Pervasive and Mobile Computing 8 (2012) 597–613

Contents lists available at SciVerse ScienceDirect

Pervasive and Mobile Computing

journal homepage: www.elsevier.com/locate/pmc

Egocentric interaction as a tool for designing ambient ecologies—The case of the easy ADL ecology

Dipak Surie [a,*], Lars-Erik Janlert [a], Thomas Pederson [b], Dilip Roy [a]

[a] *Department of Computing Science, Umeå University, S-90187 Umeå, Sweden*

[b] *Innovative Communication Group, IT University of Copenhagen, Rued Langgaards Vej 7, 2300 Copenhagen S, Denmark*

ARTICLE INFO

Article history:
Available online 8 January 2012

Keywords:
Ambient intelligence
Ambient ecology
Context-aware computing
Mixed-reality
Human–computer interaction

ABSTRACT

The visions of ambient intelligence demand novel interaction paradigms that enable designers and system developers to frame and manage the dynamic and complex interaction between humans and environments populated with physical (real) and virtual (digital) objects of interest. So far, many proposed approaches have adhered to a device-centric stance when including virtual objects into the ambient ecology; a stance inherited from existing interaction paradigms for mobile and stationary interactive devices. In this article, we introduce *egocentric interaction* as an alternative approach, taking the human agent's body and mind as the center of reference. We show how this interaction paradigm has influenced both the conception and implementation of the easy ADL ecology, comprising of smart objects, a personal activity-centric middleware attempting to simplify interaction given available resources, ambient intelligence applications aimed at everyday activity support, and a human agent literally in the middle of it all.

1. Introduction

With the advancement in computing, communication, sensing, actuation, interface and interaction technologies, the boundaries between physical and virtual (or computing) worlds are becoming transparent with the possibilities of seamless integration of the two worlds. Ambient intelligence [1] refers to the vision of integrating computational intelligence within human environments and the artifacts that it contains, such that, human-centered services are offered to satisfy human agents' immediate needs. Ambient ecology [2] could be considered as the infrastructure through which ambient intelligence can be realized. It comprises of an inter-connected collection of heterogeneous components like smart objects [3], middleware components, and ambient intelligence applications, virtual objects that are part of such applications and are accessed through the smart objects, artificial agents, and human agents with a collective goal of supporting human agents' activities, lifestyle and well-being.

The visions of ambient intelligence introduce several challenges in framing and managing human interaction with computing systems. The computing systems are expected to be context-aware [4], knowing a human agent's current situation, activity, and interaction context especially in the physical world in facilitating human–computer interaction. Traditional interaction approaches like the WIMP (windows, icons, menus and pointing devices) interaction paradigm is not aware of a human agent's physical context. It offers interaction through limited information channels and is overall obsolete for ambient ecologies. Several approaches like tangible user interfaces [5], ambient displays [6], surface computing [7], etc. are device-centric, while approaches like perceptual user interfaces [8], attentive user interfaces [9], affective user

* Corresponding author. Tel.: +46 704396048.
E-mail address: dipak@cs.umu.se (D. Surie).

1574-1192/$ – see front matter
doi:10.1016/j.pmcj.2011.12.004

interfaces [10], activity-based computing [11], etc. focus on specific aspects of a human agent. Embodied interaction [12] focuses on human embodiment in an environment in facilitating interaction while reality-based interaction [13] provides a comprehensive framework for emerging trends in human–computer interaction inspired from the functioning of the real world. While such approaches are valuable from an ambient ecology context, they do not consider a human agent as a whole within an ambient ecology. This conceptual article introduces *egocentric interaction* as an alternative approach toward the modeling of ambient ecologies with the unique feature of taking the human agent's body and mind as center of reference, as opposed to more common device-centric approaches in facilitating human–environment interaction.

A first assessment of the viability of the concepts is made through our experience in building a prototypical ambient ecology, named the easy ADL ecology, based on egocentric interaction principles. Outline of the article: Section 2 describes egocentric interaction and its principles; Section 3 presents the situative space model; Section 4 describes the easy ADL ecology comprising of smart objects, a personal activity-centric middleware, and ambient intelligence applications providing everyday activity support; Section 5 outlines the interaction management rules and techniques for facilitating human interaction within the ambient ecology; and Section 6, finally, makes some brief concluding remarks.

2. Principles of egocentric interaction

The easy ADL ecology to be presented in Section 4 is designed and implemented based on the principles of egocentric interaction. There are many research efforts that can be seen as working toward a new way of understanding and framing interaction between a human agent and their environment. The reason for this interest and direction in the research society is not hard to see. Important new advances in technology — communication, sensor, presentation, and actuation technology, together with the "ordinary" advances in computational power, storage and communication capacity, and cost efficiency — now allow the formerly chair ridden "user" to stand up, use the body and move around as a free agent in the real world.

The notion of "ambient ecologies" is a good example in itself, taking inspiration from nature's ability of maintaining a balanced functioning system of entities while still being able to adapt to external changes: something every ambitious system engineer is looking to achieve. Relevant research ranges from very theoretical and general with theories from psychology, philosophy and cognitive science (e.g. activity theory [14], ecological psychology [15], distributed cognition [16], situated cognition [17], extended mind hypothesis [18], embodied cognition [19], and more), to research that is practically oriented and in particular tries to apply and adapt such theories to human–computer interaction (e.g. reality-based interaction [13], embodied interaction [12], activity-based computing [11], augmented reality [20], context-aware computing [4]). Egocentric interaction is one such research effort that takes inspiration from psychology, philosophy and cognitive science in forming a set of concepts to describe a human agent's interaction in an environment, in ways useful for facilitating human interaction with computing systems.

The term "egocentric" has been chosen to signal that it is the human body and mind of a specific human individual that (sometimes literally, as will be shown later) acts as center of reference to which all interaction modeling is anchored. In the context of this article, the term should not be taken to refer to selfishness or other personality traits but to the basic human condition of dealing with the world based on their own senses, actions and cognitive abilities, even when working in groups and with shared goals.

What follows is a description of some basic principles and assumptions that underlie egocentric interaction.

2.1. Terminological shift from "user" to "agent"

The perspective in human–computer interaction research is shifting toward conceiving human beings as mobile agents in a dynamically changing environment populated by physical and virtual objects alike [21], instead of conceiving them as a "user" performing a dialog with a computer. Within egocentric interaction, the term "human agent" is used instead of the term "user" for several reasons. Egocentric interaction models interaction occurring between a specific human being and potentially several interactive devices at the same time, as well as the interaction between the human being and the physical environment. In such a wide-ranging variety of interactive situations, it is more natural to regard the modeled individual as an *agent* in a mixed-reality environment [21]. This shift in terminology, from "user" to "agent" is also a consequence of letting go of "the application" as the tacit focus: applications are "used" by someone and the "user" is defined by the artifact, whereas agents do actions and are engaged in activities that typically involve a number of objects. The notion of "agent" also implies that not all human actions need to be (explicitly or implicitly) directed toward a computing system. Egocentric interaction is imagined to be part of activity scenarios where only parts of the activity involve events that a computer system can or should care for. Compare the notions of "in-band" and "out-of-band" in [5] for distinguishing between events that are part of a human–computer dialog and those that are not.

2.2. Merging virtuality, mobility and ubiquity paradigms

Egocentric interaction could be considered as a novel interaction paradigm where the term "paradigm" is used in a sense inspired by Thomas Kuhn's influential notion of scientific paradigms [22]. In analogy with scientific paradigms, we understand an *interaction paradigm* to include components such as: important design examples and use scenarios, important techniques and technologies, key problems and challenges, articulations of ideals and goals to pursue, interpretations of key concepts, such as 'user,' 'interface,' 'interaction,' and finally, groups or communities of people (researchers and interaction designers) developing and defending the paradigm. Unlike Kuhn's paradigms, different interaction paradigms may peacefully coexist; an older paradigm may find niches to survive the rise of a new dominating paradigm.

From an ambient ecology perspective, there are three major interaction paradigms in operation (or *use*paradigms as referred to in [23]). In the *virtuality*paradigm, users interact with virtual objects, i.e., data objects, which are accessed through an interface to a virtual, symbolic world, usually with very little relation to the particular physical, real, use situation. In the *ubiquity*paradigm, human agents interact with multiple real objects augmented and interfaced by computer technology, which pervade the real world as researched in the field of ubiquitous computing. In the *mobility*paradigm, human agents move through the real world, while staying in contact and interacting with the virtual world through mobile and wearable devices. In some mobility scenarios, human agents are using mobility primarily for remote access and operation, *independent* of the situation while in other mobility scenarios, mobility is used for in situ application, *dependent* of the situation; i.e., bringing the computational and informational resources to bear directly on the very situation of use in the real world [23]. Egocentric interaction is in effect merging the virtuality paradigm and its preoccupation with virtual objects with the mobility paradigm's emphasis on movement and the human agent as a moving point-of-view in the physical world, and with the ubiquity paradigm's acknowledgment of the importance of physical objects and the richness of the physical environment.

2.3. Different from classical user-centered approaches

Egocentric interaction is centered on the "user", or the (human) *agent*, not only in the sense that interaction design should bow to the user's needs and preferences, should ultimately satisfy the user, but in the sense that the agent should be understood as being in the (moving) center of the world. This distinguishes egocentric interaction from earlier "user-centered" approaches in HCI such as user-centered design [24] which largely ignored an agent's current bodily situation (which is fixed by the stationary desktop computer), egocentric interaction acknowledges the primacy of the agent's current bodily situation in the environment at each point in time in guiding and constraining the agent's behavior (which is precisely what "egocentric" refers to). It assumes that the whole environment is taken into consideration, not just a single targeted artifact or system thereby facilitating human–environment interaction within an ambient ecology. It makes the assumption that proximity plays a fundamental role in determining what can be done, what events signify, and what the agent is up to. It seriously considers mobility and takes into account the agent's more or less constant movements of head, arms, hands, and body, locally and through the environment, as well as the agent's constant rearrangements and modifications of various parts of the environment (especially the manipulation of objects). It is neither mainly oriented toward the interaction with "virtual", immaterial data objects (classical HCI), nor predominantly toward the interaction with physical objects and machines (classical ergonomics and HMI), but pays equal attention to virtual and physical objects, circumstances, and agents, and their interrelations. It recognizes that the agent typically will have multiple ongoing activities within an ambient ecology, often in parallel, some of which may have little relation to each other, and that activities are started, put on hold, interrupted, resumed, and finished in a never ceasing flow. It makes no pretense and no tacit assumption that the ambient ecology or the agent has full information of the situation, because situations are in principle open-ended and ever changing.

It should be noted that there were and still are user-centered approaches to interaction design that do take some (but never quite all) of these factors into account, a prime example being participatory design [25]. However, these approaches are focused on methods and ideals for the design and development *process*, and on basically work-oriented and fixed, well-defined and stable tailor-made systems while egocentric interaction focuses on the design and modeling process of ambient ecologies which is beyond work-oriented systems.

Also, interaction paradigms that are centered on interaction devices will have to change frequently to keep pace with new technological advancements, whereas interaction paradigms that are human-centric like egocentric interaction are more stable over time: basic human characteristics change very slowly.

2.4. Situatedness: the view from here

The following is a commonsensical summary of the *situatedness* of agents with regard to interactions and activities. Human beings have a self-centered perspective on their environment before any other perspective, such as second person perspective or subspecie-aeternitatis perspective that science aspires to ("the view from nowhere" [26]). The world is so large and so rich in details that any agent with limited cognitive capacity must necessarily narrow its focus in some manner. Human beings have physical bodies that are located at a single particular place and oriented together with their limbs and sense organs in particular directions at any particular time. That gives them a natural primary vantage point for selecting which details and aspects to attend to: *the view from here*, what a particular embodied agent can perceive given its current bodily situation in the world. It also gives them a focus for actions: anything they can do is limited by what is within their reach given their current bodily situation in the world.

To control one's actions and the effects of their actions it is particularly helpful to be able to perceive the part of the environment in which that person is acting and affecting. Vice versa, to perceive, to pick up information from a certain part of the environment it is helpful to be able to do actions that change one's angle of perception and the parts that person currently perceive because that can assist and improve the information extraction. *Natural* agents are usually concerned about aspects and details of the world that have relevance for themselves, and naturally perform actions that are relevant for their own existence and relation to the environment.

Human agents are where they are, and that determines what they can do. If they want to do something that cannot be done right here right now, they will have to do something about it, and they will always have to do *that* starting from their very situation here and now. By doing things with what is within their reach, by moving themselves in the world, and by moving other objects in the world, they can change what they can do. The further removed from the current situation what they want to do is, the longer chains of action they need to be able to contemplate and follow to reach their goal. To their advantage, at each point in a long chain of actions, their current situation in the world, what they currently perceive, can help them recall and guide their next action; they may also be able to specifically arrange the environment around them and their relation to it so that their situational guidance through the process is improved.

The trajectory of a human body through space–time and the trail of environmental events and changes in its close vicinity may reveal a lot about which goals and intentions are on this human's mind, and which activities are being carried out. *Epistemic operations.*, i.e., physical operations performed to facilitate cognition rather than to further physical progress toward some external goal [27], may offer particularly transparent access into the mind of the agent. Epistemic operations make cognition transparent by instantiating (partially) external cognition in the sense of the extended-mind hypothesis [18]. They are transparent also in the practical sense of being within the scope of current tracking technology, just as other physical operations are.

The view-from-here perspective means that all information about the environment that is picked up by the human agent derives from information converging on the body and modulated by the point of view of the body wherever it is situated; and that all actions, all effects on the environment produced by the human agent have their starting point in the body and are modulated by the angle of approach of the body wherever it is situated. Most human activities involve the handling of a number of objects and actions to make them available (locomotion, object /transportation) in certain sequences and patterns that are appropriate for the objective of the activity. Human agents involved in a particular activity are at each moment guided in their actions by the array of objects in their immediate proximity, at the same time as some of their actions serve to or has as a side effect to change the array of objects in their immediate proximity. We may count on a general human tendency to minimize the effort spent, physical as well as cognitive. At any point in time, objects that are close are likely to be or become relevant for an ongoing or near future activity and objects that are relevant for an ongoing activity are likely to be or become close.

2.5. The proximity principle

Much of the enabling as well as limiting consequences of situatedness within an ambient ecology can be captured in the *proximity principle* [28]: *Things that are close tend to matter; things that matter tend to be(come) close*. Objects (and other agents, circumstances) that are close to the agent tend to matter in the sense that they have a fair chance of getting the agent's attention and figure in the agent's current cognitive processes and current activities. Sometimes such objects also trigger action, and the start of a new activity or the resumption of an activity currently on hold that directly involves the objects in case; and the closer the objects are the more likely they will play such a role, other things being equal.

Objects (and other agents, circumstances) that matter to the agent's activities will tend to either already be or soon become within close range; and the more imminent their use is, the closer they will tend to be. The rationale is obviously that if an object is needed in the activity the agent is likely to make sure that it is at hand in order to proceed with the activity, either by moving up to the object or by having moved the object to the current location in preparation (short term or long term). When a number of objects are needed simultaneously or in swift succession, we may consequently expect to observe prearranged environments for the purposed activities, as well as elaborated logistic strategies for moving objects around.

2.6. Perception and action instead of input and output

One of the strengths of direct manipulation mechanism [29] widely deployed in user interfaces for changing the state of *virtual* objects, is the fact that it makes the strong relationship between manipulation and perception prevail also in many parts of the virtual world. The classical HCI concepts of input and output need to be substituted with something that works both for physical and virtual object manipulation common in an ambient ecology. Also, the concepts of input and output are device-centric and refer to device interfaces, while egocentric interaction uses concepts that refer to an agent's interface to the ambient ecology (i.e., an *agent's* input and output). Hence, the concepts of perception and action replace the traditional concepts of input and output. Action is often inseparable from or intertwined with perception in everyday physical activities. One cannot change the state of an object (e.g., open a refrigerator door) without perceiving feedback (e.g., tactile feedback while opening the refrigerator door). Conversely, perception needs support of action (e.g., to see what is in the refrigerator, the door must be opened). This tight coupling between action and perception is the result of the way objects are designed, how we manipulate them, and how the manipulation process interplays with the laws of physics. The traditional concepts of input and output, stemming from a time when interaction typically was in terms of exchanges of language expressions, still have a bias toward a turn-taking approach where first an input is provided to a device that processes it and provides an output, whereas the concepts of perception and action are inter-coupled and take place in parallel, which is suitable for an ambient ecology context. One activity may typically involve several distinct sources (for perception) and targets (for action) at the same time, and multiple activities are often going on in parallel, more or less independent of each other.

2.7. The physical–virtual equity principle

Egocentric interaction differs from traditional interaction paradigms by explicitly ignoring the input and output devices (e.g. keyboard and display of a mobile phone or a laptop), considering them as transparent (both physically and cognitively) *mediators*for accessing virtual objects within an ambient ecology. Mediators include the sensors, actuators, input devices, output devices, user-interface software, recognition algorithms, etc., that are part of smart objects (refer to Section 4.1) and serve as tools for providing human agents access to virtual objects. Taking such a stance permits the modeling of physical and virtual objects as if they were situated in the same *physical space* (also referred to as the situative space), which is advantageous within an ambient ecology where the interaction complexity vastly surpasses what can be sufficiently described using a classical human–computer interaction dialog model.

Even though there are inherent differences between physical and virtual objects [21], the goal is *not* to make them resemble each other as much as possible. Designing virtual objects as exact copies of physical objects removes the inherent advantages of being virtual (or digital), such as being easier and cheaper to transport across space compared to transporting physical objects. Egocentric interaction proposes an approach where physical objects and virtual objects are co-located and complement each other with their inherent properties.

The point with physical–virtual equity is to handle physical and virtual objects uniformly on a high level of abstraction in order to: (a) enable seamless integration of the physical and the virtual aspects of an ambient ecology by introducing concepts like physical–virtual artifacts [21]; and (b) enable better modeling of the physical and the virtual aspects of the ambient ecology like a human agent's physical–virtual situations and their physical–virtual activities [21]. Such a view aligns well with findings in psychology indicating that expert users of tools (whether a frying pan or a computing device) tend to focus on the domain object they are working on (whether it is a vegetable or an email) and become less aware of the details of the tool (such as the door of the microwave oven or the workings of a particular input device) itself.

Traditional interaction paradigms associate virtual objects to specific devices, while egocentric interaction allows for dynamic association of virtual objects to smart objects depending on the human agent's context in the ambient ecology.

2.8. Summary of basic principles and assumptions

To summarize, the basic principles and assumptions of egocentric interaction are as follows: situatedness; attention to the local environment; the proximity principle; changeability of environment and agent–environment relationship; the physical–virtual equity principle; concept of perception and action; multiple concurrent activities; and incomplete information and control.

From the assumption of situatedness, several principles have emerged like the proximity principle and human attention to the local environment. Human agents attend to their complete local environment even when interacting with a specific physical object; however they do so with varying levels of attention. Agents perform actions and engage in activities motivated and guided by what matters to them. The proximity principle has something to say about locality and how what matters is related to the relation between the agent and the environment, and so about what drives actions and movements. From that follows, among other things, that agents do manipulate objects and rearrange the environment in the pursuit of an ongoing activity, and that agents move to get close to objects (and places, and other agents) that are needed for (an upcoming phase of) some intended ongoing or upcoming activity. That accounts for some of the changeability of the environment and the relation of the agent to the environment; other agents are another source of change.

The principle of physical–virtual equity is motivated pragmatically: our everyday interactions with virtual objects are already considerable and keep increasing. From the physical–virtual equity principle and the important fact that virtual objects are so much more mobile than physical objects (including the agent's own body), we should expect an increase in the number of concurrent activities of an agent: while moving toward a physical object is needed for one activity, it becomes possible to use the transportation time to make progress with other activities that in their current stage rely on virtual objects and actions; similarly, while waiting for the completion of an action directed toward some object involved in one activity, the waiting time can be used for virtual handling of other activities.

The incompleteness of information and control, finally, implying an element of improvisation that has been absent or suppressed in earlier paradigms, follows from the changeability of the environment and the agent–environment relationship, and from the richness of the real world.

3. A situative space model

When applying the general principles of egocentric interaction in the design of the ambient ecology, we chose to concretize and synthesize some of the principles into a *situative space model,* which could immediately be implemented into the personal activity-centric middleware. The situative space model is described in detail elsewhere [30] and only a summarized account of it is provided in this article, enough to explain the design rationale behind the development of the smart objects (described in Section 4.1) and the personal activity-centric middleware (described in Section 4.2).

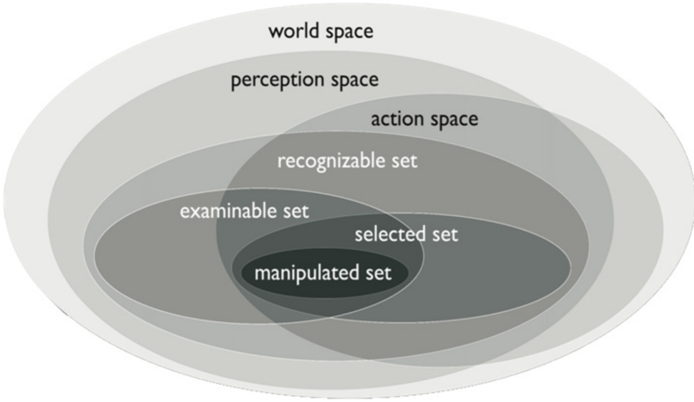

Fig. 1. A physical–virtual situative space model based on a human agent's perception and action possibilities [30].

The situative space model is centered on a human agent and varies continuously according to their movement of body and body parts: it is intended to capture what the specific human agent can perceive and not perceive, affect and not affect at a given moment in time. The situative space model is heavily inspired by situatedness (refer to Section 2.4) and the proximity principle (refer to Section 2.5), and it facilitates replacing the traditional concepts of input and output by perception and action (refer to Section 2.6). The situative space model represents a human agent's perception and action possibilities in a unified manner including: (a) both physical and virtual objects, supporting the physical–virtual equity principle (refer to Section 2.7), and (b) multiple modalities through which a human agent perceive and act.

Operational definitions of the components of the situative space model are as follows (also refer to Fig. 1):

- *World Space (WS):* A space containing the set of all physical and virtual objects to be part of the easy ADL ecology.
- *Perception Space (PS):* The part of the space around the agent that can be perceived at each moment. Perception space can be given a simple geometrical interpretation like a cone in the case of vision as a rough approximation. Objects may occlude other objects and thus create (temporary) holes in the space.
- *Recognizable Set (RS):* The set of objects currently within perception space that are within their recognition distances, i.e., it is possible for a human agent to recognize what type of objects they are.
- *Examinable Set (ES):* The set of objects currently within perception space that are within their examination distances, i.e., it is possible for the human agent to not only recognize the objects but also their states. Normally, we expect the examinable set to be a proper subset of the recognizable set.
- *Action Space (AS):* The part of the space around the agent that is currently accessible through the agent's actions. Objects within this space can be directly acted on. The outer range limit is basically determined by the physical or virtual reach of the agent, but obviously depends qualitatively also on the type of action and the properties of objects involved; e.g., a reachable object may be too heavy to handle with outstretched arms. Since many actions require perception to be efficient or even effective at all, action space is qualitatively affected also by the current shape of perception space.
- *Selected Set (SS):* The set of objects currently being physically or virtually handled (touched, gripped; or selected in the virtual sense) by the agent.
- *Manipulated Set (MS):* The set of objects whose states (external as well as internal) are currently in the process of being changed by the agent. Normally, we expect the manipulated set to be a subset of the selected set.

4. The easy ADL ecology

In this section we present an ambient ecology based on egocentric interaction: the easy ADL ecology. It comprises smart objects (i.e. physical objects augmented with ambient intelligence technology, see Section 4.1), a wearable computer running a personal activity-centric middleware (refer to Section 4.2), a set of ambient intelligence applications containing virtual objects (refer to Section 4.3), and a human agent literally in the middle of it all. The virtual objects are accessible to the human agent through the smart objects. Refer to Fig. 6 for an overview of the infrastructure of the easy ADL ecology.

4.1. Smart Objects (SOs)

Smart objects are ordinary physical objects with additional capabilities that provide complementary virtual (or computing) services without compromising on their established primary purpose [31]. Virtual services include sensing and recognizing their states and the surrounding context, communicating information with other smart objects and the personal activity-centric middleware, facilitating actuation and human interaction with ambient intelligence applications. In providing virtual services, smart objects are expected not to compromise on their appearance and their interaction metaphor, which is an interesting and largely unexplored challenge; aesthetics and other artistic, non-technical factors come

Fig. 2. Smart objects augmented with interactive mediators providing access to virtual objects within the easy ADL ecology.

into play. Smart objects are expected to assist and support human agents with their activities within an ambient ecology and are important building blocks in addressing ambient ecology challenges like recognizing human activities, localizing objects within the ambient ecology, and facilitating human interaction with ambient intelligence applications.

4.1.1. *Heterogeneity of smart objects*

Smart objects possess varied computational and technological capabilities. At one end there are everyday objects that are tagged with passive tags where the smart object's intelligence is located in an external infrastructure, at the other end there are everyday objects with built-in intelligence able to act autonomously like intelligent robots. In this article, smart objects are classified depending on the type of mediators that are embedded in it. The type of mediators to a large extent determines the underlying technologies used in building those smart objects. Smart objects equipped with interactive mediators like LCD screens, audio speakers, microphones, and touch screens communicate with the wearable computer using WLAN technology, forming a local area network for media exchanges. Refer to Fig. 2:

- Fig. 2(a) and (b) *Bedroom photo frame* smart object that usually displays photos of relatives, friends, nature, fantasy, etc. depending on the human agent's context. It also provides access to other virtual objects like alarm clock, day schedule, weather information, dressing assistant, etc. depending on the context.
- Fig. 2(c) and (d) *Bathroom mirror* smart object that usually presents latest news, personal information like day schedule, emails, etc. in parallel to reflecting the human agent's face.
- Fig. 2(e) and (f) *Refrigerator* smart object usually provides access to virtual objects that keep track of the expiry date of food items within the refrigerator, suggests recipes depending on the ingredients available within the refrigerator, etc. Virtual objects that are not associated to any particular physical object like bus timetable, news, emails, etc. are also presented and provide access to depending on the context.
- Fig. 2(g) and (h) *Showcase* smart object presents virtual objects like calendar, photos, etc.
- Fig. 2(i) and (j) *Cutting board* smart object provides access to virtual objects like recipes, cooking assistant, etc. in parallel to allowing the human agent to cut food items like vegetables and meat.
- Fig. 2(k) and (l) *Bookshelf* smart object stores books, CDs and DVDs, in parallel provides information about the latest books from a specific author available in the nearest book store, IMDB ratings and comments, and allow the human agent to access a virtual library from which books could be browsed, borrowed, ordered, etc.

Fig. 3. Smart objects augmented with sensor motes within the easy ADL ecology [32].

- Fig. 2(m) and (n) *Dining table* smart object allows the human agent to have breakfast, lunch, fika, etc. in parallel to providing access to relevant virtual objects.
- Fig. 2(o) *Entrance door* smart object usually reminds the human agent to wear proper jackets, gloves, etc. depending on the weather information, welcomes the human agent when he/she comes home, etc.

Smart objects augmented with interactive mediators present and provide access to virtual objects depending on the human agent's contextual conditions like their activity and action context; situational context described using the situative space model; etc. The interaction management rules to be described in Section 5 is used in determining if, when, where and how to present the virtual objects on smart objects in the easy ADL ecology.

Smart objects augmented with sensor motes communicate with the wearable computer using ZigBee technology forming a wireless sensor network sufficient for low data rate communication in the easy ADL ecology. Refer to Fig. 3:

- Fig. 3(a) *Refrigerator* smart object with embedded temperature sensors, light sensors, etc. presents information like the internal temperature of the refrigerator and the freezer, if their doors are open or closed, etc.
- Fig. 3(b) *Microwave oven* smart object with embedded pressure sensitive pad, on-off switches, light sensor, rotation sensors, etc. that allow for capturing the internal states of the microwave oven and also the state changes caused by a human agent's manipulation.
- Fig. 3(c) *Hand wash* smart object that detects when it is used.
- Fig. 3(d) *Coffee maker* smart object that detects the quantity of coffee beans in it and the type of coffee selected by a human agent.
- Fig. 3(e) and (g) *Stove* smart object that detects if the stove and/or the oven is on or off, the stove's temperatures at different heating plates, the oven's temperature, etc.
- Fig. 3(f) *Waste bin* smart object that detects if it is full and needs to be emptied.
- Fig. 3(h) *Dish washer* smart object that detects if it is on or off, the current program selected, if the dishwasher is open or closed, etc.
- Fig. 3(i) *Cutlery drawer* smart object that detects if it is open or closed.
- Fig. 3(j) *Toilet closet* smart object that detects when the closet is flushed.

Smart objects that are wearable and communicate with the wearable computer using Bluetooth technology and/or wired connection form a body area network in the ambient ecology. Refer to Fig. 4. The *headset* smart object facilitates speech based interaction with virtual objects in the ambient ecology, while the *bracelet* smart object on each hand provide acceleration values along the 3-axis useful in facilitating gesture based interaction with virtual objects.

The aim is to reduce the heterogeneity amongst the smart objects within the easy ADL ecology by channeling the interaction between the human agent and the objects through an activity-centric middleware running on a wearable computer (refer to Section 4.2). The middleware is built on several of the principles of egocentric interaction like situatedness, the physical–virtual equity principle, the proximity principle, replacing input and output by perception and

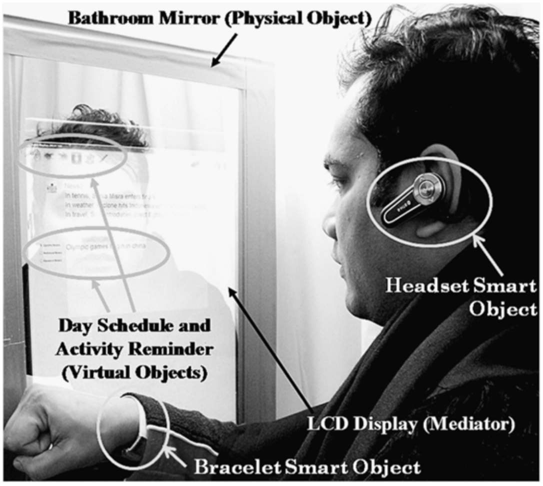

Fig. 4. A *bathroom mirror* smart object presenting day schedule and activity reminder virtual objects manipulated through speech and gestures using a *headset* smart object and the *bracelet* smart objects worn on both the hands by a human agent. The *bathroom mirror* smart object reflects the human agent's face in addition to presenting the virtual services.

action, etc., and is also inspired by Intel's personal server concept [33] where a wearable computer allows a human agent to readily store and process virtual objects while providing access to those virtual objects through the interfaces found in the smart objects. The design choice of centralizing as much as possible of the computation and sensing to a couple of wearable device(s) instead of embedding it in the smart objects is motivated by an expected reduction in design complexity (less attention needs to be directed toward mitigating the problem of uneven conditioning [34]), increased privacy, and reduced cost. The heavy calculations are performed by the wearable computer instead of expecting the smart objects to perform such calculations. There exists a continuum between purely wearable computing approaches to approaches that are purely distributed through co-operative smarts objects [35]. The approach presented in this article lies somewhere in between these two extremes, but more closer to the wearable computing approach.

In an ambient ecology inhabited by multiple human agents (undoubtedly very common), the inherent complexity of individual smart objects might be higher than for those described in this article. It is deliberately chosen to start exploring egocentric interaction using the simple case of one single human agent leaving potential extensions toward explicit support for collaborative human activities for future work.

4.1.2. Entity–relationship representation of smart objects

The entity–relationship (E–R) diagram for smart objects as modeled by the personal activity-centric middleware is presented in Fig. 5. The E–R diagram represents many of the principles of egocentric interaction like situatedness, the physical–virtual equity principle, replacing input and output by perception and action, etc. described in Section 2. Smart objects that are wearable and the ones that are in the environment are treated alike in their representation even though there are inherent differences in their properties, usability, etc. Their differences are known (based on their identity information) and used by the personal activity-centric middleware in its decision making process.

As shown in the E–R diagram, a smart object is considered to be a physical object that provides access to virtual objects through its embedded mediator(s). Physical objects, virtual objects and mediators within the easy ADL ecology are entities that are modeled according to the physical–virtual equity principle (refer to Section 2.7) supporting the concepts of physical–virtual artifacts, physical–virtual situations and physical–virtual activities [21]. In Fig. 5, |p|, |m| and |v| refer to the cardinality of the relationship between physical objects, mediators and virtual objects associated with a smart object, respectively. For a smart object,|p| is always 1 since the smart object has to be a physical object. A purely virtual object is not considered to be a smart object even though it might offer smart services without being associated to specific physical objects (or) mediators in the easy ADL ecology. From the implementation perspective, such virtual objects are associated to all smart objects in the ambient ecology. A smart object should contain at least one mediator and |m| can be greater than or equal to 1. If |m| is equal to 0 then such objects are referred to as "plain" objects requiring external mediators to enable them to be a part of the ambient ecology. A smart object may in some cases not provide access to virtual objects leading to $|v| = 0$.

In the implemented ambient ecology, a smart object has an object identity ObjID and an object type ObjType. A smart object is a physical object with characteristics and functionalities of the physical object dependent on the physical object type PhyObjType. For instance, a *refrigerator* smart object may represent a refrigerator physical object with a freezer or without a freezer depending on PhyObjType. The refrigerator physical object has PhyObj internal state attributes including *temperature*, *on-off*, etc. while the LCD display mediator has mediator state attributes like *display brightness, screen size,*etc. The loudspeaker mediator has mediator state attributes like *loudness, on-off,*etc. Virtual objects are dependent on the

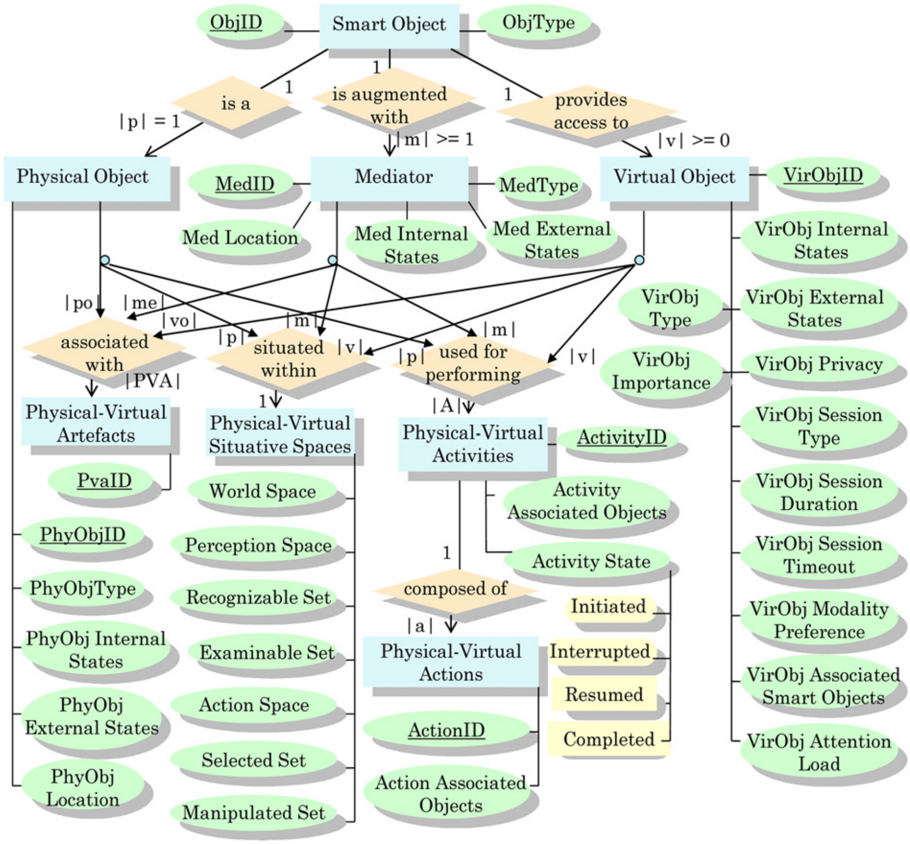

Fig. 5. Entity–relationship diagram of a smart object within the easy ADL ecology.

mediators (usually interactive devices) to enable human agents to access them and interact with them. Physical objects on the other hand, depend on mediators (usually sensors) for allowing the personal activity-centric middleware to keep track of their internal states and external states (i.e., states with reference to other physical objects in a physical space).

When a human agent moves around and acts in physical space, the situative space model associated with them is subject to dynamic changes. Smart objects including mediators enter and leave the situative spaces and sets as an effect. Smart objects that are wearable are usually more stable within the situative spaces and they automatically become a part of the situative spaces the moment a human agent wears them. In environments with a limited amount of stationary (embedded) mediators, wearable mediators are useful in making sure that the human agent is pervasively connected to the virtual world, reducing the problem of uneven conditioning [34].

In Fig. 5, |po|, |me| and |vo| refer to the cardinality of a particular set of physical objects, mediators and virtual objects that are associated with each other to form physical–virtual artifacts (PVA) [21]. |PVA| refers to the set of physical–virtual artifacts currently known to the personal activity-centric middleware. Note that the associative bonds between physical objects, virtual objects and mediators that form PVAs can be both strong and permanent (e.g. created by the designer of the PVA) or weak and temporary (e.g. created dynamically by the activity-centric middleware depending on the human agent's context). Also, note that there is just one situative space model (refer to Section 3) for each individual human agent within the easy ADL ecology and that all the physical objects, virtual objects and mediators known to the personal activity-centric middleware are represented in this model. Physical objects, virtual objects and mediators play a role in physical–virtual activities [21] and are associated to human activities and actions as described in the E–R diagram. Since the smart objects that are part of the easy ADL ecology are managed by the personal activity-centric middleware, they function in an activity-centric manner while offering services to the human agent. Human activities are regarded to be in one of four possible states: initiated, interrupted, resumed, and completed. The activity recognizer to be described in Section 4.2.3 keeps track of these activity states.

4.2. A personal activity-centric middleware

A middleware is useful for handling the heterogeneity of smart objects; it is also useful for dealing with other sources of complexity involved in ambient intelligence applications by encapsulating the low-level sensing and networking tasks, context recognition, and supporting human interactions. The personal activity-centric middleware we propose is intended to

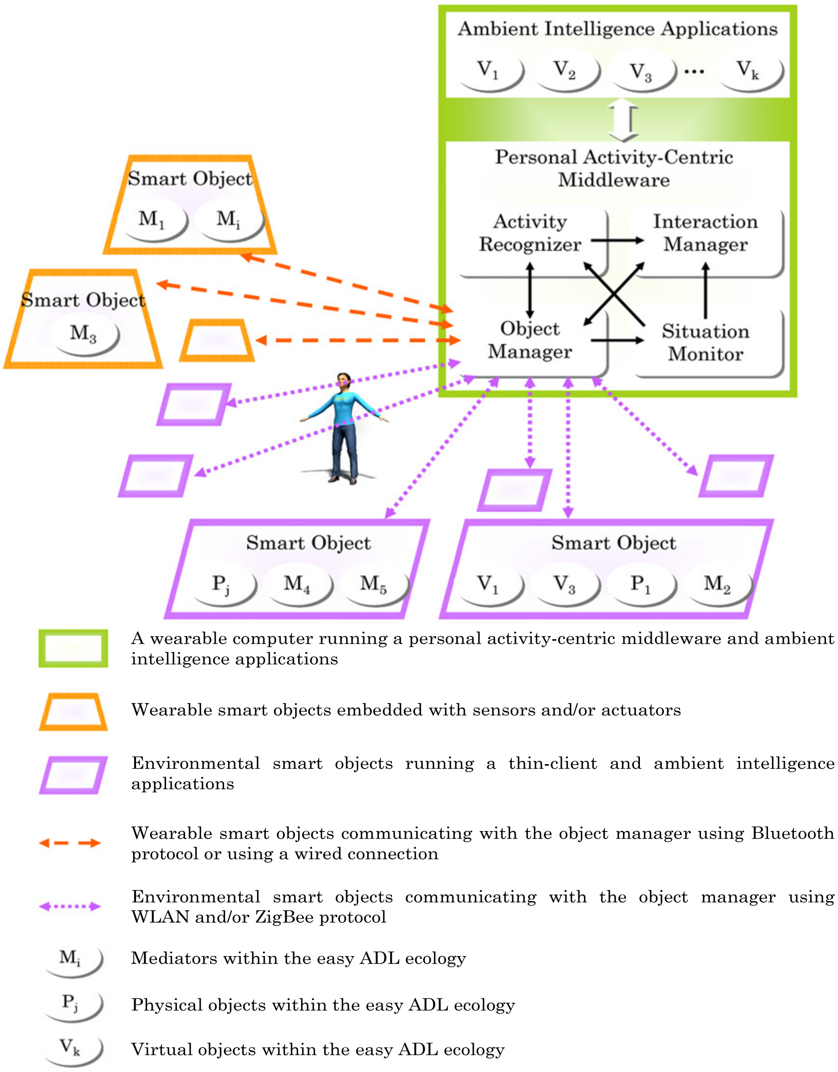

Fig. 6. An infrastructure for ambient ecology comprising of smart objects, a personal activity-centric middleware, ambient intelligence applications and a human agent in the middle of it all.

facilitate the handling of human-centered parameters inspired by egocentric interaction such as perception, action, intention and attention in a unified manner.

Considering the number and diversity of smart objects in the easy ADL ecology, the middleware has to address many of the requirements of traditional distributed systems such as interoperability, scalability, security, and tolerance for component failures and disconnections [36]. Egocentric interaction introduces additional requirements including the modeling and tracking of the smart objects, the situative spaces, and the agent's activities and actions — information which is used to facilitate human interaction with the ambient intelligence applications.

The personal activity-centric middleware is composed of four components, namely the object manager, the situation monitor, the activity recognizer, and the interaction manager, as shown in Fig. 6. The ambient intelligence applications run on top of the middleware, and their synchronized copies run on specific smart objects to reduce the amount data communication between the wearable computer and the smart objects with some processing done locally on the smart objects. The middleware components run as independent modules and offer a common interface to communicate with the ambient intelligence applications. Microsoft Message Queues (MSMQ) are used for exchanging data objects and strings among the middleware components, the ambient intelligence applications and the smart objects. It is important to keep down the number of queues since each queue added introduces additional overhead in exchanging messages. Instead of one queue for each individual communication path, which might also reduce flexibility when new components are added, a single queue for each individual component is used. The queues work like a mailbox where sent messages are collected and can be read periodically. During the initialization phase, the existing queues are emptied and newly arriving messages are listened to. The middleware is implemented in C# while some modules of the object manager are implemented in C++ and some modules of the activity recognizer in Matlab. In the future, the aim is to convert the code written in Matlab to C#.

4.2.1. *Object manager*

The object manager is responsible for the following:

- To manage the smart objects within the easy ADL ecology.
- To initiate and manage wireless communication (and wired communication) with smart objects.
- When a new smart object is detected in the ambient ecology, the object manager attempts to obtain more information about the smart object from its manufacturer's database [37]. A mock-up database is used with an assumption that in the future, smart object manufacturers will maintain such a database online considering the interest shown by the industry and academia in the *Internet of Things* [38].
- Since smart objects usually contain physical objects, virtual objects and/or mediators, real-time information about them are also maintained in an object-centric manner by the object manager. Maintaining information includes not only to identify the objects, but also to store up-to-date information about their attributes and current values. For instance, information about an object's association to relevant activities and actions (from the activity recognizer), the object's position in the situative spaces at a particular moment in time (from the situation monitor), an object's current internal states, external states, etc., are maintained. Refer to Fig. 5 for the E–R diagram.
- The object manager keeps track of the physical objects and mediators by sensing and communicating with the smart objects in the ambient ecology. It keeps track of the virtual objects by communicating with the interaction manager and the ambient intelligence applications running on the smart objects.
- The object manager keeps track of the symbolic location of physical objects and mediators (like "in the bathroom", "in the kitchen", etc.) by associating the individual objects to co-located stationary objects with known symbolic location. Since indoor environments contain walls that often act as a boundary for the perception and action spaces, objects within these spaces often (but not always) can be located to a specific room.
- To exchange information with other middleware components and the ambient intelligence applications, and offer them services related to smart objects. For instance, the interaction manager presents and provides access to virtual objects on smart objects by the mediation of the object manager. Also, to maintain and exchange information about the physical–virtual artifacts formed by the association of physical objects, virtual objects and mediators.

4.2.2. *Situation monitor*

The situation monitor is responsible for the following:

- To keep track of the identity and location of the physical objects, virtual objects and mediators within the human agent's situative spaces (refer to Section 3) by communicating with the object manager and the smart objects mediated by the object manager. Also, to keep track of the states and state changes of these objects (important in keeping the examinable set and the manipulated set updated).
- To exchange information with other middleware components and the ambient intelligence applications, and offer them services related to the agent's situation. For instance, the interaction manager can decide if, when, where and how to position virtual objects within the individual situative spaces by getting up-to-date information about the content of the situative spaces while the activity recognizer recognizes activities and actions using the content of the situative spaces [39].

The primary approach taken for situative space tracking is to use wireless LAN signal-strength-based localization as described in [40]. ZigBee-based wireless sensor networking of smart objects is used to keep track of the state changes (examinable set and manipulated set) to physical objects and mediators as described in [32].

4.2.3. *Activity recognizer*

The activity recognizer is responsible for the following:

- To model and recognize the agent's activities and actions using information channels provided by the situation manager. The different information channels contain information about the content of the different situative spaces and sets over time. Two different activity recognition systems have been built based on these information channels, for details refer to [39,41].
- Modeling human activities and actions are done offline (i.e. while not performing the task of activity and action recognition) and in a supervised manner during an explicit training phase. In a future implementation, the aim is to model human activities using online learning algorithms. Supervised learning is preferred to unsupervised learning because of the complexities involved in modeling and recognizing everyday human activities and actions. Supervised learning allows for personalized modeling of human activities and actions, and facilitates their recognition with higher accuracy.
- To derive activity-centric information including: (a) the set of objects that are associated to individual activities and actions; (b) the set of actions and operations that are mandatory for individual activities; and (c) important events that determine the initiation, interruption, resumption and completion of individual activities.

- To exchange information with other middleware components and the ambient intelligence applications, and offer them services related to human activities and actions. For instance, the interaction manager communicates with the activity recognizer to know the human agent's current activity and action context for deciding if, when, where and how to present and provide access to virtual objects within the ambient ecology. Other services include providing activity-centric information to: the object manager that uses this information in associating individual objects known to the object manager with their respective activities and actions; and the cognitive prosthesis application that attempts to provide activity support for mild-dementia patients within an ambient ecology.

4.2.4. Interaction manager

The interaction manager is responsible for the following:

- To facilitate interaction with virtual objects (part of the ambient intelligence applications) based on the principles of egocentric interaction (refer to Section 2) and the situative space model (refer to Section 3). Even though the focus is on facilitating interaction with virtual objects, the interaction manager handles both physical and virtual objects, physical and virtual situations, and physical and virtual activities from an egalitarian perspective as per the physical–virtual equity principle.
- The interaction manager facilitates situated interaction by considering the current content of the situative spaces (from the situation monitor) and making decisions concerning if, when, where and how to present virtual objects so that they can be perceived and/or acted upon by the human agent.
- The interaction manager facilitates activity-centered interaction by considering the human agent's activity and action context (from the activity recognizer), and the virtual objects' association to the current activity and action context (from the object manager).
- The interaction manager facilitates multimodal interaction by considering the mediators available within the situative spaces (from the situation monitor), their current states (from the object manager) and the virtual objects' mediator and modality preferences (information from the ambient intelligence applications). Multimodal interaction improves: (a) the bandwidth of interaction both in terms of quantity and quality; and (b) the usability by offsetting the weakness of one modality by the strengths of another in various contexts.
- The interaction manager facilitates mixed-initiative interaction by allowing both the human agent and the ambient intelligence applications to initiate interaction sessions. The interaction manager uses separate sets of interaction management rules (refer to Section 5) for handling the virtual objects in human-initiated interaction sessions and ambient intelligence application-initiated interaction sessions. Mixed-initiative interaction enables the agent to exert control in parallel with proactively operating ambient intelligence applications in the background to better satisfy the agent's computing needs.
- The interaction manager facilitates ambient interaction by considering the human agent's peripheral and central attention capabilities, and the virtual objects' attention requirements (by communicating with the ambient intelligence applications). Human attention is modeled at two levels of abstraction: (a) at a low-level, human attention is driven by objects within the situative spaces (obtained by communicating with the situation monitor); and (b) at a high-level, human attention is driven by human intentions (obtained by communicating with the activity recognizer).
- To handle additional virtual object attributes like privacy, importance, etc. that are important to be considered and is established by the ambient intelligence applications. Privacy is addressed by considering the symbolic location of the mediators in the ambient ecology and augmented within smart objects.
- To exchange information with other middleware components and the ambient intelligence applications, and offer them services related to human interaction with virtual objects within an ambient ecology.

4.3. Ambient intelligence applications

Ambient intelligence applications run on top of the personal activity-centric middleware and synchronized copies run on selected smart objects for the reason mentioned earlier. The applications on the smart objects are synchronized, managed and controlled by the middleware in a centralized manner. Ambient intelligence applications communicate with other middleware components to obtain relevant contextual information about the easy ADL ecology including the human agent in it. The applications contain domain specific knowledge and by using the additional contextual information, they become context-aware thereby satisfying the general requirements for being a part of an ambient ecology. Ambient intelligence applications communicate with the human agent by presenting virtual objects that are filtered and managed by the interaction manager according to the interaction management rules to be described in Section 5. These virtual objects maintained by the ambient intelligence applications possess many attributes that are relevant within the ambient ecology (e.g. activity association, attention load, associated smart objects, etc.) which traditional virtual objects belonging to typical desktop applications (accessed through mouse, keyboard and a screen) do not possess.

In the specific design case reported in this article, the ambient intelligence applications have been designed to provide physical–virtual activity [21] support to people in the early stages of dementia. The easy ADL ecology (where ADL stands for activities of daily living as commonly understood within the medical field) is intended to provide human agents with more independence and satisfaction than a normal, non-augmented environment can. The research effort is long-term and the work presented here includes 19 mock-up (limited functionality and/or partially implemented) applications introducing more than 100 virtual objects. Applications include medicine reminder, old-food-in-the-refrigerator reminder, food recipe provider, diet controller, shopping assistant, safety and security manager, weather information provider, clothing assistant, news provider, transportation information provider, day scheduler, etc.

5. Interaction management rules and techniques

The interaction manager provides access to virtual objects within the easy ADL ecology on the basis of interaction management rules that answer the important questions of *if*, *when*, *where* and *how* a virtual object should be made present and accessible on request by (a) the human agent or (b) an ambient intelligence application. The rules are prioritized according to which question they address with the *if* question having the highest priority, followed by the *when* question, the *where* question, and the *how* question. In Table 1, *if* question is linked to the virtual object attributes *importance*, *activity association*, and *action association*. Similarly, the *when* question is associated with *session duration* and *session timeout*; the *where* question to *attention load*, *associated smart objects*, and *privacy*; and (finally) the *how* question to *modality preference*. The interaction techniques described in this also address the *how* question.

The virtual objects in the interaction manager either enter a *human-initiated interaction session* or an *application-initiated interaction session* depending on who initiated interaction, activating different sets of interaction management rules. The interaction manager filters the virtual objects: only virtual objects that satisfy the necessary conditions of the interaction management rules are pushed into the dispatching queue.

The dispatching queue is a simple priority queue that constantly gets updated (every 2 to 3 s) with virtual objects depending on the contextual conditions within the ambient ecology and the current needs of the ambient intelligence applications. The update rate is dependent on the low-level sensing and networking infrastructure of the easy ADL ecology (which could be improved with technological advancements in the future), and some of the personal activity-centric middleware components like the activity recognizer and the situation monitor. The update rate is not an issue since human agents usually take a few seconds to stabilize themselves in a situation before virtual objects could usefully be presented to them.

Handling conflicts among the interaction management rules is a major challenge and in the event of a conflict, the virtual objects involved in a human-initiated interaction session is given priority over virtual objects involved in an application-initiated session. It should be noted that several sessions usually run at the same time. Within a session, the *importance* attribute is given the highest priority followed by *activity association* and *action association*, *session duration*, *session timeout*, etc. *Modality preference* is given the least priority. The specific formulation of the interaction management rules in Table 1 has been guided by egocentric interaction principles but details may change in future versions of the system as the ambient ecology setup evolves. For instance, the priorities could be personalized to a human agent both implicitly by learning and modeling their interaction experience, and explicitly by allowing the human agent to modify the priorities.

Refer to Table 1. The importance attribute value of a virtual object may change while being a part of an interaction session. For instance, a virtual object with low importance can suddenly advance to medium importance and/or high importance depending on the context and the passage of time. The virtual object's attribute values are handled by the respective ambient intelligence applications and updated instantly before the virtual object is filtered by the interaction manager. Some virtual objects are associated to specific activity and action context. For instance, a diet dinner virtual object is associated to the activity of preparing dinner where the virtual object presents information when fatty ingredients are used during the meal preparation. Such virtual objects are presented by taking into consideration a human agent's current activity and action context. Also note that some virtual objects are independent of the activity and action context. Such virtual objects are handled by associating them to all activity and action contexts modeled within the ambient ecology before applying the interaction management rules described in Table 1. Similarly, some virtual objects are independent of the smart objects available in the ambient ecology and such virtual objects are considered to be associated to all smart objects in the ambient ecology. However, the virtual objects that are associated to specific smart objects like the old food indicator virtual object which is associated to the refrigerator smart object, the virtual objects are presented considering their association to currently available smart objects.

Note that when a human agent changes location or if the mediator selected by the human agent is not available for mediation anymore within a human-initiated interaction session, then the new mediator, the new smart object and/or the new modality for presenting and providing access to virtual objects are handled by the interaction manager. Then some of the virtual object's attribute values like attention load, association to smart objects, privacy attribute, and modality preference, which were previously ignored, are considered before presenting the virtual object in an appropriate manner.

In Table 1, the interaction management rules consider the mediators that support demand-agent-response to also include mediators that support central and peripheral attention, while mediators that support central attention to also include mediators that support peripheral attention. Mediators that can draw on a human agent's peripheral attention are referred to as peripheral mediators; mediators that demand central attention are referred to as central mediators in Table 1. Mediators that present personal virtual objects (like emails) are referred to as personal mediators; mediators that present personal and private virtual objects (like shopping lists) are referred to as private mediators; and other mediators handling public virtual objects (like morning news) are referred to as public mediators.

Table 1
Interaction management rules within the easy ADL ecology.

Virtual Object Attributes {*values*}	**Human-Initiated Interaction Session** **Initiation**: Selection of a virtual object by a human agent. **Termination**: De-selection of a virtual object by a human agent.	**Application-Initiated Interaction Session** **Initiation**: Selection of a virtual object by an application. **Termination**: Virtual object reaches the session timeout or the end of its session duration.
Importance {*high, medium, low*}	Importance value is *ignored*. Virtual object is presented immediately.	Importance value is *considered*. **High**: Virtual object is presented immediately. **Medium**: Virtual object is presented if no high importance virtual objects are in the dispatching queue. **Low**: Virtual object is presented if no high and medium importance virtual objects are in the dispatching queue.
Activity association {*set of activities*} Action association {*set of actions*}	Activity and action association is *ignored*. Virtual object is presented immediately.	Activity and action association is *considered*. **Activity and action context matches**: Virtual object is presented immediately. **Activity context alone matches**: Virtual object waits in the dispatching queue for the action context to match. **Activity and action context does not match**: Virtual object waits in the dispatching queue for the activity and action context to match.
Session duration {*duration in seconds*}	Session duration value is *ignored*. Virtual object presentation continues until the human agent terminates it.	Session duration value is *considered*. **Session duration is not reached**: Virtual object presentation continues. **Session duration is reached**: Virtual object presentation terminates.
Session timeout {*timeout in seconds*}	Session timeout value is *ignored*. Virtual object waits in the dispatching queue until the human agent terminates it.	Session timeout value is *considered*. **Session timeout is not reached**: Virtual object waits in the dispatching queue. **Session timeout is reached**: The ambient intelligence application is informed about the failure in presenting the virtual object.
Attention load {*peripheral attention, central attention, demand agent response*}	**Same mediator as selected by the human agent**: Attention load is *ignored*. Virtual object is presented through that mediator. **Different mediator selected by the interaction manager**: Attention load is considered. Rules 1–3 applies.	**Mediator is selected by the interaction manager**: Attention load is *considered*. **Peripheral attention**: Virtual object is presented if a peripheral mediator is available. (Rule 1) **Central attention**: Virtual object is presented if a central mediator is available. (Rule 2) **Demand agent response**: Virtual object is presented if mediators capable of allowing the human agent to manipulate the virtual object is available. (Rule 3)
Associated smart objects {*set of smart objects*}	**Same smart object selected by the human agent**: Association to that smart object is *ignored* and the virtual object is presented. **Different smart object selected by the interaction manager**: Association to that smart object is *considered*. Rules 4 and 5 applies.	**Smart object is selected by the interaction manager**: Association to that smart object is *considered*. **Associated**: Virtual object is presented. (Rule 4) **Not associated**: Virtual object waits in the dispatching queue. (Rule 5)
Privacy {*personal, private, public*}	**Same mediator selected by the human agent**: Privacy is *ignored* and the virtual object is presented. **Different mediator selected by the interaction manager**: Privacy is *considered*. Rules 6–8 applies.	**Mediator is selected by the interaction manager**: Privacy is *considered*. **Personal**: Virtual object is presented if a personal mediator is available. (Rule 6) **Private**: Virtual object is presented if a private mediator is available. (Rule 7) **Public**: Virtual object is presented if a mediator is available. (Rule 8)
Modality preference for access {*visual, audial*} Modality preference for manipulation {*gesture, speech*}	**Same modality selected by the human agent**: Modality preference is *ignored* and the virtual object is presented. **Different modality selected by the interaction manager**: Modality preference is *considered*. Rules 9 and 10 applies.	**Modality is selected by the interaction manager**: Modality preference is *considered*. **Preferred modality**: Virtual object is presented. (Rule 9) **Not the preferred modality**: Virtual object waits in the dispatching queue. (Rule 10)

Separate threads run to keep track of the virtual object's session timeout attribute value and session duration attribute value continuously. A virtual object in an application-initiated interaction session can stay in the interaction manager's dispatching queue, or can continue to be presented through smart objects as long as the virtual object's session timeout or session duration has not expired, irrespective of the other contextual conditions.

Egocentric interaction does not favor the use of any particular modality or mediator type for communication between the human agent and the surrounding ambient ecology. It only highlights the prerequisites for such communication to be at all possible: (a) mediators used for controlling or manipulating virtual objects/applications need to be reachable (inside the action space in the situative space model, Fig. 1) in the given situation; (b) mediators used for perceiving the presence and state of virtual objects need to be present themselves in the perception space of the given human agent. It is the role of the interaction manager (Fig. 6) to find the best combination of mediators given the situation. To increase the presence of mediators that enable manipulation of virtual objects, wearable mediators were relied on: accelerometer-fitted bracelets for identifying hand gesture commands, and a headset microphone for speech command reception. For the presentation of the virtual objects, mediators on smart objects both in the environment (visual displays for visual presentation and loudspeakers for audial presentation) and the wearable ones (headphone for audial presentation) were used. Interaction techniques were developed which address the question of *how* to present and provide access to virtual objects in the ambient ecology. These interaction techniques were designed to accommodate for limitations of the technology that was at hand with future improvement possibilities. The overall goal has been to enable interaction with multiple objects at the same time over multiple interaction modalities distributed within the situative spaces.

Personalization of the situative spaces surrounding a human agent is an important challenge to address within an ambient ecology. Even though the problem is simplified in single human agent scenarios, it is still important to identify the parts of the situative spaces that are personal and the ones that are anonymous since ambient ecologies may contain multiple human agents. The selected and the manipulated sets are directly associated to object selections and manipulations respectively where the identity of the human agent causing such object manipulation are identified based on: the mediators worn by the agent (accelerometers for hand gesture recognition), or based on profile selection (speech recognition engine trained for specific user profiles). The action space and the perception space of a specific human agent are shared spaces that changes dynamically. Personalization of such spaces might form intersected volumes with other human agent's situative spaces introducing a need to handle shared personalization.

6. Concluding remarks

In this article, we have presented the concepts of egocentric interaction, a novel interaction paradigm for the design and use of ambient intelligence applications. Egocentric interaction is a paradigm that is literally human-centered and based on a mobile human agent's dynamically changing possibilities and propensities to perceive and act in an environment. To capture this flow of changes, a simplified, implementable situative space model has been developed in conjunction with an ambient ecology, the easy ADL ecology, giving us an infrastructure for further explorations. The constituents of the easy ADL ecology have been detailed and specific interaction management rules presented. A user experience evaluation of the easy ADL prototype system using these rules can be found in [42]. The work reported here can be considered as an initial, yet concrete step toward the exploration and development of a human-centered interaction paradigm for ambient ecologies.

References

[1] E. Aarts, R. Wichert, Ambient Intelligence, Technology Guide, Springer, Berlin Heidelberg, 2009, 244–249.
[2] A. Alvarez, A. Arcos, N. Herrasti, AMEC — facilitating a user centred evolution to ambient intelligent environments, 2nd IET International Conference on Intelligent Environments, Vol. 2, 2007, pp. 61–67.
[3] K. Fujinami, F. Kawsar, T. Nakajima, AwareMirror: A personalized display using a mirror, Pervasive Computing, LNCS 3468, 2005, pp. 137-150.
[4] A. Dey, Understanding and using context, Personal and Ubiquitous Computing 5 (1) (2001).
[5] B. Ullmer, H. Ishii, R.J.K. Jacob, Token+constraint systems for tangible interaction with digital information, ACM Transactions on Computer-Human Interaction 12 (1) (2005) 81–118.
[6] C. Wisneski, H. Ishii, A. Dahley, M. Gorbet, S. Brave, B. Ullmer, P. Yarin, Ambient displays: turning architectural space into an interface between people and digital information, in: Proceedings of the First International Workshop on Cooperative Buildings, Integrating Information, Organization, and Architecture, Springer-Verlag, 1998.
[7] J. Wall, Demo I microsoft surface and the single view platform, in: Proceedings of the 2009 International Symposium on Collaborative Technologies and Systems, IEEE Computer Society, 2009.
[8] M. Turk, G. Robertson, Perceptual user interfaces, Communications of the ACM 43 (3) (2000) 33–34.
[9] R. Vertegaal, J. Shell, D. Chen, A. Mamuji, Designing for augmented attention: towards a framework for attentive user interfaces, Computers in Human Behavior 22 (2006) 771–789.
[10] H. Gunes, M. Piccardi, Automatic temporal segment detection and affect recognition from face and body display, Trans. Sys. Man Cyber. Part B 39 (2009) 64–84.
[11] Y. Li, J. Landay, Activity-based prototyping of ubicomp applications for long-lived, everyday human activities, in: Proc. of CHIí08, 2008, pp. 1303–1312.
[12] P. Dourish, Where The Action Is: The Foundations of Embodied Interaction, MIT Press, Cambridge, MA, 2001.
[13] R. Jacob, A. Girouard, L. Hirshfield, M. Horn, O. Shaer, E. Solovey, J. Zigelbaum, Reality-based interaction: a framework for post-WIMP interfaces, in: Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems, Florence, Italy, April 2008, pp. 201–210.
[14] K. Kuutti, and B. Nardi (Eds.), Activity theory as a potential framework for human-computer interaction research, Context and Consciousness. Chapter 2, 1995, pp. 17–44.
[15] J.J. Gibson, The Ecological Approach to Visual Perception, Houghton Mifflin, Boston, 1979.
[16] E. Hutchins, Cognition in the Wild, MIT Press, Cambridge, 1995.
[17] L. Suchman, Plans and Situated Actions: The Problem of Human Machine Interaction, Cambridge University Press, Cambridge, 1987.

[18] A. Clark, D. Chalmers, The extended mind, Analysis 58 (1) (1998) 7–19.
[19] G. Lakoff, M. Johnson, Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought, Basic Books, New York, 1999.
[20] W. Mackay, A.-L. Fayard, L. Frobert, L. Médini, Reinventing the familiar: exploring an augmented reality design space for air traffic control, in: Proceedings of ACM CHI 1998, Los Angeles, ACM Press, New York, 1998, pp. 558–565.
[21] T. Pederson, From conceptual links to causal relations — physical-virtual artefacts in mixed-reality space, Ph.D Thesis, Dept. of Computing Science, Umeå University, report UMINF-03.14, 2003.
[22] T.S. Kuhn, The Structure of Scientific Revolutions, second ed., The University of Chicago Press, Chicago, 1970.
[23] L.-E. Janlert, J. Jacko (Eds.), The Evasive Interface — the changing concept of interface and the varying role of symbols in human–computer interaction, in: Human–Computer Interaction, Part I, HCII, LNCS 4550, 2007, pp. 117–126.
[24] D. Norman, S. Draper (Eds.), User Centered System Design, Erlbaum, Hillsdale, NJ, 1986.
[25] P. Ehn, Work-Oriented Design of Computer Artifacts, Lawrence Erlbaum Associates, Hillsdale, NJ, 1990.
[26] T. Nagel, The View From Nowhere, Oxford Univ. Press, New York, 1986.
[27] D. Kirsh, P. Maglio, On distinguishing epistemic from pragmatic action, Cognitive Science 18 (1994) 513–549.
[28] L.-E. Janlert, Contextual strategies — notes for a theory of context. Technical report UMINF 02.23, Umeå University, ISSN-0348-0542, 2003.
[29] B. Shneiderman, Direct manipulation: a step beyond programming languages, Computer 16 (1983) 57–69.
[30] T. Pederson, L.-E. Janlert, D. Surie, Setting the stage for mobile mixed-reality computing — a situative space model based on human perception, IEEE Pervasive Computing 10 (4) (2011).
[31] F. Kawsar, K. Fujinami, T. Nakajima, Prottoy middleware platform for smart object systems, International Journal of Smart Home 2 (2008).
[32] D. Surie, O. Laguionie, T. Pederson, Wireless sensor networking of everyday objects in a smart home environment, International Conference on Intelligent Sensors, Sensor Networks, and Information Processing, Sydney, Australia, IEEE CFP08842-PRT, 2008.
[33] R. Want, T. Pering, G. Danneels, M. Kumar, M. Sundar, J. Light, The personal server: changing the way we think about ubiquitous computing, in: G. Borriello, L.E. Holmquist (Eds.), UbiComp 2002, in: LNCS, vol. 2498, Springer, Heidelberg, 2002.
[34] M. Satyanarayanan, Pervasive computing: vision and challenges, IEEE Personal Communications (2001).
[35] B.J. Rhodes, N. Minar, J. Weaver, Wearable computing meets ubiquitous computing: reaping the best of both worlds, Third International Symposium on Wearable Computers (ISWC'99), 1999.
[36] K. Henricksen, J. Indulska, T. McFadden, S. Balasubramanian, Middleware for distributed context-aware systems, International Symposium on Distributed Objects and Applications (DOA), 2005.
[37] U. Kulkarni, J. Vadavi, S. Joshi, K. Sekaran, A. Yardi, Ubiquitous object categorization and identity, International Conference on Advanced Computing and Communications, IEEE Computer Society, Issue (20–23), 2006, pp. 585-588.
[38] N. Gershenfeld, R. Krikorian, D. Cohen, The internet of things, Scientific American 291 (4) (2004) 76–81.
[39] D. Surie, T. Pederson, F. Lagriffoul, L.-E. Janlert, D. Sjölie, Activity recognition using an egocentric perspective of everyday objects, Ubiquitous Intelligence and Computing, LNCS 4611, 2007.
[40] D. Surie, F. Jäckel, L.-E. Janlert, T. Pederson, Situative Space Tracking within Smart Environments, Intelligent Environments, IEEE Computer Society, 2010, 152–157.
[41] D. Surie, F. Lagriffoul, T. Pederson, D. Sjölie, Activity recognition based on intra and extra manipulation of everyday objects, Ubiquitous Computing Systems, LNCS 4836, 2007.
[42] D. Surie, Egocentric interaction for ambient intelligence, Ph.D Thesis, Department of Computing Science, Umeå University, Sweden, UMINF 12.01, 2012.