Senseye in Depth

1

Why is Predictive Maintenance so hard?

Over 150 person-years of experience, research and development exposed.

senseye

Introduction

Most maintenance practices are based on service-interval schedules, a preventative and reactive approach that doesn't take actual machine usage and health into account, and which is questionable in substantially reducing unplanned downtime and improving maintenance effectiveness. Predictive maintenance (PdM), on the other hand, is a proactive approach, enabling machine failures to be dealt with before they stop production. PdM is achieved by analyzing huge volumes of machine and maintenance data to decode the health of machines and enable maintenance staff to optimize their activities.

By predicting when machines will break down, companies can therefore eliminate sudden failures, reduce unplanned downtime, increase machine life, optimize scheduled maintenance, and reduce the routine replacement of parts that may, in fact, be perfectly healthy.

Why, then, is the implementation and deployment of PdM so misunderstood and littered with failures (over 80 percent according to one study)?

In this paper, we'll look at the origins and evolution of PdM maintenance, our years of experience and what we've learned, and how all of this informed our own Senseye PdM solution.

PART 1

Industry 4.0 and the challenges of Predictive Maintenance

PdM - a longer history than you might imagine

B-24 "Liberator" bomber.

Predictive maintenance is not new. With its roots in observations made by scientist CH Waddington concerning the maintenance of Royal Air Force bombers during the Second World War, PdM has, in fact, been around for over 75 years.

Observations that a plane's rate of failure or repair tended to be highest immediately after an inspection or maintenance session - a phenomenon later dubbed the "Waddington Effect" - led to the first development in condition-based maintenance. Maintenance processes were adjusted to correspond with the equipment's physical condition and the frequency of its use, and the resulting data compiled and analyzed to form the basis of adjusted inspection cycles, thereby heralding the beginnings of PdM.

Years later, it could be argued that Aerospace remains the leading sector for PdM, successfully deploying techniques such as condition monitoring, diagnostics, and prognostics. The proven success of Integrated Vehicle Health Management (IVHM) and Health and Usage Monitoring Systems (HUMS) in helicopter maintenance over the last 25 years are testament to this. Indeed, HUMS serves as an effective illustration of a successful PdM system.

To monitor the health of a helicopter's rotating components - its gearboxes, engines, and rotors - subtle changes in vibration are recorded during flight, visualized on the HUMS ground station computer, and evaluated by technicians. The resulting intelligence allows fleet operators and aircraft maintainers to make informed decisions about flying and maintaining their aircraft - enhancing safety, decreasing maintenance burden, increasing availability and readiness, and reducing operation and support costs[1].

1. https://publications.parliament.uk/pa/cm200809/cmselect/cmdfence/434/434we09.htm.

Its capacity to inform decisions is key. The typical user of a PdM system is a maintenance professional. Extremely busy, and concerned with overseeing many machines across a site, they'll often have little time (or patience) to understand the workings of their company's latest IT solution.

But, in providing them with a better understanding of the ongoing health of their machines, a PdM system enables them to decrease the chaotic nature of their jobs. Through the use of PdM, downtime, scheduled maintenance and routine replacement of parts are all reduced, removing surprise and lowering the burden for the maintainer.

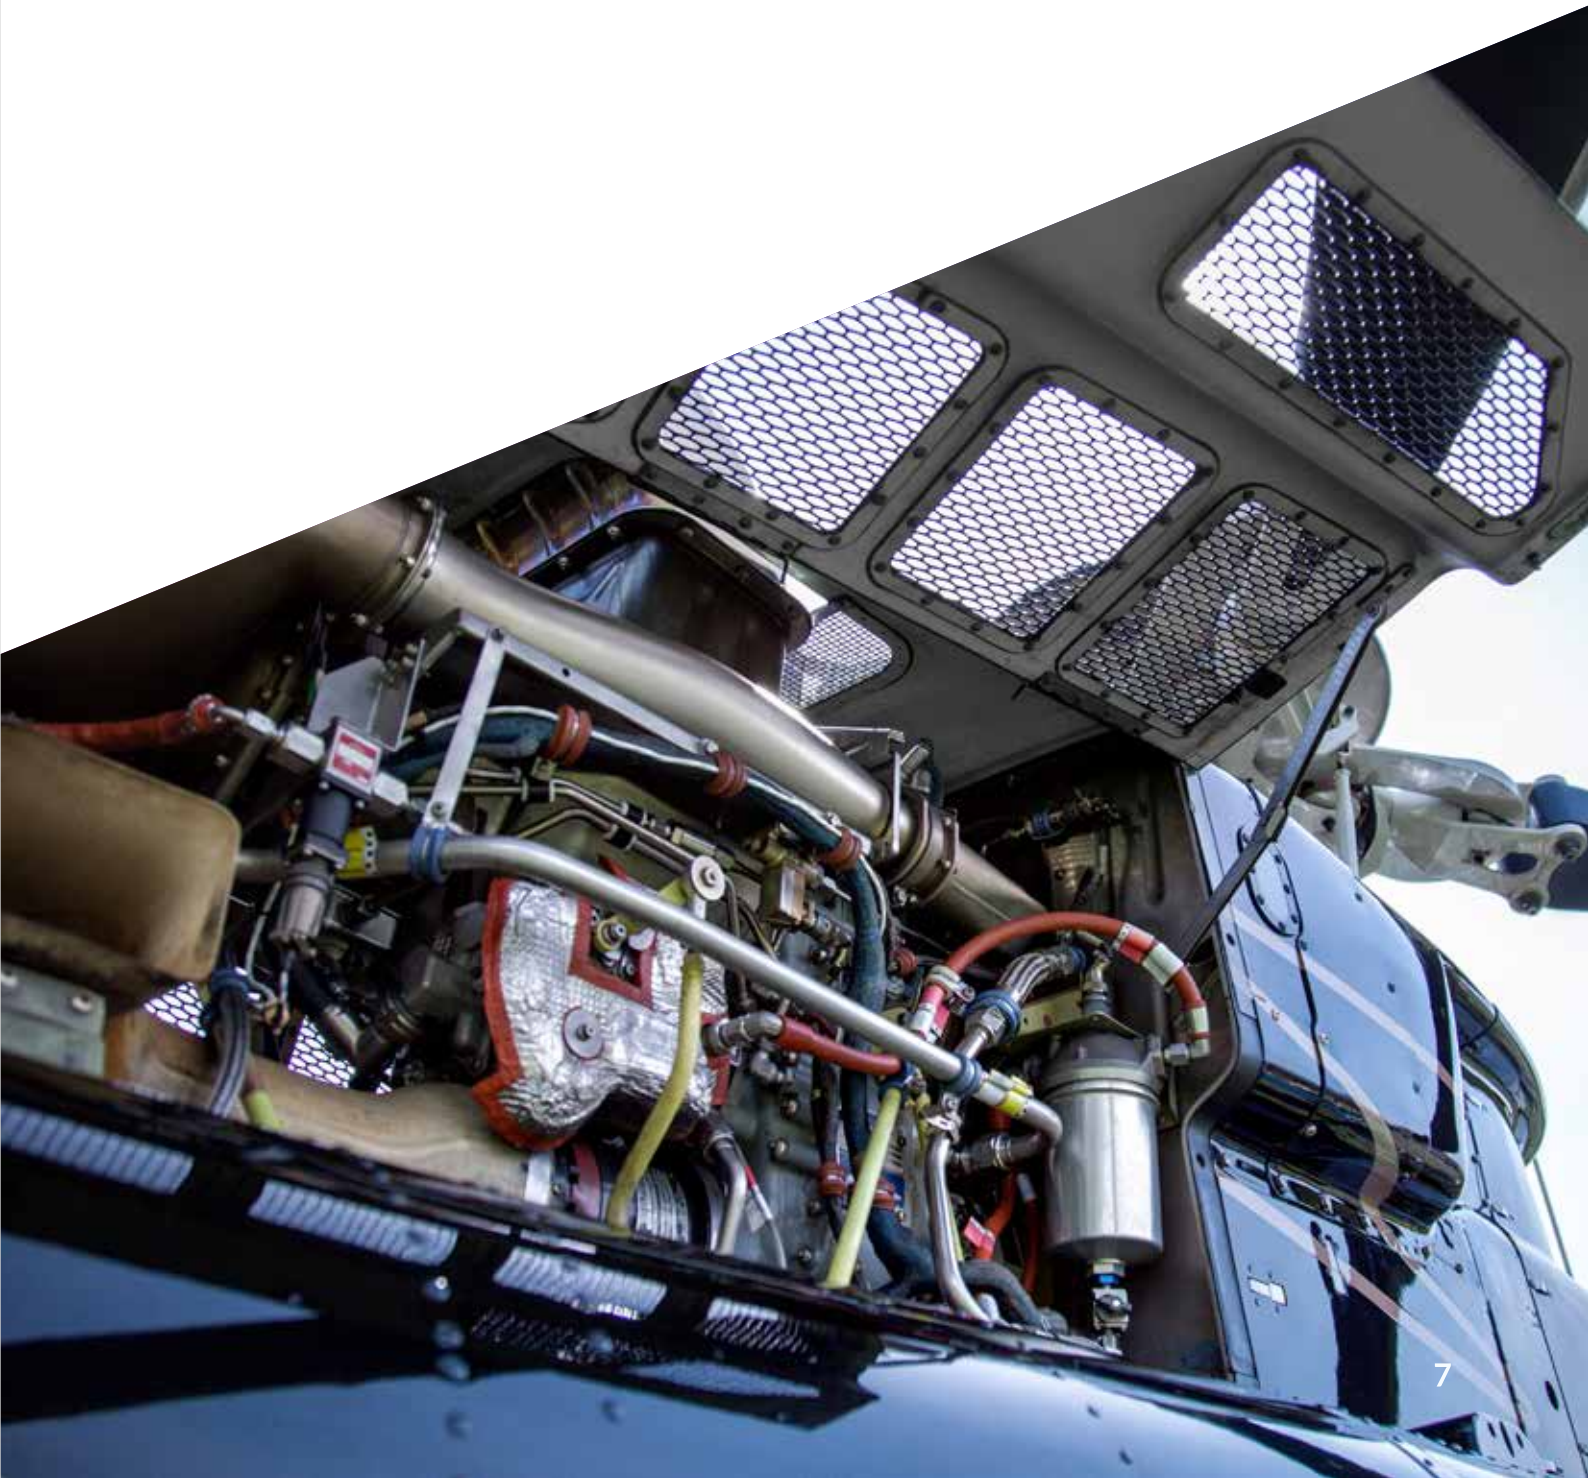

Fast-Forward: Enter Industry 4.0

Industry 4.0 and the Fourth Industrial Revolution, terms often used interchangeably, are having a massive impact on the direction of our manufacturing industries.

According to our friends at the World Economic Forum:

First Industrial Revolution		The First Industrial Revolution used water and steam power to mechanize production.
Second Industrial Revolution		The Second used electric power to create mass production.
Third Industrial Revolution	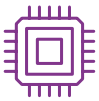	The Third used electronics and information technology to automate production.
Fourth Industrial Revolution		Now a Fourth Industrial is building on the Third, the digital revolution that has been occurring since the middle of the last century. It's characterized by a fusion of technologies that is blurring the lines between the physical, digital, and biological spheres.

The advent of Industry 4.0 has changed everything. Klaus Schwab, who coined the term "the Fourth Industrial Revolution", describes the speed of current technological breakthroughs as having "no historical precedent" and that it is "disrupting almost every industry in every country".

This has spurred significant improvements in sensor, network, data acquisition, and storage technologies. This unprecedented disruption is further augmented by the recent progress in Artificial Intelligence (AI) made possible by exponential increases in computing power and access to a previously unimaginable wealth of data. The result is that advances in PdM are increasingly being applied to wider industry, not just solely to highly critical assets in industries such as Aerospace and Oil & Gas.

Industry 4.0 can therefore also be viewed as a technology, making powerful technologies available to more than just those with deep pockets. This has, unfortunately, come with its own difficulties.

Industry 4.0 can therefore also be viewed as a technology, making powerful technologies available to more than just those with deep pockets. This has, unfortunately, come with its own difficulties.

So why has PdM success been so hard to achieve for many?

Pilot purgatory is one of industrial IoT's biggest concerns. IoT World Today.

Looking to capitalize on these benefits, many PdM vendors have rushed into an increasingly crowded marketplace. They often appear to have little appreciation of what is a very unique domain.

Some simply try to 'super-charge' legacy monitoring tools without taking advantage of the automation and advanced features that an AI system can deliver. Others apply conventional data science approaches to a problem space that is distinctly different from other "Big Data" challenges as found in finance, retail, and health.

Indeed, IDC states that 50 percent of all AI projects fail mostly due to data challenges. More alarmingly, 80 percent of AI initiatives are project-based and are struck at the PoC stage, unable to scale out. Similarly, a McKinsey survey found that less than 30 percent of pilots are starting to scale. 84 percent of companies were stuck in pilot mode for over a year and 28 percent for over two years.

There is, therefore, a wide lack of understanding - both by manufacturers and by vendors - as to how to deliver a successful PdM system. Without that understanding, many of these new solutions are destined to languish in pilot purgatory.

So, despite the advances and energy behind the desire to deploy and benefit from PdM, few businesses have achieved genuine success at scale.

Why? We repeatedly witness the same three key mistakes that vendors and buyers make:

1. Although the term PdM may be relatively new (and has certainly captured the imagination), the concept itself isn't. As pointed out earlier, it's been successfully employed to optimize and improve operational efficiency in the Aerospace industry alone for at least 30 years. Techniques such as condition monitoring, fleet management, diagnostics, maintenance credits, and prognostics have reached a mature state. The key issue has always been scaling these techniques beyond critical machines, which has made them unviable in the wider manufacturing space. That's where AI can help, by automating core activities. However, it cannot be at the expense of ignoring fundamental principles from the last 30 years.

2. PdM is far more than just a matter of data science. It's not, as some believe, a Big Data problem in which there are millions of data points and labels on which to train models. The fact is, failures in a factory environment are actually quite rare (Gartner refers to this type of challenge as Little Data). And the environment itself is highly dynamic and noisy, with a range of variables such as machine maintenance, different behaviors

of machine operators, changes in ambient temperature and humidity, and different production speeds. Moreover, every machine is unique, even if it's the same brand and model doing the same job. Yet, time and time again, this classic Big Data approach is what we see many organizations attempt.

3. A PdM system needs to carefully address the needs of maintenance professionals. The user experience must therefore be spot on. Otherwise, there's a risk that all the information, intelligence, and insight will be disregarded. All of an organization's investment in technology and data will count for nothing if its output isn't clearly and succinctly communicated.

We've seen so many in-house PdM projects fail where data scientists have created some algorithms, connected them to an off-the-shelf dashboarding tool, and are then surprised when user engagement is lacking. The unsatisfactory user experience offered by their off-the-shelf dashboard has failed to engage or meet the needs of the system's end-user - the maintenance professional - and the system has simply been ignored.

At Senseye, we've spent over 150 person-years of research and development time exclusively on PdM. This, combined with our experience of working in the Aerospace industry over the last 30 years, means we know what works and what doesn't. It's why we know that building a successful PdM system isn't as simple as just hiring some data scientists and setting them to work for six months.

Read the next section to learn how we've applied these lessons to everything we do.

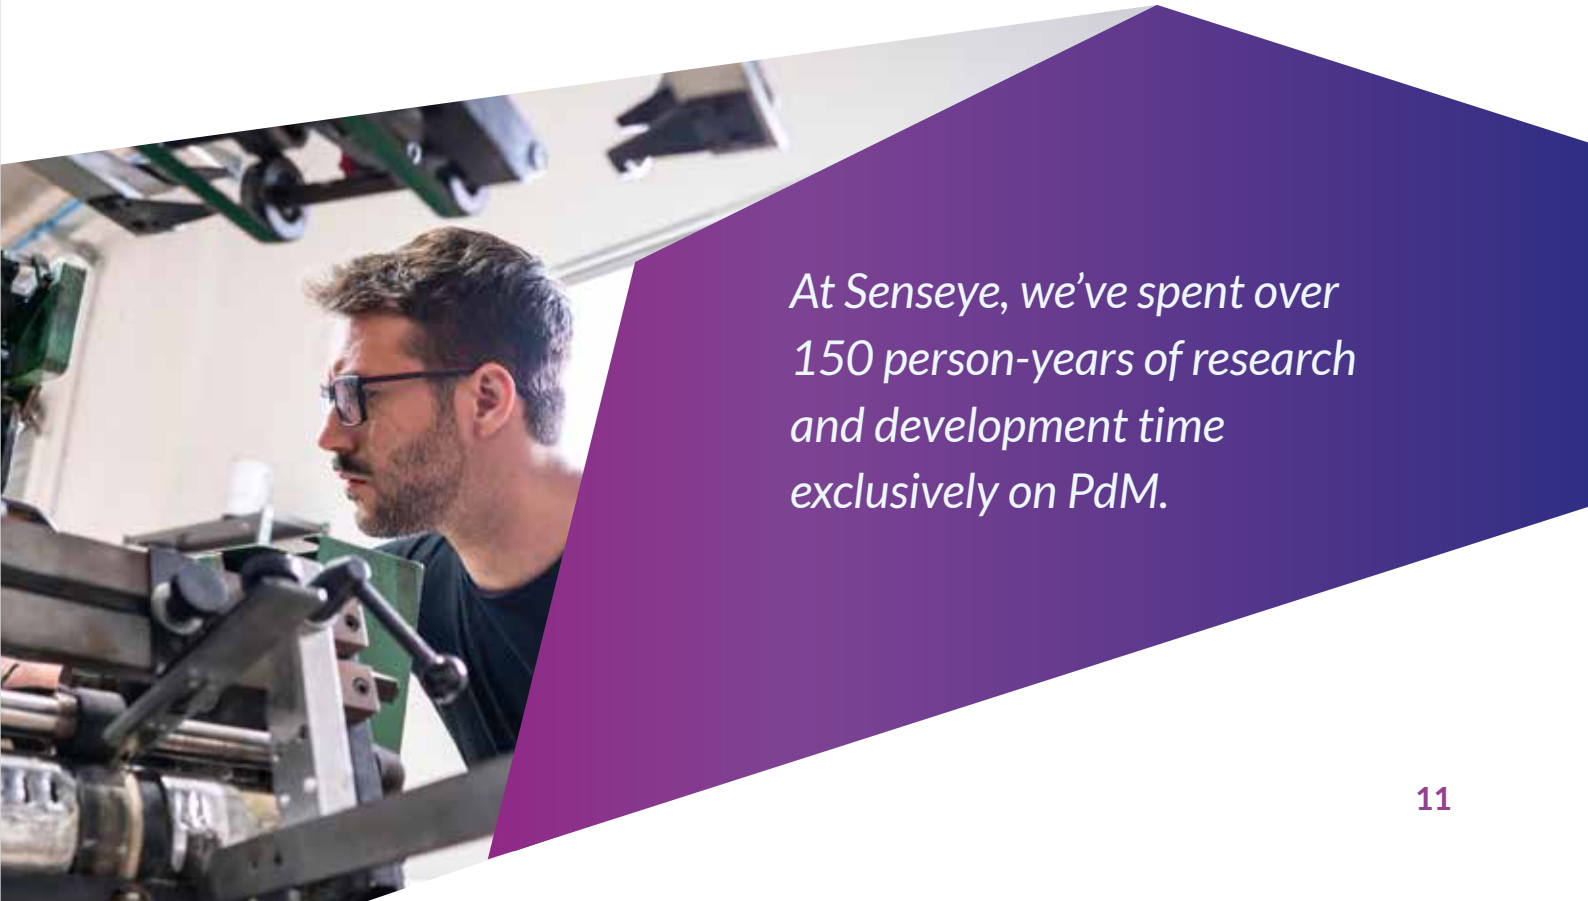

At Senseye, we've spent over 150 person-years of research and development time exclusively on PdM.

PART 2

Succeeding at Predictive Maintenance at Scale

So, what are our key takeaways?

We've learned a lot about deploying predictive maintenance and related technologies across a variety of sectors. It has been - and continues to be - a learning experience, especially as we encounter different sectors and different levels of customer maturity. Predictive Maintenance is hard.

It's important to appreciate not only what's needed for PdM to work well, but why it matters. We've learned that it can be hard to explain the value of PdM, even when it's delivering a substantial return on investment. After all, implementing a PdM strategy requires a whole business transformation. Involving a shift in mindset from everyone from the board to the shop floor, it's a move that can't be undertaken lightly, and can often require some serious justification.

Don't treat Predictive Maintenance as just another Data Science use case!

That PdM is not a typical Big Data problem is something that's become increasingly apparent as we engage with new customers that share their experiences about previous failed projects. Big Data solutions work best in context-rich environments - which are severely lacking in PdM. Moreover, many machine failure modes are subtle in nature and, if they don't follow a clear pattern, can be difficult for established machine learning algorithms to accurately predict.

The number of times we're asked how accurate our algorithms are (not even constraining it to a specific machine type or failure mode), with expectations of an answer in the high 90%'s, is telling. While it's a reasonable question to be asked when deploying many Data Science solutions, it exposes a fundamental misunderstanding of the problem.

It's true that, in a lab environment with some high quality and curated test data, you will achieve such percentages for specific machines and failure modes. But, if you then factor in that each machine and each instance of failure mode is unique, along with the severe data quality issues faced in a live environment, the type of sensor data available (you can't expect the same level of accuracy when using temperature data versus vibration), the highly dynamic nature of factories, and the dearth of available crucial context information, you'll quickly appreciate the inappropriateness of such a generalized question.

You should ask tough questions of anyone that tries to attach a single accuracy figure to their solution.

Know your users!

We appreciate that busy maintenance teams will typically have only a few minutes at the start of each shift to identify which among their thousands of machines most need their attention. Keeping the design of the software and its output incredibly simple and intuitive saves valuable time.

Maintenance professionals do not want to spend hours diagnosing graphs and mining raw data for valuable insights.

Yet, many vendors supply traditional dashboarding or business intelligence type interfaces that are generic and do not factor in the uniqueness of the user's workflow.

Perhaps even more importantly, maintenance professionals are a valuable source of knowledge and experience and a successful PdM solution must harness and engage with that experience.

At Senseye, we've always aimed at following an approach to designing beautiful products based on the principle of simplicity and the seamless integration of analytics and the user interface. That's why Senseye PdM presents information to users in the form of a prioritized list, sorted by the Attention Engine, a proprietary algorithm that estimates an Attention Index® for each of an organization's machines, based on machine data, maintenance data, and operator behavior.

Know your users' place on the digital journey

We've learned, too, that businesses - and often areas within the same business - will be at different levels of data and cultural readiness. At the lower end, a company may do little more than carry out periodic route-based condition monitoring checks while, at the other, a company will combine robust, automated condition monitoring with the right PdM solution to give accurate predictions of time-to-failure and mode of failure for each of its machines. Each type of customer needs a completely different support and deployment package.

Most businesses will be somewhere between the two extremes. Increasing their maturity level relies on a greater understanding by their management and buy-in from their IT team.

In the next section, we'll explore how all of this experience and understanding has turned Senseye PdM into the leading PdM product on the market.

PART 3

150 person-years of R&D exposed

(Introduction)

Achieving a Scalable PdM solution

Senseye PdM is the result of over 150 person-years of research and development, leveraging skills from industry specialists, condition monitoring experts, and mechanical engineers, in addition to a team of leading data scientists. Senseye also draws upon its founders' extensive knowledge and experience in predictive maintenance developed during their time in the aerospace and defense industries - industries that lead the world in the culture of safety, maintenance practices, and predictive maintenance technologies.

This combination of heritage and deep experience, together with advanced and proprietary AI and machine learning, results in a truly unique product, which is further bolstered by the company's capability to provide deep domain support and consultancy.

This means that, unlike other solutions, Senseye PdM doesn't require ongoing manual reviews of sensor data or the development of custom models for each type of machine. Instead, high fidelity models are constructed automatically and uniquely for every machine, without the need for user intervention.

This results in predictive maintenance that can be rapidly applied to all machines, including those of lower criticality, thereby easily covering the entire balance of plant.

The system uses a range of machine learning, AI, statistical modeling, prognostic, and data mining approaches to automatically analyze the data, while canceling out the huge quantity of environmental noise, to establish machine health and a forecast of degradation. This is fed into Senseye's unique Attention Engine - a proprietary algorithm that estimates an Attention Index® for each machine. If the Attention Index® is sufficiently high, the Attention Engine creates a Case to carefully direct the user's attention to the machine in question. All this is delivered in an easily understandable manner so maintenance teams can react with increased speed and accuracy.

High fidelity models are constructed automatically and uniquely for every machine, without the need for user intervention.

Our Guiding Philosophy

Everything we do is supported by our underlying philosophy of how PdM needs to function. At its most fundamental, Senseye PdM is a decision support system that helps maintenance professionals take care of their machines. There are three principles that have guided the design of the application and the analytics that underpins it:

Guide attention:
The aim of the application is to focus the user's attention on machines that need it. The analytics is designed to support this aim.

Focus on meaning:
Analytics systems require good data, plentiful context, and high-quality algorithms. In the environments where we work, context is a scarce resource and is the thing that places the fundamental limit on how well machine learning can perform. As such, we focus on making our analytics work with the minimum amount of context, using the context we do have as efficiently as possible, and using the application to gather as much additional context from the user as possible.

Model the user:
We make optimal use of scarce context by focusing the analytics on the user. Instead of just predicting health, as a property of the machine, we predict user interest, a property of the user. And we can do this effectively because we have access to the user and can use their feedback to optimize our predictions.

Senseye is unique in applying these three principles to the PdM problem.

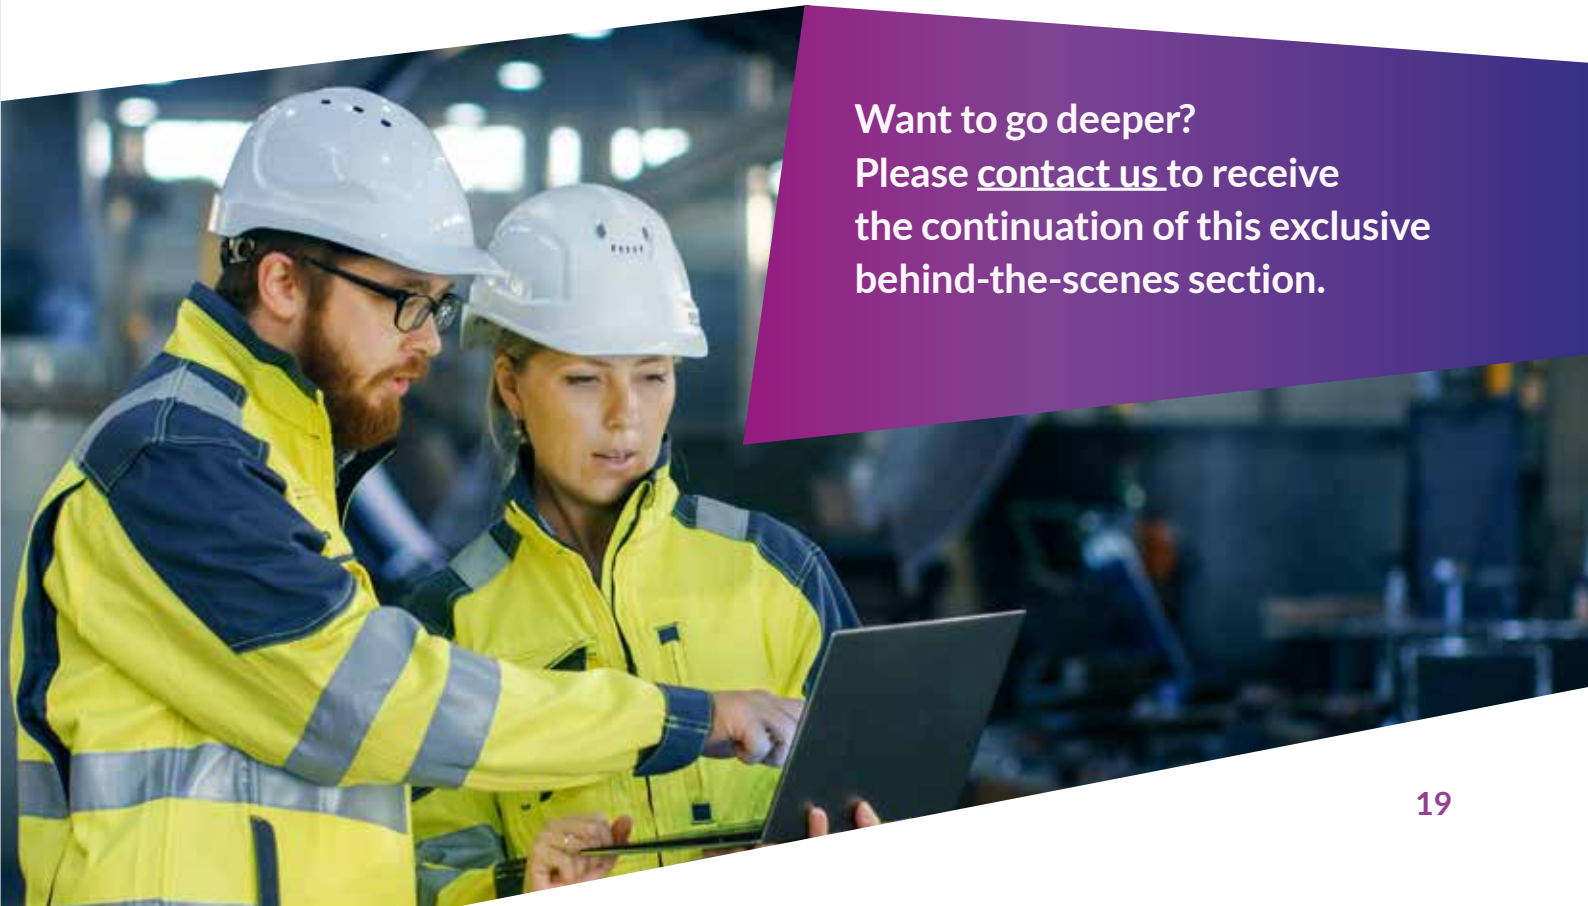

Want to go deeper?
Please contact us to receive the continuation of this exclusive behind-the-scenes section.

PART 4

What does the future hold?

The global predictive maintenance market is maturing rapidly, largely driven by a rise in the need to improve the uptime of equipment. Worth an estimated $4bn in 2020, its value is forecast to reach over $12bn by 2025. So what needs to happen to achieve that level of growth?

Clearly the market remains fragmented, with difficulty in separating the wheat from the chaff. This is going to take some time to resolve itself. We believe that solutions that take an integrated, user-centric, and holistic perspective will prevail.

A solution that's integrated with respect to broad sources of information, user-centric in the sense of having 'conversations' between users and the system to capture key knowledge and experience, and holistic to algorithms and models; there is no single 'master algorithm' so there will always be a need for data scientists and custom models. The future belongs to a system that can elegantly integrate all of this variety into a single product.

What else? We'll see an acceleration of the ownership and responsibility of deploying PdM from OT teams to IT teams. This is crucial in ensuring a successful scale-up and integration in the company's wider digital eco-system.

There'll be a continued acceptance and shift to cloud computing, supported by edge devices to carry out the heavy lifting on the shop floor. This is key to gaining rapid scale and truly benefiting from a pan-business view of operations. On-premise, on the other hand, has a more limited role to play in the future, and will be largely restricted to more sensitive sectors.

We'll also see the continued introduction of more diverse and lower-cost sensors, especially with regard to vibration, although other sources, such as current, will be harnessed more often than they are today.

PdM is, of course, only one building block to a company's wider Industry 4.0 initiative. As such, it's important to think big: IoT platforms such as PTC's Thingworx, Siemens Mindsphere, and Schneider Electric's EcoStruxure will play an increasingly pivotal role in enabling this.

But growth may yet be hampered by a dearth of skilled staff, concerns around data privacy and security, and difficulties in deploying at scale. Many businesses are still nervous and confused by the marketing hype and their companies' own internal direction around PdM. It is these issues that will take longer to resolve but have the biggest impact on how steep that growth curve will be.

EcoStruxure Innovation At Every Level

thingworx

MindSphere

IoT platforms

Predictive Maintenance may not be new, but it's continuously and rapidly evolving. It's underpinning a fundamental shift in the drive to become ever more efficient. It's the darling of Industry 4.0.

As the market leader, Senseye is proud to be a key player in directing this exciting evolution.

We believe that solutions that take an integrated, user-centric, and holistic perspective will prevail.

That's integrated with respect to broad sources of information, user-centric in the sense of having 'conversations' between users and the system to capture key knowledge and experience, and holistic to algorithms and models; there is no single 'master algorithm' so there will always be a need for data scientists and custom models.

The future belongs to a system that can elegantly integrate all of this variety into a single product.

senseye

UK (HQ)
Epsilon House
Enterprise Road
Southampton Science Park
Southampton
SO16 7NS
United Kingdom

+44 (0) 845 838 8615

UK
90 Long Acre
Covent Garden
London
WC2E 9RZ
United Kingdom

+44 (0) 845 838 8615

sales@senseye.io

USA
100 Madison Street
Unit 308
Nashville
Tennessee
37208
USA

+1 415 523 0447

Please visit www.senseye.io to see our other regional office locations.

Follow us:

www.senseye.io