


Deep Learning Myths, Lies, and Videotape – Part 2: Balderdash!

Ken Steinhardt

 <p>Ken Steinhardt Field Chief Technology Officer Infinidat</p>	<p>Biography</p> <p><i>Ken Steinhardt serves as Field CTO at Infinidat (https://www.infinidat.com). He has held many different roles in the IT industry over the last 47 years, having started as a mainframe applications programmer and systems analyst, part-time when he was still in high school.</i></p> <p><i>After receiving his BS degree in Computer Science from WPI, Ken was at Digital Equipment Corporation for 16 years, EMC Corporation for 20 years, and joined Infinidat from Pure Storage.</i></p> <p><i>His longest-tenured role was as EMC's Vice President of Enterprise Storage & CTO for Global Product Sales.</i></p> <p><i>Ken regularly blogs at https://www.infinidat.com/en/blog</i></p>
<p>Keywords Deep learning, AI, Enterprise Storage, Machine learning, Performance, Computer science, Data storage Paper type Research</p>	

Abstract

In Part 1 of this article¹, the author discussed the history and definitions of Artificial Intelligence (AI), Machine Learning (ML) and Deep Learning (DL), as well as Infinidat's use of true Deep Learning in our Neural Cache software. As the author of this article explains, while there are no enterprise storage systems other than InfiniBox that claim to use DL (and none do), there are many that claim to use AI/ML. So let's take a closer look.


Introduction

For most vendors that claim to have Artificial Intelligence (AI) or Machine Learning (ML) associated with their storage. However, AI or ML is not commonly a native part of the storage system at all. Typically, there is telemetry data gathered from systems that are sent off to a cloud-based data lake, where independently run external AI/ML applications perform predictive analytics. Typical uses are pre-emptive system maintenance, repair, and problem avoidance, capacity planning, and performance planning. These are very good and useful things, and most of the better storage vendors do it. At Infinidat we do this with InfiniVerse².

There is one very large storage vendor in particular that aggressively markets some of their native product features using both the terms AI and ML very freely. In my opinion their use of these terms is usually inaccurate.

Here is an examination of one prominent example. Last year that major storage vendor introduced a completely new storage system targeted to replace multiple incompatible storage system families that they had acquired over the years. As part of their announcement, and ongoing marketing, they continue to claim that this new system has “integrated machine learning” and an “AI/ML-based Resource Balancer.” In December 2020, its President of Global Sales and Customer Operations touted this feature in a press article, stating that their product “has a machine-learning engine to optimize performance and reduce cost by automating labour-intensive processes like initial volume placement.”

The feature that is being referred to appears to be called a resource balancer. After the storage administrator creates a new storage volume, instead of them having to manually determine which one of up to a maximum of four appliances in the system to place the new volume on, the system will choose which appliance to place the volume on, and then automatically complete the process of volume placement.



As I believe any good storage admin will attest, when trying to determine where to manually place a volume, there are many important factors to take into consideration before making their decision. These include questions such as:

- What are the performance requirements for the volume?
- What are the data access patterns for the volume?
- What are the data protection and replication requirements for the volume?

- What are the performance profiles for each of the appliances where it might be placed?
- What are the performance profiles for the other volumes on those appliances?
- How much capacity is already on each appliance?
- What percentage of total capacity is being consumed on each appliance?
- Should data reduction options be enabled for the volume?
- Are there any special security of access concerns?

Unfortunately, the “machine-learning engine” for the product in question doesn’t seem to consider these kinds of important things. The product’s technical primer states that the placement is based on which appliance has the most unused capacity. That’s it.

Let that sink in. Since this feature appears to completely ignore all the real-world important factors of volume placement (mentioned above) there is a fair probability that it may create problems rather than solve any. Fortunately for any users of that platform, it also offers the ability to override the resource balancer and place volumes manually. I can’t imagine any good storage admin that wouldn’t choose to do it manually on that platform, since doing it manually allows their human intelligence to consider all the issues above, of which the resource balancer seems to be unaware.

An analogy might be that you have just returned from shopping at a grocery store and have different groceries to put away in your home. There are four places that you can put them: a refrigerator, a freezer, a storage shelf, or an oven heated to 450 degrees Fahrenheit.

A truly intelligent approach would be to put ice cream in the freezer, milk in the refrigerator, tortilla chips on the storage shelf, and a pizza in the oven. Based upon the kind of “intelligence” in the product described above however, any of these products could wind up in any of the locations, simply based upon how much available free space they have – like putting the ice cream into the oven, or the milk into the freezer.

So how does this product’s AI/ML claim hold up against the AI/ML/DL definitions from Part 1 of this article¹? It doesn’t. Any regular human knows not to put milk in the freezer simply because it can physically fit into a larger space. Just as any good storage admin knows not to put a performance-critical volume on the slowest appliance simply because it has the most available free capacity. There is no need to even consider testing this resource balancer as ML, since it doesn’t pass the more basic test to be AI. In consideration of the Turing Test, it is very easy to tell the difference between what a storage admin (human) would do, and what the machine (the resource balancer) might do. The resource balancer isn’t “intelligent” by definition.

But there is an accepted industry term for what it actually is. First, let me summarize what a basic software flowchart for the resource balancer probably would look like, based upon its described function:

- There is a new volume to place, open the resource balancing module, if enabled.
- The module queries each appliance visible to the server that created the volume, and requests the available free capacity from each.
- The module compares and sorts the responses from highest to lowest (there are no more than at most four different numbers to sort, so that shouldn't take very long for either a machine OR a human!)
- The volume is placed on the appliance that had the highest amount of free capacity.

The general computer science term for this is **programming**, and what appears to be particularly simple programming at that. It doesn't pass the sniff tests to be AI or ML. In spite of this, every public mention of this feature that I have seen from independent sources has simply repeated the vendor's claim that it's "AI/ML," without questioning the accuracy of the claim nor the merit of the feature.

Speaking of merit, there is a more basic issue. This feature may not be AI/ML, but is it even useful? The claimed benefit as was stated above is "automating labour-intensive processes like initial volume placement." So how labour intensive is volume placement?



On our InfiniBox, this process is done automatically by the system, without ever requiring any human intervention. It is impossible to improve storage admin productivity relative to how InfiniBox has always done volume placement, since humans don't ever need to do it.

In practice, even the entire end-to-end larger process of both volume definition and the (fully automated) volume placement can take less than a minute with just a couple of clicks on the GUI – not just on InfiniBox, but also on some of the better storage products from other storage vendors in today's market. We don't call this particular individual feature AI or ML on our system because it isn't – it's just simply good programming.



As I was getting ready to publish the blog³ that this article is based upon, the very same vendor mentioned in the example above has done it again, publicly proclaiming “machine learning” for a new feature that again appears to not even be AI. This new feature is referred to under the category of “Intelligent data reduction enhancements (DRR).”

Prior to this new release of software, that storage system's “deduplication and compression are always engaged inline and cannot be bypassed.” With this new updated version of the storage O/S, the system “will dynamically prioritize workload performance and defer deduplication operations to a period in time with a less extreme I/O demand.”



In simple terms, if performance demand rises to a challenging level, the system will tell the deduplication process to turn itself off. Then, if and when a lower level of performance demand has been achieved, it will tell the deduplication process to turn itself back on.

Here's an analogy: I have a thermostat in my home that senses when the temperature gets too cold and automatically turns the heat on. Then when the temperature rises to the desired temperature, it turns the heat off. I do not consider that to be AI or ML.

Like the data placement example above, one might also ask if this is actually a useful feature. There is only one logical reason why such a feature would be introduced, and that is that the system may experience serious performance problems if attempting to continuously provide data deduplication. So, it is a bit surprising that what appears to merely be a valid and noble attempt to mitigate a problem that was not understood or foreseen when the product was first introduced is now being intentionally highlighted as a "feature."

All IT vendors make mid-course product corrections to address newly discovered problems all the time, but most don't try to sell them as differentiated value-added features simply because they are new – and I still think that this kind of thing is just programming rather than AI, let alone ML.

In summary

The terms AI, ML, and DL have become popular because real AI, ML, and DL have proven to be able to add significant value – and the deeper you go from AI to ML to DL usually the greater the value. At Infinidat, while some of our software we have simply call programming, we pride ourselves on AI/ML/DL in our neural cache

software being the core essence of what makes our systems unique, and more importantly their ability to provide unique value for our customers. We don't take these terms lightly.

When vendors attempt to exploit popular terms – like riding the coattails of AI/ML/DL value by claiming to offer AI, ML, or DL for features that don't meet the accepted definitions for those concepts – they do our entire industry a disservice. When they push features that upon examination might create unwanted problems rather than solve any, they do their users a disservice.

Reference

- ¹ Steinhardt, K. "Deep learning myths, lies and videotape (1: Verities!)", *IT for CEOs & CFOs*. Available at: <https://www.itceoscfo.com/deep-learning-1>
- ² Data Sheet: Cloud-Based Predictive Analytics, Support, Monitoring, and Reporting <https://www.infinidat.com/sites/default/files/resource-pdfs/InfiniVerse%20Data%20Sheet.pdf>
- ³ Steinhardt, K. (4 May 2021), "Deep Learning Myths, Lies, and Videotape - Part 2: Balderdash!", Infinidat. Available at: <https://www.infinidat.com/en/blog/deep-learning-myths-lies-and-videotape-part-2-balderdash>