


COMPUTER VISION APPLICATIONS WITH POWERSCALE



Sneha Venkatesh

Associate Sales Engineer Analyst
Dell Technologies
Sneha_venkatesh@dell.com

Parikshit Patil

Associate Sales Engineer Analyst
Dell Technologies
Parikshit_p@dell.com

Sai Ganesh

Specialist 2, Inside Product
Dell Technologies
Saiganesh27.is@gmail.com



The Dell Technologies Proven Professional Certification program validates a wide range of skills and competencies across multiple technologies and products.

From Associate, entry-level courses to Expert-level, experience-based exams, all professionals in or looking to begin a career in IT benefit from industry-leading training and certification paths from one of the world's most trusted technology partners.

Proven Professional certifications include:

- Cloud
- Converged/Hyperconverged Infrastructure
- Data Protection
- Data Science
- Networking
- Security
- Servers
- Storage
- Enterprise Architect

Courses are offered to meet different learning styles and schedules, including self-paced On Demand, remote-based Virtual Instructor-Led and in-person Classrooms.

Whether you are an experienced IT professional or just getting started, Dell Technologies Proven Professional certifications are designed to clearly signal proficiency to colleagues and employers.

[Learn more at www.dell.com/certification](http://www.dell.com/certification)

Table of Contents

Abstract	4
Introduction	6
Machine Learning	6
Deep Learning	7
Reinforcement Learning	7
History of Deep Reinforcement Learning	9
Reinforcement Learning Simplified	10
Applications of Deep Reinforcement Learning	10
Applications of Deep Reinforcement Learning in computer vision	12
Unlocking Computer Vision with Dell Technologies	15
Conclusion	17
References	18

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect Dell Technologies' views, processes or methodologies.

Abstract

Deep reinforcement learning is a subfield of machine learning that combines reinforcement learning (RL) and deep learning and works on large data sets. Neural networks became an area of interest during the mid-1980s and deep reinforcement learning gained popularity among researchers.

Deep reinforcement learning uses neural networks in reinforcement learning to represent value functions. It is also an emerging field in Artificial Intelligence, and it fuses the power and the ability of deep neural networks to symbolize and understand the world.

The architecture of Deep Reinforcement Learning consists of:

- an agent and an environment
- A reward system that aims at maximizing the rewards.
- prioritizing near future reward than a futuristic reward

Deep reinforcement learning has the potential to solve complex problems. It has scope in:

- Medicine
- Robotics (MELA – a four-legged robot that learns adaptive behavior)
- Finance
- Smart grids
- Computer vision

Artificial Neural networks can process different types of data like a human brain and with reinforcement learning combined, we can tap into its full potential.

Deep reinforcement learning is used in computer vision which is in turn used in the safety and security vertical. Computer Vision is an emerging field that is important to every business as safety is a priority for all. PowerScale offers an excellent solution for storage of these large amounts of unstructured data that computer vision generates. The benefit of this solution is that PowerScale comes with very powerful processors that can handle large data sets.

A glance at storage and infrastructure requirements

Due to geometrically increasing unstructured data sets, there is a surge in machine learning (ML) and deep learning (DL) research, which requires more powerful hardware to handle both ML and DL parallelly to support enterprise AI applications. Enterprises will prefer private and dedicated hardware rather than a cloud to minimize costs pertaining to the movement of data (training data sets) in and out of the cloud.

AI applications require massive amounts of data and this increases the storage capacity required for these applications. Critical components required are:

- CPU
- GPU
- Memory
- Network
- Storage IOPS

In this article, we will explore deep reinforcement learning, its architecture, and its emergence. We dig deep into the various applications and scopes of deep reinforcement learning and computer vision and try to relate deep reinforcement learning to AI solutions to what Dell has to offer. Some of the challenges include handling very large unstructured data and analyzing it in real-time and storing huge data sets which we will address.

Moreover, we present how Deep Reinforcement Learning can be used in computer vision which falls under the safety and security vertical and examine the storage requirements Dell offers for applications utilizing Deep Reinforcement Learning and unstructured data.

Introduction

The capabilities of the human brain are fascinating; how it can learn from experiences, make decisions based on a situation, and much more. Scientific advances in technology have made it possible to inculcate this decision-making ability to machines (Artificial Intelligence) and the study of computer algorithms used in AI is known as machine learning.

Reinforcement Learning (RL) means reinforcing or training the existing ML models so that they may produce a sequence of decisions. The decisions made generate various types of results which can be classified into two categories – Positive Reinforcement Learning and Negative Reinforcement Learning. In Positive Reinforcement learning, rewards are added to the existing ML model so that they are more likely to generate this result again. Whereas, in Negative Reinforcement learning, punishments are added as a negative behavior so that the model does not choose this path again and it encourages them to perform better.

Deep reinforcement learning (DRL) is a subfield of machine learning that combines reinforcement learning (RL) and deep learning and works on large data sets. The architecture of Deep Reinforcement Learning combines pre-learned skills to create new sets of skills on the fly.

Machine Learning

Machine Learning (ML) is a major subject of study and development in and of itself. ML is a branch of AI that enables machines and computers to learn from their environment without needing to be explicitly programmed. In simple terms, a computer automatically makes decisions learning from its surroundings and previous experiences without being coded to perform a certain way.

Machine learning is a broad discipline with a complicated classification of algorithms that can be divided into three categories:

- **Supervised Learning** is learning from a tagged data set and the goal of this is to generalize. The learning here is supervised and the desired outcome of a problem statement – called a “label” – is already known. Popular ML algorithms in this category are linear regression, logistic regression, support vector machines, decision trees, random forest, and neural networks.
- **Unsupervised Learning** is from unlabeled data and the goal is to compress. When the outcome of a problem is not known, i.e unlabeled data, unsupervised learning will try to classify the information by itself. Well-known algorithms in this class are clustering (K-means) and principal component analysis (PCA).
- **Reinforcement Learning** is learning through trial and error and its goal is to take actions that will increase the rewards. This learning category can be combined with other categories, and it is now a very active research area, as we will see in this series.

Deep Learning

Deep learning (DL) is a type of machine learning that uses an artificial neural network to transform a set of inputs into a set of outputs. DL methods, which often employ supervised learning with labeled datasets, have been shown to solve tasks involving complex, high-dimensional raw input data, such as images, with less manual feature engineering than previous methods, enabling significant progress in fields such as computer vision and natural language processing (NLP).

Reinforcement Learning

Reinforcement learning (RL) refers to goal-oriented algorithms that learn how to attain a difficult objective (goal) or maximize along a particular dimension over many steps. For example, they can maximize the points won in a game over several moves. RL algorithms can start anew and achieve superhuman performance under the correct circumstances.

There are two main threads in the history of reinforcement learning, both of which are long and rich, and were followed individually until coming together in modern reinforcement learning. One thread began in the psychology of animal learning and covers learning through trial and error. This thread runs through some of the earliest artificial intelligence research and contributed to the early 1980s rebirth of reinforcement learning. The other thread is the problem of optimum control and how value functions and dynamic programming can be used to solve it. This thread, for the most part, was not about learning.

Optimal control problems have been addressed by methods of dynamic programming (Bellmann 1957) which is a large scientific area in its own right. Classical conditioning and instrumental conditioning are two types of trial-and-error learning that have their origins in psychology. As a result, the first thread (optimal control) was driven from the start by very algorithmical/mathematical approaches, but the first, still more qualitative, mathematical models for the second thread (animal learning) took much longer to build (see, for example, the Rescorla-Wagner Model). Closed-loop control problems are addressed via optimal control and instrumental conditioning. Classical conditioning, on the other hand, deals with a prediction-only problem because the animal's response has no bearing on the experiment, or, to put it another way, on the environment.

In Reinforcement Learning there are two core components:

- An Agent is computer software that has the sole job of making decisions (take actions) in order to tackle complicated decision-making problems under uncertainty.
- An Environment is a representation of a "problem," which is everything that follows the Agent's decision. The environment reacts to such behaviors in the form of observations or states, as well as rewards, also known as costs.

These two core components continuously interact so that the Agent attempts to influence the Environment through actions, and the Environment reacts to the Agent's actions. This circumstance is referred to as model-based Reinforcement Learning when the Agent is aware of the model. When we have a complete understanding of the environment, we may use Dynamic Programming to identify the best solution.

When the Agent does not know the model, it must make judgments based on inadequate information, perform model-free Reinforcement learning, or attempt to explicitly learn the model as part of the algorithm.

The terms "agents," "environments," "states," "actions," and "rewards" are all used to describe reinforcement learning, which we'll detail below:

- **Agent:** An agent takes action. A drone making a delivery or Super Mario navigating a video game are examples of agents. The agent is the algorithm. It could be beneficial to remember that the agent in your life is you.
- **Action (A):** The collection of all feasible moves that the agent can make is called A. Although the action is almost self-explanatory, agents normally select from a list of discrete, viable actions.
- **Discount factor:** To reduce the effect of future benefits on the agent's choice of action, the discount factor is multiplied by future rewards as discovered by the agent. Why? Its purpose is to make future benefits less valuable than present rewards.
- **Environment:** The region in which the agent moves and is affected by it. The environment receives the agent's current state and activity as input and returns the reward and future state as output. If you're the agent, the environment could be the physical laws and social regulations that process your activities and decide their outcomes.
- **State (S):** A state is a precise location and time in which the agent finds itself; it is an instantaneous configuration that places the agent in relation to other significant objects such as tools, barriers, opponents, or prizes. It can be the current circumstance or any future situation as determined by the surroundings. Have you ever found yourself in an awkward situation? That's a state.
- **Reward (R):** A reward is the form of feedback that we use to determine if an agent's activities in a particular state were successful or unsuccessful. When Mario, for example, touches a coin in a video game, he earns points.
- **Policy (π):** The policy is the method by which the agent decides what to do next based on the present situation. It connects states to behaviors, focusing on those with the greatest potential for reward.
- **Value (V):** In contrast to the short-term reward R, the expected long-term return with discount.
- **Q-value or action-value (Q):** Q-value is similar to Value, except that it takes an extra parameter, the current action A. $Q\pi(s, a)$ refers to the long-term return of an action taking action 'a' under policy ' π ' from the current state 's'. Q maps state-action pairs to rewards. Note the difference between Q and policy.
- **Trajectory:** A set of states and actions that have an impact on those states. "To throw across" comes from the Latin. Much like humans in the modern world, an agent's life is like a ball hurled high and arching across space-time unmoored.




Image Source - <https://wiki.pathmind.com/deep-reinforcement-learning#define>

History of Deep Reinforcement Learning

Deep reinforcement learning has been a topic of interest and in works from the 1980s. Interest grew in deep reinforcement learning, where a neural network is used in reinforcement learning to represent policies or value functions. Because the entire decision-making process in such a system, from sensors to motors in a robot or agent, is handled by a single neural network, it's also known as end-to-end reinforcement learning.

One of the first successful applications of reinforcement learning with neural networks was TD-Gammon, a computer program developed in 1992 for playing backgammon. In 2013, DeepMind performed extremely well in Atari video games by using Deep Reinforcement Learning. The neural network used a combination of deep reinforcement learning and deep Q-networks (DQN) and performed equal or superior to a professional human game tester with limited prior knowledge. In 2015, AlphaGo, a computer program schooled with deep reinforcement learning to play Go, became the first computer Go program to defeat a human professional Go player without handicap on a full-sized 1919 board.

Beyond games, deep reinforcement learning has been applied to a variety of fields. It has been utilized in robotics to enable robots to conduct simple household duties and solve a Rubik's cube with the help of a robot hand. Deep Reinforcement Learning has also found uses in sustainability, i.e. lowering data center energy consumption. Deep reinforcement learning for autonomous driving is a hot topic in academia and industry.




Image Source - <https://towardsdatascience.com/drl-01-a-gentle-introduction-to-deep-reinforcement-learning-405b79866bf4>

Reinforcement Learning Simplified

- **Input:** The input should be a starting state for the model to work from.
- **Output:** There are several alternative outputs, just as there are numerous solutions to a given issue.
- **Training:** The model will return a state depending on the input, and the user will determine whether to reward or penalize the model based on its output.
- The model is always evolving.
- The optimal option is chosen based on the highest possible reward.

Applications of Deep Reinforcement Learning

Reinforcement learning is making a series of decisions in order to achieve a goal over a long period of time. Reinforcement learning conducts tactical and strategic tasks, while other types of AI perform perceptual tasks, such as recognizing the content of a picture. While games are a good proxy for the kinds of problems that reinforcement learning can resolve, it's also being used in real-world processes in the corporate and governmental sectors.

- Robotics
- Industrial Operations
- Supply Chain & Logistics

- Traffic Control
- Bidding & Advertising
- Recommender Systems
- Load Balancing
- Augmented NLP

AlphaZero, a single system that essentially taught itself how to play and master chess from scratch, has been formally tested by chess masters and has consistently won. Traditional chess engines, such as Stockfish and IBM's Deep Blue, plan their strategy based on thousands of rules and situations created by expert human players to anticipate every potential scenario.



Image Source - <https://tinyurl.com/mtn923yu>

Shell is guiding gas drills into the subsurface using deep learning algorithms that have been taught using historical drilling data as well as data from simulations. Mechanical data from the drill bit, such as pressure and bit temperature, are also included in the Deep Reinforcement Learning technology, as well as seismic survey data pertinent to the subsurface. As a result, the drilling machine's human operator has greater awareness of the environment in which they're working, which leads to faster results and less wear and tear – or damage – on pricey drilling machinery.

Deep reinforcement learning was employed in Chinese retail to improve the online retail environment of Taobao, an online shopping website owned by Alibaba which is one of the world's largest e-commerce websites.

- **Deep Reinforcement Learning in landmark detection**

In recent years, autonomous landmark detection has gotten a lot of interest. The rise of automation for data evaluation is one of the key causes for this increased propensity. The reason for utilizing an algorithm instead of a person for landmarking is that manual annotation is time-consuming, tedious, and prone to errors.

- **Deep Reinforcement Learning in object detection**

The goal of object detection demands the algorithm to find bounding boxes for all objects in a given image. Object detection has been attempted in a variety of ways, including focal loss and fast yolo (you only look once). It performs object detection in video on embedded devices.

- **Deep Reinforcement Learning in object tracking**

Real-time object tracking has a wide range of applications in autonomous driving, robotics, security, and even sports, where the umpire requires accurate ball movement estimation to make choices. Over the last few years, a deformable face tracking method that could predict bounding box along with facial landmarks in real-time, an end-to-end active object tracker using Deep Reinforcement Learning, and an end-to-end method for SOT using sequential search strategy and Deep Reinforcement Learning have all been proposed, with promising results.

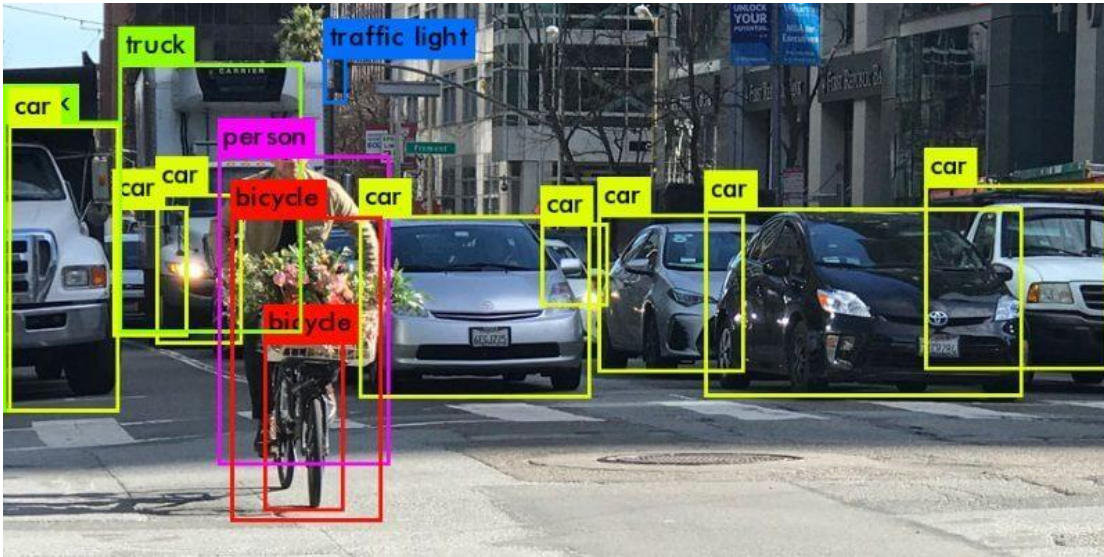


Image Source - <https://azati.ai/wp-content/uploads/2020/04/object-detection-800x400-1.jpg>

Applications of Deep Reinforcement Learning in computer vision

One of the most important areas of Artificial Intelligence is computer vision, which allows computers to detect and analyze images and scenes. Let's take a look at the safety and security vertical. Every organization irrespective of its size considers computer vision a very important aspect affecting their business. Facial recognition and road assessment for autonomous cars are two fields where Deep Reinforcement Learning can be helpful.

What is Computer Vision?

Computer vision technology is designed to work in the same way as the human brain. How does our brain, on the other hand, deal with visual object recognition? Our brains rely on patterns to decode individual objects, according to one prominent theory. Computer vision systems are built using this principle.

Pattern recognition is the foundation of today's computer vision algorithms. We train computers on a vast amount of visual data – images are processed, objects are labeled, and patterns are found in those objects. For example, if we submit a million photographs of flowers to the computer, it will analyze them, find patterns that are common to all flowers, and produce a model "flower" at the end of the process. As a result, every time we submit photographs, the computer will be able to precisely recognize whether a certain image is a flower.

Artificial intelligence has made significant strides in recent years, surpassing humans in several tasks relating to object detection and classification, thanks to breakthroughs in artificial intelligence and innovations in deep learning and neural networks.

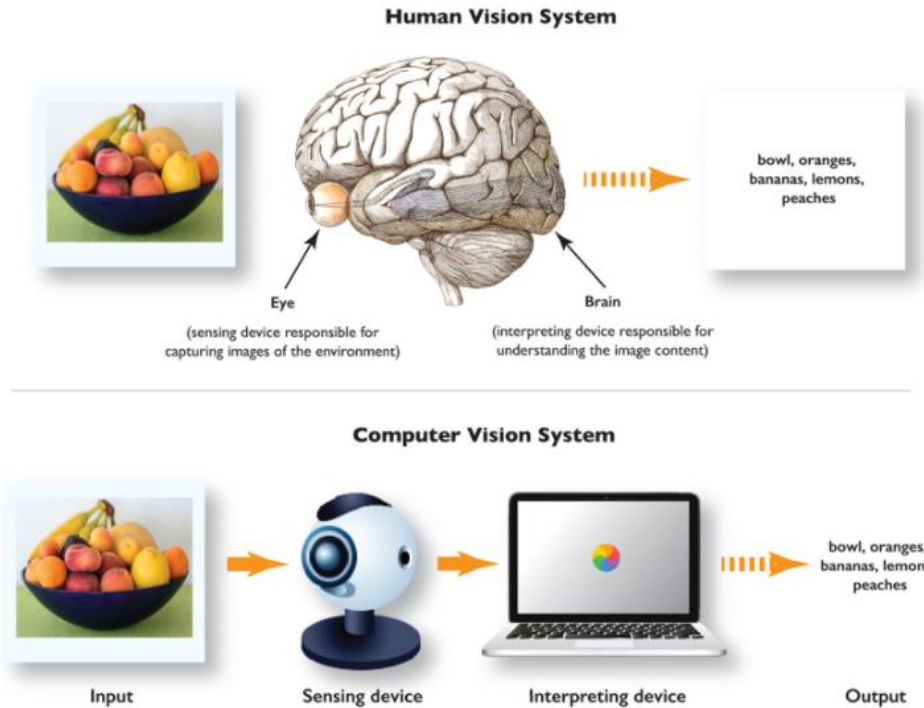


Image Source - <https://xd.adobe.com/ideas/principles/emerging-technology/what-is-computer-vision-how-does-it-work/>

Computer vision is not a new technology; the first computer vision tests took place in the 1950s, and it was used to analyze typewritten and handwritten text at the time. Computer vision analysis processes were basic at the time and required much effort from human operators who had to manually supply data samples for analysis. As you may expect, manually providing a large amount of data was difficult. Plus, the computational power wasn't good enough, so the error margin for this analysis was high.

To comprehend the most recent developments in computer vision technology, we must first examine the algorithms that this technique is based on. Deep learning, a subset of machine learning that uses algorithms to extract insights from data, is used in modern computer vision. On the other hand, machine learning is based on artificial intelligence, which serves as a basis for both technologies.

Computer Vision Applications

- **Content organization**

We already use computer vision tools to help us organize our content. Apple Photos is a great illustration of this. The software has access to our photo collections, and it tags photos automatically, allowing us to view a better-organized collection of images. Apple Photos is fantastic since it curates a curated perspective of your favorite memories for you.

- **Facial recognition**

Photos of people's faces are matched to their identities using facial recognition technology. This technology is built into many of the products we use daily. Facebook, for example, uses machine vision to recognize people in images.

Facial recognition is an important biometric authentication method. Many smartphones today allow users to unlock their phones by displaying their faces. For facial recognition, a front-facing camera is employed; mobile devices scan this image and, depending on analysis, can determine whether the person holding the device is permitted to use it. The beauty of this technology is how quickly it works.

- **Augmented reality (AR)**

Augmented reality apps rely heavily on computer vision. This technology enables AR apps to detect physical things in real-time (both surfaces and individual objects inside a physical location) and use that data to position virtual objects within the physical environment.

Cars can make sense of their environment thanks to computer vision. A smart car contains a number of cameras that capture video from various angles and feed it to computer vision software as an input signal. The system analyses the video in real-time and detects objects such as road markings, objects in the immediate vicinity of the vehicle (such as pedestrians or other vehicles), traffic signals, and so on. Autopilot on Tesla automobiles is one of the most well-known examples of this technology in action.

- **Self-driving cars**

The amount of data we collect today, which is subsequently used to train and improve computer vision, is one of the driving forces behind of self-driving cars.

- **Video Analysis**

In the realm of computer vision, object segmentation in videos is tremendously valuable but a difficult task. A framework for video object segmentation based on Deep Reinforcement Learning has been suggested.

- **Health**

Because image data accounts for 90% of all medical data, it is a critical component for diagnosis in medicine. Image processing is used in many health diagnoses, including X-rays, MRIs, and mammography, to mention a few. During the study of medical scans, picture segmentation proved to be beneficial. Computer vision algorithms, for example, may detect diabetic retinopathy, the leading cause of blindness. Computer vision can analyze and grade images of the back of the eye for disease prevalence and severity.

- **Agriculture**

Many agricultural companies use computer vision to monitor harvests and handle common agricultural issues like weed growth and nutrient deficiency. Computer vision systems analyze photos from satellites, drones, and planes to discover problems early on that could lead to financial losses.

Unlocking Computer Vision with Dell Technologies

Let us take a look at computer vision in safety and security, one of the leading verticals.

The following elements are included in an optimal computer vision workstream:

1. **Data capture:** For all sensor data, the data capture layer is the initial edge aggregation point. This layer's edge platform gathers, formats, and stores the massive volumes of data coming in from cameras and IoT devices. Depending on the requirements and architectural design, a certain amount of deduplication can be done at this layer.
2. **Data curation, movement, storage, and preparation:** Intelligent networking and security procedures, as well as a centralized data warehouse for safe, scalable access and curation of camera data at the edge, data center, and cloud, are used to transfer data created at the edge to the development environments that are required. The next crucial stages in the process to train machine learning models include creating labeled data sets and performing the appropriate preprocessing of the data, such as face redaction.
3. **Computer Vision App Store:** To train in-house computer vision models and/or integrate with third-party computer vision applications, an open platform is required to aid integration and eliminate analytical mistakes caused by out-of-sync isolated copies of data.
4. **Deploy:** Computer vision systems should be installed to the edge to run with minimal latency, relying on scalable edge compute equipment.
5. **Consume:** Real-time integration with corporate APIs and business apps generates real-time insights to offer the operational efficiencies and timely business insights you need.




Image source - <https://www.delltechnologies.com/asset/en-ca/solutions/business-solutions/briefs-summaries/dell-technologies-computer-vision-solution-brief.pdf>

Computer vision use cases extend across industries, including:

- **Healthcare:** Medical Device Management, Digital Pathology, Extended Patient Care
- **Energy:** Grid Modernization, Capacity Balancing, Vegetation Control, Maintenance
- **Retail:** Loss Prevention, Inventory Management, Smart Store Operations, Customer Personalization
- **Manufacturing:** Plant Visibility, Equipment Effectiveness, Quality Assurance, Predictive Maintenance, Connected Worker
- **Transportation:** Capacity Balancing, Optimized Fleet Management, Cargo Tracking, Intelligent Driving
- **Smart Cities:** Smart Buildings, Traffic Pattern Analysis, Environmental Quality & Management Campuses, Pedestrian Safety
- **Events & Environment:** Safety & Security, Queue Management, Smart Facility, Social Distancing, Mask Adherence

Organizations today are flooded in data, but the technology to process and analyze it sometimes struggles to keep up with the deluge of telemetry emitted by every machine, application, and sensor. A surge in unstructured data has proven particularly difficult for traditional information systems based on organized databases, prompting the development of new machine learning and deep learning techniques. As a result, businesses must either buy or construct systems and infrastructure to support machine learning, deep learning, and AI workloads.

That's because the confluence of geometrically growing unstructured data sets, a surge in machine learning and deep learning research, and exponentially more powerful hardware designed to parallelize and accelerate ML and DL workloads has sparked a surge in interest in enterprise AI applications. IDC predicts AI will become widespread by 2024, used by three-quarters of all organizations, with 20% of workloads and 15% of enterprise infrastructure devoted to AI-based applications.

AI applications necessitate a large amount of data, which raises the storage requirements for these applications. The following are the system components that are most important for AI performance:

- **CPU.** Operating the VM or container subsystem, dispatching code to GPUs, and handling I/O are all responsibilities. Although systems using second-generation (Rome) AMD Epyc CPUs are becoming increasingly prevalent, current offerings use a second-generation Xeon Scalable Platinum or Gold processor. Current-generation CPUs include characteristics that considerably speed up ML and DL inference procedures, making them appropriate for production AI workloads that use models trained on GPUs.
- **GPU.** Handles machine learning (ML) or deep learning (DL) training and (often) inference, which is the capacity to categorize data automatically based on learning, and is often an Nvidia P100 (Pascal), V100 (Volta), or A100 (Ampere) GPU for training and a V100, A100, or T4 (Turing) GPU for inference. Though AMD's Instinct (Vega) GPUs have yet to gain widespread adoption among system makers, some OEMs now offer solutions in 1U-4U or Open Compute Project 21-inch form sizes.
- **Memory.** System memory isn't frequently a limitation because AI functions happen from GPU memory, and servers typically have 128 to 512 GB of DRAM. Current GPUs feature embedded high-bandwidth memory (HBM) modules that are substantially quicker than traditional DDR4 or

GDDR5 DRAM (16 or 32 GB for the Nvidia V100, 40 GB for the A100). For AI operations, a system with 8 GPUs can have a total of 256 GB or 320 GB of HBM.

- **Network.** Multiple 10 Gbps or higher Ethernet ports are used because AI systems are frequently clustered together to scale performance. For intracluster communication, some feature InfiniBand or dedicated GPU (NVLink) ports.
- **Storage IOPS.** Another performance bottleneck for AI applications is data movement between storage and compute subsystems. As a result, most systems employ local NVMe SSDs rather than SATA SSDs.

Massive amount of data utilized for training sets, as well as data created by Deep Reinforcement Learning networks, necessitate a storage system capable of handling big data and unstructured data. For this use case, Dell Technologies PowerScale is the best option.

Conclusion

While reinforcement learning remains a hot topic in academia, great progress has been made in using it in the real world. Still, Reinforcement Learning has a lot of issues and is difficult to use. If more effort was invested towards tackling the problems, it would be influential and effective in the following ways:

- **Assisting humans:** Perhaps it is a stretch to predict that Reinforcement Learning will one day grow into artificial general intelligence (AGI), but it does have the capability to assist and collaborate with humans. Consider the possibility of a robot or virtual assistant collaborating with you and taking your actions into account to reach a common goal. Wouldn't that be fantastic?
- **Understanding the consequences of different strategies:** Because time does not go backward and things only happen once, life is truly great. However, there are moments when we wish to know how things could be different.

In writings on new technology, computer vision is a prominent topic. This technology differs from others in that it takes a distinct approach to data. Massive amounts of data that we generate daily, which some perceive as our generation's curse, are also used to our advantage – the data can educate computers to see and understand objects.

Computer vision is one of the most advanced and constantly expanding areas and is positioned to execute a wider range of tasks in the future. Computer vision technologies will not only be easier to train, but they will also be able to distinguish more from visuals than they do now. Dell Technologies is dedicated to continuous innovation as this field advances, with major R&D expenditures in computer vision. This commitment paves the way for game-changing, actionable intelligence to revolutionize an organization.

This technology also represents a significant step forward in our civilization's efforts to develop artificial intelligence that is as intelligent as humans. In terms of reinforcement learning and computer vision applications, we have barely scratched the surface in this article.

References

- [A Beginner's Guide to Deep Reinforcement Learning](#) – [Pathmind]
- [A gentle introduction to Deep Reinforcement Learning](#) – [Towards Data Science]
- [Deep reinforcement learning](#) – [Wikipedia]
- [Reinforcement learning](#) – [Wikipedia]
- [What is Computer Vision & How Does it Work? An Introduction](#) – [Adobe XD Ideas]
- [Infrastructure for machine learning, AI requirements, examples \(techtarget.com\)](#) – [Tech target]
- [The Applications of Deep Reinforcement Learning | GetSmarter Blog](#) – [GetSmarter Blog]
- [Reinforcement learning – GeeksforGeeks](#) – [GeeksforGeeks]
- [\(PDF\) Deep Reinforcement Learning in Computer Vision: A Comprehensive Survey \(researchgate.net\)](#) – [Research gate]
- [Everything You Ever Wanted To Know About Computer Vision. Here's A Look Why It's So Awesome](#) – [Towards Data Science]
- [Dell Technologies Computer Vision Solutions](#) – [Dell Technologies]
- [Infrastructure for machine learning, AI requirements, examples](#) – [Search Converged Infrastructure]
- [Applications of Reinforcement Learning in Real World](#) – [Towards Data Science]
- [Present and future of computer vision](#) – [Forbes]

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license.

Copyright © 2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.