Goal Processing and Semantic Matchmaking in Opportunistic Activity and Context Recognition Systems

Gerold Hoelzl, Marc Kurz, Alois Ferscha
Institute for Pervasive Computing
Johannes Kepler University Linz
Linz, Austria
surname@pervasive.jku.at

***Abstract*—The Opportunistic Sensing Paradigm shifts away the configuration of activity- and context recognition systems from design time of the system to a dynamic runtime configuration. Systems following this new paradigm have to autonomously configure themselves during runtime according to a stated recognition goal and the available sensing infrastructure. A crucial point therefore is to find a methodology to express the recognition goal itself and to map the requirements specified by the stated recognition goal to the capabilities of the available sensors in the ecosystem of the users. This paper presents an approach for a semantic engine capable of processing and matching a given recognition goal to the semantically described sensing infrastructure. This matching allows to autonomously spot and quantify the best set of sensors available at any point in time that can contribute to the stated recognition goal and therefore being configured to a sensing ensemble.**

***Index Terms*—Goal Processing; Sensor Networks; Activity and Context Recognition; Opportunistic Sensing**

I. Introduction

Sensor networks enable the sensing of the physical world and detect context and activity, which is a key to build and develop intelligent environments [1]. Different Frameworks exist that handle inferring context and activity out of raw sensor data. Their common disadvantage is that the used sensing infrastructure and the context that has to be recognized must be defined at the design time of the system [2]. A dynamic change in the sensing infrastructure in terms of the spontaneous availability of new sensors or the change of the recognition purpose respectively the recognition goal [3] is not supported by these systems. Due to the emerging plethora of sensing devices (e.g., smart phones or smart watches) with integrated sensing capabilities in the ecosystem surrounding the user, the necessity arises to federate these devices in an autonomous, adaptive and goal oriented manner to relieve developers from low level, platform specific concerns. Kurz et al. present in [3] methodologies to semantically describe the recognition capabilities of sensors in terms of labels reflecting meaning in the real world (e.g., *Walk* or *Drinking-Coffee*). They propose *ExperienceItems* that encapsulate the necessary stages of the *Activity Recognition Chain* [4] to infer activity and context information out of the raw sensor data readings. In addition, a metric, the *Degree of Fulfillment (DoF)* is defined, that quantifies how good a sensor can be used to detect a specific label (i.e., activity class) as a value between $[0, 1]$. Furthermore, a way to abstract from the low level access details yielding to a common accessible interface for different sensing devices using *Sensor Abstractions* is described and evaluated. Using these presented methodologies, we can see a sensor as a label delivering entity that can be queried for its capabilities, and the needed machine learning techniques can be instantiated during runtime to infer activity and context information. To direct the autonomous planning process [5] in selecting the best suitable sensing devices, we propose a goal oriented approach that defines at a high semantic level how the system should behave. Goal oriented approaches have proven their flexibility and effectiveness in various research domains like *Operations Research* [6], *Requirements Engineering* [7] and *Mobile Agent Systems* [8]. The common objective is that a goal specifies at an abstract level what the system should achieve but not how. Multiple paths exist at each point in time on how the goal can be achieved. In this work we adapt the goal methodology and use a goal oriented approach to direct the configuration of an activity and context recognition system. This is a novel approach as the sensing infrastructure needs not to be defined at design time of the system and can change even during its runtime. It is autonomously configured and adapted according to the stated recognition goal. In detail, the hypotheses valid within this paper can be formulated as follows:

(H) *The explicit formulation of a recognition goal is a suitable way of directing the autonomous configuration of an activity and context recognition system.*

The remainder of the paper is structured as follows: Section II proposes a methodology on how activity can be understood and modeled. This is a crucial point to define the necessary terms a recognition goal can be composed of. Based upon our activity modeling approach that encapsulates our understanding of what an activity is, Section III discusses a technique on how to formulate a recognition goal based on this understanding. Section IV shows how a formulated

 ISBN: 978-1-61208-257-8

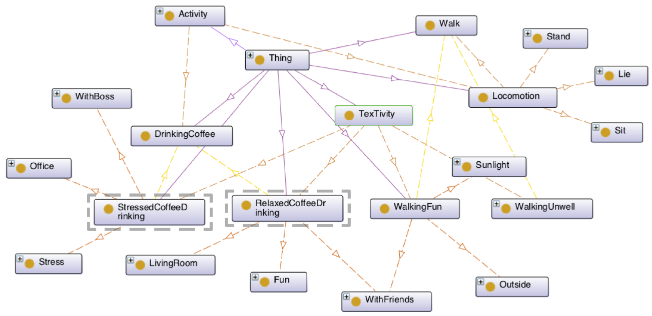

Fig. 1: Partial view of the approximately 200 modeled *Context* and *Activity Relations* gathered during a large scale kitchen experiment presented in [11] using the Protege-OWL modeling tool. Activities and Context are modeled using the relations $\Re_{isRelated}$, $\Re_{Combination}$, and $\Re_{has\_A}$.

recognition goal is processed, refined, semantically reasoned and quantified, and finally mapped to the available sensing ecosystem. Section V evaluates the proposed approach in a real-time setting and Section VI closes with a discussion of the hypotheses and a conclusion.

II. Activity Relations Modeling

To formulate a goal that describes the recognition purpose of the system at an abstract, semantic level, the necessity evolves to understand, define, model, and relate *Activities* and *Context* that can be stated to the system. We see and define *Activities* as a much more complex thing rather than a simple *"sequence of actions performed"* [9] or *"hierarchies of activities of daily living"* [10]. Its meaning is not only defined by the performed physical activity itself. It is defined in combination with a situation respectively the context. To model these relations we use an *Ontology* that builds the semantic knowledge base and represent the *Activity Relations* gathered during a large scale kitchen experiment presented in [11].

In each context, a physical activity can have a complete different meaning. From our point of view *Activity* can be seen as any information that characterizes the behavior of an individual interweaved with additional context [12]. We model these relations using the *TexTivity*-Concept (as shown in (1)). This concept allows to relate activities to different contexts and assign different meaning to these relations. *TexTivity* therefore is an amalgam of *context* and *activity* highlighting the close relation of both concepts. Fig. 1 shows a partial view of the modeled relations. Two *TexTivities* namely *StressedCoffeeDrinking*, and *RelaxedCoffeeDrinking* are modeled. Both *TexTivities* are related to same physical activity *DrinkingCoffee* but assign a different meaning of this activity in relation to the two contexts *[WithBoss, Office, and Stress]→StressedCoffeeDrinking* and *[WithFriends, Outside, and Fun]→RelaxedCoffeeDrinking*. According to the definition of *TexTivities* we named the Ontology to model these relations *TexTivity-Ontology*. Using the definition of *TexTivities* (as defined in (1)) allows us to link a set of context literals (without activity) [$\Re_{Combination}$] to a set of activity literals [$\Re_{isRelated}$] and give this relation a "meaning". Furthermore, by knowing the relations of activities [$\Re_{has\_A}$], the substitution and reasoning of activities is possible by combining activities at different levels of abstraction. This is especially useful if no sensor can be found in the surrounding of the user that is explicitly designed to detect the stated recognition goal. Using the modeled relations of activities, we can semantically reason which activities the stated recognition goal is composed of, and use sensors that can detect these semantically related activities instead. This allows the dynamic configuration of sensing ensembles and exploits the full flexibility of the opportunistic sensing approach, as the recognition goal can be processed, reasoned and matching methodologies can be applied to substitute activities out of related ones. How this is done is described in detail in Sections III and IV.

$$TexTivity \rightarrow \{Activity\} \odot \{Context_{\setminus Activity}\} \qquad (1)$$

We see *Activities* as a multidimensional labyrinth of complex actions. They can happen interleaved, concurrent or in any other temporal or causal relationship. Making it even more complex, *activities* can happen in combination with different contexts that influence their meaning. Consider the activity *"Talking"* either in the context of watching a football match or improperly in a class room. Both activities *"Talking"* refer to the same *"physical"* action but with a complete different meaning in the particular context.

III. Recognition Goal Formulation

One of the crucial things when developing an opportunistic activity and context recognition system that adapts autonomously during runtime is to find a way to formulate the recognition goal explicitly to relieve developers from low level, platform specific concerns while federating the sensing infrastructure. The explicitly stated recognition goal has to be reason- and processable by the opportunistic system. This enables the autonomous selection of the best set of available sensors, called *ensemble* [3], according to the stated recognition goal. We see a sensor as a label delivering entity at a semantic level reflecting meaning in the real world that can be used and queried for its capabilities, can be quantified during runtime using the *Degree of Fulfillment (DoF)* metric, and can be dynamically instantiated at any point in time. How the sensor ensemble is configured and the corresponding activity recognition chains are instantiated by using sensor abstractions, sensor-self-descriptions, and experience-items is described in [3].

To formulate the recognition goal we use a derivation of the *Context Predicates* presented in [13]. There, Stevenson and Dobson present Context Predicates to reason information in smart environments (e.g., smart homes). The Context Predicates are used as entry points into the modeled domain knowledge captured in form of an Ontology. Furthermore, they are used to reason about the modeled knowledge. If for example, one wants to know if Bob is at home, and the information that Bob is in the living room is given (by a particular sensor), the domain knowledge can be used to reason that the living room is located in the house and to infer that Bob is at home. The reasoning capabilities can be used to infer and refine knowledge that is not explicitly given by a particular sensor. Even if we do not have the explicit information that Bob is at home (e.g., no sensor is available that can deliver this information explicitly), we can infer this knowledge by using other sensors that detect that Bob is in the living room. Having the semantic modeled relation of *Home* ↔ *Living Room* we can infer that Bob is at home without having the explicit knowledge of a particular sensor. The syntax of the *Context Predicate* is defined in (2).

$$< CP >=< subject, predicate, object > \tag{2}$$

The following examples clarify the use of Context Predicates. A Context Predicate can be for example $< \{Bob\}, located, livingRoom >$, querying if Bob is located in the living room. This query can be evaluated to *true* or *false*. Another example is $<?, located, livingRoom >$. Using the question mark, querying for all subjects that are located in the living room is possible. The last example $< Person, ?, livingRoom >$ queries for all predicates (activities) persons are performing in the living room.

Following the approach of using Context Predicates, that showed to be a promising solution as a goal description using domain knowledge and reasoning, we enhanced them to more closer fit our proposed *TexTivity*-Ontology to enable a quick and precise formulation of the recognition purpose of the system. As we talk about activity and context recognition, we renamed the *predicate* into *Activity* and made the term *object* more generic in calling it *Context*. Furthermore we added the entity *TexTivity* to allow its direct statement as a goal. We permit to state each item in form of a set to allow multiple objectives per item being stated in a single goal and define a *TexTivity Predicate* as follows (3):

$$\begin{aligned} < TP >=< \{Subject\}, \{Activity\}, \{Context\}, \\ \{TexTivity\}, \{Constraint\} > \end{aligned} \tag{3}$$

The definition of a TexTivity Predicate can be used to formulate a goal like $< \{Bob, Paul\}, ?, \{LivingRoom, BathRoom\}, ?, \{\} >$. The formulated goal configures a system that detects all activities and TexTivities (as indicated by the question marks) of Bob and Paul in the living- and bathroom. As from the systems point of view, the best ensemble is always defined by having the highest recognition accuracy concerning the stated context and activity items, the necessity may evolve to further optimize such an ensemble in terms of energy consumption, size and weight, infrastructure or body-worn sensors, or the maximum numbers of sensors that are configured to an ensemble. As these constraints cannot be foreseen now, the TexTivity predicates allow to state them in an open and extendable way. In the previous example, no constraints were defined. So the best available sensor ensemble is configured regardless its quantitative properties. If one wants to formulate the same recognition goal, but assuring that the configured ensemble is above a stated recognition accuracy, e.g., with a *Degree of Fulfillment (DoF)* above 89%, it has to be added as constraint ($< \{Bob, Paul\}, ?, \{LivingRoom, BathRoom\}, ?, \{[DoF] : 0.89\} >$) instructing the system to configure a sensing ensemble with a DoF above 89% if possible, or fail otherwise. Formulating subjects, activities, context, TexTivities and constraints in form of sets is an elegant and fast way to state the recognition purpose of the system for multiple objectives.

Activity and Context is not limited to one, synchronized stream of information that can be described and formulated using one single TexTivity predicate. To meet these concerns, we identified two ways of combining TexTivity Predicates for activity and context recognition systems: (i) a logical combination and (ii) a temporal combination. The logical combination can be done e.g., with an *AND* or *OR* connector, allowing to state multiple recognition purposes to the system that are handled in parallel. An example is shown in (4) where the logical combination *AND* is used to combine two TexTivity Predicates.

$$\begin{aligned} < \{Bob\}, ?, \{LivingRoom\}, ?, \{\} > AND \\ < \{Paul\}, ?, \{BathRoom\}, ?, \{\} > \end{aligned} \tag{4}$$

The purpose of the recognition goal in (4) is to detect all Activities and TexTivities that are performed by Bob in the livingroom *AND* all Activities and TexTivities that are performed by Paul in the bathroom. Therefore, two TexTivity Predicates are defined that formulate the recognition goals. These two TexTivity Predicates are then combined using the logic *AND* operation to get the overall recognition goal for the activity and context recognition system.

Allen defined in [14] temporal operators like *Overlaps, Meets, Equal, Before, During, Starts and Finishes*. By using this set of operators the whole space of possible temporal relationships can be represented. Using these operators to combine TexTivity Predicates allows to formulate timing behavior between multiple TexTivity Predicates. We think of recognition goals as shown in (5) and (6). In (5) the purpose of the recognition goal is to detect that Bob brushed his teeth *Before* he was having breakfast that could be an interesting context information for an child education system. In (6) the recognition goal is the detection of Paul reading his emails on Apple's IPad while having lunch in the office, which might also be useful information for systems helping to avoid the burn-out syndrome.

$$\begin{aligned} < \{Bob\}, \{BrushTeeth\}, \{\}, \{\}, \{\} > BEFORE \\ < \{Bob\}, \{Breakfast\}, \{\}, \{\}, \{\} > \end{aligned} \tag{5}$$

$$\begin{aligned} < \{Paul\}, \{ReadingMail\}, \{IPad\}, \{\}, \{\} > \\ DURING \\ < \{Paul\}, \{\}, \{\}, \{OfficeLunch\}, \{\} > \end{aligned} \tag{6}$$

 ISBN: 978-1-61208-257-8

TABLE I: SIMILARITY SCORES OF ONTOLOGICAL RELATIONS TO MATCH ADVERTISEMENTS OF RESOURCES (E.G., SENSOR) TO REQUESTS [15].

Similarity-Score(C_R, C_A)	*Taxonomic-Ontological Relation*	*Computation*
1.0	$C_A \equiv C_R$	1.0
1.0	$C_A \subseteq C_R$	1.0
[0,1]	$C_R \subseteq C_A$	$\frac{\lvert A(C_A)\rvert}{\lvert A(C_R)\rvert}$
[0,1]	$\neg(C_R \sqcap C_A \subseteq \perp)$	$\frac{\lvert A(C_A) \cap A(C_R)\rvert}{\lvert A(C_R)\rvert}$
0.0	$(C_R \sqcap C_A \subseteq \perp)$	0.0

The use of *TexTivity*-Predicates to define the recognition purpose of an activity and context recognition system in form of an semantic abstracted, high level goal, is a flexible way to direct the autonomous and dynamic configuration of such a system. Having the option of combining multiple TexTivity Predicates either with logical or temporal operators makes this approach even more powerful and allows to represent the nature of activities more clearly.

IV. SEMANTIC ENGINE

Using the semantic information modeled in the *TexTivity*-Ontology is the key aspect to utilize the full power and flexibility of the opportunistic approach. To configure sensor ensembles according to the formulated high level recognition goal, we use the semantic information modeled in the *TexTivity Ontology*. Using the modeled *Context*, *Activity* and *TexTivity* relations of a domain allow to reason which Context- and Activity attributes can be substituted by related ones. As the *TexTivity Ontology* models the concepts and relations of a domain, each possible sensor output in terms of labels is an entry point in the modeled domain knowledge. The labels that can be delivered by sensors are the connection point between the *"semantic"-sensors*, the *recognition goal* and the *TexTivity Ontology*. This allows the combination of sensors according to the defined relations in the *TexTivity Ontology* and to reason knowledge at a semantic level. If there is, for example, no sensor in the environment that can be explicitly used to detect a stated recognition goal, we can search for sensors that can be used instead and therefore substituting the not available sensor entity. According to the defined semantic relations in the *TexTivity Ontology*, we can reason how the formulated recognition goal can be substituted by related contextual information. In [15] Bandara et al. propose a way based on the semantic similarity concepts presented by Lin [16] on how to match the advertisements of a resource (e.g., a (self-described) sensor) to a request (e.g., the recognition goal respectively the recognition purpose). The matching scores between the concepts are shown in Table I.

To find the best sensor ensemble according to the stated recognition goal, we define a *Goal Function* (as shown in (7)) that uses the semantic similarity concepts presented in [16] and matches the requested capabilities of the recognition goal to the advertised recognition capabilities of the sensors (as described conceptual in Table I). The function weights the resulting *DoF* according to the number of labels (and their single DoFs) that can be satisfied by the sensor ensemble and returns a value between [0,1] that reaches a maximum for the best set of sensors.

$$DoF = \frac{1}{LtR} \sum_{i=0}^{n-1} DoF_i \tag{7}$$

LtR=number of *Labels* the sensor is queried for
n=number of *Labels* satisfied by the sensor ($n \leq LtR$)
$\boldsymbol{DoF_i}$=DoF of a single label with index i

In case where no sensor is found in the ecosystem of the user that can explicitly be used to detect the stated recognition goal, *Goal Substitution* and *Reasoning* using the captured *Domain Knowledge Relations* takes place. By knowing the semantic relations of context and activities, we can reason which sub-context can be used to reason context at a higher semantic level. This is an extremely powerful approach as we can use sensors that were not explicitly designed to detect the stated recognition purpose. An example is the recognition goal <{},{*Locomotion*},{},{},{}> that formulates to detect *Locomotion*. If we can not find a sensor in the environment, that can explicitly be used to detect *Locomotion* we use the semantic information modeled in the TexTivity Ontology. We refine *Locomotion* into its related sub-activities ($\Re_{has\_A}$) $\rightarrow$ *Walk, Stand, Lie* and *Sit* as shown in Fig. 1 and search for sensors that are capable of recognizing these activities. We rank the found sensors according to the *Goal Function* defined in (7) and select the highest ranked as ensemble. The example highlights, the use of sensors to detect *Walk, Stand, Lie* and *Sit* that were not designed to detect *Locomotion* explicitly. Only because the *semantic relations* of *Locomotion* to these four activities are known, and modeled in the *TexTivity Ontology*, we can combine them to detect *Locomotion*.

As the reasoning capabilities are not limited to one refinement step $\{A \rightarrow \{B, C\}\}$, the refinement of activities and context may involve the recursive evaluation of an entire tree of related context $\{A \rightarrow \{B, C\}; B \rightarrow \{U, V, W\}, C \rightarrow \{X, Y\}\} \rightarrow \{\ldots\}$. This can take as many steps as necessary to decide if the formulated recognition purpose is satisfiable or not. Furthermore, reasoning of a recognition goal, expressed in form of a TexTivity Predicate, is a reasoning in multiple dimensions as it can be done for all elements contained in the TexTivity Predicate as exemplarily shown in Fig. 2. The two presented methodologies, (i) the *semantic matchmaking* of the recognition goal expressed using TexTivity Predicates to the sensing infrastructure and (ii) the *goal splitting*, *substitution* and *quantification* according to the defined *Goal Function* (7) are implemented in a *Semantic Engine* that handles the *Semantic Concept Matching* between the sensing ecosystem and the stated recognition goal as shown in Figure 3. In the following Section V we evaluate the presented methodologies using a reference implementation of an opportunistic activity and context recognition system called the OPPORTUNITY Framework [3].

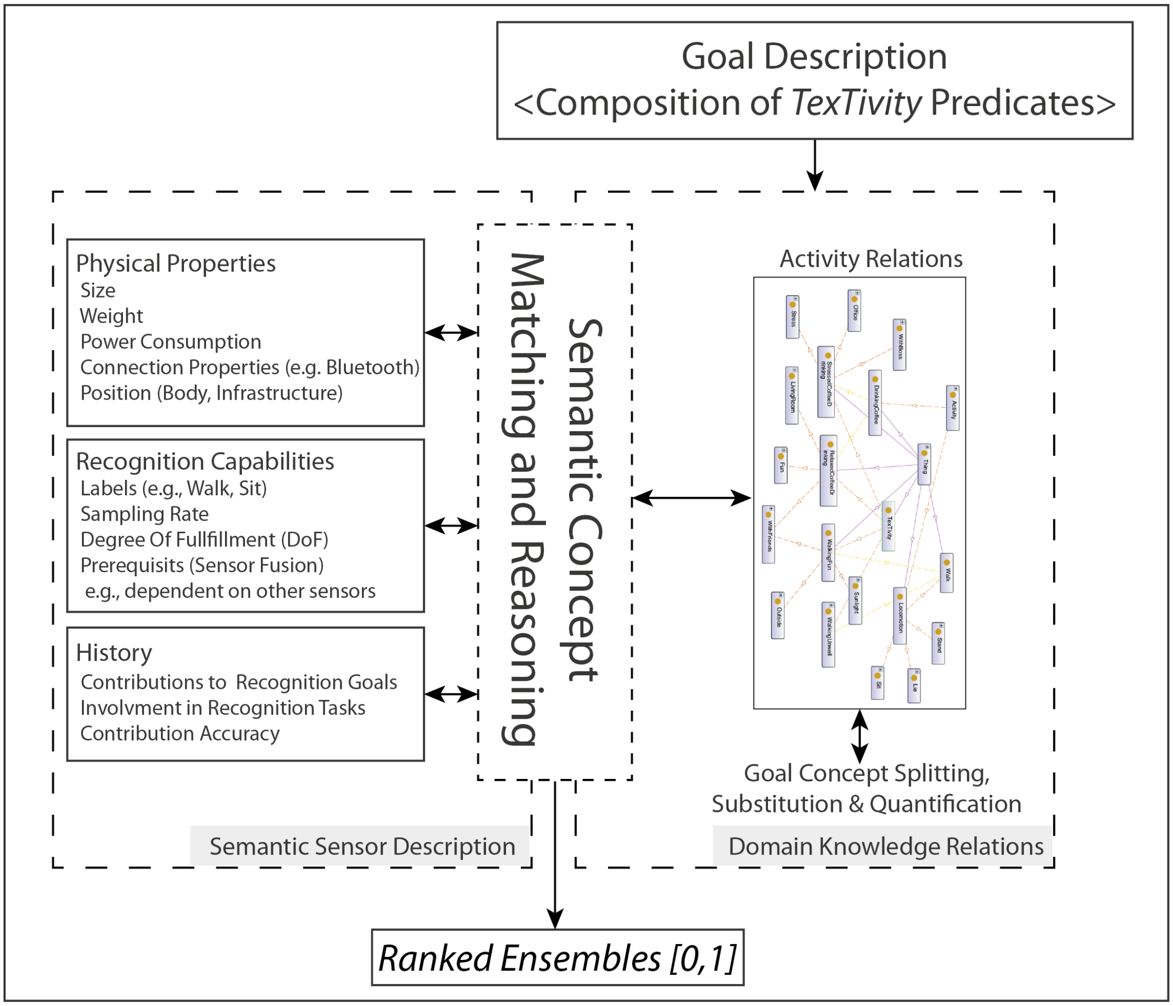

Fig. 3: Semantic Engine used for Processing the stated Recognition Goal in mapping, reasoning and substituting it to the available sensors in the ecosystem of the user.

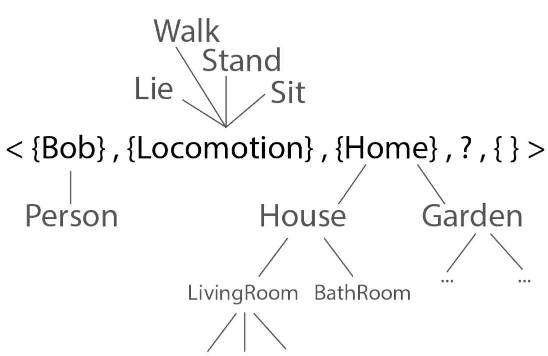

Fig. 2: Reasoning of TexTivity Predicates in multiple dimensions for the recognition goal <{Bob},{Locomotion},{Home},?,{}> according to the defined relations in the TexTivity Ontology.

V. Evaluation

To evaluate the approach of using high level recognition goals to direct the dynamic configuration of an activity and context aware system, we developed the OPPORTUNITY Framework [3]. The framework is a reference implementation of an opportunistic context and activity recognition system. It allows to state high level recognition goals during runtime in form of TexTivity predicates that direct the dynamic and autonomous configuration of the system. It uses the goal processing, reasoning and semantic matchmaking capabilities of the *Semantic Engine* described in Sections III and IV to configure sensor ensembles to fulfill the stated high level recognition goal. A similar setup was already successfully used by our group to test the configuration of sensor ensembles [17]. For the evaluation we set up a real time scenario with body worn sensors to test and evaluate the presented approach. We picked a rather easy activity set for evaluating, as the goal is not to work with highly sophisticated and complex activity classes. The handling of complex activities using dynamically configured Hidden Markov Models with a similar sensor setup is presented in [17]. For testing we used the activity classes of the *modes of locomotion* $\rightarrow$ (Walk, Sit, Stand, Lie). We used four XSense-Mti-sensors that deliver

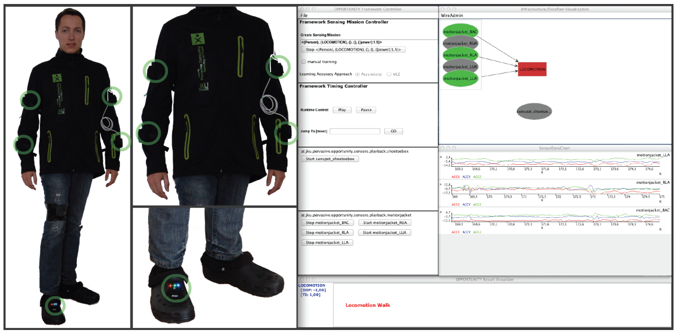

Fig. 4: Experimental setup of the on-body sensors to test and evaluate the autonomous configuration of an Activity and Context recognition system in a realtime setting.

triaxial acceleration data. The sensors were mounted on the left and right upper-/lower arm and are named according to their position *motionjacket_RUA* (right-upper-arm), *motionjacket_RLA* (right-lower-arm), *motionjacket_LUA* (left-upper-arm) and *motionjacket_LLA* (left-lower-arm). Additionally, one SunSPOT accelerometer sensor was attached to the right shoe (*sunspot_shoetoebox*). Fig. 4 shows the sensor setup and the corresponding realtime schematic of the runtime configured sensor ensembles using the OPPORTUNITY Framework. The four *motionjacket* sensors were trained to detect *Walk*, *Sit*, *Stand*, and *Lie* using the OPPORTUNITY Dataset [11]. The *sunspot_shoetoebox* was trained to detect *Locomotion*. In the following two scenarios the TexTivity predicate <{},{*Locomotion*},{},{}, {} >, formulating the *recognition goal* to detect the activity *Locomotion*, was stated to the system. In the first scenario, all sensors were available (Fig. 5a) as indicated by the green bubbles. The OPPORTUNITY Framework queried the available sensors for their recognition capabilities. One sensor, the *sunspot_shoetoebox*, was found that can explicitly be used to recognize *Locomotion*. The goal refinement process stopped as a sensor was found that is explicitly dedicated to detect Locomotion. So this sensor is selected as ensemble (as indicated by the black arrow), the corresponding recognition chain is dynamically instantiated, and the recognition process is initiated.

In the second scenario, we turned off the *sunspot_shoetoebox* sensor as indicated by the grey bubble (Fig. 5b). The four motionjacket sensors were still available. In this case, no sensor is available that can explicitly be used to detect *Locomotion*, as the motionjacket sensors are trained to detect *Walk, Sit, Stand* and *Lie*. Again, the OPPORTUNITY Framework queried the available sensors for their recognition capabilities. As no sensors were found by the framework that can explicitly be used to detect *Locomotion*, the *Semantic Engine* initialized the *Goal Refinement* process according to the semantically modeled relations of the TexTivity Ontology. The OPPORTUNITY Framework autonomously reasoned the related activities of *Locomotion* ($\Re_{has\_A} \rightarrow$) {*Walk, Stand, Lie, Sit*} (Fig. 1). Knowing the related activities, the framework queried for sensors that can detect the related activities. Four sensors, the *motionjacket_RUA*, *motionjacket_RLA*, *motionjacket_LUA* and *motionjacket_LLA* were found that can be used. The *Goal Function* (7) was evaluated for each sensor to quantify its capabilities according to the reasoned recognition goal as shown in Table II. After evaluating the *Goal Function*, the four *motionjacket*-sensors were configured to an ensemble as the *Condorcet's jury theorem* [18] states that if the probability of a single vote is grater than 0.5 (*DoF/Score*), then the probability that the majority decision is correct is increased. The corresponding recognition chains of the single sensors are dynamically instantiated and are fused using *majority voting* (Fig. 5b). Out of the four delivered labels (*Walk, Stand, Lie, Sit*) we can reason their relation on a higher semantic level inferring *Locomotion* using the modeled activity relations in the *TexTivity Ontology* that are utilized by the *Semantic Engine*. The evaluation scenario showed that it is possible to combine sensors according to semantically modeled information in a realtime system. We showed that the system can autonomously configure and adapt an activity and context recognition system dependent on the available sensing infrastructure and the stated recognition goal during runtime. Even if no sensors are available to detect the stated recognition goal explicitly, the system autonomously utilizes the semantically modeled information in the *Semantic Engine* in two ways. First, in a top-down fashion, it queries for related activity and context according to the stated *Recognition Goal* and then it quantifies sensors using a Goal Function (as shown in (7)) that can be used to detect them (e.g., *Locomotion* $\rightarrow${*Walk, Stand, Lie, Sit*}). Second,

TABLE II: SEMANTIC SCORES FOR THE SEMANTICALLY REASONED RECOGNITION GOAL "LOCOMOTION" ACCORDING TO THE GOAL FUNCTION (7).

Sensor	*Walk*	*Stand*	*Lie*	*Sit*	*Score*
motionjacket_RUA	0.700	0.770	0.980	0.860	0.828
motionjacket_RLA	0.690	0.380	0.900	0.950	0.730
motionjacket_LUA	0.780	0.670	0.970	0.930	0.838
motionjacket_LLA	0.790	0.320	0.480	0.730	0.580

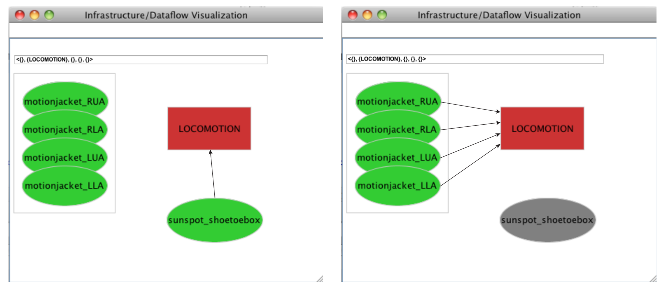

(a) (b)

Fig. 5: Visual, real-time representation of the physical sensors used in the scenario to evaluate the *Goal Description and Semantic Matchmaking* in a varying sensor ecosystem.

the system uses the semantically modeled information in a bottom-up fashion, to reason context at a higher semantic level out of the semantic data delivered by the sensors (e.g., {*Walk, Stand, Lie, Sit*}→*Locomotion*).

VI. Conclusion

Sensors providing information to detect human activities are embedded in more and more artifacts of everyday living (e.g., smartphones, watches, cameras, clothing, etc.). Using these artifacts motivates the shift away from deploying "application specific" sensors and to support developers in relieving them from low level, platform specific concerns. We use sensors that just happen to be available and utilize them in an goal oriented, opportunistic way for activity and context recognition. We can handle and exploit heterogeneous resources using *Sensor Abstractions* and dynamically instantiate the necessary stages of the *Activity Recognition Chain* to infer activity and context information out of the raw sensor data readings using *ExperienceItems*. We see each sensor as a label delivering entity that can be queried for its capabilities.

As central contribution we propose a *Semantic Engine* that matches the stated recognition goal in form of *TexTivity Predicates* to the available, semantically described sensor ecosystem. Therefore it utilizes modeled *context* and *activity relations* in the *TexTivity Ontology* that can be reasoned and inferred to combine sensors at different semantic levels. To direct the autonomous and adaptive configuration, we use a goal-oriented approach. The goal encapsulates a high level directive that defines, at a semantic level, what the system should do and how it should behave. As syntactical formalism, to describe and formulate the recognition goal we propose the use of *TexTivity-Predicates* as a precise, flexible, and reasonable way of formulating the recognition purpose of the Activity and Context recognition system.

The autonomous and adaptable combination and use of sensors at a semantic level provide the possibility to use sensors that are already out there in the environment and eliminates the need to deploy more and more application specific sensors. The fact that the system can react on changes in the sensing infrastructure (e.g., due to sensor failures) increases its stability and robustness compared to traditional systems. An essential point for the scalability of the opportunistic sensing approach is the fact that sensors, processing, and communication resources are only used and configured to ensembles, if they are needed to fulfill the stated recognition goal.

The evaluation clearly highlights that the research hypotheses holds that a goal oriented, top-down configuration approach for autonomously adapting an activity and context recognition system during runtime offers a high flexibility in terms of combining sensors at a semantical level resulting in a stable, and accurate activity and context recognition system envisioning the autonomous and dynamic creation of sensing infrastructures of massive scale during runtime.

Acknowledgment

The project OPPORTUNITY acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET-Open grant number: 225938.

The project PowerIT acknowledges the financial support of the FFG FIT-IT under grant number: 830.605.

References

[1] L. Benini, E. Farella, and C. Guiducci, "Wireless sensor networks: Enabling technology for ambient intelligence," *Microelectron. J.*, vol. 37, no. 12, pp. 1639–1649, 2006.

[2] D. Roggen *et al.*, "Opportunity: Towards opportunistic activity and context recognition systems," in *Proceedings of the 3rd IEEE WoWMoM Workshop on Autonomic and Opportunistic Communications (AOC 2009)*. Kos, Greece: IEEE CS Press, June 2009, pp. 1–6.

[3] M. Kurz *et al.*, "The opportunity framework and data processing ecosystem for opportunistic activity and context recognition," *International Journal of Sensors, Wireless Communications and Control, Special Issue on Autonomic and Opportunistic Communications*, pp. 102–125, December 2011.

[4] D. Roggen, K. Förster, A. Calatroni, and G. Tröster, "The adarc pattern analysis architecture for adaptive human activity recognition systems," *Journal of Ambient Intelligence and Humanized Computing*, pp. 1–18, 2011.

[5] J. O. Kephart and D. M. Chess, "The vision of autonomic computing," *Computer*, vol. 36, no. 1, pp. 41–50, Jan. 2003.

[6] T. Ellinger, G. Beuermann, and R. Leisten, *Operations Research*. Springer, 2003.

[7] A. Van Lamsweerde, "Goal-oriented requirements engineering: A guided tour," in *Proceedings of the Fifth IEEE International Symposium on Requirements Engineering*, ser. RE '01. Washington, DC, USA: IEEE Computer Society, 2001, pp. 249–263.

[8] A. S. Rao and M. P. Georgeff, "BDI-agents: from theory to practice," in *Proceedings of the First Intl. Conf. on Multiagent Systems*, San Francisco, 1995, pp. 312–319.

[9] P. Turaga, R. Chellappa, V. S. Subrahmanian, and O. Udrea, "Machine Recognition of Human Activities: A Survey," *Circuits and Systems for Video Technology, IEEE Transactions on*, vol. 18, no. 11, pp. 1473–1488, Sep. 2008.

[10] X. Hong, C. Nugent, M. Mulvenna, S. McClean, B. Scotney, and S. Devlin, "Evidential fusion of sensor data for activity recognition in smart homes," *Pervasive Mob. Comput.*, vol. 5, pp. 236–252, June 2009.

[11] D. Roggen *et al.*, "Collecting complex activity data sets in highly rich networked sensor environments," in *Proceedings of the Seventh International Conference on Networked Sensing Systems (INSS), Kassel, Germany*. IEEE Computer Society Press, June 2010, pp. 233–240.

[12] A. K. Dey, "Understanding and using context," *Personal and Ubiquitous Computing*, vol. 5, pp. 4–7, 2001.

[13] J. Ye, G. Stevenson, and S. Dobson, "A top-level ontology for smart environments," *Pervasive Mob. Comput.*, vol. 7, pp. 359–378, June 2011.

[14] J. F. Allen, "Maintaining knowledge about temporal intervals," *Commun. ACM*, vol. 26, pp. 832–843, Nov. 1983.

[15] A. Bandara, T. Payne, D. De Roure, N. Gibbins, and T. Lewis, "A pragmatic approach for the semantic description and matching of pervasive resources," in *Proceedings of the 3rd international conference on Advances in grid and pervasive computing*, ser. GPC'08, Springer Berlin, Heidelberg, 2008, pp. 434–446.

[16] D. Lin, "An information-theoretic definition of similarity," in *Proceedings of the Fifteenth International Conference on Machine Learning*, ser. ICML '98. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1998, pp. 296–304.

[17] G. Hölzl, M. Kurz, and A. Ferscha, "Goal oriented opportunistic recognition of high-level composed activities using dynamically configured hidden markov models," in *The 3rd International Conference on Ambient Systems, Networks and Technologies (ANT2012)*, August 2012, pp. 308–315.

[18] M. Condorcet, *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix.* Paris: Imprimerie Royale, 1785.