CHAPTER 7

The Web of Smart Entities—Aspects of a Theory of the Next Generation of the Internet of Things

Michael Wollowski*, **John McDonald**†
*Rose-Hulman Institute of Technology, Terre Haute, IN, United States
†ClearObject, Fishers, IN, United States

7.1 INTRODUCTION

We argue that the next generation of the Internet of Things (IoT) is about a web of smart entities (WSE). We define smart entities as software applications that build real-time models that are informed by real-time data. Smart entities are authorized to act and will manage routine behavior. Software applications in WSE will interact with each other to regulate behavior so as to satisfy certain goals. This interaction will lead to as yet unforeseen levels of automation. We see smart entities as polite assistants, designed to make our lives more convenient; something that will gracefully bow out, when asked to do so. We will address several modes in which to interact with and control the resulting automation.

Gubbi, Buyya, Marusic, and Palaniswami (2013) present a vision of IoT in which they emphasize the importance of cloud computing; we agree with their assessment. On page 1646, the authors state that "This platform [i.e. cloud computing] acts as a receiver of data from ubiquitous sensors; as a computer to analyze and interpret the data; as well as providing the user with easy to understand web based visualization. The ubiquitous sensing and processing works in the background, *hidden* from the user." Again, we could not agree more and explain in detail what sort of processing may take place in the background.

Artificial Intelligence for the Internet of Everything
https://doi.org/10.1016/B978-0-12-817636-8.00007-7

Weiser, Gold, and Brown (1999) defines a smart environment as "the physical world that is richly and invisibly interwoven with sensors, actuators, displays, and computational elements, embedded seamlessly in the everyday objects of our lives, and connected through a continuous network." We will generalize this portrayal to emphasize real-time data that enables one to build real-time models. In this context, we will argue that there is real-time data that comes from sources other than sensors.

Stankovic (2014) sees a "... significant qualitative change in how we work and live." We will expose some of those changes and further refine his assessment. He continues by stating that "We will truly have systems-of-systems that synergistically interact to form totally new and unpredictable services." We agree with this assessment and shed light on the kinds of services we may expect.

This chapter continues to develop the themes of the book *The Internet of Things*, by Greengard (2015), *Precision*, by Chou (2016), and the paper "Network of 'Things'," by Voas (2016). From a perspective of analyzing the impact of IoT, this paper continues to refine the ideas presented in the book *How IoT Is Made* by McDonald, Pietrocarlo, and Goldman (2015).

Greengard (2015) is focused on a contemporary version of IoT. In particular he focuses on automation that results from real-time data. This automation is true even in his extended example entitled *2025: A Day in the Life*, pp. 180–186.

McDonald et al. (2015) argue that it is pertinent for companies to join the IoT space as it offers vast new opportunities for revenue streams and for optimizing operations. It furthermore exposes what the authors call the "democratization" of information. This book does not address the bigger picture that evolves when IoT devices act and interact. We go beyond this book with a nuanced discussion of how, where, and by whom data is generated, where it is stored, and who ought to own it.

Chou (2016), similar to McDonald et al. (2015), is focused on IoT for industry and makes a case for companies to join the IoT to develop new business models and revenue streams that take advantage of the data that is generated by smart devices. This book does not address the bigger picture that evolves when IoT devices act and interact.

Tucker (2014) and Siegel (2016) focus on big-data and predictive analysis. Predictive analysis can reveal things that may be shocking to individuals (see Duhigg, 2012). While predictive analysis will lead to automation, we focus on the automation that results when models that learn specifics about someone or something's behavior are empowered to act.

7.2 SMART THINGS

It has been argued that IoT has a PR problem (see Eberle, 2016). Eberle argues that rather than talking about IoT, we should be talking about smart things, such as smart cars or smart cities, which are powered by IoT. We agree with this assessment and so do others (Bassi et al., 2013; Willems, 2016). At the most basic, IoT is about connecting all sorts of things to the internet. Those things, whether washing machines, cars, our bodies, or our food, produce data, in particular real-time data (see Heikell, 2016). Often this data is useful on its own; however, we are interested in what we can do when those devices interact.

In addition to producing, processing, and reporting data from internal sensors, IoT devices may also receive input from entities external to them. Consider Google's "Nest" thermostat, which may receive weather information from a website in addition to data from internal sensors. As such people consider Nest to be a smart thermostat. Taking several devices inside the home and programming them so that they communicate with each other leads to a smart home.

While often data collected and processed by a smart device is useful on its own, and while connecting smart devices together is useful too, more value can be generated by building models of the data available to them. At the most basic, a model of a sensor may be used to interpolate missing data or determine whether data is out of an expected range and as such may be faulty. At a higher level, models of data can be used to produce considerable value. Cummins Engines, the largest independent manufactures of diesel engines, uses telematics, i.e., real-time engine data to build real-time models of how their engines actually perform. These models are then used by Cummins in several ways. By running live engine data against the model, they can ascertain the general health of a particular engine. By using predictive analysis, Cummins is able to predict various scenarios ruinous to an engine and as such is able to alert fleet operators, in real time, about fault-codes and their significance on the continued operation of the engine (see Cummins, 2016).

Moving a step further, one can authorize a model to act. While the model of a Cummins engine alerts an operator at Cummins, consider the Nest thermostat; it builds a model of the comfort preferences throughout a week and then enforces the preferences by turning on and off the air conditioner and heater.

We consider Google's Nest to be the state-of-the-art with regard to current practice for IoT, in the sense that robust and repeatable solutions in this mold exist. This state-of-the-art is captured in Fig. 7.1.

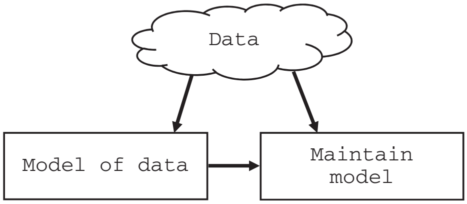

Fig. 7.1 Current state-of-the-art in data processing for the Internet of Things.

7.3 A VISION OF THE NEXT GENERATION OF THE IOT

As mentioned in the prior section, the Google Nest thermostat represents the current state-of-the-art in advanced use of IoT technology: it uses data from several internal sensors and from the web and it plays well with other IoT devices, such as mobile phones and IoT devices found in the home. The Nest thermostat develops a model that is authorized to act: it learns the resident's temperature preferences and maintains the temperature according to those learned specifications. In many ways the Nest thermostat incorporates key properties we wish to formalize. We feel that it represents a glimpse into what the future might bring.

In this section, we paint a broader picture of a likely future in which smart entities in the form of software applications interact with each other. We show that those smart entities rely on data from sensors but also from data compiled and processed by each other. As such, some of the data is fairly far removed from sensors. We show that some of the data is produced and processed continuously and some is produced in an irregular fashion. In the next generation of IoT, we see many different systems interacting to produce data and information. They will be used to seamlessly manage many aspects of businesses and of people's lives.

Perhaps the best way to characterize the next generation is by describing a rich extended example. We pick the domain of personal health. We portray a future in which a person's health is maintained at an optimal level, expressing the sort of systems that we wish to formalize. While the next generation of IoT will impact all aspects of people's lives, this domain is sufficiently complex to expose pertinent aspects of WSE. We should point out that the future of IoT cannot be seen in isolation; it is imperative that advances in IoT be seen in the larger context of advances in technology, such as predictive analysis (see Siegel, 2016; Tucker, 2014) and automation, such as smart factories (see Wikipedia, 2018), an example of which is the Daimler's Factory 56 (see Daimler, 2018).

Exercise. IoT has made great strides in measuring physical exercise activities. Many wearables can synchronize exercise data to various websites. It is fair to state that a small set of wearables enables a typical user to record an accurate picture of their exercise activities. In this context, we would like to point out that in most people's lives, there are clearly identifiable periods when meaningful exercise takes place. As such, for large portions of the day, these sensors do not produce meaningful data.

Diet. In most people's lives there are identifiable events when food and drinks are consumed. Just as with exercise data, we are interested in developing a picture of when, how much, and what kind of nourishment a person consumes. Unlike exercise data, when it comes to entering diet information, much of the data entry is manual at this time. Similar to exercise data, diet information comes in bursts. Even if we were to read off data continuously, the data is meaningful only during certain times of the day, i.e., when people actually consume food.

Websites such as "myfitnesspal.com" take advantage of the fact that many people are creatures of habit. They simplify the data entry process by giving the user the ability to select from prior entries rather than having to re-enter detailed information about a food dish. Another way to automate the process of maintaining diet information is by tying a meal planner to a site that maintains information about a person's diet. Websites such as "yummly.com" offer diet information associated with a recipe. We imagine that restaurants, by way of an itemized bill augmented by nutrition information, will soon enable the automatic entering of diet information by uploading it to diet management software. For this to occur, think of augmenting "expensify.com" with diet information and a plug-in for your "myfitnesspal.com" account.

Fitness. Given diet and exercise data, one can now track whether a targeted balance of exercise and diet has been reached (Fig. 7.2). Websites such as "myfitnesspal.com" keep track of past exercise and diet activities and use various graphics to indicate the degree to which exercise and diet are balanced. While one can create a basic model of a person's physical fitness, these models are passive; they merely report fitness data.

We believe that in the future, we will see applications that, in addition to compiling an accurate real-time model of a person's fitness, are authorized to act to maintain it. For example, in an increasingly wired world, a fitness application could refuse to pre-approve a meal in a restaurant that is judged as not fulfilling set dietary goals. Alternatively, the fitness application may suggest a walk or bike ride instead of the use of either a car or public transportation.

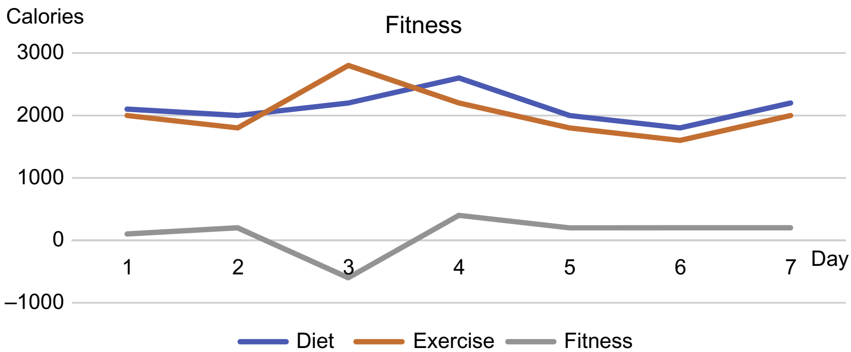

Fig. 7.2 Measuring fitness through diet and exercise data.

For habitual offenders we imagine that such an app may schedule an appointment with a physician. Some insurance companies already tie their rates to their clients' fitness data; as such, insurance rates are, in some cases, already tied to fitness. In general, we imagine that many people wishing to lead healthy lives will appreciate an application that helps them maintain their fitness.

7.3.1 Interlude

So far, we have seen that meaningful data may be generated and uploaded continuously. However, we have also seen cases in which data is generated and uploaded sporadically. We consider both to be real-time data. Additionally, we have seen applications that model behavior and are authorized to act, while enforcing certain constraints. We will now continue to weave a larger web of interconnected applications that manage additional aspects of our lives. In this context we will move further away from sensor data. We will argue that data generated by applications are to be considered part of the next generation of IoT.

Mental health. Mental health is equally important to physical health. IoT and derived applications will enable us to monitor and gauge mental health as well. We know we will soon have mirrors that are equipped with cameras that can interpret a person's mood. Certainly, the same software can be installed on cameras of various computing devices that people use on a daily basis. We envision that someone will soon develop a working laugh-o-meter app for smartphones, providing useful information about a person's mental health. These are but two examples; we mention them to express our vision that some of the IoT data will require sophisticated processing to derive desirable information.

For those people who maintain precise calendars of most of their daily activities, one could determine the kinds and duration of their mental activities. By consulting a person's calendar, one could determine whether someone reads books, completes puzzles, engages in social activities, or has other creative pursuits. This kind of information, while not currently derived from sensors, provides useful information that we feel belongs in the space of smart entities and applications.

Physical health. We have already addressed fitness. While obtaining reliable and complete exercise data for healthy people seems fine, there are other aspects that also ought to be measured directly rather than inferred, especially for people with chronic illnesses. There are already medical devices that people use, such as pulse monitors, blood-pressure monitors, and wireless scales. If we included implanted devices, such as defibrillators, pace makers, and blood glucose monitors, a good picture of physical health emerges even for people with major illnesses.

An exciting future development will be the use of nano-bots (Akyildiz, Jornet, & Pierobon, 2017), which, when placed in the body, can provide more fine-grained monitoring of a person's health or can be used to treat diseases such as cancer (Gaudin, 2009).

Automatic scheduling of doctor visits. Combining a real-time accurate model of physical health with best practices in health care, we imagine that the model will be empowered to make appointments with various health-care professionals as necessary. There are several immediate benefits to such a system: it will likely reduce the number of frivolous office visits, it will likely provide health care for people who are unwilling to see their doctor, and it will provide for a fast response to an emerging illness. Some office visits will likely be eliminated entirely. For example, often when our children are ill we know that they need an antibiotic. Perhaps the systems and the regulations about prescribing medication will change so that some medication can be prescribed based on real-time data and best practices.

Another form of real-time data is input by a health-care provider. We imagine that visits with health-care providers will remain, except that the role of health-care providers will change. People are not often good diagnosticians of their own mental or physical states. We believe that it takes an independent expert to recognize and enter some health information. Notice that while the data provided by a health-care provider is not as frequent as that of, say, a wearable device, it nevertheless is real-time data. Another kind of data may come in the form of revised nutrition or exercise guidelines, such as those issued by the US Department of Health and Human Services.

An interesting side effect of this scenario is the effect it would have on how doctors and health-care professionals spend their time. According to the *New York Times*, doctors find it hard to spend more than 8 min per patient visit (Chen, 2013). With the ability to measure blood pressure and weight, run blood tests, and conduct other simple tests by connected devices, there will likely be a drop-off in patient visits. This reduction in office visits will allow doctors to spend more time with those patients who need it. More importantly it will change the role of a health-care professional. We believe that the role of health-care professionals will transform into that of a health coach or advocate.

With real-time data, emergency responses can be automated with great benefits; see Lange (2013) for an insightful use case. Consider a car crash; based on data from wearables as well as telematics of all of the involved parties, the severity of a crash can be assessed and the need for medical assistance evaluated. If emergency assistance is deemed necessary, controlling for privacy, pertinent information about the patient should be sent to the attending paramedics, and the person's physical health records should interact with the assigned hospital's scheduling system. Finally, if appropriate, the model could alert family members and coworkers. Notice that the data is sourced from wearable devices as well as from multiple devices external to us.

In today's healthcare world, patients and physicians are seen as partners. Many patients want to know more about their conditions or feel that they are in charge of their own health care. As such, we imagine that if a model determines a person has a certain illness, it may make information about that condition available to that person in a way that appeals to their background knowledge.

Mens sana in corpore sano. With an adequate model of a person's mental and physical health, one can now develop a more complete model of a person's overall health and automate the model to maintain overall health to specifications that will likely include competing parameters. This automation may be as simple as dynamically injecting physical or recreational mental exercises into a person's calendar, based on real-time data of a person's mental or physical state. Perhaps a system may decide to send an employee home at an earlier time or assign them different work so as to alleviate stress.

7.4 THE USE OF ARTIFICIAL INTELLIGENCE IN THE WEB OF SMART ENTITIES

Processing sensor data to elicit higher levels of information, such as might be seen in smart mirrors or laugh-o-meter applications, requires advanced artificial intelligence (AI) techniques. We imagine that when gathering data from different scenarios to form an overarching model there will be inconsistencies. Detecting and possibly resolving inconsistencies or conflicts can be accomplished with AI techniques such as proof checkers. The connected nature of WSE requires further, perhaps more mundane uses of AI techniques. In this section, we will highlight some of these, as they suggest additional benefits from WSE.

Constraint satisfaction. The most obvious use of constraint satisfaction is when more than one person occupies the same space. Consider temperature settings, light settings, or entertainment choices that need to be resolved. A more sophisticated example involves regulating sleep. With the creation of smart beds and wearables, it is possible to monitor people's sleeping patterns. A model of sleeping patterns informs whether one is getting enough sleep each night. The sleep model can interact with several systems in an attempt to regulate sleep. For example, it could be empowered to regulate the temperature in the bedroom. It could interact with the meal planner to detect foods or drinks that are not conducive to sleep. It could be empowered to remove or rescheduled these items to earlier in the day. The sleep model could interact with the calendar to reschedule certain kinds of physical exercises that are detrimental to sleep.

Recommender system. Given models of people's behavior, we are in a position to make recommendations. For example, the "yummly.com" website makes recommendations based on the preferences entered by a user. We imagine that in the future recommendations can be made based on matching a user's meal-time recipe usage to those of others. This matching would be similar to how Netflix and Amazon.com recommend movies and goods. Similarly, based on a user's exercise patterns, we imagine recommendations for modifications, additions, or substitutions of exercise regimes.

Epidemics. Automatic collection and consolidation of health data will enable public agencies to detect developing trends in real-time (Jalali, Olabode, & Bell, 2012). Since time is of the essence in formulating a response, the more real-time data that is available, the faster one can detect trends. On a more local scale, it will help health-care providers in a given community to determine what sort of illness is afflicting their patients, enabling them to act accordingly.

Cognitive assistants. Cognitive assistants, as proposed by IBM (Kelly, 2015), are aimed at digesting vetted data to provide additional information to health-care providers. IBM sees cognitive assistants as "wise counselors" (IBM Watson, 2012). As IBM sees it, "IBM Watson, through its use of information retrieval and natural language processing, draws from an impressive corpus of information, including MSK [Memorial Sloan-Kettering] curated literature and rationales, as well as over 290 medical journals, over 200 textbooks, and 12 million pages of text. Watson for Oncology also supplies for consideration supporting evidence in the form of administration information, as well as warnings and toxicities for each drug" (IBM Watson, 2016). In essence, cognitive assistants data-mine the results of research. In the context of this chapter we see cognitive assistants used to provide additional inputs to models.

7.5 TOWARDS A THEORY OF THE WEB OF SMART ENTITIES

In this section, we develop a theory of WSE. We use the examples described in the prior section to justify the components of the WSE theory. We show that this use of the web is about real-time data, real-time models that capture routine behavior, and models that are authorized to act. We show the effects of this automation. We will end this section by highlighting the changing roles of established stakeholders and practices.

7.5.1 Real-Time Data

Smart and not so smart devices already generate data. While data on IoT comes from "things," in the extended scenario we described earlier, we demonstrated that data originates not only from things, even if they are *everythings*, but also from software applications that are not directly connected to things and, as a matter of fact, can be quite removed from the data produced by devices. We additionally exposed the applications to the readers that collect real-time data in a noncontinuous fashion.

Definition 1. *Real-time data* originates from different kinds of sources and is reported with different kinds of frequencies.

Let us consider some of the different kinds of data sources and frequencies under consideration.

Sensor data. Without a doubt, a key aspect of IoT and, by extension WSE, is real-time data obtained from sensors. Typically this data is reported continuously.

Manually entered data. If we look at how a person's diet data is entered into a system, it is currently not generated by sensors. If a meal planner is used, controlling for portion size, then some of the data is known and can be entered automatically. No matter how the data is entered, whether manually or automatically, it still is real-time data. It is just that most people do not eat continuously. While continued automation and perhaps video analysis will eventually enable the automatic generation of diet data, we believe that there will always be cases in which data will need to be entered manually. We would like to point out that, in the case of video recognition, the data, while technically coming from a sensor, requires sophisticated image processing.

Aggregated data. If we look at how "Google maps" ascertains traffic data, it is simply the aggregate of data from cell phones in cars. There is certainly a good amount of processing necessary to produce useful data about the movement of phones in vehicles. Notice that "Google maps" uses this data to eventually produce a model of congestion. However, before doing so, "Google maps" does produce aggregate data.

Other models. We have seen several examples in which data from models feed into other models and, as such, generate useful data for these other models. For example, a model that is designed to balance fitness will need access to the data from a model capturing diet data as well as a model capturing exercise data. We imagine that a model that balances fitness would furthermore interact with other models, such as calendars, vehicles, public transportation and restaurants.

Aggregate models. Just as Google aggregates data from individual phones in cars to construct a model of traffic flow, we can imagine cases in which we wish to aggregate models. Consider models of exercise data. If we were interested in simply ascertaining the overall exercise activities of a firm's employees, we would only need to gather a single data point from each employee. However, if we wish to ascertain exercise patterns, perhaps in the context of scheduling gym hours or to determine how big of a gym to build, then models of exercise patterns are necessary.

Feedback loop. A feedback loop of a model to itself enables monitoring and reflection on the workings of the model. Suppose a model of a person's food preferences is matched to someone else's model. A recipe may be returned that is deemed to match a person's preferences. In case the person does not like the recipe, or perhaps the matching parameters are insufficient or were weighted improperly, we would like to adjust the model. We then think of how case-based reasoning matches new cases to an existing case-base (see Wikipedia, 2016).

7.5.2 Real-Time Models

A good number of smart devices already maintain real-time models. Consider a Nest thermostat; it builds a model of a user's heating and cooling preferences. In particular it builds a real-time model as it constantly learns from real-time data. Similarly a Cummins Engine is processing sensor data from an engine to produce a model that reflects the performance and health of an engine, another prime example of a real-time model.

Definition 2. *Real-time models* represent aspects of the world that are continuously updated by real-time data.

We use the term "model" as shorthand for applications that maintain an underlying model of the data available to them. Fig. 7.3 captures the discussion so far to show potential inputs to a model.

7.5.3 Automation

If we look at the Nest thermostat, in addition to building a model it acts on data by turning on and off the air-conditioner or the heater. Cummins Engines analytics at this point in time notifies an operator who will then act on the information provided to them. A key effect of automation is that smart entities will learn routine behavior and automate it. In many instances, such routine behavior is not very exciting, but is rather considered a "nuisance" activity.

Definition 3. *Automation* results from real-time models that are authorized to act.

Automation takes on several forms and we list some of them in the following section.

Managing learned behavior. Suppose a model learned that every Tuesday evening is pizza night. Suppose it also learned that a given family always orders the same pizza. In that case the model can order the same pizza to arrive at the usual time. To look at a more complex case, suppose that the model also learned that the given family never orders pizza twice in a row and that this family had pizza the night before. In that case the model could ask for input, or perhaps act on some other learned behavior. Notice that in this case the model acts on learned behavior as well as real-time data.

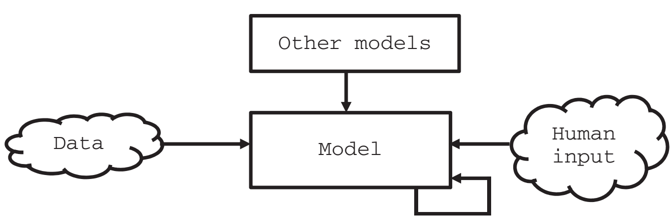

Fig. 7.3 A model and its potential inputs.

Smart substitutions. The use of AI technologies and the use of ontologies such as used in the context of the semantic web enable smart substitutions. We see examples of this substitution when, based on dietary restrictions, alternate meals may be suggested, or when certain kinds of exercises are recommended based on availability or opportunity.

7.5.4 Web of Smart Entities

Consider Google's Nest thermostat; in addition to processing data from its internal sensors, it can process data about the weather communicated to it by a weather app. We see Google's Nest as highlighting the beginnings of a richly interwoven fabric of applications that are directly or indirectly informed by sensor data.

Definition 4. WSE consists of a highly connected web of software applications that manage and automate routine behavior.

A few representative tasks for these smart applications are listed in the following section.

Balancing. If an application that manages a person's exercise activities interacts with an application that manages a person's dietary intake, physical fitness can be balanced to specifications. If we empower the fitness model to make the relevant decisions, we can dynamically adjust a person's fitness. For example, the fitness model may encourage a walk or bike ride rather than the use of a car or public transportation. Perhaps together they recommend a dish that lowers a person's caloric intake at a restaurant within walking distance.

Seamlessness. Given the proliferation of data, it is likely that models will gather data about particular activities in different contexts. For example, food preferences will likely be gathered not just from meals prepared at home, but also from meals ordered at restaurants or consumed in other settings. This way an overarching and more informed model can be built. Seamlessness comes about when an overarching model is applied in different contexts. If the model learned that someone likes their coffee black, then this is how it should be prepared, whether at home, at work, or by a coffee shop.

Recommendations. Models of a person's behavior can be used to make recommendations based on matching to like models. For example, diet preferences, just as preferences that Netflix and Amazon gather about their customers, can be used to match to similar models and, based on those matches, recommendations may be made.

7.5.5 Changing Roles of Stakeholders

We expect that the large-scale automation described in this chapter will have a significant impact on the participants of WSE.

Prediction 1. The *web of smart entities* will have a transformative effect on its stakeholders.

Consider an application that manages a person's health. It ensures that we live our lives within scientifically based parameters. One may wish to call such an application the "guardian angel" app. Knowing that such an application provides a kind of safety net, it is not unreasonable to assume that many people will live their lives to the fullest; i.e., they will "die with their boots on." At the very least, automating the management of health will enable people to live longer, more productive and, hopefully, happier lives. In this context, such health management applications would be able to make the necessary health-care appointments for those people who are reluctant to visit doctors, and as such may bring about a situation in which illnesses are diagnosed early, before they become terminal. Equally beneficial, such applications may be able to identify mentally disturbed people and offer or make them seek help long before they become a danger to themselves or society.

Health-care providers, such as general practitioners, will likely see their roles transform from a service provider that patients seek to individuals who will manage and fine-tune a patient's health. Similarly, people will likely have personal trainers who fine-tune their exercise regimens and personal dietitians who fine-tune their diets beyond what big-data might do for them. On the subject of diets, we imagine that cook-book authors may transform from writers who cook to consultants for people who like to cook. In order to better manage mental health, we see life coaches as becoming a staple in people's lives, someone who will not just give advice on living life to the fullest, but who may fine-tune personal calendars to eliminate stresses and replace them by leisure activities.

We can see insurance companies as transforming into businesses that ultimately manage and determine what people can and cannot do for some cost. Perhaps it is not a black-and-white decision, rather a spectrum of choices that people may make. Perhaps it depends on agreed-upon standards of care or even agreed-upon risk a person wishes to assume.

In this context, we hope that we have outlined scenarios that either change people's jobs for the better or generate additional forms of employment.

7.6 INTERACTING WITH AUTOMATION

We described a highly automated world that is built on and derived from real-time data and a world in which models of routine behavior are authorized to act for the benefits of their users. It might be daunting to know that various computing systems record our every activity and build various models about us, constructing a kind of a virtual alter ego. It is not unreasonable to assume that various computing systems know aspects of a person's live better than the person knows him of herself. To some, this may be exciting, but to others, this scenario may be frightening. How will this affect the way people conduct their lives? Will it be liberating, as our own personal systems watch over us? Will people live more vicarious lives as they know the system will intervene when necessary? Will people feel watched? Will they feel "verklemmt"? Will people hide things from the model or purposefully engage in activities to deceive it, as described in Orwell (1950)? Will people get used to "big brother" watching them? Will the automation limit what we can do, a point made by Agamben (2010), or will it liberate us to live life to the fullest?

We attempted to give a reasonable view of the future, which we see as largely positive. We see the WSE as inhabited by polite assistants, designed to make our lives more convenient. We envision automated assistants that gracefully bow out, when asked to do so. As such, we envision, perhaps too hopefully, a future in which people can choose and change, at a moment's notice, the level of interaction with the WSE. In particular we would argue that the ability to choose the degree of automation should be a design feature, something that the user can explicitly manage and, to a certain degree, something that the model anticipates. In the same context, users should be able to control what information is gathered about them and who has access to it.

We now describe three points across a spectrum of interactions with automation: autonomous, semiautonomous, and manual interaction. Among others, a model authorized to act will seamlessly switch between modes, or, better yet, move across the spectrum of automation. A smart system will learn when to bow out, when to step in and at what level to take over.

7.6.1 Fully Autonomous

In this mode of interacting with automation the system makes all of the decisions. For example, as already mentioned, some people eat the same dish on specific days of the week. This stability is behavior that can be quickly learned.

The meal planner can be authorized to order dishes or the ingredients for them and arrange for delivery at desired times (another learned behavior). Similarly, some people always order the same dish at a particular restaurant. This behavior, too, can be quickly learned and applied appropriately. There are many other components of our lives that have little to no variation. Many people order the same toiletries, clothes, cars, take the same route to drive to work, have the same weekly work schedule, and engage in the same sort of recreational activities on a weekly basis. It is not unreasonable to assume that large swatches of our lives can be automated. The benefit of this mode is that it would take care of routine activities.

On a side note, we recall a time when people first attempted to "live off" the world-wide web for a given period of time. In the same vein, it might be asked whether people would be able to live in a fully autonomous mode. Many people are creatures of habit. We believe that people can live in fully autonomous mode. Whether such a life is interesting is another question.

7.6.2 Semiautonomous

In this mode the user gives some input to the model. In some cases information will be requested, in other's the user will simply override certain inputs or parameters. The override may be as innocuous as not following the directions of a navigation system. For a more concrete example, suppose a cook heard about substituting riced cauliflower for rice in stir-fry dishes. The cook may simply ask the recipe manager to use the new ingredient. If there is a recipe in some user-permitted or accessible data base that already accounts for the new ingredient, then it can be consulted. The automated pantry would be authorized to purchase the new ingredient, if necessary. If the system is sufficiently knowledgeable, it may inform the cook that they may first have to obtain an appropriate device to turn cauliflower into riced cauliflower.

When operating in this mode, we imagine that the input range will be limited to acceptable operating parameters. Examples of this are Airbus airplanes; they are designed not to be placed in a stall situation, no matter what input a pilot gives.

7.6.3 Manual

In this mode, the user acts without the assistance of automation, but the system will likely continue to record information. In this mode, the system will enforce certain boundary conditions. For example, for a logger, a square donut burger with bacon may be fine. For someone who spends most of their time in an office, a burger may still be fine if consumed within reason. For people with high cholesterol, a burger may not be an option at all and they may not be authorized to purchase it.

This brings up the issue of abilities. This system would disable some of the choices available to users and as such there will be certain things users cannot do, a concern raised by Agamben (2010). While such a system would take choice away from us, on the flipside, it may encourage us to live life to the fullest. Just as technologies like engine rev-limiters take choices away from us, there certainly are people who take advantage of technology to push their cars to the limit without reproach.

7.6.4 Extent of Automation

Shall there be limits to the hyper-automation we have described? Consider the following example. Suppose someone is in a car accident. Certainly emergency response should be scheduled immediately. With real-time data and models, a system may select a hospital based on distance, the availability of medical personnel with the necessary skills to treat the given injuries once known, especially in the context of a given health history. Obviously pertinent health data will be made available to approved providers to ensure proper and expedited care. In addition, the health insurance company, loved ones, colleagues, and superiors will be informed.

However, the automation does not have to stop there. After a car accident, in addition to the health insurance company and the car insurance company, advanced telematics will likely have been informed of the crash too. It could then arrange for a rental car to be delivered to the customer at a time when the injured person is expected to be released from the hospital, or for an autonomous car if the client is impaired. In the same context, the car insurance company can and will likely arrange for the damaged car to be repaired. If the car is considered a total loss, something that, based on telematics, additional sensors, and big data, can likely be determined automatically, should the car insurance company purchase a new car? To many, purchasing a car is not a pleasant experience. This experience is not made more pleasant when conducted from a hospital bed. So anticipated, the automation described in this example may be much appreciated.

Suppose the injury requires a longer-lasting recuperation period. We can imagine that short-term disability insurance will be activated automatically. However, what sort of response should an employer automate? An employer could automatically reassign others to cover the duties of the injured colleague or they could automatically hire a temporary employee. If the disability is judged to be longer lasting or permanent, would the employee be automatically terminated? Would some system automatically find the ex-employee a new job, based on skills and disability? What if the new job pays less? Would some system automatically sell the house and purchase a cheaper one? All of this automation can be seen as useful. However, at what point are we just along for the ride?

7.7 DEPTH OF WSE

We argued that the WSE will consist of many applications generating and processing data; applications that will interact with each other to produce an unseen level of automation.

Some people have expressed concern about designing applications for trillions of devices (Sangiovanni-Vincentelli, 2015). We submit that based on our analysis this problem may be quite manageable. In particular it is unlikely that any application will directly interact with three trillion devices. Based on our theory, the WSE will be compartmentalized so that many applications will process fairly local data. If we look at the dependencies of the models from our extended example about a person's health, we see a fairly low depth, where depth is measured by the number of applications that depend on crucial data from those applications that report to them.

Consider Fig. 7.4, in which we portray this scenario. It should be noted that we only included a small subset of the applications that were mentioned in the health scenario. The figure suggests that the complexity of the WSE, as judged by the depth of it, might grow approximately in a logarithmic fashion in relationship to the number of linked IoT devices. To be clear, while we believe that there will be an exponential growth in the number of applications, we think that the WSE will be wide rather than deep, with depth as defined above and where width is measured by applications that loosely depend on data from other applications.

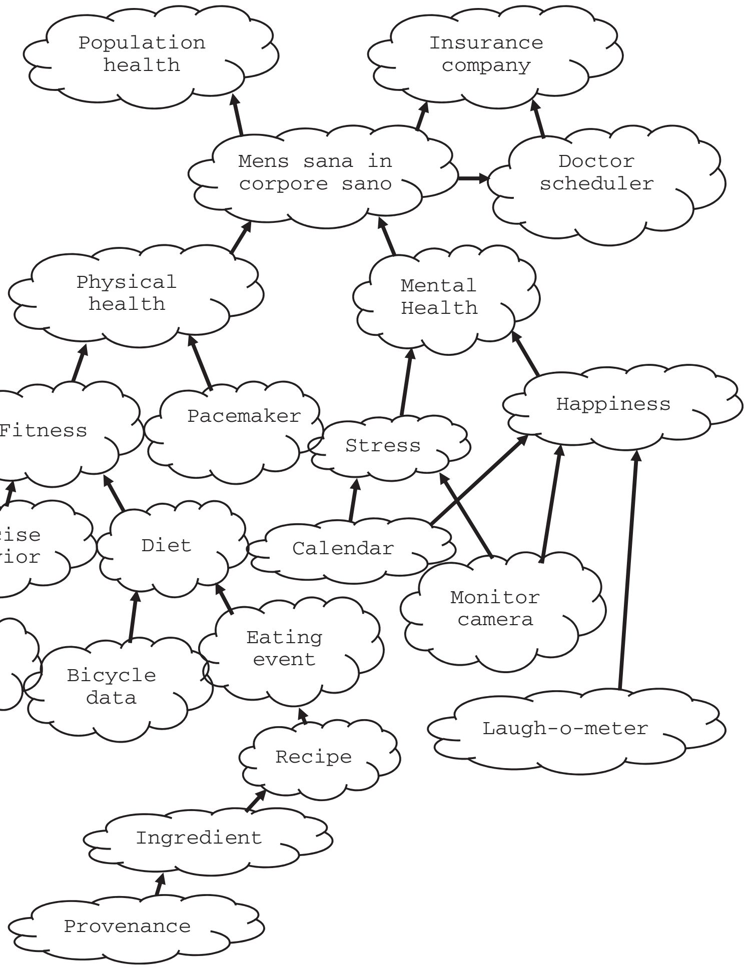

Fig. 7.4 Notional depth of dependencies of WSE in health.

7.8 CONCLUSIONS

In this chapter we described a likely future scenario in which IoT maintains people's health. It is a fascinating world in which software applications manage health based on real-time data and to scientific specifications.

We defined the next generation of IoT as a WSE. We argued that this web is about real-time data that originates from many sources at varying frequency, but where only some of the sources are sensors. We argued that a defining characteristic of the WSE is the development of accurate real-time models that capture and model the data. We argued that when models are empowered to act, an unprecedented level of automation will result. We depicted a world in which this automation will manage and arrange many routine activities.

We discussed the effects of this automation on several stakeholders. We believe that the hyper-automation described in this chapter will enable people to live life to the fullest. We portrayed three principle ways of interacting with models: fully autonomous, semiautonomous, and manual.

We believe that IoT is an exponential technology and that it is crucial that we consider and debate its likely future developments so that we can create an environment that brings to fruition a positive future. We believe that developers of this technology, stakeholders, customers, and regulatory agencies need to work together to define standards, best practices, and a legal framework for the vision to become a reality.

ACKNOWLEDGMENTS

This work was completed while the first author was on sabbatical at Clear Object. The authors would like to thank Ben Chodroff and Vishal Kapashi who provided input on an earlier version of this chapter.

REFERENCES

Agamben, G. (2010). On what we can not do. In G. Agamben (Ed.), *Nudities*. Stanford, CA: Stanford University Press.

Akyildiz, I. F., Jornet, J. M., & Pierobon, M. (2017). Nanonetworks: a new frontier in communications. *Communications of the ACM*, *54*, 84–89.

Bassi, A., Bauer, M., Fiedler, M., Kramp, T., van Kranenburg, R., Lange, S., et al. (2013). *Enabling things to talk—Designing IoT solutions with the IoT architectural reference model*. Cham, Switzerland: Springer Verlag.

Chen, P. (2013). For new doctors, 8 minutes per patient. Retrieved from http://well.blogs.nytimes.com/2013/05/30/for-new-doctors-8-minutes-per-patient/?_r=0.

Chou, T. (2016). *Precision: Principles, practices and solutions for the internet of things*. CrowdStory Publishing.

Cummins (2016). *Connected diagnostics—the lifeline for your engine*. Retrieved from https://cumminsengines.com/connected-diagnostics.

Daimler (2018). *Factory 56: the inventor of the car re-invents production*. Retrieved from https://blog.daimler.com/2018/02/20/factory-56/.

Duhigg, C. (2012). How companies learn your secrets. The New York Times, February 19. Retrieved from http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html?pagewanted=6\&_r=2\&hp.

Eberle, R. (2016). The internet of things has a vision problem. Retrieved from http://www.cio.com/article/3028054/internet-of-things/the-internet-of-things-has-a-vision-problem.html.

Gaudin, S. (2009). Nanotech could make humans immortal by 2040. Retrieved from http://www.computerworld.com/article/2528330/app-development/nanotech-could-make-humans-immortal-by-2040–futurist-says.html.
Greengard, S. (2015). *The internet of things*. Cambridge, MA: The MIT Press.
Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): a vision, architectural elements, and future directions. *Future Generation Computer Systems*, *29*, 1645–1660.
Heikell, L. (2016). Connected cows help farms keep up with the herd. Retrieved from https://news.microsoft.com/features/connected-cows-help-farms-keep-up-with-the-herd/\#sm.001npdttm13z6dn2spb2ce2sm2jay.
IBM Watson (2012). Assisting oncologists with evidence-based diagnosis and treatment. Retrieved from https://www.ibm.com/developerworks/community/blogs/efc1d8f5-72e5-4c4f-99df-e74fccea10ca/resource/Case\%20Studies/IBMWatsonCaseStudy-MemorialSloan-KettingCancerCenter.pdf?lang=en.
IBM Watson (2016). *IBM Watson platform helps fight cancer with evidence-based diagnosis and treatment suggestions*. Retrieved from http://www.ibm.com/watson/watson-oncology.html.
Jalali, A., Olabode, O. A., & Bell, C. M. (2012). Leveraging cloud computing to address public health disparities: an analysis of the SPHPS. *Online Journal of Public Health Informatics*, *4*(3).
Kelly III, J. (2015). Computing, cognition and the future of knowing. Retrieved from http://www.research.ibm.com/software/IBMResearch/multimedia/Computing_Cognition_WhitePaper.pdf.
Lange, S. (2013). *The internet of things architecture, IoT-A*. Retrieved from https://www.youtube.com/watch?v=nEVatZruJ7k.
McDonald, J., Pietrocarlo, J., & Goldman, J. (2015). *How IoT is made*. (n.p.): Author.
Orwell, G. (1950). *1984*. New York, NY: Signet Classics.
Sangiovanni-Vincentelli, A. (2015). *Design tools for the trillion-device future*. Retrieved from https://www.youtube.com/watch?v=ViJ3SH5t4Ys&feature=youtu.be.
Siegel, E. (2016). *Predictive analytics: The power to predict who will click, buy, lie, or die* (2nd ed.). Hoboken, NJ: Wiley.
Stankovic, J. (2014). Research directions for the Internet of Things. *IEEE Internet of Things Journal*, *1*(1), 3–9.
Tucker, P. (2014). *The naked future—What happens in a world that anticipates your every move*. New York, NY: Current Publishers.
Voas, J. (2016). Networks of 'Things'. NIST Special Publication 800-183. Retrieved from https://doi.org/10.6028/NIST.SP.800-183.
Weiser, M., Gold, R., & Brown, J. (1999). The origins of ubiquitous computing research at PARC in the late 1980s. *IBM Systems Journal*, *38*(4).
Wikipedia (2016). Case-based reasoning. Retrieved from https://en.wikipedia.org/wiki/Case-based_reasoning.
Wikipedia (2018). *Industry 4.0*. Retrieved from https://en.wikipedia.org/wiki/Industry_4.0.
Willems, C. (2016). *Cruising to safer, smarter street*. Retrieved from https://blogs.cisco.com/government/cruising-to-safer-smarter-streets.