



Data Mining in Equipment Maintenance and Repair: Augmenting PM with AM

Exclusive Ore Inc.¹
PO Box 1024, Blue Bell, PA
(215) 643-3110

Equipment failure is expensive. Unplanned downtime of a commercial airplane, such as a Boeing 757, has been estimated to exceed \$300,000 per year for each plane². Maintenance of aircraft, chemical and petroleum processing plants, weapon systems, early warning systems, and FAA control systems are examples of life threatening potentials if a failure were to occur. Even when not life threatening, equipment failure frequently results in expensive down-time.

Preventative Maintenance (PM) has long been used to replace parts at scheduled intervals. The PM approach has been based primarily on lifetime estimates for particular parts, and then replacing those parts at scheduled intervals before they exceed their lifetime estimate.

Anticipatory Maintenance (AM) is a new concept in maintenance, developed by Exclusive Ore, and now available as XAAM™, XAffinity® Anticipatory Maintenance. XAAM augments traditional PM by using data mining to predict parts that are likely to fail. XAAM uses Exclusive Ore's XAffinity®, a data mining tool that finds affinities between repairs or affinities between trouble reports and subsequent repairs.

There are some similarities between Anticipatory Maintenance (AM) and Condition Monitoring (CM). Both attempt to predict failures before they occur. The principal difference is that CM requires the installation of probes and monitors, and associated data collection and transmission devices, to collect information on operating conditions. AM uses existing maintenance and repair records. It is therefore much less expensive to implement AM than it is to implement CM.

PM, AM and CM are complementary. Each method can be used exclusively or in conjunction with the others to reduce overall maintenance costs. In general, the payoff from AM and CM will be higher for equipment and industries where the costs associated with unplanned down-time are highest.

¹ To find out more about Exclusive Ore Inc., and its custom solutions to data warehouse, database and data mining problems, please go to www.exclusiveore.com or click [here](#).

² Casebank Technologies, Inc. <http://www.casebank.com/about/aerospace.asp>

The in-depth analysis afforded through affinity data mining can discover hard-to-find incident affinities that can reduce down time from failures and may save lives in the extreme case.

There are other uses for these affinities in addition to the money saving opportunity described above. For example, unexpected affinity between two repairs might present an opportunity to improve repair procedures. Maybe the first repair involves replacement of an access panel that causes contamination leading to the subsequent repair. The affinity analysis can't tell you that contamination is the cause for the failure, but it can generate alerts for engineers who can then further examine the situations that have been brought to their attention by the analysis.

A Brief Example of Cost Savings from AM

Consider the following situation, involving the classic b-school parts – widgets and frammys. Let's say that it costs \$1,000 to replace a widget, but if a widget breaks unexpectedly, we also incur downtime costs of \$49,000. So it is 50 times less expensive to replace this part if the equipment is already down for some other reason. Let's say we've also discovered that during any given year there is a 1% probability that the widget will break. Issues of expected product life aside, with these numbers it doesn't make sense to replace the widget whenever the equipment is down because the \$1,000 cost is higher than the \$500 expected value cost (1% of \$50,000) of an unexpected break during the next year.

An affinity analysis of maintenance logs might lead to the discovery that the probability of a widget failure is much higher after a frammy failure. In fact, let's say the analysis finds that this probability is 20%. Now we have an obvious money saving opportunity:

Whenever the equipment is down for a frammy failure, also replace the widget. This Anticipatory Maintenance has an expected savings of \$9,000 (20% of \$50,000 minus \$1,000 actual cost).

What an affinity analysis does is find sequences of repairs (e.g., frammy followed by widget) along with measures to indicate how interesting or important each sequence might be.

Each repair sequence, along with its measures, is called a *rule*. Key measures for rules are the *confidence*, or conditional probability of the subsequent repair within the context of the rule, the *expected* probability, the prior probability of the subsequent repair, and the *lift*, the ratio between the two. In the above example, the confidence is 20%, the expected probability is 1% and the lift is 20 (20% divided by 1%). Rules with the highest lift flag the cases with the biggest surprise factor – where the probability differs the most from the prior or expected probability. Now, whenever a frammy is being repaired you can look in the rules to find other candidates (there could be more than just the widget)

for simultaneous AM. You'd select those repairs that have a confidence that is higher than their cost ratio.

Finding Noteworthy Sequences

Affinity data mining of equipment maintenance and repair data finds sequences of repairs (e.g., repair "A" followed by repair "B"), sequences of repairs with trouble reports (e.g., repair "B" followed by trouble report 103) or sequences of trouble reports followed by repairs (e.g., trouble report 428 followed by repair "C") that are *noteworthy*. More complicated sequences can also be found. For example, XAffinity can find sequences involving three repairs, e.g., repairs "A" and "B" followed by repair "C".

What constitutes a noteworthy sequence? There are various characteristics that might make a sequence noteworthy. Here is an example.

Let's say that the prior probability of repair "B" is 2%. (*Prior probability* is a term used by statisticians; you could think of it as the *normal* or *usual* probability, instead.) What this probability means is that two out of every 100 objects has had repair "B" during the past year (or the past month, or some other fixed time period). Well, if we find that repair "A" is followed by repair "B" 20% of the time, again within the same fixed time period, then we might have found a noteworthy sequence. There's something about repair "A" that multiplies the likelihood of repair "B" by a factor of 10. This multiplier is called the *lift*.

Affinity data mining with, for example, XAAM™ will find such noteworthy sequences for you, simply by analyzing a database of repair records.

Case Study - Locomotive Repair

The customer is a large locomotive manufacturer. Fault or incident reports are logged into a database as are the subsequent repair activities. The goal of this study was to find significant interactions between fault reports and repairs.

The customer posed specific questions as follows:

- Do certain faults lead to specific repairs?
- Do certain repairs produce subsequent faults?
- Are there repairs that lead to other repairs?

Locomotive Repair Data

The customer provided us with a sample of fault and repair data from 114 locomotives spanning one year. While this is sufficient data for a *proof of concept*, a study from which real decisions will be made would require more data.

An example of the data is shown below. The fault and repair data are intermixed into a single table; a specific row is either for a fault or a repair but not both. Each repair code has a corresponding description. Fault codes also have descriptions, but these descriptions were not provided to us. Although the term *fault* is used here, a better term might be *incident*, since fault implies the existence of a problem and many of the so-

called faults are merely observations of unusual conditions (e.g., a gauge approaching a warning level).

Unit ID	LocoType	Event Time	Repair Code	Repair Desc	Fault
106	2	7/22/2001 5:40:48 PM	<NULL>		7485
106	2	7/26/2001	4300	Toilet/Urinal/Electric Hopper Assembly	
106	2	8/1/2001 8:48:36 PM	<NULL>		7A50
106	2	8/2/2001 12:33:00 AM	<NULL>		76D5
106	2	9/2/2001 11:54:36 AM	<NULL>		7A2B
106	2	9/2/2001 7:15:36 PM	<NULL>		76D5
106	2	10/27/2001 10:42:00 PM	<NULL>		7154
106	2	10/27/2001 10:42:00 PM	<NULL>		7154
106	2	10/27/2001 10:42:00 PM	<NULL>		7154
106	2	10/27/2001 10:42:36 PM	<NULL>		7154
106	2	1/9/2002	2706	SS_X - Traction Motor Speed Sensor, X=#	
107	2	1/22/2001 4:01:48 PM	<NULL>		7140
107	2	2/1/2001	3740	Cooler, Lube Oil	
107	2	2/1/2001 12:07:48 PM	<NULL>		7A68
107	2	2/1/2001 12:09:36 PM	<NULL>		76D5
107	2	2/4/2001	0410	Scheduled Maintenance - RR Inspection	
107	2	2/16/2001 2:09:00 PM	<NULL>		7084
107	2	2/16/2001 2:12:36 PM	<NULL>		7084
107	2	2/16/2001 2:15:00 PM	<NULL>		7084
107	2	2/16/2001 2:16:48 PM	<NULL>		7084

In the data we see, for example, that locomotive #106 (Unit ID is the unique locomotive identifier) experienced fault code 7485 on July 22, 2001. Four days later, on July 26 this locomotive was subjected to repair code 4300 (Toilet/Urinal/Electric Hopper Assembly).

For any kind of sequencing analysis, one needs to identify the entity(s) to be analyzed, the order in which things happen and the event(s) that occurred. In this case the entities being analyzed are the locomotives, which are uniquely identified by Unit ID. The column that sequences the events and tells us the order in which things happened is called Event Time. Finally, one or more columns are needed to describe the events. In this case, because there are two columns that each describe the repair (Repair Code and Repair Desc) we can choose either the pair (Repair Code, Fault) or the pair (Repair Desc, Fault).

Note that it is not necessary to have both fault and repair data for a useful analysis. If we had only repair data, we could still analyze for interactions between repairs (something we will do below). And, if we had only fault data we could analyze for interactions between faults. Because both are available we can do three types of analyses:

- Between faults
- Between repairs
- Between faults and repairs

Although possible, between faults analysis was not done in this case study.

Faults to Subsequent Repairs - Highest Lift

The first analysis we performed examined the relationship between faults and subsequent repairs. Some of the rules generated are shown below. Each rule details a relationship

between a fault and a repair. To do this analysis, we used Repair Desc instead of Repair Code.

LeftSide	RightSide	Expected	Confidence	Lift	Support	LH Freq	
Fault = 7294	Repair Desc = 17FB135,17FB148 - System CPU	2.63%	100.00%	38.00	1.75%	1.75%	
Fault = 7292	Repair Desc = 17FB135,17FB148 - System CPU	2.63%	66.67%	25.33	1.75%	2.63%	
Fault = 7130	Repair Desc = Report No Dynamic Braking	2.63%	50.00%	19.00	1.75%	3.51%	
Fault = 729B	Repair Desc = 17FB135,17FB148 - System CPU	2.63%	50.00%	19.00	1.75%	3.51%	
Fault = 7400	Repair Desc = 17FB135,17FB148 - System CPU	2.63%	40.00%	15.20	1.75%	4.39%	
Fault = 7C34	Repair Desc = Camshaft Assemblies - Eng	3.51%	50.00%	14.25	1.75%	3.51%	
Fault = 7C31	Repair Desc = Camshaft Assemblies - Eng	3.51%	50.00%	14.25	1.75%	3.51%	
Fault = 7A40	Repair Desc = Coalescer, Crankcase	2.63%	33.33%	12.67	1.75%	5.26%	
Fault = 76D6	Repair Desc = Failure on Line	1.75%	22.22%	12.67	1.75%	7.89%	
Fault = 7A42	Repair Desc = 17FB152, 17FB158, 17FB163	1.75%	20.00%	11.40	1.75%	8.77%	
Fault = 7401	Repair Desc = 17FB135,17FB148 - System CPU	2.63%	28.57%	10.86	1.75%	6.14%	
Fault = 7300	Repair Desc = Covers, Cylinder - Eng	1.75%	18.18%	10.36	1.75%	9.65%	
Fault = 7310	Repair Desc = 17FB129, 17FB171	1.75%	18.18%	10.36	1.75%	9.65%	
Fault = 7335	Repair Desc = 17FB138 - INV CPU	10.53%	100.00%	9.50	1.75%	1.75%	
Fault = 7094	Repair Desc = Wing Plate Assembly	7.89%	75.00%	9.50	2.63%	3.51%	
Fault = 73A5	Repair Desc = LPS - Logic Power Supply	5.26%	50.00%	9.50	2.63%	5.26%	

The rules were sorted by descending lift, with those rules with highest lift coming first. Remember from our earlier discussion that lift is essentially a measure of the surprise factor inherent in the rule.

The first rule, with a lift of 38, says that Fault 7294 multiplies the probability of the “17FB135, 17FB148 – System CPU” repair by a factor of 38, raising it from a 2.63% expected probability to a certainty (100% probability). Recall that *Expected* is the prior probability of the repair; *Confidence* is the conditional probability of the repair given that the indicated fault occurred.

Rules Help Focus Further Investigation

Before using rules to change operations (for example, to automatically schedule the System CPU repair as soon as fault 7294 occurs), it is important to examine the rules. Rules produced by associations and sequencing show correlation, not cause. For this a domain expert is required who must then determine, for example:

- Is it appropriate that faults 7294, 7292, 729B, 7400 greatly increase probability of System CPU repair?
- Is it appropriate that fault 7294 ALWAYS leads to a subsequent System CPU repair?

Any such questions that get the answer “No” from the domain expert will need to be further investigated. For example, if the answer to the second question is “No” it may lead to the discovery of an undocumented regulation being informally enforced, e.g., “Always replace the System CPU after Fault 7294.”

Unexpected rules that don't make sense may also signal other, more nefarious, activity, such as pilfering.

Repair to Subsequent Repair

The data provided by the customer also supports an analysis of interactions between repairs (to answer customer's third question). Shown below is the result of that analysis. Again, the resulting rules were sorted by descending lift.

LeftSide	RightSide	Expected	Confidence	Lift	Support	LH Freq
Turbocharger Assembly - General - Eng	17FB129, 17FB171	1.75%	66.67%	38.00	1.75%	2.63%
IFC - IFC Integrated Function Computer	17FB129, 17FB171	1.75%	50.00%	28.50	1.75%	3.51%
Filters, Baggie Element - Sec Eng Air	PSC - Propulsion System Controller	4.39%	100.00%	22.80	1.75%	1.75%
MC - Master Controller	IMC_X - Inverter Motor Controller X=1,2	6.14%	100.00%	16.29	1.75%	1.75%
Cable, Traction Motor Cable or Connector	17FB135, 17FB148 - System CPU	2.63%	40.00%	15.20	1.75%	4.39%
PCU - 13 Pipe Control MV13S, 13CD, MV13E	Piping & Fittings - Eng Fuel Injection	3.51%	50.00%	14.25	1.75%	3.51%
Load, Not Loading, Load Limited	Handbrake Mechanism	4.39%	60.00%	13.68	2.63%	4.39%
DID - Diagnostic Information Display Pnl	17FB134, 17FB159 - Inv I/O	7.89%	100.00%	12.67	1.75%	1.75%
Cable, Traction Motor Cable or Connector	Piping & Fittings - Eng Fuel Injection	3.51%	40.00%	11.40	1.75%	4.39%
DID - Diagnostic Information Display Pnl	17FB138 - INV CPU	10.53%	100.00%	9.50	1.75%	1.75%
Brake Shoes - Truck	Brake Shoe Keys - Truck	2.63%	25.00%	9.50	1.75%	7.02%

The third rule tells us that “Filters, Baggie Element – Sec Eng Air” is always followed by “PSC - Propulsion System Controller”. However, the PSC repair normally occurs in only 4.4% of locomotives.

Again, we need to consult a domain expert and ask the question – is this reasonable? If it isn't, then it raises the possibility that “Filters, Baggie ...” was the wrong repair – i.e., a misdiagnosis? Or, perhaps the “Filters, Baggie ...” repair is causal for the “PSC...” repair. As its name suggests, the “Filters, Baggie ...” repair may involve removing a filter, suggesting this may be a prime example of a repair that causes contamination leading to subsequent repairs. If that is the case, then modifying the “Filters, Baggie ...” repair, for example, by adding a cleaning step or by spraying the filter with a fixative before removal, may result in a reduction of the need for a subsequent PSC repair.

If neither of the above is applicable, then the organization might consider a repair policy change – e.g., always do the “PSC” repair when doing “Filters, Baggie...” repair. This would eliminate having to bring the locomotive back into the shop at a future time to do the PSC repair.

Note that in all of these cases, the analysis results are alerting us to potential problems, but an eventual determination of whether there really is a problem requires a further analysis by a domain expert.

Time Issues in Sequence Analysis

XAAM supports analysis of three types of sequences (most competing products for sequence analysis support only the first type):

- Later (any time later)
- Next Visit
- Between (time window, e.g. 1-5 days)

The *Later* algorithm (which was used in the analysis described above) is the most inclusive and simplest to understand, as it considers all sequences. For Locomotive 106 (see illustration above) there are 7 faults occurring at distinct times, resulting in a total of 13 distinct fault sequences, as shown below:

- | | |
|----------------|----------------|
| a) 7485 → 7A50 | h) 76D5 → 7A2B |
| b) 7485 → 76D5 | i) 76D5 → 76D5 |
| c) 7485 → 7A2B | j) 76D5 → 7154 |
| d) 7485 → 7154 | k) 7A2B → 76D5 |
| e) 7A50 → 76D5 | l) 7A2B → 7154 |
| f) 7A50 → 7A2B | m) 7154 → 7154 |
| g) 7A50 → 7154 | |

The *Next Visit* algorithm would, as its name suggests, only consider events that immediately follow one another. In this case that would include the following reduced set of fault sequences for Locomotive 106:

- | | |
|----------------|----------------|
| a) 7485 → 7A50 | d) 76D5 → 7154 |
| b) 7A50 → 76D5 | e) 7A2B → 76D5 |
| c) 76D5 → 7A2B | f) 7154 → 7154 |

Between requires the additional specification of a time window (or time interval), e.g., "between 0 and 5 days". With this time interval, the fault sequences for Locomotive 106 are further reduced to:

- | | |
|----------------|----------------|
| a) 7A50 → 76D5 | c) 7154 → 7154 |
| b) 7A2B → 76D5 | |

Selecting which time interval to use for an analysis depends on the application. Since it is unlikely that locomotive repairs and faults interact with each other over a time span longer than about 30 days, and with the data spanning a year, the *Later* algorithm is probably not appropriate for this data. On the other hand, it might be worth an exploration with the *Between* algorithm to see if any high lift sequences show up that span long intervals. Such a discovery analysis might invalidate the assumption that there are no interactions beyond 30 days and unearth some problems that no one suspects.

How Can One Repair Predict Another?

Why does AM work? There are numerous reasons why repairs are related in time, in other words, why a second specific repair frequently follows on the heels of another

specific repair. One category is called *downstream* failure. In complex machinery and electronics, many parts are downstream from each other. For example, a valve might be downstream from a pump. If the pump has deteriorated but not yet failed, it is quite possible that it is no longer operating at the called for level. This might result in lower or higher pressures at the valve, increasing the wear on the valve. Downstream effects include changes in pressure, temperature, volume, voltage, current, vibration, etc. A part that is failing might also release debris while it is failing that can have downstream effects. Parts that are downstream from the failing part will be operating in an out-of-specs environment until the upstream failure is detected and corrected. This out-of-specs operation can cause wear and subsequent downstream failure.

When a part is failing, it may no longer be doing its share of the work. For example, when a brake is failing, other brakes in the system may be asked to do more work. This could cause extra wear in the other brakes.

The act of replacing a part can also contribute to failures in other parts. Primarily this will be due to the release of contaminants or debris that is disturbed when the failed part is replaced.

Is LocoType a Factor?

LocoType is an attribute that distinguishes between types of locomotive. The data that the customer provided only contained faults and repairs for two types of locomotives, type 1 and type 2. We might be interested in repair sequence rules that highlight differences between locomotive types. Such an analysis might help us find maintenance problems that are unique to a particular locomotive type.

Partitioned analysis is a unique feature in XAAM. Partitioning allows you to analyze the data separately for each different value in the partitioning column. In this case, we used LocoType as the partitioning column, producing rules which can be compared across LocoType to identify similarities and dissimilarities across types of locomotives.

Another unique tool in XAAM is the *discriminator*. The discriminator helps you find rules that have the largest variations across the partitioning column in rules measures such as Lift or Confidence.

Shown below is discriminator output from a repair-to-repair sequence analysis that was partitioned on LocoType. Discriminator output summarizes the rules. The output has been sorted in descending order on the standard deviation in lift (LiftStdDev). As a result, those rules that have the largest difference in lift are at the top of the list.

LeftSide	RightSide	Count	ConfStdDev	LiftStdDev	SuppStdDev
PM_XYZ - Phase Module,X=Inv#,Y=Ph	PM_XYZ - Phase Module,X=Inv#,Y=Ph	2	56.5685%	9.2065	3.2855%
IFD - IFC Interactive Display Unit (3)	Coupling, Eng Pumps Drive (Rubber Bo	2	16.4992%	5.2326	3.9998%
Electronic Air Brake General - Knorr	17FB132 - Digital I/O	2	27.4986%	5.1265	8.7138%
Gear Drive, Case	ACHU,HVAC - Air Cond/Heater Unit,Op	2	61.2826%	4.4548	3.2855%
IDU - Input/Display Unit (TIBS)	ACHU,HVAC - Air Cond/Heater Unit,Op	2	56.5685%	4.0305	3.9998%
Load, Not Loading, Load Limited	BATT - Battery (Locomotive Batteries)	2	11.7851%	3.8596	3.2855%
Handbrake Mechanism	BATT - Battery (Locomotive Batteries)	2	11.7851%	3.8596	3.2855%
Wheels - Truck	Gear Drive, Case	2	21.2892%	3.3850	3.9998%
17FB138 - INV CPU	Wheels - Truck	2	64.2824%	3.3302	3.9998%
Handbrake Mechanism	Cylinder Assembly - General - Eng	2	11.7851%	3.1820	3.9998%

The first rule, which has the same value on the left and right hand side, identifies a repair that apparently needs to be repeated. Because there is a very large difference in lift, apparently the repair needs to be repeated much more frequently for one locomotive type than for the other.

A single mouse click lets us review the detailed rules that correspond to the first summary, as shown below. This shows that there is indeed a remarkable difference between locomotive types – locomotive type 1 has a 100% re-repair incidence!

LeftSide	RightSide	LocoT	Expecte	Confident	Lift	Suppor	SupCou	LH	Fre
PM_XYZ - Phase Module,X=Inv#,Y=Ph	PM_XYZ - Phase Module,X=Inv#	1	6.67%	100.00%	15.00	6.67%	1	6.67%	
PM_XYZ - Phase Module,X=Inv#,Y=Ph	PM_XYZ - Phase Module,X=Inv#	2	10.10%	20.00%	1.98	2.02%	2	10.10%	

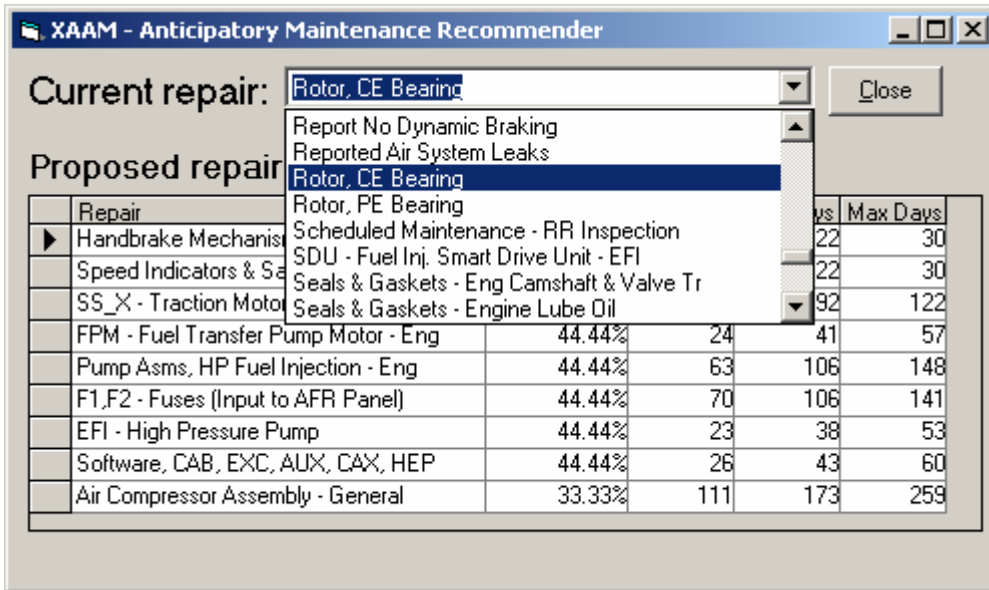
The discriminator can be used in the same way to see if there are other important factors, such as repair depot (if multiple repair facilities are in use), day-of-week of the repair, training level or seniority of the repair person, etc.

Deploying AM Rules on the Shop Floor


XAAM includes a deployment tool called the Anticipatory Maintenance Recommender. This program can be run on any computer with Microsoft Internet Explorer.

As the name suggests, the Recommender proposes other items which should also be considered for repair in conjunction with an item currently being repaired. As we shall see later, selecting one or more of the proposed repairs can be based on the likelihood of failure, the historical time interval until failure occurs, and/or potential cost savings.

For example, the CE Bearing Rotor is being replaced because it is defective. The maintenance engineer starts the Recommender on any network connected computer, and selects the current repair from a drop down list, as shown below.



The Recommender then lists suggested repairs. By default, the repairs are listed in descending order of propensity, with the most likely subsequent failures first and the least likely last, as shown below. Handbrake Mechanism and Speed Indicators ... are listed first because their likelihood of failure, subsequent to the CE Bearing Rotor repair, is just over 55%.



As an alternative to viewing the list of recommended repairs by propensity, they can be listed in order by Min Days, as shown below. Min Days shows the minimum historical time lapse between the two repairs.

Repair	Propensity	Min Days	Avg Days	Max Days
▶ Handbrake Mechanism	55.56%	13	22	30
Speed Indicators & Safety Control - Gen	55.56%	13	22	30
EFI - High Pressure Pump	44.44%	23	38	53
FPM - Fuel Transfer Pump Motor - Eng	44.44%	24	41	57
Software, CAB, EXC, AUX, CAX, HEP	44.44%	26	43	60
SS_X - Traction Motor Speed Sensor, X=#	44.44%	61	92	122
Pump Asms, HP Fuel Injection - Eng	44.44%	63	106	148
F1,F2 - Fuses (Input to AFR Panel)	44.44%	70	106	141
Air Compressor Assembly - General	33.33%	111	173	259

If costs are known, they can be used to prioritize repairs. There is a need to know both actual repair costs (parts and labor) as well as the costs associated with the interruption of service that would occur if the part failed unexpectedly.

We call the latter the **Unscheduled Cost**. Unscheduled costs are typically estimated and tiered, based on the nature of the failure and the expected time out of service due to an unexpected failure. If a part is not a critical part, that is, if operation of the equipment can continue for an indeterminate time, even after the part fails, then the **Unscheduled Cost** is zero or close to zero. If failure of the part causes the equipment to stop functioning altogether, then the **Unscheduled Cost** is very high. If failure requires taking the equipment out of service, but some operation for a short time or at reduced levels is possible, then the **Unscheduled Cost** is at an intermediate level. For our example, we elected to use four tiers for **Unscheduled Cost**, as follows:

1. Non-critical failure. **Unscheduled Cost** is \$0.
2. Operation generally can continue to next station, but then equipment will be out of service for at least one hour and up to 8 hours. **Unscheduled Cost** is \$1,000.
3. Critical failure. Equipment stops working and/or equipment will be out of service for 8-24 hours. **Unscheduled Cost** is \$2,000.
4. Critical failure. Equipment stops working and/or equipment will be out of service for 24 hours or more. **Unscheduled Cost** is \$3,000.

Shown below is the Recommender with costs. The recommended repairs are listed in order by overall expected savings if the recommended part is replaced. For example, the expected savings of replacing the Handbrake Mechanism is \$470. These savings have been calculated by comparing the cost of replacing the handbrake now (\$192) to the expected cost of replacing it when it fails. If the handbrake is not replaced, and it subsequently fails, then the cost to replace it will be \$1,192 - the cost of parts and labor (\$192) plus the **Unscheduled Cost** (\$1,000). Since the probability of future failure is

55.56%, the expected cost of future failure is $0.5556 * \$1,192$, or \$662. The difference between the expected cost and the actual cost is \$470 (\$662 minus \$192). Hence, replacing the handbrake has an expected savings of \$470.

Some of the parts on the recommended list have negative savings, and generally such parts should not be replaced. For example, replacing the Fuel Transfer Pump Motor has an indicated savings of minus \$47. This indicates that it's less expensive to wait for the part to break than to repair it now.

Repair	Propensity	Min Days	Avg Days	Max Days	RepairCost	UnSchedCost	ExpectedCost	Savings
SS_X - Traction Motor Speed Sensor, X=#	44.44%	61	92	122	\$157	\$2,000	\$959	\$802
Handbrake Mechanism	55.56%	13	22	30	\$192	\$1,000	\$662	\$470
Air Compressor Assembly - General	33.33%	111	173	259	\$668	\$2,000	\$889	\$221
FPM - Fuel Transfer Pump Motor - Eng	44.44%	24	41	57	\$884	\$1,000	\$837	-\$47
Pump Asms, HP Fuel Injection - Eng	44.44%	63	106	148	\$167	\$0	\$74	-\$93
Speed Indicators & Safety Control - Gen	55.56%	13	22	30	\$385	\$0	\$214	-\$171
F1,F2 - Fuses (Input to AFR Panel)	44.44%	70	106	141	\$1,178	\$1,000	\$968	-\$210
EFI - High Pressure Pump	44.44%	23	38	53	\$780	\$0	\$347	-\$433
Software, CAB, EXC, ALD, CAX, HEP	44.44%	26	43	60	\$832	\$0	\$370	-\$462

Conclusion

Data mining of equipment maintenance and repair data can help discover:

- Anticipatory maintenance (AM) procedures that reduce future failures
- Repairs (or maintenance operations) that are being done improperly
- Ways to improve repairs that reduces subsequent down time
- Undocumented methods being used by experienced personnel that result in reduced down time
- Advance notice of likely failures *before* failures occur

Such discoveries can be used to modify your maintenance and repair procedures thereby reducing downtime, increasing uptime, and significantly reducing the costs of maintenance and repair.

To find out more about Exclusive Ore Inc. and its custom solutions to data warehouse, database and data mining problems go to www.exclusiveore.com or click [here](#).