Ontology Engineering for the Design and Implementation of Personal Pervasive Lifestyle Support

Michael A. van Bekkum
TNO
Kampweg 5,
3769 DE, Soesterberg, Netherlands
michael.vanbekkum@tno.nl

Hans-Ulrich Krieger
DFKI GmbH
Stuhlsatzenhausweg 3
D-66123 Saarbrücken, Germany
krieger@dfki.de

Mark A. Neerincx
TNO
Kampweg 5,
3769 DE, Soesterberg, Netherlands
mark.neerincx@tno.nl

Frank Kaptein
Delft University of Technology
P.O. Box 5,
2600 AA, Delft, Netherlands
f.c.a.kaptein@tudelft.nl

Bernd Kiefer
DFKI GmbH
Stuhlsatzenhausweg 3
D-66123 Saarbrücken, Germany
bernd.kiefer@dfki.de

Rifca Peters
Delft University of Technology
P.O. Box 5,
2600 AA, Delft, Netherlands
r.m.peters@tudelft.nl

Stefania Racioppa
DFKI GmbH
Stuhlsatzenhausweg 3
D-66123 Saarbrücken, Germany
stefania.racioppa@dfki.de

ABSTRACT
The PAL project[1] is developing an *embodied conversational agent* (robot and its avatar), and applications for child-agent activities that help children from 8 to 14 years old to acquire the required *knowledge, skills, and attitude* for adequate diabetes self-management. Formal and informal caregivers can use the PAL system to enhance their supportive role for this self-management learning process. We are developing a **common ontology** (i) to support normative behavior in a flexible way, (ii) to establish mutual understanding in the human-agent system, (iii) to integrate and utilize knowledge from the application and scientific domains, and (iv) to produce sensible human-agent dialogues. The common ontology is constructed by relating and integrating partly existing separate ontologies that are specific to certain contexts or domains. This paper presents the general vision, approach, and state of the art.

CCS Concepts
Computing methodologies → Artificial intelligence → Knowledge representation and reasoning → Ontology engineering

Keywords
Ontology engineering; common ontology; embodied conversational agent; DIT++ standard; HFC inference engine; human-agent dialogue.

1. INTRODUCTION
In Europe, an increasing number of about 140,000 children (<14 year) have Type 1 Diabetes Mellitus (T1DM) [4]. The **PAL** project develops an *Embodied Conversational Agent* (**ECA**: robot and its avatar) and several applications for child-agent activities (e.g., playing a quiz and maintaining a timeline with the agent) that help these children to enhance their self-management .

PAL is part of a joint cognitive system in which humans and agents share information and learn to improve self-management. The required sharing of (evolving) knowledge has four important challenges:

1. To address the values and norms of both the caregivers and the caretakers in their different contexts (e.g., diabetes regimes, privacy).
2. To establish mutual understanding between the different human stakeholders of the PAL system, e.g., the end-users (children, caregivers), researchers and developers (e.g. academics, engineers).
3. To acquire, utilize, and deploy knowledge about child's self-management support.
4. To support natural and personalized interaction between the humans and PAL system agents.

In PAL, we are trying to meet these four challenges by developing a **common ontology** as an integrated part of the system development. The ontology addresses the aforementioned challenges by (1) serving as a knowledge basis for requirements analysis, (2) providing an unambiguous vocabulary and communication between stakeholders, (3) supporting system implementation of knowledge-based reasoning functionalities and (4) serving as a basis for interoperability in human-agent interaction. Engineering this ontology is a systematic, iterative, and incremental development process. Firstly, available ontologies and approaches are assessed on relevance and, possibly, adapted and integrated for our purposes (cf. Section 2). Secondly, relevant theories and models of the concerning scientific research fields are identified and formalized for adoption in the ontology (cf. Section 3). Thirdly, the ontology is implemented in an artefact or prototype, and subsequently, tested and refined (cf. Sections 4 and 5).

---

[1] **PAL**, *Personal Assistant for healthy Lifestyle*, is an European Horizon-2020 project; http://www.pal4u.eu

2. Engineering PAL Ontologies for Diabetes Self-Management Support

Because PAL covers a large domain of interest, we have developed separate ontology models as high-level building blocks for smaller, more specific areas of interest (frames). We have subsequently modeled each frame by either developing a new ontology or by selecting relevant, already existing models from (global) libraries that are similar in scope to the frame of interest.

The frames we have identified and modeled so far are among others (1) human/machine roles/actors involved in self-management, (2) task/goal/activity that includes self-management activities, tasks, associated goals, and results and the setting they take place in, (3) diabetes self-management activities and games, (4) issues related to medical examinations (e.g., lab values), and (5) dialogue management through a combination of dialogue acts and shallow semantic frames. A more elaborate PAL ontology will also include interaction and behavior models of robot and avatar, a model for privacy of information of self-management activities, and a model to cover the agreements and social contracts between patient and avatar/robot and a model for emotion and sentiment that covers the emotional responses of both robot and child to interaction as well as the general state of mind of the child.

As a modelling strategy, we have turned to existing (global) libraries to cover the various frames. Although our frames of interest are typically generic in nature, pre-existing models for these frames may differ (slightly) in scope and/or intention and may thus be a partial fit to the intended scope of PAL. Whereas e.g. the self-management part of diabetes is a relevant topic, the entire professional medical diagnosis and treatment model of diabetes may not be relevant here. We have therefore adapted these models whenever required by either extending them when concepts are missing or by selectively downsizing them when there are details/concepts in the model that are irrelevant to the scope of PAL. An example of reuse is displayed in the adoption of the well-known ontology for task world models [10] in the frame for tasks/goals. We have used this model at the core, but extended it with the Group concept, as a collection of Agents. At the same time, the notion of (external) Events triggering Tasks, has been discarded.

In the PAL project, dedicated studies of models in the scientific research areas concerned are also being conducted in order to adequately represent the frames of interest. For supporting the social processes that are involved in self-management learning, PAL models relationships in terms of familiarity or intimacy, liking, attitude, and benevolence [1]. Cognitive processes, the diabetes knowledge and corresponding learning goals have been explicitly modeled for purposes of monitoring and reasoning (aiming at personalized feedback by the ECA).

The affective process and state of a child are represented by a child ontology that allows the PAL system to estimate emotions experienced by the child resulting from activities proposed by the ECA. For example, the ECA can propose to play a quiz with the child, and predict joy as the emotional state of the child when the child did well during the game. This requires a complex affective state to be stored that contains all the affective information, be they, e.g., emotions (short, intense episodes) or moods (prolonged period of time). Emotions in this case need to be related to both child and the activity that had this emotion as a consequence. Moods need to contain a timestamp, indicating during which period it was measured. This representation makes it possible to find correlations between activities and affect over a prolonged period of time.

3. Extended Representation

The goal of the PAL ontology is norm-compliancy, shared understanding, interpretation, reasoning, and the generation of speech acts (e.g. verbal utterances). The ontology is based on a *uniform representation* of an application semantics that uses *dialogue acts* and *shallow semantic frames*, being represented by an extension of the RDF/OWL format (triples/binary relations) [9]. In addition, all user and other data that influence multimodal generation are specified in the ontology which facilitates access and combination of the different bits of information. We have extended existing processing components, e.g., the reasoning engine *HFC* [5] from DFKI and its database layer which makes information available to the interaction management and analysis.

One important part of the PAL ontology combines dialogue acts utilizing the DIT++ standard [2] and semantic frames, loosely based on thematic relations [3], used in today's frameworks VerbNet, VerbOcean, or FrameNet. Here is a simplified version of the combined representation that will be built for the sentence *Would you like to play a quiz*?

```
Offer[sender=MYSELF, addressee=ROBOT, ... ,
      frame=(Asking,agent=MYSELF,
             patient=ROBOT, theme=Quiz]]
```

Dialogue acts as well as semantic frames may contain further properties not depicted here; e.g., to represent the *continuation of dialogue acts* via `follows` or to model `indirect speech` through `refersTo`. Both properties map, again, to dialogue acts that have been introduced earlier in the conversation. The seemingly redundant specification of both dialogue act `Offer` and frame `Asking` is motivated by the fact that a *positive* answer to the question (= `AcceptOffer`) still refers to the `Asking` frame (*I'm accepting the offer you had asked for = yes*).

We have also defined a new way to marry the RDF-based triple representation with **transaction time** [6], as known from temporal databases [8]. This is possible because the inference engine HFC is based on *general tuples (n-ary relations)*, instead of restricting itself to triple-based representations. In the implementation of *HFC*, we employ 8-byte long integers (XSD datatype `long`) to encode *milli* or even *nano* seconds w.r.t. a fixed starting point (viz., Unix Epoch time). As a consequence, given a time point *t*, the next *smallest* or *successor* time point would then be *t+1*, thus our time is *discrete*. Like in *valid time*, the original approach to *transaction time* makes use of temporal intervals in order to represent the time during which a fact is stored in the database, even though the ending time is not known in advance. This is indicated by the wildcard **?** which will later be overwritten by the concrete ending time. We deviate here from the interval view by specifying both the *starting time* when an ABox statement is *entered* to an ontology, and, via a separate statement, the *ending time* when the statement is *invalidated*. For this, we exploit the propositional modals ⊤ and ⊥ (see [6]).

This idea is shown in the following picture for a binary relation P. We write P(c,d,b,e) to denote the row <c,d,b,e> in the database table P for relation P.

Time	Database View	Ontology View
⋮	⋮	⋮
t_1	add: P(c,d,t_1,?)	add: $\top$P(c,d)@t_1
⋮	⋮	⋮
t_2	overwrite: P(c,d,t_1,t_2)	——
t_2+1	——	add: $\bot$P(c,d)@t_2+1
⋮	⋮	⋮

Figure 1: Tuple representation with transaction time

As we see from this picture, the invalidation in the ontology happens at t_2+1, whereas [t_1, t_2] specifies the transaction time in the database. Clearly, the same transaction time interval for P(c,d) in the ontology can be derived from the two statements $\top$P(c,d)@ t_1 and $\bot$P(c,d)@ t_2+1, assuming that there does not exist a $\bot$P(c,d)@*t*, such that $t_1 \leq t \leq t_2$ (we can effectively query for this by employing an ValidInBetween test in our SPARQL-like queries). Extending ontologies by transaction time the way we proceed here gives us a means to easily encode *time series data*, i.e., allows us to record the *history* of data that changes over time. Time-stamped data such as x $\top$ P(c,d)@t is represented in *HFC* by a quintuple: `true c P d t`.

4. Application and Implementation

The implementation of the ontology models is done in a PAL system that consists of several modules. A *dialogue manager* is, e.g., responsible for conversations between child and avatar/robot, an *action-selection module* decides what the best actions are at a particular moment (e.g., when playing a game), while a *child model module* is able to reason on the (emotional) state of the child.

This modular setup of the system requires clearly defined knowledge representations for each of the modules: the set of PAL ontologies for diabetes self-management support. To this end, the system has all individual ontologies defined in the extended *HFC* reasoner and has connected semantically-equivalent concepts, properties, and individuals through OWL interface axioms, utilizing standard constructors, such as `rdfs:subClassOf` or `owl:equivalentProperty`, or by posing domain and range restrictions (e.g., `rdfs:domain`).

The ontology engineering in PAL is meant to be an iterative and incremental process, with continuous refinement and extension of the involved ontologies. The models are adjusted according to new insights and continue to be aligned with sources of information in the entire project. The development of the ontology thus runs in parallel with the system design it is supposed to support and the modular approach allows for testing and refining these incrementally.

The PAL system is currently used in hospitals, diabetes camps, and at home. The analysis of the first recorded data sets for children and caregivers that have used the PAL system from a few days to four weeks in Italy and the Netherlands, is underway. Based on the ontological concepts, we will be able to identify meaningful patterns in the data that will be used to improve the intelligence of PAL, both in the knowledge base with refined ontologies, as well as its associated reasoning mechanisms. E.g., for the provisions of personalized feedback based on the identified user profiles, addressing cultural differences [7].

5. Discussion

This paper presents the development of a common ontology that underpins the design and implementation of an ECA for children, that help these children to acquire the required knowledge, skills, and attitude for adequate diabetes self-management. This ontology is used to establish mutual understanding in the human-agent system, to integrate and utilize knowledge from the application and scientific domains, and to produce natural human-agent dialogues. A set of interconnected ontologies ("frames") have been constructed, each consisting of general concepts and their relations: (1) roles and actors, (2) task, goal, activity and context, (3) diabetes self-management, and (4) dialogue management. We developed the first version of an ontology which species the data structures to be used by the dialogue specifications, dialogue history, and information state, and adapted our reasoning components, so that this knowledge source can be used efficiently once the formalism specification is fully implemented. The current version of the ontology is available at the PAL ontology website (http://www.dfki.de/lt/onto/pal/).

6. Acknowledgements

The research described in this paper has been funded by the Horizon 2020 Framework Programme of the European Union within the project PAL (Personal Assistant for healthy Lifestyle) under Grant agreement no. 643783.

7. REFERENCES

[1] Altman, I. and Taylor, D. Social penetration theory. Holt, Rinehart & Mnston, New York. (1973).

[2] Bunt, H., Alexandersson, J., Choe, J.-W., Fang, A.C., Hasida, K., Petukhova, V., Popescu-Belis, A., and Traum, D. ISO 24617-2: A semantically-based standard for dialogue annotation. In *Proceedings of the 8th International Conference on Language Resources and Evaluation (LREC)*. (2012).

[3] Fillmore, C.J. (1977). The Case for Case Reopened. In *Grammatical Relations*. Syntax & Semantics. Academic Press.

[4] D. Freeborn, T. Dyches, S. O. Roper, and B. Mandleco. Identifying challenges of living with type 1 diabetes: child and youth perspectives. *Journal of clinical nursing, vol. 22, no. 13-14*. (2013), 1890–1898.

[5] Krieger, H.-U. An Efficient Implementation of Equivalence Relations in OWL via Rule and Query Rewriting. *Proceedings of the 7th International Conference on Semantic Computing (ICSC).* (2013).

[6] Krieger, H.-U. Integrating Graded Knowledge and Temporal Change in a Modal Fragment of OWL. In *Agents and Artificial Intelligence, edited by Jaap van den Herik & Joaquim Filipe. Lecture Notes in Computer Science.* Springer. To appear. (2016).

[7] Neerincx, A., Sacchitelli, F., Kaptein, R., van der Pal, S., Oleari, E., and Neerincx, M. A. (2016). Child's Culture-related Experiences with a Social Robot at Diabetes Camps. *Eleventh ACM/IEEE International Conference on Human Robot Interaction* IEEE Press. (2016), 485-486.

[8] Snodgrass, R.T. Developing Time-Oriented Database Applications in SQL. Morgan Kaufmann, San Francisco, CA. (2000).

[9] ter Horst, H. J. Completeness, decidability and complexity of entailment for RDF Schema and a semantic extension involving the OWL vocabulary. Journal of Web Semantics, 3 (2005), 79–115.

[10] Welie, M. Van, Veer, G. C. Van Der, & Eliëns, A. An ontology for task world models. *Design, Specification and Verification of Interactive Systems '98*. (1998), 57–70.