Predictive Maintenance Beyond Prediction of Failures

Francesco Mari
Vice President Business Innovation – Internet of Things
Products and Innovation – SAP SE
September 2016

Prediction of Rare Events is Very Difficult Indeed
Chapter 1: Fighting Overfitting

Trying to predict which animals are mammals with a dataset with too few positive cases...

Name	Body Temperature	Gives Birth	Four- Legged	Hibernates	Class Label
salamander	cold-blooded	no	yes	yes	no
guppy	cold-blooded	yes	no	no	no
eagle	warm-blooded	no	no	no	no
poorwill	warm-blooded	no	no	yes	no
platypus	warm-blooded	no	yes	yes	yes

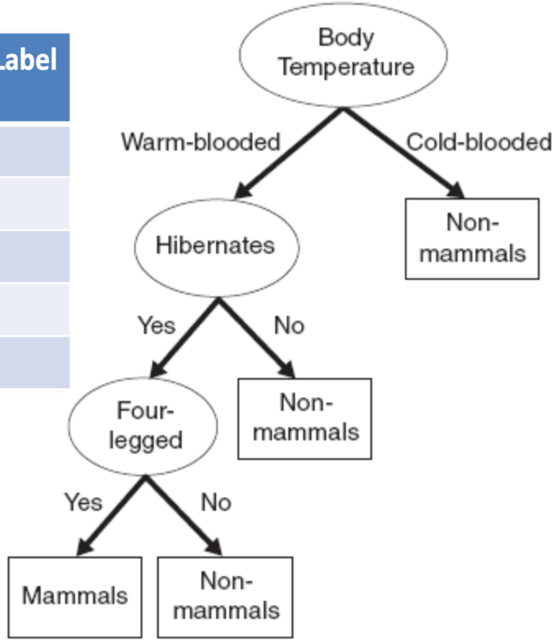

…one could end up with a model indicating that humans are not mammals!

This is a well-known problem called **overfitting** (confusion between noise and meaningful relationships)

Prediction of rare events is an elusive field for data science – we still can't predict earthquakes, and still use the Gutenberg – Richter power distribution law formalized 60 years ago

Prediction of Rare Events is Very Difficult Indeed
Chapter 2: Bayes' Revenge

And even if you can create an unrealistically good model…

Let's imagine that:

- You have 100 machines operating 24*365
- You are trying to predict a type of failure that happened 50 times last year - discretizing time in hours this means MBTF = 17K
- You were able to build a model with 90% True positive rate 99% accuracy

Looks like you are in a great shape…

		Predicted		
		YES	NO	
Actual	YES	45	5	**50**
	NO	8760	867190	**875950**
		8805	**867195**	**876000**

Yes, you could prevent 45 out of the 50 failures before they happen

But the real failures are drowned among 8805 positive predictions (*and keep in mind that in the real world you don't have models with this level of accuracy…)*

How can you rely on these results?

The Core Tenet of Predictive Maintenance

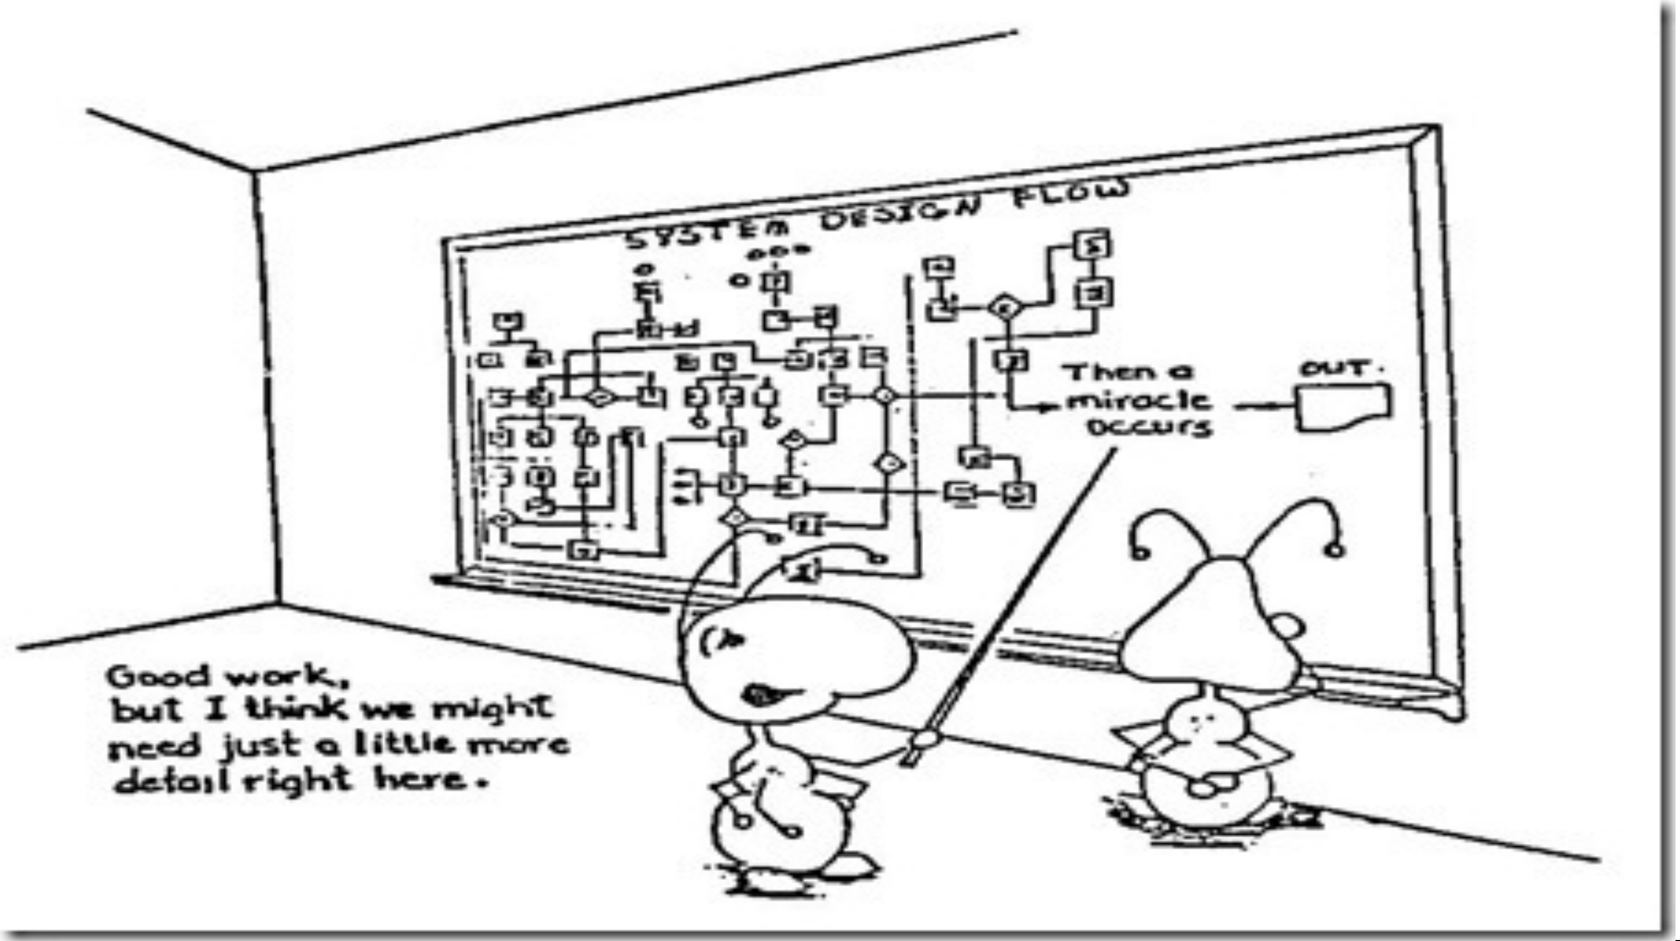

The Nature of the Asset Makes a Difference

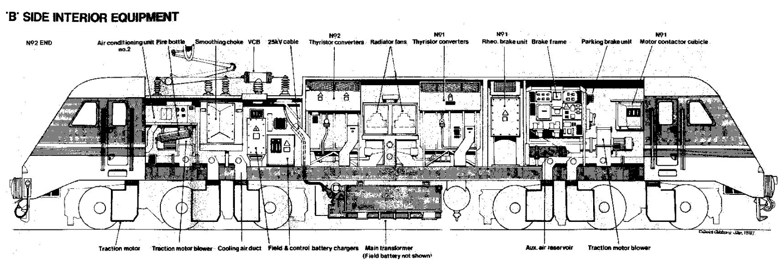

← Simple ———————————————— Complex →

Low	Cost of the asset	*High*
Limited	Number of interdependent components	*Very high*
High	Possibility of redundancy	*Limited / Close to nothing*
Manageable	Cost of failure	*Massive*
Simple and uniform	Operational conditions	*Extremely diverse*
Simple and deterministic	Failure modes	*Complex and stochastic*
Limited	**Value of predictions**	***Very high***
Limited	**Complexity of prediction**	***Very high***

Let’s Score This Model

		Predicted		
		YES	NO	
Actual	YES	45	5	**50**
	NO	8760	867190	**875950**
		8805	**867195**	**876000**

Accuracy	99%	(TP + TN) / Total
Misclassification	1%	(FP + FN) / Total
True Positive Rate	90%	TP / Actual Yes
False Positive Rate	1%	FP / Actual No
Specificity	99%	TN / Actual No
Precision	1%	TP / Predicted Yes
Prevalence	0,0%	Actual Yes / Total
Negative Predictive Value	100,00%	TN/Predicted No

- While Precision is typically very low, the Negative Predictive Value is close to 100%, reflecting an almost perfect homogeneity of results for the negative outcomes of the model // **When the model says “no”, failures almost never happen**
- *What to do with this result?*

Strategies to Cope with Rare Events

Earthquakes and Maintenance of Complex Assets Diverge

	Earthquakes	Maintenance of complex assets
Actual = YES	Emergency management	Reactive and corrective maintenance
Actual = NO	//	Preventative maintenance

- Due to the extreme costs and penalties associated with failures of complex assets, significant resources are often invested in preventative maintenance aimed at structurally reducing the risk of breakdowns
- Traditional planning for preventative maintenance does not always consider neither the different operational conditions of the assets, nor their actual health status
- As a consequence, **complex assets are often over-maintained**. The same predictive models that fail to reliably predict failures can on the contrary effectively address this issue starting from the very high Negative Predictive Value

Anomaly Detection Analysis on Turnouts

- 19 turnouts analyzed over a period of around 17 months
- Diagnostic equipment records tension and current of the actions
- Over the period there were only 11 actual failures, corresponding to an MTBF of around 21,000 hours
- Turnouts are inspected on the average once a month

Data seem to indicate a situation of over-maintenance, driven by the strong priority of avoiding failures almost at any cost. This is typical of critical assets where failures come with big penalties

Analysis on Turnout 1

Thresholds on anomalies defined as a combination of severity and number of instances in a rolling period, and tuned to match violations and prevalence of failures, at least in terms of order of magnitude

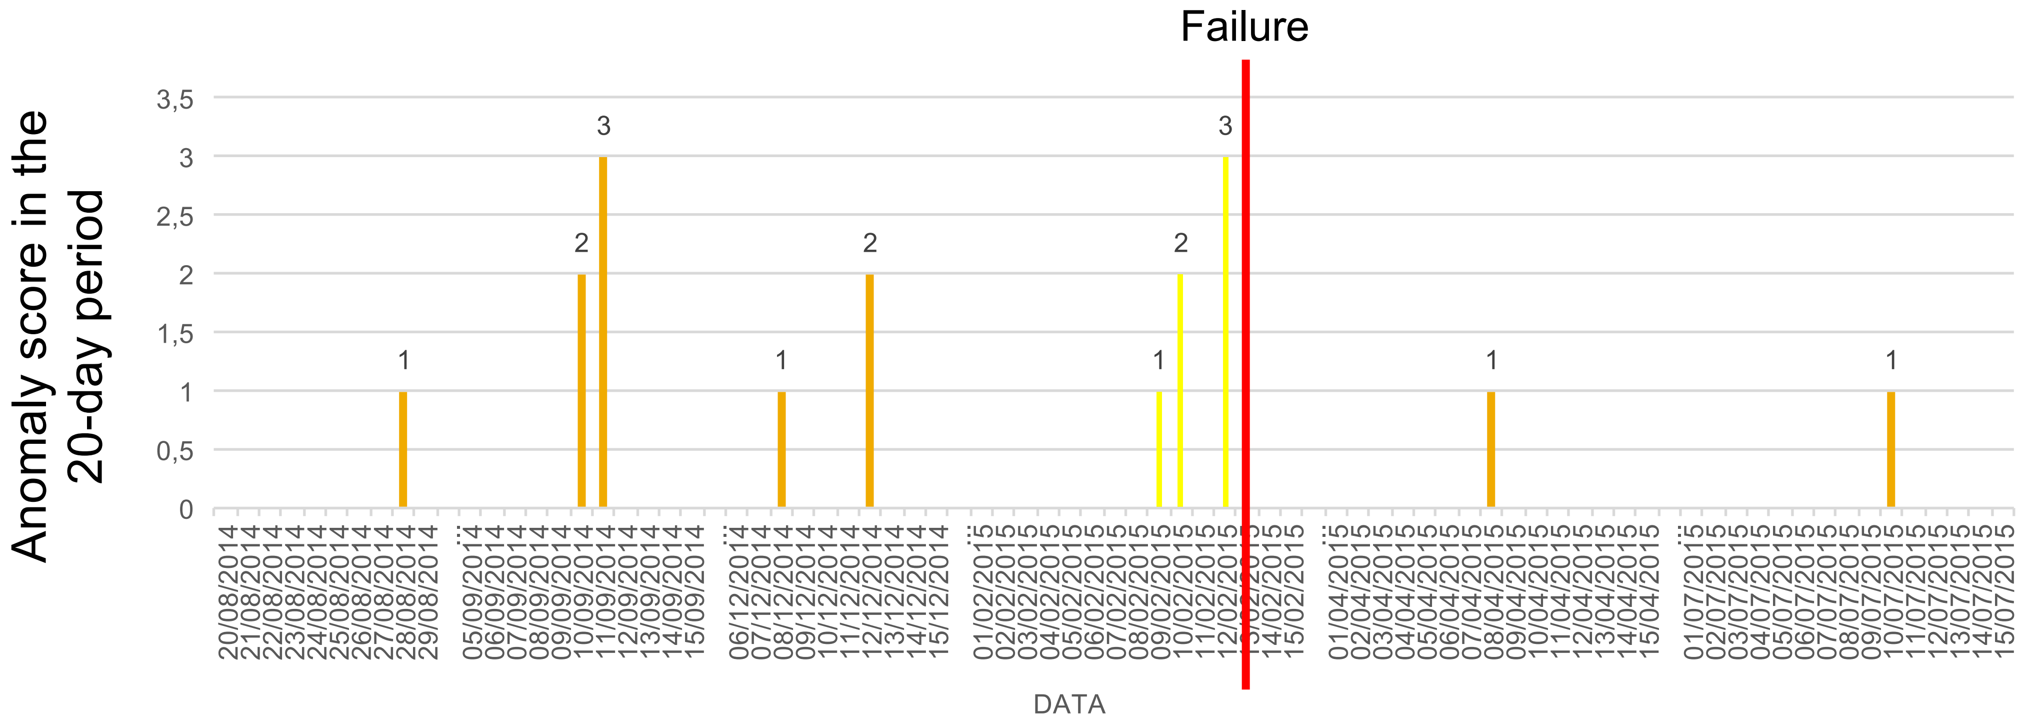

Results of Anomaly Detection Analysis on Turnouts

		Predicted		
		YES	NO	
Actual	YES	5	6	**11**
	NO	69	1276	**1345**
		74	**1282**	**1356**

Accuracy	94%	(TP + TN) / Total
Misclassification	6%	(FP + FN) / Total
True Positive Rate	45%	TP / Actual Yes - AKA Sensitivity / Recall
False Positive Rate	5%	FP / Actual No
Specificity	95%	TN / Actual No
Precision	7%	TP / Predicted Yes
Prevalence	0,8%	Actual Yes / Total
Negative Predictive Value	99,53%	TN/Predicted No

Data are analyzed in time series, but presented discretized by week

- 5 out of the 11 failures were predicted by the algorithm a few days in advance
- 6 failures were not predicted, best hypothesis is that around half could be identified with better diagnostic equipment, half is highly intractable from a data science perspective
- 69 false positives were also generated

→ Very high accuracy and specificity

→ Very low precision

→ Prevalence (number of weeks in which a failure happens / total number of weeks analyzed) is below 1%

In rare event prediction case, achieving high accuracy, precision and sensitivity at the same time is extremely challenging

The real actionable insight often sits in the Negative Predictive Value

Turning the Insight into Action, and Outcome

Due to low precision, Condition-based / Predictive Maintenance tend to be additive strategies, to be combined with the existing approaches dramatically improving the business performance (risks vs. cost balance)

As-is Maintenance Policy	
Cost of a visit	1
Penalty for a failure	20
Failures over the period	11
As-is cost over 17 months	535

New Policy with Application of Anomaly Detection	
Frequency of visits for turnouts without anomalies	1.5 months
Cost of emergency visits	2
Failures over the period	8
To-be cost over 17 months	444

- Significant reduction in terms of failures, around 30% as a starting point, with the opportunity of increasing it with better diagnostic equipment and next iterations in the data science algorithms
- Sizable improvement in the overall efficiency of the maintenance system, with potential savings exceeding 15% of resources invested, with further opportunities to be analyzed

The Problem with Preventative Maintenance Planning

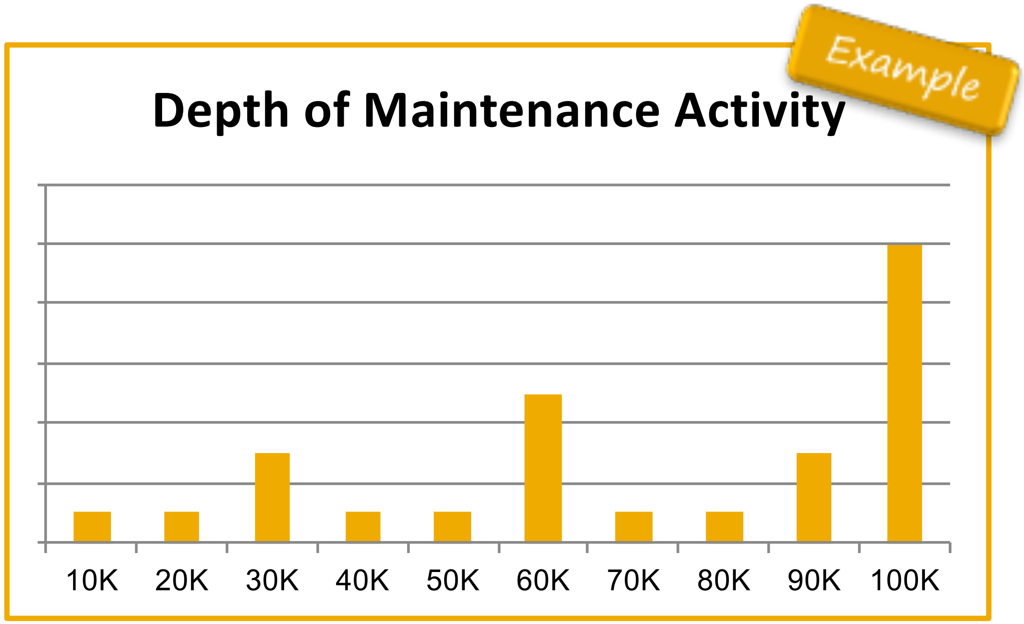

- Work done by simple assets can be described by a simple indicator such as time of operations or units of throughput, but if we consider a complex machine like a train, **the specific operational conditions have a huge impact on how its various components wear out**
- These operational conditions are not known a priori by the manufacturer of the machine, and can change dramatically in the day-by-day work – **this creates the need of operating under extremely conservative assumptions**
- The only real value of these plans is that they are logistically very simple to execute and don't require much information

Increase Planning Efficiency and Relevance With Life and Health Indicators

Life

Measure as precisely as possible the expected wear of the part by counting and projecting a set of relevant parameter (e.g. cycles, hours of operation, kilometers, energy etc.)

Maintenance is performed when predefined thresholds are reached

Health

Takes into account the actual status of operation by measuring physical parameters (e.g. closing time for a door, temperatures of cooling systems) or relevant combinations of indicators

Maintenance is performed when the parameter goes out of the normal range, indicating a probable deterioration of condition

Life Indicators in Action: Braking System

Objective and Approach

- Calculation of the life indicator → **energy dissipation** by friction braking systems, with separated analysis for locomotives and coaches
- Development and test of calculation algorithms for all the possible cases identified

Results Achieved

- Monitoring of the effective usage and level of wear-out for every single component of the braking system against the risk thresholds identified

Pressione in Condotta Generale e Cilindro Freno

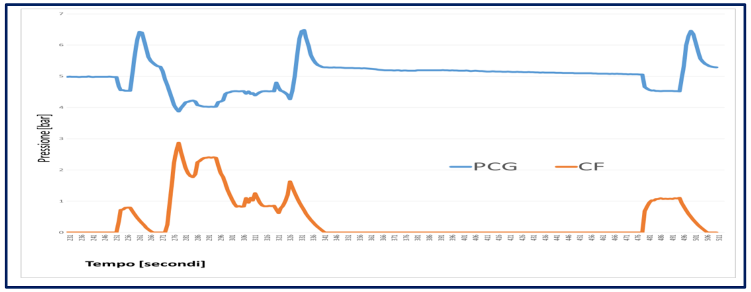

Energia dissipata cumulata per le carrozze

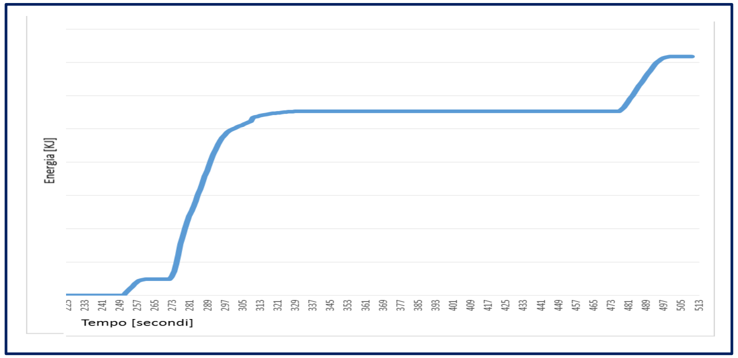

Life Indicators in Action: Braking System

MATERIAL_ID	CALENDAR_DATE	DEPARTURE	ARRIVAL	Measures: ENERGY/KM	KM	ENERGY
E464003	02/05/2015	CATANIA CENTRALE	MESSINA CENT.	0,03	92.597,35	2.912,53
	04/05/2015	CATANIA CENTRALE	MESSINA CENT.	0,05	92.839,25	4.221,51
	06/05/2015	CATANIA CENTRALE	MESSINA CENT.	0,06	92.713,60	5.382,53
	07/05/2015	CATANIA CENTRALE	MESSINA CENT.	0,04	93.789,17	4.022,59

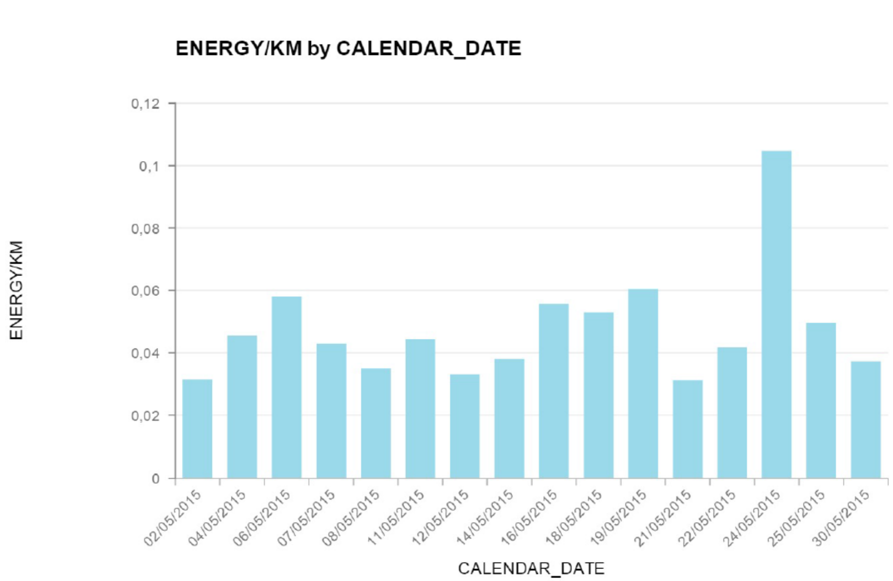

COMPONENT_ID
- E464003

MATERIAL_ID
- E464003

BRAKE_PAD_ID
- P-0002-02

DEPARTURE
- CATANIA CENTRALE

ARRIVAL
- MESSINA CENT.

- High variability of the energy dissipated per km clearly indicates that distance is not a good indicator of consumption for braking systems
- Comparison between traditional indicators such as km and more precise life indicators highlights the significant opportunity for optimization of maintenance operations

Life Indicators in Action: Braking System

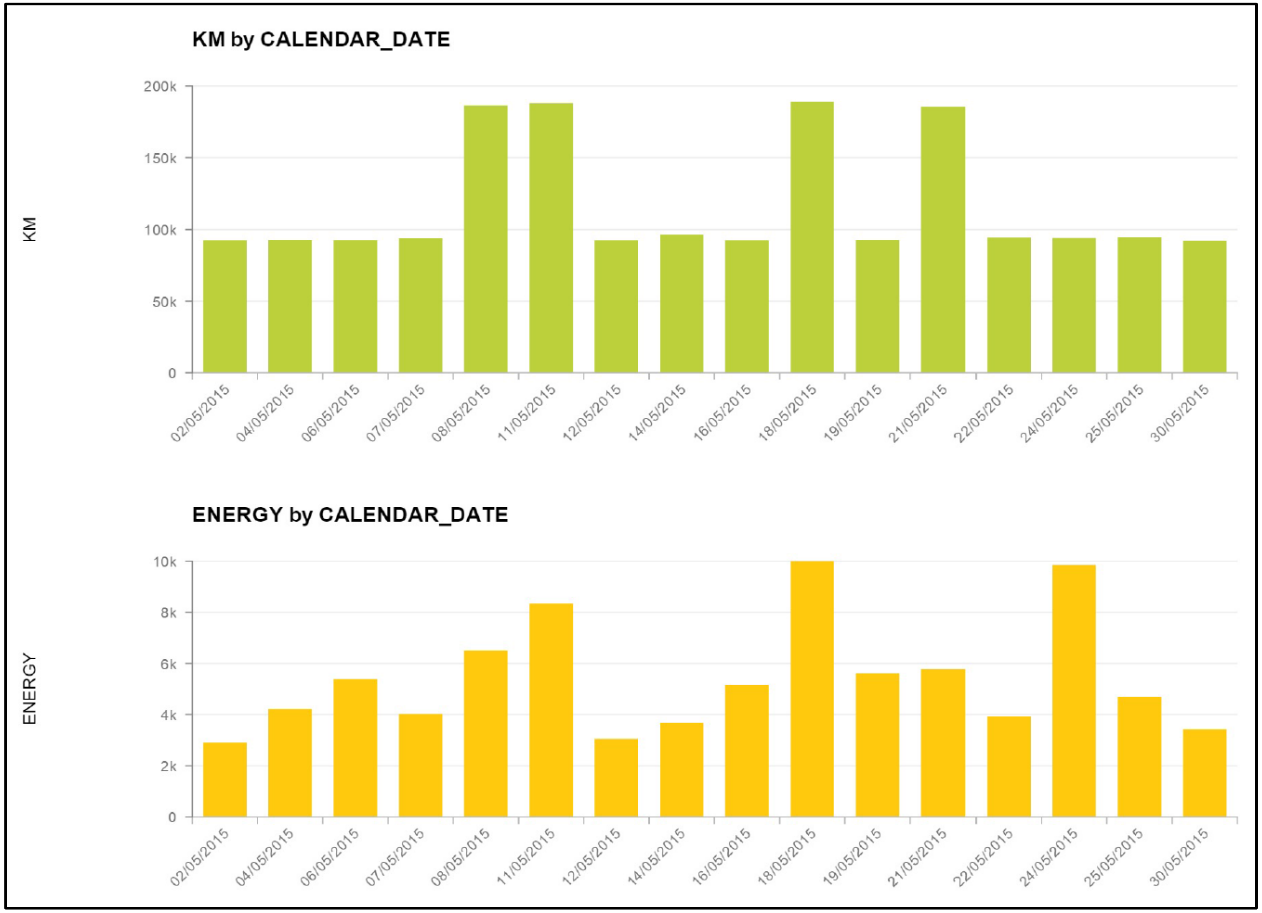

- High variability of the energy dissipated per km clearly indicates that distance is not a good indicator of consumption for braking systems
- Comparison between traditional indicators such as km and more precise life indicators highlights the significant opportunity for optimization of maintenance operations

What a Difference a Good Indicator Can Do!

Measuring the actual consumption of components when maintained

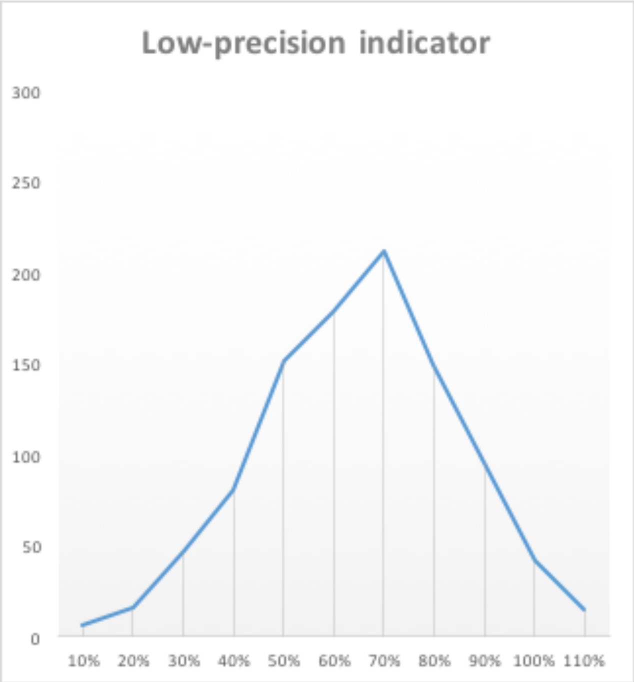

Average consumption	65%
Standard deviation	20%
Number of activities	1000

Moving from a low-precision (such as distance) to a high-precision indicator (e.g.: energy dissipated) allows to maintain parts when more consumed.
The two solutions here illustrated have the same maintenance value, and demonstrate a **huge opportunity for savings without impact on safety and reliability of assets**

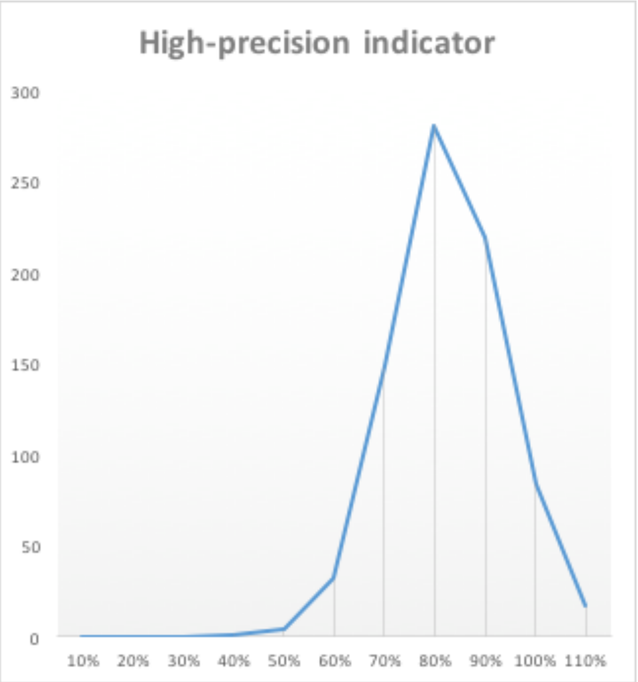

Average consumption	83%
Standard deviation	11%
Number of activities	788
Saving	**21%**

Some Tentative Conclusions

- While maintenance, especially when looking at complex, critical assets, has been largely left behind in the last decades by technology-driven disruption, huge innovation opportunities are now open by the Internet of Things and specifically by the convergence of IT and OT
- “Pure” prediction of failures is viable only in restricted cases, and possibly mostly a distraction
- New maintenance strategies will not substitute the traditional corrective + preventative approach, but rather extend and complement, with huge opportunities of increased efficiency and effectiveness – more an evolution than a revolution
  - Put the maintenance money where it really matters
  - Don’t touch what works fine
  - Reduce over-maintenance of critical assets without increasing failures

Thank you

Contact information:

Francesco Mari
francesco.mari@sap.com