

A close-up, low-angle photograph of the spines of several books. The books are arranged in a row, receding into the distance. The spine of the book in the foreground is painted a vibrant red, while the others are a light, natural wood color. The books are resting on a reflective surface, creating a subtle reflection. The background is a soft, out-of-focus light blue.

DETECTING
THE BUSINESS
INCIDENTS
THAT MATTER
**WITH ANOMALY
DETECTION**

Remedy urgent problems
faster and capture
opportunities sooner

WHAT IS AN ANOMALY?

With all the different systems going into managing a business and keeping track of every aspect of business activity, there is more data available than ever before. This includes the operational performance of applications and infrastructure components as well as key performance indicators (KPIs) that evaluate the success of the organization. With millions of metrics that can be measured, companies tend to end up with quite an impressive dataset to explore the performance of their business.

Within this dataset are data patterns that represent, basically, business as usual. An unexpected change within these data patterns, or an event that does not conform to the expected data pattern, is considered an anomaly. In other words, an anomaly is a deviation from business as usual.

So what do we mean by “business as usual” when it comes to business metrics? Surely we don’t mean “unchanging” or “constant;” there’s nothing unusual about an ecommerce website collecting an out-of-the-ordinary amount of revenue in a single day – certainly if that day is Cyber Monday. That’s not unusual because a high volume of sales on Cyber Monday is a well-established peak in the natural business cycle of any business with a web storefront.

Indeed, it would be an anomaly if such a company didn’t have high sales volume on Cyber Monday, especially if Cyber Monday sales volumes for previous years were very high. The absence of change can be an anomaly if it breaks a pattern that is normal for the data from that particular metric. Anomalies aren’t categorically good or bad, they’re just deviations from the expected value for a metric at a given point in time.

Anomalies aren’t categorically good or bad, they’re just deviations from the expected value for a metric at any given point in time.

WHY DO COMPANIES NEED ANOMALY DETECTION?

Outliers in the data are caused by business incidents in the real world: a new successful marketing campaign which increases leads, a promotional discount that drives up sales, the launch of a much improved module that reduces load time of key pages, lifting conversion rates, or a software update which breaks the localization code for an ecommerce website, tanking online sales in Asia.

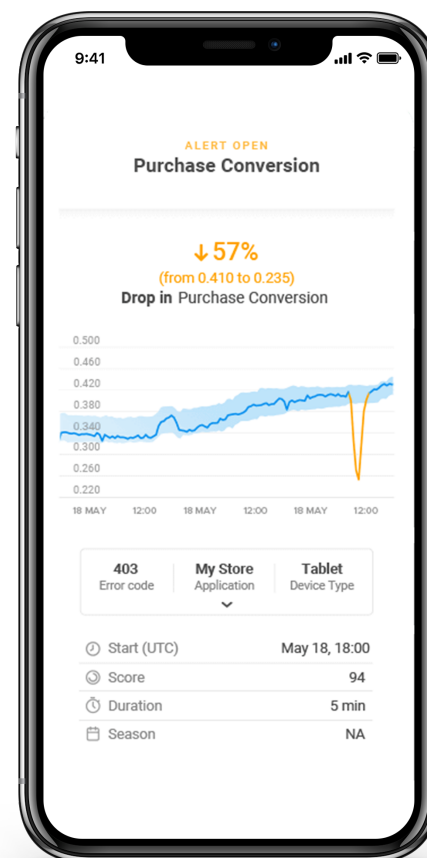
Business incidents are the real-world causes, the anomalies in KPIs are the effects.

In the case of a new marketing campaign, the quantifiable increase in leads above the norm is the anomaly which tells you that the campaign was successful. When the anomaly is a good thing, as it is in this case, we want to accurately attribute it to the right business incident so that we can repeat that success and increase the effect. Maybe the marketing team behind that campaign needs a bigger budget for the next fiscal year, along with a bigger headcount?

You can't correctly attribute a specific anomaly to the underlying business incident if you don't know about anomalies to begin with. And that's one of the main reasons companies need anomaly detection: to get accurate feedback on the effectiveness of business initiatives so that money and manpower can be utilized much more efficiently and to greater impact for a company's bottom line.

Furthermore, would you rather find out instantly from your own data metrics that something is amiss with your online sales or from angry (and now lost) customers on social media? Anomaly detection can point to positive business incidents as well as to potential disasters.

Would you rather find out instantly from your own data metrics that something is amiss with your online sales or from angry (and now lost) customers on social media?





HOW ARE ANOMALIES DETECTED?

As we've explained above, anomalies are deviations from expected values for the time series of that particular metric. Those expectations are derived from a model built and continuously updated by the same data we're checking for outliers. The reasoning here may seem circular: we're comparing a sequence of values against a known standard (in this case, a model) which itself was created from and is being updated by those same values.

Imagine you've recently moved into a new neighborhood where you don't know any of your new neighbors. Now, also imagine one of these neighbors leaves his house for his morning jog every morning at seven—rain, wind or shine. Months later, one morning, you don't see him leave. Nor do you ever see him return home, confirming your suspicion that he never left.

Your months of observing your neighbor jogging every morning is a pretty weird thing to do, but technically speaking, it has created and then reinforced a mental model (every morning at seven he starts jogging), which in turn created an expectation (every morning I will see him start his jog at or near seven) which isn't met in one particular instance (I expected to see him start his jog at seven this morning, but didn't).

Human brains naturally focus on and flag anomalies, the rare exceptions to the rule. An anomaly detection system is a piece of software written to do the same thing, but with data.

Continuing with the neighbor analogy, after seeing this anomaly you might speculate as to the cause: maybe he's too sick to exercise right now, maybe he's out of town. And it will probably stop there because, aside from neighborly curiosity, it's really none of your business.

When it comes to your organization's KPIs, however, anomalies literally are your business.

THE NECESSITY OF REAL-TIME ANOMALY DETECTION

Earlier, we gave an example of a software update causing online sales from Asia to plummet. Obviously, an anomaly in online sales volume for any specific region or device type needs to be detected immediately, and the same is true for other anomalies. This is because many real-life business anomalies *require* immediate action. That bad software update is causing you to lose a lot of money every second. And since discovering the problem is the first step in resolving it, eliminating the delay between when the problem occurs and when the problem is detected immediately brings you one crucial step closer to rolling back that update and restoring revenue flow from Asia.

This is also true for anomalies which aren't problems to be solved, but opportunities to be seized. For example, an unusual uptick in mobile app installations from a specific geographical area may be due to a celebrity share on social media that has gone viral in that region. Given the short lifespan of such surges, your business has a limited time window in which to capitalize on this popularity and turn all those shares, likes and tweets into sales.

Real-time anomaly detection is advantageous even when the detected anomalies include ones which don't require an immediate response. This is because you can always choose to postpone action on an instant alert, but you can never react in real-time to a delayed alert.

Real-time anomaly detection is advantageous even when the detected anomalies include ones which don't require an immediate response.

But let's think about it – what kind of anomaly detection systems can provide this type of real-time notification? For only one or a few KPIs, a human monitoring a dashboard may work. This *manual* approach, however is not scalable to thousands or millions of metrics while maintaining real-time responsiveness. Beyond the mere number of metrics in many businesses, is the complexity of each individual metric: different metrics have different patterns (or no patterns at all) and different amounts of variability in the values of the sampled data. In addition, the metrics themselves are often changing, often exhibiting different patterns as the data exhibits a new “normal.”



MANUAL VS. AUTOMATED ANOMALY DETECTION

If manual anomaly detection is inadequate, then automated anomaly detection must be used to achieve real-time anomaly detection at large scale, and it must be sophisticated enough to handle all the complexity described above at the scale of millions of data points or more, updating every second. The only way to do this is through machine learning.

The [machine learning algorithms](#) that power Anodot's automated anomaly detection system utilize the latest in AI research to meet this task. Anodot's patented machine learning algorithms fall under the "online" category. This means that each data point in the sequence is processed only once and then never considered again. One key reason for this is that online machine learning applications have the required scalability to deal massive amount of business data.

As each data point in the time series is processed, the online machine learning algorithms work similarly to the human brain in the jogger example:

1. A model that fits the data is created.
2. This model, in turn, is used to predict the value of the next data point.
3. If the next data point differs significantly from what the model predicted, that data point is flagged as a potential anomaly.
4. Anodot's machine learning algorithms use each new data point to intelligently update the model.

AI ANOMALY DETECTION IN THE REAL WORLD

The power of AI to spot anomalies and the opportunities they present, beyond any human capability, has already been used to great scientific success. For example, an AI system developed by NASA's Jet Propulsion Laboratory was able to [detect and command an orbital satellite to image a rare volcanic event in Ethiopia](#) – before volcanologists even asked NASA for that satellite to take images of the eruption.

When working with thousands or millions of metrics, real-time decision making requires online machine learning algorithms. Whether it's saving your business money or gleaning scientific insights from a brief volcanic eruption, real-time anomaly detection catches the important deviations in the data.

Since business metrics report data relating to some facet of a business, an anomaly in the data can reflect an important event or change in the business, possibly affecting the revenue stream. The anomaly in the data could be a sign of an opportunity to earn more money – or of a problem that's causing you to lose it.

Either way, companies need to know what their data is trying to tell them right away in order to take advantage of opportunities or fix costly problems. Simple, manually configured monitoring and traditional BI is insufficient to actively track thousands or even millions of metrics. Automated real-time anomaly detection methods must be used when working at this scale.

ANODOT: REAL-TIME AI ANALYTICS

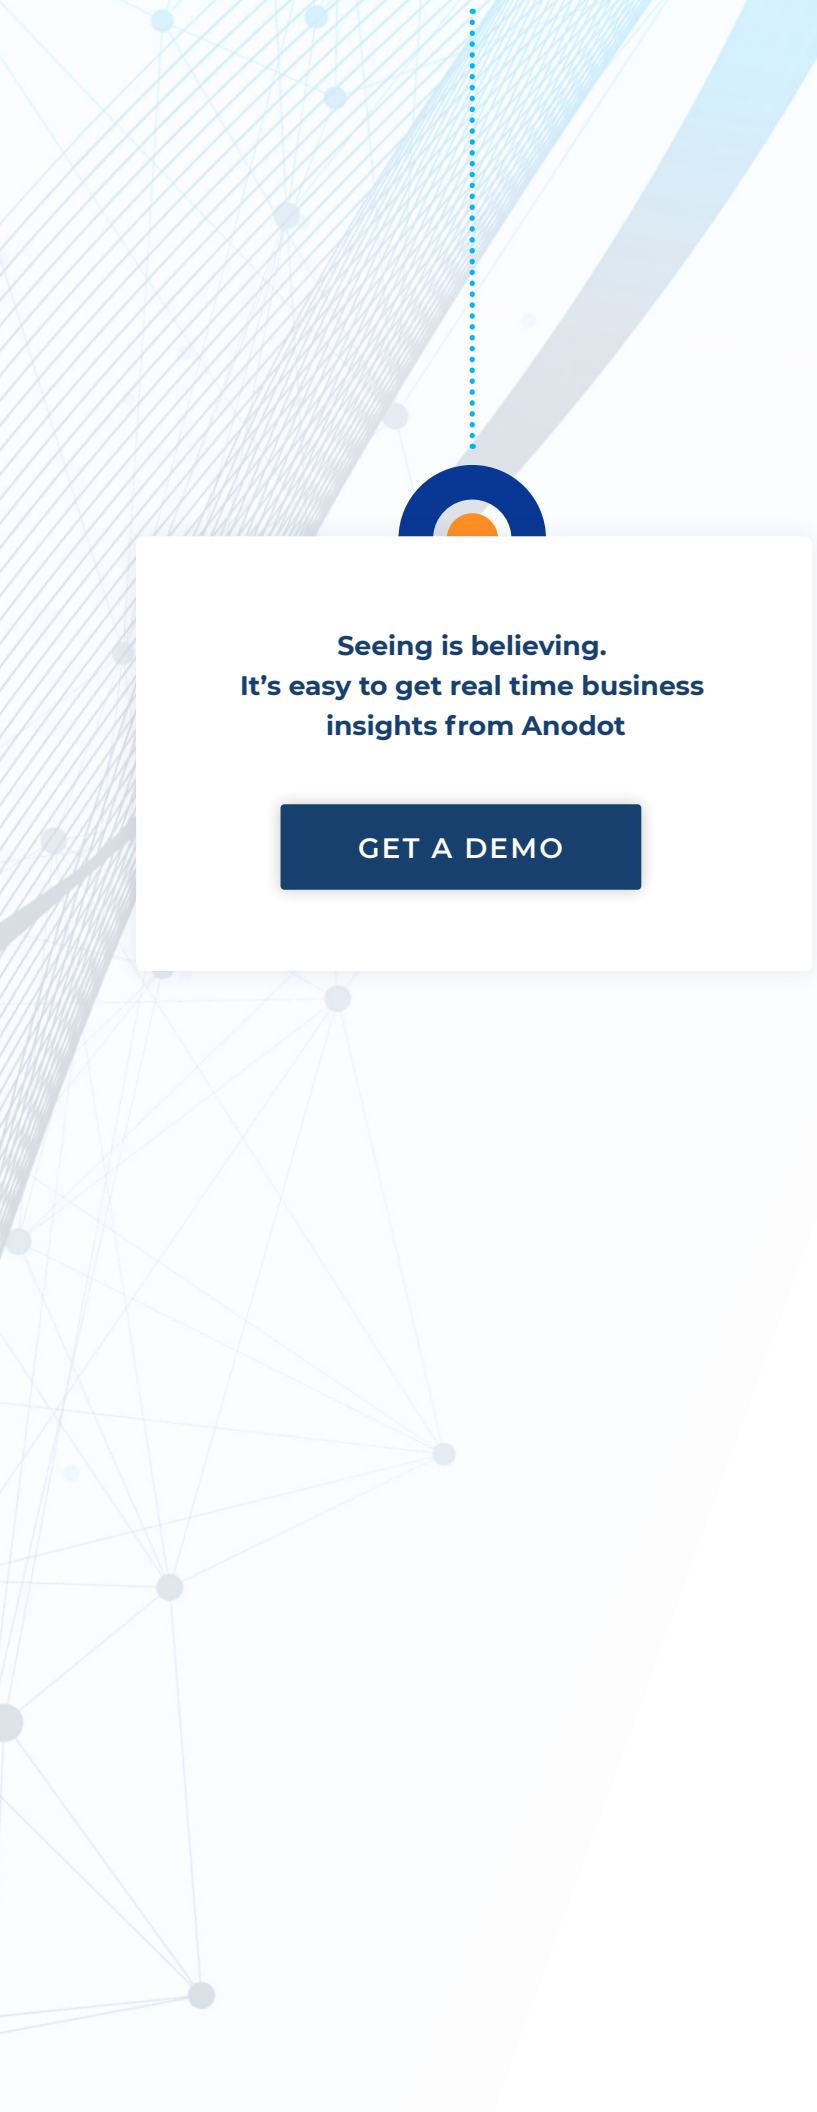
Anodot provides real-time anomaly detection using what are known as “unsupervised” machine learning algorithms. This means the algorithms learn what normal is, and then apply a statistical test to determine if a specific data point is an anomaly. This is crucial for detecting “unknown unknowns,” that is, detecting anomalies which have never been seen before.

Once the anomalies are found by these multiple algorithms, a whole additional layer of machine learning works to discover the relationships between metrics so that the flood of discovered anomalies can be distilled down to a much more manageable number of correlated incidents.

The anomalies correlated across multiple data sources tell a clear story of what is happening with the data, which can then either be investigated by human experts, or used as triggers for automated business actions.

The end result is a powerful analytics solution based on AI, which empowers users to remedy urgent problems faster and capitalize on opportunities sooner.

[Anodot's AI analytics solution](#) is ideal for companies in various industries (e-commerce, online retail, software, adtech, digital entertainment, fintech and more). It automatically illuminates data blind spots so companies never miss another brand damaging incident or business opportunity due to being overwhelmed with dashboards and alerts. Its automated machine learning algorithms continuously analyze all business data and alert in real time whenever an incident occurs, even for questions that were never asked. Its built-in data science means that any user can easily gain actionable insights, even without any data science knowledge. Over 40% of Anodot's customers are publicly traded companies, including Microsoft, Waze (a Google company), AppNexus, Comcast and many others.



**Seeing is believing.
It's easy to get real time business
insights from Anodot**

GET A DEMO

For more information,
please contact Anodot:

North America

669-600-3120

info.us@anodot.com

International

+972-9-7718707

info@anodot.com



anodot

Anodot's autonomous business monitoring platform leverages advanced machine learning techniques to constantly analyze and correlate every business parameter, providing real-time alerts and forecasts, in their context, lowering time to detection and resolution.

The company's leading-edge, patented technology is trusted by clients such as **Facebook, Microsoft, Lyft, Waze, Pandora, Appnexus, Wix** and **King**, in industries ranging from eCommerce to finserv, adtech, telco, gaming and more.

Anodot is headquartered in Silicon Valley and Israel, with sales offices worldwide.

Visit www.anodot.com to learn more.

© Copyright 2019, Anodot. All trademarks, service marks and trade names referenced in this material are the property of their respective owners.