Towards a formalisation of Action Identification Hierarchies*

Pietro Pasotti, Catholijn M. Jonker, M. Birna van Riemsdijk
Delft University of Technology, The Netherlands

Abstract

Our goal is to develop a formal knowledge representation (KR) language oriented to capturing in a natural way the subjective view people have of their behaviour. We draw together Action Identification Theory (AIT) from social psychology and research on cognitive agent programming languages to derive a set of meta-level requirements that the KR language should satisfy. We show that the existing approaches must be extended to suit our purposes, and we propose a general solution: a novel KR language able to express Action Identification Hierarchies (AIHs). Then, we use AIT to give a preliminary discussion of a set of rationality constraints that we can use to tell pathological AIHs from 'good' ones.

1 Introduction

Interactive behaviour support systems (BSS) employ various personalisation techniques in order to be more effective in supporting the user. However, personalisation of models of user behaviour underlying these systems is currently limited. User models are generally constructed at design time, meaning that the basic structure of these models cannot evolve at run-time to better fit the users or to reflect changes in them (cf. [17; 13]).

We know from research in psychology, in particular Action Identification Theory (AIT) [22], that the way people conceptualize their behaviour is subjective, i.e., differs from person to person, and changes over time. This means that in current behaviour support systems there may be a mismatch between the representation of user behaviour in the system and the way people think about their behaviour. We conjecture that a representation (and corresponding behaviour support) that is more in line with people's mental model of their behaviour will help the common grounding between user and system. This contributes to transparency and trust [7] and supports human-machine teamwork [12]. We envisage that this can be done by allowing the system to construct a model of user behaviour incrementally at run-time in dialogue with the user.

*M. Birna van Riemsdijk was supported by the NWO Vidi project CoreSAEP.

Therefore we need a generic knowledge representation structure (KRS) that allows the representation of a wide variety of human behaviour. In this paper we discuss the requirements for this KRS based on AIT (section 2) and analyze to what extent existing work already satisfies these (section 3). Since end users rather than system developers construct these representations, it is important that the system understands which representations "make sense", i.e., which can be regarded as *rational* descriptions of behaviour. We formally define a KRS based on the identified requirements (section 4) and introduce informally a number of rationality constraints to identify which behaviour models can be considered 'good'. We show how these rationality constraints can be motivated by AIT, resulting in a KRS for human behaviour that is grounded in psychological research. Finally in section 5 we wrap up our findings and discuss some future work.

1.1 Motivating example

Our lead example concerns a patient suffering from mild intellectual disability, Pedro. He has difficulty with making sound decisions, judging the time or sequence of steps needed to complete a complex task, thinking things through, predicting how his day will unfold or taking new events into account. For example, he might forget to make a sandwich and bring it along if he goes to work. In winter he may forget to dress warmly, or to bring a coat when travelling. If he is late, he fails to inform his boss.

These issues, in most cases, mean that he needs more-or-less constant supervision from a caregiver throughout the day. We do not aim to help with the mysteries of life or romantic struggles. We aim to capturing such knowledge as the little things of routine and daily life. We want to formalise daily activities so as to understand with sufficient detail what is going on. What is 'sufficient', will depend on the specific application context.

Suppose that Pedro, the user of our system, says:

> Every day I *go home*. I can either go gome *stopping for groceries*, or *without doing the groceries*. If I go home stopping for groceries, that simply means that I "*leave*", "*stop at the grocery store*", and "*arrive late*". On the other hand, if I go home without doing groceries, then I simply leave and then "*arrive early*".

Now suppose that Pedro's (human) caregiver adds the following:

> Going home has three parts: leaving, arriving, and, but only sometimes, doing the groceries. Pedro either arrives early, which happens in most cases, or late, which happens rarely.

By the end of this paper we will have the tools to encode this information and dependency relations among action identifications, enabling reasoning about the relation between identifications and the frequency with which they occur.

2 KR Requirements

Before we can discuss solutions, we need to take a closer look at the problem. We are going to do it through the lens of AIT. We choose AIT because it has proven useful for conceptualising how people think and talk about their behaviour [22].

2.1 Action Identifications

In order to explain the concept of action identification, we use the following metaphor. Imagine yesterday you had been followed around every corner by a cameraman. Now, imagine you were asked to go through the film depicting your whole day and now and then you were asked "what are you doing here"? The words with which you describe the activity/ies you are carrying out in that bit becomes a label, that the cameraman uses to annotate that clip of tape. Then, you would go procedurally over the previously tagged clips and name smaller fragments of those clips with other tags.

The tags we attach to clips are known as 'action identifications' in the context of AIT. AIT puts forward the idea that the understanding of how we identify our activity is crucial for determining the *subjective* view on what it is that we are doing. As such, understanding it is a fundamental step towards interactive systems able to handle the different factors that determine how we address the issues we face and how we behave. For example, When Pedro is going home, he could call what he is doing in a variety of ways, e.g., as "*leaving office*", "*going to rest*", "*getting scolded*" (if he expects he will be late for supper) or "*doing the groceries*".

According to AIT actions are naturally identified at different levels of *concreteness*; the more familiar the action, the more abstract the label. Therefore, we can use the label concreteness (its degree in the hierarchy) to predict how familiar the action is to the actor or whether something is out of place. Furthermore, tracking the transformation patterns of labels, a support system can gain insight into the (evolution of the) degree of familiarity of a behaviour. This can be used to determine the type and amount of support that the user may need with that activity.

One could say that identifications act as *motivation* for their concretisations, and vice versa, they act as specifications for the identifications they are a concretisation of. Why does Pedro arrive home early? Because he didn't stop for groceries. How can Pedro go home? By stopping for groceries (or not). We can ask 'why' and 'how' indefinitely, regressing to evermore abstract or concrete identifications, until we run out of words or fail to find an answer.

2.2 Derivation of Requirements

AIT and the application domain (interactive behaviour support agents) impose requirements on the KR from two sides. In this section we introduce these requirements based on our discussion of AIT, and taking into account a number of additional considerations.

First, we need the language to represent action identifications explicitly so as to facilitate user interaction, and to be amenable to formalisation and have a sound interpretation so that it can be used for KR and automated reasoning.

Moreover, activities are not only characterised by the various ways in which one can identify them through different levels of concretisation as highlighted in AIT, but also by the many things (actions) one can do as part of the process of (plan for) doing them.

For example, "*going home stopping for groceries*" may involve "*leaving*" and "*stopping for groceries*", and in turn "*stopping for groceries*" may be broken down to smaller pieces (entering the shop, paying...). We aim for a KRS that can represent both concretisations over action identifications, following AIT, as well as actions that are part of doing other things, so we can express identifications in terms of their parts (cfr. [22, p.48], 'action scripts'). We envisage that the resulting structure forms a graph whose nodes are action identifications and whose edges determine the relations (concretisation and part-of) between them.

While in AIT the part-of relation is seen as a special case of the concretisation relation, we keep them apart for two main reasons:

1. Many existing approaches from the AI and agent programming literature do (see, e.g., [20]). Aligning our KRS with this work will allow us to import their results or use their reasoning algorithms.
2. Representing complex activities may require sequentially breaking them down into sub-action scripts, especially for applications where the actor is not fully familiar with them or tends to lose track of the progress made so far.

Next to representing action identifications and their relations, many application domains (e.g. behaviour support) require some representation of action frequencies. This, so as to reason about the future given what has happened in the past, and to model and provide support for adopting desired habitual (i.e., frequent) behaviour or changing undesired habits. The idea is that we represent this by attaching probabilities to both part-of and concretisation relations to encode the probability that actions are executed in some way rather than some other.

According to AIT, moreover, when conceptualizing their activities different people will stop at different points in the 'why' and 'how' game of identifying more abstract and more concrete actions, respectively. In practice, this means that the user could add a concretisation or an abstraction of an identification anytime. Consequently our KRS cannot be limited in size a priori, but needs to accommodate structures of any depth.

Finally, since we encode both part-of and concretisation relations, we aim for a KRS that allows them to be combined in different ways. This, together with the other requirements thus discussed, is required to ensure that the language can express sufficiently complex behaviour identification structures, such as those represented in [22, fig. 3.1, p. 46] or [22, fig. 3.5, p. 51]

Summing up, the KRS needs to:

(1) be symbolic and explicit.

(2) have a formal semantics.

(3) model concretisation.

(4) model a part-of relation, which specialises concretisation.

(5) model behaviour likelihoods, both with respect to concretisation and part-of relations.

(6) allow for arbitrarily deep (finite) hierarchies of behaviour identifications.

(7) allow each identification to simultaneously have both parts and concretisations.

(8) allow each identification to simultaneously be a part of other identifications and a concretisation of other identifications.

Requirements (1), (2), (4) (in part) and (5) are technical requirements: we need them to make something that can be usable in practice for behaviour support. The remaining requirements are prescribed by AIT.

3 Related work

In this section we discuss the knowledge representation structures of existing models for representing human behaviour, in particular Activity and Intent Recognition approaches, and for representing agent behaviour, in particular the agent programming literature. We evaluate to what extent they meet our requirements. The results are summarised in Table 1 below.

Activity Recognition The field of Activity Recognition (AR) has developed machine learning approaches to deduce what a human being is doing based on sensor data. A few of those approaches employ ontologies to structure their KR, and we will treat those separately below.

In [14] a review of AR papers is presented. All of the techniques discussed therein are sub-symbolic, most of them have a clear formal definition and all of them are based on a two-layer behaviour/features KRS. Typically, the KRS consists of two layers: a bottom layer of features such as heart rate, accelerometer status, and other sensor measurements; and a top layer of activities such as running, sitting, eating, etc. The layers are connected by a dependency relation that we may conceive of as a part-of relation. This means that these approaches do not encode concretisation relations.

In [8; 9] a sub-symbolic approach is also used, with a formal semantics, but their KRSs are somewhat more expressive than other AR papers. This is because they focus on the hierarchical decomposition of activities in parts, which also means that they do not encode concretisation. They model behaviour likelihoods and potentially they can encode arbitrarily deep part-of structures. The structures are tree-like, which means that a behaviour cannot be part of multiple behaviours (even though the same identification could occur in multiple trees, but it is then unclear how one could merge them).

Ontology-enhanced AR In [10] an AR framework is presented based on log-linear description logics. The KRS they use includes both part-of and concretisation relations. They assume all activities can be categorised in one of four levels of abstraction, and consequently their KR only features at most 4-deep concretisation structures. However, within each class the part-of relation can encode arbitrarily deep structures. So while their structures can be arbitrarily deep, this only concerns part-of relationships and not concretisation, and only the behaviours belonging to the top classes can have both parts and concretisations.

In [19] a KRS is proposed including a representation of sensors, devices, activities and atomic actions. Their approach shows a way of merging DL with probabilistic reasoning. However, the probability values they use are not founded in the semantics and are given by prior manual labelling. Secondly, this approach models concretisation but not part-of relations.

In [5] an approach is described in which DL ontologies are used to provide a semantic characterisation of behaviours in an assisted living scenario. Their KRS is similar to the one presented in [19], with the difference that [5] assumes full certainty and thus does not model behaviour likelihoods.

In [16; 6] an approach is proposed where [5] is one of the steps of an iterative process where a limited set of high-level activities' DL descriptions is given in advance, and the low-level ones are learned through user interaction. Many of the limitations of [5] for our purposes are overcome here, but some still remain. For one, the concretisation structure can be finitely deep, but the part-of relation is restricted to the leaves of the concretisation tree and is not allowed to branch further. This results in limited flexibility in how part-of and concretisation relations can be interwoven.

Goal and Intent Recognition In [24], the authors argue for the need to recognise the goals of humans. It is part of a research line that tries to recognise more abstract aspects of behaviour such as goals and intentions. It uses plan and activity recognisers (see above), and maps goals to observable activities that come from such recognisers. In that sense they already do part of what we suggested above. The KRS are based on goals, observations, timeframes, and plans. The following approaches fall in this category and share one more common attribute: they use sub-symbolic techniques.

Suspicious and anomalous activity recognition is related to intent recognition, and is studied in e.g., [1; 3]. Relevant for us is that these approaches can recognise types of behaviour that cannot be classified by observing a single event, but require instead the observation of multiple events, just like most complex behaviours.

Similar approaches include [15], which presents a supervised machine learning approach based on Markov Models that can be trained to classify raw data into labelled sets that correspond to activities at a few levels of concretisation. The depth of their concretisation hierarchy, however, has to be determined before the training of the model and is hence not quite dynamic enough for our purposes.

From our discussion we conclude that some of the available activity and goal or intent recognition techniques can be used in the monitoring component. In particular, that of [23], and [2] have KRS that can be (partially) mapped to the hierarchical identification structures we develop. Their approaches can furthermore be used to gather information on the frequency of the recognised behaviours.

However the fact that these approaches are sub-symbolic means they are unsuitable for doing KR in our domain, where user interaction is fundamental and it is precisely symbols that are central for effective user interaction and support.

Agent Programming Knowledge representation structures for goals, plans, and actions have been extensively investigated by the agent-programming community, cfr. [4] for a comprehensive overview. In particular Goal-Plan Trees (GPTs) [20] are similar to our KRSs. A GPT is a hierarchical tree of goals and known plans to achieve those goals. The knowledge structures can be arbitrarily deep, but branching is restricted: goals can only concretise to plans and each plan can only be decomposed into a sequence of goals (part of). We relax this constraint and allow indiscriminate chaining of part of links and/or concretisation links. Overall, the KR structure we propose is similar to GPTs, but with less structural constraints so we can be flexible enough to match the user's behaviour identification structure.

In our previous work [18] a basic definition of behaviours is provided in terms of goals and activities that aims at capturing the subjective view of people on their behaviour. The KRS represents concretisation and part-of relations, however these cannot be combined in the expressive way we envisage here, i.e., concretisation and part-of relations cannot be chained freely. Moreover, the paper does not comprise a discussion of rationality constraints, nor does it ground the KRS in psychological literature as we do here using AIT.

We summarize this discussion of related work in Table 1.

	requirements							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
[14]	✗	✓	✗	✓	✓†	✗	✗	✗
[8; 9]	✗	✓	?	✓	✓†	✓†	✗	✗
[10]	✓	✓	✓	✓	✓	✓†	✗$^-$	✗$^-$
[19]	✓	✓	✓	✗	✓‡	✓‡	✓‡	✗
[5]	✓	✓	✓	✗	✗	✓‡	✗	✗
[16; 6]	✓	✓	✓	✓	✓	✓‡	✗	✗
[1; 2; 3] [15; 23]	✗	✓	✗	✓	✗	✗	✗	✗
[20]	✓	✓	✓	✓	✗	✓	✗	✗
[18]	✓	✗	✓	✓	✓	✓‡	✓‡	✓‡
us	✓	✗^{f}	✓	✓	✓	✓	✓	✓

†= only part-of, ‡= only concretisation ?= unclear, $^-$= only for some identifications, ^{f} = future work

Table 1: Summary of relevant literature matched against our requirements.

To conclude, we note that most existing work covers the first four requirements; the main limitations are in the latter ones, i.e. the expressivity of the KR, which is limited by the freedom with which one can interweave concretisation and part-of relations.

4 Action Identification Hierarchies

In this section we present the knowledge structure we propose, called Action Identification Hierarchies (AIHs) and rationality constraints for characterizing when these can be considered rational representations of behaviour.

4.1 Definition

The central component of the structure is precisely *identifications*. We introduce two relations over identifications. First, the $\multimap$ relation represents the 'is a concretisation/way of' relation between identifications. For example, "*go home doing groceries*" is a concretisation of "*go home*": going home *doing anything* is a way of going home. The $\multimap$ relation, as we have seen, orders identifications by generality/specificity.

Adding probabilistic information to this, a user might go home (h) without doing groceries (j) nearly every day, but still go home doing the groceries (f) once in a while. In this case the knowledge structure representing user behaviour will contain $h \overset{x}{\multimap} j$ and $h \overset{z}{\multimap} f$ for some $x > z$. The intuition is that these numbers represent the ratio of h-clips to j- and f-clips respectively, where h-clips are a subset of the j-clips and the f-clips.

Second, the $\mathrel{-\!\lbrack}$ relation expresses the part-of relation between an identification and its parts. More specifically, $\mathrel{-\!\lbrack}$ specifies a (non-necessary) strict 'part of' relation: $a \mathrel{-\!\lbrack} b$ holds iff b *can be executed as part of carrying out* a[1], but carrying our b does not completely carry out a too. So not necessarily b is a required step towards a (there might be other ways to achieve a without going through b); furthermore, a cannot be equal to b. Similarly to $\multimap$, $h \overset{p}{\mathrel{-\!\lbrack}} j$ means that p of all h-movies have some j-clip as a part.

A way of going home, for example, may include the subplan of stopping at the grocery store. In this case, "*stop at grocery store*" is *part of* the "*go home*" identification. Still, it is *possible* to go home without doing the groceries (Pedro maybe does not need them every day). Adding probabilistic information to this relation is to specify the ratio of clips in which the user goes home stopping at the store to the total number of clips depicting him going home (in any way).

We formalize AIHs as follows:

Definition 1 (Action Identification Hierarchies (AIHs))**.** *Let B be a set of action identifications. Then an* AIH *is a structure* $\langle B, \multimap, \mathrel{-\!\lbrack} \rangle$ *where:*

$$\multimap \; := B^2 \times [0,1]$$
$$\mathrel{-\!\lbrack} \; := B^2 \times [0,1]$$

Furthermore, we postulate that no pair b, b' occurs more than once in $\multimap$ or $\mathrel{-\!\lbrack}$. $\forall b, b' \in B$:

$$\exists x.\langle b, b', x\rangle \in \multimap \;\Rightarrow \exists! x.\langle b, b', x\rangle \in \multimap$$
$$\exists x.\langle b, b', x\rangle \in \mathrel{-\!\lbrack} \;\Rightarrow \exists! x.\langle b, b', x\rangle \in \mathrel{-\!\lbrack}$$

And we introduce the following short-hand notation:

$$b \overset{x}{\multimap} b' := \langle b, b', x\rangle \in \multimap$$

and $\overset{+}{\multimap} := \overset{x}{\multimap}$ *where $x > 0$; and the same for* $\mathrel{-\!\lbrack}$ *.*

[1] Can also read: *while* carrying out a.

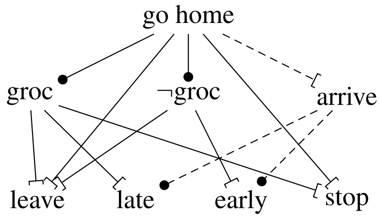

Figure 1: Graph of the ‘get up’ identification. groc = go home doing the groceries, late = arrive late, stop = stop at the grocery store, and so on.

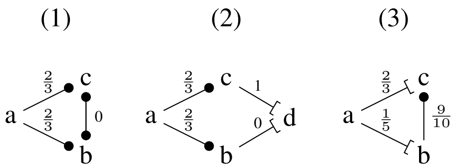

Figure 2: Irrational AIHs.

The AIH of the example presented above can be represented as the graph depicted in fig. 1 (probability values omitted)[2]. Each node is an identification. The information provided by the caregiver is highlighted by means of dashed lines.

4.2 Rationality constraints

In contrast with most other approaches, where the behaviour model is built at design time by the system developer, or constructed by labelling human-generated data, the intent of our KRSs is to let the model be built interactively with the user. However, users can make mistakes and say irrational things. So we need a way of characterising what a rational AIH is, so that we can support the user in constructing a model that can be a representation of actual behaviour.

We introduce a number of *rationality constraints* (in the sense of [11]), that are derived from AIT and probability theory (for the probabilistic part of the framework). We informally describe each constraint, provide some explanation and illustrate it using an example.

rc1 The $\multimap$ relation is reflexive, and the weight of the reflexive edge is always 1.

Motivation: Every identification, even though in an abstract sense, is a way of doing itself.

Example: By going home, you are going home. $b \overset{0.5}{\multimap} b$ is irrational: it is impossible that only 50% of the times you go home you go home: it cannot be but 100%.

rc2 No identification is both a concretisation and a part of some other identification.

Motivation: The concretisation relation links identifications that represent the same underlying action. The part-of relation $\dashv$ represents a strict part-of, i.e., refers to a different action. Hence the $\multimap$ and $\dashv$ relations must be disjoint.

[2]We also omit the redundant links, that can be obtained by closing the graph under the rationality constraints to be explained below.

Example: Going home doing groceries can also be identified as ‘going home’. They are in some sense the same action, as one is a way of achieving (the whole of) the other. However, ‘leaving’ is not the same as ‘going home’ (if not in a metonymy sense), but part of it.

Remarks: As we will see below (cf. **rc**8), $\dashv$ is such that $b \overset{+}{\dashv} b' \wedge b'' \overset{+}{\multimap} b \Rightarrow b'' \overset{+}{\dashv} b'$. So, to ensure that $b \overset{+}{\multimap} b' \Rightarrow b \overset{0}{\dashv} b'$ it suffices to ensure that $b \overset{x}{\dashv} b'$ holds for *some* x, and then by **rc**8 we obtain that if $x = +$, then $b \overset{+}{\dashv} b$, which contradicts **rc**4. Hence it suffices to require that $b \overset{+}{\dashv} b' \Rightarrow b \overset{x}{\multimap} b' \wedge b \overset{+}{\multimap} b' \Rightarrow b \overset{x}{\dashv} b'$. The other constraints will ensure that such x cannot be other than 0.

rc3 Two identifications can be equivalent (i.e. $b \overset{+}{\multimap} b' \wedge b' \overset{+}{\multimap} b$) if and only if they are the same identification.

Motivation: In AIT, low-rank identifications in the hierarchy detail how the identification is done, while high-rank identifications detail why it is done, or what its implications are. It has been shown in the context of AIT that the concretisation relation is commonly understood by people to be asymmetric (cfr. [21, p.4], [22, p.45]).

Example: If doing groceries were a way of going home then going home cannot be a way of doing groceries.

Remark: This constraint can in fact be derived from **rc**1 and the transitivity of $\multimap$.

rc4 No identification is part of (some identification which is part of(...)) itself.

Motivation: This property is a standard assumption in the literature on activity representation. This intuition is also supported by the clip metaphor we introduced earlier: the $\dashv$ relation links clips to their subclips (proper parts of). The $\dashv$ relation is a strict ordering of identifications, hence it is acyclic.

Example: If $b \overset{+}{\dashv} b' \overset{+}{\dashv} b$, this means that b' is a part of b and vice versa. Because of the strict part-of interpretation of $\dashv$, we would reach the paradoxical conclusion that b is a part of itself (and is thus infinitely complex or infinitely simple).

rc5 If two identifications are concretisations of the same identification, then the one cannot be part of the other.

Motivation: If two identifications are concretisations of something, then executing either will achieve that thing (fully). So it makes no sense to say that either one is part of the other, as we know that they are the same thing.

Example: Suppose “*going home by train*” was another way of going home for Pedro. Since going home doing the groceries and going home by train are both ways of going home, it is impossible that going home doing the groceries is part of going home by train, since if you go home doing the groceries that means that by the time you are done you are home already, and it makes no sense to take the train to go home (again). The same applies for the converse.

Remarks: **rc**1, in conjunction with the fact that by definition (of AIH) no pair of identifications can occur twice in $\relbar\!\bullet$ or $\relbar\!\!\sqsubset$, plus $\mathbf{rc}2$, entail that no identification can be part of itself. It is not enough, however, to ensure $\relbar\!\!\sqsubset$-acyclicity.

What is missing is to require that for all identifications b, b' with a common parent it holds that $b \overset{x}{\relbar\!\!\sqsubset} b' \wedge b' \overset{y}{\relbar\!\!\sqsubset} b$. Then $\mathbf{rc}3$, $\mathbf{rc}8$ and $\mathbf{rc}9$ will ensure that $x = y = 0$.

$\mathbf{rc}6$ The probability values on $\relbar\!\bullet$ arrows must be consistent.

Motivation: The probability values on $\relbar\!\bullet$ links are interdependent, because not all identifications can refer simultaneously to the same underlying activity ($\overset{0}{\relbar\!\bullet}$ tells us which ones cannot).

We regret that $\mathbf{rc}6$ and $\mathbf{rc}7$ are too technically complex to discuss here in full detail.

Example: Assuming one cannot do and not do the groceries, it is impossible that $\frac{2}{3}$ of the times Pedro goes home doing groceries, and $\frac{2}{3}$ of the times he does not do groceries; cf. fig. 2.(1).

Remark: In fact, the bounds on the probabilities of each $b \overset{x}{\relbar\!\bullet} b'$ are also dependent on $\relbar\!\!\sqsubset$: you can see it in fig. 2.(2), where it is implied that, paradoxically, $\frac{2}{3}$ of the a-clips have a d-part, and $\frac{2}{3}$ do not.

$\mathbf{rc}7$ The probability values on $\relbar\!\!\sqsubset$ links must be consistent.

Motivation: Similarly as we saw in $\mathbf{rc}6$, the interaction between $\relbar\!\bullet$ and $\relbar\!\!\sqsubset$ and their probability values means that the bounds on each $\overset{x}{\relbar\!\!\sqsubset}$ link depend in a complex way from other information.

Example: Suppose that $\frac{9}{10}$ of the times Pedro goes home he does the groceries. Then suppose that both going home and going home stopping for groceries were in Pedro's mind *parts of* some other identification, say, "*spending the evening with grandma*". It is impossible then that $\frac{2}{3}$ of the times Pedro spends the evening with grandma he goes home stopping for groceries, but only $\frac{1}{5}$ of the times he spends the evening with grandma he *goes home* (no matter how). Cf. fig. 2.(3).

$\mathbf{rc}8$ If an identification b' is a way of b, and b' has a part b'', then b'' is also part of b'.

Motivation: In AIT the part-of relation is a specialisation (a 'subset', cf. [22, p. 49]) of the concretisation relation. This is because achieving a part of something is, stretching a bit, a *way of* achieving that thing (albeit only part of it). Consequently each identification inherits the parts of its children.

Example: If as part of, say, going to the cinema, sometimes you buy popcorn, then also as part of whatever going to the cinema is a way of doing, you sometimes (proportionately less times) buy popcorn.

$\mathbf{rc}9$ If an identification b' is a way of b, and b' is a part of b'', then also b is a part of b''.

Motivation: This is another consequence of $\relbar\!\!\sqsubset$ being a specialisation of $\relbar\!\bullet$.

Example: If buying popcorn is a way of buying food, then definitely as part of going to the cinema you sometimes buy food (specifically, popcorn).

Remarks: While $\relbar\!\!\sqsubset$ specialises $\relbar\!\bullet$, the opposite is not true. Therefore $a \overset{+}{\relbar\!\bullet} b \overset{+}{\relbar\!\!\sqsubset} c$ does not imply that $a \overset{+}{\relbar\!\bullet} c$, and in fact it implies that $a \overset{0}{\relbar\!\bullet} c$ in our formalisation.

$\mathbf{rc}10$ $\relbar\!\bullet$ is transitive, and the probability of the transitive links is the product of the probabilities of the constituting links.

Motivation: Studies in AIT show that $\relbar\!\bullet$ is transitive (cfr. [22, p.51]). As the probability values encode dependency among identifications, we expect transitive links to have probabilities equal to the product of those of the intermediate links. This is a consequence of the probabilistic interpretation of the weights over the links.

Example: If 0.2 of the times Pedro goes home he does the groceries, and, suppose, 0.2 of the times he does the groceries he does the groceries in a hurry, then exactly $0.2 \cdot 0.2$ of the times Pedro goes home he does groceries in a hurry.

$\mathbf{rc}11$ If two identifications are known to have a common concretisation, then they cannot be disjoint.

Motivation: If b, b' have a common concretisation b'', that means that b'' is a set of clips labelled with both parent identifications. The parents can be $\overset{0}{\relbar\!\bullet}$ -related only if they label entirely disjoint sets of clips, and b'' is in fact witnessing the contrary. So neither $b \overset{0}{\relbar\!\bullet} b'$ nor $b' \overset{0}{\relbar\!\bullet} b$ can be the case.

Example: Suppose doing the groceries is, beside being a way of going home, also a way of "*taking care of grandma*". Then it is impossible that there is no way in which you can go home *and* take care of grandma at the same time.

In summary, each constraint is motivated by AIT or probability theory. This makes each one of them necessary for determining when an AIH is rational. Once formalised, these constraints can be used for automatically checking at run-time whether an AIH that is being constructed in interaction with the user is good or not.

5 Discussion and Future Work

We have presented a KRS called AIH based on Action Identification Theory. AIT, probability theory and technical necessity led us to define rationality constraints that each rational AIH should satisfy.

We have argued that these constraints are necessary, however, the question remains whether they are sufficient for adequately capturing what we mean by rational descriptions of behaviour. In future work we will formally specify these constraints and develop a formal semantics for AIHs. We will use this to prove that the constraints we provide are sufficient to characterise rational AIHs. Specifically, we envisage translation of our KRS to Description Logic (DL) and proving that AIHs satisfying our rationality constraints yields a consistent DL theory. Moreover, since action identifications are contextual and evolve over time (cf. [22]), in future work we plan on showing how AIHs can be combined with contextual and temporal information representation. We plan to evaluate the resulting framework against existing data on human action identification hierarchies.

References

[1] *Towards Detection of Suspicious Behavior from Multiple Observations*, Plan, Activity, and Intent Recognition, 2011.

[2] Liliana Ardissono and Dario Sestero. Using dynamic user models in the recognition of the plans of the user. *User Modeling and User-Adapted Interaction*, 5:157–190, 1996.

[3] Dorit Avrahami-Zilberbrand and Gal A. Kaminka. Keyhole adversarial plan recognition for recognition of suspicious and anomalous behavior. In Gita Sukthankar, Robert P. Goldman, Christopher Geib, David V. Pynadath, and Hung Bui, editors, *Plan, Activity, and Intent Recognition*. Morgan Kaufmann, 2014. In press.

[4] Rafael H. Bordini, Mehdi Dastani, Jürgen Dix, and Amal El Fallah Seghrouchni. *Multi-Agent Programming: Languages, Tools and Applications*. Springer, Berlin, 2009.

[5] Liming Chen, Chris Nugent, and Ahmad Al-Bashrawi. Semantic data management for situation-aware assistance in ambient assisted living. *Proceedings of the 11th International Conference on Information Integration and Web-based Applications & Services*, 2009.

[6] Liming Chen, George Okeyo, and Hui Wang. An ontology-based hybrid approach to activity modeling for smart homes. *IEEE Transactions on Human-Machine Systems*, 44(1), 2014.

[7] Andy Clark and David Chalmers. The extended mind. *Analysis*, 58(1):7–19, 1998.

[8] Thi V. Duong, Hung H. Bui, Dinh Q. Phung, and Svetha Venkatesh. Activity recognition and abnormality detection with the switching hidden semi-markov model. In *CVPR 2005 : Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 838–845, 2005.

[9] Shai Fine, Yoram Singer, and Naftali Tishby. The hierarchical hidden markov model: Analysis and applications. *Machine Learning*, 1998.

[10] Rim Helaoui, Daniele Riboni, and Heiner Stuckenschmidt. A probabilistic ontological framework for the recognition of multilevel human activities. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '13, pages 345–354, New York, NY, USA, 2013. ACM.

[11] Koen V. Hindriks. *Multi-Agent Programming*, chapter Programming Rational Agents in GOAL, pages 119–157. Springer, 2009.

[12] Catholijn M. Jonker, M.Birna van Riemsdijk, and Bas Vermeulen. Shared mental models - a conceptual analysis. In van der Hoek, editor, *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, 2010.

[13] Catriona M. Kennedy, John Powell, Thomas H. Payne, John Ainsworth, Alan Boyd, and Iain Buchan. Active assistance technology for health-related behavior change: An interdisciplinary review. *Journal of medical internet research*, 14(3), 2012.

[14] Óscar D. Lara and Miguel A. Labrador. A survey on human activity recognition using wearable sensors. *IEEE Communications Surveys and Tutorials*, 15(3):1192–1209, 2013.

[15] Sebastian Lühr, Hung H. Bui, Svetha Venkatesh, and Geoff A.W. West. Recognition of human activity through hierarchical stochastic learning. *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications (PerCom'03)*, 2003.

[16] George Okeyo, Liming Chen, and Hui Wang. Combining ontological and temporal formalisms for composite activity modelling and recognition in smart homes. *Future Generation Computer Systems*, 39:29–43, 2014.

[17] Harm op den Akker, Valerie M. Jones, and Hermie J. Hermens. Tailoring real-time physical activity coaching systems: a literature survey and model. *User Modeling and User-Adapted Interaction*, 24(5):351–392, 2014.

[18] Pietro Pasotti, M.Birna van Riemsdijk, and Catholijn M. Jonker. Representing human habits: towards a habit support agent. In *Proceedings of the 10th International workshop on Normative Multiagent Systems (NorMAS'16)*, LNCS. Springer, 2016. To appear.

[19] Daniele Riboni, Timo Sztyler, Gabriele Civitarese, and Heiner Stuckenschmidt. Unsupervised recognition of interleaved activities of daily living through ontological and probabilistic reasoning. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2016.

[20] John Thangarajah. *Managing the Concurrent Execution of Goals in Intelligent Agents*. PhD thesis, Royal Melbourne Institute of Technology, 2005.

[21] Robin R. Vallacher and Daniel M. Wegner. What do people think they're doing? action identification and human behavior. *Psychological Review*, 94(1):3–15, 1987.

[22] Robin R. Vallacher and Daniel M. Wegner. *A theory of action identification*. Psychology Press, 2014.

[23] Mor Vered and Gal A. Kaminka. A computational cognitive model of mirroring processes: A position statement. In *Plan, Activity, and Intent Recognition*, 2013.

[24] Mor Vered, Gal A. Kaminka, and Sivan Biham. Online goal recognition through mirroring: Humans and agents. In *Proceedings of the Annual Conference on Advances in Cognitive Systems*, 2016. A slightly modified version appears in Proceedings of the IJCAI 2016 workshop on Human-Agent Interaction Design and Models (HAIDM).