Ambient Intelligence Implies Responsibility

Christian A. Detweiler[a,b,1] , Francien Dechesne[a], Koen V. Hindriks[b] and Catholijn M. Jonker[b]
[a] *Philosophy Section, Faculty of Technology, Policy and Management, Delft University of Technology*
[b] *Man-Machine Interaction Group, Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology*
{C.A.Detweiler, F.Dechesne, K.V.Hindriks, C.M.Jonker}@tudelft.nl

Abstract. Ambient Intelligence Systems are complex, knowledge providing systems. The group of intended users typically comprises both various kinds of professionals and of laymen. In this chapter we address responsibility issues associated with the use and design of AmISs. We explain the connection between knowledge and responsibility, and use that connection to show that the usage of Ambient Intelligence Systems leads to moral responsibilities for all types of users (professionals and otherwise). It follows from the same reasoning scheme applied to the role of the AmIS developer, that she is responsible to design the system for responsible use. We give some initial criteria for such *design for responsibility* of AmISs.

Keywords. Ambient Intelligence, Agent Technology, Responsibility

Introduction

Ambient Intelligence Systems (AmISs) are being developed to improve the quality of life. However, as is typically the case with new technology, AmISs also introduce some new concerns. This chapter first argues that using an AmIS will burden every type of user with responsibilities and then studies the possible impact of these responsibilities on the types of users. This has an effect on the design responsibilities of the developer of AmISs. The argument is based on the idea that knowledge implies responsibilities; see e.g., [1,2]. As human users of Knowledge Based Systems (KBSs) in general, and AmISs in particular, are equipped with more knowledge by using such systems, we address the question whether this leads to new or more responsibilities.

Ambient Intelligence Systems can be seen as a form of Knowledge-Based Systems (see e.g., [3]), where AmISs distinguish themselves e.g. by their distributed nature, and the way they are used (in time, place, and in numbers of users). AmISs have in common with KBSs that they provide the user with information and knowledge, though AmISs, where such systems are ubiquitous and invisible, can support users without users knowing. From the starting point that with knowledge comes responsibility, AmISs raise a number of questions: What kind of responsibilities come with the use of AmISs? For which users? Under which conditions? Could the impact be potentially

[1] Corresponding Author.

more substantial than that of traditional Knowledge Based Systems? Can we single out aspects that are due to the inclusion of (Wireless) Sensor Networks and Intelligent Agents in AmISs?

AmISs typically are designed for a combination of both professional users and layperson users. For example, the CCF2 system to support independent living for the elderly (see e.g., [4]), has professional nursing staff, the elderly client, and family members of the client as users. These user groups together form the group of the intended users. To add to the variety within the user group, an AmIS may also have unintended users, i.e., people that the stakeholders and designers of the system did not envisage as users when they designed the system. Although issues of security (which are tied to unintended use) fall out of the scope of this chapter, we find it important to address the difference between intended and unintended users in terms of responsibilities.

The various user types, such as physicians, nurses and family members, and their differences in training and background puts the different types of users in different positions with respect to responsible use of the system. An important question is what these different responsibilities are for the different user types, and whether all users are aware of their responsibilities. In this chapter, we will argue that an important responsibility of the designer is to at least make all types of *intended* users of AmISs aware of the responsibilities that come with the use of the system, so that they can accept them (*informed consent*).

Another important question with respect to responsibility in AmISs is the role of intelligent software agents [5], though the question of the moral agenthood of intelligent software agents is beyond the scope of this work. Intelligent software agents typically play a role as representatives of human stakeholders, for example software agents that manage users' personal data or an agent that plans appointments with a physician. Should we regard these agents as part of the design of the AmIS, or as users of it (as the human parties they represent)? If seen as users, they would inherit the moral responsibilities that come with the use of the system. If seen as part of the AmIS, responsibility would be attributed to their human designers.

The previous remark points at an important discussion this chapter intends to initiate, concerning the responsibilities of the AmIS developers. If responsibilities can be associated with the use of AmISs, then we claim this should in itself impact the design of AmISs. The question is, how? What are the criteria for designing a system for responsible usage?

The focus of this chapter is on how ambient intelligence systems affect their stakeholders' and designers' responsibilities through the knowledge these systems produce. Before we can examine these relationships, we must explain the relationship between knowledge and responsibility. This requires an explanation of what responsibility is; how it affects people and how it is ascribed to them. Therefore, this chapter is organized as follows. Section 1 discusses the notion of (moral) responsibility. Various types of responsibilities are discussed as well as when people can be held responsible (the so-called knowledge condition of responsibility). Section 2 relates knowledge-based systems to responsibilities. A case study is discussed in Section 3 with an emphasis on the stakeholders and their responsibilities. Section 4 discusses the knowledge produced by AmIS and the responsibilities that brings to the various user types. Section 5 reflects on the previous sections of the chapter and distills design responsibilities for the designers of Ambient Intelligence Systems. Throughout the chapter, the CCF2 system for Independent Living is used as an example AmIS.

1. Responsibility

Responsibility affects people in different ways. It can dictate how someone should behave; for example, if a nurse is responsible for the stability of a patient's condition, she should check on that patient frequently. Responsibility also co-determines how others should behave in light of someone's actions, e.g. whether they should (also) check the patient, or whether they should check the nurse. There are many different ways in which the concept of responsibility is used, and it is hard to define clearly under which circumstances and how people "get" responsibilities, i.e. how responsibility is ascribed. This section provides some forms of responsibility and a definition of fair ascription of responsibility.

1.1. Forms of Responsibility

There is no unified understanding of responsibility. In fact, there are many subtly different senses in which the word "responsibility" is used, as this anecdote illustrates [6, p. 96] (the anecdote originates from [7]):

> *(1) Smith had always been an exceedingly responsible person, (2) and as captain of the ship he was responsible for the safety of his passengers and crew. But on his last voyage he drank himself into a stupor, (3) and he was responsible for the loss of his ship and many lives. (4) Smith's defense attorney argued that the alcohol and his transient depression were responsible for his misconduct, (5) but the prosecution's medical experts confirmed that he was fully responsible when he started drinking since he was not suffering from depression at that time. (6) Smith should take responsibility for his victims' families' losses, but his employer will probably be held responsible for them, as Smith is insolvent and uninsured.*

The six different forms of responsibility used in this example, are what Vincent calls (1) virtue, (2) role, (3) outcome, (4) causal, (5) capacity, and (6) liability responsibility.

Virtue responsibility (1) refers to having the characteristic of being the type of person who does the right thing. Role responsibility (2) refers to what someone ought to do, i.e., the duties that come with a role; the captain ought to have seen to the safety of his passengers and crew. Outcome responsibility (3) refers to rightfully being identifiable as the cause of some situation *X* and deserving of praise or blame for it; in the example the captain is responsible for the loss of the ship and of many lives. This concept is closely related to causal responsibility (4) that refers to a causal chain that leads to the outcome that is situation X; in this case the alcohol that the captain should not have taken. Capacity responsibility (5) refers to the fact the captain had the capacity to realize that he shouldn't drink so much alcohol. Being capacity responsible means having the capability to understand what behavior relevant norms require, to reason about what one ought to do in light of such requirements, and to control one's behavior accordingly [8]. Liability responsibility refers to the responsibility of the captain for the consequences of the outcome, in particular for the financial losses of the families.

Accountability is closely related to the notion of liability. Whereas liability involves the question of who will pay for damages done or who will set things straight, accountability involves answering for one's actions (of which damages were an outcome) [9]. This concept is also related to Vincent's concept of causal responsibility, in that in answering for one's actions, there is assumingly a causal connection between

those actions and the outcome. Also, there is a connection between the concept of outcome responsibility and accountability, in that "[a]n important set of cases in which one may reasonably expect accountability for a harm is that in which an analysis points to an individual (or group of individuals) who are morally blameworthy for it" [9, p. 28].

For other views of responsibility, such as those based on merits, rights or consequences, see [10]. However, this chapter focuses on those notions of responsibility that are most relevant for AmISs during design time and when starting to use the functionalities of AmISs.

For analysis of responsibility also the notions of "forward-looking" and "backward-looking" responsibility are relevant, see, e.g., [11]. Forward-looking responsibility refers to being expected or required to behave or act in a certain way. In other words if a person has a forward-looking responsibility, she has a duty to do *X*, she is bound to do X, and she ought to do *X*.

There is a relationship between forward-looking and backward-looking forms of responsibility, in that, under certain circumstances forward-looking responsibility implies backward-looking responsibility [11] or forward-looking responsibilities can be translated into backward-looking responsibilities and vice versa [10]. That is, if one has a duty to ensure some state of affairs, *X*, and one does in fact ensure that *X* is the state of affairs, one would be morally responsible for *X* and praised for it. If one failed to ensure this state of affairs (and the situation was within one's control), one would be morally responsible and blamed for *X* not being the state of affairs.

Liability responsibility, i.e., being identified as the person (or legal entity) that will pay for the harm caused, involves both backward-looking and forward-looking elements. It is backward-looking in the sense that it is attributed on the basis of past causal contributions to outcomes (e.g., the sinking of the ship) but implies a duty to rectify matters (e.g., paying for the harm caused) [12].

In this chapter all notions of responsibility are relevant that can be ascribed to users of AmISs and the designers of AmISs, whether forward- or backward-looking. The next section discusses when responsibility can be ascribed.

1.2. Ascribing responsibility

As this chapter focuses on the responsibilities that users and designers of AmISs can have, it is important to realize under which circumstances responsibility can be ascribed, which differs for different forms of responsibility.

Ascribing backward-looking responsibility (e.g., outcome or liability responsibility in Vincent's taxonomy), involves praising or blaming someone for a rightful or wrongful outcome to which her actions contributed. As such praise or blame can have consequences for the praised or blamed party, the ascription of this form of responsibility should be fair, cf. e.g. [10]. Doorn lists five conditions to be fulfilled in order for ascription of backward-looking responsibility to be fair (on a merit-based view of responsibility ascription), namely moral agency, voluntariness or freedom, knowledge of the consequences, causality, and transgression of a norm [10].

The moral agency condition is understood as the condition that the responsible actor possesses adequate mental faculties at the moment of engaging in an action, and thus is related to capacity responsibility.

The voluntariness condition is the condition that the actor voluntarily performed the action that led to the outcome for which she is responsible. Taking this condition into consideration, actions performed under compulsion or other circumstances outside of the actors control are reasons to refrain from ascribing backward-looking responsibility to that actor.

The knowledge condition of responsibility as understood by Doorn is the condition that the actor knew or could have known the outcome of her actions [10], again related to capacity responsibility. We will discuss the knowledge condition of responsibility in more detail in the next section.

The causality condition means that the actor's action or inaction is the cause or contributed causally to the outcome, i.e., causal responsibility.

Finally, the condition that there is a transgression of a norm entails that responsibility can only be ascribed fairly if the actor's action that causally contributed to the harmful outcome infringed on a relevant norm. The source of norms to be considered comes from either the role someone plays, or from the cultural setting, or from the law. Relevancy of norms is situational and is hard to describe in general terms.

The ascription of forward-looking responsibility is somewhat different. As forward-looking responsibility is related to role and liability responsibility, identifying conditions for fair ascription of forward-looking responsibility is a question of identifying sources of duties and obligations. [13] mentions four criteria for fair attribution of forward-looking responsibility: the agent does what she ought to do; the agent is able to do what she ought to do; the agent knows what she ought to do; the norm-violation is not unintentional.

Most of Mastop's conditions are intuitively clear; it seems unfair to obligate someone to do something if she is unable to do that thing. Similarly, it seems unfair to obligate someone to do something if she does not know she ought to do that thing. The not unintentional norm-violation condition seems to be the forward-looking dual of the voluntariness condition of [10].

However, the first criterion for ascribing forward-looking responsibility fairly, i.e. the agent does what she ought to do, requires some discussion. This criterion seems to suggest that forward-looking responsibility be ascribed at the moment the agent is performing the action that the responsibility requires. We argue that forward-looking responsibility refers to the obligation of behaving and acting in a responsible way.

Note that the definitions of ascribing both backward-looking and forward-looking responsibility contain a knowledge condition. This is the fundamental reference as to why knowledge implies responsibility. As the focus of this chapter is on the relationship between knowledge and responsibilities, the knowledge condition and its limits are discussed in more detail in the next section. In this discussion, the other conditions are assumed to have been fulfilled.

1.3. The Knowledge Condition of Responsibility

Both ascriptions of backward-looking and forward-looking responsibility involve a knowledge condition. The previous section described under which conditions responsibility can be ascribed. Ascription can be done by others, but of course also by the person herself. When considering users and designers of AmISs the question is which are the responsibilities that can be ascribed to them? We use the following argument:

If people know that certain responsibilities result from having a piece of knowledge, then if people have that piece of knowledge, the resulting responsibilities are also theirs, and they should be aware of the fact that they have those responsibilities. So, assuming all other ascription conditions have been met, we can schematize the argument referring informally to a logic including knowledge and responsibilities.[2] The first step is to conclude that a person having a piece of knowledge, in fact should carry the corresponding responsibility:

$\forall i\colon K_i\,(k) \rightarrow\ R_i(r)$	Knowing k implies carrying responsibility r
$K_a(k)$	a knows k

$R_a(r)$	a carries responsibility r

In order to be held responsible, as we argued, the person should also be aware of her responsibility. Above, we assumed that the people know that certain responsibilities result from having a certain piece of knowledge, this adds an extra layer in the assumptions:

$K_a(\ \forall i\ K_i\,(k) \rightarrow\ R_i(r)\ )$	a knows that knowing k implies responsibility r
$K_a(k)$	a knows k

$K_a\ (R_a(r))$	a is aware that she carries responsibility r

The question now is under which circumstances the preconditions of the above arguments can be taken to hold. Knowledge is a condition for responsibility in that responsibility can only be ascribed (assuming other conditions have been fulfilled) if the actor to which we seek to ascribe responsibility knew that his act would lead to a certain outcome and knew that this outcome was right or wrong. The first part of George Sher's Full Epistemic Condition (FEC) of responsibility is equivalent to this view. The Full Epistemic Condition of responsibility states that [14 p. 143]:

> When someone performs an act in a way that satisfies the voluntariness condition, and when he also satisfies any other conditions for responsibility that are independent of the epistemic condition, he is responsible for his act's morally or prudentially relevant feature if, but only if, he either
>
> 1. is consciously aware that the act has that feature (i.e., is wrong or foolish or right or prudent) when he performs it; or else
> 2. is unaware that the act is wrong or foolish despite having evidence for its wrongness or foolishness his failure to recognize which
>   a) falls below some applicable standard, and
>   b) is caused by the interaction of some combination of his constitutive attitudes, dispositions, and traits; or else
> 3. is unaware that the act is right or prudent despite having made enough cognitive contact with the evidence for its rightness or prudence to enable him to perform the act on that basis.

[2] This is meant as a way of schematically presenting our argument, rather than as a fully formal account. We do assume the knowledge modality to be closed under elementary logical rules like modus ponens and universal instantiation, as is the case in standard epistemic logic. The R-operator serves as a predicate rather than a modality.

For example, a nurse in a care home, Donna, empties a bottle of rat poison into the ice machine in the recreation room. Donna knows that the air conditioning in the care home is faulty, so residents are likely to put ice in their beverages. Furthermore, Donna has read and understood the label of the bottle of rat poison, which clearly states that under no circumstances should human beings ingest the contents of the bottle. Over time, residents start to become ill (not dying, as the rat poison was diluted in the ice machine's water supply). Eventually, the care home physician finds out that this occurred as a result of poisoning. Most people would be inclined to blame Donna for the residents' illness. Assuming other conditions for ascribing responsibility have been met, Donna would be blamed because she knew that adding rat poison to something people would ingest would lead to their becoming ill. Furthermore, she knew that the outcome of such an act (and the act itself) was wrong.

Donna's actions exemplify the knowledge condition at work. Knowing that one's actions could lead to a certain outcome and knowing that that outcome is morally wrong are grounds for ascribing responsibility. It could be the case that one does not possess such knowledge in performing a certain act. In cases where one lacks the knowledge that one's actions will lead to a certain outcome and/or lacks the knowledge that this outcome is wrong, we would intuitively excuse this person from responsibility on the grounds of ignorance. We take ignorance to mean "not knowing". For example, if another nurse, Horatio, brought a resident a pitcher of water filled with ice cubes, and later that day the resident died (after drinking the entire pitcher), we would not blame Horatio for the resident's death, as Horatio was ignorant of the fact that the ice cubes were poisoned.

There are, however, cases in which responsibility can fairly be ascribed to an actor, despite that actor's ignorance of the wrongness of her actions or the outcomes thereof. According to Sher's FEC, the actor who is unaware of the wrongness of her act is still responsible if she is unaware despite having evidence for its wrongness and this unawareness falls below some applicable standard and is caused by some interaction of her constitutive attitudes, dispositions and traits [14]. On a slightly different view, one can do some wrong act *A*, ignorant of the fact that it is wrong, and still be responsible for that act and/or its outcomes if at some earlier point one knowingly performed another wrongful act or omission of which the ignorance of the wrongness of *A* is a consequence. In Zimmerman's words, "culpability for ignorant behavior must be rooted in culpability that involves no ignorance" [2, p. 417]. In other words, the person to whom we attribute responsibility was at some point aware of wrongdoing.

One source for such ignorance is negligence. Zimmerman defines a situation of negligence as a situation in which a person believes she is able to perform a certain set of precautionary measures against a negative outcome of an event; believes that not performing these events yields a possibility of a negative event occurring; unjustifiably decides not to do any of a subset of the set of precautionary measures; does an act *A* without doing or attempting to do any of a subset of the precautionary measures [1].

Furthermore, ascription of moral responsibility for negligence requires not just that one actually foresee the action and its consequences (currently believe that there is a possibility the event in question will occur at some later time), but that these are *foreseeable*. Zimmerman defines this foreseeability as follows. "[An event] *e*'s occurring at [time] *t2* is foreseeable (to some degree) by *P* if and only if it is possible that *P* foresees at *t1* (to some degree) that *e* will occur at *t2*" [1, p. 207].

For example, a third nurse, Patsy, is bringing the residents their medication, and pouring each of them a glass of water with which to take their medication. Patsy stays

with each patient to ensure that he finishes his medication. Upon drinking his glass of water, the first patient complains that he feels unwell. Patsy dismisses his complaint and proceeds to the second patient. The second patient has the same complaint, which Patsy also dismisses. Later that day, when all the patients Patsy saw are ill and it is discovered that they were poisoned, Patsy is blamed for ignoring patients' specific complaints and not recognizing the emerging pattern, both of which were signs that Patsy was giving them poisoned water.

One might argue that ascribing responsibility to an actor despite the fact that that the actor was ignorant of or did not advert to the risk created by her act or omission depends on claiming that that actor *could have* known (or *could have* done something). On this view, the claim that actor could have done something, but actually did not, is a form of counterfactual reasoning and as such is problematic.

Such capacity judgments ("she could have known" or "she had the capacity to know") are indeed counterfactual, in that they 'suppose a world different than the actual world and ask, "what would have happened in that possible world?"' [15, p. 165] However, we agree with Moore and Hurd that such counterfactual reasoning is inevitable in capacity judgments in that there are no 'unchanged-actual-world' capacity judgments with the exception of the judgment that "we all had the capacity to do just what we did and nothing else" (Ibid.).

This issue can be circumvented, to some extent, by focusing on the morally relevant factors that the type of capacity judgments of interest here depend on in explaining why an actor did not advert to a risk, such as "she would have adverted to the risk of her omission if she were not so selfishly indifferent to the welfare of other human beings and not so hurried to finish her rounds that she did not attend to signs that something was wrong" [3]. If the actor did not advert to a risk because, in this example, she was selfish, and if she can be blamed for being selfish, then we can construct counterfactuals in which the actor was not selfish and in those possible worlds would have adverted to the relevant risks. We can then infer that she *could have* adverted to the risk in the relevant sense [15]. The actor is then responsible based on the blameworthiness of her trait (selfishness) rather than a general capacity or incapacity to know her behavior was risky.

This, in turn, raises the question of whether an actor can be blamed for being stupid or having other flaws. Moore and Hurd discuss the blameworthiness of being clumsy, stupid, weak, and indifferent and draw a number of conclusions that are relevant in this context [15]. We argue with Moore and Hurd that while the actor's flaws might not be blameworthy in themselves, the actor's failure to take precautions against the harms that her flaws are predicted to cause or to change those flaws in ways that reduce risk to others *is* blameworthy. In such cases, the actor has failed to care sufficiently about the welfare of others in light of traits that have a known propensity to cause harm.

For forward-looking responsibility, matters are slightly different. Above, we argued that a criterion for fair attribution of forward-looking responsibility is that the responsible actor knows what she ought to do. So, the relationship between knowledge and forward-looking responsibility is that if an obligation is unknown to the actor to whom we wish to apply it, then that actor has no obligation (or forward-looking responsibility). Consider that the head nurse in the care home, Ellen, expects a new nurse, Django, to change all resident's bed sheets every morning at 11:00. However,

[3] This example is based on an example provided in [15].

Ellen has not told Django that he has this obligation. In fact, no one has told Django that he has this obligation. In this case, Django does not have the obligation to change all residents' bed sheets every morning at 11:00.

2. Knowledge Based Systems Imply Responsibilities

The principle underlying Section 1.3, viz. that knowledge implies responsibility, has far reaching consequences for the use of information and communication technology in general, and more specifically for knowledge-based systems (KBSs). With the introduction and use of knowledge-based systems, the user has an additional source of knowledge and information. Thus, the use of a knowledge-based system brings more responsibilities to the user. In particular the user might have more obligations, and might also become liable for some possible situations. The obligations refer to actions that people should take (are expected to take) and liability is in light of actions they should have taken or should not have taken. Here actions refer not only to actions that affect the physical state of the world, but also to speech acts [16,17,18]. Speech acts such as informing and requesting can set things in motion in a way that is deemed responsible by the society.

Traditional KBSs, like medical expert systems, are typically designed to be used by (single) educated professional users in a clearly specified context of application. In contrast, Ambient Intelligence Systems (AmISs) are designed for more complex situations. AmISs have a number of features, such as sensitivity, responsiveness, adaptivity, transparency, ubiquity and intelligence [19]. These features have an impact on the relationship between AmISs and knowledge. Sensitivity, responsiveness, adaptivity and intelligence imply that AmISs reason about and act upon information they gather from their users and environment. Ubiquity implies that AmISs are applied in a diverse and possibly dynamic range of contexts. Transparency implies that (intended but also unintended) users will not always be aware of the systems they are subject to.

These features are knowledge-related and show that the responsibility implications of AmISs are more intricate than for traditional KBSs. First, there are more and more diverse users involved. For example, we can distinguish between the experts involved in the use of the system (*professional users,* e.g. doctors), and the lay-people (*non-professional users,* e.g. patients that are being assisted and monitored). Second, the *ambient (or ubiquitous)* feature of the technology, implies more flexibility in the contexts of application. Thus, chances are higher that there will be unintended users and usages of the system. It is up to the designer to see to it that the technology is good for the intended users and usages, and anticipate as much as possible the responsibilities that might ensue from using the system.

This means, that different responsibilities follow for the different types of users of an AmIS. Each type of users, has a different knowledge base. To simplify we distinguish simply between a professional user (e.g. the medical specialist) and a layman user (e.g. the patient). How the combination of this personal knowledge base with the knowledge base provided by the AmIS will also be different for different users, as their ways of combining the knowledge elements to a new combined knowledge base will be different: a professional can be expected to have a more sophisticated set of reasoning rules. Hence, there is a dependency of the combined knowledge base on

both AmIS and user (U) as follows, as witnessed in the subscripts of the following scheme:[4]

$$\overline{k_{AmIS}}, \overline{k_U} \quad |-_{AmIS,U} \quad \overline{k_{U,AmIS}}$$

By combining the knowledge of the AmIS with the personal knowledge, information might be deduced, that is not deducible from either the knowledge base $\overline{k_{AmIS}}$ alone or from $\overline{k_U}$ alone. Thus, new responsibilities can arise that are neither the responsibility of the AmIS (alone), nor that of the user alone (i.e., when not using AmIS). As the human user is the only one (of $AmIS$ and U) who is capable of inferring the new piece of combined knowledge, the user is the one that gains the responsibilities that come with inferred knowledge. So the use of AmISs leads to new responsibilities for its user.

It is important as a developer to be aware of this dependency when designing in the reasoning capacities of the AmIS: it becomes the responsibility of the developer to design the reasoning capacities fit for the right knowledge, reasoning capacities and intended use for each user type. Furthermore, the combined knowledge of the user, the knowledge of the developer of the knowledge base and reasoning capacity of the AmIS, give a combined responsibility for the responsible use of the AmIS. The designer is responsible for informing all types of users about the way the AmIS is designed, and what responsibilities follow from that knowledge for each (intended) user type.

Sher's Full Epistemic Condition, as discussed in Section 1, has an important implication for the designers of AmISs. With this chapter, we aim to make the designers aware of the responsibilities that the users of their AmISs will carry, and what this implies for the responsibility they themselves carry. In particular: by reading this chapter, designers will become consciously aware of the fact that the users of their AmISs will carry responsibilities, and thereby satisfy FEC. Remember the second argumentation scheme we presented in Section 1:

$K_d(\ \forall i\ K_i\ (k_d) \rightarrow\ R_i(r_d)\ )$
$K_d(k_d)$

$K_d\ (R_d(r_d))$

Taking the designer as d, the contents of this chapter as $k \epsilon k_d$, and the responsibilites of the designer for making users aware of their respective responsibilities as $r \epsilon r_d$, we may conclude that we may ascribe the responsibility r to the designer of AmISs.

For those designers who do not read this chapter, one could argue that prudence associated with their profession, requires them to read chapters like ours. This implies that designers reading this chapter can reasonably foresee that the use of the knowledge system they construct leads to responsibilities for themselves, and for the users of their knowledge system. Thereby they have the responsibility to design the AmISs in such way that (at least the intended) users will be made aware of the responsibilities that come with the use of the system (and that the user herself might not anticipate).

[4] We use the overline to indicate sets rather than single pieces of knowledge.

3. Independent Living for the Elderly: A Case Study

A prominent application domain for Ambient Intelligence is health care. The leading case study presented in the remainder of this chapter is based on the CCF2 project [4] that one of the authors was involved in, in which an AmIS was developed to support independent living for the elderly, and in particular to stimulate social and physical activity of elderly residents of a care home. This section describes the stakeholders in the CCF2 project, their knowledge and responsibilities prior to the introduction of the AmIS.

The group of professional caregivers consists of general practitioners, nurses, occupational therapists, supervisors, and assistants to occupational therapists. Each professional is expected to know enough about the residents under their care to be able to carry out their profession. Of course they can be expected to have the types of knowledge required to practice their profession. Furthermore, they have their professional responsibilities for every resident (even those not directly under their supervision). Finally, as members of a social community and being of general good health and mental capacity they have general responsibilities for everyone in their surroundings.

From most general to most specific these responsibilities can be formalized in schemes. Starting with the general responsibilities of members of a social community, a scheme is set up that reflects general norms and their associated responsibilities.

Note that this is an operationalization of the definitions in Section 1. The rules can be applied in a forward-looking way to determine someone's responsibilities for now and the future. Also rule 2 can be used in backward-looking way to determine whether or not a person should have taken responsible actions at a time in the past. That means that current_state(S) should refer to what was current at that time in the past.

For every person A that is a member of a social community C, if V is a value in the community, A knows (or should know) that V is value in C, and R is a responsibility associated with that value, and A knows (or should know) that R is a responsibility for value V, then A has that responsibility R.

Member (A,C)
Value_in(V,C)
K_A(value(V,C))
Responsibility_for_Value(R,V)
K_A(Responsibility_for_Value(R,V))
---------- (Rule 1)
Has_responsibility(A, R)

The following applies for any responsibility of any person: In a current situation S, if a person A knows (or should know that) some responsibility R applies in situation S, and P knows (or should know) that action α might be an appropriate action in light of responsibility R, then A is obliged to perform action α.

Current_state(S)
Has_responsibility(A,R)
K_A(applies(R,S))
K_A(appropriate(α,R))
---------- (Rule 2)
O_P(α)

To make it concrete, suppose a person p is a member of the social community of the care home, and consider the general value of fairness, for which a corresponding responsibility is to treat people fairly. Then person p has the responsibility to treat people fairly. Thus, if a situation arises where some resident is excluded from a game (without sufficient cause), and if p knows that speaking up for the excluded resident is appropriate, then p should speak up for the excluded resident.

The general responsibilities of professionals are more specific than the above. Every agent A having profession P has the role responsibilities R that come with her profession:

Is_a_profession(P)
Has_profession(A,P)
Role_responsibility_for(R, P)
---------- (Rule 3)
Has_responsibility(A,R)

For example, a nurse has the responsibility to provide medical assistance to a person in need for as far as her capabilities go, and call for help if needed. This can be put into effect by applying Rules 2 and 3 in combination with the following facts:

Is_a_profession(nurse);
Role_responsibility_for(medically_assist_those_in_need, nurse);
Role_responsibility_for(call_medical_assistance_if_needed, nurse);

For example, if a nurse sees that a resident has trouble drinking, and no other help is provided (which corresponds to the knowledge condition that she knows that her helping the resident to drink is an appropriate action), then she knows she should help the resident to drink.

Furthermore, a professional might have more specific responsibilities for those patients assigned to her care. For example, among the role responsibilities for occupational therapists is the duty to study the patient record of each resident H assigned to them, to set up a partial care plan for that resident H with respect to the professional's expertise and taking into account the partial care plans of other professionals for that resident, and to explain that partial care plan to the assistant S assigned to H so that assistant S can carry it out. Some of this is easily formalized according to the earlier schemes:

Role_responsibility_for(study_patient_record, occupational_therapist);
Role_responsibility_for(set_up_partial_care_plan, occupational_therapist);
Role_responsibility_for(explain_partial_care_plan, occupational_therapist;

However, per role responsibility more knowledge is required to adequately perform that role. For example, the role responsibilities to study a patient record, set up a partial care plan, and to explain a partial care plan should be performed only for the patients H assigned to the occupational therapist A.

Is_assigned_to(H, A)
Has_responsibility(A, R)

Has_responsibility(A, R(H))

More specific knowledge is required to deal with studying a patient record (e.g., only study what you need to know to perform your role), or explaining a partial care plan. As an illustration:

```
Is_assistant_of_for(S,T,H)
Has_responsibility(T, explain_partial_care_plan(H))
---------------------------------------------------------------------------
Has_responsibility(T, explain_partial_care_plan_of_to(H,S))
```

The main tasks of an assistant of an occupational therapist (o_t_assistant), as can be seen in the scenarios is receiving notifications of residents' level of physical activity and acting on this if necessary. For these tasks, the o_t_assistant has to monitor the physical activities, execute the partial care plan, and notifying the occupational therapist if the resident is not active enough for some time. Furthermore, the o_t_assistant is to stimulate active behavior by visiting the resident at least once a week and motivate the resident to schedule and participate in some activities.

```
Role_responsibility_for(monitor_physical_activities, o_t_assistant);
Role_responsibility_for(execute_care_plan, o_t_assistant);
Role_responsibility_for(notify_o_t, o_t_assistant);
Role_responsibility_for(stimulate_physical_activities, o_t_assistant);
```

The group of activity organizers is another group closely related to the group of professional caregivers in the CCF2 case. Their professional knowledge and skills include and pertain to organizing activities for individual needs and for groups. Their responsibilities include maintaining and making available a list of upcoming activities, organizing those events, and keeping track of which residents participate.

Finally, of non-professional caregivers only family members of residents were considered in the case study. Family members cannot be expected to have professional knowledge of care giving, although it cannot be excluded that they have such knowledge. Family members are attributed the general responsibilities that are attributed to any person of general good health and mental capacity. Furthermore, family members are expected to contribute and use their knowledge about the resident in terms of general health issues, preferences and values. The main duty of this group, in the case of elderly care, is to tend to the perceived needs of an old parent [20]. This can include assisting with activities of daily life (ADLs), such as bathing, dressing, moving around indoors, transferring from bed to chair, using the toilet or eating, and/or instrumental activities of daily life (IADLs) such as light housekeeping, meal preparation, grocery shopping, laundry, taking medications, managing money, telephoning, outdoor mobility and transportation. This is formalized directly in terms of having responsibilities, for example person A is a family member of resident H that has accepted the responsibility for bathing H:

```
Family_member_of(A, H)
Has_responsibility(A, bathing(H))
```

In the CCF2 case, (non-adult) grandchildren are specifically included in the family group of stakeholders. While not considered caregivers they interact with the system

nonetheless, gaining and giving insight into their grandparents' preferences and daily activities.

4. Knowledge and Responsibility in Ambient Intelligence Systems

Ambient Intelligence Systems make use of and produce knowledge in to serve a variety of users. Their key features are sensitivity, responsiveness, adaptivity, transparency, ubiquity and intelligence [19]. The intelligence is typically spread over a number of contributing agents that by working together form an AmIS.

AmISs are sensitive in that they sense the environment in which they are situated. That is, AmISs obtain information about their operational environment. AmISs are responsive in the sense that they respond to (the presence of their) users, which is achieved, in part by modeling users, recognizing patterns, making decisions, and reasoning [19]. Adaptivity similarly involves reacting to information, in this case by changing (e.g., user models in the system). Transparency refers to the notion of the disappearing computer, that is, the idea that computing technology should weave itself into and thus become part of the environment [21]. The ubiquity of an AmIS means that it is present everywhere in the area where the AmIS is active. This ubiquity and the required intelligence make agent technology a driver for AmISs.

In all cases the AmIS is created for the use of some people, i.e., the direct stakeholders and might (inadvertently or intentionally) affect others, i.e., the indirect stakeholders.

4.1. Stakeholders, their knowledge and responsibilities

Though stakeholders of AmISs differ from domain to domain and given a domain from application to application, in health care the primary stakeholders are arguably always patients. Patients can be users of AmISs (e.g., in applications for assisted living), or indirect stakeholders as data subjects who do not interact with the system directly (e.g., in monitoring applications). In the case of smart homes for the elderly and people with physical disabilities, the group of patients can include non-disabled users, disabled users, elderly people, people with diminished vision, hearing impaired people, and cognitively impaired people [22]. The diversity of these types of patients makes it impossible to make general claims about the knowledge this group of stakeholders possesses. As a result, identification of (forward-looking) role responsibilities of this group of stakeholders is impossible. In the CCF2 case, the patients are the elderly residents of a care home that are not cognitively impaired. They are attributed a responsibility for their personal health (i.e., to look after their own health), although this responsibility is shared with the caregivers.

The second main group of stakeholders consists of the caregivers that can be sub-typed depending on the context of care and whether or not they are professional care givers. In a hospital or care home setting, caregivers include physicians of various specialisms, nurses (again in various types), and (non-professional) volunteers. As care shifts away from centralized, expert-driven models to one in personal living spaces in response to an ageing population, informal caregivers, such as family, friends, and members of the community are involved [23].

Professional caregivers can be expected to have the necessary knowledge and skills for their profession and are expected to behave according to the responsibilities

that come with their profession. In general terms for care giving professionals are expected to have the knowledge and skills to perform medical diagnosis (e.g., knowledge of anatomy, physiology, pathology, as well as operational knowledge about how to conduct diagnostic tests, how to interpret results of such tests) and/or the knowledge and skills how to treat what has been diagnosed (e.g., administer painkillers to a patient who is in pain).

Non-professional caregivers can be subdivided into those coming from the social network of the patient and unrelated volunteers. The basic difference between them is the reason why they have become caregivers. Family members and friends might take on care giving responsibilities because of social norms indicating the responsibility of taking care of family and those close to you [20]. Even so, out of a social network some people are more ready to take on these responsibilities than others. General volunteers might lack the social bond with the patient, but typically feel they contribute to society and the individual patients by offering their help.

The knowledge and skills for taking care of others of the group of non-professional caregivers is diverse. A young person, inexperienced in taking care of others, might not realize that she needs to check the temperature of a bath before placing someone in that bath. People that took care of their own small children can be expected to know about checking the temperature of the bath. As the background and intelligence of the large group of non-professional caregivers cannot be established upfront, only the general responsibilities that correspond to common sense can be attributed to the group as a whole. However, their individual background and intelligence might place additional individual responsibilities on them.

When introducing AmISs the knowledge on how to use the technology becomes relevant. Professionals are expected to make sure they know how to use the system, either through training or by studying its use themselves [23], depending on the use complexity of the system. However, who will make sure that non-professional caregivers that are allowed to use the system get the necessary training or instructions for properly usage? Fundamental to this chapter is the question: how do AmISs affect the knowledge condition of stakeholders and thus their responsibilities?

4.2. AmIS Architecture for Health Care Applications

AmISs are equipped with a set of sensors to gather data from the physical environment, and consist of intelligent agents to reason about that data, and to act upon that data to achieve the goals of the AmIS [19].

Sensors are devices that produce a measurable response to changes in their physical environment, e.g., temperature, humidity, movement, which electrical apparatus are in operation. Sensors can be arranged in wireless sensor networks (WSNs) to enable spatio-temporal sampling of physical, physiological, cognitive, and behavioral processes with a density that was previously impractical [23]. For health care monitoring, five basic subsystems have been identified [24]:

1. Body Area Network (WSN)
2. Personal Area Network (WSN)
3. Gateway to the Wide Area Networks
4. Wide Area Networks
5. End-user health care monitoring applications

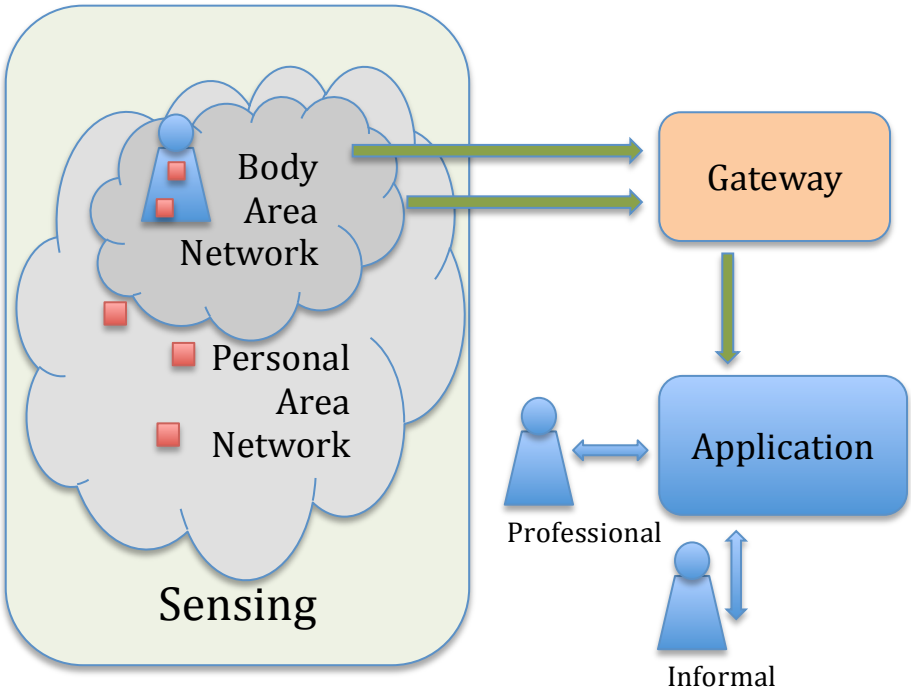

Figure 1. Monitoring subsystems (based on [24]).

Examples of components of Body Area Network (BAN) subsystems include RFID tags, electrocardiogram (ECG) sensors, and accelerometers worn by the patient [24]. Personal Area Networks (PAN) can consist of environmental sensors such as RFID readers, pressure sensors, luminosity sensors, and temperature sensors that gather contextual information about the person (or people) being sensed. Raw data of changes in one aspect of the environment, such as temperature, might not be useful in isolation. Therefore, data is fused. Data fusion is the process of putting together information obtained from many diverse sensors into a single representation of the environment [25]. These data can then be reasoned about by agents, e.g., by modeling the user, recognizing and predicting activity, and making decisions [19].

The gateway subsystem connects the BAN and PAN subsystems to wide area networks (WANs). The WANs comprise the infrastructure that distributes information from the BAN and PAN subsystems to the end-user applications. Agents can be employed at the sensor side, but also at the end-user application side. Finally, the collected information comes together in end-user applications (typically agents) that process and interpret collected data, and trigger actions. Additionally, (in the case of health care monitoring application) graphical user interfaces enable users to monitor vital signs in real time and receive alerts in case of emergencies [24].

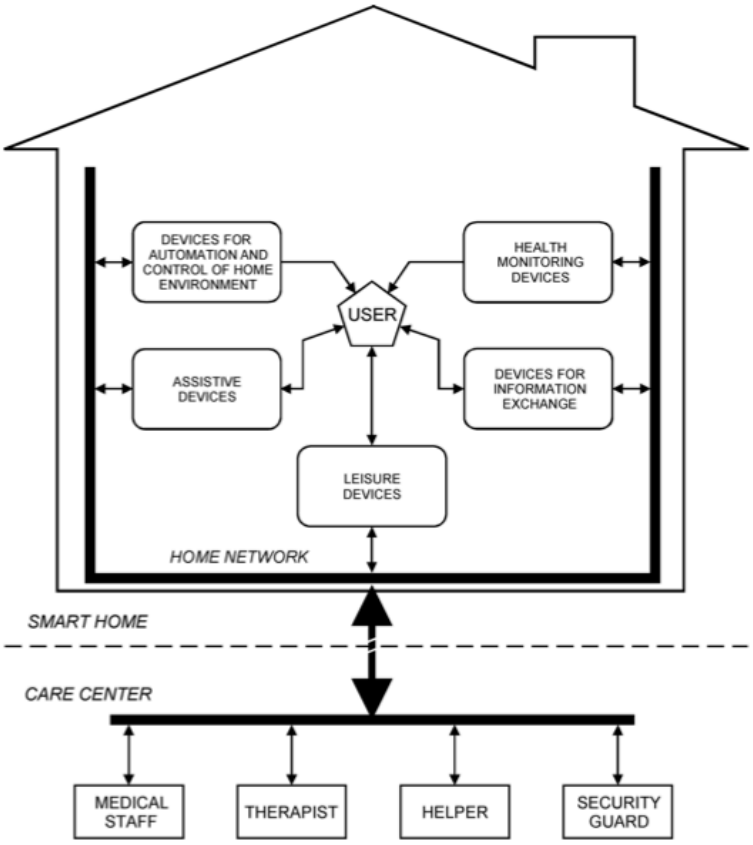

Figure 2. Smart home for people with disabilities [22].

Stefanov and colleagues [22] identify a number of types of home-installed devices for smart homes for the elderly and people with disabilities. Among those mentioned are devices for health monitoring and devices for information exchange. Devices in the former category include devices for monitoring vital signs, posture monitoring, behavior monitoring, recognition of facial expressions, and advanced chemical analysis. Devices in the latter category include systems for information access and telecommunication, systems for telemonitoring, teleinspection and remote control, and home network [22].

4.3. An AmIS for Independent Living for the Elderly

In the CCF2 project an AmIS was developed to support independent living for the elderly, and in particular to stimulate social and physical activity of elderly residents of a care home. The stakeholders identified for the system are described in Section 3.

Three scenarios were the point of departure: maintaining a social network, coaching, and intergenerational communication and engagement, see [4]. Support for these scenarios can be constructed from the basic pervasive health care monitoring applications identified in [24]: monitoring activities of daily living, fall and movement

detection, location tracking, medication intake monitoring, and medical status monitoring.
In the first scenario, the system informs elderly residents of upcoming activities (such as group walks) and of which other residents intend to participate in these activities, and allows them to share their preferences for participation with other users. By gathering sensor proximity information, the system can infer how frequently elderly residents have been near other residents, and hence how socially active they have been. The system also informs residents' caregivers of residents' participation in activities.

In the second scenario, agents in the system monitor residents' (levels of) physical activity and quality of sleep (inferred from the amount of movement registered by bed-mounted sensors). The agents notify the residents' caregivers if residents are not active enough, so that they can urge them to be more active. Personal coaching agents provide elderly residents with feedback, so that they can be urged to be more active or, if they are too active, they can be persuaded to reduce their level of activity. Furthermore, by aggregating information from all residents, the AmIS lets residents compare their health and physical activity with that of their peers.

The third scenario is about monitoring the communication between elderly people and their grandchildren. The system supports elderly residents in keeping track of their grandchildren's activities and vice versa, and in keeping each other up to date on daily achievements. The system encourages and enables elderly residents to remotely participate in their grandchildren's games.

To realize the functionality sketched in the scenarios the CCF2 AmIS was equipped with fixed and mobile sensor nodes. The fixed sensor nodes are deployed in fixed locations throughout a care home to form a wireless sensor network. For example, the resident's beds are fitted with accelerometers.

The mobile sensor nodes worn by elderly residents contain accelerometers, photometers that measure the amount of light reaching the sensor, and a thermometer. Some patients wear a heartbeat sensor.

Data from the sensors is used to derive all kinds of information. For example, data from the accelerometers can be used to determine the activity and mobility of the person wearing the sensor node. Combining this with data from the heartbeat sensor can provide information on how the amount a resident is moving affects her heart rate. The photometer and thermometer give general information about the wearer's environment that, in combination with other sensor data can be used to detect the activity of the wearer and even whether the user is indoors or outdoors (in case localization fails). The accelerometers in residents' beds measure how much the resident moves in bed which is an indicator of the resident's quality of sleep. The fixed sensor nodes can by cooperation determine the location of mobile sensor nodes. If multiple mobile sensor nodes are detected in each other's proximity, this provides an indication that the wearers of the mobile sensor nodes are engaged in a social activity. If there is a sudden acceleration detected by the accelerometer, and localization indicates that the mobile sensor is no longer changing location, this could indicate that the person wearing the mobile sensor node has fallen (or the mobile sensor node has fallen).

The CCF2 system contains an Activator agent that aims to persuade elderly residents to plan and participate in social activities [4]. The Activator agent notifies residents of activities that are going to take place and supports residents in self-monitoring their activities by providing them insight into how much physical and social activity they have had in the recent past.

Information on upcoming activities the care home is organizing can be entered into an activity database by activity organizers. A resident database contains information about whether or not residents will participate. When residents accept or decline an activity, this gets registered in the resident database. The Activator pulls the number of people that are joining an activity from this database, and presents this to residents and activity organizers.

A Communal agent was developed to support activities in communal areas of the care home. Like the personal Activator, the Communal agent also displays upcoming activities, and measures the actual participation and activity of an ongoing event.

So, the CCF2 AmIS involves a combination of monitoring daily activities, detecting movement, and tracking location. These techniques are used for giving elderly residents of a care home insight into the social and physical activity levels of themselves and other residents. Also, professional caregivers as well as the social network consisting of some family members and friends can monitor activity levels of residents. In an abstract sense, the CCF2 AmIS aims at self-awareness (to persuade elderly users to adopt healthy behavior), and keeping professional and informal care networks aware. Hence, different types of knowledge are produced for different stakeholders.

4.4. Knowledge Production and Responsibilities in Ambient Intelligence

AmISs make various kinds of information available to users. Combined with the personal knowledge of their users, information can be deduced that is not deducible from the information produced by the system or users' knowledge alone.

In this section, we will examine how each stakeholder group we discussed above accesses information (made available by the underlying technology described above) and what knowledge each group gains based on the prior knowledge each group has. We will limit our discussion to the CCF2 case.

4.4.1. General practitioners

4.4.1.1. Pre-existing knowledge

General practitioners acquire much of their knowledge through their studies and in practice. This includes knowledge of standards, for example, the standards of the Dutch GP association (NHG) in the Netherlands.

Also, GPs must have knowledge of the International Classification of Primary Care (ICPC), provided by the World Health Organization [26]. The ICPC is a method for the classification of the treatment required for a patient presenting herself with a request for primary care. For this aim, the ICPC contains 17 chapters corresponding to different aspects of the human body, each divided into 7 components corresponding to different aspects of medical treatment.

Besides this general medical knowledge, GPs have knowledge of specific patients, consisting of a medical history and clinical signs. A medical history, is information gained through asking specific questions about demographics; chief complaint; history of present illness; past medical history; review of organ systems; family diseases; childhood diseases; social history; regular and acute medications; allergies; sexual history. Clinical signs are gained through clinical examination. These can include vital signs such as temperature, blood pressure, and pulse.

Finally, GPs have knowledge of the responsibilities that come with their profession, which they gain from the Hippocratic oath and applicable codes of behavior and/or ethics.

4.4.1.2. Pre-existing responsibilities

The professional responsibilities of GPs consist of responsibilities derived from Hippocratic oath, including beneficence, non-maleficence and confidentiality. Also, GPs have responsibilities that are derived from codes of ethics, such as the World Medical Association's (WMA) International Code of Ethics, which includes responsibilities such as bearing in mind the obligation to respect human life.

Further, GPs have a responsibility to keep professional knowledge up to date. For example, in the Netherlands, GPs must take an exam twice a year testing their knowledge of the ICPC and of theoretical aspects of their profession.

4.4.1.3. New knowledge

Through use of the system, GPs gain detailed knowledge of patients' physical and social activity. This includes physiological data, location data, proximity to other people, and activities participated in. Before the introduction of the AmIS, GPs had to rely on patient's accounts of activities, complaints, and previous illness, and so on. That is, GPs had to take a medical history and rely on this to form a diagnosis and provide medical care. Additionally, GPs could perform tests (such as taking blood pressure) on site or order them elsewhere (e.g., blood tests).

The type of AmIS described here makes detailed information continuously available, providing a detailed picture of patients' condition over time. From the GPs point of view, it is as if the patient is undergoing continuous clinical examination and the GP can look into the resulting clinical signs at a given moment or over a longer period, gaining knowledge of trends and developments in a way that previously was not feasible.

4.4.1.4. Improved fulfillment of responsibilities

Due to the type of knowledge the system makes available, GPs are better able to diagnose and treat their patients, and to evaluate the outcomes of that treatment.

4.4.1.5. New responsibilities

The continuous availability clinical signals could lead to new responsibilities for GPs. Specifically, the availability of this information, combined with GPs' duty of beneficence, could lead to a responsibility to constantly act based on the information that is available to them. Not doing so could be considered a wrongful omission, as the GP could have known she was required to act. Fulfilling such a responsibility to act (at all times) on the continuous stream of clinical signs is not feasible, but this does raise an important question of who should monitor this information and who should act on it.

4.4.2. Nurses

4.4.2.1. Pre-existing knowledge

As a result of their training, nurses can be expected to have basic medical knowledge. Knowledge of care practice is part of this knowledge and includes:

- How to assess and diagnose needs
- How to plan outcomes, interventions
- How to implement interventions
- How to evaluate outcomes of care provided

- Knowledge of drug interactions
- Knowledge of specific patients' needs
- Knowledge of specific patients' care plans

More specifically, geriatric care nurses will have knowledge of special needs of elderly and afflictions particular to the elderly

4.4.2.2. Pre-existing responsibilities

Nurses have a responsibility to assess and diagnose patient needs. Based on this, they have a further responsibility plan outcomes and interventions. They must implement interventions. Furthermore, they must evaluate outcomes of care. They also have a duty to inform patient of risks involved in various treatment options. As with other professional caregivers, it is essential that nurses keep their professional knowledge up to date.

4.4.2.3. New knowledge

The information made available by the system gives nurses detailed knowledge of patients' physical and social activity. Regarding physical activity, sensors continuously and ubiquitously collect information on location, amount (and speed) of movement, and pulse of the patient, as well as lighting conditions and temperature in the patient's environment. This basic information gives detailed knowledge of location, physical properties of location, and patients' physiological state through the course of the day (or any other time interval). Information on proximity to others and participation in social activities provides knowledge of levels of social activity.

Combined with the knowledge nurses already have of individual patients, their medical histories, their needs, and their care plans, knowledge emerges of care needs and outcomes (in different settings) at an entirely new level of detail and frequency.

4.4.2.4. Improved responsibility fulfillment

The knowledge that nurses can gain from the system makes it easier for them to fulfill their responsibility to assess the care needs of their patients. More specifically, by gaining insight into how every moment of the day is spent and how patients' physiological condition changes throughout the day, nurses can better assess what kind of care is needed in certain circumstances. This knowledge also makes it easier for nurses to evaluate outcomes of the provided care.

4.4.2.5. New responsibilities

With new detailed continuous and ubiquitous knowledge of patients' activities and resulting detailed knowledge of care needs and care outcomes, it could be argued that a new responsibility arises to adjust care plans (and practice) whenever new information on needs and outcomes is available, which, with the introduction of the AmIS, is always. However, a nurse cannot realistically carry out the continuous adaptation of care plans for all of her patients. Hence, the introduction of the AmIS could potentially lead to a responsibility that nurses could not fulfill.

If a nurse were to know that a resident's behavior is unhealthy and the nurse, knowing that she should generally look after residents' health, decided not to act on this knowledge, then backward-looking responsibility for this wrongful act (omitting to look after the elderly resident's health) could fairly be attributed to the nurse. As for forward-looking responsibility, we argued that nurses have a general duty (or there is a norm) to look after elderly resident's health regardless of system use. Again, what the system changes, is the frequency with which the specific duty can arise. That is,

whenever a nurse gains knowledge that a resident is acting in an unhealthy ways, a specific duty arises to ensure that that behavior is changed.

4.4.3. Family members

4.4.3.1. Pre-existing knowledge

Family members have intimate knowledge of elderly person, such as knowledge of the elderly person's identity and preferences. They can also be expected to have some form of knowledge of the elderly family member's basic physical condition, which they gain through direct observation during visits, from what their loved ones tell them, and from what staff members report.

Beyond knowledge of the elderly family member, family members with a caregiver role also have knowledge of "best care", that is, a vision of what criteria care should meet. As these family members monitor care provision to their loved ones (through direct observation and querying elderly loved one), they will have some knowledge of how care is conducted. Finally, family members will have some knowledge of responsibilities of care home staff and of their own responsibilities.

4.4.3.2. Pre-existing responsibilities

One of the primary responsibilities family members see themselves as having is overseeing the care of the loved one, representing the resident's perspective and history, and keeping the family connections [27]. Further responsibilities perceived by family members include:

- maintaining continuity (based on relative's intimate knowledge of the older person), which involves helping the older person to maintain their sense of identity through the continuation of loving family relationships and through helping staff to get to know the resident as an individual;
- keeping an eye, by monitoring the care received, providing feedback to staff and filling any gaps; and
- contributing to community through interacting with other residents, relatives and staff, taking part in social events and generally providing a link with the Outside world [28].

4.4.3.3. New knowledge

With the introduction of an AmIS, family members gain knowledge of how their loved ones spend their days in terms of physical and social activities. This goes beyond the knowledge family members normally gain from direct observation in visits, from what their loved ones tell them, and from what staff reports. It is different in that it consists of all activities engaged in throughout the day, with additional knowledge of their loved one's specific locations, physical properties of those locations (e.g., temperature), and their loved ones' physiological responses to those activities and environments. This knowledge then gives some insight into the effects of the care the loved one is receiving, though family members will not necessarily have the knowledge required to interpret these effects.

4.4.3.4. Improved fulfillment of existing responsibilities

The new knowledge for family members that the AmIS produces improves family members' ability to carry out their duty perceived duty of keeping an eye, in that it allows them to scrutinize in greater detail and with greater frequency how care is being provided.

4.4.3.5. New responsibilities

With the introduction of the AmIS, family members potentially have knowledge of their elderly family members' activities, levels of activity, and physical condition at all times. In light of their responsibility to keep an eye on the care received family members might feel a needs to interpret the knowledge available to them in order to scrutinize care. In light of their perceived responsibility to maintain continuity by helping staff get to know the elderly family member as an individual and in light of their perceived responsibility to represent the elderly person's perspective, they might feel the need to explain the elderly person's actions and responses to activities to staff. Considering the continuous availability of information on activities and the like, such explanations could burden the family members giving them, as well interfere with caregivers' work.

4.4.4. Elderly residents

4.4.4.1. Pre-existing knowledge

Elderly residents of care homes can be expected to have knowledge of their own physical condition. The level of detail of this knowledge will vary. The most important source of this knowledge is introspection – people know something about their own physical condition because they know how they feel. A further source of this knowledge is caregivers. For example, an elderly person's GP might have told her she has diabetes. In that case, she will know she has diabetes (though, as a layperson, she might not understand the full implications of that diagnosis).

Furthermore, assuming they are not (significantly) cognitively impaired, the elderly can be expected to have knowledge of their own activities of daily life and other activities they engage in. For example, on might know (based on recollection) that this morning she woke up, had breakfast alone in her room, went downstairs to play chess, and went up to take a nap afterwards, from which she has just awakened.

Finally, the elderly should have some knowledge of what behavior is good for their health and what behavior is bad for their health. For example, an elderly person should be aware that going out in the cold without warm clothes could lead to a cold.

4.4.4.2. Pre-existing responsibilities

One primary (though debatable) responsibility for the elderly (or for any person with adequate capacities) is the individual responsibility for their own health. That is, elderly people should, to some extent, do things that are good for their health and avoid those things that are not. As residents of a care home, elderly people also have a responsibility to respect other residents and staff, and follow facility rules.

4.4.4.3. New knowledge

Information collected from sensors is presented to elderly residents through the performance state (or view) of the Activator interface. This gives them insight into how physically and socially active they have been. It also allows them to compare this to the average levels of the other users.

The combination with their prior knowledge of which activities they participated in (and with whom), how physically active they were, and how much activity is healthy for them, allows users to gain more detailed information on their levels activity than could be gained from introspection alone. For example, they could come to know how their levels of activity affected their heartbeat (though we can not expect them to have the knowledge required to properly interpret information such as heartbeat).

Furthermore, if the system shows that their activity levels have been higher or lower than the elderly residents know they should be, then they come to know that their level of activity has not been healthy and that they should be more or less active.

4.4.4.4. Improved fulfillment of existing responsibilities
Feedback on how activities influence pulse and health helps elderly better look after their health by adopting healthy levels of activity.

4.4.4.5. New responsibilities
In the CCF2 case, elderly users gain detailed knowledge about their social and physical performance. This is not just knowledge of their actual activity levels, but also knowledge of how their activity levels compare to some standard. That is, they gain knowledge that they are not as active as or more active than they *should be*. By knowing this, and deciding not to change their behavior, we could say they are responsible for the state of their (social and/or physical) well-being. If the other conditions for responsibility are met, these stakeholders' use of the system could contribute to (backward-looking) responsibility being ascribed to this group of stakeholders. Also, gaining knowledge that they are not active enough or too active could give them a forward-looking responsibility to be more or less active.

Table 1. Summary of stakeholders, their knowledge and responsibilities

Stakeholder	**Knowledge**	**Responsibilities**
General Practitioner	Standards	Hippocratic Oath (Forward-looking)
	Classification of primary care encounters	Codes of Ethics (Forward-looking)
	Medical history of specific patient	Keep professional knowledge up to date (Forward-looking)
	Clinical signs of specific patient	Monitor information (New, forward-looking)
	Knowledge of professional responsibilities (Hippocratic Oath and applicable codes of ethics)	Constantly act on information (New, forward-looking)
	Specific patient's physical activity (New)	
	Specific patient's social activity (New)	
	Environmental conditions (New)	
	Trends (New)	
Nurse	How to assess and diagnose needs	Assess and diagnose patient needs (Forward-looking)
	How to plan outcomes and interventions	Plan outcomes and interventions (Forward-looking)
	How to implement interventions	Implement interventions (Forward-looking)
	How to evaluate outcomes of care	Evaluate outcomes of care (Forward-looking)
	Knowledge of drug interactions	Inform patient of risks involved in treatment options (Forward-looking)
	Knowledge of specific patients' needs	Keep professional knowledge up to date (Forward-looking)
	Knowledge of specific patients' care plans	Continuously adjust patients' care plans (New, forward-looking)
	Knowledge of special needs of elderly	Ensure unhealthy behavior is changed (New, forward-looking)

	Knowledge of afflictions particular to the elderly	Monitor information (New, forward-looking)
	Specific patients' physical activity (new)	
	Specific patients' social activity (new)	
	Current location of patients	
	Current physiological state of patients	
	Level of social activity	
Family members	Elderly person's identity and preferences	Overseeing care of loved one (Forward-looking)
	Elderly person's basic physical condition	Representing loved one's perspective and history (Forward-looking)
	Knowledge of "best care" for family member	Keeping family connections (Forward-looking)
	How care is conducted	Maintaining continuity (Forward-looking)
	Responsibilities of care homes staff	Keeping an eye (Forward-looking)
	Own responsibilities	Contributing to community
	Physical activity (including location, physical properties of locations, physiological information (New)	Interpret available information to scrutinize care (New, forward-looking)
		Explain elder's actions to care professionals (New, forward-looking)
Elderly	Own physical condition	Taking care of own health (Forward-looking)
	Own activities of daily living	Respect other residents and staff
	Other activities	State of (social and/or physical) well-being (New, Forward-looking/Backward-looking)
	Knowledge of what behavior is good for health	Be more or less (socially or physically) active (New, forward-looking)
	Detailed physical activity (new)	
	Detailed social activity (new)	
	Average level of activity of others (new)	
	Effect of behavior (new)	

5. The Moral Implications for AmIS Developers

In the previous section we argued that the use of AmIS leads to certain knowledge in users, which, through the knowledge condition of responsibility, leads to responsibilities for those users. So, by using the type of system described here, people could gain duties and could be praised or blamed from their actions or omissions. So, these systems could affect people's lives in important ways. We argue that designers of such systems should take this into account when designing such systems. In this section, we offer initial guidelines on how to do so.

One of the main issues with the way responsibilities arise through the use of these systems is that there is a possibility that users do not realize that responsibilities arise

from using the system. We cannot expect all users to know that by using these systems they gain specific types of knowledge which, if the users act in certain ways, could contribute to them being blamed (or praised) for the outcome of their acts. Similarly, we cannot expect all users to know that by having access to certain information, they acquire certain duties. As users cannot be expected to foresee these implications of using these systems, designers should at the very least ensure that users understand, or are given the possibility of understanding, the implications of being a user.

Informing the user of the implications of use could be taken a step further by requiring users' informed consent in order for them to use the system. A problem with this is that giving the user only two options, consent (with taking on responsibilities) and use, or refrain from use altogether approaches blackmailing [29]. Especially in the case of the type of system described here, which is made part of the infrastructure of the care home and used for health care, this all-or-nothing approach is undesirable. A better approach would be to offer users insight into the implications of using particular parts of the system or accessing particular kinds of information and letting them make (informed) decisions about whether or not to use each function based on that insight. This would involve explaining the implications of each function to users.

Discussion so far has been about informing users so that they can decide whether or not they want to take on the responsibilities that come with certain types of knowledge that can be acquired by using the system. This does not take into account the stakeholders about whom the system provides information, that is, the stakeholders whom the system monitors. These stakeholders could take issue with the fact that by being monitored by the system, they are making monitoring information available to family members and others, and in doing so they are in a sense imposing responsibilities on these family members and others. As with other users of the system, these 'data subjects' should be given the opportunity to understand what responsibilities they are creating by being monitored. Further, they should be given the ability to choose whether or not they want to make available the different kinds of information that can lead to responsibilities, and they should be able to do so in an informed way.

For designers to be able to inform users how use of their systems can lead to responsibilities, designers must first understand how use of the system leads to responsibilities. To do so, designers should conduct (or have access to) stakeholder analyses. They should find out and be aware of who the direct and indirect stakeholders of the envisioned system are. Furthermore, they should analyze which responsibilities these stakeholders have prior to using the envisioned system, which knowledge these stakeholders can be expected to have, and which knowledge use of the system results in for each of these stakeholder groups.

The AmIS should support the users in fulfilling their responsibilities. Ensuring that they do belongs to the responsibilities of the AmIS developer. For example, the developer could build in artificial agents into the AmIS, supporting the responsibilities of the human users by alerting them, preventing and/or reporting misuse of the system, alerting officials or care organizations, refraining from observing when not necessary for the purpose of the system, etc. We consider that the idea of artificial agents supporting ethical behavior (by supporting human users' in fulfilling their responsibilities) falls within the realm of machine ethics, which, in one sense, deals with questions of whether ethics can be "put into" a machine [30].

One way of doing so, according to Moor, would be to constrain the machine's actions to avoid unethical outcomes, making it an implicit ethical agent [30].

According to Moor, a stronger sense of moral agency could be achieved by creating explicit ethical agents that can represent ethical categories and perform analysis. One example of work in this direction is Deontic Epistemic Action Logic (DEAL), which models statements on permission and obligation, beliefs and knowledge, and actions [31,32]. A final type of ethical agent that Moor describes is the full ethical agent, although it is debatable whether a machine can be one. An explicit ethical agent approach such as DEAL could help realize systems that support users in fulfilling their responsibilities, by having the analyzing (and reasoning about) responsibilities and acting appropriately in light of these.

More generally, we conclude that the responsibility of the AmIS-developer to design for responsible use of his system, involves at least the following criteria: 1) Developers need to be aware of foreseeable responsibilities of the usage of their system (intended or unintended). 2) The system needs to prepare intended users for their moral responsibilities before the user is accepted as a user. 3) The intended user needs to accept these responsibilities when accepting to use the system. 4) With respect to professional usage and anticipated responsibilities for non-professionals the system and its agents should have a regular reminder of moral responsibilities.

6. Conclusion

In this chapter, we aimed to address the responsibilities associated with Ambient Intelligence. Central to our arguments was the fact that knowledge leads to responsibilities. Because Ambient Intelligence Systems (AmISs) provide their users with information and support the users in using that information to extend their knowledge, the use of AmISs leads to responsibilities for both intended and unintended users. By making developers aware of this, through the current chapter, a reflective argument leads to the conclusion that developers of AmISs have the responsibility to design for the responsibilities of the users of theirs systems. We have formulated an initial list of criteria for how the developer can take up this responsibility.

Many questions regarding responsibility in the use and development of AmISs deserve further investigation. For example: What is the relation (or what should the relation be) between our conclusions and legal obligations, for the developer, for the users, for other stakeholders? Can we attribute responsibility to artificial agents? Would a “license to use” agreement be an adequate measure? And on a higher level: does the access to a tool that has certain information (regarding someone’s health) bring or imply the moral obligation to use the tool?

We hope this chapter contributes to putting such questions on the agenda of the AmIS research community.

Acknowledgements

We would like to thank Jeroen van den Hoven for many thoughtful comments and discussions that helped shape earlier drafts of this chapter. We would also like to thank the anonymous reviewers for their helpful comments.

References

[1] M. J. Zimmerman (1986). Negligence and Moral Responsibility. *Noûs*. **20**(2), pp. 199-218.
[2] M. J. Zimmerman (1997). Moral Responsibility and Ignorance. *Ethics* **107**(3), pp. 410-426.
[3] R. Davis, (1986). Knowledge-Based Systems. In: *Science*, **231**(4741), pp. 957-962.
[4] N. Romero, J. Sturm, M. Bekker, L. de Valk, L., and S. Kruitwagen, (2010) Playful Persuasion to Support Older Adults' Social and Physical Activities. *Interacting with Computers* **22(6)**, pp. 485-495.
[5] M. Wooldridge and N.R. Jennings, (1995). Intelligent Agents: Theory and Practice. *The Knowledge Engineering Review* **10**, pp. 115-152.
[6] N.A. Vincent, (2011). A structured taxonomy of responsibility concepts. In: N.A. Vincent, I. van de Poel & M.J. van den Hoven (eds), *Moral responsibility: Beyond free will and determinism*. New York: Springer.
[7] H. L. A. Hart (1968). IX. Postscript: Responsibility and Retribution. *Punishment and Responsibility.* 210-237. Oxford: Clarendon Press.
[8] K.G. Coleman, (2011). Computing and moral responsibility. In Edward N. Zalta (ed.), The Stanford Encyclopedia of Philosophy. Summer 2011 edition. http://plato.stanford.edu/archives/sum2011/entries/computing-responsibility/
[9] H. Nissenbaum (1996). Accountability in a computerized society. *Science and Engineering Ethics*. **2**(1), pp. 25-42. Springer Netherlands.
[10] N. Doorn, (2009). Responsibility Ascriptions in Technology Development and Engineering: Three Perspectives. *Science and Engineering Ethics*, pp. 1-22. Springer Netherlands.
[11] T. De Lima, Royakkers, L., Dignum, F. (2010) A logic for reasoning about Responsibility. *Logic Journal of the IGPL*, **18(1)**, pp. 99-117
[12] N. Doorn and I. van de Poel. (2011). Editors' Overview: Moral Responsibility in Technology and Engineering. *Science and Engineering Ethics*, pp. 1-11. Springer Netherlands.
[13] R. Mastop, (2010). Characterising Responsibility in Organisational Structures: The Problem of Many Hands. In: G. Governatori and G. Sartor (eds.), *Deontic Logic in Computer Science*. Heidelberg: Springer.
[14] G. Sher, (2009). *Who knew?: Responsibility without awareness*. Oxford: Oxford University Press.
[15] M. S. Moore and H. M. Hurd, (2011). Punishing the Awkward, the Stupid, the Weak, and the Selfish: The Culpability of Negligence. *Criminal Law and Philosophy* **5(2)**, pp. 147-198. Springer Netherlands.
[16] J.L. Austin, (1962). How To Do Things With Words, 2nd Edition, ed. J.O. Urmson and M. Sbisá. Cambridge, MA: Harvard University Press.
[17] J. Searle, (1968) 'Austin on locutionary and illocutionary acts'. *The Philosophical Review* 77: 405–424.
[18] J. Searle, (1969). *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.
[19] D. J. Cook, J. C. Augusto, V. R. Jakkula (2009). Ambient intelligence: Technologies, applications, and opportunities. *Pervasive and Mobile Computing*, **5**(4), pp. 277-298
[20] S. H. Matthews and T. Tarler Rosner (1988). Shared Filial Responsibility: The Family as the Primary Caregiver. *Journal of Marriage and Family*, **50**(1), pp. 185-195
[21] M. Weiser, (1991). The computer for the 21st century, *Scientific American* **265**, pp. 94–104
[22] D. H. Stefanov, B, Zeungnam, W.-C. Bang (2004). The smart house for older persons and persons with physical disabilities: structure, technology arrangements, and perspectives. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, **12**(2), pp. 228-250
[23] J. Ko, C. Lu, M.B. Srivastava, J.A. Stankovic, A. Terzis, & M. Welsh, (2010). Wireless Sensor Networks for Healthcare, *Proceedings of the IEEE* **98**, 1940–1960.
[24] H. Alemdar and C. Ersoy, (2010). Wireless sensor networks for healthcare: A survey. *Computer Networks*, **54**, 2688–2710.
[25] H. Lee, K. Park, B. Lee, J. Choi, & R. Elmasri, (2008). Issues in data fusion for healthcare monitoring. In: *Proceedings of the 1st international conference on Pervasive Technologies Related to Assistive Environments*, pp. 1–8.
[26] World Health Organization, (2011). *International Classification of Primary Care, Second edition*. Available online, retrieved July 28, 2011 at http://www.who.int/classifications/icd/adaptations/icpc2/en/
[27] M. Bern-Klug, Forbes-Thompson, S., (2008). Family Members' Responsibilities to Nursing Home Residents: "She Is the Only Mother I Got". Journal of Gerontological Nursing, **34**(2):45-52. SLACK INC, NJ.
[28] S. Davies and Nolan, M. "Making it better": Self-perceived roles of family caregivers of older people living in care homes: A qualitative study. International Journal of Nursing Studies **43**(3), pp. 281-291. Pergamon-Elsevier Science LTD, Oxford.

[29] M. Langheinrich, (2001). Privacy by Design – Principles of Privacy-Aware Ubiquitous Systems. In: Gregory Abowd, Barry Brumitt, Steven Shafer, (eds.), *Ubicomp 2001: Ubiquitous Computing*, Springer, Heidelberg, pp. 273 – 291.

[30] J. H. Moor, (2006). The nature, importance, and difficulty of machine ethics. *IEEE Intelligent Systems*, **21 (4)**, pp. 18-21.

[31] J. van den Hoven and G. J. C. Lokhorst, (2002). Deontic Logic and Computer-Supported Computer Ethics. In: J.H. Moor and T.W. Bynum (eds.), *Cyberphilosophy: The Intersection of Computing and Philosophy*, pp. 280–289. Blackwell.

[32] V. Wiegel, M. J. Van den Hoven and G. J. C. Lokhorst, (2005). Privacy, Deontic Epistemic Action Logic and Software Agents: An Executable Approach to Modeling Moral Constraints in Complex Informational Relationships. *Ethics and Information Technology* **7**(4), pp. 251-264. Springer Netherlands.