

Sequential Lifted Bayesian Filtering in Multiset Rewriting Systems

Max Schröder, Stefan Lüdtkke, Sebastian Bader, Frank Krüger and Thomas Kirste

Mobile Multimedia Information Systems Group, Institute of Computer Science
University of Rostock, 18051 Rostock, Germany

{max.schroeder, stefan.luedtke2, sebastian.bader, frank.krueger2, thomas.kirste}@uni-rostock.de

Abstract

Bayesian Filtering for plan and activity recognition is challenging for scenarios that contain many observation equivalent entities (i. e. entities that produce the same observations). This is due to the combinatorial explosion in the number of hypotheses that need to be tracked. However, this class of problems exhibits a certain symmetry that can be exploited for state space representation and inference. We analyze current state of the art methods and find that none of them completely fits the requirements arising in this problem class. We sketch a novel inference algorithm that provides a solution by incorporating concepts from Lifted Inference algorithms, Probabilistic Multiset Rewriting Systems, and Computational State Space Models. Two experiments confirm that this novel algorithm has the potential to perform efficient probabilistic inference on this problem class.

1 Introduction

Bayesian Filtering algorithms are often applied in the domain of plan and activity recognition (Sukthankar and Sycara, 2008; Huang, Guestring, and Guibas, 2009; Ramírez and Geffner, 2011). These algorithms typically suffer from an exponential state space growth when reasoning about multiple entities that cannot be distinguished from observations. This is for instance the case in multiple person tracking by use of anonymous sensors (Fox et al., 2003) or assisted manufacturing, where the identity of different tools and parts is uncertain. Currently, no algorithm exists that exploits these symmetrical parts of the state space and that is applicable to all instances of the underlying problem class.

Consider the following multiple person tracking example:

Example 1. *Two persons “A” and “B” are moving within an environment observed by anonymous presence sensors. Both are at the same room. When the presence sensor indicates that one person left this room, the following two hypotheses have to be tracked: (1) person “A” left the room and person “B” stayed in the room, and (2) the other way round – “B” left, “A” stayed.*

This example is one instance of the *data association problem* sometimes referred to as track confusion problem (Fox et al., 2003; Wilson and Atkeson, 2005; Huang, Guestring,

and Guibas, 2009). The data association problem occurs if two (or more) entities (e. g. persons) are recognized by the same anonymous sensor and thus cannot be distinguished anymore. In other words, the person’s identity associated with the movement track is lost. As a consequence, all possible resulting hypotheses have to be tracked. This leads to a combinatorial explosion with respect to the number of hypotheses that have to be considered.

Tracking single persons or very small groups is feasible using state of the art techniques. However, tracking more than 10 persons is not computationally feasible without abstraction. With respect to modeling chemical reaction, Multiset Rewriting Systems (MRS) allow to reason over thousands of atoms by abstracting from individual atoms. However, applications of activity and plan recognition require the identification of selected individuals. Thus, MRS are not immediately applicable.

The contribution of this paper is threefold: (1) By providing two example scenarios, we analyze the underlying characteristics and specify a general problem class. (2) We analyze state of the art methods using six evaluation criteria derived from the class. (3) We sketch a novel inference algorithm that combines ideas from different state of the art methods and present first evaluation results based on our prototypical implementation.

Next, we further analyze the data association problem and derive evaluation criteria for state of the art methods (Sec. 2). In Sec. 3 we analyze state of the art methods with respect to the evaluation criteria. As none of the methods is capable of computing all instances of the problem class, we sketch a novel combination of these methods in Sec. 4. An evaluation of our mechanism using two scenarios is described in Sec. 5 and our conclusion and future work in Sec. 6.

2 Problem Specification

The data association problem can be found in many domains, e. g. monitoring and assistance at manufacturing (Aehnelt and Bader, 2015), assistance of elderly (Müller and Hein, 2015), or monitoring of buildings to support disaster management (Krüger et al., 2014a). An intuition for the characteristics of the underlying problem class is given by the following two scenarios.

Scenario 1 (Warehouse). *Ten forklifts are used in a warehouse, consisting of three main storage rooms, a service sta-*




Figure 1: Overview of the warehouse scenario. Gray areas indicate where multiple forklifts are observation equivalent. White areas are equipped with identifying sensors.

tion where the forklifts get refueled, and a room where all forklifts are parked at night. The room layout is depicted in Fig. 1. Forklifts can move between two connected locations at each time step. Each room is equipped with an anonymous sensor that delivers binary information about whether at least one forklift is present at the room. When a forklift gets refueled, it is identified. Our goal is to track the position of the forklifts over time to optimize storage management.

Another scenario that does not only track the position of several entities, but also maintains additional information for each entity has been adapted from Bui (2003):

Scenario 2 (Office). Three persons act in an office environment which is observed by anonymous sensors. The office contains a coffee machine, a coffee jar and a water tap as well as a printer and a paper stack. The goal of the agents is to print a number of documents and to get a coffee. To do so, the coffee machine needs to be filled with water and ground coffee, and paper must be provided for the printer. Whenever a person is operating the printer, she has to authenticate herself and thus can be identified. Our goal is to estimate the items a specific person is carrying over time in order to provide assisting operations such as opening doors.

2.1 General Problem Class

The previously described scenarios entail the characteristics of the following general problem class:

- C1. Multiple entities (possibly of the same kind) are interacting simultaneously to reach their individual goals. Entities, here, include passive objects as well as active ones like persons or robots. An interaction can be any type of creation, manipulation or deletion of some entities. Entities have an identity that might be of interest to answer the application specific question.
- C2. The scenario is partially observed given a set of (noisy) sensors. In the case of anonymous sensors, these do not allow to distinguish between multiple entities as they potentially produce the same sequence of sensor data, i. e. they are *observation equivalent*. However, there may also occasionally appear sensor measurements that reveal the identity of some entity.
- C3. Given a situation, future trajectories, i. e. sequences of entity interactions towards a future situation, need to be predicted in order to recognize the pursued goal or to provide assistance.

In the warehouse scenario, the entities are forklifts which move simultaneously through the environment and that act

autonomously (C1). Their movement is partially observed by binary sensors indicating that one or more forklifts are present at a specific location. These sensors do not allow to distinguish individual forklifts. However, identification is provided at the refuel station (C2). This scenario does not include the need for predicting future action courses.

Considering Scenario 2, the entities are persons moving in the environment as well as devices and consumables. The entities possibly interact with each other (e. g. a person holding paper replenishes the printer) or act independently (e. g. by moving to another place). This interaction is partially observed by anonymous sensors at all locations and identifying sensors at the printer. In contrast to anonymous sensors, identifying sensors enable the identification of individual entities. This allows to prune impossible hypotheses (C2). Predicting future courses of the persons' actions would allow for a situation aware assistance, e. g. by opening doors if agents are carrying coffee (C3) and printed documents.

The general problem class includes also other instances that are encoding only some characteristics such as the predator-prey population scenario (Gordon et al., 2014) or the reaction trajectory of chemical elements (Barbuti et al., 2011).

2.2 Evaluation Criteria

In Sec. 3 we review state of the art approaches which are applicable to these problems. To evaluate the capabilities of these algorithms, we use the following evaluation criteria arising from C1–C3:

- E1. Can an observation sequence be processed sequentially? *An observation sequence can be indefinitely long and not be known completely in advance, therefore the algorithm should not depend on the complete sequence.*
- E2. Can the structure of the state space change over time? *As interactions of entities may change the set of entities that exist (C1), the resulting state space must change accordingly.*
- E3. Is the approach capable of handling nondeterminism resulting from actions and observations? *Given a prior distribution, the sensor measurements need to be taken into account for updating this belief state (C2).*
- E4. Can the approach predict future states, based on the current belief state? *This is necessary e. g. for goal recognition or to provide assistance.*
- E5. Are state transitions able to represent multiple parallel actions of individual entities? *Entities are interacting simultaneously (C1), i. e. multiple entities may perform different actions at the same state.*
- E6. Can the approach represent equivalent (e. g. observation equivalent) aspects of a state compactly when possible and ground this representation when necessary? *If parts of the state space cannot be distinguished (C2), the algorithm should combine these to increase efficiency. If an identifying observation is made (C2) or if we are interested in the identity of a specific entity (C1), this compact representation needs to be grounded to some extent.*

3 Related Work and Preliminaries

Different research domains target different aspects of the problem class. Here, we give an overview of techniques related to the problem domain and analyze their strengths and weaknesses with respect to this class.

3.1 Identity Management (IM)

Identity Management (Kondor, Howard, and Jebara, 2007; Huang, Guestring, and Guibas, 2009) is concerned with multiple-object tracking in case of anonymous observations such as radar measurements. Objective is to associate observed tracks with concrete entities, that is maintaining a distribution over permutations of entities. Identity management is concerned with finding a compact representation of these distributions by approximation employing the first Fourier coefficients of the group of permutations.

This approach allows to perform Bayesian Filtering (BF) (E1, E3, and E4) to solve the data association problem (E6). It is very efficient and able to handle a large number of entities, e. g. up to 100 in Huang et al. (2009).

This ability comes at the cost of a limited expressiveness, as only problems of observing tracks of concrete entities can be modeled. The approach can not be used in scenarios where additional information about these objects have to be maintained as for instance in the office scenario (Sec. 2). IM offers no mechanism to model dynamically changing states (E2) and parallel actions (E5).

3.2 Computational State Space Models (CSSM)

CSSMs allow the knowledge-based construction of state spaces for BF. They are for instance used for human behavior and goal recognition (Baker, Saxe, and Tenenbaum, 2009; Ramírez and Geffner, 2011). The transition model is described by a computable function by means of preconditions and effects. This allows the compact representation of potentially infinite state spaces by avoiding explicit state enumeration, as in Hidden Markov Models.

CSSMs allow to handle large, even infinite, state spaces (Krüger et al., 2014b). By employing standard methods for BF (e. g. Particle Filters), online inference with identifying observations can be performed (E1, E3, and E4). A dynamically changing state space (E2) can be modeled by state predicates with infinite domain. CSSMs allow the parallel application of actions of multiple entities (E5).

CSSMs perform inference in grounded state spaces (i. e. concrete values are assigned to all state variables). With respect to observation equivalent states (E6), CSSM's do not allow any abstraction but rather require to track all hypotheses. This leads to a combinatorial explosion of the number of states that have to be tracked.

3.3 Lifted Inference (LI)

Lifted Inference is concerned with efficient inference in relational graphical models (RGMs). An RGM is a compact representation of a graphical model that has certain symmetries: random variables (RVs) that have similar (conditional) distributions are represented as a single *parametric* random variable (par-RV). Different exact and approximate methods

exist (Braz, Amir, and Roth, 2005; Kersting, Ahmadi, and Natarajan, 2009; Gogate and Domingos, 2016).

The framework of LI allows the abstract representation of observation equivalent entities via par-RVs (E6). *Split* operations allow to incorporate evidence about particular RVs, i. e. dividing a par-RV into the RV that we have evidence about and the remaining RVs. It is possible to perform probabilistic inference in models with very large domain sizes (Poole, 2003; Van Den Broeck et al., 2011).

LI is typically applied to Bayesian Networks, which require the complete observation sequence to be present at the start time. Thus, they do not allow for *online inference* — the recursive computation of the belief state for each time step, without knowing the complete observation sequence in advance (E1, E3, E4, and E5). Furthermore, inference is not guaranteed to be linear in the number of time steps. The algorithms require a fixed network structure, therefore they can not model a dynamically changing state space (E2).

3.4 Lifted Bayesian Filtering (LBF)

Logical particle filters represent multiple states within one particle. This is done by leaving some state predicates undefined (Zettlemoyer, Pasula, and Pack Kaelbling, 2008) or by defining a distribution for some of the predicates, instead of a concrete value. This implicitly defines a distribution over the states represented by the particle (Nitti, De Laet, and De Raedt, 2013). It is related to the Rao-Blackwellized particle filter (Doucet et al., 2000), as parts of the state are represented by samples and others by sufficient statistics. Multiple approaches combine LI and BF. For example, the interface algorithm (Murphy, 2002) has been extended by using a LI algorithm at each time step (Kersting, Ahmadi, and Natarajan, 2009; Vlasselaer et al., 2016).

By using the framework of BF, Logical Particle Filters satisfy E1, E3 and E4. Similar states are grouped by providing distributions over uncertain predicates to implement state space abstraction (E6).

Although similar states are grouped by providing distributions over uncertain predicates, logical particle filters do not dynamically group similar variables (E6): Multiple (observation equivalent) entities are always represented by different variables for each entity, that potentially have identical distributions. Furthermore, none of the models can handle a dynamically changing state structure (E2) or has been designed to support parallel actions for state transitions (E5).

3.5 Multiset Rewriting Systems (MRS)

Multiset Rewriting Systems are an established formalism for modeling systems with many equal objects. They are used to model chemical reactions (Berry and Boudol, 1990) or cell interactions (Bistarelli et al., 2003). The state is described as a multiset of entities, where each entity is an instance of one of finitely many *species*. The reactions of entities are modeled as *multiset rewriting rules* (E4) that have preconditions (entities are consumed by the reaction) and effects (entities are created by the reaction). Typically, rewriting rules are executed in parallel (E5), that is, the maximal set of parallel applicable rewriting rules (*maximal parallel step*) is deter-

Method	E1	E2	E3	E4	E5	E6
IM	●	○	●	●	○	●
CSSM	●	●	●	●	●	○
LI	○	○	●	○	○	●
LBF	●	○	●	●	○	◐
MRS	●	◐	○	●	●	◐

Table 1: Evaluation of inference methods based on criteria from Sec. 2. ●: satisfies criterium; ○: does not satisfy criterium; ◐: satisfies criterium partially

mined. There may be more than one maximal parallel step, leading to multiple possible execution paths.

Probabilistic MRS (Barbuti et al., 2011) are concerned with determining the probability of different execution paths. This is done by assigning *rates* to each rewriting rule, which leads to a probability of actions and thus, for each maximal parallel step. A set of rewriting rules and an initial multiset defines a probabilistic state space.

Probabilistic MRS allow an abstract representation of a state space via the multiset representation (E6). Furthermore, MRSs provide a semantics for parallel action execution (E5) and a process for drawing samples (E4) of the sequential process (E1). The state space is latently infinite, as an arbitrary number of entities can exist (E2).

BF algorithms for MRSs that incorporate observations (E3) are not available. MRSs do not maintain identities for entities (E6). Furthermore, despite the latently infinite state space, the structure of the state space cannot be extended arbitrarily, as the existing species are defined beforehand (E2).

3.6 Summary

As evident from the discussion above, none of the methods satisfies all of the requirements E1–E6 (see Tab. 1). To address the problems described above, a combination of CSSMs, LI and Probabilistic MRSs seems reasonable.

4 Lifted Marginal Filtering

Our approach – called *Lifted Marginal Filtering* (LiMa) – combines different ideas from the approaches described above. The general concept is inspired by Lifted Inference algorithms, that is, finding a parametric representation for equivalent aspects of the state instead of explicitly enumerating all possibilities. It is based on defining an update operation to Probabilistic MRSs, or, from a different point of view, introduces a multiset-based state description to CSSMs. Furthermore, the approach defines *structured* entities, as opposed to atomic entities in MRSs, to be able to incorporate additional information about entities such as the identity.

4.1 Abstract State Representation

Like in MRS, similar entities are grouped to overcome the combinatorial explosion. However, in order to group entities that are similar but not completely equal, the plain multiset representation known from MRS has to be extended to cover

the uncertainty in the entity properties. To this end, the representation of entities is separated into the internal structure of the entity and the actual values. Entities with the equal structure are maintained in a multiset (in the following called *state formula*). The *context* describes distributions over values within these structures. Thus, we get a parametrized description of property value maps. This concept is similar to that of lifting RV into par-RV in LI. Next, we will refine the concepts of state formulae and contexts.

A state formula is a multiset of entities. An entity, is a map of slots to distribution labels. Thus, labels are used to map properties to possible property values. Note that slots of an entity are assumed to be independent from each other.

Example 2. *An entity modeling the forklift with two properties location loc and unique identifier ID pointing onto value labels $LStor1$ respectively LID is denoted as:*

$$\langle loc: LStor1, ID: LID \rangle$$

Example 3. *A state formula with 9 entities as in Ex. 2 and another entity in storage room 2 can be denoted as:*

$$\llbracket 9 \langle loc: LStor1, ID: LID \rangle, 1 \langle loc: LStor2, ID: LID \rangle \rrbracket$$

The state formula is representing the structure of the entities only. It is connected to the context via distribution labels. A context is a map from distribution labels to distribution representations. In order to employ the same distribution in different entities, the same distribution label has to be referred to. Note that this allows us to represent slot values of multiple entities that depend on each other. This is for instance done in Ex. 3, where the identifier (ID) of all entities is distributed according to LID (thus making sure that no two entities have the same name).

Example 4. *Let $\mathcal{U}(fl1, fl2, \dots, fl10)$ be the representation of the distribution that represents an urn without replacement and contains ten identifiers: $fl1, fl2, \dots, fl10$. Furthermore let $\delta(storage1)$ represent the distribution which is non-zero for $storage1$ only. A context that maps the label LID to the urn of identifiers and $LStor1$ to $\delta(storage1)$ is denoted as follows:*

$$\{ LID \mapsto \mathcal{U}(fl1, fl2, \dots, fl10), LStor1 \mapsto \delta(storage1) \}$$

A pair of state formula and valid context forms the basis for a compact abstract state representation (E6 and E2), which we call *lifted state*. A context is valid with respect to a state formula, if and only if (1) each distribution that is referred to within the state formula is defined in the context and (2) the distribution allows to sample at least as many values as the sum of the cardinalities of the entities. A state formula represents a set of grounded states.

Example 5. *The state formula from Example 3 can be connected to the following context to represent all situations (i. e. all permutations of ID s) in that one forklift is at storage room 2 and the other forklifts are in storage room 1:*

$$\llbracket 9 \langle loc: LStor1, ID: LID \rangle, 1 \langle loc: LStor2, ID: LID \rangle \rrbracket \\ \{ LID \mapsto \mathcal{U}(fl1, fl2, \dots, fl10), LStor1 \mapsto \delta(storage1), \\ LStor2 \mapsto \delta(storage2) \}$$

Obviously, distributions are not restricted to discrete domains. The same abstraction can be performed in continuous domains, where lifted states represent infinitely many grounded states.

4.2 Transition Model

Adapting the idea of computational action languages from CSSMs, our algorithm uses *action schemas* to describe transitions between states (E4 and E2). Furthermore, as multiple actions can be executed in parallel (E5), the transitions in the state space actually involve possibly multiple action schemas that are combined to compound actions, similar to those in MRS.

An action schema in LiMa is a triple of (1) the *action name* which is necessary for plan recognition, (2) a *tuple of preconditions* that is used to filter entities that satisfy these conditions, and (3) an *effect function* mapping entities to a resulting set of entities.

Preconditions are modeled by constraints over slots that must hold in order to apply the action. If an action is applied to a set of entities that satisfy the preconditions, the entities within this set are consumed and replaced by the effect of the action.

Example 6. *The following action schema captures the lifting of a heavy weight box which can only be conducted by fl1. The precondition makes sure that the value of the slot ID is fl1. In the effect, a new slot load is created and the value set to heavy:*

```
(fl1_lift_heavy',                               /* name */
  ({ID ↦ (λ v ↦ v ≡ fl1)}),                       /* precondition */
  (E) ↦ [ [1E{load ↦ heavy} ] ]                   /* effect */
```

An action is applicable in a given lifted state if and only if the state contains entities which satisfy the actions preconditions, (i. e., all constraints need to be satisfied by the entity, respectively, the lifted state). However, if a lifted state contains an entity with uncertain property values, e. g. an urn, the preconditions are indeterminate. Similar to LI, the lifted state is split to decide on the preconditions. A split is the partition of a lifted state into two lifted states: (1) a lifted state that encodes all situations satisfying the preconditions (action can be executed), and (2) a lifted state encoding all remaining situations (action cannot be executed). As part of a split, both the state formula and the context of the lifted state have to be adjusted. Thus, splits increase the number of hypotheses. However, to keep the number of hypotheses small, indistinguishable states can be merged. Merging is the inverse operation of a split. Note that both concepts are known from the domain of LI.

Example 7. *Considering the lifted state in Example 5 and the action schema in Example 6, we need fl1 to satisfy the precondition. The state has to be split into two lifted states: (1) fl1 being at storage room 1 and (2) fl1 being at storage room 2. In order to perform this split, the urn LID is divided into two distributions. The first being LID', which represents LID without fl1 and the second LID'' which represents fl1. The resulting states are as follows:*

```
(1) [ [8⟨loc: LStor1, ID: LID'⟩, 1⟨loc: LStor1, ID: LID''⟩,
      1⟨loc: LStor2, ID: LID'⟩ ]
  { LID' ↦ U(fl2, ..., fl10), LID'' ↦ δ(fl1),
    LStor1 ↦ δ(storage1), LStor2 ↦ δ(storage2) }
```

```
(2) [ [9⟨loc: LStor1, ID: LID'⟩, 1⟨loc: LStor2, ID: LID''⟩ ]
  { LID' ↦ U(fl2, ..., fl10), LID'' ↦ δ(fl1),
    LStor1 ↦ δ(storage1), LStor2 ↦ δ(storage2) }
```

Action application is similar as in MRSs. Action schemas can be applied to different subsets of entities within the same lifted state. A set of action schemas can be applied in parallel. Similarly, compound actions are used to describe sets of actions that are executed in parallel. Here, we use the concept of maximal parallel compound actions (Barbuti et al., 2011), i. e. no further action can be applied in a lifted state.

4.3 Inference

A lifted state models multiple situations. However, as there are several sources of uncertainty (observations, non-deterministic actions, ...), we will consider not just a single lifted state, but a probability distribution over lifted states, which we call *lifted belief state*.

Example 8. *As an example, the following lifted belief state describes two hypotheses. One reflects Example 5 with the probability of 0.75 and the other encodes that all forklifts are in storage room 1 with probability 0.25.*

```
(3) 0.75 × [ [9⟨loc: LStor1, ID: LID⟩, 1⟨loc: LStor2, ID: LID⟩ ]
  { LID ↦ U(fl1, fl2, ..., fl10), LStor1 ↦ δ(storage1),
    LStor2 ↦ δ(storage2) }
(4) 0.25 × [ [10⟨loc: LStor1, ID: LID⟩ ]
  { LID ↦ U(fl1, fl2, ..., fl10), LStor1 ↦ δ(storage1) }
```

Starting with an initial lifted belief state, in the framework of BF (E1), we need to perform consecutively (1) updating the belief state according to the observations (E3), (2) answering of application specific questions, and (3) predicting the next lifted belief state by applying the compound actions. Note that all updates and manipulations of the belief states are performed in the lifted domain without grounding.

While applying the observation update, splits are required if identifying observations are made. Application specific questions can be answered on every updated posterior distribution over lifted states.

Example 9. *We are interested in the location of fl1 for the lifted belief state in Example 8. To determine the location of fl1 we need to split both lifted states:*

1. *Splitting state (3) leads to two lifted states (1) and (2) from Example 7. Whereas the first has 9 different instantiations, the second has a single one, which naturally lead to the probabilities of 0.9 and 0.1, respectively. Weighted by the probability of the original lifted state (3) this results in 0.675 and 0.075.*
2. *From (4) we know that fl1 is at location storage room 1 with probability 0.25.*

Summing the probabilities, leads to a probability of 0.675 + 0.25 = 0.925 that fl1 is at storage room 1 and 0.075 that it is at storage room 2.

5 Experiments

In this section, we evaluate the Lifted Marginal Filtering algorithm on the two scenarios (Sec. 2). Scenario 1 is a location tracking task, Scenario 2 has a more structured state




Figure 2: Results for the warehouse scenario. Top: Estimation of the location of *f1*, black diamonds: true location. Bottom: Number of hypotheses during inference. Notice that the number of hypotheses gets small when all forklifts cluster at few locations (time 10, 20, 30). Identifying observations are made at the refill location. Grounded inference was infeasible after timestep 2.




Figure 3: Results for the office scenario. Top: Estimation of the object person 1 is holding, black diamonds: true object. Bottom: Number of hypotheses during inference. An identifying observation has been made at timestep 10.

space with additional entity information. Both scenarios are used to compare our approach with a grounded inference approach. For our experiments we randomly sampled observation sequences from a HMM. Performance is assessed by counting the number of hypotheses during inference. Furthermore, application specific questions are answered.

5.1 Warehouse Scenario

As can be seen from upper part of Fig. 2, the location of forklift *f1* can be determined exactly when either an identifying observation was made (time 4, 12, 16, and 28), or all forklifts are at the same location (time 10, 20, and 30). Knowledge about such situations is incorporated in preceding timesteps, as for instance one timestep after identification, *f1* cannot be at storage room 1, because it takes two time steps to get there (cf. time 4 and 5 in Fig. 2).

Fig. 2 (lower part) shows the number of hypotheses tracked during inference. At most, 630 hypotheses had to be tracked simultaneously in our algorithm. The corresponding grounded inference algorithm was infeasible after timestep 2 because of the large number of hypotheses. Specifically, the algorithm exceeded the available 128 GB memory. The number of hypotheses decreases one timestep before all forklifts meet at the same location (time 10, 20, 30), and one timestep after this event. The reason is that there are only two possible locations for the forklifts (parking room and

storage room 2). This situation can be represented by very few hypotheses in our algorithm.

This experiment shows that LiMa can efficiently perform BF in situations with mixed anonymous and identifying observations. Grounded inference, in contrast, is infeasible. It is worth mentioning that LiMa performs *exact* inference, as opposed to IM approaches that use approximation.

5.2 Office Scenario

The office scenario requires to model more complex situations that go beyond location tracking tasks, since different entity properties (what a person holds) and different types of entities (persons and printjobs) are involved. Figure 3 depicts the number of hypotheses during inference and as application specific questions, the object that person 1 holds. The maximum number of hypotheses of our approach is 16, compared to 46 hypotheses tracked by the grounded approach.

We restricted the number of persons to 3 to enable the comparison with a grounded version. Since LiMa uses an abstract state representation, it would have been able to handle larger numbers.

This scenario shows that LiMa is not only capable of tracking object locations, but can also handle scenarios that contain much more internal structure, e.g. multiple entity properties and multiple different actions.

6 Conclusion

We introduced a general version of the *data association problem* and showed that it occurs in many real world scenarios. In this problem class, multiple similar entities act (or get manipulated) simultaneously, while the observations allow only incomplete knowledge about their identity. This leads to a combinatorial explosion of the number of hypotheses. We showed that none of the state of the art approaches satisfies all of the requirements resulting from the generalized version of this problem. However, ideas from different state of the art methods can be combined to a novel filtering algorithm that can handle this problem, which we call *Lifted Marginal Filtering*. The state space representation, and the description of state transitions is based on Multiset Rewriting. Furthermore, the entities are structured, such that additional properties of an entity, like its location or name, can be modeled. Based on the presence of certain information of entity properties, entities are either represented explicitly or grouped together by representing their properties parametrically. We can dynamically modify the granularity of these representations by *splitting* and *merging* abstract states. Two scenarios were used to illustrate the superior state space representation.

Future work includes the automatic switching from exact to approximate inference if the state space gets to large. Another aspect is the investigation of approximate merging, that is, grouping multiple states even if the resulting state does not exactly represent the same distribution. A further research goal is to investigate whether entity multiplicities can be represented parametrically.

References

- Aehnel, M., and Bader, S. 2015. Information Assistance for Smart Assembly Stations. In *ICAART*, 143–150. Lisbon, Portugal: SCITEPRESS.
- Baker, C. L.; Saxe, R.; and Tenenbaum, J. B. 2009. Action understanding as inverse planning. *Cognition* 113(3):329–349.
- Barbuti, R.; Levi, F.; Milazzo, P.; and Scatena, G. 2011. Maximally Parallel Probabilistic Semantics for Multiset Rewriting. *Fundamenta Informaticae* 112(1):1–17.
- Berry, G., and Boudol, G. 1990. The chemical abstract machine. In *POPL*, 81–94. San Francisco, USA: ACM.
- Bistarelli, S.; Cervesato, I.; Lenzini, G.; Marangoni, R.; and Martinelli, F. 2003. On representing biological systems through multiset rewriting. In *EUROCAST*, 415–426. Las Palmas de Gran Canaria, Spain: Springer.
- Braz, R.; Amir, E.; and Roth, D. 2005. Lifted first-order probabilistic inference. In *IJCAI*, 1319–1325.
- Bui, H. H. 2003. A General Model for Online Probabilistic Plan Recognition. In *IJCAI*, 1309–1315.
- Doucet, A.; De Freitas, N.; Murphy, K.; and Russell, S. 2000. Rao-blackwellised particle filtering for dynamic bayesian networks. In *UAI*, 176–183. Stanford, USA: Morgan Kaufmann Publishers Inc.
- Fox, V.; Hightower, J.; Liao, L.; Schulz, D.; and Borriello, G. 2003. Bayesian filtering for location estimation. *IEEE Pervasive Computing* 2(3):24–33.
- Gogate, V., and Domingos, P. 2016. Probabilistic Theorem Proving. *Commun. ACM* 59(7):107–115.
- Gordon, A. D.; Henzinger, T. A.; Nori, A. V.; and Rajamani, S. K. 2014. Probabilistic programming. In *ICSE*, 167–181. Hyderabad, India: ACM.
- Huang, J.; Guestrin, C.; Jiang, X.; and Guibas, L. 2009. Exploiting Probabilistic Independence for Permutations. In *AISTATS*, 248–255.
- Huang, J.; Guestrin, C.; and Guibas, L. 2009. Fourier theoretic probabilistic inference over permutations. *JMLR* 10:997–1070.
- Kersting, K.; Ahmadi, B.; and Natarajan, S. 2009. Counting belief propagation. In *UAI*, 277–284.
- Kondor, R.; Howard, A.; and Jebara, T. 2007. Multi-object tracking with representations of the symmetric group. In *AISTATS*, volume 2, 211–218.
- Krüger, F.; Kasparick, M.; Mundt, T.; and Kirste, T. 2014a. Where are my colleagues and why? Tracking multiple persons in indoor environments. In *IE*.
- Krüger, F.; Nyolt, M.; Yordanova, K.; Hein, A.; and Kirste, T. 2014b. Computational State Space Models for Activity and Intention Recognition. A Feasibility Study. *PLOS ONE* 9(11):e109381.
- Müller, S. M., and Hein, A. 2015. Multi-target data association in binary sensor networks. In *Proceedings of the Fifth International Conference on Ambient Computing, Applications, Services and Technologies (AMBIENT’15)*, 79–84.
- Murphy, K. P. 2002. *Dynamic Bayesian Networks: Representation, Inference and Learning*. Ph.D. Dissertation, University of California, Berkeley.
- Nitti, D.; De Laet, T.; and De Raedt, L. 2013. A particle filter for hybrid relational domains. In *IROS*, 2764–2771. IEEE.
- Poole, D. 2003. First-order probabilistic inference. In *IJCAI*, 985–991.
- Ramírez, M., and Geffner, H. 2011. Goal Recognition over POMDPs: Inferring the Intention of a POMDP Agent. In *IJCAI*, 2009–2014. Barcelona, Catalonia, Spain: AAAI.
- Sukthankar, G., and Sycara, K. 2008. Hypothesis Pruning and Ranking for Large Plan Recognition Problems. In *AAAI*, 998–1003. Chicago, Illinois, USA: AAAI.
- Van Den Broeck, G.; Taghipour, N.; Meert, W.; Davis, J.; and De Raedt, L. 2011. Lifted probabilistic inference by first-order knowledge compilation. In *IJCAI*, 2178–2185.
- Vlasselaer, J.; Meert, W.; Van den Broeck, G.; and De Raedt, L. 2016. Exploiting local and repeated structure in Dynamic Bayesian Networks. *Artificial Intelligence* 232:43–53.
- Wilson, D. H., and Atkeson, C. 2005. Simultaneous Tracking and Activity Recognition (STAR) Using Many Anonymous, Binary Sensors. In *Pervasive*. Springer. 62–79. DOI: 10.1007/11428572_5.
- Zettlemoyer, L. S.; Pasula, H. M.; and Pack Kaelbling, L. 2008. Logical particle filtering. In *Dagstuhl Proceedings*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.