

AIOPS MACHINE LEARNING: SUPERVISED VS UNSUPERVISED



This post is intended to provide a short explanation of the difference between supervised and unsupervised machine learning (ML) and offer some simple examples of how we use them in TrueSight AIOps. I am not suggesting that you must have ML skills in your IT organization; rather that an understanding of how ML functions for IT Operations will help you evaluate AIOps strategy and vendors.

For a deeper discussion of evolving IT skill sets, see my other [post](#).

What is “machine learning”?

Machine learning refers to a process of 'training' a machine to execute a task and is differentiated from writing software code to 'program' a machine to execute the same task.

In [software programming](#), you tell a machine every specific action to take and in what order. You let it know in advance what outcomes to expect and how to deal to them. A software application is just a set of instructions to a machine about what to do and how to react to what happens (user input, feedback, data, etc.) “Bugs” are cases where the programmers failed to put in an instruction, did it incorrectly, didn't account for output or user response, etc. That's why good programming is hard: you have to anticipate every possibility and eventuality.

In [machine learning](#), you are concerned with the 'what' the program should accomplish, but not the 'how'. You don't give a specific set of instructions to the machine to execute in order. You don't tell it what data or input to expect and how to respond. You let the machine figure it out. This is a broad

generalization but you get the point.

- Obviously, you can't start out from nothing. Machine learning requires the user to make design decisions about what analytics and algorithms the machine should use to learn.
- Also obviously, the more complex and abstract your task is, the more complicated machine learning becomes. Therefore, machine learning is generally implemented to solve very specific problems.
- Additionally, there are different ways to train the machine based on the desired task. These different approaches are captured in the terms 'supervised' and 'unsupervised' machine learning.

Supervised Machine Learning

If a machine needs to learn a task using sample data ("input") and an expected outcome ("output"), then the learning is supervised. Supervised machine learning gives the machine a starting point – the *input* – and an end point – the *output*. The job of the machine is to infer how to get from input to output.

The machine must be told the 'what', it has to figure out the 'how'. In supervised machine learning:

- The 'supervisor' must make decisions about what sample data will best train the machine.
- The supervisor must determine what learning algorithm should be used.
- The supervisor must verify the accuracy of the machine output.

Once the machine can accurately give the expected output from the sample data, it can be considered 'trained'. It can then be applied to input data that has not previously been analyzed. This type of machine learning is best used on data that is *labeled* (in the IT world = "structured") to solve [classification](#) problems like 'spam/not spam' or 'threat/not threat' and [regression](#) problems like 'when will metric X hit 90%'.

Unsupervised Machine Learning

Machine learning is unsupervised when you have input data but no expected outcome. With no outcome, you can't train the machine so the input data cannot be used as a sample. Instead, the machine is tasked to learn from the data itself. There are no correct answers and no supervisor.

Unsupervised machine learning is used to look at the structure of the data or the distribution of elements in the data set. It is used for [clustering](#) to identify inherent groupings like common phrases in logs/events, or associations, like the frequency when X failure occurs, failure Y also occurs.

Machine Learning Considerations and BMC Implementations

What type of machine learning should be used depends on the data available and the problem you are trying to solve. No one approach works for everything, and even within the same area, different approaches have tradeoffs. Some considerations for machine learning in AIOps:

- Whether you pick supervised or unsupervised learning depends on the problem. IT has problems that fit both profiles. There is no one single "correct" approach: different IT problems require different approaches and multiple different approaches can be used to solve a specific problem.

- Someone must make design decisions about which algorithms are used for machine learning and in the case of supervision, what data is used to train the system and what constitutes "correct". If not done by a vendor, customers must supply that knowledge themselves.
 - **NOTE:** If you don't know what 'good' looks like, you can't supervise machine learning.
- A lot of enterprise IT data is similarly structured regardless of industry or application. E.g. CPU utilization - as a data input - is highly structured and follows general patterns regardless of what workload is running on the server you are monitoring. This means for many use cases, vendors can build machine learning analytics that will be broadly applicable across different IT environments - which BMC in fact does.

Some examples of machine learning in TrueSight AIOps

Here are some examples of machine learning analytics that BMC has implemented and to which products they apply. For each one I indicate whether we have added proprietary BMC IT domain knowledge (e.g. IT data model output for supervised learning) and what value the analytics provide.

Forecasting

Forecasting is determining when metrics will hit thresholds and performing "what if?" scenarios

- **Algorithms:** Proprietary combination of multiple techniques including [Linear regression](#), [Regime change detection](#), [Seasonality decomposition](#), [Box and Jenkins method](#), and more.
- **BMC IT domain knowledge added?** Yes
- **Type of Machine Learning:** Supervised
 - The system is already trained by BMC, so you benefit without having to actively supervise, but you can modify parameters.
- **Products:** [TrueSight Capacity](#)
- **Value:**
 - Reduce on-premises cost up to 30% by optimization of IT resources
 - Reduce or eliminate infrastructure related application failures
 - Eliminate surprise infrastructure expenditures and budget over-runs
 - Plan for upcoming resources needs, budget and expenses

Dynamic Baselining

Determine future behavior of a metric based on that metric's past behavior. Dynamic baselining incorporates seasonality.

- **Algorithms:** [Poisson](#) and [normal linear](#) regression
- **BMC domain knowledge added?** Yes
- **Type of Machine Learning:** Unsupervised
 - Historical data from metric used without training or specific output)
- **Products:** [TrueSight Capacity](#), [TrueSight Operations Management](#)
- **Value:**
 - Reduce event noise up to 90% and improve productivity
 - Reduce the number of incidents generated from events up to 40%
 - Proactively remediate issues before any service impact to meet SLAs

Clustering

Find similarities and frequency distributions of word pairings in unstructured data (logs, notes, etc.).

- **Algorithms:** [Levenshtein](#) (logs), [Latent Dirichlet Allocation](#) (events)
- **BMC domain knowledge added?** Yes
- **Type of Machine Learning:** Unsupervised
 - Data from logs or events used without specific outcome known in advance
- **Products:** IT Data Analytics
- **Value:** Reduce time to identify root cause up to 60%

Some concluding thoughts on machine learning in AIOps

All AIOps platforms use machine learning in some capacity to solve specific IT domain problems on specific data sets. Whether it is clustering on events, pattern matching on logs, modeling and forecasting on metrics or something else – someone has done the hard work of looking at what algorithms are best suited to the data and what approach to machine learning using those algorithms fits the desired outcome. If needed, they have also put in the hours to train the system. The value proposition of an AIOps platform is that IT operators are buying that expertise and research in addition to the monitoring or aggregating functions of the solution.

The ultimate benefit to the customer is removing the need for a user to have the appropriate analytic skill set, build and configure analytics and machine learning technology, execute analysis, modeling and system training and then implement it against their domain data. Ideally, the user can focus on operational tasks leveraging their IT and specific ecosystem domain knowledge, trusting the system to provide desired outcomes for decision or automation.

Organizations implementing AIOps platforms should do due diligence to understand as thoroughly as possible the data sets they need to analyze and the outcomes they want to achieve. They can then use those specific use cases to vet potential vendors through a proof of value. For a broader roadmap to implementing AIOps, please see my other [post](#).