



**HAL**  
open science

# Smart Home Futures: Algorithmic Challenges and Opportunities

Hussain Kazmi, Manar Amayri, Fahad Mehmood

► **To cite this version:**

Hussain Kazmi, Manar Amayri, Fahad Mehmood. Smart Home Futures: Algorithmic Challenges and Opportunities. 2017 14th International Symposium on Pervasive Systems, Algorithms and Networks & 2017 11th International Conference on Frontier of Computer Science and Technology & 2017 Third International Symposium of Creative Computing (ISPAN-FCST-ISCC), Jun 2017, Exeter, United Kingdom. 10.1109/ISPAN-FCST-ISCC.2017.60 . hal-01864974

**HAL Id: hal-01864974**

**<https://hal.science/hal-01864974>**

Submitted on 30 Aug 2018

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/321403107>

# Smart Home Futures: Algorithmic Challenges and Opportunities

Conference Paper · June 2017

DOI: 10.1109/SPAN-FCST-ISCC.2017.60

CITATIONS

3

READS

57

3 authors, including:



Hussain Kazmi

KU Leuven

15 PUBLICATIONS 30 CITATIONS

SEE PROFILE



Fahad Mehmood

Lahore University of Management Sciences

10 PUBLICATIONS 22 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Residential energy efficiency and demand response [View project](#)



3D shape recovery from image focus [View project](#)

# Smart Home Futures

## Algorithmic Opportunities and Challenges

Hussain Kazmi

ELECTA & Data Science Group  
KU Leuven & Enervalis,  
Belgium  
[Hussain.kazmi@enervalis.com](mailto:Hussain.kazmi@enervalis.com)

Manar Amayri

G-SCOP Lab  
Grenoble Institute of Technology,  
France  
[manar.amayri@grenoble-inp.fr](mailto:manar.amayri@grenoble-inp.fr)

Fahad Mehmood

Lahore University of Management  
Sciences (LUMS),  
Pakistan  
[Fahad.mehmood@lums.edu.pk](mailto:Fahad.mehmood@lums.edu.pk)

**Abstract**—Humans are increasingly spending their time indoors. This, along with higher wealth levels and rise of internet of things, has provided designers and planners the opportunity to reimagine living spaces. Smart homes come in many different shapes, but to gain widespread acceptance they have to increase the utility of building occupants in some meaningful way. The most straightforward way of creating these smart homes is assumed to be through artificial intelligence. In this paper, we take a critical look at some algorithmic approaches that have been formulated to do so and the opportunities they will create in the short term. We also present some key challenges that must be overcome before these opportunities can be realized in practice.

**Keywords**— smart homes, control, automation, reasoning, challenges, opportunities

### I. INTRODUCTION (HEADING 1)

Imagine waking up in the morning to sunlight streaming into the living space, the temperature is just right and the home has just drawn a warm bath. Later, LED's light up the path to the living room where a hot mug of coffee sits, freshly brewed, next to the breakfast. Having run out of groceries, the refrigerator has already ordered more for dinner. The autonomous electrical car in the garage is fully charged and ready to drive to work. As it pulls out of the garage, lights in the living room dim and fade out, the spatial heating and ventilation too are placed on hold.

This is one vision of a more synergistic future, enhancing the relationship between humans and their homes – entirely plausible in a few decades, if not sooner. There is no single monolithic vision for smart homes however, and they come in many different forms and futurisms. Ranging from the readily habitable to the wildly experimental, research has led us to diverging visions for the future of smart homes [1]. On the one hand is the fully connected, automated home which can anticipate every single one of its occupant's needs and desires. On the other lies reality which has to grapple with technical, legal and economic challenges to making such a complex endeavor work. The unifying theme across all these visions, despite numerous challenges, is the use of artificial intelligence for determining the practical appeal of smart homes.

Recent advances in artificial intelligence and more specifically machine learning have created a new hype cycle around possible applications to smart homes. Amongst the

most popular of these lies in creating a connected intelligence that can, over time, learn the building occupants' behavior and moods. By tighter integration with voice activated systems and assistive robots, this has the potential to revolutionize modern living spaces, even to the extent of exceeding the comfort levels pictured earlier in ways we can't imagine at present.

With an increasingly urban population spending most of its time indoors [1], such smart homes can considerably improve the utility of residents by providing an array of cosmetic and pragmatic benefits. Cosmetic applications of smart homes range from scheduling robot vacuums [3] to setting up mood lighting [4] etc. Integration with voice activated systems [5] has greatly aided in the popularization of such systems and will undoubtedly continue to keep doing so in the future [6]. Living in such a smart home can greatly enhance the experience and free up the residents to spend their time pursuing more productive tasks. However, in the near term, it is the more cost-effective pragmatic benefits of smart homes that is more likely to drive growth. These arise from smart lighting [7] as well as intelligent control of heating, ventilation and air conditioning (HVAC) systems [8] etc. Such 'smartening' can lead to both financial as well as comfort or health gains for the occupant. Obvious applications include higher energy efficiency and better indoor air quality etc.

These pathways to smart homes are enabled by internet of things [9], [10]. Smartphones, wearable devices and connected sensors allow service providers to gather and process data at unprecedented scales. However, even in the presence of ubiquitous sensors, the gathered data might be insufficient to perform the required tasks. Additional challenges remain as well, these include possibly storage, bandwidth and processing power requirements for transferring the generated data to a central server where it is processed. The alternative of distributed computation [11], while theoretically attractive, requires stringent hardware constraints placed at the end users premises. Large scale collection and storage of occupant data also creates data privacy and security concerns [12], either through negligence or malicious attacks by hackers. Furthermore, many of the sensory requirements placed on data collection for reasoning in smart homes are simply not possible at the moment.

Before we can arrive at fully connected homes which are able to anticipate occupant's whims, a better understanding of

the artificial intelligence that will control our environment is required. This paper presents a brief overview of some of the advances that will one day make this goal possible, the opportunities that will be created and some of the many obstacles we face today. The focus of this work will be algorithmic and it will only allude to some relevant regulatory and economic challenges that need to be overcome for the vision of smart homes to become a reality.

## II. ALGORITHMIC FRAMEWORKS

The concept of smart homes, bottom-up, can be divided into providing innovative services to the end user. A non-exhaustive list of practical use cases includes (1) automatically adjusting the temperature and lighting in the building as desired by a user, (2) maintaining air quality in the building at appropriate levels by e.g. forced or natural ventilation, and (3) ensuring adequate supply of hot water when necessary. Other more futuristic use cases are also possible, however these require even higher levels of abstraction for control.

Towards this end, a detailed understanding of the behavior of the building, its occupants and the environment is required. As a concrete example, the façade of the building, the presence (or absence) and nature of HVAC systems and the choice of lighting all affect the choice of intelligence in a building. The behavior of the occupant can only be modelled when it is sufficiently predictable (i.e. demonstrating high levels of cyclical patterns such as diurnal, annual etc.). An improved understanding of occupant behavior can considerably improve the quality of decisions made by the artificial intelligence system. Finally, the environment has a huge influence on living spaces and whether a building exists in one geographic location or the other has a critical effect.

In this section, we explain three increasingly automated strategies of reasoning in smarter living spaces before explaining some abstract objectives these reasoning frameworks can optimize towards.

### a. Rule-based reasoning systems

The most straightforward example of reasoning and control implemented in most real world systems is rule-based control [13]. While this might be archaic, it remains popular even now because it is easy to implement, understand and predict. Such systems simply implement rules based on the beliefs of the programmer. Reasoning in such systems follows logical patterns and it is straightforward to explain decision processes [14]. For example, a rule based controller for an HVAC system with a thermostat will try to keep the temperature being certain bounds. Lighting can be always on during certain hours of the day or as soon as an ambient sensor detects motion. Such control does not necessarily have to be trivial however. IFTTT controllers, a popular platform for smart homes, offer users the chance to create their own rules to better meet their needs [16]. While an improvement on the original formulation, this requires considerable user involvement and is ultimately constrained by the controls made available by the product creators.

More sophisticated examples of rule based systems also exist which learn rules from observational data. These are

hybrids between reinforcement learning agents and rule based systems, and we defer their discussion to the section on reinforcement learning systems.

### A. Model Predictive Control (MPC)

Historically, creating a model for the building (and its components such as the thermal system etc.) was done using offline methods. This meant significant effort had to be expended to gather the necessary data which would then be fit laboriously by a human expert to create a model that could approximate a building and its behavior. For instance, such a model could be created to explain the building temperature dynamics. Another example is reasoning about the ventilation system controlling air quality in confined living spaces. This would take the form of predicting future temperatures based on current observations. Once a model of the environment and the building is available, it is relatively straightforward to plan optimal actions for the future. This is usually done by maximizing (or minimizing) some reward function which is the ‘reasoning’ part [15]. The same logic can be applied to every aspect of a building’s operation from heating, ventilation and cooling to lighting.

This method creates a huge reliance on the human modeler and the quality of control becomes a function of the effort put into creating this initial dynamics model. This is obviously not a practical solution since millions upon millions of houses exist worldwide and it would require an impossible amount of work to model each of them. It also does not account for possible non-stationarities in the system, since while the thermal equipment or even the façade of a building might change, it is not necessary that the ‘reasoning’ artificial intelligence will be informed of these changes. In the absence of proper sensory data and dynamic learning models, this would create sub-optimal control behavior, to the extent of adversely affecting occupant utility. A related drawback of such reasoning techniques is that occupant influence is usually relegated to a secondary concern because it is not available offline. Some work in using occupancy forecasts has mitigated this to a large extent however [17].

### B. Reinforcement Learning

A longstanding alternative to model predictive control that has gained practical traction recently is in using reinforcement learning methods to reason in smart homes. The intent here is to create online, dynamic representations of the three key components forming a smart home: the environment, the occupant and the home itself [18], [19].

In reinforcement learning, two strands of research exist, model-free and model-based reinforcement learning. Model-free learning usually pertains to an agent that translates its state observations directly to control actions. In the context of smart homes, this means agents which observe the state of the smart home, the building and the occupant and then reason about possible next actions which can then be executed (a practical example of such an approach appears in [20]). Model-based reinforcement learning includes an additional step, whereby the controller first learns a representation of the environment, which it then uses to reason about next actions (a practical

example of such an approach appears in [21]). These differences are elaborated on in Fig. 1.

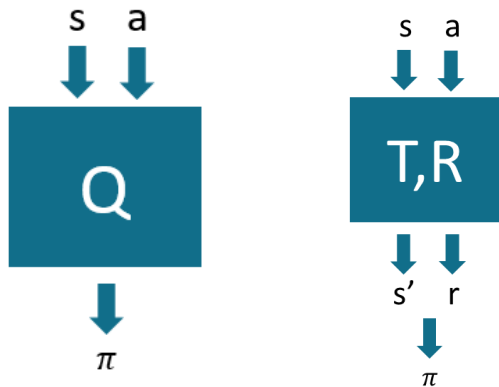


Figure 1: (a) Model-free reinforcement learning translates states ( $s$ ) and actions ( $a$ ) to a  $Q$ -value (the ‘goodness’ or desirability of a certain state-action pair), which is then mapped to optimal control actions in new (possibly unseen) states; (b) Model-based reinforcement learning learns a mapping from states and actions to a dynamic transition model ( $T$ ) and the immediate rewards ( $R$ ) which are then used to find the optimal control action given a certain state

Hybrids of the two approaches are possible and some potential applications include separating the three components into individual models, some of which are learnt online and some are derived from elsewhere (an example is the Dyna family of reinforcement learning algorithms [22]). These hybrids frequently appear in practice, especially in cases where inadequate sensing or considerable prior knowledge exists. The most obvious example of these is to use weather predictions for a certain geographic location from an online repository and only learn how these weather variations influence the system under consideration.

The action chosen by the reinforcement learner can have two competing objectives, often referred to as the exploration-exploitation dilemma [23], [24]. On the one hand, the agent can choose to accept its representation of the universe as ground truth and reason using this model. At the same time, if it has never observed the system in such a state before, it stands to reason that its beliefs about the system will be incorrect. This is true for the obvious case of seasonality (e.g. a learner trained using data only from the winter months will have poor generalization to summer months). More subtle cases exist as well, these include a sudden change in occupancy patterns (e.g. occupants leaving the building on vacation) and special events (such as unforeseen holidays which disrupt the normal occupancy patterns).

A representation of uncertainty in the reasoning agent’s knowledge can be critical in performing risk-aware control of buildings [25], [26]. Such uncertainty can also be used to integrate the information arising from sensors (which are themselves plagued with noise – quantization, sampling etc.) with the model predictions. When the model is sure about its prediction, it can be weighted higher and vice versa. The Kalman filter offers a theoretically sound way of doing this computation in real time [27]. Just as in autonomous driving of

cars, it is better to be safe than sorry in the case of controlling smart homes. Such representations also help in alleviating the exploration-exploitation dilemma.

By virtue of being data-driven, the approach solves many of the problems which plague model predictive control. Some challenges remain however. Key amongst these is the requirement of sufficient data collection for the learner to create an accurate representation of its operating conditions. Black swan type events might lead to catastrophic failure in such situations. Secondly, the concept of responsibility becomes diffuse in such autonomous systems when things go wrong with many different parties involved in different parts of the lifecycle. Finally, if the occupant demonstrates extremely erratic behavior, it is often impossible for a reasoning agent to anticipate the occupant’s next move.

### III. OPPORTUNITIES

In this section, we discuss some of the higher level objectives a reasoning smart home can work towards achieving. It is possible to do this with all three frameworks presented in the last section, albeit with varying success.

It is possible to see these opportunities as bi-level. On the one hand, smart homes raise the standard of living of their occupants. On the other, additional opportunities can be created when many smart homes are aggregated into a societal optimization problem forming smart cities.

#### A. Smart Homes

##### 1) Resource efficiency

There are two aspects to the utilization of resources optimally. The first is by automatically consuming in such a way that it is tailored to occupant behavior, the second is by informing the occupant of choices that are harming their own utility. The former falls under the ambit of the reasoning systems presented in the last section; the latter is more akin to a recommender system. In conversations on smart homes, it is often one or the other but the two form a natural feedback loop which can be leveraged to arrive at the optimal consumption of resources.

The primary example of such resource efficiency is in energy consumption. The energy consumption can be for any end draw, e.g. for thermal conditioning of the building or to providing adequate lighting under all possible conditions [15], [18], [21]. It can also be extended to other use cases such as minimizing the amount of water being consumed by informing the user of any leaks or the fact that a dishwasher might save them additional water. These are only representative examples and plenty more crop up in practice.

##### 2) Cost minimization

While making efficient use of resources invariably leads to a minimization of costs, there can be additional factors to be taken into consideration while explicitly minimizing operational costs [30]. Foremost amongst these is time-of-use tariffs which brings homeowners a step closer to the real market dynamics. Under these tariff schemes, electricity is more expensive during daytime hours reflecting the general situation of the electricity grid (based on supply and demand

principles). Minimizing for costs under such a pricing structure would therefore prioritize electricity consumption during night hours. While this can lead to peak shaving and valley filling globally, it can also lead to an increase in the overall energy consumption both directly and through indirect, unforeseen effects. An example of such an unforeseen effect is the efficiency of heat pump based systems which have usually higher efficiency when the ambient temperature is higher. Since the temperature is higher during the day, efficiency is likely to be decreased by prioritizing electricity consumption during the night (when ambient temperatures are lower).

Another example of a cost minimization strategy in a smart home can be to prioritize local consumption of energy generated by rooftop or building integrated solar PV panels [29]. This is becoming increasingly attractive from an economic perspective as governments rollback renewable subsidies and slash or debate feed-in tariffs and net metering rules.

### 3) *Comfort maximization*

Arguably the most important purpose of a smart home is to maximize (or improve in a quantifiable manner) the occupants' utility [31]. This takes on multiple forms from never being left in the dark to the temperature in the living room being just right without manually tweaking the thermostat. Beyond these smart lighting and thermostat functionalities are more complex and upcoming value propositions. Among these are smart charging of electric vehicles in a way that respects the users' wishes and smart reheating of water for showers etc. The latter is seen as a practical alternative to buffer-less heating systems which usually consume enormous amounts of electric power.

Furthermore, in the medium term, assistive technologies, possibly in the forms of robots, are expected to understand and anticipate the needs of the users, based both on historic patterns and stated commands. This will find increasing utility in caring for elderly people in the years to come with aging populations in many developed parts of the world [32]. For this to mature, natural language technologies and actuators need to improve drastically. While deep learning (a form of machine learning using neural networks) has meant that natural language processing has improved by leaps and bounds recently, the gains have not been uniform across all languages spoken by the people of the world. Smart homes should not be constrained by the occupants' proficiency in English (or a language which is not their own).

## B. *Smart Cities*

### 1) *Demand response*

So far, we have discussed smart home concepts in isolation. With increasing electrification and greater proliferation of solar PVs, individual homes will wield significant influence on especially the distribution electricity grid [33], [34], [35]. Peak shaving and valley filling will thus become important concepts in the pursuit of smart homes which act in an environmentally responsible way. The most representative example of this is perhaps the feed-in phenomenon where individual houses offload their excess solar production to the electric grid. This might be undesirable

behavior, especially if the grid has already a surplus of electricity. A context aware smart home is one which isn't just cognizant with the needs of its occupants but also of the environment it is operating in. Overloading local low voltage distribution networks could be locally optimal, but a community wide black-out is in no one's best interest.

### 2) *Ancillary services*

In addition to providing automatic demand response capabilities to the electricity grid, smart homes in an aggregated form can also provide ancillary services. By ancillary services, we mean reacting to supply and demand imbalances at (extremely) short time intervals [36]. While an individual home has negligible impact on the overall machinations of the electric grid, aggregated residential clusters can provide this functionality in an effective manner, in theory at least. The effort to realize these ancillary services using residential homes has been elusive historically, but there is no reason to assume that with advances in smart home reasoning, this will stay the case in the years to follow.

## IV. CHALLENGES

Despite technological gains and the allure of commercially available systems, substantial challenges persist. Primarily, there are questions of technical feasibility and whether artificial intelligence research has advanced to the point where such an endeavor is practical in the general sense. At the same time, economic questions linger relating to whether this can ever be more than a passing fad. The economic and technical questions are intricately tied together because eventually algorithms rely on input provided by sensors and output to actuators. If these are of insufficient quality because of economic choices, the reasoning component can't perform its job well.

In the following, we describe some of the many technical challenges faced while reasoning in smart homes.

### A. *Limited sensing capabilities*

A key factor limiting real time robust reasoning in smart homes is the limited sensing capabilities available in most houses. While internet of things has translated into an abundance of data, often data can't be reliably retrieved in a timely manner. This is exacerbated further when communication creates a bottleneck and control is implemented in a centralized manner.

A concrete example of limited sensing capabilities is controlling the HVAC of a building. While the thermostat might create a temporal mapping for the temperature for one particular location, it has no way of knowing the spatial temperature distribution in the entire room or building. It might be located right next to a window, heating radiator or cooling duct. This means that the temperature it senses won't be representative of the room, or home, in general. Using this sensor measurement to reason about the state of the home will therefore lead to erroneous control actions. The obvious solution of increasing the number of sensors also poses multiple problems: (1) it raises sensor costs, (2) it increases problem complexity considerably to learn spatiotemporal

temperature mappings rather than just temporal ones, and (3) it creates more points of failure, i.e. unless properly designed, a single temperature breaking down can cause the entire system to crash.

This matches well with the partially observable formulation of reinforcement learning problems. However, in many cases, this partial observability might inconvenience the human users. In cases where feedback to the control system is not directly possible (or even possible but not implemented in a very straightforward manner), this might well lead to disillusionment and the system being disabled eventually.

### *B. Disaggregation*

A special case of sensing limitations is disaggregation of consumption data. Disaggregation refers to deconstructing aggregated data into its individual sources. In the context of smart homes, disaggregation usually refers to decomposing aggregated electricity smart meter data. Such knowledge can provide insights into occupancy patterns as well as preferences. Occupancy profiles can be built from appliance usage and individual consumption profiles for separate appliances can be learnt over time.

In practical settings, overall household electricity consumption is disaggregated into individual draws such as for heating, ventilation, lighting and other appliances etc. This is an extremely ill-posed problem. Multiple devices can have similar consumption profiles and similar devices can have different consumption profiles based on usage patterns. Furthermore, it is impossible to enumerate all possible devices and their behavior in different settings. While supervised learning has been used both in research and commercially as a pattern recognition technique, this makes far too many assumptions and relies on extensive occupant feedback which can usually not be counted upon.

Another practical limitation of disaggregation techniques is that they require high frequency consumption data at the smart meters. This frequency can often be in the range of hundreds or thousands of hertz (many thousands times higher than what commercial smart meters capture data at). In addition to creating huge amounts of data which is problematic for storage and communication, it is also often physically impossible for commercial smart meters to log data at such high frequencies. Similar problems are also encountered for disaggregating water and natural gas consumption.

### *C. Difficulty in estimating occupant behavior*

As mentioned before, design and operation of buildings should not only consider building physics and HVAC systems, but also human behavior. This largely depends on energy monitoring and management systems (EMMS), which are a framework to collect sensor data from within a building. More concretely such a system can be used to (1) estimate the number of people residing in a living space and taking appropriate control actions (e.g. to turn on HVAC or not), (2) identify the person occupying the space and provide customized services and / or feedback, and (3) create future forecasts for occupancy profiles (of individual or aggregated occupants). Another key purpose of including occupant

behavior in considerations early on is to estimate the relevance of their activities in building simulation in order to reduce the so called performance gap with reality [40].

Living area systems are highly human-machine cooperative systems. Indeed, the purpose of developing a smart home is to create additional value for building occupants but occupants are also part of the system and influence the available possibilities with their own behavior. Thus, in reasoning about occupants and their behavior it is important to include contextual information. This is composed of: (1) context related to time, weather conditions, energy costs, heat gains per zone but also occupant current positions and activities, (2) controls related to doors, windows, flaps and shutters positions, configurations of the HVAC system and other electric appliances, and (3) reactions related to indoor temperature and air quality, and to satisfaction of occupants regarding services provided by electric appliances.

The problem is to identify and calculate features that could be used in a classification model for identifying various activities happening in a building space. The features must provide rich context for the learning system to classify various states of interest. Since, the use of video cameras and audio recorders is a problem for most residential spaces, the solution must keep privacy issues as well as cultural sensitivities and rely largely on non-intrusive sensors. These include electricity and hot water consumption, CO<sub>2</sub> measurements as well as motion sensors and door / window contact sensors [41]. These come with additional complications since they increase initial installation costs and bring only marginal improvement in occupancy detection. It is important to stress here that even with all these sensors, it is very difficult to estimate the exact number of occupants, let alone the identity of individual occupants.

### *D. Robustness to broken sensors or missing inputs*

Creating reasoning systems dependent on multiple relevant sensors can improve the learning performance of a task substantially. This follows intuitions from the human experience whereby we rely on multiple senses to gather environmental input which we then combine to form decisions for subsequent actions. Humans can however adapt to their context and learn to compensate for loss of one sensory source by focusing on a different one. As alluded to earlier, unless explicitly designed for, reasoning in smart homes is unable to follow the same principles. This creates fragile systems which rely on their individual components to work well to create the desired output. Likewise, in the case of missing inputs from a sensor, the decision making process can be impeded which can lead to loss of occupant comfort etc.

### *E. Robustness to actuators*

In addition to broken sensors, it is not guaranteed that the control commands will be mapped out as expected. Unexpected behavior can occur at any given time, especially in complex systems such as buildings. The simplest example of this is hardware implemented overrides, a common 'last defense' mechanism in many systems. These have the ability to overrule any commands generated by a reasoning system based on what they perceive as preserving occupant comfort. A



possible fix is human intervention to configure the device correctly, but this is not always practical.

#### *F. Robustness to goal oriented approaches and perverse decision making*

As discussed earlier, reasoning in smart homes is usually to maximize some preconceived notion of utility. If care is not taken in formulating this notion, it can have disastrous unintended consequences. An early example of such an awry system is a robot vacuum which is designed to collect reward by gathering garbage. A robot with such a reward framework might, unless explicitly designed otherwise, repeatedly dump and collect garbage from the same location instead of cleaning the entire space [38]. These kinds of unintended consequences can be spotted in the testing phase and can be corrected by tweaking the reward streams a reasoning agent receives.

A second more complex example is of a perverse spatial heating controller. Assume a heating system which reheats a room and gathers reward to keep the temperature between certain boundaries and is penalized for the amount of electricity it consumes from the grid. At the same time, it can consume electricity from solar PV panels; such consumption is usually incentivized because of changes (or upcoming changes) in regulation surrounding feed-in tariffs so the heating system is rewarded for heating the building when there is surplus solar electricity available. Now, also assume a cooling system which is being controlled by a different reasoning agent, except operating in the opposite direction. The heating agent will soon come to realize that it should always reheat the building whenever solar power is available, prompting the cooling system to also come on to keep the temperature within the specified bounds. Both systems collect rewards while doing something that is obviously not the design intent. This is just one example of learning agents learning perverse policies which don't reflect the original intent of the system at all.

While it is easy to think that such problems are only theoretical, these concerns will become increasingly real. As technologies mature, rooting out these inconsistencies will take on greater importance.

#### *G. Centralized vs. distributed architectures*

In addition to the question of devising robust reasoning strategies, there are questions pertaining to the choice of centralized and distributed architectures. Both have their advantages and disadvantages [39]. The distinction between centralized and distributed systems can span many levels of abstraction. At the most fundamental, a distributed system can involve different sensors directly taking control actions to maximize some reward signal. On the other hand, a centralized system would combine information from multiple sensors to create a more holistic picture, which can then be used to make decisions. Centralized systems can have far superior computational resources and can thus solve more complicated reasoning tasks. However, complexity can grow quickly in the presence of multiple sensors. Furthermore, data privacy issues might necessitate distributed architectures.

#### *H. Security and privacy*

As mentioned before, data security and privacy are key concerns in how effective automated reasoning can be in smart homes [37]. Data security means that the data has to be communicated and stored in a secure manner. Usually, this involves setting up of encrypted protocols. The primary threat in this domain comes from hackers who can gather insights into occupancy patterns from historic consumption profiles. Data privacy, on the other hand, refers to sharing of data unnecessarily. The need to know principle has to be applied in production settings to preserve user data privacy. A sensor that doesn't aid in the reasoning process is not only a useless expenditure, it is a needless liability.

The EU directive on General Data Protection (GDPR) lays down clearly defined rules for gathering and processing of data in an automated manner [28]. Automated decisions have to be explainable under the purview of this law and the user has a right to know what data was used to make a certain decision and how it affected them. This can complicate the implementation of many black-box reasoning systems but will also bring transparency to the decision making process.

## V. DISCUSSION

In this paper, we have highlighted three major frameworks used to reason in smart home settings. The focus of this paper has been on realistic applications such as smart lighting and thermostat functionality. For residential settings, rule-based and model predictive control might be too expensive to set up and generalize. This is because the initial high cost of creating a prior model that effectively captures building physics in a detailed manner requires significant manual effort. Reinforcement learning presents a low cost alternative that has been demonstrated to work in a cost effective manner. Hybrids between different formulations offer an attractive way to design the best solution for a problem on a case by case basis.

Some of the most obvious opportunities arising from such a control framework would be to achieve increased resource efficiency, either in the form of reduced electricity (or water / natural gas) consumption or reduced costs. While largely aligned today, these might well diverge in the future with time-of-use tariffs and rapidly changing renewable integration subsidy schemes. Another explicit objective to optimize towards is simply maximizing the occupant comfort regardless of cost. At the same time, more global objectives are being formulated as well, these will require many smart homes to be connected together in a microgrid or grid.

Questions remain around whether the presented framework will be sufficient to realize the functionality envisaged in smart homes. Sensors form the critical backbone of this enterprise, and without sufficient sensing reinforcement learning algorithms can't operate reliably. The question of limited sensing includes both temporal and spatial resolution, and draws realistic questions about something even as basic as estimating occupant presence, much less detecting occupant identity and offering customized services. This also opens the door to questions on data security and privacy and the possibility for abuse and surveillance. While distributed frameworks have been touted as possible solutions, it is not



entirely clear if these will have the necessary computational resources to perform complex reasoning tasks.

In the end, despite considerable challenges, there is much to be optimistic about. The progress made in just the past few decades has been astounding. As the sphere of technological influence grows, it is not unimaginable that sensing will become even more pervasive. Combined with gains in available computational resources and more efficient artificial intelligence algorithms, this has the potential to completely transform our living spaces. Whether the end result is the hedonistic vision of a smart home painted at the beginning of this paper is something that remains to be seen.

#### ACKNOWLEDGMENT

This work was done with support from VLAIO (formerly IWT) and InnoEnergy.

#### REFERENCES

- [1] Lauri Kainulainen, "Reasoning in a Smart Home", AIOME 2006
- [2] Ericsson AB, "Optimizing the indoor experience," 2013.
- [3] J. Forlizzi, C. DiSalvo, "Service robots in the domestic environment: a study of the roomba vacuum in the home", Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction, pp. 258-265, 2006.
- [4] Jeffrey J. Wolf, Scott D. Walter, Scott W. Demarest, Cory J. Nelson, Steven B. Mineau, "Color changing light devices with active ingredient and sound emission for mood enhancement", US7476002 B2, 2003.
- [5] Mital, Dinesh P., and Goh Wee Leng. "A voice-activated robot with artificial intelligence." *Robotics and Autonomous Systems* 4.4 (1989): 339-344.
- [6] Hirschberg, Julia, and Christopher D. Manning. "Advances in natural language processing." *Science* 349.6245 (2015): 261-266.
- [7] Garg, Vishal, and N. K. Bansal. "Smart occupancy sensors to reduce energy consumption." *Energy and Buildings* 32.1 (2000): 81-87.
- [8] Agarwal, Yuvraj, et al. "Occupancy-driven energy management for smart building automation." Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building. ACM, 2010.
- [9] Gubbi, Jayavardhana, et al. "Internet of Things (IoT): A vision, architectural elements, and future directions." *Future generation computer systems* 29.7 (2013): 1645-1660.
- [10] Piyare, Rajeev. "Internet of things: ubiquitous home control and monitoring system using android based smart phone." *International Journal of Internet of Things* 2.1 (2013): 5-11.
- [11] Camponogara, Eduardo, et al. "Distributed model predictive control." *IEEE Control Systems* 22.1 (2002): 44-52.
- [12] McDaniel, Patrick, and Stephen McLaughlin. "Security and privacy challenges in the smart grid." *IEEE Security & Privacy* 7.3 (2009).
- [13] Hayes-Roth, Frederick. "Rule-based systems." *Communications of the ACM* 28.9 (1985): 921-932.
- [14] Doukas, Haris, et al. "Intelligent building energy management system using rule sets." *Building and environment* 42.10 (2007): 3562-3569.
- [15] Oldewurtel, Frauke, et al. "Energy efficient building climate control using stochastic model predictive control and weather predictions." *American control conference (ACC)*, 2010. IEEE, 2010.
- [16] Dash, Suvendu Kumar, Amit Kumar Sinha, and Sudhir Kamble. "Design And Development of Mobile Application Based Wireless Controlled Energy efficient Automation System." *International Journal of Applied Engineering Research* 10.9: 2015.
- [17] Gyalistras, D., et al. "Use of weather and occupancy forecasts for optimal building climate control." *Technical Report* (2010).
- [18] Cook, Diane J., et al. "Learning to control a smart home environment." *Innovative applications of artificial intelligence*. 2003.
- [19] Rashidi, Parisa, and Diane J. Cook. "Keeping the resident in the loop: Adapting the smart home to the user." *IEEE Transactions on systems, man, and cybernetics-part A: systems and humans* 39.5 (2009): 949-959.
- [20] Ruelens, Frederik, et al. "Learning agent for a heat-pump thermostat with a set-back strategy using model-free reinforcement learning." *Energies* 8.8 (2015): 8300-8318.
- [21] Kazmi, H., et al. "Generalizable occupant-driven optimization model for domestic hot water production in NZEB." *Applied Energy* 175 (2016): 1-15.
- [22] Peng, Jing, and Ronald J. Williams. "Efficient learning and planning within the Dyna framework." *Adaptive Behavior* 1.4 (1993): 437-454.
- [23] Kaelbling, Leslie Pack, Michael L. Littman, and Andrew W. Moore. "Reinforcement learning: A survey." *Journal of artificial intelligence research* 4 (1996): 237-285.
- [24] Thrun, Sebastian B. "Efficient exploration in reinforcement learning." (1992).
- [25] Deisenroth, Marc, and Carl E. Rasmussen. "PILCO: A model-based and data-efficient approach to policy search." *Proceedings of the 28th International Conference on machine learning (ICML-11)*. 2011.
- [26] Hester, Todd, and Peter Stone. "TEXPLORE: real-time sample-efficient reinforcement learning for robots." *Machine learning* 90.3 (2013): 385-429.
- [27] Welch, Greg, and Gary Bishop. "An introduction to the Kalman filter." (1995).
- [28] European Parliament, "Regulation (EU) 2016/679 27", 2016.
- [29] Castillo-Cagigal, Manuel, et al. "PV self-consumption optimization with storage and active DSM for the residential sector." *Solar Energy* 85.9 (2011): 2338-2348.
- [30] Missaoui, Rim, et al. "Managing energy smart homes according to energy prices: Analysis of a building energy management system." *Energy and Buildings* 71 (2014): 155-167.
- [31] Kastner, Wolfgang, Mario J. Kofler, and Christian Reinisch. "Using AI to realize energy efficient yet comfortable smart homes." *Factory Communication Systems (WFCS)*, 2010 8th IEEE International Workshop on. IEEE, 2010.
- [32] Homes, Smart, Beyond C. Nugent, and J. C. Augusto. "Human-robot user studies in eldercare: Lessons learned." *Smart Homes and Beyond: ICOST 2006: 4th International Conference on Smart Homes and Health Telematics*. Vol. 19. IOS Press, 2006.
- [33] Palensky, Peter, and Dietmar Dietrich. "Demand side management: Demand response, intelligent energy systems, and smart loads." *IEEE transactions on industrial informatics* 7.3 (2011): 381-388.
- [34] Vivekananthan, Cynthujah, et al. "Demand response for residential appliances via customer reward scheme." *IEEE transactions on smart grid* 5.2 (2014): 809-820.
- [35] Dupont, Benjamin, et al. "Impact of residential demand response on power system operation: A Belgian case study." *Applied Energy* 122 (2014): 1-10.
- [36] Ma, Ookie, et al. "Demand response for ancillary services." *IEEE Transactions on Smart Grid* 4.4 (2013): 1988-1995.
- [37] McDaniel, Patrick, and Stephen McLaughlin. "Security and privacy challenges in the smart grid." *IEEE Security & Privacy* 7.3 (2009).
- [38] Russell, Stuart J., and Peter Norvig. "Artificial intelligence: a modern approach (International Edition)." (2002).
- [39] Bernstein, Daniel S., et al. "The complexity of decentralized control of Markov decision processes." *Mathematics of operations research* 27.4 (2002): 819-840.
- [40] Majcen, Daša, Laure Itard, and Henk Visscher. "Actual and theoretical gas consumption in Dutch dwellings: What causes the differences?." *Energy Policy* 61 (2013): 460-471.
- [41] Amayri, Manar, et al. "Estimating occupancy in heterogeneous sensor environment." *Energy and Buildings* 129 (2016): 46-58.