



# Artificial Intelligence and the IoT Connected Home

A CABA WHITE PAPER


L. Anne Breene  
ArcoLogix

Lawrence Silverman  
ArcoLogix



Connect to what's next™

[www.caba.org](http://www.caba.org)





## Artificial Intelligence and the IoT Connected Home

A CABA White Paper

### Authors

**L. Anne Breene**

ArcoLogix

**Lawrence Silverman**

ArcoLogix

### Working Group

**L. Anne Breene**

ArcoLogix

**Lawrence Silverman**

ArcoLogix

**John Feland**

Argus Insights, Inc.

**David Rogers**

BC Hydro

**Konkana Khaund**

Frost and Sullivan

**Ken Wacks**

Ken Wacks Associates

**Derek Cowburn**

Luman Cache, Inc.

**David Katz**

Sustainable Resources  
Management

Working Group: Individuals who either contributed ideas and input into the direction of paper or reviewed the final draft.

### Sub-Committee

**Name**

Company

Sub-Committee: Under the direction of the Sub-Committee Chair, this formal committee reviewed and approved both the initial white paper proposal and final draft.

## ABOUT CABA

The Continental Automated Buildings Association (CABA) is an international not-for-profit industry association, founded in 1988, and dedicated to the advancement of intelligent home and intelligent building technologies. The organization is supported by an international membership of over 365 organizations involved in the design, manufacture, installation and retailing of products relating to “Internet of Things, M2M, home automation and intelligent buildings”. Public organizations, including utilities and government are also members. CABA's mandate includes providing its members with networking and market research opportunities. CABA also encourages the development of industry standards and protocols, and leads cross-industry initiatives. CABA's collaborative research scope evolved and expanded into the CABA Research Program, which is directed by the CABA Board of Directors. The CABA Research Program's scope includes white papers and multi-client market research in both the Intelligent Buildings and Connected Home sectors. [www.caba.org](http://www.caba.org)

## ABOUT CABA'S CONNECTED HOME COUNCIL (CHC)

Established in 2004, the CABA Connected Home Council initiates and reviews projects that relate to connected home and multiple dwelling unit technologies and applications. Connected homes intelligently access wide area network services such as television and radio programming, data and voice communications, life safety and energy management/control information and distribute them throughout the home for convenient use by consumers. The Council also examines industry opportunities that can accelerate the adoption of new technologies, consumer electronics and broadband services within the burgeoning connected home market. [www.caba.org/chc](http://www.caba.org/chc)

## DISCLAIMER

This white paper was developed and published by CABA for the industry with permission from the authors. CABA expresses its appreciation to the authors and contributors for making this white paper available to be included as part of CABA's Members Library and CABA's Public Library. CABA, nor any other person acting on their behalf of CABA assumes any liability with respect to: the use of, or for damages resulting from the use of, any information, equipment, product, method or process disclosed in this white paper.

This CABA White Paper and other industry research reports can be found in CABA's Members Library and CABA's Public Library at: [www.caba.org](http://www.caba.org). This information is also keyword searchable. Contact the CABA office if you do not have the passwords to access this material by email [caba@caba.org](mailto:caba@caba.org) or phone 888.798.CABA [2222] or 613.686.1814 (x228). CABA encourages you to share this white paper with others in your organization and the industry. Permission is not required from CABA to share this white paper, as long as proper acknowledgment is provided to CABA.

## PUBLISHED

May 2018

## TABLE OF CONTENTS

Artificial Intelligence and the IoT Connected Home .....	1
Authors.....	1
Working Group.....	1
Sub-Committee .....	1
ABOUT CABA .....	2
ABOUT CABA’S CONNECTED HOME COUNCIL (CHC) .....	2
DISCLAIMER .....	2
PUBLISHED.....	2
<b>1. INTRODUCTION .....</b>	<b>4</b>
<b>2. TWO AI-IOT CONNECTED HOME WALKTHROUGHS.....</b>	<b>7</b>
2.1. Scenario: AI in the AI-IoT Connected Home is working as it should .....	7
2.2. Scenario: AI is wreaking havoc .....	9
<b>3. ARTIFICIAL INTELLIGENCE &amp; MACHINE LEARNING.....</b>	<b>11</b>
3.1. Overview .....	11
3.2. Machine Learning: Supervised and Unsupervised .....	12
3.3. XAI – Artificial Intelligence Explained, Ethics and the EU Regulations .....	13
3.4. Advantages and Challenges to AI in the Connected Home .....	15
<b>4. TECHNICAL DETAILS – AI AND THE IOT.....</b>	<b>16</b>
<b>5. BUSINESS MODELS.....</b>	<b>17</b>
<b>6. CONCLUSIONS .....</b>	<b>18</b>
<b>REFERENCES AND NOTES.....</b>	<b>20</b>

## 1. INTRODUCTION

In the months leading up to the final test of the first atomic bomb in Los Alamos near the end of WWII, a years-long debate continued among the scientists as to whether the heat of the nuclear explosion would ignite the atmosphere and destroy all life on earth.

One group argued that this was a distinct possibility,<sup>1</sup> the other denied it. Edward Teller and his colleagues made some calculations that indicated the chance was reportedly<sup>2</sup> less than three in one million. Hans Bette termed this infinitesimally small<sup>3</sup>. Bette's conclusion was accepted by Arthur Compton, director of the project, and the Trinity test went ahead - despite any lingering doubts. In fact, on the day of the test Enrico Fermi half-jokingly took bets that this would be the end of life as they knew it.

If three in one million seems like very small odds, consider that the odds of dying in a shark attack - 1 in 4 million,<sup>4</sup> of getting killed in a plane crash are 1 in 11 million,<sup>5</sup> and of winning the lottery - 1 in 292 million<sup>6</sup>.

Scientists and engineers sometimes do things to prove that they can, rather than because they should.

Today the scientific and engineering communities are abuzz with the idea of implementing Artificial Intelligence (AI) over the Internet of Things (IoT) - connecting not just things but connecting everything (which may or may not be wise). The world's most critical systems will be operated by the fastest computers known, communicating over fiber at 2/3 the speed of light - gathering information and making decisions more than 1 million times as fast as a human can think. By the time something seems to be wrong, it will likely be too late to do anything about it.

Can an imperfect human really design a perfect system that will never ever make a mistake?

---

Industry insiders have been predicting that "the Automated Home is 5 years away" for more than 40 years. Today, advances in miniaturizing computers, sensor technology and wireless communications seem to have finally made it possible that the dream will be fulfilled.

"In 1966, ECHO IV, the first home automation machine, was built to compute shopping lists, control a home's temperature and turn appliances on/off, although there were complaints about whether the device was successful. Another downside: it was huge."<sup>7</sup>

In the 1970's, Pico Electronics in Scotland introduced the first series of home automation products, known as X10, for plug-in control of a selected group of existing electrical devices in the home. X-10 used Radio Frequency (RF) bursts over residential power lines for communication and control. While it was not truly "automated" since there was no computer control yet available, still it provided a glimpse into what the future might hold. Initially, however, it was popular only among home hobbyists, who would tolerate its instability and susceptibility to electrical interference.<sup>8</sup> Over the years, the technology was improved, and, although it

never moved into the mainstream, there are still millions of X-10 units in use today around the world.<sup>9</sup> However, X-10 was useful for a limited product set, did not allow true interoperability as it communicated only in one direction, and operated with its own closed, proprietary protocol.

In 1984, a collaboration between IBM and the Electronics Industries Association<sup>10</sup> (EIA) sought to create what was intended to be the first open, interoperable standard for home automation. It was called the Consumer Electronics Bus, or “CEBus”. In 1985, the National Association of Home Builders (“NAHB”) agreed to adopt the standard when it was completed and funded an industry group recruited to write a complete specification for CEBus. The planned CEBus chips would allow products to communicate interactively over power lines, low voltage twisted pair wire, coaxial cables, infrared, RF, and fiber optics. It appeared that home automation was finally just over the horizon. Over the next 6 years, the CEBus group developed the standard, and it was released in 1992. However, the complex requirements and high cost of actual implementation failed to find support among manufacturers, and the EIA abandoned the effort and closed its doors in 2011.<sup>11</sup> One small company in Canada did implement the CEBus protocol, and its products are available today on a very limited basis.

While the CEBus group was laboring to define its standard, a competing effort was initiated by a start-up in Silicon Valley. It was founded by a computer industry pioneer, Mike Markkula, who had provided the initial funding and formulated the original business plan for Apple Computer. Markkula had the vision of a chip that could be inserted into any electrical product to make it “smart” and enable it to interoperate with other similarly-equipped products on a local network, which he called a “Local Operating Network” or “LON™”. Markkula named the company Echelon. He then recruited another industry heavyweight, Ken Oshman, who was one of the four founders of ROLM, principally known for its innovation of computerized PBXs for business (the company was named using the founders’ initials - he was the “O” in ROLM). ROLM was ultimately sold to IBM for over \$1-billion.

Markkula and Oshman fast-tracked their private development program, and in 1991-2 introduced the first chips, manufactured by Motorola, that used their new standard “LONWorks™”. LONWorks was a sophisticated protocol that used either twisted pair or powerline carrier to send, receive and confirm messages bidirectionally between products, for the first time enabling true interoperability among a wide range of electric and electronic devices. While it was originally intended for use in the home, the high cost and difficult implementation process caused the technology to be largely adopted in the commercial and military industries, where it found wide acceptance over the following ten years. Today, LONWorks products are used in millions of high-reliability applications around the world, including aircraft, ships, industrial and building automation, and electric meters. Approximately 90 million devices are installed with LONWorks technology in the US, EU, and China.<sup>12</sup> However, LONWorks never made it into the home.<sup>13</sup>

Since 2010, three new technologies have entered the home that seemingly can provide the infrastructure for a new generation of practical and affordable home automation products; these are the Internet, wireless networks and the smart phone. Together, the reliability, convenience, capabilities and benefits of these technologies are starting to change the attitude of consumers to one where they now may accept the idea of living in a home that is “smarter” than its predecessors.

In this paper, we will explore the idea that AI may provide the means to make the connected and truly automated home a practical and beneficial reality for the average home owner at long last. We will also examine the possible pros and cons of such an outcome.

Today new “intelligent” devices are arriving at an astonishing rate. Hundreds were displayed at the Consumer Electronics Show (“CES”) in January 2018 in the new Internet of Things (IoT) Pavilion and the Show Stoppers Hot Product Cool Companies Expo. While many at the show were excited by the possibility that these products could be interconnected over the Internet (or other networks), others expressed disappointment that the IoT lacked the compelling applications and customer service capabilities that the average consumer required, especially in the home. Implementing the missing applications and services will complete the end-to-end operational infrastructure required to bind it all together, as an integrated-platform for AI.<sup>14</sup>

Homes are places where people expect to feel safe. To bring AI into the IoT Connected Home, security, privacy and safety are paramount. To many familiar with the technology, AI promises to be a good solution. AI is the new buzzword across nearly every industry around the world. And while many AI solutions appear innocuous, we are just at the beginning of this exciting new field. As we extend AI beyond the separated domains in which it is now used, so that the Internet of Things becomes the Internet of Everything, there will be unseen dangers and potential disasters that arise. We are entering entirely uncharted territory, and we must ensure that our solutions include both extensive testing and intensive ethical scrutiny.

In this white paper, we will discuss the capabilities, potential benefits and risks we must protect against from the convergence of the IoT and AI in the home. We will also point to areas in which further research must be undertaken. In Section 2, we take an imaginary tour through an AI enhanced IoT Connected Home of the not too distant future. We will describe various aspects of the technology. We will do this by imagining two scenarios: one in which the technology performs to our best expectations, and the other in which we envision ways it fails. In Section 3, we will examine the present state of AI to give the reader a feeling for which aspects of the technology as applied to the AI-IoT Connected Home are simply innocuous, and which others might be downright dangerous. We will also address some of the ethical issues involved, including privacy, security and safety. We review current regulations and indicate the directions that future regulations are preparing to follow. We will examine the lifestyle benefits and economic value delivered to homeowners themselves that will justify the purchase of products and services that the AI-IoT Connected Home can offer as well as the challenges and pitfalls of using AI methods in our homes as the Internet of Things becomes the Internet of Everything (if in fact it does<sup>15</sup>). In Section 4, we will briefly summarize the future technology that must be developed if the AI-IoT Connected Home is to be truly accepted by the mass of consumers in their homes. These include innovative new approaches to human interface, new types of connected products, advanced sensors, appropriate infrastructure such as networks and clouds, and software applications that will tie them all together. In Section 5, we survey business models that are currently in use and can encourage companies to make the vast investments necessary to support mass deployment. Finally, in Section 6 we indicate the further research and development that is required to ensure the home continues to remain safe and private.

## 2. TWO AI-IOT CONNECTED HOME WALKTHROUGHS

In the following two scenarios we take an imaginary walkthrough by a member of the household of an IoT Connected Home outfitted with Artificial Intelligence. In the first walkthrough, the home performs seamlessly giving a range of examples of AI. In the second, there are serious problems.

### 2.1. Scenario: AI in the AI-IoT Connected Home is working as it should

You are driving home from work. The car, which has identified you before allowing you to drive it, has signaled your AI-IoT Connected Home that you will arrive in 30 minutes, so it sets the air-conditioning to your preferred comfort level. The home also checks the present electricity rate offered by your utility and decides whether to continue selling power from your solar system into the grid, or to retain it to operate the increased air-conditioning load. As you come up your driveway, the car confirms your identity to the house and signals your garage door to open. After you get out of your car and leave the garage, an RFID sensor tracks your movement, closes the garage door and turns on the lights in the hallway leading inside.

Your phone has already informed you that a package was delivered to your AI-IoT Connected Home, having taken a picture of the deliverer, and using facial recognition to find out whether they have been there before, the AI-IoT Connected Home authorizes them to place the package in a secure area on the front porch.<sup>16 17</sup> You go out the front door, which automatically unlocks for you since your identity has already been confirmed and retrieve the package. However, your neighbor is outside, and you walk over to speak with her. Since you spend more than 5 minutes in conversation, you must reauthenticate yourself by standing in front of the camera and sensor in your front entrance, which checks your RFID badge and performs facial recognition on your image. The door then unlocks, and you re-enter your home.

You walk into your home office to leave your package and briefcase and tell the AI-IoT Connected Home control to turn on the TV in the living room. As a result, the computer does not adjust the lighting to your work level and leaves the computer and monitor in sleep mode. It also does not release the electronic locks on your file cabinets, although you could do this from your smart phone if you decided to take some papers with you to read while watching your favorite news program, which is already running as you enter the room. Your home has observed that you always turn to this channel when you first get home, so it has taken over that task for you.

As you are walking into the living room, you are suddenly hungry. You change direction, and the sensors brighten the lights in the hall to the kitchen, while a speaker in the ceiling asks you what you would like to eat. You say, “just a snack”, and the system responds with a list of your favorite snacks that are in the refrigerator. As you open the refrigerator door, however, it also reminds you that your nutritionist has instructed you to cut down on your sugar and recommends a vegetable and dip that is not on your list of favorites. You decide to have the donut anyway, and the refrigerator says “tsk, tsk” as you remove it past the scanner in the refrigerator door. You figure your spouse will scold you, so you tell the home to forget your snack.

Your AI-IoT Connected Home, watches all your behaviors and those of your family members, with your express permission given through the system preferences and applied only to your selected areas of the house. This enables your home to observe and record your behavior, and after a certain number of repetitions, recognize it as a familiar habit, remember it and respond accordingly. It does this independently for each member of the family, since it has learned their faces and voices, and can recognize them locally<sup>17</sup>.



After you have settled back in the living room, the temperature has already been adjusted to your preference. When no one is in the home, the temperature is maintained at a preferred conservation level that you have set in your household preferences. The temperature in each room is adjusted in real time according to the number of current occupants and their preferences. The system mediates any conflicts between individual preferences, based on the priority of the occupants. Grandma always gets the warmest setting when she visits, since she is very sensitive to the cold.

The kitchen has many connected components including the refrigerator, range/oven and microwave, as well as other smaller appliances like coffee maker, blender etc. When you are too busy to push the buttons yourself, or you are not at home and want to start the crock pot, you can control them all remotely, using your voice, hand gestures or smart phone. The preferences have been set so that only certain members of the family access or use remote activation, so that, for example, your small children cannot accidentally turn on the oven. You just tell the stove what you want to cook, put it in the oven, and it does the rest. If you need to turn the food over, the stove sends a message to your phone.

The fridge has a screen on the front showing you what is inside, and as an item is removed, a scanner records it and puts it as a candidate in your next shopping list. Knowing what is in the fridge reduces the time and energy wasted looking around with the door open and knowing what has been removed gives you a head start in assembling your shopping list. Your screen brings up a list of items available from your grocery store, and you tap the ones you want to add to your list. Then you go back to the list, swipe across any you want to remove, and click on “buy”. Within two hours, the groceries arrive at your door, paid through your credit card on file. You can set your preferences so that certain items in the home inventory are reordered automatically, and whether they should be delivered or picked up the next time you’re in the store. Scanners in the trash bins can also keep track of items to be ordered for your inventory, and they will also tell you if you put something in the trash that should be recycled.

Cleaning the AI-IoT Connected Home can be orchestrated by the home itself. Automatic vacuuming can be routinely scheduled, and any spills can be cleaned up when they appear. Some researchers are training robots to do the daily clean up by watching how the household does it themselves. The robot can collect laundry, sorting the loads by color and starting the machines with the proper settings.

If there are elderly in the home, they can be monitored opening the fridge, or entering the kitchen, to make sure they are eating, and entering the bathroom, to make sure they are using the facilities. For those of any age who are more seriously ill, they can be monitored closely, which can include a heartbeat sensor, etc. Others can check on them remotely via the AI-IoT Connected Home. If they ask for assistance, or the monitoring senses a crisis, a call to emergency services or to another person or organization can be placed automatically by the home.

Fire and smoke alarms alerts are sent to the local Fire Company and to selected members of the household when an event is sensed. Windows and doors are equipped with sensors that trip alarms if entrance is gained by unrecognized parties while a call is placed to the local police. If maintenance is required on any of the connected devices, alerts are sent to the household and, depending on the preferences, messages are sent to specified service personnel and, depending on the preferences selected, to members of the household as well. Water and energy use are monitored and recorded and can be managed by the AI-IoT Connected Home to optimize comfort vs. cost, depending on the occupancy and preferences specified.

Cameras, microphones, recorders, thermostats, and control for lights, media and other devices are distributed throughout the house. There is a hub in each room that will operate devices in that room in response to voice commands, gestures, and instructions you send from your phone. You may choose to set your preferences when the AI-IoT Connected Home is first installed, but whether you do it then or not, the home will continually observe the actions of each of its identified occupants and suggest updates for your approval. You can also put your home in “training” mode, and then perform specific actions and instruct the home how to respond in each case. As behavior changes, the preferences can adapt and learn, either automatically or with a prompt for permission to make a change. But no matter how much you are comfortable with a home that thinks for you, you should always have manual controls and overrides - just in case. Power outages, lightning strikes other unexpected emergencies can wreak havoc with any electronic system, no matter how secure you think it may be.

Lastly, security from bad actors is constantly maintained by the home using AI to identify, quarantine and issue an alarm for threats of all kinds – from computer malware that might come in over your network to malicious vandals who walk up your sidewalk. The AI system that safeguards your home will be redundant – it lives in computers and gateways in your house as well as systems and software in the cloud. It becomes your best line of defense.

Even with an extremely well-designed system AI-IoT Connected Home System, the unexpected can occur. You are again driving home from work. Yesterday you were in a car accident, you were not at fault. The insurance company has given you a rental car to use until yours is repaired. While you were not badly injured, your face is heavily bandaged because you hit it pretty hard on the side window when the airbag inflated. Your mouth is swollen so your speech is somewhat slurred. Also, you lost your keys in the accident and you have not had a chance to get replacements made. As you pull into the driveway, the house does not recognize your rental car. You get out and stand in front of the facial recognition camera, but the bandages keep the system from recognizing you. You say a command into the microphone, but your voice is distorted by the swelling and bandages, so you are unrecognized. You get out your RFID badge and the AI-IoT Connected Home finally opens the door. As you walk in the system updates the images of you, asking you to speak and updating your current voice print. You tell the system to store the updates as “temporary.” You sit down on the couch in the living room to assess the situation. Suddenly a bell sounds informing you that someone is entering your driveway. It must be the insurance adjuster. You ask the home to show you the wide-angle image from the front door to see who it is. You recognize the adjuster. She is in a wheel chair. You get up to open the door for her. The camera won’t be able to see her. The agent is, in fact, a wounded veteran in a wheel chair. Both the facial-recognition camera and the badge ID scanner are located too high up to identify her. You walk to the door and manually let her in, making a note to yourself to call customer service to discuss and solve the problem that just presented itself.

## 2.2. Scenario: AI is wreaking havoc

Here we see the result of an AI-IoT Connected Home system that is not as well-designed, is buggy, and susceptible to outside threats. Any sort of system flaws of this kind can cause real problems.

You are driving home from work. The car, which has identified you before it will start, has signaled your home that you will arrive in 30 minutes – you receive a short message from your AI-IoT Connected Home, saying that it is on emergency power and cannot let you into the house, that the power is out, call the electric company. You call the electric company to find out what the problem is. They say they have not received your payment and sent you a disconnect notice 10 days ago with no response. They have therefore turned off the power. You discover that your on-line payment was intercepted by hackers, your bank account has been drained, your payment history was corrupted in the electric company’s system and the emailed shut off notice was deleted when it came in. You

check into a nearby hotel, using a credit account you keep separate from your other accounts since you have been targeted before. You call your financial monitoring company to let them know of the hack. You access your private databases in which you retain your financial records. You make hard copies of your electric transactions for the last year to take with you on a visit to the electric company in the morning. You have been through this before. The first time it took weeks to straighten this out. Now you are better prepared.

You are again driving home from work. There is no response from the AI-IoT Connected Home when your car signals to let it know you are 30 minutes away. You park on the street and walk to the front door. You turn your face to the camera – facial recognition, not done locally and taking a long time with occasional quirks, as in not recognizing you, verifies it is you, opens the front door, turns on the safety lights throughout the house and asks if you would like to review your messages. You tell it not now, and that it has messed up the safety lights again and to reset them. It asks you to repeat this. You repeat this several times before it recognizes what you are saying and correctly responds.

It feels very cold to you. You ask what the temperature is. You are told something that seems very off. You walk over to the nearest thermostat and find a very different number, lower than the minimum you have set for the home. Low enough to freeze the pipes if left as it is. You reboot the AI-IoT Connected Home system and reload your stored preferences, hoping for the best. It is possible that hackers have gotten into your system and corrupted it with malware, or perhaps it is simply a fault in the system hardware or software.

As the temperature returns to normal, you are bombarded by waiting communications from the home. There are too many messages to answer immediately and there is no way to turn them off or prioritize them. You are shouting at the home by this time and you tell it to turn itself off.

You walk to your home office to leave your briefcase, telling the room that you are not staying but the AI-IoT Connected Home turned on your preferred lighting, computers and entertainment center anyway, unlocking your filing cabinets as well, anyway.

You shake your head, leave the office and walk through the AI-IoT Connected Home to the kitchen. The kitchen confuses your preferences with your daughter's and her lighting, radio channel and volume almost blast you out of the room. You verbally countermand these choices. Now the fridge will not let you in. The home has placed orders and scheduled deliveries. They form an unruly pile outside the front door.

Fire and smoke alarm alerts are sent to the local Fire Company when an event is sensed. Windows and doors are equipped with sensors that trip alarms if entrance is gained by unrecognized parties, and a call is placed to the local police. You have been fined more than once in recent months for false alarms and baseless calls to the police.

The company that installed the hardware says that it is a software problem; the software company is blaming the hardware manufacturer. You are at your wits end and have no idea what to do.

### 3. ARTIFICIAL INTELLIGENCE & MACHINE LEARNING

#### 3.1. Overview

Most people do not realize how pervasive AI and Machine Learning have become already. Applications we use every day that incorporate AI include: email categorization and smart reply, mobile check deposits, fraud prevention, credit decisions, individual risk assessment, personalization of news feeds and advertisements, facial and voice recognition and text understanding.<sup>18</sup> This trend toward the increasing use of AI is steadily growing – in business, infrastructure and even in our homes. Some industry experts are becoming concerned about this unsupervised growth, and they should be. In the US, AI is almost completely unregulated by any government agency or industry organization. AI can already cause harm that is not officially recognized because it is not appropriately monitored and studied.

Let's begin with the definition of AI proposed by Poole and Markham:<sup>19</sup>

“Artificial intelligence, or AI, is the field that studies the synthesis and analysis of computational agents that act intelligently. Let us examine each part of this definition.

An **agent** is something that acts in an environment – it does something. Agents include worms, dogs, thermostats, airplanes, robots, humans, companies, and countries. We are interested in what an agent does; that is how it **acts**. We judge an agent by its actions.

An agent acts **intelligently** when

- what it does is appropriate for its circumstances and its goals,
- it is flexible to changing environments & goals,
- it learns from experience,
- and it makes appropriate choices given its perceptual and computational limitations. An agent typically cannot observe the state of the world directly; it has only a finite memory and it does not have unlimited time to act.

A **computational** agent is an agent whose decisions about its actions can be explained in terms of computation.”

The above definition sidesteps any concentration on *Human Intelligence*, which is where we often get bogged down in philosophical discussions on the nature of intelligence and do not recognize that intelligence is not solely a characteristic of humans but includes many other non-computational elements – such as inspiration, intuition and emotion, to name a few.

In the last several years, most, if not all, of the advances in what we call AI have come from the field of Machine Learning (ML).<sup>20 21</sup> One is hard pressed to find an example of AI that does not rely on ML. Some examples of pure AI are Knowledge Representation and Reasoning<sup>22</sup> which have morphed into Semantic Web Research<sup>23</sup> both of which work with ML to orient computational agents to their environments. Also included in AI are future algorithms and areas of AI that have not been discovered yet. Leaving these aside we will focus here on ML and Deep Learning (DL)<sup>24</sup> and give an overview of how this processing works. The currently accepted relationships between AI, ML and DL are illustrated in Figure 1.




Figure 1 – DL is contained in ML is contained in AI.

### 3.2. Machine Learning: Supervised and Unsupervised

Supervised learning is by far the most commonly used ML today.<sup>25</sup> Supervised learning includes parametric<sup>26</sup> and non-parametric algorithms,<sup>27</sup> support vector machines,<sup>28</sup> kernels<sup>29</sup> and neural networks.<sup>30</sup> In parametric supervised learning we use a training set made up of inputs associated with outputs. We choose a ‘hypothesis function’, which in linear regression is a function of two parameters. The training set is generally much larger than the number of parameters. We look for the straight line that best fits the training data. Minimizing the collected error between the predicted training outputs and the straight-line modeling them will give us the two parameter values. If the training set generalizes to give good predictions on new data, we have succeeded in training the ML algorithm. One example of such a trained algorithm is the prediction of housing prices.

Logistic regression<sup>31</sup> is similar and is useful for classification, for example, classification of tumors as benign or not, according to tumor size. “Application-specific similarity representations are also popular in supervised learning applications typically grouped under the name kernel function.”<sup>32</sup> The support vector machine falls in this class and is used for both regression and classification.<sup>33</sup>

Back-propagation in neural networks is also used to achieve supervised learning. In this case, for each training example, the input propagates through the network and the output is checked against its input. The connection weights are then slightly changed to decrease the error so that the estimated output will be closer to its assigned value. After a reasonable number of iterations, if the error reaches an acceptable minimum, the neural network is trained. These algorithms are founded on statistics and the outcomes or decisions are probabilities. In a classification, the doctor may say that you have a 70% chance that your tumor is benign, this implies that the tumor was classified as benign.

Unsupervised learning includes clustering,<sup>34</sup> dimensionality reduction,<sup>35</sup> recommender systems,<sup>36</sup> Information Filtering Systems<sup>37</sup> and Deep Learning(DL). Clustering is used to find structure in a set of data, to determine its probability density. There is no mapping to correct output and hence no supervision. In clustering the aim of the learning algorithm is to find groupings or classes of the data, in statistics - mixture models. These are used as recommender systems to provide customer groupings, for instance, to allow a business to profile its customers as well as to find outliers.

Dimensionality reduction is a method that reduces the size of the input data set by removing noise and maximizing information, so that the data set is easier to process.

An autoencoder, which is unsupervised, “is a type of neural network that is trained to reconstruct its input at its output. Because there are fewer intermediary hidden units than inputs, the network is forced to learn a short, compressed representation at the hidden units, which can be interpreted as a process of abstraction.”<sup>38</sup> This is an example of DL.




Figure 2 – Neural Network Autoencoder

Finally, an autonomous agent by means of ML, must be able to make decisions in the process of performing any given task. For example, an autonomous vehicle must decide on an alternate route if the planned route is somehow unavailable while travelling from location A to location B. A credit application recommender system to determine the credit worthiness of an applicant will make the credit decision.

“In order to successfully complete tasks, autonomous agents require the capacity to reason about their environment and the consequences of their actions, as well as the desirability of those consequences. The field of decision theory uses probabilistic models of the environment, called decision problems, to formalize

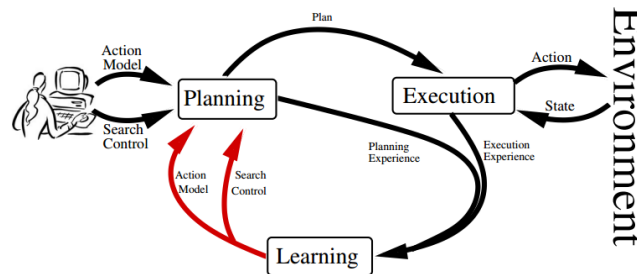


Figure 3 – Planning Process for Autonomous Agents<sup>39</sup>

the tasks about which such agents reason. Furthermore, the desirability of actions and their effects are modeled as numerical feedback signals. These feedback signals are typically referred to as reward, utility, payoff, or cost functions. Solving a decision problem consists of finding a policy, i.e., rules for how to behave in each state, that is optimal in some sense with respect to these feedback signals.”<sup>40</sup>

### 3.3. XAI – Artificial Intelligence Explained, Ethics and the EU Regulations

“The goal of Explainable Artificial Intelligence (XAI) is to create a suite of new or modified machine learning techniques that produce explainable models that, when combined with effective explanation techniques, enable end users to understand, appropriately trust, and effectively manage the emerging generation of Artificial Intelligence (AI) systems. Proposed research should investigate innovative approaches that enable revolutionary advances in science, or systems.”<sup>41</sup> This initiative - announced in 2016 by the Defense Advanced Research Projects Agency (DARPA) has been put in place in response to what is ‘often called the “black box” problem — the inability to discern exactly what machines are doing when they’re teaching themselves novel skills — and it has become a

central concern in artificial-intelligence research.<sup>42</sup> The Defense Department clearly sees a critical need for the solution of this problem in its own application space.

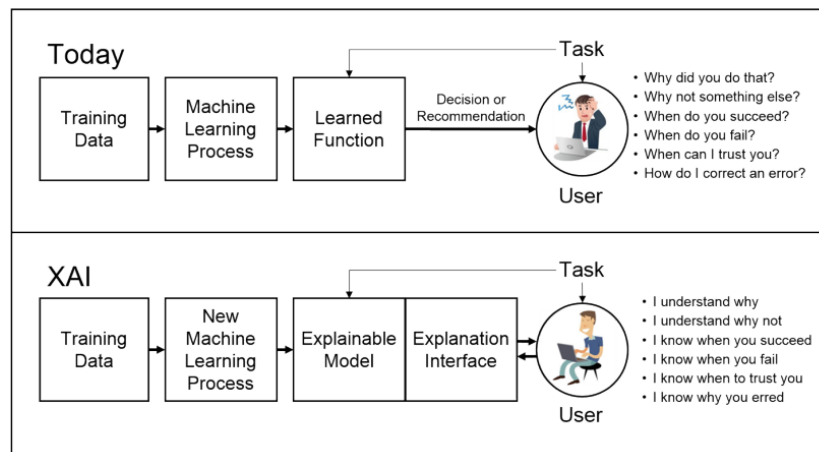


Figure 4 – XAI, Explainable AI

This is not just a matter of unsupervised learning, although that is where the bulk of the black box problem lies. It can also happen in supervised algorithms as the following excerpt shows.

'Rich Caruana, an academic who works at Microsoft Research, has spent almost his entire career in the shadow of this problem. When he was earning his Ph.D. at Carnegie Mellon University in the 1990s, his thesis adviser asked him and a group of others to train a neural net — a forerunner of the deep neural net — to help evaluate risks for patients with pneumonia. Between 10 and 11 percent of cases would be fatal; others would be less urgent, with some percentage of patients recovering just fine without a great deal of medical attention. The problem was figuring out which cases were which — a high-stakes question in, say, an emergency room, where doctors have to make quick decisions about what kind of care to offer. Of all the machine-learning techniques students applied to this question, Caruana's neural net was the most effective. But when someone on the staff of the University of Pittsburgh Medical Center asked him if they should start using his algorithm, "I said no," Caruana recalls. "I said we don't understand what it does inside. I said I was afraid." The problem was in the algorithm's design. Classical neural nets focus only on whether the prediction they gave is right or wrong, tweaking and weighing and recombining all available morsels of data into a tangled web of inferences that seems to get the job done. But some of these inferences could be terrifically wrong. Caruana was particularly concerned by something another graduate student noticed about the data they were handling: It seemed to show that asthmatics with pneumonia fared better than the typical patient. This correlation was real, but the data masked its true cause. Asthmatic patients who contract pneumonia are immediately flagged as dangerous cases; if they tended to fare better, it was because they got the best care the hospital could offer.'<sup>43</sup>

In this case, it was not that the learning was unsupervised, but that the data was not clearly understood.

In 2016, out of concern for what is coming out of the field and perhaps somewhat late, the Institute of Electrical and Electronics Engineers (IEEE) initiated a 'TechEthics' program to foster consideration of Ethics in AI. The goal is to ensure that “ethical and societal implications are considered in AI design and deployment.”<sup>44</sup> Some of the issues brought up at their last convention were:

- Who should be held responsible for the harm an application causes by its actions.
- Researchers placed 4 black and white stickers on a stop sign. A self-driving car interpreted the sign to be a speed limit sign and sped up. Insufficient testing?
- A Microsoft bot named Tay began learning to engage in pleasant and playful conversations on Twitter and within 24 hours was tweeting misogynist and racist comments it picked up from other Twitter users.

Another issue that fits into the IEEE’s initiative, is that a commonly used ML technique leveraging public Facebook data to identify gay individuals had better accuracy than a human’s. The researcher, Michal Kosinski, by publishing his findings was raising questions about privacy and the potential for discrimination in the digital age. The learning was unsupervised, and the rationale was not explainable.<sup>45</sup> Since, in some countries this identification could bring a death sentence, this is a result that should be subject to ethical consideration, which was the author’s intention.

An additional step forward in this area was that the new National Institute of Standards and Technology(NIST) draft embeds privacy into US government security for the first time, expanding its scope to include the internet of things and smart home technology.<sup>46</sup>

Those worried that regulation might stifle progress in AI have loud voices. However, the DARPA, IEEE and NIST initiatives are a welcome, and perhaps overdue, response.

Even more compelling however are the European Union’s General Data Protection Regulations (GDPR) that go into effect on May 25, 2018. Although, “the bulk of the language deals with how data is collected and stored, the regulation contains Article 22: Automated individual decision making, including profiling, potentially prohibiting a wide swath of algorithms currently in use in recommendation systems, credit and insurance risk assessments, computational advertising, and social networks, for example.”<sup>47</sup> As well, the GDPR asserts that “Citizens have the right to receive an explanation for algorithmic decisions.”<sup>48</sup> Since, “ML depends upon data that has been collected from society, and to the extent that society contains inequality, exclusion, or other traces of discrimination, so too will the data.”<sup>49</sup> This boils down to the following: “black box” techniques as well as some supervised techniques that are using data collected from society and for which no measures have been taken that provably remove any discrimination, will not be allowed. Any US company with EU clients will be obliged to honor these regulations.

### 3.4. Advantages and Challenges to AI in the Connected Home

There are many advantages, but there are also challenges. First some of the advantages:

- Energy usage can be monitored for either cost or energy saving or both. Multiple sources of energy can be optimized.
- Lighting and other common items can be automatically turned on or off by schedule or by room population.
- Diet can be supervised or not depending on personal preferences.
- Home maintenance can be monitored and calls for service can be communicated to the home owners or service people.



- Fire, toxic air quality and intrusion can be monitored, and alerts sent to the home owner and/or the fire and police.
- The sick or elderly can be monitored with cameras, listening devices, and medical instruments such as heartbeat monitors, etc.

Now for the challenges:

- Privacy is an important concern. New regulations are being instituted by the National Institute to ensure privacy.<sup>50</sup>
- Safety is also an important concern and new laws must be put in place in case a safety incident causes one of the household harm. California has just mandated a year-long human observed testing period for autonomous cars.<sup>51</sup>
- In smart speaker systems incorporating voice recognition, computer science researchers have figured out ways to completely fool these systems by slightly altering pixels in a photo or adding faint noises to audio files. These minute tweaks are undetectable by humans, but completely alter what an A.I. hears or sees.<sup>52</sup>
- Voice interfaces are a way Amazon, Google and Apple can gather information by “eavesdropping” —not just on shopping preferences and other internet activities—but also about how they behave and interact with one another in the home itself.<sup>53</sup>
- With the techniques of 3 & 4 above, hackers or other bad actors can set off fire and police alarms, thermostats can be lowered or raised, voice archives can be subpoenaed in court cases, medical devices can be breached, etc.

#### 4. TECHNICAL DETAILS – AI AND THE IOT

To merge AI and IoT in the Connected Home, we need to discuss the infrastructure required to do so. There are several critical areas. These include the cloud, fog computing, security, safety, regulation, semantic web and resource scheduling, among others. Today, we have minimal infrastructure supporting the AI that is already in the home. Alexa, Siri, and Cortana are minimal, but do serve to introduce people to the possibilities to come from the IoT and AI.

The Internet of Things quite literally means ‘things’ or ‘objects’ that connect to the Internet — and each other. This could be almost anything — a computer, a smartphone, a door lock, a crock pot, an HVAC unit, ...to name a few. Communication typically takes place in three different ways: “machine to machine (M2M), human to machine (H2M), and machine to smartphone (M2S)”<sup>54</sup> or another device. When controlling a more complicated system or infrastructure, the smartphone will probably be augmented by more complicated units. “The sum of all this is the Internet of Everything (IoE).”<sup>55</sup> More research and experience is needed before we see a seasoned AI-IoT Connected Home System.

Cloud computing, due to its on-demand processing and storage capabilities, can be used to analyze data generated by IoT objects, for example sensor outputs used in ML algorithms, in batch or stream format. A pay-as-you-go model adopted by all Cloud providers has reduced the price of computing, data storage and data analysis, creating a streamlined process for building IoT applications. Proposed solutions that “only utilize cloud computing as a processing or storage backbone are not scalable and cannot address the latency constraints of real-time applications.”<sup>56</sup> Real-time processing requirements and the increase in computational power of edge devices such as routers, switches, and access points lead to the emergence of the Edge Computing paradigm. The Edge layer contains the devices that are in closer proximity to the end user than the application servers, and can include

smartphones, room control units, smart TVs, network routers, data storage, simplified ML analysis on smaller data sets that predict and optimize energy usage, small set facial and voice recognition, preference storage and so forth. Processing and storage capability of these devices can be utilized to extend the advantages of using Cloud computing by creating another cloud, known as the Edge Cloud, near application consumers, to decrease networking delays, save processing or storage cost, perform data aggregation, and prevent sensitive data from leaving the local network.<sup>57</sup>

Fog Computing is an extension of cloud computing that aims to keep the same features of Cloud, such as networking, computation, virtualization, and storage, but also meets the requirements of applications that demand low latency<sup>58</sup> and foster privacy protection. Fog Computing is a form of Edge Cloud that utilizes distributed processing at the network edge.

Recent innovations in AI, both supervised and unsupervised, as applied to intruder protection and to virus quarantine at the edge have already improved threat elimination, and will increase in their ability to eliminate threats, although this is still a worry.<sup>59</sup> Another option to consider in eliminating threats is the use of non-Wi-Fi networks such as Bluetooth and Zigbee, among others, within the home. Systems that provide multiple network access are already available. As well there are several vendors now offering routers that can monitor and protect all connected devices within the Connected Home. New Operating Systems are being developed specifically for the IoT which provide resource management, workload balancing, task scheduling, etc. As adoption of IoT continues to grow, attackers and malicious users are shifting their target from servers to end devices.<sup>60</sup>

Developers and business managers are advised to focus on developing and sharing APIs from the early stage of their application development lifecycle, so that eventually, by properly exposing data to other developers and end users, an open-data environment is created that facilitates collaborative information gathering, sharing, and updating.<sup>61</sup> Even though more organizations and industries make themselves ready to embrace and incorporate IoT, increase in IoT growth rate will cause difficulties for standardization. Strict regulations must be put into place about accessing radio frequency levels, creating a sufficient level of interoperability among different devices, authentication, identification, authorization, and communication protocols are all open challenges facing IoT standardization.<sup>62</sup>

## 5. BUSINESS MODELS

Depending on the type of computing resources delivered via the cloud, cloud services take different forms, such as Infrastructure as a service (IaaS), Platform as a service (PaaS), Software as a service (SaaS), Storage as a service (STaaS), and more.<sup>63</sup> These models have been working for the cloud, and they are even more appropriate for the IoT and AI.

The above service models are defined by the National Institute of Standards and Technology (NIST) as follows.<sup>64</sup> Also included in the cited NIST document are definitions of Cloud Infrastructure and their essential characteristics.

Infrastructure as a Service (IaaS). “The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and

possibly limited control of select networking components (e.g., host firewalls).”<sup>65</sup> The IaaS service definition includes Storage as a Service (STaaS).

An example of the use of IaaS in the AI-IoT Connected Home might be to store large amounts of sensor data and ML algorithms in the cloud and run analytics there. Fog computing and an Edge Cloud might be used to recognize faces, voices and haptics for a household, with more general facial recognition being done in the cloud where access to larger databases would allow for wider recognition. You would pay as you go depending on what software and storage you use as well as processing time. You might have a monthly cost for the home automation infrastructure (bundled).

Platform as a Service (PaaS). “The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.”<sup>66</sup>

These sorts of service models would be used by developers, either fixing bugs, enhancing currently running software or developing new software. Analysts would be here working with customers to determine where their problems lie and how to fix them, i.e., providing customer service.

Software as a Service (SaaS). “The capability provided to the consumer is to use the provider’s applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, apart from limited user-specific application configuration settings.”<sup>67</sup>

Here the user is using software, for example facial recognition or energy usage planning, i.e., do I use the solar unit’s energy or that from the electric company – this would depend on how much solar I can expect to store and use today vs. at what price could I sell it to the electric company vs. the current rate the electric company is charging. Other examples include security software for threat elimination, elder care, progressive health monitoring and archiving of all visitor images and conversations.

## 6. CONCLUSIONS

We see that the AI-IoT Connected Home can provide some real benefits to its occupants – but it is not without potential dangers.

Homeowners, in general, are not technologists. They expect things to work almost perfectly and are quick to abandon new products if they fail to meet expectations. They have little patience for systems that can malfunction and cause them embarrassment in front of friends, family and neighbors.

The challenge we face is produce AI-IoT systems for the Connected Home that are intelligent, learn quickly and operate flawlessly in all situations. Most important, they must not allow threats or perceived threats to a home’s occupants in any way or intrude on their privacy.

To accomplish this, further work and development is required in the following areas:

- **User Interface** – the AI-IoT Connected Home must understand the needs and desires of its occupant in natural ways, however they choose to express them. Homeowner Focus groups must meet to suggest ways they would like to see the user interface designs go forward as well as the customer support is interlaced with their systems.
- **Threat Elimination** – using AI techniques to discover and neutralize threats and to quarantine malware at the network edge. These products are already available using these techniques.
- **Ethical Considerations** – as we have seen, there are serious ethical concerns involved with people in their homes. People expect an exceptionally high degree of privacy, safety and convenience in their homes. These and other factors must be extensively analyzed with respect to the operation of the AI-IoT Connected Home.
- **Ease of Understanding and Intuitive Operation** – the average person does want to work to understand what their home system is doing and why it is doing it. All the operations must make sense to the homeowner, and the system should clearly inform them what it is doing – before, during and after it does it. Operational approvals and manual overrides are a must to provide confidence to the residents.
- **Convenience of Service** – trained service personnel must be available to quickly address any failures or malfunctions – ideally before the homeowner knows they exist. Hardware systems must be assembled with “hot-pluggable,” or “AI-IoT plug and play,” components that can be exchanged without requiring a home shutdown.
- **Ability to Upgrade as Technology Advances** – while consumers generally want the “latest and greatest”, they are also extremely sensitive to price and convenience. New products and software upgrades must be inexpensive and easy to install.

As AI-IoT technology for the Connected Home evolves, we must remain diligent in analyzing the implications of each new step in the light of the above considerations. Only by doing so will we create and deliver products and systems that achieve broad public acceptance and enthusiastic adoption.

## REFERENCES AND GLOSSARY

- 
- <sup>1</sup> Hogan, J. 2015. Bette, Teller, Trinity and the End of the Earth. *Scientific American*. Aug.4.
- <sup>2</sup> Dudley, H.C. 1975. The Ultimate Catastrophe. *Bulletin of the Atomic Scientists*. Nov.
- <sup>3</sup> Bette, H. 1976. The Ultimate Catastrophe. *Bulletin of the Atomic Scientists*. Nov.
- <sup>4</sup> Florida Museum. International Shark File. Available from: <https://www.floridamuseum.ufl.edu/shark-attacks/odds/compare-risk/death/>
- <sup>5</sup> Haltiwanger, J. 2015. Elite Daily. Available from: <https://www.elitedaily.com/news/world/people-terrified-plane-crashes-even-though-rare/977885>.
- <sup>6</sup> CNBC. Finance. 2018. Available from: <https://www.cnbc.com/2018/01/05/odds-of-winning-a-lottery-jackpot-are-worse-than-you-expect.html>.
- <sup>7</sup> Rothfeld, L. 2015. TECH Time Machine: The Smart Home. Available from: <https://mashable.com/2015/01/08/smart-home-tech-CES/#nhVFP8VgGkqqa>.
- <sup>8</sup> Ibid.
- <sup>9</sup> Wikipedia contributors. X10 (industry standard) [Internet]. Wikipedia, The Free Encyclopedia; 2018 Feb 1, 17:56 UTC [cited 2018 Feb 24]. Available from: [https://en.wikipedia.org/w/index.php?title=X10\\_\(industry\\_standard\)&oldid=823506497](https://en.wikipedia.org/w/index.php?title=X10_(industry_standard)&oldid=823506497).
- <sup>10</sup> Wikipedia contributors. Electronic Industries Alliance [Internet]. Wikipedia, The Free Encyclopedia; 2018 Jan 16, 07:00 UTC [cited 2018 Feb 25].
- <sup>11</sup> Wikipedia contributors. CEBus [Internet]. Wikipedia, The Free Encyclopedia; 2017 Jan 20, 22:05 UTC [cited 2018 Feb 24]. Available from: <https://en.wikipedia.org/w/index.php?title=CEBus&oldid=761100655>.
- <sup>12</sup> Wikipedia contributors. LonWorks [Internet]. Wikipedia, The Free Encyclopedia; 2017 Dec 4, 15:44 UTC [cited 2018 Feb 24]. Available from: <https://en.wikipedia.org/w/index.php?title=LonWorks&oldid=813656245>.
- <sup>13</sup> Silverman, L. Personal recollections.
- <sup>14</sup> Feland, J. 2017. *State of the Smart Home Market*. Webinar. Argus Insights. Q2.
- <sup>15</sup> Breene, L. I remember one of the first lessons I was taught in computer science – one need not computerize everything.
- <sup>16</sup> The facial recognition would in this case be done in the cloud, since fog computing (see 17) would not be appropriate here because of the larger population it would have to recognize.
- <sup>17</sup> If the delivery person has not been identified previously, or worse, is a bad actor, you are given a call, the situation is explained to you, and you are connected to that person. You can then ask them to leave the package or ask them to redeliver at a time that you are at home. – this could be one response, or the response could vary according to the preferences established for this scenario.
- <sup>18</sup> Narula, G. 2018. *Everyday Examples of Artificial Intelligence and Machine Learning*. Available from: <https://www.techemergence.com/everyday-examples-of-ai>.
- <sup>19</sup> Poole, D. L. and Mackworth, A. K. *Artificial Intelligence*. Cambridge University Press. 2010. (kindle location 257).
- <sup>20</sup> Alpaydin, E. *Machine Learning: The New AI*. (The MIT Press Essential Knowledge series). The MIT Press. (kindle location 147).
- <sup>21</sup> Machine learning is a field of computer science that gives computer systems the ability to "learn" (i.e. progressively improve performance on a specific task) with data, without being explicitly programmed. Supposedly paraphrased from: Samuel, Arthur (1959). "Some Studies in Machine Learning Using the Game of Checkers". *IBM Journal of Research and Development*. 3 (3).
- <sup>22</sup> Examples of knowledge representation formalisms include semantic nets, systems architecture, frames, rules, and ontologies. Wikipedia contributors. Knowledge representation and reasoning [Internet]. Wikipedia, The Free Encyclopedia; 2018 Feb 15, 07:06 UTC [cited 2018 Mar 10]. Available from: [https://en.wikipedia.org/w/index.php?title=Knowledge\\_representation\\_and\\_reasoning&oldid=825761182](https://en.wikipedia.org/w/index.php?title=Knowledge_representation_and_reasoning&oldid=825761182).
- <sup>23</sup> 'The Semantic Web is an extension of the World Wide Web through standards by the World Wide Web Consortium (W3C). The standards promote common data formats and exchange protocols on the Web, most fundamentally the Resource Description Framework (RDF). According to the W3C, "The Semantic Web provides a common framework that allows data to be shared and reused across application, enterprise, and community boundaries". The Semantic Web is therefore regarded as an integrator across different content, information applications and systems.' Wikipedia contributors. Semantic Web [Internet]. Wikipedia, The Free Encyclopedia; 2018 Feb 27, 21:17 UTC [cited 2018 Mar 11]. Available from: [https://en.wikipedia.org/w/index.php?title=Semantic\\_Web&oldid=827978771](https://en.wikipedia.org/w/index.php?title=Semantic_Web&oldid=827978771).

<sup>24</sup> “Deep learning is a type of machine learning in which a model learns to perform classification tasks directly from images, text, or sound. Deep learning is usually implemented using a neural network architecture. The term “deep” refers to the number of layers in the network—the more layers, the deeper the network. Traditional neural networks contain only 2 or 3 layers, while deep networks can have hundreds.” Mathworks. Deep\_Learning\_ebook. 2017. Available from:

[https://www.mathworks.com/content/dam/mathworks/tagteam/Objects/d/80879v00\\_Deep\\_Learning\\_ebook.pdf](https://www.mathworks.com/content/dam/mathworks/tagteam/Objects/d/80879v00_Deep_Learning_ebook.pdf).

<sup>25</sup> Alpaydin. (kindle location 177).

<sup>26</sup> “A learning model that summarizes data with a set of parameters of fixed size (independent of the number of training examples) is called a parametric model. No matter how much data you throw at a parametric model, it won’t change its mind about how many parameters it needs.” Russell, S. and Norvig, P. 2009. *Artificial Intelligence: A Modern Approach* (3rd Edition). Pearson. p. 737.

<sup>27</sup> “Nonparametric methods are good when you have a lot of data and no prior knowledge, and when you don’t want to worry too much about choosing just the right features.” Russell, S. and Norvig, P. 2009. *Artificial Intelligence: A Modern Approach* (3rd Edition). Pearson. p. 757.

<sup>28</sup> “Suppose some given data points each belong to one of two classes, and the goal is to decide which class a *new data point* will be in. In the case of support vector machines, a data point is viewed as a p-dimensional vector (a list of p numbers), and we want to know whether we can separate such points with a (p-1)-dimensional hyperplane. This is called a linear classifier.” Wikipedia contributors. Support vector machine [Internet]. Wikipedia, The Free Encyclopedia; 2018 Mar 1, 18:53 UTC [cited 2018 Mar 7]. Available from:

[https://en.wikipedia.org/w/index.php?title=Support\\_vector\\_machine&oldid=828295144](https://en.wikipedia.org/w/index.php?title=Support_vector_machine&oldid=828295144).

<sup>29</sup> “In machine learning, kernel methods are a class of algorithms for pattern analysis, whose best known member is the support vector machine (SVM).” Wikipedia contributors. Kernel method [Internet]. Wikipedia, The Free Encyclopedia; 2018 Jan 7, 19:15 UTC [cited 2018 Mar 7].

Available from: [https://en.wikipedia.org/w/index.php?title=Kernel\\_method&oldid=819149994](https://en.wikipedia.org/w/index.php?title=Kernel_method&oldid=819149994).

<sup>30</sup> ‘Artificial neural networks (ANNs) or connectionist systems are computing systems vaguely inspired by the biological neural networks that constitute animal brains. Such systems “learn” (i.e. progressively improve performance on) tasks by considering examples, generally without task-specific programming.’ Wikipedia contributors. Artificial neural network [Internet]. Wikipedia, The Free Encyclopedia; 2018 Mar 1, 17:55 UTC [cited 2018 Mar 7]. Available from: [https://en.wikipedia.org/w/index.php?title=Artificial\\_neural\\_network&oldid=828287834](https://en.wikipedia.org/w/index.php?title=Artificial_neural_network&oldid=828287834).

<sup>31</sup> Logistic regression uses the sigmoid function, a non-linear function that can be used to divide classes into two parts, in supervised classification by regression.

<sup>32</sup> Alpaydin. pp. 116-117.

<sup>33</sup> Vapnik, V. 1998. *Statistical Learning Theory*. New York. Wiley.

<sup>34</sup> A loose definition of clustering could be “the process of organizing objects into groups whose members are similar in some way”. A **cluster** is therefore a collection of objects which are “similar” between them and are “dissimilar” to the objects belonging to other clusters. These techniques are unsupervised. Available from: [https://home.deib.polimi.it/matteucc/Clustering/tutorial\\_html/index.html](https://home.deib.polimi.it/matteucc/Clustering/tutorial_html/index.html).

<sup>35</sup> Dimensionality reduction is an unsupervised technique that uses mathematical transformations of the data space to reduce noise and allow for better feature recognition.

<sup>36</sup> ‘A recommender system or a recommendation system (sometimes replacing “system” with a synonym such as platform or engine) is a subclass of information filtering system that seeks to predict the “rating” or “preference” a user would give to an item.’ Wikipedia contributors. Recommender system [Internet]. Wikipedia, The Free Encyclopedia; 2018 Mar 2, 14:32 UTC [cited 2018 Mar 11]. Available from:

[https://en.wikipedia.org/w/index.php?title=Recommender\\_system&oldid=828433000](https://en.wikipedia.org/w/index.php?title=Recommender_system&oldid=828433000).

<sup>37</sup> Hanani, U., Shapira, B. and Shoal, P. 2001. Information Filtering: Overview of Issues, Research and Systems. *User Modeling and User-Adapted Interaction*. 11: pp. 203-209.

<sup>38</sup> Alpaydin, E. p. 169.

<sup>39</sup> Jimenez, S. et al. A Review of Machine Learning for Automated Planning. *The Knowledge Engineering Review*. Vol. 00:0. Pp. 1–24. c 2009. Cambridge University Press.

<sup>40</sup> Roijers, D. and Whiteson, S. 2017. *Multi-Objective Decision Making*. (Synthesis Lectures on Artificial Intelligence and Machine Learning). Elsevier. p.1.

<sup>41</sup> Broad Agency Announcement. 8/10/2016. Explainable Artificial Intelligence (XAI). DARPA-BAA-16-53.

<sup>42</sup> Huang, C. 2017. Can AI be Taught to Explain itself? *New York Times Magazine*. p. 4. Available from: <https://www.nytimes.com/2017/11/21/magazine/can-ai-be-taught-to-explain-itself.html>.

<sup>43</sup> Ibid. p. 5.

- 
- <sup>44</sup> The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. Available from: [https://standards.ieee.org/develop/indconn/ec/autonomous\\_systems.html](https://standards.ieee.org/develop/indconn/ec/autonomous_systems.html).
- <sup>45</sup> Huang. pp. 2-3.
- <sup>46</sup> McCarthy, K. 2017. New NIST draft embeds privacy into U.S. government security for the first time. *The Register*. Aug 24. Available from: [https://www.theregister.co.uk/2017/08/18/new\\_nist\\_draft\\_embeds\\_privacy\\_into\\_security\\_for\\_the\\_first\\_time/](https://www.theregister.co.uk/2017/08/18/new_nist_draft_embeds_privacy_into_security_for_the_first_time/)
- <sup>47</sup> Goodman, B. and Flaxman, S. 2017. European Union Regulations on Algorithmic Decision Making and a “Right to Explanation”. *AI Magazine* 38(3): 50-57, *Association for the Advancement of Artificial Intelligence*. p. 51.
- <sup>48</sup> Ibid, p. 51.
- <sup>49</sup> Ibid, p. 53.
- <sup>50</sup> McCarthy. 2017.
- <sup>51</sup> California Department of Motor Vehicles. 2017. Adopted Regulations for Testing of Autonomous Vehicles by Manufacturers. Available from: <https://www.dmv.ca.gov/portal/dmv/detail/vr/autonomous/testing>.
- <sup>52</sup> Paez, D. 2018. Smart Speakers Can be hacked with sounds, say researchers out to stop it. Available from: <https://www.inverse.com/article/40367-researchers-can-fool-speech-recognition-a-i-with-this-trick>.
- <sup>53</sup> Wueest, C. 2017. A guide to the security of voice-activated smart speakers. Symantec Special Report. Available from: <https://www.symantec.com/content/dam/symantec/docs/security-center/white-papers/istr-security-voice-activated-smart-speakers-en.pdf>
- <sup>54</sup> Greengard, S. 2015. *The Internet of Things*. The MIT Press. Kindle location 304.
- <sup>55</sup> Ibid. Kindle location 304.
- <sup>56</sup> Buyya, R. and Dastjerdi. A.V. Eds. 2016. *Internet of Things: Principles and Paradigms. Elsevier Science*. p 13.
- <sup>57</sup> Ibid. p. 14.
- <sup>58</sup> Ibid. p. 14.
- <sup>59</sup> For example, Threat Stack uses unsupervised learning to stop threats inside and outside the network edge. Bluvector uses supervised learning to quarantine malware at the network edge. Norton offers the Core router to protect up to 50 devices within the home network and provides deep web blocking and quarantine.
- <sup>60</sup> Buyya & Dastjerdi. p. 19.
- <sup>61</sup> Buyya & Dastjerdi. pp. 9-10.
- <sup>62</sup> Buyya & Dastjerdi. p. 22.
- <sup>63</sup> Buyya & Dastjerdi. p. 29.
- <sup>64</sup> Mell, P. and Grance, T. (September 2011). *The NIST Definition of Cloud Computing (Technical report)*. National Institute of Standards and Technology: U.S. Department of Commerce. [doi:10.6028/NIST.SP.800-145](https://doi.org/10.6028/NIST.SP.800-145). *Special publication 800-145*. pp. 1-3.
- <sup>65</sup> Ibid. p. 3.
- <sup>66</sup> Ibid. pp. 2-3.
- <sup>67</sup> Ibid. p. 2.

---

## GLOSSARY

### **Artificial neural networks (ANNs)**

– are computing systems vaguely inspired by the biological neural networks that constitute animal brains. Such systems "learn" (i.e. progressively improve performance on) tasks by considering examples, generally without task-specific programming. They are also known as connectionist systems.

### **Clustering**

- the process of organizing objects into groups whose members are similar in some way. A **cluster** is therefore a collection of objects which are "similar" between them and are "dissimilar" to the objects belonging to other clusters. These techniques are unsupervised.

### **Deep Learning**

- a type of machine learning in which a model learns to perform classification tasks directly from images, text, or sound. Deep learning is usually implemented using a neural network architecture. The term "deep" refers to the number of layers in the network—the more layers, the deeper the network. Traditional neural networks contain only 2 or 3 layers, while deep networks can have hundreds.

### **Dimensionality Reduction**

- an unsupervised technique that uses transformations of the data space to reduce noise and allow for better feature recognition.

### **Information Filtering Systems**

– to expose users only to information that is relevant to them. Many IF systems have been developed in recent years for various application domains. Some examples of filtering applications are: filters for search results on the internet that are employed in the Internet software, personal e-mail filters based on personal profiles, listservers or newsgroups filters for groups or individuals, browser filters that block non-valuable information, filters designed to give children access only to suitable pages, filters for e-commerce applications that address products and promotions to potential customers only, and many more.

### **Kernel Methods**

- are a class of machine learning algorithms for pattern analysis, whose best known member is the support vector machine.

### **Logistic regression**

- uses the sigmoid function, a non-linear function that can be used to divide classes into two parts, in supervised classification by regression.

### **Machine Learning**

- a field of computer science that gives computer systems the ability to "learn" (i.e. progressively improve performance on a specific task) with data, without being explicitly programmed.

### **Nonparametric Model**

- used when you have a lot of data and no prior knowledge, and when you don't want to worry too much about choosing just the right features.

### **Parametric Model**

- a learning model that summarizes data with a set of parameters of fixed size (independent of the number of training examples). No matter how much data you throw at a parametric model, it won't change its mind about how many parameters it needs.

### **Recommender System**

- is a subclass of information filtering systems that seek to predict the "rating" or "preference" a user would give to an item, also called a recommendation system (sometimes replacing "system" with a synonym such as platform or engine).

### **Semantic Web**

- an extension of the World Wide Web through standards by the World Wide Web Consortium (W3C).<sup>11</sup> The standards promote common data formats and exchange protocols on the Web, most fundamentally the Resource Description Framework (RDF). According to the W3C, "The Semantic Web provides a common framework that allows data to be shared and reused across application, enterprise, and community boundaries." The Semantic Web is therefore regarded as an integrator across different content, information applications and systems.'

### **Support Vector Machines (SVM)**

– a set of data points each belong to one of two classes, and the goal is to decide which class a *new data point* will be in. In this case a data point is viewed as a p-dimensional vector (a list of p numbers), and we want to know whether we can separate these points with a (p-1)-dimensional hyperplane. This is called a linear classifier (since a (p-1)-dimensional hyperplane is similar to a straight line in p-dimensional space).





© CABA 2018

888.798.CABA (2222)

613.686.1814

Connect to what's next™

[www.caba.org](http://www.caba.org)

