hard" constraints, i.e. physical constraints that cannot be violated by the target. The uncertainty is introduced through the lack of perfect knowledge about the presence of hard constraints at a specific location.

978-0-9964527-1-7©2015 ISIF

II. Context Driven Tracking

This section introduces an approach that supports a sound incorporation of uncertain CD into tracking systems based on Particle Filters. We first discuss the modelling and inference approaches used for tracking and estimation of uncertain constraints based on fusion of CD. This is followed by a derivation of a method for sound integration of uncertain constraints into a PF algorithm.

A. Modelling Techniques

The presented approach combines Particle Filters and Bayesian Networks, two related types of theoretically sound approaches to Bayesian modelling and inference. A PF approximates a probability density over a set of hypotheses about the states of a tracked object by maintaining a set of particles. Each particle represents a possible assignment of the random variables that may include location, speed, target type, etc. In this paper we assume PFs where the posterior over the possible states, the so-called target distribution, is obtained through sequences of Sampling Importance Resampling steps, approximating sequential Bayesian updating shown in (1). PFs are used when exact calculation of the target density is intractable, as is usually the case with distributions in continuous state spaces.

$$P(x_{1:t}|z_{1:t}) = \frac{P(z_t|x_t)P(x_t|x_{t-1})}{P(z_t|z_{1:t-1})}P(x_{1:t-1}|z_{1:t-1}) \quad (1)$$

A PF algorithm uses a transition model $p(x_t|x_{t-1})$ to sample the next state of each particle, given its previous state. A sensor model $p(z_t|x_t)$ assigns an importance weight to a particle. The particle's importance weight represents the likelihood of the particle being a correct estimate given the available evidence, i.e. observations of the target. When tracking a target, these steps can be considered as predicting the target's next location and checking how likely the estimated target's state is, given the incoming observations, respectively.

CD, on the other hand, is used to compute probability distributions $P(m_{k,t}|\epsilon_{k,t})$, where $m_{k,t} = true$ means that a specific type of a target can move at a specific location A_k while $m_{k,t} = false$ means that an object cannot move at the location. $\epsilon_{k,t}$ denotes the entire data about the mobility collected at this location up to time t (i.e. CD). In this paper $P(m_{k,t}|\epsilon_{k,t})$ is called a Corrective Factor (CF) for all particles located at A_k. Thus, each particle is assigned a CF corresponding to the location A_k the particle is currently in, after the prediction step. In Section II-B, we show that the CD can influence the filtering process in a theoretically sound way by a simple modification of the importance weights: CF $P(m_{k,t}|\epsilon_{k,t})$ influences the re-sampling step in PF by using $P(z_t|x_t)P(m_{k,t}|\epsilon_{k,t})$ as importance weights.

In the targeted domains the CFs can be inferred by fusing different types of correlated CD. BNs allow efficient handling of correlations between such diverse CD and computation of $P(m_{k,t}|\epsilon_{k,t})$. Figure 1 shows the qualitative part of an example BN. The model captures relations between the mobility $m_{k,t}$ and other types of CD represented by the following discrete random variables:

- $s_{k,t}$ is a multi-state variable that represents the type of ground surface in area A_k at time t. The states of the variable can represent a road, flat hard surface, swamp, forest, water, etc. The prior distribution over the states of this variable can be obtained from a GIS map.
- $si_{k,t}$ is a multi-state variable representing any available intelligence about the surface, such as roadblocks, floods, etc.
- v is a multi-state variable that represents the type of the target vehicle, such as cars, boats, etc.

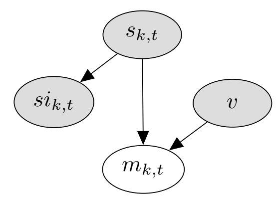

Fig. 1: A causal model describing a BN used for computing CFs. Variables s, si and v represent surface, surface intel and vehicle type respectively and constitute the context data variables (colored in grey) in area A_k at time t. The variable m represent the derived mobility.

Moreover, the CFs are organized in the Context Grid (CG), a framework that supports continuous updating of CFs and simplifies their use in the PF process. The framework implements a discrete, grid-like data structure. Each cell in the Context Grid has correspondence in the physical world. It represents an area A_k (or a part of it) associated with specific mobility conditions. Each area A_k represents a set of geographical locations at which the mobility conditions are constant, i.e. A_k is defined in such a way that the collection of CD at any point in A_k would yield the same values for all variables representing the CD. In general, however, such areas have irregular shapes which can be captured in the grid in two different ways:

1) Each grid cell has an irregular shape that encloses the actual surface of the modelled area A_k.
2) Area A_k is represented by a cluster of fine grained cells with identical shape and size. Each cell in such a cluster is associated with identical CF, as they all represent the same area A_k.

Overall, the CG associates each particle at location A_k with a CF, derived from all available CD collected within A_k. A new piece of CD about mobility results in an instantaneous recalculation of the corresponding CF, that can be used by the PF algorithm at the next resampling step. In this way, new CD immediately influences the tracking process.

B. Corrective Factors in Filtering Processes

This section discusses incorporation of Corrective Factors of the form $P(m_{k,t}|\epsilon_{k,t})$ into the PF process. We show that this type of derived CD can easily be incorporated into the PF process and supports sound approximation of posterior distributions, which converge to the correct posteriors in the limit as the number of particles approaches infinity.

First, we formulate a motion model in area A_k conditioned on the entire set of context data $\epsilon_{k,t}$ (i.e. CD collected in A_k up to time t). Such context data is relevant for the estimation of the mobility $m_{k,t}$ in this area. For the sake of clarity, we simplify the notation as follows: $\epsilon_{k,t}$ and $m_{k,t}$ are replaced by m_t and ϵ_t, respectively. In the following text it is implicitly assumed that these variables correspond to the data associated with the A_k in which the target is assumed to be located. Accordingly, we rewrite $P(m_{k,t}|\epsilon_{k,t})$ as $P(m_t|\epsilon_t)$ and formulate the motion model:

$$P(x_t|x_{t-1}, \epsilon_t) = \sum_{m_t} P(x_t|x_{t-1}, m_t)P(m_t|\epsilon_t), \qquad (2)$$

In the simplest case m_t is a binary variable: $m_t = true$ represents full mobility while $m_t = false$ represents zero mobility. In such a case the motion model $P(x_t|x_{t-1}, m_t)$ is equivalent to:

$$\begin{aligned} P(x_t|x_{t-1}, m_t = true) &= P(x_t|x_{t-1}) \\ P(x_t|x_{t-1}, m_t = false) &= 0, \end{aligned} \qquad (3)$$

where $P(x_t|x_{t-1})$ is the usual motion model assuming full mobility. With this definition, the motion model from (2) is reduced to:

$$\begin{aligned} &P(x_t|x_{t-1}, \epsilon_t) \\ &\quad = P(x_t|x_{t-1}, m_t = true)P(m_t = true|\epsilon_t) \\ &\quad = P(x_t|x_{t-1})P(m_t = true|\epsilon_t) \end{aligned} \qquad (4)$$

This factorized representation of $P(x_t|x_{t-1}, \epsilon_t)$ facilitates the derivation of a simple extension for importance weights that support theoretically sound approximation of the posterior distributions in the PF process. The derivation is based on the factorized representation of $P(x_{1:t}|z_{1:t}, \epsilon_{1:t})$, the posterior probability of an entire state sequence $x_{1:t}$ conditioned on all target observations $z_{1:t}$ and a CD set $\epsilon_{1:t}$ collected up to the t^{th} timestep. This approach is based on the derivation of a basic PF in [2]. Moreover, for the sake of clarity, at this stage we assume that with each timestep, the target moves to a different area A_k associated with a specific mobility m_t and evidence set ϵ_t. Later we will discuss a solution for more general situations where a target can remain within the same area A_k for a longer sequence of timesteps.

We first express the proposal distribution:

$$\begin{aligned} &P(x_{1:t}|z_{1:t-1}, \epsilon_{1:t-1}) \\ &\quad = P(x_t|x_{t-1})P(x_{1:t-1}|z_{1:t-1}, \epsilon_{1:t-1}) \end{aligned} \qquad (5)$$

Note that this distribution does not take into account the mobility m_t at time step t. It uses the normal motion model $P(x_t|x_{t-1})$ and takes the state sequence estimate from step $t-1$. Furthermore, the target distribution $P(x_{1:t}|z_{1:t}, \epsilon_{1:t})$ of a sequence of states can be expressed as follows:

$$\begin{aligned} &P(x_{1:t}|z_{1:t}, \epsilon_{1:t}) \\ &\quad = \eta P(z_t|x_{1:t}, z_{1:t-1}, \epsilon_{1:t})P(x_{1:t}|z_{1:t-1}, \epsilon_{1:t}), \end{aligned} \qquad (6)$$

where η is the normalizing constant. This expression can be rewritten as a product of simpler terms if the domain is Markovian. In such a case the following equalities hold for the two terms in (6):

$$P(z_t|x_{1:t}, z_{1:t-1}, \epsilon_{1:t}) = P(z_t|x_t) \qquad (7)$$

$$\begin{aligned} P(x_{1:t}|z_{1:t-1}, \epsilon_{1:t}) &= P(x_t|x_{1:t-1}, z_{1:t-1}, \epsilon_{1:t-1}, \epsilon_t) \\ &\quad \cdot\ P(x_{1:t-1}|z_{1:t-1}, \epsilon_{1:t-1}) \end{aligned} \qquad (8)$$

Decomposition in (8) is possible because of the following identities:

$$\begin{aligned} P(x_{1:t-1}|z_{1:t-1}, \epsilon_{1:t-1}) &= P(x_{1:t-1}|z_{1:t-1}, \epsilon_{1:t-1}, \epsilon_t) \\ &= P(x_{1:t-1}|z_{1:t-1}, \epsilon_{1:t}) \end{aligned} \qquad (9)$$

Namely, the mobility m_t at the new location does not have any influence on the sequence of target states up to the $t-1^{\text{th}}$ step. In other words, we would not know more about the state $x_{1:t-1}$ if we obtained the evidence ϵ_t about the mobility at the location that would be reached with the t^{th} step. Moreover, because of the Markovian assumptions the following equality holds for the first term in (8):

$$P(x_t|x_{1:t-1}, z_{1:t-1}, \epsilon_{1:t-1}, \epsilon_t) = P(x_t|x_{t-1}, \epsilon_t) \qquad (10)$$

This equality is justified, as the transition between the states x_{t-1} and x_t does not depend on the previous observations and the mobility at previous steps. The last assumption is justified, because the target could not reach x_{t-1} if $m_{t-1} = false$; i.e. this is correct if the target cannot get stuck (e.g. a ship will not run aground).

By using (7), (8), (9) and (10), we can rewrite the target distribution of (6) as:

$$\begin{aligned} P(x_{1:t}|z_{1:t}, \epsilon_{1:t}) &= \eta P(z_t|x_t)P(x_t|x_{1:t-1}, \epsilon_t) \\ &\quad \cdot\ P(x_{1:t-1}|z_{1:t-1}, \epsilon_{1:t-1}) \end{aligned} \qquad (11)$$

This equation expresses simple recursions between the subsequent posterior probabilities. Given the formulation of $P(x_t|x_{1:t-1}, \epsilon_t)$ in (2), we can further rewrite the target distribution:

$$\begin{aligned} &P(x_{1:t}|z_{1:t}, \epsilon_{1:t}) \\ &\quad = \eta P(z_t|x_t)\left[\sum_{m_t} P(x_t|x_{t-1}, m_t)P(m_t|\epsilon_t)\right] \\ &\quad\quad \cdot\ P(x_{1:t-1}|z_{1:t-1}, \epsilon_{1:t-1}) \end{aligned} \qquad (12)$$

The marginalization step in this expression can further be simplified as a consequence of (3), yielding:

$$\begin{aligned} P(x_{1:t}|z_{1:t}, \epsilon_{1:t}) = \eta P(z_i|x_t) P(x_t|x_{t-1}) \\ \cdot\ P(m_t = true|\epsilon_t) \\ \cdot\ P(x_{1:t-1}|z_{1:t-1}, \epsilon_{1:t-1}) \end{aligned} \tag{13}$$

This factorization introduces a CF $P(m_t = true|\epsilon_t)$ in the estimation process. By dividing (13) with (5), the terms $P(x_t|x_{t-1})$ and $P(x_{1:t-1}|z_{1:t-1}, \epsilon_{1:t-1})$ in the nominator and the denominator cancel out, yielding the following expression for the importance weight:

$$w_i = \eta P(z_i|x_t) P(m_t = true|\epsilon_t). \tag{14}$$

Thus, the CF $P(m_t = true|\epsilon_t)$ computed with the Bayesian mobility model is included directly into the posterior estimation process in a straightforward way, as an additional factor of the importance weights. If this estimate is correct, then the overall solution correctly approximates the posterior distribution over the state sequences $P(x_{1:t}|z_{1:t}, \epsilon_{1:t})$. Consequently, the sample for x_t is distributed according to the correct $P(x_t|z_{1:t}, \epsilon_{1:t})$ (see [2]). Note, however, that updating with a CF at each timestep is correct only if the target moves to a new area A_k with different mobility at each timestep. The following section discusses this issue in more detail.

C. Updating Strategies

In this section we take a closer look at the meaning of the cells in the Context Grid and the associated CFs $P(m_t|\epsilon_t)$ to derive sound rules for updating with uncertain context data.

Key to such a derivation is understanding of the relations between the states of a tracked object and the physical phenomena that influence the states of such an object[2]. Such relations are not deterministic, therefore we describe them with the help of causal probabilistic models. The graph in Figure 2 is a qualitative representation of the causal dependencies between the subsequent target states, observations of the target and physical phenomena influencing the target's states. For the sake of clarity, the notation in Figure 2 is simplified by replacing $s_{k,t}$ and $si_{k,t}$ with s_t and si_t, respectively; it is implicitly assumed that variables correspond to the data associated with A_k that contains a specific state estimate x_t.

The graph in Figure 2 corresponds to a scenario where the target enters a different area A_k at each time step. Subscripts of the variables in the graph explicitly indicate time slices corresponding to discrete time intervals. The graph reflects the fact that the physical constraints influence the state of the tracked object via the mobility m_i, that directly depends on the surface s_i and the vehicle type v. Stochastic relations between subsequent target states x_{i-1} and x_i are given by (3) while the observation models relating x_i and z_i are given by (7). The BN shown in Figure 1 defines the relations between the variables influencing mobility m_i. The model captures a factorized representation of the joint probability distribution over all variables that represent phenomena relevant for the inference about the dynamic process [12], [4]. By using an exact inference algorithm, variables s_i and m_i can be marginalized out for all time-slices resulting in the following factorized posterior of a target state sequence $P(x_{1:t}|z_{1:t}, \epsilon_{1:t})$:

$$\begin{aligned} &P(x_{1:t}|z_{1:t}, \epsilon_{1:t}) \\ &= \eta P(m_1 = true|\epsilon_1) P(x_1|m_1 = true) P(z_1|x_1) \\ &\quad \cdot\ P(m_2 = true|\epsilon_2) P(x_2|x_1) P(z_2|x_2) \\ &\quad \cdot\ \ldots \\ &\quad \cdot\ P(m_t = true|\epsilon_t) P(x_t|x_{t-1}) P(z_t|x_t) \end{aligned} \tag{15}$$

This equation reflects the fact that $P(x_t|x_{t-1}, m_t = true) = P(x_t|x_{t-1})$ and $P(x_t|x_{t-1}, m_t = false) = 0$ (see Section II-B), which means that in the marginalization over the states of m_t, all terms with $P(m_t = false|\epsilon_t)$ cancel out. In addition, we assume that each tracking process is carried out for a specific type of vehicle, which means that the state of variable v (the type of the target vehicle) in a particular sampling process is known, i.e. it is part of the evidence set ϵ_t. Therefore, variable v does not introduce additional dependencies that would require marginalization steps in the inference process.

Equation (15) is obtained also by expanding (13) implemented by the modified PF process, where CF $P(m_t = true|\epsilon_t)$ is used at each updating step. Each line in this equation corresponds to a new iteration based on an update step using the modified weight from (14) in the PF. The PF algorithm supports inference over variables x_i, z_i and m_i while the distribution $P(m_i = true|\epsilon_i)$ is obtained by running exact inference algorithm on the BN relating basic CD and the mobility m_i for each A_k. In other words, the modified PF algorithm in combination with BNs correctly considers the dependencies between physical phenomena relevant for the state of the target if it moves to a new area A_k at each time step.

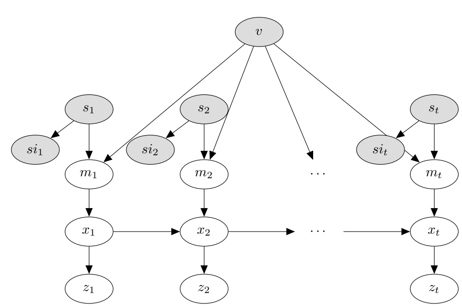

Fig. 2: A discrete causal model describing a dynamical stochastic process where the tracked object enters an area A_i with a different mobility at each timestep.

Clearly, the assumption that the tracked object enters a new area A_k at each timestep (i.e. sampling iteration) is not realistic in many cases as A_k can have significant dimensions. Then it

[2]By fusing CD we can derive knowledge about such physical phenomena

is likely that the target will remain in the same area A_k over a number of subsequent steps. Such a situation corresponds to dependencies between the target states and the physical conditions influencing the target's mobility that are different than the dependencies assumed by the factorization in (15). Namely, at each timestep within the same A_k, the transitions are influenced by the same mobility conditions. While we do not know whether the mobility in this area is $m_t = true$ or $m_t = false$, we know that if the actual $m_t = true$ at the entry into area A_k, then the object can move anywhere in this area; i.e. $m_t = m_{t+1} = \ldots = m_{t+n} = true$ as long as the object remains in A_k. The qualitative representation of such physical dependencies is captured by the directed graph in Figure 3, which represents a process spanning two areas A_q and A_r, respectively. The tracked object entered area A_q at time step 1 and transitioned to A_r at time step n. To simplify the discussion we rename $s_{q,1}$, $si_{q,1}$, $m_{q,1}$, $\epsilon_{q,1}$ to s_1, si_1, m_1, ϵ_1, respectively. Likewise, we rename $s_{r,1}$, $si_{r,1}$, $m_{r,1}$, $\epsilon_{r,1}$ to s_n, si_n, m_n, ϵ_n, respectively. Given this notation, we can write the factorization of the posterior $P(x_{1:t}|z_{1:t}, \epsilon_{1:t})$ corresponding to the graph in Figure 3 as follows:

$$
\begin{aligned}
&P(x_{1:t}|z_{1:t}, \epsilon_{1:t}) \\
&= \eta P(m_1 = true|\epsilon_1) P(x_1|m_1 = true) P(z_1|x_1) \\
&\quad \cdot\ P(x_2|x_1) P(z_2|x_2) \\
&\quad \cdot\ \ldots \\
&\quad \cdot\ P(x_{n-1}|x_{n-2}) P(z_{n-1}|x_{n-1}) \\
&\quad \cdot\ P(m_n = true|\epsilon_n) P(x_n|x_{n-1}) P(z_n|x_n) \\
&\quad \cdot\ P(x_{n+1}|x_n) P(z_{n+1}|x_{n+1}) \\
&\quad \cdot\ \ldots \\
&\quad \cdot\ P(x_t|x_{t-1}) P(z_t|x_t)
\end{aligned} \tag{16}
$$

Note that this equation contains only two CFs, $P(m_1 = true|\epsilon_1)$ and $P(m_n = true|\epsilon_n)$ corresponding to the time steps when the target entered new areas, at the first and the n^{th} step, respectively. Namely, between steps 1 and $n-1$ all states are conditioned on the same context (i.e. mobility in A_q) and the same is true for the iterations between steps n and t, when the target would move in area A_r. Consequently, for $2 \leq t < n$ and for $t > n$, $P(x_t|x_{t-1}, m_t = true) = P(x_t|x_{t-1})$.

The factorization in (16) provides a guidance for the use of CFs in a PF approximating this posterior. In this example, the correct factorization is approximated by the PF if CF is used for the modification of importance weights according to (14) only twice, at the first and the n^{th} step. In any other iteration step the importance weights should be identical to the observation model, i.e. $w_i = \eta P(z_i|x_t)$.

This can be generalized in the following CD updating rule for the PF process.

CD update rule: *the updating step uses importance weights as defined in (14) if a particle enters a new area A_k* [3] *and a new observation was made at that time step. Otherwise, $w_i = \eta P(z_i|x_t)$ is used as importance weight. If a particle enters a new area and no observation is obtained in the same time step, then the particle is resampled by using the CF-value.*

As it is demonstrated by the experimental results, a sound strategy for using importance weights defined in (14) is indispensable. The naive updating approach using importance weight defined in (14) at each updating step can result in severe estimation errors, as the processing does not correspond to the underlying dependencies between the physical phenomena. Namely, the usage of CF at each time step without moving to a different A_k corresponds to reusing the same knowledge (i.e. data) about the physical constraints multiple times. Thus, naive use of CD can have similar effects as data-incest.

It should be emphasized that the presented causal graphical models relating different environmental phenomena with the target states are not directly used in the estimation process. They are used merely for the analysis of dependencies and the derivation of the updating rules for uncertain constraints. However, the presented combination of PF and BNs implements inference that is equivalent to approximate reasoning with such graphical models. A PF implements approximate probabilistic spatio-temporal inference about target states while BNs efficiently exploit the knowledge of the correlations between heterogeneous CD for the exact estimation of uncertain constraints.

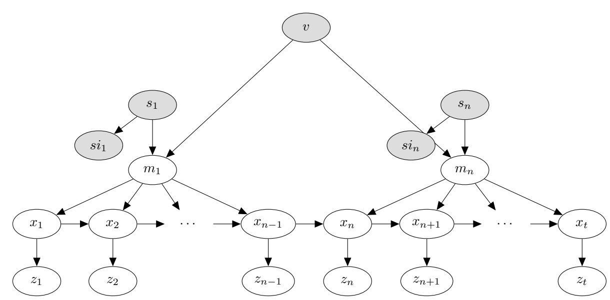

Fig. 3: A model of a dynamical stochastic process where the tracked object remains in the same area with a constant mobility over a number of timesteps.

III. Experiments

The performance of the tracking algorithm with and without context data was investigated in a series of experiments making use of simulated targets and sensors. In all experiments, a synthetic map was used as a basis for the derivation of a Context Grid to a PF process. In each experiment the target

[3] This is equivalent to a particle transitioning between two cells in the CG, each representing a different A_k, and consequently associated with a different CF.

starts around the coordinates (5,75), follows a path through a terrain with different levels of certainty about the mobility and ends in the north-western part of the map, around coordinates (2,12). The results are shown in two plots for each experiment. The first plot shows the CF values (1 - white, 0.5 - light-gray, 0 - dark-gray), the target's true track, all observations and the estimated tracks. In the second plot, the root-mean-square error (RMSE) of the experiment is shown, which is calculated as: $RMSE = \sqrt{\frac{\sum_i (x_t^i - x_t^*)^2}{N}}$, where x_t^* is the true location of the target and x_t^i the location of particle i.

In the first two experiments, there is a radar-like sensor that can observe the whole area. The timespan of the experiment is 3000 timesteps and the sensors provide a measurement every 50 timesteps. The measurement model assumes a Gaussian function with the noise defined by $\sigma = 10$.

A. CD at every timestep

In the first experiment (Figure 4) we compare a standard PF tracking algorithm with a tracking algorithm applying CD in each updating step. This was a naive updating strategy where at each update the PF process used the corresponding CF. Consequently, the inference process violated the dependencies in the underlying simulated physical world, as the target moved through significant areas associated with constant, smaller likelihood of mobility (i.e. constant CF).

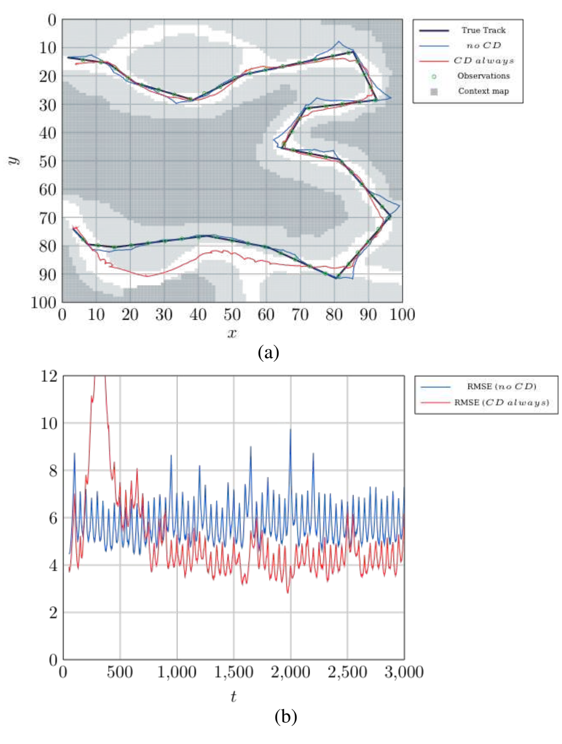

Fig. 4: The results of the experiment showing the impact of the CD update strategy where the CF is used at every timestep.

The resulting tracks show clearly that when using CD, the algorithm is able to compensate for overshoots that occur if no CD is used and the target makes sharp turns that are not adequately captured by the transition models. This also results in an overall lower RMSE throughout the experiment if CD is used. However, a naive approach to updating resulted in severe estimation errors in the south-western part of the map, where the target moved through an area with lower CF values (this corresponds to an area the mobility constraints are uncertain at this moment in time). The derived context data in form of CF values was used too often and biased the PF. In this case, the PF algorithm implemented computation based on posterior factorization (15) while the underlying dependencies in the simulated physical world corresponded to a factorization similar to (16). The estimation errors due to this violation are also very visible in the RMSE, between steps 0 and 700. This experiment illustrates the impact of the modeling violation introduced through naive use of CF discussed in section II-C.

B. CD at mobility change

In the next experiment, we compare the standard PF tracking algorithm with the improved updating rules correctly considering the underlying dependencies (see Section II-C). A particle's CF was only used for the modification of the importance weight if the particle moved into a new area in which the CF value is different than at the previous location. In contrast to the naive updating with CD shown in the previous experiment, tracking in areas of lower mobility likelihood is significantly more accurate (see the South-Western part of the map in (Figure 5). For the remainder of the target's route, a similar performance as in the first experiment is observed. When using available CD, the target's estimated location has a lower RMSE in general than the algorithm that is not using CD at all.

C. Partial sensor coverage

In this experiment, we show the impact of using CD in a tracking algorithm when the sensors provide only partial coverage of the area. In this experiment, there are three sensors, each providing a measurement every 50 timesteps. The range of each sensor is shown in Figure 6a as yellow dotted circles.

Compared to tracking without CD, we observe a significant increase in performance when a target is outside of the range of available sensors, visible in the RMSE curves. The estimated track is much closer to the true track than in the case when no CD is used. Especially interesting parts of this experiment are those where the target changes direction outside of the area in range of one of the sensors. Similarly, the uncertainty about the target's locations are smaller. The resulting particle cloud is more focused (see Figure 7), as impossible locations are ruled out from the set of hypotheses through resampling. The use of context data compensates the lack of observations and the lack of knowledge about the dynamic properties of the target.

IV. Discussion

We present a novel approach to incorporating Context Data in a PF-based tracking algorithm. The overall solution imple-

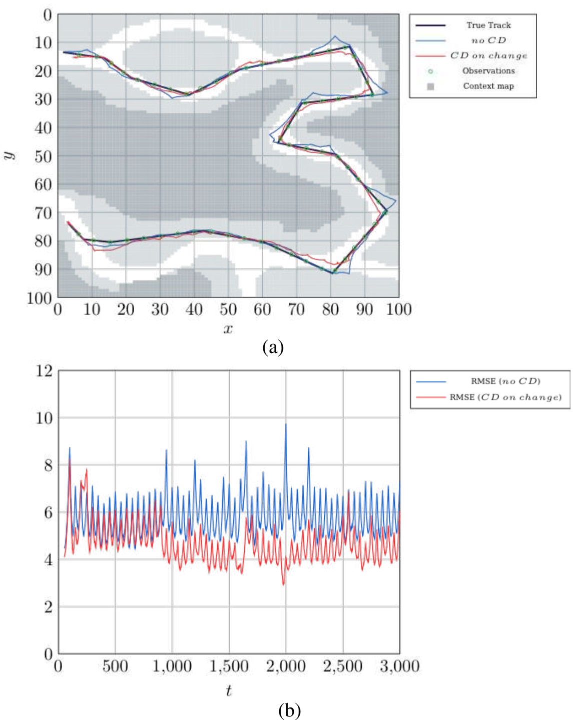

Fig. 5: The results of the experiment showing the impact of the CD update strategy where the CF is used when the particle enters a new cell.

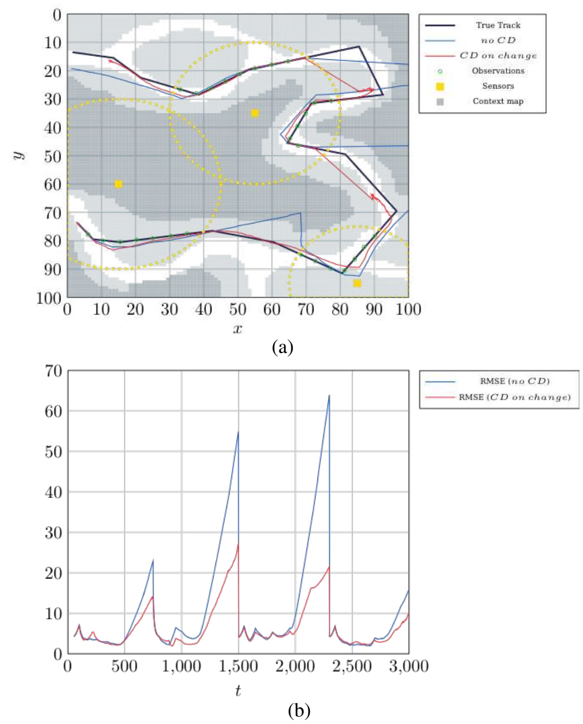

Fig. 6: The results of the experiment showing the impact of using CD with only partial sensor coverage of the area in which the target is moving in.

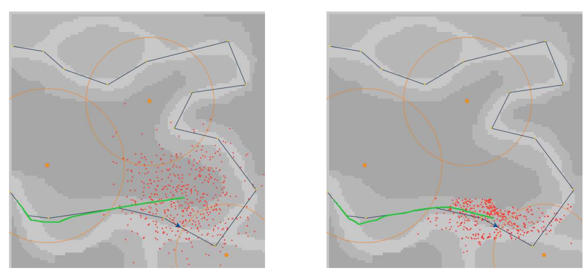

(a) No CD (b) CD on change

Fig. 7: Two snapshots from the partial sensor coverage experiment showing the spread of the particles (in red), the true track (dark blue line), the estimated track (in green) and the sensors and their range (in yellow).

ments inference about complex dynamical systems (see, for example, Figure 3) by seamlessly combining two related types of theoretically sound approaches to exact and approximate Bayesian inference. By using the theory of Particle Filters and Bayesian Networks, we show that uncertain CD can be introduced into a PF process through a straightforward modification of importance weights used in the updating step.

The resulting approach allows incorporation of different types of CD without requiring any special adaptation of the PF algorithm. The CD is integrated into the PF with the help of Bayesian Networks that systematically capture causal relations between heterogeneous types of CD and the possible target states. Bayesian Networks describe how the environment affects the mobility of the target in some area of interest. Bayesian Networks compute the probability distribution that a certain type of targets is mobile at different locations/areas. This probability distribution is called the Corrective Factor, representing uncertain constraints. The CF values are organized and maintained in a Context Grid, a data structure that associates different locations with specific CFs. Different CFs are updated with BNs dedicated to different locations, as soon as new CD becomes available for a specific location (e.g. an intelligence report about a roadblock) and the results are instantaneously made available to the PF via updated CFs.

By using the theory of graphical causal models, we show that naive use of uncertain CD in the updating process in PF can result in significant estimation errors. With the help of graphical models we make the factorization of the posterior probability distribution over the target's states explicit. We show that the tracking process using uncertain CD in a naive way corresponds to different factorizations of the estimated posterior than the factorization associated with the physical processes. Moreover, by using the analysis of the factorized posterior, we derive simple updating rules, that alleviate the problems of the naive updating method while the CD is fully considered in the reduction of uncertainties. Note that the graphical models relating different environmental phenomena with the target states are not used in the tracking process itself. They are merely used to derive the updating strategy, to facilitate the analysis of the dependencies between the phenomena considered by the presented solution.

The experimental results clearly support the theoretical discussion on the rules for uncertain constraints derived from CD. The proposed updating rules improve the tracking accuracy significantly.

The theoretically predicted and experimentally confirmed properties of the proposed method have important practical implications. Namely, a robust approach to using CD enables tracking and the estimation of whereabouts in a relevant class of problems characterized through partial observability[4] and low updating frequencies. Examples of such applications are tracking in urban or littoral environments, wildlife protection etc. The final experiment illustrates the effects of using CD in such environments.

Another important property of the presented method is that the computation of CFs in the Context Grid is carried out in a system that is separated from the PF processes. The BNs used to compute CFs can be updated any time, without requiring any changes to the PF algorithm, as the fusion is reduced to the computation of CFs, a uniform representation of the knowledge about the mobility of a target.

The main limitation of the current approach is discretization of the space. Currently, the CG supports uniform cells. While this already proved to be effective, more advanced representations for the CG are being investigated, such as hierarchical grid representations and vector-based approaches that avoid discretization altogether. In this manner the CG will be more scalable and will allow a better resolution. Moreover, in the future work we will look into combining CD with more advanced PF approaches. A special focus will be on extending the proposed method to blending fusion of CD with PF approaches that also support target classification and estimation of motion modes, such as for example [11]. Furthermore, in the future research, we will investigate the extension of the CD fusion to concepts that go beyond mere mobility. For example, we will incorporate models of intentions, the most likely behaviour, etc.

ACKNOWLEDGEMENTS

The work in this paper was partly supported by WBSO+ program (project number 114340) and FP7 Athena (project number: 313220). The authors would like to thank Hans Driessen and Martin Podt from Thales Nederland B.V. for their guidance, valuable ideas and feedback on the work described in this paper.

REFERENCES

[1] A. Doucet, "On sequential simulation-based methods for bayesian filtering," Tech. Rep., 1998.

[2] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005.

[3] F. Gustafsson, F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P.-J. Nordlund, "Particle Filters for Positioning, Navigation, and Tracking," *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 425–437, Feb. 2002.

[4] F. V. Jensen and T. D. Nielsen, *Bayesian Networks and Decision Graphs*, 2nd ed., ser. Statistics for Engineering and Information Science. Springer Publishing Company, Incorporated, 2007.

[5] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.

[6] M. Podt, M. Bootsveld, Y. Boers, and P. F., "Exploiting Imprecise Constraints in Particle Filtering Based Target Tracking," in *Proceedings of the 17th International Conference on Information Fusion, Volume I*, 2014.

[7] F. Papi, M. Podt, Y. Boers, G. Battistello, and M. Ulmke, "On Constraints Exploitation for Particle Filtering Based Target Tracking," in *Proceedings of the 15th International Conference on Information Fusion*, Jul 2012, pp. 455–462.

[8] M. Ulmke and W. Koch, "Road-map assisted ground moving target tracking," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 42, no. 4, pp. 1264–1274, October 2006.

[9] Y. Cheng and T. Singh, "Efficient Particle Filtering for Road-Constrained Target Tracking," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 43, no. 4, pp. 1454–1469, October 2007.

[10] C. S. Agate and K. J. Sullivan, "Road-Constrained Target Tracking and Identification using a Particle Filter," vol. 5204, 2003, pp. 532–543.

[11] J. Murphy and S. J. Godsill, "Road-assisted multiple target tracking in clutter," in *17th International Conference on Information Fusion, FUSION 2014, Salamanca, Spain, July 7-10, 2014*, 2014, pp. 1–8.

[12] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1988.

[4]Sensors do not cover all parts of the target's path.