TRANSFORMING IT OPS WITH MACHINE LEARNING? APPLY CONTEXT.

Charles Araujo Principal Analyst, Intellyx September 5, 2018







The transformation of IT operations from back-office to business catalyst isn't about machine learning, and it isn't about data. This transformation is about finding contextbased insights that enable the IT organization to focus on and prioritize those things that are most important to the enterprise.

It has become an axiom: you cannot manage what you do not measure.

Data, the thinking goes, is good. And more data is even better.

As a result, IT organizations have invested significant amounts of money and effort instrumenting nearly every facet of their operational technology stack.

The problem is that the haphazard collection of mountains of data doesn't generate insights — it just creates noise.

If you read the industry press, a solution to the noise problem has appeared on the horizon in the form of machine learning. But as organizations attempt to use this emerging technology to transform all of the noise into insights, they realize that it's not as simple or as magical as it sounds.

The real solution, however, does not require magic. In fact, the path from noise to insights may require just one thing: *context*.



IT'S NOISE PROBLEM

To become both more proactive and more efficient, enterprise IT organizations initially turned to monitoring to gain the upper hand on the challenges endemic to managing the modern technology stack.

In those early days, IT suddenly had empirical data that identified the source of problems, could predict impending failures (with varying degrees of accuracy), and help set and manage priorities based on business impact and criticality.

But two things conspired to make these early successes unsustainable.

The first problem was the early success itself. As organizations gained efficiencies and accolades from their initial use of data to manage IT operations, they clamored for more of the same. As they instrumented more of the technology stack and began to capture more data, however, they started to create noise.

Had that been the only cause, organizations could have simply stopped collecting so much data. But there was a second, more significant challenge leading to the surge of noise: the growing complexity of the enterprise technology stack.

The exponentially growing need for automation and ever-increasing interconnectedness of applications has led to both more technology in the stack, as well as more connections to track between applications and their underlying infrastructure. The result is a tangled, endlessly shifting technology architecture where a neat and tidy architectural diagram once existed.

As this complexity arose, it demanded more and more instrumentation. The inevitable result was a massive amount of noise.

The static that all of this complexity generates, however, now threatens to undermine IT's ability to operate and function — and with it, threatens the very viability of the enterprises IT organizations serve.

MACHINE LEARNING TO THE RESCUE?

The rising pervasiveness of IT's noise problem was not lost on the technology companies that make management tools for IT operations.

The static that complexity generates now threatens to undermine IT's ability to operate and function.

New technology has emerged that is giving hope to technology vendors and the IT operations teams that need more advanced tools to solve the noise problem: *machine learning*.

The promise of machine learning, a form of artificial intelligence, is that it can sort through the mountains of data that the enterprise technology stack generates and find patterns in the data that will make it all make sense.

The promise of machine learning to solve IT's noise problem is real. There is, however, one little fly in the ointment: the data itself.

The dirty little secret about machine learning is that to work it not only requires lots of data (at least that part is not a problem for IT!), but also that it must be good data.

Machine learning does not work well if data is inaccurate or fragmented.

The problem is that most of the data that organizations generate is incomplete — fragmented and difficult to understand in an abstract form. While the data may make perfect sense in the context of what is happening in the moment, once a system extracts that data, it becomes difficult to use.

It turns out, therefore, that the key to being able to trust and understand data — and, likewise, the key to making machine learning work most effectively — is data context.

UNDERSTANDING CONTEXT

Imagine that you were tasked with creating a card catalog for a library, and were told that all of the information about every book already exists. All you have to do is put it together in a way that makes sense.

With an evil grin on his face, however, the head librarian dumps a massive box of small pieces of paper on the floor. Each piece of paper has just a word or small collection of words on it. Some of the words represent book titles. Others represent author names and others still represent names of publishers or some other small piece of information related to a book.

The problem, of course, is that you don't know which is which.

It turns out that the key to trusting and understanding data and to making machine learning work - is data context.



Yep, you have all of the information you need, but it is going to take a lot of time and energy to make sense of it — and there's a very high chance that a lot of what you come up with will be completely wrong.

Now imagine the same situation with one small difference. Each piece of paper also contains a description of what type of data it represents and its relationship to other parts of data in the pile. How much more quickly and accurately could you put together that card catalog?

That is the power of context.

Most systems that organizations use to instrument and collect data from the technology stack do so in a way very similar to a bunch of huge boxes, each filled with lots of little snippets of data. Sure, you can use machine learning to try to make sense of it, and it may provide some value, but it will take a long time, cost a lot of money, and much of what it finds will be wrong.

This is where the idea that just having lots of data goes wrong, even when — and perhaps, especially when — you apply machine learning to it.

The lack of context makes it much harder to transform all that data into meaningful insights. The information that describes the data and its relationship to other data is called metadata — and it's the secret sauce that can unleash the power of machine learning and allow it to deliver on its promise to IT operations teams.

PUTTING CONTEXT INTO CONTEXT

Machine learning is, admittedly, a powerful tool.

Unlike traditional algorithms in which developers pre-program a specific set of parameters into the system, machine learning systems "learn" and create their own algorithms by analyzing massive volumes of data.

While there is a lot of technical nuance beneath the covers that we will not address here, the result can appear almost magical. Plug it in, turn it on, and it spits out magical insights.

On the one hand, as we've already addressed, it's not that simple. Organizations need lots of good data — with context — to make it work.

It's also important, however, to realize that making machine learning work for your organization is, in fact, no magic trick.

Like most technology projects, there is a systematic approach to applying machine learning that, if followed, will significantly increase the likelihood of a good outcome.

The first step in this transformation is to recognize that creating context is really about creating a common data model that connects the technical data to a business context and provides contextual consistency within the data.

As IT leaders have recognized the importance of modernizing their IT operations, they have sought ways to automate and apply machine learning and various forms of artificial intelligence (AI) to their operations. Unfortunately, many of these IT leaders have been led to believe that they could simply collect data, turn on some AI, and by doing so would magically transform their IT operational model.

The hard reality is that it just doesn't work that way.

To transform their operations, IT leaders must commit not only to collecting data, but to preparing the data by putting practices in place that ensure data quality and enrich that data with context.

Most importantly, they must recognize that the data isn't actually the starting point for this process. Instead, the process must begin with a clear set of objectives and desired outcomes that will drive what data they collect, and how they collect, manage, and enrich it.

THE CONTEXT-POWERED TRANSFORMATION OF IT OPERATIONS

As an IT leader, there are several new realities that you must face and address. The first is that the newfound complexity of the technology stack is here to stay — and this complexity will only continue to increase over time.

Second, managing this complexity will require that you apply data and automation effectively and relentlessly if you want to have any hope of keeping it all running smoothly. Put this together, and you realize that this represents a fundamental transformation in the way you manage IT operations.

The good news is that the instincts to instrument your environment, collect data, and even to look at machine learning as a tool to help you make sense of it were all right-minded. As we've discussed, however, the missing piece has been context. And it is the addition of context, which will be the ultimate enabler of IT operations' transformation.

This context-powered transformation of IT operations requires, however, that IT leaders take a fresh look at how they approach IT operations management itself.

IT has historically taken a process-centric view of managing operations. As data became more available, organizations used it to merely augment this process-based approach.

To effectively manage the complexity of the modern stack, however, IT organizations must reorient their operational models around this context-infused data.

The first step in this transformation is to recognize that creating context is really about creating a common data model that connects the technical data to a business context and provides contextual consistency within the data. Once you have structured the mountains of data in this way, you can then apply machine learning and other AI technologies to make sense of it.

In doing so, however, it will be essential that you focus on the stories the data tells you — the insights it delivers — that will enable you to focus on the right things at the right time.

THE INTELLYX TAKE

In the end, the most important thing for IT leaders to realize is that the transformation of IT operations isn't about machine learning, and it isn't about data.

This transformation is about finding context-based insights that enable the IT organization to focus on and prioritize those things that are most important to the enterprise. That's no small feat amid all of the noise and static that the modern enterprise stack generates.

It's tempting to think that just collecting data is enough. But this transformation demands a reorientation in how organizations collect, enrich, and use that data.

It also demands that IT organizations change the way they look at the tooling they use to support this transformation. Tools that abstractly collect data, without context, will become increasingly less useful.

Instead, IT leaders should seek out next-generation platforms that put data — and context — at the center of the operational paradigm.

It is both a challenging and incredible time to be an IT leader. On the one hand, the complexity of the enterprise IT stack makes the job of managing it nearly impossible using traditional approaches.

At the same time, however, the challenge that this complexity creates has also created an opportunity to transform IT operations at its core, using context-infused data to tame the complexity and make IT an enabler of business value.

ABOUT THE AUTHOR

Charles Araujo is an industry analyst, internationally recognized authority on the Digital Enterprise and author of The *Quantum* Age of *IT*: Why Everything You Know About IT is About to Change.

He is Principal Analyst with Intellyx, the first and only industry analyst firm focused on agile digital transformation. He has authored three books and published over 100 articles. He has been a regular contributor to both InformationWeek and CIO Insight Magazine and has been quoted or published in magazines, blogs and websites including Time, CIO, CIO & Leader, IT Business Edge, TechRepublic, Computerworld, USA Today and Forbes.



He is the founder of The Institute for Digital Transformation and a sought after keynote speaker having addressed over 10,000 business and IT leaders in 10 countries over the last several years. He is passionate about the power of technology to deliver competitive and transformational advantage to organizations and in the critical need to develop next generation "digital leaders" that can transform their organizations into Digital Enterprises. He is presently at work on a new book entitled, *Thinking Digital: How to Thrive and Win in the Digital Era*, which will explore this topic in detail.

Prior to joining Intellyx, Charles served as an advisor and consultant for nearly twenty years, leading numerous large scale transformation programs for Fortune 1000 organizations and government institutions involving as many as 10,000 program participants. In his early career, he spent many years working in and with IT organizations in the healthcare, financial services and aerospace industries, directly leading teams of more than 100 members.

ABOUT INTELLYX

Intellyx is the first and only industry analysis, advisory, and training firm focused on agile digital transformation.

Intellyx works with enterprise digital professionals to cut through technology buzzwords and connect the dots between the customer and the technology – to provide the vision, the business case, and the architecture for agile digital transformation initiatives.

ABOUT SCIENCELOGIC

ScienceLogic is a leader in IT Operations Management, providing modern IT operations with actionable insights to predict and resolve problems faster in a digital, ephemeral world. Its solution sees everything across cloud and distributed architectures, contextualizes data through relationship mapping, and acts on this insight through integration and automation.

Trusted by thousands of organizations, ScienceLogic's technology was designed for the rigorous security requirements of United States Department of Defense, proven for scale by the world's largest service providers, and optimized for the needs of large enterprises. https://sciencelogic.com/

Copyright @ Intellyx LLC. As of the time of writing, ScienceLogic is an Intellyx customer. Intellyx retains final editorial control of this paper. Photo credit: Colleen Cavendish.

9