



# GBT Technologies, Inc.

Danny Rittman PhD - CTO | Jan 1, 2021

Copyright © 2021 GBT Technologies, Inc.



*“Computers are not intelligent, they just think they are.”*  
– *Computer Symposium, 1979*

## Introduction

The field of Artificial Intelligence (AI) is growing aggressively. It's influencing our lives and societies more than ever before and will continue to do so at a breathtaking pace. Areas of application are diverse, the possibilities far-reaching if not limitless. Thanks to recent improvements in hardware and software, many AI algorithms already surpass the capacities of human experts. Algorithms will soon start optimizing themselves to an ever greater degree and may one day attain superhuman levels of intelligence.

Our species dominates the planet and all other species on it because we have the highest intelligence. Scientist believe that by the end of our century, AI will be to humans what we now are to chimpanzees. Moreover, AI may one day develop phenomenal states such as self-consciousness, subjective preferences, and the capacity for suffering. This will confront us with new challenges both in and beyond the realm of science.

AI ranges from simple search algorithms to machines capable of true thinking. In certain domain-specific areas, AI has reached and even overtaken human abilities. Machines are beating chess grandmasters, quiz show champs, and poker greats. Back in 1994, a self-learning backgammon program found useful strategies no human had considered. There now are algorithms that can independently learn games from scratch, then reach and even surpass human capacities.

It's not just fun and games. Artificial neural networks approach human levels in recognizing handwritten Chinese characters. They vie with human experts in diagnosing cancer and other illnesses. We are getting closer to creating a general intelligence which at least in principle can solve problems of all sorts, and do so independently.

Today's AI is more a form of "cognitive computing," which is true machine learning. Cognitive computing was born from the fusion of cognitive science (the study of the human brain) and computer science. It's based on self-learning systems that use machine-learning techniques to perform specific, human-like tasks in an intelligent way. IBM describes it as "Systems that learn at scale, reason with purpose and interact with humans naturally." According to Big Blue, while cognitive computing shares many attributes with AI, it differs by the complex interplay of disparate components, each of which comprises its own mature disciplines. The sheer volume of data being generated in the world is creating cognitive overload for us. These systems excel at gathering data and making it useful.

So here we are at a new age. But a word of caution is in order. Almost all progress poses risks and that's certainly the case with the bold new age of AI. Some of the problems will be vexing ethical ones. How will AI affect individual lives and whole civilizations around the world? Studies conclude that even though dire scenarios are unlikely, maybe even highly so, the potential for grave damage must be taken seriously. We at GBP certainly do.

With that in mind, let's look at what lies ahead and what role GBT will play in the AI age.

## The term AI

The term “Artificial Intelligence” was coined by John McCarthy in 1956 when he held the first academic conference on the subject. But the journey to understand if machines can ever truly think began before that. Vannevar Bush’s seminal 1945 work *As We May Think* advanced the idea of a system that amplifies human knowledge and understanding. Five years later Alan Turing wrote a paper on machines simulating human beings and being able to do intelligent things, such as play chess.

No one denies a computer’s ability to process logic. But it’s not clear a machine actually thinks. The precise definition of thinking is important and there’s been strong opposition to the very idea of the possibility. For example, there’s the so-called “Chinese room” argument. Imagine someone locked in a room and passed notes in Chinese. Using an entire library of rules and look-up tables, they could produce valid responses in Chinese, but would they really “understand” the language? The argument is that since computers would always be applying rote fact checks, they will never truly “understand” a subject. This objection has been refuted in numerous ways, but it does not undermine confidence in machines and so-called expert systems in life-critical applications.

The idea of inanimate objects becoming intelligent life forms has been with us for a long time. The ancient Greeks had myths about such transformations, and Chinese and Egyptian engineers built automatons. In the previous century, movies familiarized us with intelligent machines. There was the humanoid that impersonated Maria in *Metropolis* (1927), the Tin Man in *The Wizard of Oz* (1939), and Robby the Robot in *Forbidden Planet* (1956). Fascinating characters all. By the 1950s, the AI concept was studied by countless scientists, mathematicians, and philosophers who assimilated it into their thinking, lectures, and books.


One such person was Alan Turing, a young British polymath who explored the mathematical possibility of artificial intelligence. Turing noted that humans use information and reason to solve problems and make decisions, then went on to ask why can’t machines do the same thing? This was the logical framework of his 1950 paper “Computing Machinery and Intelligence,” which we’ll discuss a bit more presently.

## AI: a brief history

The beginnings of modern AI can be traced to classical philosophers’ attempts to describe human thinking as a symbolic system. But the field of AI wasn’t formally founded until 1956, at a conference at Dartmouth College. MIT cognitive scientist Marvin Minsky and others at the conference were buoyant about AI’s future. “Within a generation the problem of creating ‘artificial intelligence’ will substantially be solved,” Minsky concluded.

However, creating an artificially-intelligent entity proved not to be as simple as Minsky thought. After several reports criticizing progress in the effort, interest and government funding dropped off, beginning a stretch from 1974-80 that became known as the “AI winter.” A thaw came in the 1980s when the British government began to fund it again, in part to compete with efforts in Japan. Yes, world politics is part of the AI story.

Backpropagation, sometimes abbreviated as “backprop,” is the single most important algorithm in machine learning. The idea behind it was first proposed in 1969, although it only became a staple in the field in the mid-1980s. Backprop allows a neural network to adjust its hidden layers if output doesn’t match what the creator expected. The creator trains networks to perform better by correcting them. Backprop modifies neural-network connections to make sure a correct answer is arrived at next time.



Backprop Algorithm - Image source: Wikipedia

Another winter set in from 1987 to 1993, coinciding with the collapse of the market for some early general-purpose computers and reduced government funding. But another thaw came in 1997. We’ve already mentioned IBM’s Deep Blue and Watson.



IBM's Deep Blue. Image source: Wikipedia

This year, the talking computer "chatbot" Eugene Goostman captured headlines by tricking judges into thinking it was a flesh-and-blood human during a Turing Test, a procedure developed by the famed Alan Turing in 1950 to assess whether a machine is intelligent. Of course, the results are controversial. AI experts say only a third of the judges were fooled and add that the bot dodged a few questions by claiming it was an adolescent who spoke English as a second language.

Until 2000 or so, most AI systems were limited and based on heuristics. In the new millennium, a novel type of universal AI is gaining momentum. It's mathematically sound, combining theoretical

computer science and probability theory to derive optimal behavior for robots and other embedded agents. And deep learning is driving modern AI applications.

AI has captured the public’s eye in the entertainment realm. In 1997 IBM’s Deep Blue defeated chess grandmaster Garry Kasparov. In 2011 another IBM creation, Watson, beat the two best human players on the language-based quiz show *Jeopardy*. In 2015 the first variant of poker, Heads-Up Limit Hold’Em (HULHE), was theoretically fully solved by Cepheus. In March 2016, Google DeepMind’s AlphaGo AI decisively bested Lee Sedol, the world champion of Go, four games to one. Almost a sweep. The eagerly anticipated contest was watched by 60 million people around the world. The contest was remarkable because the number of allowable board positions was more than the total number of atoms in the universe. This might be AI’s most noteworthy feat to date.



Google DeepMind AI – Image source: Google

## AI themes

Although the game competitions made AI prominent, the main advances of the past 60 years have been out of public sight. Search algorithms, machine-learning algorithms, and integrating statistical analysis into understanding the world have all made great strides, but they have not boldly gone out and explored new worlds. Hence, their relative obscurity.

Advances are nonetheless remarkable. AI examines our purchase histories and influences marketing decisions. Amazon and eBay are using it today. Perhaps you’ve noticed a certain intelligence at play in the emails they send and what’s on their home pages when you log in. Perhaps you’ve noticed them becoming cleverer. Maybe even more helpful?

A common theme in AI is overestimating the difficulty of foundational problems. Some people in the field like to quip that significant AI breakthroughs in the next 10 years have been promised for the past 60 years! Furthermore, there's a tendency to redefine what "intelligent" means after machines have mastered an area or problem. This "AI Effect" contributed to the downfall of US-based AI research back in the 80s.

Expectations seem to always outpace reality. Expert Systems have grown but not become as common as human experts. And while we've built software that can beat humans at some games, open-ended ones remain beyond a computer's capacities. Is the problem simply that we haven't focused enough resources on basic research, as was the case with AI winters? Or is the complexity of AI something we haven't yet truly grasped? Instead, as in the case of computer chess, we may be focusing on specialized problems rather than on grasping what constitutes "understanding."



## The Turing Test


Passing the test Alan Turing gave us is a long-term goal in AI research. Will we ever build a computer that can imitate a human such that a judge cannot tell the difference? Passing it initially looked daunting but possible, at least once hardware technology reached a certain level. However, research found the task far more complicated than thought. Despite decades of research and great technological advances, many wondered if the work was only revealing one vexing problem after the other. Some began to think the goal was unattainable. Others stayed the course.



Turing’s 1950 paper “Computing Machinery and Intelligence” opened the door to what would be called AI. (The actual term was coined years later by John McCarthy.) Turing began his paper with a simple question, “Can machines think?” He went on to propose a method for answering his question. The “Imitation Game,” as he called it in his paper, takes a simple pragmatic approach: if a computer cannot be distinguished from an intelligent human, then this demonstrates that machines can think.

The idea of such a long-term, difficult problem was key to defining the field of AI because it cut to the heart of the matter. It defined an end goal that can drive research down many paths.

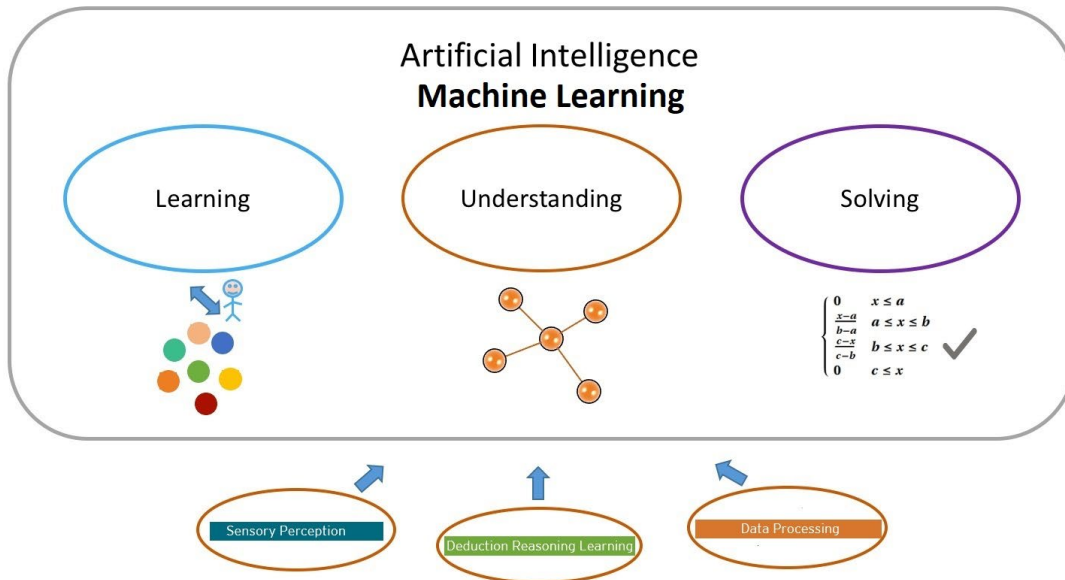
Without a vision of what AI could achieve, the field might never have formed, or it would have remained an obscure branch of mathematics and philosophy. The fact that the Turing Test is still discussed, and that researchers attempt to produce software capable of passing it, indicate that Turing gave us a useful idea. Others, however, feel that given our technology and approaches, the Test is not as important as once thought. Nonetheless, research goes on and strides are made.



## Cognitive computing

Cognitive computing is based on self-learning systems that use machine-learning techniques to perform specific, human-like tasks in an intelligent way. It’s already augmenting and accelerating

## Cognitive Computing



human capabilities by mimicking how we learn, think, and adapt. Cognitive computing will one day replicate human sensory perception, deduction, learning, thinking, and decision-making. Harnessing vast amounts of computing power will take machines beyond simply replicating and into distinguishing patterns and providing solutions. The jump in machine capacity will augment human potentials. It will inspire and increase individual creativity and create waves of innovation. The next figure identifies the three converging areas of cognitive computing capabilities:

*Sensory perception* – simulating human senses of sight, hearing, smell, touch, and taste. The most developed in terms of machine simulation are visual and auditory perception.

*Deduction* – simulating reasoning, learning, and decision making. Machine learning, deep learning, and neural networks are the most prominent of these technologies. They are already deployed as systems of intelligence to derive meaning from information and apply judgments.

*Data processing* – processing huge data sets to provide insights and accelerate decision making. Hyper-scale computing, knowledge representation and ontologies, high-performance searches, and Natural Language Processing (NLP) are leading technologies here. They provide the processing power to ensure systems work in real time.

Cognitive computing is not a single technology. It makes use of multiple technologies that allow it to infer, predict, understand, and make sense of information. These technologies include AI and machine-learning algorithms that train the system to recognize images, understand speech, recognize patterns, and, via repetition and training, produce more accurate results. Through the use of NLP systems based on semantic technology, cognitive systems can understand the meaning and context of words. This allows a deeper, more intuitive level of discovery – and even interaction.

## Features of a cognitive system

In order to implement cognitive computing into commercial and individual use, a cognitive computing system must be:

*Adaptive* – The system must be able to adapt, as a brain does, to any surrounding. It needs to be dynamic in gathering data and understanding goals and requirements.

*Interactive* – The system must be able to work easily with users so they can define needs readily. Similarly, it must also work with other processors, devices, and cloud services.


*Iterative and “stateful”* – The system must carefully apply data quality and validation methodologies to ensure the system has sufficient information and its data sources deliver reliable, up-to-date input.

*Contextual* – The system must be able to understand, identify, and extract contextual elements such as meaning, syntax, time, location, appropriate domain, regulations, user’s profile, process, task, and goal. It must draw on multiple sources of information, including both structured and unstructured digital information.

In regard to deduction, reasoning, and learning systems, technologies such as machine learning, deep learning, and neural networks are enabling machines to simulate and augment human thinking at higher and higher levels. Let’s look at them.

## Machine learning

Machine-learning capabilities and techniques help applications identify patterns in large amounts of information, classify the information, make predictions, and detect anomalies. This helps organizations build systems to process vast amounts of data, while applying human-like thinking. They then classify and correlate disparate pieces of information, make more informed decisions, and trigger actions downstream. Such systems are also capable of learning over time, without needing to be explicitly programmed to do so.




## Deep learning

A branch of machine learning, deep learning attempts to model high-level abstractions. Deep-learning methods are based on assimilating representations of data. They apply multiple layers of processing units, with each successive layer using the output of the previous layer as input. Each layer corresponds to different layers of abstraction, forming a hierarchy of concepts. Deep-learning solutions help create applications that can be trained in both supervised and unsupervised manners. (More on this soon.) Numerous techniques, such as neural networks, are under development in the deep-learning world. They'll soon be part of our daily lives.



## Neural networks

Neural networks are among the most used deep-learning methods to create learning and reasoning systems. The application of backprop to neural networks helps generate continuous improvements in deep-learning algorithms and also train multi-layered deep-learning systems to enhance their knowledge bases. Neural networks will analyze audio, video, and images as capably as any human expert, but at exponentially faster speeds, and at much greater levels of quantity and depth.



## Cognitive systems

There are many examples of remarkable progress in cognitive computing. The accuracy of Google's voice recognition technology, for instance, improved from 84% in 2012 to 98% two years later. Facebook's DeepFace technology now recognizes faces with 97% accuracy. From 2011–2014, over \$2 billion in venture capital went into firms developing cognitive-computing products and services.

The cognitive-computing landscape is presently dominated by large players such as IBM, Microsoft, and Google. Big Blue, the pioneer in the field, has invested \$26 billion in big data and analytics. It now spends close to one-third of its formidable R&D budget in cognitive computing.

Many new companies are investing heavily in machine learning and AI to better their products. Some of the key players in this sector are:

*IBM Watson* – The famed question-answering supercomputer combines AI and analytical software for optimal performance. It leverages deep-content analysis and evidence-based reasoning. Combined with massive, probabilistic processing techniques, Watson can improve decision making, reduce costs, and optimize outcomes.



IBM Watson. Image source: Wikipedia

*Microsoft Cognitive Services* – This set of APIs, SDKs, and cognitive services allows developers to make their applications more intelligent. Developers can easily add intelligent features (e.g., emotion and sentiment detection, vision and speech recognition, knowledge, search and language understanding) into applications.




Image source: Microsoft

*Google DeepMind* – In the past few years, Google has acquired many AI and machine-learning-related startups and rivals, moving the field towards consolidation. The company has created a neural network that learns to play video games just as we do.






Image source: Google

We at GBT are making a mark too, learning from available research and going on from there. Our chatbot uses the Microsoft Bot Framework to improve our engineering team designs. Then there's Avant! which we'll say a lot more about soon.

Cognitive computing is transforming how we live, work, and think. Little wonder cognitive computing is, and will continue to be, an amazingly dynamic field.

Cognitive computing is a powerful tool, but only a tool. The humans with the tool must decide how best to use it. And that's up to us.

## Introducing GBT's thinking entity – Avant!



GBT is proud to announce Avant! – our exciting and innovative contribution to cognitive computing. It provides answers and insights for a wide range of information. It can put a vast volume of information into context, derive value from it, and increase human expertise and capabilities. Avant! actually reasons.

Avant! can understand unstructured data, which is how most data exists today. Furthermore, most data comes from a wide range of sources – professional articles, research papers, blogs, and

human's input. Avant! relies on natural language, and obeys the rules of grammar. Avant! breaks down a sentence grammatically and structurally, then extracts its meaning. When Avant! works on a field, it searches for thousands of articles in this domain. Next, Avant! is narrowing down to only few hundreds documents that are topic related. Finally, it will conclude the best answer and deliver it. The process is governed by sets of neural network algorithms that work cognitively to learn the topic and respond accordingly.

Avant! is trained to respond to questions about highly-complex situations and quickly provide a range of responses and recommendations, all backed by evidence. Avant! is using statistical modeling and scores a viable solution then estimates assurance. Avant! improves its expertise by learning from its own experience. Over time, it gains robust knowledge learning from experience and from its own successes and failures, exactly as humans do. Avant! gets wiser and knowledgeable over time.

When a query is executed Avant!, cognitive computing relies on its own vast knowledge, the available human data, and the query conditions. A huge amount of data is searched, both structured and unstructured. Then an analysis is performed. Then elimination process takes place to sort out the best answer, logical validation and finally answer's delivery.

For the user Avant! provides almost instantaneous response. It rapidly processing large data and by using its cognitive computing capabilities provides efficient answer within seconds.

## First: Avant! “eyes” - semantic-segmentation system

Avant! is equipped with a semantic-segmentation engine that enables it to see and identify objects. The engine enables understanding an image at the pixel level. A computer image is taken and then “understood.” For example, the next image was pixelated, analyzed, and understood.



Image source: [https://sthalles.github.io/deep\\_segmentation\\_network/](https://sthalles.github.io/deep_segmentation_network/)

The more advanced the semantic-segmentation system, the better the object-identification ability in the AI system. For example, in the next images, the system has identified objects according to their nature and category:





Image source: <https://blog.goodaudience.com/using-convolutional-neural-networks-for-image-segmentation-a-quick-intro-75bd68779225>

Here's another example:





Semantic segmentation assigns per-pixel predictions of object categories. This provides a comprehensive description that includes information on object category, location, and shape. State-of-the-art semantic segmentation is typically based on a Fully Convolutional Network (FCN) framework. Deep Convolutional Neural Networks (CNNs) benefit from the rich information of object categories and scene semantics learned from diverse sets of images.


Avant! features a different approach to semantic segmentation. Because classic vision approaches have the advantage of capturing a scene's semantic context, Avant! sees the image context as a visual vocabulary. Hand-engineered features of an image are densely extracted by our NativeEXT algorithm that uses its own encoder to produce a unique dataset of the image. The representation encodes global contextual information by capturing feature statistics. The image feature is greatly improved by our CNN-based methods. The encoding process is a powerful break from existing methods.

The Avant! Semantic Segmentation Approach is CNN-based and enables highly accurate computer vision that's implemented within our GBT technology. The system is a core IP within Avant! AI that enables accurate seeing and recognition. This proprietary system is one of Avant! AI's remarkable innovations.

## Avant! – artificial neural networks

One of the many remarkable aspects of Avanti! is its Artificial Neural Network (ANN). An ANN, as we've seen, is an information-processing paradigm inspired by the way nervous systems, such as our brains, process information. The key element of this paradigm is the novel structure of the information-processing system. It's composed of a large number of highly-interconnected processing elements (neurons) working in harmony to solve specific problems.

ANNs, like humans, learn by example. An ANN is configured for a specific application (e.g., pattern recognition or data classification) through a learning process. Learning in biological systems involves adjustments to the synaptic connections between neurons.



ANN – Source: engineering.com

Each neuron in a neural structure is connected with varying strength to other neurons. Based on inputs from the output of other neurons, and on connection strength, an output is generated that can be reused as input by other neurons.

This biological process has been translated into an artificial neural network by using “weights” to indicate the strength of the connection between neurons. Furthermore, each neuron takes the output from connected neurons as input, then uses a math function to determine its output. The output is then used by other neurons.

Neural networks, with their remarkable ability to derive meaning from complicated or imprecise data, can be used to extract patterns and detect trends that are just too complex for humans or most computers. A trained neural network is an "expert" in its assigned field. We also call this an Expert System. Our expert can then provide projections and answer "what if" questions. Avant! AI includes ANN as a base for its deep-learning technology.

The advantages:

- *Adaptive learning* – Avant! ANN learns to perform tasks based on data given from training or initial experience.
- *Self-organization* – Avant! ANN creates its own organization or representation of the information during learning time.
- *Real Time Operation* – Avant! ANN computations are carried out in parallel processing.
- *Cognitive learning* – Avant! AI ANN is a cognitive system that learns from its own experience and that of others too.

Avant! ANN takes a different approach to problem solving. Its part of a network comprising a large number of processing elements that work in parallel to solve problems. Avant! ANN learns by experience or by instruction. It's not programmed simply to perform a specific task.

Avant! ANN and its mathematical algorithms complement each other. Tasks are evaluated for logic/arithmetic operations and for neural networks. Furthermore, a large number of tasks require systems using combinations of both in order to perform at maximum efficiency and at performance levels that Avant! AI requires.

A key feature of Avant! ANN is an iterative learning process in which data cases are presented to the network, many at a time (parallelism). The weights associated with input values are adjusted each time. After all cases are processed in parallel, a conclusion is drawn. Typically, an ANN is trained one case at a time but Avant! processes several cases in parallel. It uses its own expert system to split the data among machine processes. Astonishingly quick responses are achieved.

## Avant! AI apps extend what cars see and recognize


Our AI goes well beyond sensor-based radar, LiDAR, and camera technologies. Avant! offers true, safe, autonomous operations through a shared knowledge bank. Machines (vehicles, drones, robots) communicate among themselves over a multi-agent, secure, real-time protocol. Avant! AI controls vehicle-to-everything (V2X) communications for optimal safety and efficiency.

Avant! semantic-segmentation technology enables vehicles to “see” more than what’s nearby. It peers through dense urban environments and even looks many miles all around. The shared databank, connected through our gNET, private, secure protocol, enables autonomous vehicles to communicate not only with cars but with everything that’s IoT-based. Cars talk to other cars, motorcycles, emergency vehicles, traffic lights, digital road signs, and pedestrians, even if they are not directly within the car’s line-of-sight.

Our constantly growing, cloud-based databank holds information on thousands of square miles of surface area, and we make it available to all network members. The Avant! AI engine supervises V2X communication to ensure real-time responses and maximum safety. And safety, after all, must always be paramount.

## GBT's microchip does it all

Avant! AI is driven by the GopherInsight™ microchip. Our chip enhances capabilities, enabling features never seen before.



Our chip contains circuitry specially designed for important emerging technologies. The advanced manufacturing node of our integrated circuits (10 nm) enables both low-power consumption and a range of onboard sub-circuits. It handles a wealth of complex tasks.

1. *Microprocessor* – This is the brain and controller of the entire system. Our GopherInsight™ chip works in parallel to deliver astonishingly powerful processing.
2. *Proprietary tracking system* – One of the most advanced units in our chip is its tracking system – essential for autonomous driving. This enables global tracking, with or without GPS services.
3. *Onboard flash memory* – Robust memory is needed for internal operations and data storage. GopherInsight™ has it.
4. *Power management* – This unit ensures efficient power consumption in all internal sub-units and also in external chips controlled by the main chip. Only necessary units receive power, others are allowed to sleep.
5. *Arithmetic logic unit (ALU)* – Our onboard mathematician performs complex calculations and problem-solving tasks. Highly advanced, our ALU uses integral high- and low-level software code.

6. *Machine-learning unit* – In conjunction with integral software, this unit collects data, stores it, and makes conclusions. It's the cognitive part of the system that performs the learn and play feature.
7. *I/O unit* – This is in charge of external input/output sub-systems.
8. *Monitoring unit* – It supervises all activities and warns of boundary cases. We also call it an “error correction unit,” as it corrects unexpected conditions from internal or external sources.
9. *Audio/video unit* – This provides audio and video support with ground control or any connected devices.
10. *Registers, PORTS, EPROM, misc.* – These subunits assist with overall system operation.



## Avant! - machine-learning wisdom

Avant! AI is a machine-learning system – essential to intelligent machines. Deep understanding helps make optimized, efficient decisions. Instead of building heavy machines with explicit programming, a new way of doing things is here. Sophisticated algorithms let machines understand the virtual environment and make decisions on their own. This will in time decrease the number of programming concepts. Designing these algorithms, and using them in the most appropriate ways, is one of the great challenges in the field.


Pattern recognition is key to machine learning. Most algorithms use pattern recognition to make optimized decisions. Consequently, we are heading into a new era in statistical and functional approximation techniques.

Our proprietary gNETCar technology, controlled by Avant! AI, enables autonomous cars to learn and adapt. Machine-learning algorithms often, though not always, look to nature for analogy and inspiration. Fundamentally and scientifically, these algorithms depend on data structures and theories of cognitive and genetic structures. We include a heuristic approach as a natural learning



procedure, with additional supervisory algorithms to ensure privacy and safety. This makes our gNETCar highly intelligent. It learns and adapts as it goes.

As with other machine-learning systems of today, Avant! algorithms are based on cognitive thinking and guided by Artificial Neural Networks (ANN). Avant! learning is measured in terms of experience and improved performance.



**Avant! “sees” the road**


Let’s look at some machine-learning algorithms.

## Avant! techniques - pattern recognition

For operations requiring sensing, identification, and decision making, Avant! AI is right there. For example, our gNETCar technology is supervised by the Avant! machine-learning system, which is based on “*pattern recognition* algorithms.” They provide a reasonable answer to all possible inputs and perform “closest to” matching of inputs, taking into account statistical variations.

This is different from “*pattern-matching* algorithms” which pair up exact values and dimensions. AI algorithms have well-defined values for mathematical models and for shapes such as rectangles, squares, and circles. It becomes difficult for machines to process those inputs with different values.

Consider a ball. The shape and pattern can be recognized, but various inflation levels change the shape. That’s a major problem with most machine-learning processes and algorithms. It can bring the entire process to a halt. That’s why GBT adds a supervisory algorithm. Our gNETCar must quickly identify objects along a route, then learn, analyze, and respond in real time. Efficient pattern recognition is absolutely essential.



Pattern recognition must be advanced in autonomous cars

## Supervised vs unsupervised learning

Machine-learning algorithms are either “supervised” or “unsupervised.” The distinction is drawn from how the learner classifies data.

In supervised algorithms, classes are a finite set, determined by a human. The machine learner’s task is to search for patterns then construct mathematical models. These are then evaluated for predictive capacity in relation to measures of variance in the data. Examples of supervised learning techniques include Naive Bayes and decision-tree induction.

Unsupervised learners, by contrast, are not provided with classifications. In fact, the task is to *develop* classification labels. Unsupervised algorithms seek out similarity between pieces of data and determine whether they form a group, or “cluster.”

## Our supervised learning and logical validation

Avant! supervised learning is implemented within GBT’s gNETCar system. The algorithm generalizes to correctly respond to all possible inputs, then produces correct output for inputs not encountered in training.

Avant! AI constantly uses logical cross-validation to analyze and evaluate conditions, analysis and conclusions. In so doing, it learns from its experience. Conclusions are made and actions taken. In autonomous driving, for example, Avant! relies on inputs from the tracking system, environmental sensors, camera input, radar, LiDAR, audio, and more. Information is shared with all network members via gNET, our private, secure communication protocol. All network members have access to the huge and ever growing database of surroundings, weather, and road conditions.

## Natural Language Processing (NLP)

Avant! AI includes Natural Language Processing (NLP). This handles textual/speech recognition and language generation, each requiring different complex techniques. Avant! uses a few methods to identify speech and textual content. Avant! is equipped with advanced NLP module in order to identify and understand speech and textual content.

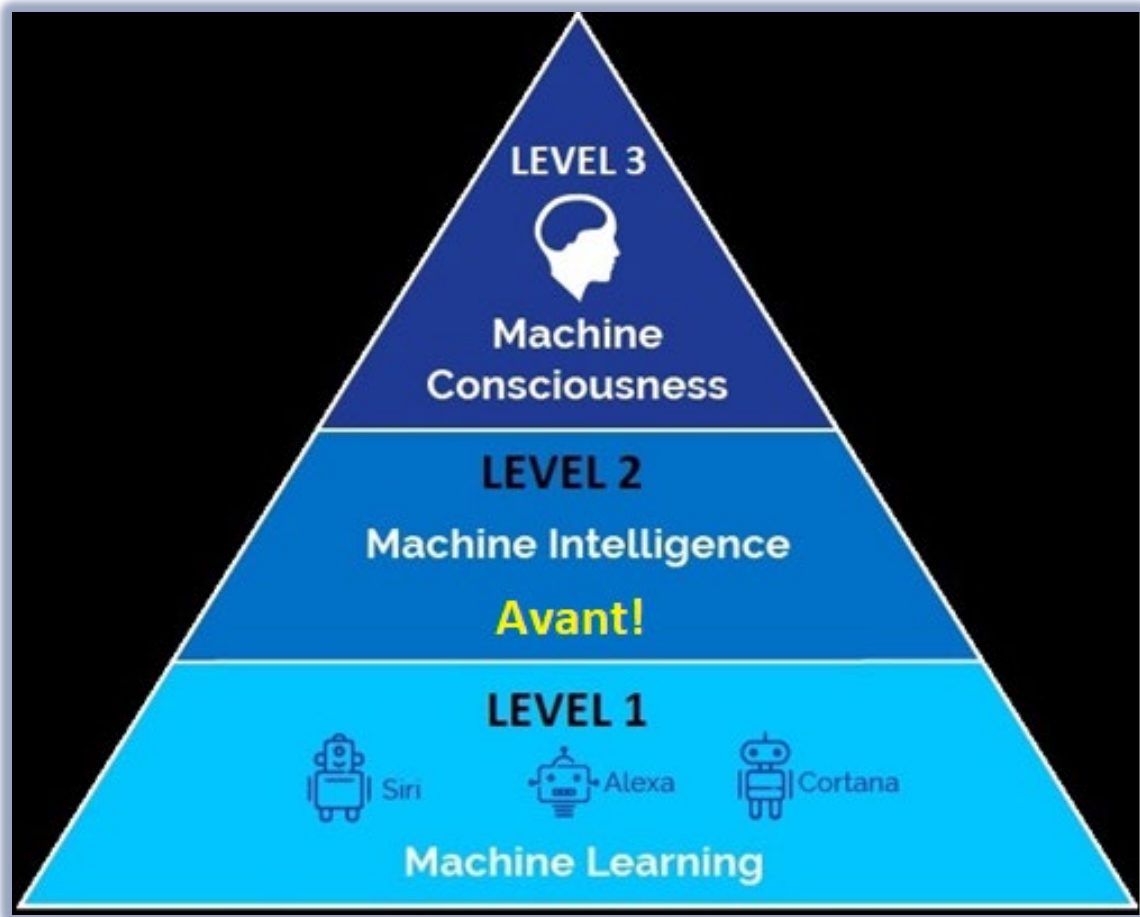
## Avant! and the AI hierarchy

AI systems are categorized by capabilities and features. There are three levels:

Level 1 – *Machine learning*: a set of algorithms used by intelligent systems to learn from experience. For example, Siri, Alexa, MS Cortana.

Level 2 – *Machine intelligence*: a set of algorithms used by machines to learn from experience. For example, Watson, Google DeepMind. Avant! is at this level – for now.

Level 3 – *Machine Consciousness*: self-learning from experience, without the need for external data.



## Avant! AI is embedded within GBT Technology

Avant! AI is central to GBT technology. For example, each gNETCar is connected with all other gNETCars through our mesh system controlled by Avant! AI. Another example is our Guardian Network agent. All Guardian network members communicate via a private, secure communication protocol called gNET whose protocol is managed and controlled by Avant! AI. To ensure confidentiality and privacy, every Guardian Tracking vehicle communicates with peers over a private network, creating enormous computing power and an immense database accessible across the globe.

Once a Guardian device sends data, the closest Guardian device receives it and passes it on to others. Information is shared in real time. Millions of units in vehicles around the world provide a truly immense and constantly growing databank. This promises fast responses and safe driving.

Avant! AI is currently in LEVEL 2 in the AI hierarchy, right up there with IBM Watson and Google DeepMind. Within two years we will bring Avant! to LEVEL 3 – Machine Consciousness. Watch for it!



## Avant! includes a self-testing redundancy and auto-correction system

Because our Avant! AI system is responsible for so many crucial functions such as autonomous machines, drones, tracking, and the gNET communication protocol, we've implemented an integrated redundancy and auto-correction system. One of the most advanced features of Avant! is its self-check scan that constantly searches for problems in hardware, network, and software. If one crops up, the user is alerted to problems via the control board.

If a problem *looms*, Avant! provides, via the mobile app and software, an estimated time until malfunction and also corrective methods. If a problem is *imminent*, the system will first use alternative redundant units to ensure continued safe operation. These sub-units automatically switch on and the vehicle continues normal operation without interruption.

## Avant! AI Cyber Security – gEYE system



The UN forecasts that by 2050 the world will have nearly 10 billion people, with perhaps three-quarters of them residing in cities. If everyone drove manually-controlled, unconnected vehicles, congestion would be worse than it is today. Far worse. Our lifestyles would suffer tremendously. New transportation methods are needed. Fortunately, technologies in development today address the problem.

Electrification, connectivity, automation, and cloud services have been evolving somewhat independently over the years. More recently, however, these technologies are converging to bring on-demand autonomous vehicles. When widely adopted, congestion will be much less fearsome, resources will be used more efficiently, and travel will be safer.

Autonomous vehicles will bring security risks though. Whether they are abused by criminals, vandals, or terrorist rings, the same technologies that promise so many benefits can also sow chaos and destruction. Manufacturers must try to minimize the possibility of attacks on the mobility ecosystem and to mitigate the impact of any attacks that do occur. Security solutions must be implemented quickly, reliably, and without harming performance. Protecting vehicles, occupants, and bystanders will require holistic approaches to design, implementation, and response.

The automobile today is an isolated mechanical device. It can be tampered with, but only if someone has direct physical access. Even then, only one vehicle can be modified at a time, so attacks are narrow. However, we are entering an era in which technologically-savvy attackers can target millions of vehicles – simultaneously and remotely over a wireless network.

If attackers were to discover a vulnerability and strike, results could vary. A ransomware attack would be the least destructive. Malware could disable anywhere from a handful of vehicles to millions of them, releasing them only after a sum of money is paid out. This would be costly and disruptive, but might not endanger lives. In the worst-case scenario, an attack could manipulate control systems, causing crashes and numerous injuries – maybe even catastrophic loss of life.

Access to vehicle control systems could undermine the very safety the technology was designed to provide. There is no internet-connected system that can build a wall high enough or wide enough to thwart a sophisticated criminal group or state. Fortunately, no one has yet demonstrated the possibility of such an attack, let alone actually executed one. Nonetheless, it's imperative that automotive OEMs, suppliers, and regulators do everything possible to make vehicles as secure as possible.

We at GBT are on the case. And of course we've built a solution into Avant! AI.

## Our security system - gEYE

Our gEYE is a multi-layered security solution that is embedded within Avant!. gEYE is a comprehensive cyber and network security system that is constantly monitoring all modules on multiple hardware and software levels for data breach or malicious activities.



## The future of AI


In the last few years, the field has made astonishing breakthroughs especially in the field of machine learning and neural networks. AI systems of today are performing complex functions like

facial/speech recognition, huge data analysis and even medical diagnostics and treatments. We predict that AI will become more dominant in our daily lives especially in the next domains:

1. *Autonomous machines* – The technology isn't perfect yet and it will take a while until autonomous cars will become large scale operational on a daily basis. The main problem is how to keep our existing infrastructure for the use of robotic cars. Also the technology is not there yet. Plenty of work is still ahead when it comes to large data handling, in real time
2. *Robotics* – Robots are on high development pace for the past decade and will continue to be a major subject for AI implementation. We are bound to witness robotics technology to significantly enhance our lives in the fields of medicine, space exploration, transportation, delivery services, entertainment, and personal assistance.
3. *Security/Military* – Robots are already playing major role within security and military applications and will perform further complex tasks during the next decade. Drones and robots assist national defense. They patrol frontiers, battle lines, and insurgent regions efficiently, effectively, and without risking lives. Drones already deliver lethal force and the day may come when AI takes humans out of the operation. Some warn of AI turning on its creators, as in the old Mary Shelley tale. Lawmakers and publics will have to grapple with this someday.
5. *Medicine* – IBM Watson is already examining symptoms, performing diagnostics, and recommending treatments. Many medical procedures are already done by robots, though they're directed by humans. That may change one day – through AI.
6. *Personal/Entertainment* – The age of robots manifesting human qualities, including emotions, is dawning. They already perform some basic operations in the manner of a personal assistant, friend, or relative.

SoftBank Robotics, a Japanese firm, has made a robot companion – one who can understand and feel emotions. “Pepper” reads human feelings, develops its own feelings, and helps its human companion stay happy. Pepper went on sale in 2015 and all 1,000 units sold out in one minute. It goes on sale in the US this year and more sophisticated friends are sure to follow soon. Isaac Asimov's vision of tomorrow isn't so far off.





## Conclusions

GBT Protocol's Avant! AI is a remarkable system developed especially for but not limited to managing and controlling our technology. Avant! is a new generation of AI, a highly sophisticated one. It can detect, analyze, and learn from experience. The system is cognitive-based and handles huge amounts of data in real time. It's ideally suited for autonomous machines, military/security drones, GEO tracking/location expert systems, and many other uses.

Avant! AI manages and controls our private, secure communication protocol, gNET, to the great benefit of network users. Finally, its constantly-growing "wisdom" can be used for general purposes as well.

Within the next few months the public will be able to take a good look at Avant! through our web-based interface. We hope you will find it interesting, intriguing, and even fun. And of course we welcome comments and suggestions. Avant! is an ongoing project. We will expand its features and capabilities over time with the ultimate goal of AI's highest aspiration – machine consciousness.

We have long-term plans for our jewel. We intend to run Avant!-managed applications on smartphones via its deep neural networks to enable exciting new capabilities. We'll also develop Avant! to control personal assistant robots who will link vision, language, and speech in such a way that the difference between humans and robots will diminish.

Personal robots will get smarter and more humanlike. We intend to develop a “nanny robot” to watch over children with dedication and devotion. Other models will learn our daily routines and habits, then take care of chores like housework and cooking.

Avant! will revolutionize engineering by designing new circuits and technologies. Firms will be able to define a product’s features and characteristics and Avant! will provide an architecture and low-level designs, including simulations.

Avant! AI will bring important changes to medicine. It will diagnose and recommend treatments. Based on accumulated experience, Avant! will become a valuable assistant to nurses and doctors – probably even an indispensable one.

These changes are not off on a distant horizon. They’re far more than the fanciful dreams of sci-fi enthusiasts. The world will undoubtedly adopt more and more AI in the next decade – with greater expansion in ensuing years. Avant! AI will join the new age within the next year.

On behalf of the GBT team, we welcome you to the new age!



## References

1. L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam. "Rethinking atrous convolution for semantic image segmentation."
2. M. Cimpoi, S. Maji, and A. Vedaldi. "Deep filter banks for texture recognition and segmentation." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2015) pp. 3828-3836.
3. G. Csurka, C. Dance, L. Fan, J. Willamowski, and C. Bray. "Visual categorization with bags of keypoints." *Workshop on statistical learning in computer vision, ECCV*, vol. 1, pp. 1-2.
4. J. Dai, K. He, and J. Sun. "Boxsup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation." *Proceedings of the IEEE International Conference on Computer Vision* (2015) pp. 1635-1643.
5. J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. "ImageNet: A Large-Scale Hierarchical Image Database," in *CVPR* (2009).
6. V. Dumoulin, J. Shlens, M. Kudlur, A. Behboodi, F. Lemic, A. Wolsiz, M. Molinaro, C. Hirche, M. Hayashi, E. Bagan, et al. "A learned representation for artistic style."
7. M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. "The Pascal visual object classes (voc) challenge." *International Journal of Computer Vision* (2010) pp. 303-338.
8. L. Fei-Fei and P. Perona. "A Bayesian hierarchical model for learning natural scene categories." *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 2, pp. 524-531. Shelley Row (2013), "The Future of Transportation: Connected Vehicles to Driverless Vehicles...What Does It Mean To Me?" *ITE Journal* ([www.ite.org](http://www.ite.org)), vol. 83, No. 10, pp. 24-25.
9. "Levels of Driving Automation Are Defined In New SAE International Standard." Society of Automotive Engineers ([www.sae.org](http://www.sae.org)).
10. "Shared Mobility Principles for Livable Cities." ([www.sharedmobilityprinciples.org](http://www.sharedmobilityprinciples.org)).
11. Brand, M. "Charting a manifold," in Becker, S., Thrun, S., and Obermayer, K. (eds.), *Advances in Neural Information Processing Systems* MIT, 2003, pp. 961-968.
12. Breiman, L., Friedman, J. H., Olshen, R. A., and Stone, C. J. *Classification and Regression Trees*. Wadsworth International Group, Belmont, CA, 1984.
13. Breiman, L. (2001). "Random forests." *Machine Learning* 45 (1), pp. 5-32.

14. Brown, L. D. “Fundamentals of Statistical Exponential Families.” Vol. 9 (1986) *Inst. of Math. Statist. Lecture Notes Monograph Series*.
15. Candes, E., and Tao, T. “Decoding by linear programming.” *IEEE Transactions on Information Theory*, 51 (2012) pp. 4203-4215.
16. Carreira-Perpinan, M. A., and Hinton, G. E. “On contrastive divergence learning.” In Cowell, R. G., and Ghahramani, Z. (eds.), *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics* (2005) pp. 33-40.
17. Caruana, R. “Multitask connectionist learning,” in *Proceedings of the 1993 Connectionist Models Summer School*, (2003) pp. 372-379.
18. Clifford, P. “Markov random fields in statistics,” in Grimmett, G., and Welsh, D. (eds.), *Disorder in Physical Systems: A Volume in Honour of John M. Hammersley*, Oxford University Press, 1990, pp. 19-32.
19. Cohn, D., Ghahramani, Z., and Jordan, M. I. “Active learning with statistical models,” in Tesauro, G., Touretzky, D., and Leen, T. (eds.), *Advances in Neural Information Processing Systems 7*, Cambridge, MA: MIT Press, 1994, pp. 705-712.
20. Coleman, T. F., and Wu, Z. “Parallel continuation-based global optimization for molecular conformation and protein folding,” Technical Report, Cornell University, Dept. of Computer Science, 1994.
21. Collobert, R., and Bengio, S. “Links between perceptrons, MLPs and SVMs,” in Brodley, C. E. (ed.), *Proceedings of the Twenty-first International Conference on Machine Learning* New York, NY, 2004.
22. Collobert, R., and Weston, J. “A unified architecture for natural language processing: Deep neural networks with multitask learning,” in Cohen, W. W., McCallum, A., and Roweis, S. T. (eds.), *Proceedings of the Twenty-fifth International Conference on Machine Learning* 2008, pp. 160-167.
23. Cortes, C., Haffner, P., and Mohri, M. “Rational kernels: Theory and algorithms.” *Journal of Machine Learning Research*, 5 (2004) pp. 1035-1062.
24. Cortes, C., & Vapnik, V. “Support vector networks.” *Machine Learning*, 20 (1995), pp. 273-297.
25. Cristianini, N., Shawe-Taylor, J., Elisseeff, A., and Kandola, J. “On kernel-target alignment,” in Dietterich, T., Becker, S., and Ghahramani, Z. (eds.), *Advances in Neural Information Processing Systems 14* (NIPS’01), vol. 14, 2002, pp. 367-373.
26. Cucker, F., and Grigoriev, D. “Complexity lower bounds for approximation algebraic computation trees.” *Journal of Complexity*, 15 (4), (1999) pp. 499-512.
27. Dayan, P., Hinton, G. E., Neal, R., and Zemel, R. “The Helmholtz machine.” *Neural Computation*, 7 (1995) pp. 889-904.

28. Japkowicz, N., Hanson, S. J., and Gluck, M. A. "Nonlinear auto-association is not equivalent to PCA." *Neural Computation*, 12 (3), 2000, pp. 531-545.
29. Jordan, M. I. *Learning in Graphical Models*. Dordrecht, Netherlands, Kluwer, 1998.
30. Kavukcuoglu, K., Ranzato, M., and LeCun, Y. "Fast inference in sparse coding algorithms with applications to object recognition. Technical Report, Computational and Biological Learning Lab, Courant Institute, NYU, 2008.
31. Kirkpatrick, S., Jr., C. D. G., and Vecchi, M. P. (1983). "Optimization by simulated annealing." *Science*, 220, pp. 671-680.
32. Köster, U., and Hyvärinen, A. (2007). "A two-layer ICA-like model estimated by Score Matching," in *Int. Conf. Artificial Neural Networks*, 2007, pp. 798-807.
33. Krueger, K. A., and Dayan. "Flexible shaping: how learning in small steps helps." *Cognition*, 110, 2009, pp. 380-394.
34. Lanckriet, G., Cristianini, N., Bartlett, P., El Gahoui, L., and Jordan, M. "Learning the kernel matrix with semi-definite programming," in Sammut, C., and Hoffmann, A. G. (eds.), *Proceedings of the Nineteenth International Conference on Machine Learning (ICML'02)*, 2009, pp. 323-330.
35. Larochelle, H., and Bengio, Y. "Classification using discriminative restricted Boltzmann machines," in Cohen, W. W., McCallum, A., and Roweis, S. T. (eds.), *Proceedings of the Twenty-fifth International Conference on Machine Learning (ICML'08)*, 2008, pp. 536-543.
36. Larochelle, H., Bengio, Y., Louradour, J., and Lamblin, P. "Exploring strategies for training deep neural networks." *Journal of Machine Learning Research*, 10, 2009, pp. 1-40.
37. Larochelle, H., Erhan, D., Courville, A., Bergstra, J., and Bengio, Y. "An empirical evaluation of deep architectures on problems with many factors of variation," in Ghahramani, Z. (ed.), *Proceedings of the Twenty-fourth International Conference on Machine Learning 2007*, pp. 473-480. ACM.
38. Lasserre, J. A., Bishop, C. M., and Minka, T. P. "Principled hybrids of generative and discriminative models," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2006, pp. 87-94.
39. Le Cun, Y., Bottou, L., Bengio, Y., and Haffner, P. "Gradient-based learning applied to document recognition." *Proceedings of the IEEE*, 86(11), 1998, pp. 2278-2324.
40. Intrator, N., and Edelman, S. "How to make a low-dimensional representation suitable for diverse tasks." *Connection Science, Special issue on Transfer in Neural Networks*, 8, 1996, pp. 205-224.
41. Quote in Crevier, D. *AI: The Tumultuous Search for Artificial Intelligence*, NY: Basic Books, 199e).