



# Deep Learning Myths, Lies, and Videotape – Part 1: Verities!

Ken Steinhardt



**Ken Steinhardt**  
Field Chief Technology  
Officer  
Infinidat

## Biography

*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.*

*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.*

*His longest-tenured role was as EMC's Vice President of Enterprise Storage & CTO for Global Product Sales.*

*Ken regularly blogs at <https://www.infinidat.com/en/blog>*

**Keywords** Deep learning, AI, Enterprise Storage, Machine learning, Performance, Computer science, Data storage  
**Paper type** Research

## Abstract

*When it comes to marketing storage products, IT vendors often misuse the terms Artificial Intelligence (AI), Machine Learning (ML) and Deep Learning (DL). Now is the time to set the record straight. Drawing on his expertise as a computer scientist (who fancies himself as a student of the history of computer science), the author of this article separates the truth of AI, ML and DL from the myths and lies surrounding storage.*

## Introduction

Over the history of IT there always seem to be certain high-profile popular terms and concepts that dominate discussion in different eras and capture the imagination, often filling the vocabulary of IT people in those eras. I believe that in recent years one such subject is Artificial Intelligence (AI). Unlike most popular IT terms that come and go, AI has shown that it really has some serious legs, by extending its reign of popularity with two additional hot terms associated with its derivative forms: Machine Learning (ML) and Deep Learning (DL).

Alan Turing was the first to discuss the concept of “intelligence” relative to computers in his famous 1950 paper<sup>1</sup>, “Computing Machinery and Intelligence,” which opens with the question, “Can machines think?” and where he introduces The Turing Test (aka: The Imitation Game)<sup>2</sup>.



In my opinion, and as a computer scientist that fancies himself a student of the history of computer science (and in my deep past has written applications in LISP and Algol that my professors called “AI”), I believe that in recent years the terms AI, ML, and DL are often misused by IT vendors rather than applied accurately. This seems to be particularly true when it comes to marketing storage products. So, let’s separate the truth from the myths and the lies, and start with some definitions.

### **Definitions**

The foundational definition for the entire subject was established by Alan Turing himself: “If a human could not distinguish between responses from a machine and a human, the machine could be considered intelligent.”

The term **artificial intelligence** was first coined by Professor John McCarthy<sup>3</sup> in 1956, where he defined it as “Getting a computer to do things which, when done by people, are said to involve intelligence.”



The term **machine learning** was introduced by Arthur Samuel in 1959 while at IBM, where he developed the first intelligent computer game (checkers)<sup>4</sup>. Machine learning is a direct subset of AI. All ML is also AI, but not all AI is ML. Samuel's definition is: "Machine learning is an application of artificial intelligence (AI) that provides systems the ability to automatically learn and improve from experience without being explicitly programmed."



Playing checkers on the 701

On February 24, 1956, Arthur Samuel's Checkers program, which was developed for play on the IBM 701, was demonstrated to the public on television. In 1962, self-proclaimed checkers master Robert Nealey played the game on an IBM 7094 computer. The computer won. Other games resulted in losses for the Samuel Checkers program, but it is still considered a milestone for artificial intelligence, and offered the public in the early 1960s an example of the capabilities of an electronic computer.

Source: IBM. (<https://www.ibm.com/ibm/history/ibm100/us/en/icons/ibm700series/impacts/>)  
Photo reproduced with kind permission from IBM.

The fundamental difference between AI and ML is that while an AI function may demonstrate intelligence, all of that intelligence was directed explicitly by humans that defined the programming. In ML, the computer itself is capable of making new and independent decisions based upon new input and its programming, without additional human directive.

Deep learning is a direct subset of ML. So, all DL is also ML, but not all ML is DL. Professor Rina Dechter<sup>5</sup> introduced the term in 1986, and defined it as: "Deep learning is a class of machine learning algorithms that use multiple layers to progressively extract higher level features from raw input."






### Data Centre and Virtualization

The basic difference between DL and ML is that DL analyzes **multiple options at multiple times** when making decisions. DL then analyzes, tests, and learns from the effectiveness of its previous decisions to continually and constantly improve the quality of its future decisions. DL is clearly the most intensive and complex approach of the three forms of AI, but also the one that has the most potential to provide the best and most effective results.

If something claims to be AI, ML, or DL, it follows that the “acid test” should be how well it meets the definitions described above.

### Deep Learning in Enterprise Storage today: InfiniBox<sup>6</sup>

Infinidat’s patented Neural Cache software is the “secret sauce” behind how our InfiniBox storage system is able to get faster than all-flash array performance at revolutionary lower cost than alternative all-flash arrays. It does this by offering the highest cache-hit ratios of any storage system in the history of the market (by a lot!) and thus being able to use low-cost media as its highly protected persistent storage.



When I joined Infinidat in 2019, I noticed that Neural Cache was described as “machine learning” in all the materials that I read. As I went deeper and learned more about it, I agreed that it met the definitions for both AI (the basic test to pass) and ML (a harder test to pass). But I also felt that we were not giving Neural Cache the higher level of respect that it deserved – as a true DL solution.



My case to justify Neural Cache as DL is simple. In the rare (typically on average only 5-7% of the time) situation in which data that has been requested by an application server is not already in the cache of an InfiniBox, it must then be retrieved from the back-end persistent storage.

In addition to bringing the requested data into cache to deliver it to the application server, Neural Cache software will also bring in multiple other extra pieces of data into cache. These additional pieces of data will be cached based upon previous intelligent analysis of their possible associations with the piece of data that actually was requested by the application. The Neural Cache software will then observe, analyze, and learn which of those extra pieces of data (if any) were subsequently also requested, and when. That learning will then influence the future multiple options for extra data that might be brought into cache prior to an actual request from an application.

This is true DL! The benefit for users is faster performance than all-flash arrays which derive their performance attributes from hardware components, and at a radically lower cost. This is the advantage of math over media, or software over hardware. It is uniquely enabled by DL.

In the second part of this article, I will look at what other storage systems that claim to use AI are doing.

#### Reference

- <sup>1</sup> Turing, A.M. (1950), 'Computing Machinery and Intelligence', *Mind – A Quality Review of Psychology and Philosophy*, Oxford University Press, Volume LIX, Issue 236, pp 433–460. Available at: <https://academic.oup.com/mind/article/LIX/236/433/986238>
- <sup>2</sup> 'The Turing Test', (18 August 2020), Stanford Encyclopedia of Philosophy. Available at: <https://plato.stanford.edu/entries/turing-test/>
- <sup>3</sup> 'Professor John McCarthy', Stanford University. Available at: <http://jmc.stanford.edu/artificial-intelligence/what-is-ai/index.html>
- <sup>4</sup> 'The IBM 700 Series – Computing Comes to Business', IBM. Available at: <https://www.ibm.com/ibm/history/ibm100/us/en/icons/ibm700series/impacts/>
- <sup>5</sup> 'Rina Dechter', University of California, Donald Bren School of Information and Computer Sciences. Available at: <https://www.ics.uci.edu/~dechter/index.html>
- <sup>6</sup> <https://www.infinidat.com/en/products-technology/infinibox>