Foundation and Evolution of AI + ADIN
May 15, 2023
Echo Messaging Systems, Inc.
info@echomessaging.com

Foundation and Evolution of AI + ADIN™

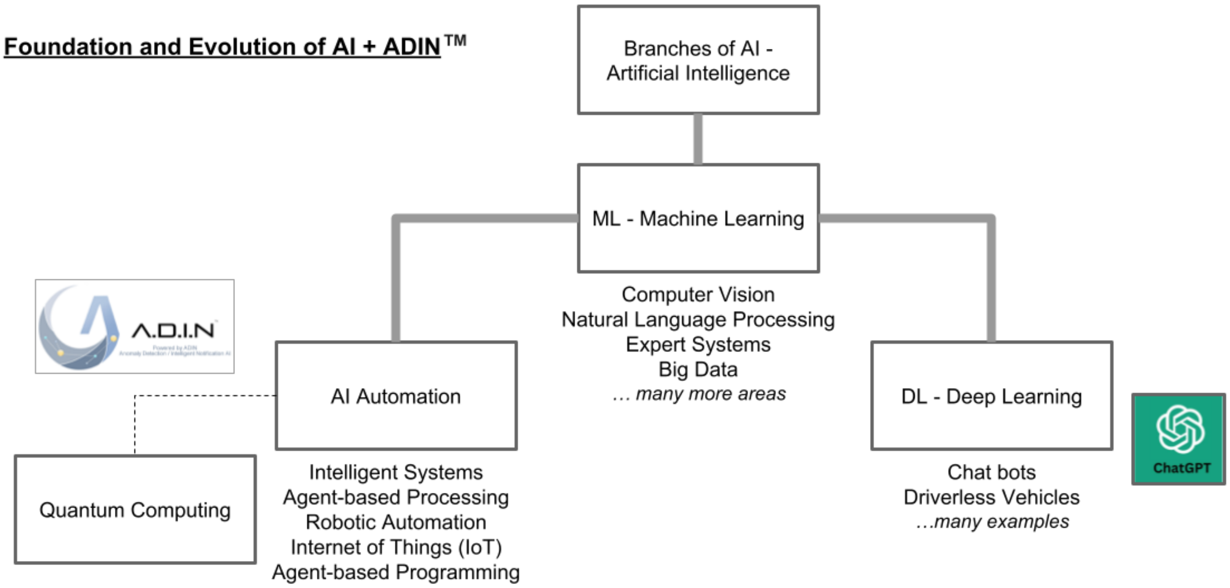

AI is a branch of computer science that reflects human intelligence. For example, computer vision is getting computers to see what's in a picture. Natural language processing is getting the computer to understand what is said or written and how to respond. Since the dawn of humanity, humans are best at adapting to their environments -- from harnessing the power of fire and tools, to exploring the globe on simple wooden vessels, to understanding, to the best of our ability, how the universe works in time and space. Software that adapts or learns is a key aspect of AI because humans are so good at evolving.

AI is software. It's different from regular software because it's complex, mimics human behavior, or is adaptive in nature. An old joke in AI is that once something's been around for long enough, it's no longer AI. Graphic artists use drawing software that have lots of features that used to be considered AI…it's just graphics software now, that is available even in your camera-phone's image editor. Clicking on an object to be removed and have the background fill in like it was never there..that's a pretty sophisticated piece of old AI. Homeowners who install heating/AC controls that remember when you leave and come back from work/school every day, but is different on the weekends..that's adaptive AI software. It's great to come home to a warm house in the winter, knowing that the heat kicks on automatically 15 minutes before you get home. Saves money and makes life comfortable..AI does good things.

Computers that see, computers that hear, computers that talk, computers that remember, computers that decide things on their own are all areas of AI. Lots of different types of software are used to make computers 'artificially intelligent'. But they always mimic how people are. This of course is both good and bad.

As you see in the figure, Machine Learning started as the main 'umbrella' branch of AI. Deep Learning branched off as it became an important topic of research and applications. Automation is the area of AI that we are focused on because we saw the power of Intelligent Systems as a platform for solving a wide variety of real-world problems. All areas of AI have experienced profound forward momentum, with Deep Learning recently getting the most attention with the breakout success on ChatGPT.

To understand how AI makes computers 'intelligent' I think of this example..when we fly in planes, the wings don't flap like when a bird flies, but flying mimics bird behavior. Some AI is written to be like the bird with the wings flapping..mirroring intelligence on a functional level, while other types of AI are like airplanes..they fly but wings don't flap..it accomplishes the goal, but does it differently from what is being modeled.

Deep Learning is based on a topic of AI called 'Neural Networks'. This is a really important area of AI that works how the brain thinks and adapts by mirroring the brain on a functional level. Humans have a brain filled with 10's of billions of neurons. A neuron is a cell with paths leading into it that connect from other neurons and a path leading out that when the neuron 'fires', it sends an electrical signal to neurons it's connected to …making a network of neurons. When a person learns something, neurons and neuron pathways are changed at a chemical level to fire more easily based on the thing just learned. It's hard to imagine this works, but with an average of 86 billion neurons in our brains, this is how we learn, how we speak, how we think, how our lungs function without us thinking about breathing, and how amazingly creative things in art and science are produced..of course this just our are general understanding. The brain is not really well understood.

Deep Learning uses deeply layered neural networks modeled on the brain, that are first trained on specific things and once trained can tell you if something new is like something it was trained on or not. Like reading a picture book to a child, where you point to the picture of a cat and say 'cat' or point to a dog and say 'dog', neural networks are presented digital information like documents, or pictures and are told what it is. Internally, neural networks get vastly huge sets of numbers that weight each neuron (modelling the chemical changes). The trained neural networks interact with new data and the networks are now deeply layered enough (hence the description 'Deep Learning') that they do a good job of identifying things or recombining things based on how they were trained to act like how humans act.

ChatGPT is doing this now. Deep Learning algorithms have a 'black box' problem..the weighted training numbers can't explain how it decides anything new. It continues to update the weighted numbers inside the neural network, and sometimes starts to veer off-course..requiring ongoing training. Deep Learning systems might be going through their teenage years and think they know things when they don't completely, but when crucially important systems are based on Deep Learning algorithms without ongoing training and monitoring, they can make decisions that weren't intended. The 'godfather of AI' Geoffrey Hinton, just recently quit Google after a long career in developing the underlying foundations to Deep Learning used by systems like ChatGPT, saying they could be trained for the wrong reasons. Reasons range the spectrum of bad outcomes like spewing hurtful and highly offensive speech, stealing original art from digital artists to existential dangers to humanity like automated killing machines. A key problem is that the reasoning of Deep Learning systems can't be explained in order to prevent negative outcomes.

Imagine this 'black box problem' continues unchecked, with more and more complex systems that increasingly have more impact on human life and safety while increasingly becoming more and more automated without any insight or ability to intervene. We found this thought quite disturbing and before we continued to create AI-based automation systems we decided to put into place, at a structural level, ways to understand and intervene to avoid a future where robots take over, like every movie involving AI fears. We fear this future too, so we created a quotient called the Deane-Fuller Quotient (DFQ) that is a number for an AI Automation system that combines the level of automation with the level of impact on human life. When the quotient gets too high, we say this is Dry AI – meaning it's too automated for the level of impact and it must include additives that will make it more 'wet' and will lower its DFQ into an acceptable range. Additives such as logging, automated notification systems, human overrides and software control points are all ways to make an automated system that is very impactful on human life require closer attention by humans. An example of a highly automated, highly impactful automated system is a remote surgical robot, with a DFQ of 2.0 – where impact score is 1.0 and automation score is 1.0 (on a scale of 0 to 1). By requiring a doctor to always fully focus on the remote surgical robot, its level of automation drops to 0.25 and the overall DFQ is now 1.25 and being closer to 1.0 DFQ, is considered in an acceptable range.

Asimov's Three Laws of Robotics are a moral code of ethics programmed into the foundation of robotic systems based on books he wrote, where robots will not harm humans. DFQ is our way of building into the DNA of AI Automation systems, a way to make sure they do the right thing even when faced with technology used for the wrong reasons. For example, imagine a robot that knows two people are going to use an advanced weapons system to kill a large number of people. DFQ overrides are setup such that the robot will disengage the weapon system based on the set of criteria it is told to consider beforehand during times of conflict, including justification for killing large numbers of people. By integrating a code of ethics that also aligns with our justice system, DFQ integrated into AI Automation systems will serve the needs of the many, over the needs of the few.

AI Automation is on a different branch from Deep Learning. We created Agent-based software as the foundation to an Intelligent System platform, called ADIN™, where we configure agents to work at a higher level than neurons. Every agent has a triggering criteria and action response that is analogous to neurons being stimulated and then firing. ADIN is more like the 'airplane' approach to flying. The goal is to build a system based on a breakdown of complex problems into small problems – a divide and conquer approach. Agents can number in the thousands and work together to solve a large goal, like bees of a hive working together, communicating what needs to be communicated, but ultimately each bee is doing the one thing it needs to, but in coordination with the hive.

Intelligent Systems, like ADIN, adapt to its environment by having agents update parameters of themselves or other agents, where decision points are always logged. When an agent fires and does its associated action, it's logged and therefore not a mystery why. If something goes wrong, the agent can be stopped, or additional agents can be added that automatically inform people of key decision points. This is fundamentally different from Deep Learning neural networks. ADIN has built-in introspection in terms of health agents that monitor how agents are processing.

Agents are small programs that are configured to respond to its particular set of triggering criteria – which is based on data, time and/or location. For example, a new record appears in a database, when a new customer job is created. The associated agent's action can be any number of preset actions from creating a record in another system, notifying an administrator of the new job, updating the status of the associated customer record. Many AI-based Automation systems have been created using ADIN.

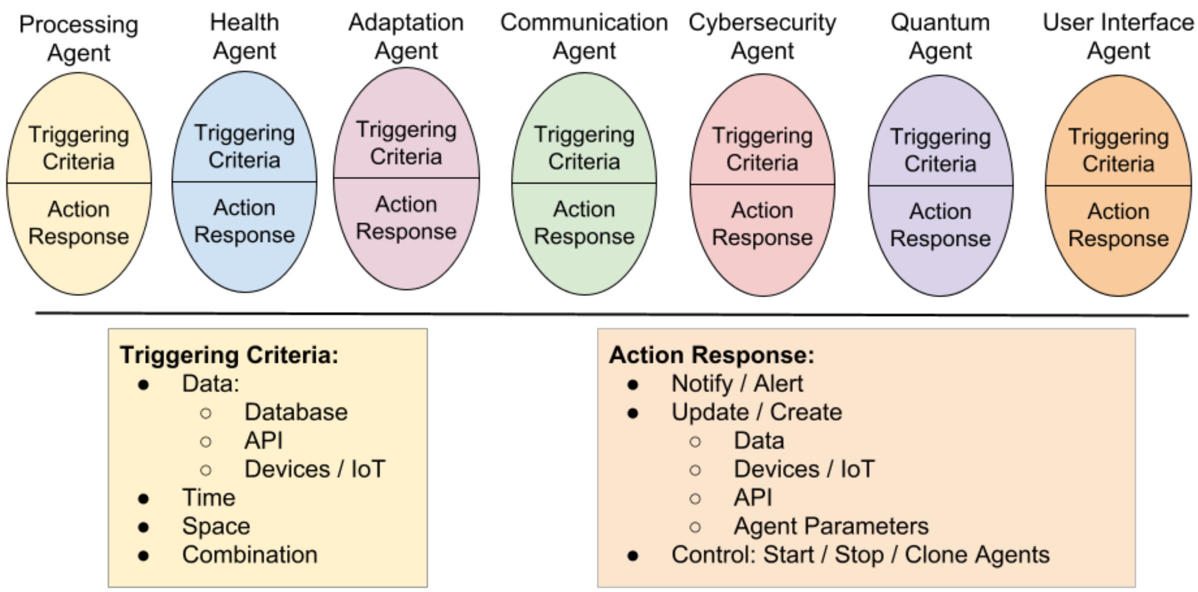

By targeting AI Automation problems using agents, we have connected to a quantum computer that quickly solves difficult problems in optimization where many constraints and preferences are involved. These types of problems are very common for applications that involve scheduling (people, equipment, or both), allocation of resources (shared resources like heavy equipment in construction, or medical devices in a hospital setting), or routing applications (you have 10 crews, 8 trucks, 300 appointments over 5 days and customers give a time slot preference). These examples are all forms of optimization. Optimization problems don't work well on computers because a globally optimized solution (meaning the "best" solution) must consider constraints and preferences, and requires every possible combination be considered which takes a long, long time. Quantum computers are different on the inside and give an optimized global solution very, very quickly. We are using ADIN's AI automation platform combined with quantum computing's ability to optimize, to solve common but difficult problems.

An example of this type of AI Automation is we connected ADIN to thousands of Internet of Things (IoT) devices to gather image and sensor data from camera-enabled insect traps as part of a large web and mobile application. The ADIN Agents continuously monitor for new IoT data as its triggering criteria.

The action response is to access the image data and associated meta-data and run a computer vision algorithm to detect and classify newly trapped insects. Another set of ADIN agents trigger off insect counts exceeding thresholds, that depend on many different things, such as insect type, location, humidity, and temperature levels, and will notify key administrators and pest control technicians that an insect infestation is imminent.

By tackling complex problems by first breaking them down into their simplest components and programming each agent to trigger and respond to a piece of the problem, with every decision being logged, we have created many AI Automation applications that appear to be widely different from each other but in terms of how the AI automation is accomplished are actually quite similar. Agents are configured for their triggering criteria as its setup phase along with its associated action. ADIN agents are highly reusable where triggers can be mixed/matched with actions. Triggers based on time mean they only happen during certain times of day or at preset intervals. Likewise, triggers based on location use GPS coordinates from devices ADIN agents monitor and when they cross into a geo-fenced area or move a certain distance, then their triggers fire.

For automated manufacturing, ADIN agents look at each manufacturing stations' preconditions analyzing sensors, and if all preconditions are met, will perform their associated manufacturing step. All the various stages of the manufacturing process have ADIN agents take on a small part of the larger goal. Using the bees in a hive example, as each bee fulfills their task, it sets the preconditions for the next task. When all ADIN agents are functioning independently, doing their one job (responding with their action to their own triggering criteria), a structured system of many concurrent processes working in harmony is achieved.

Since AI is programming and computers do what we tell them to do, it's important to develop AI-based systems that ensure artificially intelligent systems do what we want them to do, and we always know why. Devoting our time to creating AI Automation systems that take on common mundane problems that are not well solved by connecting to quantum computers, frees people to work more creatively and efficiently. There are many areas that must progress along with our growing dependence on AI. Legal liability laws when death or injury is caused by AI-based technology, widespread AI-based surveillance systems as it applies to privacy protection, invisible bias baked into decision-making systems trained on incomplete datasets, wrangling errant AI, and establishing ethical standards to promote human flourishing are but a few AI-related technology topics that must also evolve.

Tammy Fuller, VP/Chief SW Architect, tammy@echomessaging.com

Gerald Deane, CEO/President, gerald@echomessaging.com

Short Paper Series on AI Automation at www.echomessaging.com/news:

ADIN + IoT for Insects Analytics