IoT Analytics, There and Back Again

Written by Jean-Francois Martin, Head of Products at mnubo

mnubo

Machine learning? Nothing new, really

I studied artificial intelligence (AI) in the early 90's. I remember exploring concepts such as machine learning, neural networks, knowledge representation, prediction, etc. Most, if not all, of these concepts, were first defined and described between the late 50s and early 70s.

Later I had the privilege of implementing one by one, directly or indirectly, several of these AI concepts when computing power allowed for increasingly complex and CPU intensive AI algorithms to find their way into commercial applications.

Early in my career, I developed rule-based systems in a number of areas including fraud detection and risk assessment. The late 80's and early 90's were somewhat of a golden age for rule-based systems. Commercial inference engines (i.e. the AI 'mechanism' used to trigger rules against facts to deduce new facts) were sufficiently robust, and personal computers were fast enough to trigger thousands of rules in seconds to solve complex problems. As computing power grew, we could put more facts and rules in our software that could characterize the problems we wanted to solve. Eventually we encounter the main disadvantage of rule-based systems, namely manual knowledge acquisition and implementation of rules and facts. We indeed had to find the domain experts, interview them, formally extract, represent and implement their knowledge in the form of thousands of single rules and facts, and then hope for the best.

Learning from the Golden Age of Text Analytics

In the meantime, computing power continued to grow and we could envisage using machine learning and neural networks in commercial applications. For myself, and others like me, the early 2000's was the golden age of text analytics. Machine learning enabled the supervised (or not) recognition of patterns (rules) from a set of data. Whereas in the past with rule-based systems we had to manually create these rules, now with machine learning we could let the computer 'learn' these rules with the help of relevant training data that 'contained' the knowledge we were looking to extract. We could analyze posts or tweets in order to deduce their 'sentiment' (positive, negative, or neutral) and then use this information to report the effectiveness of a marketing campaign, for example. Just like rule-based systems, machine learning had its own drawbacks, the main one being the ability (or often inability) to create relevant and large enough data sets to train and then test the system.

IoT Analytics

Nowadays, the Internet of Things (IoT) is transitioning from an idea to a reality. Awareness within industries is increasing and companies are starting to accept there is value associated with the IoT's ability to provide real-time, actionable and compelling insights from sensor data. However, with the hype comes also inevitable disillusions; and the last 30 years of evolution in implementing AI techniques to solve business problems has been at the forefront of various disillusions.

Good Common Sense brings you a long way

We, at mnubo, use common sense when applying computer algorithms to the world of IoT in general and IoT analytics in particular. We believe in using the right tool to solve the right problem. That is, for IoT analytics to be relevant, it needs to leverage a number of AI techniques in conjunction with other approaches such as aggregations, statistics, linked sensor data and so on. We also believe that in order for the Internet of Things to reach its plateau of productivity, the industry needs to get back to basics and clearly state the potential ROI and business value that can be derived from IoT sensor data instead of focusing on vague, nonrelatable buzzwords such as machine learning, prescriptive heuristics, non-semantic models, and so on that are contributors to some of the currently inflated expectations.

Empowering IoT Makers to become Data-Driven

To achieve this, our IoT Analytics SaaS solution, SmartObjects, embeds out of the box generic libraries, which can be augmented for specific verticals. These libraries help answer important business questions such as:

+ Where are my connected products and how are they being used by my clients?
+ How can I bridge the gap between unconnected and connected product states?
+ What proportion of my connected products are active or inactive and why?
+ Which products are performing well, or not, and how can I benchmark them?
+ What is my connected product's average time to activation, to operation, to obsolescence, to loss of clientele, and so on?
+ How are my connected products behaving? How do they compare to the rest of my connected products?
+ Are my clients using my connected products to their full potential? If not how can I help them achieve this potential?
+ What are my product's key patterns of usage, behavior and performance? Any anomalies tracked recently?

We believe that by answering such questions, instead of focusing on technology lingo, will true business value be produced; and formal ROI will emerge in the form of product usage tracking, feedback monitoring, brand engagement, and many others.