EQUITY TECHNICAL REPORT 237 – 2025

EXPLAINABLE AI (XAI) FOR MI WITH THE EQUITY BRAIN

Charles H. Panzarella, Ph.D.
Aaron Stenta, Ph.D.

EQUITY
RESEARCH

Intentionally Left Blank

The Equity Technology Group, Inc.

20600 Chagrin Blvd.

Suite 1200

Shaker Heights, OH 44122

www.E2G.com

Intentionally Left Blank

TABLE OF CONTENTS

Intentionally Left Blank

1 EXECUTIVE SUMMARY

Equity is spearheading a world-class reliability program rooted in explainable artificial intelligence (XAI) that bridges the gap between raw data and effective decision-making in asset lifecycle management. The program is designed to transform operational, inspection, and maintenance data into actionable strategies using a hybrid of human expertise and machine intelligence. Central to this approach are probabilistic, physics-based, causal network models that are transparent, learn from field data over time, and are rooted in causality—enabling them to deliver explainable and continuously improving insights. By doing so, Equity provides a system that not only predicts asset health degradation but also recommends optimal, cost-effective actions in real time.

This technology directly addresses the two primary challenges in mechanical integrity programs: underinvestment leading to asset failure, and overinvestment resulting in unnecessary cost. Equity's patent-pending methods determine the optimal spend by balancing the benefit of risk reduction with the associated cost, maximizing return on investment. The probabilistic nature of these models enables accurate prediction of failure time distributions and associated probabilities of failure, leading to better operational, inspection, and maintenance decisions. The methodology constitutes a significant advancement over traditional risk-based inspection (RBI) methods by incorporating uncertainty from all inputs, more accurately assessing the effectiveness of inspections, and providing real-time, dynamic decision-making capabilities.

Equity's solution is implemented through a modular, cloud-based software platform called HealthSight, which connects to any asset system and live sensor infrastructure. HealthSight integrates with the Equity Engineering Cloud (eec), and key predictive engines such as DamageSight and CUISight, to provide a centralized dashboard for real-time asset health monitoring and guided recommendations, setting a new standard in mechanical integrity. DamageSight models the rate of damage accumulation from all known damage mechanisms. CUISight is the first fully implemented submodule of DamageSight focusing on one of the most widespread and costly damage mechanisms, Corrosion under Insulation (CUI), with demonstrated successful pilot implementations in predicting and managing damage from CUI, resulting in significant cost savings.

Equity's approach is backed by years of research and development supported by U.S. Department of Energy grants, including the BENGI (Bayesian Engine for Insights) platform. Equity has already secured one patent and has another pending on asset lifecycle optimization systems. By combining advanced probabilistic modeling, causal inference, and cloud-based delivery, Equity's program represents a transformative leap forward in mechanical integrity—enabling companies to make better, data-informed decisions that enhance safety, reliability, and economic performance across industrial assets.

2 INTRODUCTION

Equity is creating a world class reliability program rooted in explainable AI (XAI) that bridges the gap between data and decisions. It is a common problem that operators have all this data, but they don't know what to do with it. This world class reliability program makes full use of all data to make the best asset lifecycle decisions with regards to operation, process, inspection, and maintenance. Equity's patent pending methods and software turn data into decisions and then actions. This approach is a hybrid between human and machine intelligence that starts with improved probabilistic, physics-based, causal network models that account for all sources of uncertainty. These causal network models are built initially by human experts but then dynamically learn and get smarter over time from on-the-job experience and observations. Intelligence, whether human or machine, is the ability to make good decisions when faced with any problem and to learn and get smarter over time (make better decisions). Understanding the root causes of any problem is key to making good decisions, and this is what Equity's explainable AI approach is all about. It is rooted in causality, which makes it smarter and more explainable right from the start. It connects to all available data sources to continuously monitor the health of assets in real time by predicting all possible ways in which health degrades. It seamlessly blends together different sources of knowledge (predictive models, human experience, and data) to arrive at a single source of truth, serving as a central repository of knowledge that is shared across all assets, facilities, plants, and companies to benefit the whole industry, as well as the public, at large. This will usher in a new era of next generation Risk-Based Inspection (RBI), asset lifecycle optimization, and mechanical integrity that makes full use of this innovative technology to enable the best possible decisions in all scenarios.

<u>The overarching goal of an effective mechanical integrity (MI) program is to maximize equipment availability at the lowest cost possible.</u> In other words, assets need to be kept reasonably healthy to achieve some desired financial objective, such as maximizing the total return on investment (ROI), while also promoting safety. The tricky part of this endeavor is knowing what the right spend is on inspection and maintenance to achieve this goal – not too little and not too much. If too little is spent, then the unreliability of the asset leads to excessive failure costs and loss of safety. If too much is spent, then failure is unlikely and safety is ensured, but one is past the point of diminishing returns, where additional spending provides less benefit than it costs and is no longer worth it. When is enough, enough?

This directly leads to the two most common problems that operators face:

1. **Excess failures** due to spending too little on inspection and maintenance and/or conducting ineffective inspection and maintenance, or
2. **Excess spending** on inspection and maintenance to achieve the desired level of reliability but at a total expenditure that is not sustainable from a business economics perspective.

<u>Equity's XAI for MI approach directly addresses both of these problems by leveraging patent-pending XAI methods to find the optimal spend that maximizes ROI (i.e., finds the sweet spot that yields the most bang for your buck).</u> This is accomplished by perfectly balancing the benefit and risk reduction of any lifecycle decision strategy with the increased cost of performing it. If the risk reduction of any activity does not outweigh its cost, then it is not worth it, at least *not at that time* (it may become worth it later on as risk increases*)*. Those activities that provide the most risk reduction for the lowest possible cost are generally the most beneficial and should be prioritized. By quantifying the risk reduction of all possible activities using these new and more effective XAI methods, improved guidance is provided on where, when, how, and how much to inspect and maintain to achieve maximum risk reduction at the lowest cost possible (i.e., the maximum ROI).

<u>Key to this XAI for MI approach is improved probabilistic, physics-based, predictive models for the damage rate from all known aging asset damage mechanisms (DMs) and their inherent uncertainties.</u> These improved models are used, along with appropriate failure conditions, to predict every asset's failure time probability distribution, and its associated probability of failure (POF) and risk versus time curves, as accurately as possible. Given this failure time prediction for every asset, optimal decision strategies are recommended, in near real time, to achieve optimal outcomes, taking into account, as well, the predicted life extension of all possible actions. This is a fully probabilistic approach that, unlike traditional RBI, properly accounts for all sources of uncertainty from start to finish, including uncertainty in all model inputs – which is why this approach is also being billed as *next-generation RBI*.

<u>To fully appreciate this novel XAI for MI approach, one must first view the underlying problems of MI from a deeper, more holistic, level.</u> In essence, MI boils down to keeping an asset healthy, where health is the ability of the asset to perform its desired function. Health degrades over time due to the accumulation of damage from one or more DMs until some critical state of damage is reached, at which point failure is said to occur. Failure can be defined in many ways, but one is as the inability of the asset to perform its desired function. There are often one or more undesirable consequences accompanying failure, besides the inability to function properly, such as loss of property, loss of life, environmental impacts, legal repercussions, loss of reputation, etc. These must all be accounted for in a unified, holistic framework that is aware of all possible causes of failure. One way of unifying all this is to focus on failure time as the primary driver for all lifecycle decision strategies being considered, with the effect of every possible event or action on the failure time predicted as accurately as possible in the form of a quantified life extension or life reduction probability distribution. Once one has a complete list of all possible inspection and maintenance actions and the cost and benefit (in terms of life extension) of those actions at all possible times they could be performed, it is straightforward to find the optimal decision strategy, even in the face of uncertainty. This is why developing improved methods for predicting the failure time more accurately, under every possible circumstance, is so essential. However, a key point is that it is not necessary to predict failure time exactly, only to *properly quantify its uncertainty*, to find the best decision strategy. *Perfection is not required*. In short, we must make full use of all available data and knowledge, often no more data and knowledge than what is already readily available, to suggest the best possible actions to take under every possible circumstance. Data that is being collected now is often not being put to full use, and this needs to be corrected.

<u>All unmaintained assets eventually fail – the only question is when?</u> Predicting this time of failure is perhaps the most important requirement for making good MI decisions. If the failure time were known precisely, the optimal solution would be trivial – run the asset right up to slightly before the precisely known failure time and then replace or repair before it fails. No further inspection or maintenance would be required, and unexpected failures would never occur. It is only because the failure time is not known precisely that all complexity arises. Uncertainties in failure time necessitate a fully probabilistic approach, wherein a failure time probability distribution is obtained from some predictive model of failure, utilizing all historical and future data and knowledge to improve future predictions. Once the failure time distribution is known as best as possible, one can then make the best lifecycle decisions regarding design, operations, inspection, and maintenance to achieve the desired outcome.

<u>Even though failure time is uncertain, optimal decisions are not.</u> In other words, there are precise times at which performing activities like inspection, maintenance, replacement, or repair leads to optimal financial results, even when the failure time is uncertain. Obviously, decisions are more conservative when this uncertainty is higher, which is why better predictive models that reduce this uncertainty and its associated conservatism as much as possible lead to greater ROI. This is the primary reason for pursuing improved predictive models that make full use of all available data and knowledge, and using more predictive and expansive XAI methods is one way to achieve this.

<u>So, what is AI and how can it improve MI?</u> Before defining artificial intelligence (AI), it is important to first understand the nature of human intelligence and its role in improving mechanical integrity. At its core, human intelligence can be thought of as the ability to make sound decisions (take sound actions) when confronted by complex problems, often under conditions of great uncertainty. This process becomes especially challenging when it involves integrating multiple sources of information—some of which may be incomplete, uncertain, or even conflicting. A more intelligent individual is typically one who can navigate these challenges more effectively to arrive at better decisions and actions. Artificial intelligence, in this context, can be viewed as a computer system designed to replicate, and potentially surpass, this human decision-making capability. A well-designed AI system for mechanical integrity must be able to synthesize diverse and sometimes contradictory inputs—ranging from sensor data and inspection results to expert judgment and historical performance—and use this information to make informed, justifiable recommendations. The goal is to automate and enhance the decision-making process so that it reflects the expertise of the most capable human operators, while also ensuring consistency, scalability, and continuous improvement.

Humans solve problems fundamentally by identifying their causes and then taking steps to eliminate or mitigate those causes. This principle lies at the heart of all science and engineering—understanding why something occurs is the first step toward controlling or preventing it. When the cause of a problem such as damage or failure is understood, informed decisions can be made to prevent or reduce its impact. In situations where prevention is not feasible, understanding the underlying cause still allows for accurate prediction of how quickly the problem will progress, enabling timely interventions such as inspections, repairs, or replacements. This foundational reliance on causality is not only central to human reasoning but is also essential for any intelligent system. This is why Equity's approach to explainable AI is deeply rooted in causality. By building models that explicitly capture cause-and-effect relationships, the system provides not just predictions, but clear rationales for those predictions—ensuring that decisions are always grounded in a comprehensive understanding of how and why damage or failure occurs.

This focus on causality is the basis of Equity's novel XAI approach to improve MI and lifecycle decisions. The explainable part refers to first developing visual, probabilistic, physics-based, causal networks that clearly identify the causes of damage, failure, or anything else of importance as best as possible. These networks are inherently probabilistic, because the failure time and everything it depends on are plagued by uncertainty that cannot be ignored. These networks are physics-based, because they also take into account all that is known about the causes of damage and failure as acquired through the traditional methods of science and engineering (e.g., electrochemical corrosion models that account for reactions, diffusion, etc.). Unlike traditional "black-box" AI methods that simply process large amounts of data with the hope of finding hidden patterns by chance that may already be understood well by science and engineering, Equity's approach is a *hybrid between human and machine intelligence*, integrating the best of both worlds to come up with a holistic solution that is better than either can achieve separately. Nothing is lost.

These probabilistic, physics-based, causal networks solve MI problems better, and they learn and get smarter over time from on-the-job experience, just like human experts do. These networks serve two essential functions: prediction and inference. For prediction, they use observations related to the causes of damage to estimate how quickly damage is accumulating. When combined with a failure model, they forecast the expected time to failure. For inference, they work in reverse—starting from observed outcomes, such as inspection results or damage measurements, they probabilistically determine the most likely underlying causes and estimate parameters like damage rates based on the evidence. A defining strength of these networks is their ability to learn and improve over time, much like human experts. By continuously comparing real-world data with model predictions, they identify discrepancies and adjust accordingly. For instance, if a model predicts a certain rate of damage but field measurements reveal faster deterioration, this information is incorporated into the network by adding new observation nodes. As a result, the model learns (a defining characteristic of intelligence) and becomes more accurate for future predictions. Importantly, this learning is shared across different applications, so insights gained at one facility or under certain conditions enhance predictions elsewhere in similar environments.

Equity Research has been working for several years on building a generic software framework for solving all MI problems with these probabilistic, physics-based, causal networks. This work has been partially funded by multiple Department of Energy (DOE) grants, since 2018. A key grant was the first Phase I DOE award in 2021 for BENGI – Bayesian Engine for Insights. BENGI is a custom-built engine for solving industrial-scale, probabilistic, physics-based, causal networks – the underlying basis of Equity's XAI for MI solution. Equity was awarded a Phase II grant for BENGI in 2022, totaling about $1.5 million in grants for BENGI research alone. Many other grants were awarded for related topics.

Equity Research has already received one patent on this approach with another one pending. Equity was awarded its first patent on using Bayesian hierarchical analysis to predict corrosion rates and failure of assets in 2022 [1] and applied for a more overarching, comprehensive patent in 2025 on *Asset Lifecycle Optimization Systems and Methods* [2]. This approach, when implemented in software to solve mechanical integrity problems, is being referred to as the ***Equity BRAIN*** **– Bayesian Reasoning for Asset Integrity using Networks**. The Equity BRAIN also integrates the human intelligence of all of its experts to be even more powerful than any machine could be on its own.

HealthSight is Equity's overarching software platform that implements these methods on the cloud. The overarching cloud software platform that implements the methods of the Equity BRAIN is referred to as HealthSight, since the principal goal of asset integrity is to keep equipment healthy, and this approach provides new insights into achieving that goal. HealthSight is a layer built on top of The Equity Engineering Cloud (eec) – Equity's cloud software platform for engineering analysis consisting of dozens of modular specialty engineering calculators that have been commercially available since 2017. The primary function of HealthSight is to serve as a central hub between any asset system, any number of live sensors (DCS or in-place NDE sensors), and any calculator in the eec to provide real-time guidance on MI decisions using these methods. The centerpiece of HealthSight is a configurable live dashboard that provides real-time updates on the health status of any asset, including an alerting system that recommends actions (immediate or delayed) based upon its continuous health monitoring of the asset's state.

DamageSight is the most important eec calculator that HealthSight connects to for the purpose of predicting damage rates from all possible damage mechanisms. This is a tool for predicting the asset damage rate and failure time for all possible damage mechanisms, including, but not limited to, all the damage mechanisms in API RP 571 using probabilistic, physics-based, causal networks.

CUISight is the first successful implementation of this XAI for MI approach, focused on one of the most nuisance and costly DMs – Corrosion Under Insulation (CUI). This tool predicts the life of any asset subject to CUI damage better than any other available approach. Before any inspection data is available, its physics-based prior predictions are used to plan the first round of inspections. Once inspections are performed, the inspection results are fed back into the model to improve future predictions, which are then used to plan subsequent inspection and maintenance activities. This process is repeated cyclically throughout the entire lifecycle with the model learning and becoming smarter all the time – a true sign of intelligence! The model retains this learned knowledge and is therefore dynamic and continuously improving, as opposed to static models used now. This system is currently being piloted at several client facilities and is already demonstrating great success. Early results indicate millions of dollars in savings per client by either 1) reducing failures or 2) focusing inspection efforts more effectively to reduce or eliminate unnecessary inspections. This achieves the desired goal of maximizing equipment availability at the lowest possible cost.

Equity is seeking innovative and forward-thinking strategic partners to participate in leading-edge pilot studies to further validate this new and exciting XAI for MI technology. Equity is hoping to expand awareness of this approach and further its adoption and use, which is game-changing for the industry. When combined with traditional AI methods, like advanced computer vision and neural networks, even better predictions are possible.

3 THE PLAYING FIELD

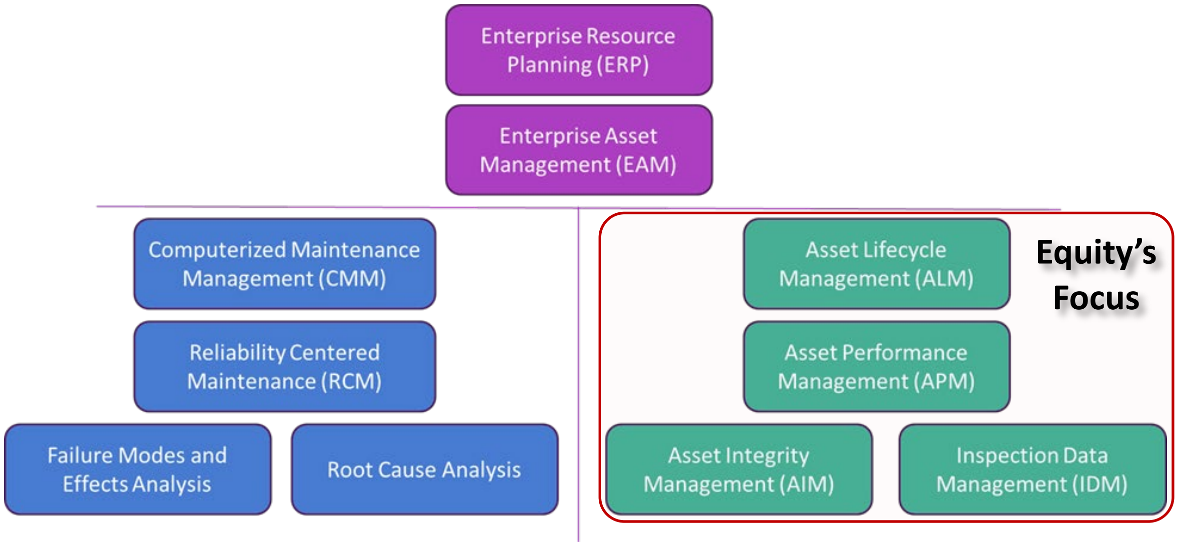

Figure 1 – Classification of existing asset management systems from the highest to lowest level. The focus of Equity's current efforts is to improve upon all of the functionality in the green section (bottom right) while integrating with other existing systems in the blue (bottom left) and purple (top) sections.

The problem area Equity is primarily focusing on is generally referred to as Asset Performance Management (APM) which fits into the hierarchy of all available systems as shown in Figure 1. The solution to problems in this area can generally be referred to as Asset Lifecycle Optimization (ALO) of decisions to receive the greatest benefit possible from operation of the asset by optimizing all lifecycle decisions. Within that field is mechanical integrity, which is more narrowly focused on maintaining the health of assets as much as possible so they can perform their function as designed at the lowest cost possible with maximum availability.

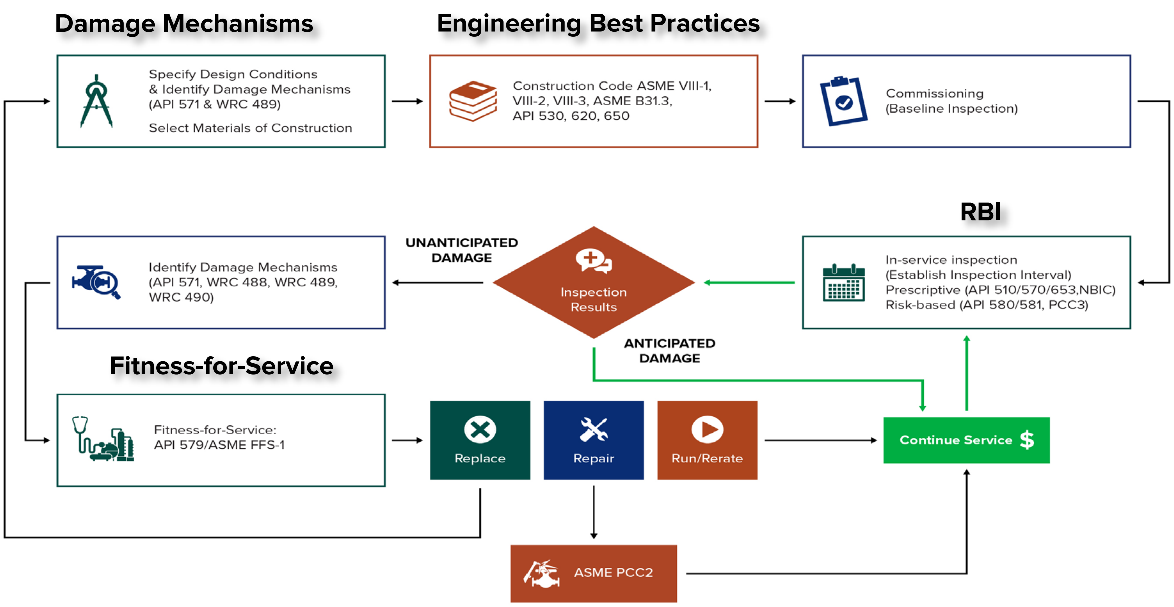

Figure 2 – Traditional asset life cycle management approach used by Equity and other companies to keep assets healthy. The four primary services are identifying damage mechanisms, engineering best practices, fitness for service, and RBI. These are conducted according to the indicated code-based procedures and practices. Equity has traditionally provided software and services in all of these areas, but new explainable-AI-based enhancements are on the horizon.

The traditional asset life cycle management approach used by Equity to serve its clients for many decades is shown in Figure 2. Equity is rooted in these classic engineering methods, but new developments are on the horizon that will enhance these by bringing in the latest in explainable-AI (and traditional AI) technology. To achieve this goal, it is useful to convert this life cycle diagram into a decision flowchart that highlights the decisions operators face and also makes explicit all costs and benefits, as shown in Figure 3.

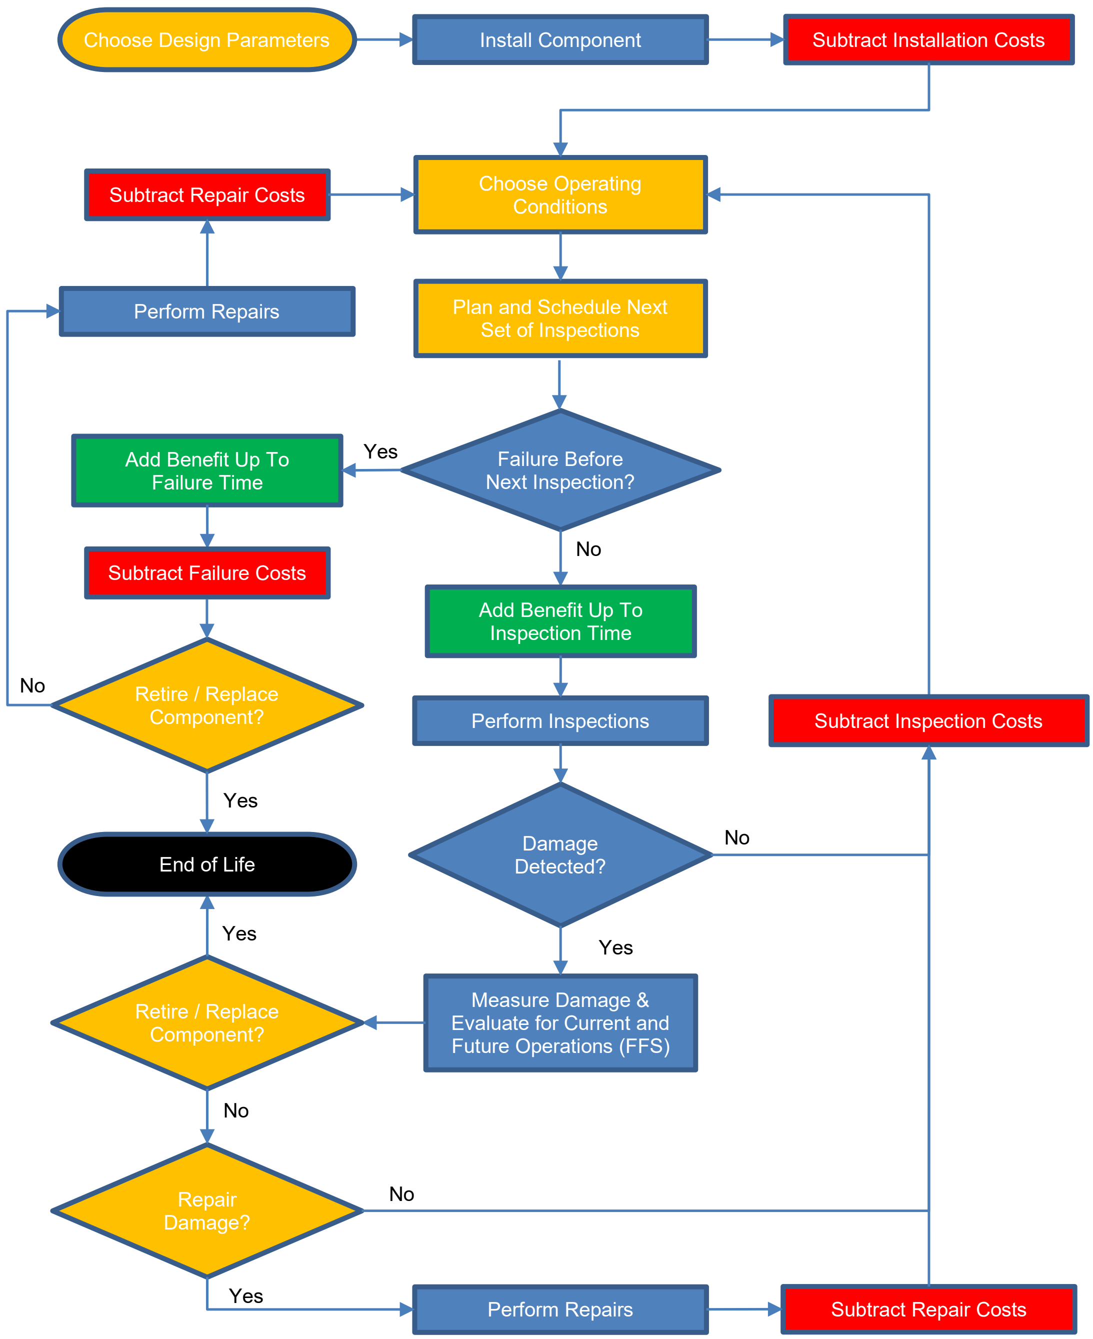

Figure 3 – Flowchart for the generic decision-making processes throughout the entire lifecycle of an aging asset from installation to end of life that highlights how costs and benefits are subtracted and added after certain decisions or events occur to get the total return on investment over the lifetime. The goal is to make the best decisions along the way so that the total lifetime-averaged return on investment is maximized (make

the most amount of profit in the shortest amount of time) while also, obviously, ensuring safety. All primary decisions are highlighted here with orange diamonds or boxes. When assessing damage that is discovered, the classic Run, Repair, Rerate, Replace, and Retire decisions (the 5 R's) need to be made, and the methods described herein can be used to help make these decisions in the best possible way (to maximize availability at the lowest possible cost).

4 IMPROVING MECHANICAL INTEGRITY

4.1 Finding the Optimal Spend on Inspection and Maintenance

The entire field of mechanical integrity exists because assets used for some beneficial purpose inevitably age and will eventually fail unless something is done to intervene. If assets did not fail, there would be no need for mechanical integrity. Predicting when an asset fails and coming up with methods to extend its usable life and prevent failure is one way to express the overarching goal of mechanical integrity. There are many other ways of framing this as well, such as maximizing equipment availability at the lowest possible cost.

At the highest level, the only reason money is spent on inspection and maintenance is to reduce the risk of failure. In general, one expects to achieve more risk reduction as more is spent, as is the case for the representative blue risk-reduction curve shown in Figure 4. From a financial perspective, the return on investment (ROI) – or the Return on Capital Employed (ROCE) – for inspection and maintenance is the benefit (risk reduction) minus its cost, which is represented by the red curve in Figure 4. One immediately notices that due to the diminishing risk reduction observed for higher and higher spends, the return reaches a maximum value here at the particular optimal spend of about $40k. This would be the recommended spend on inspection and maintenance if the desired objective were to maximize the ROI or ROCE.

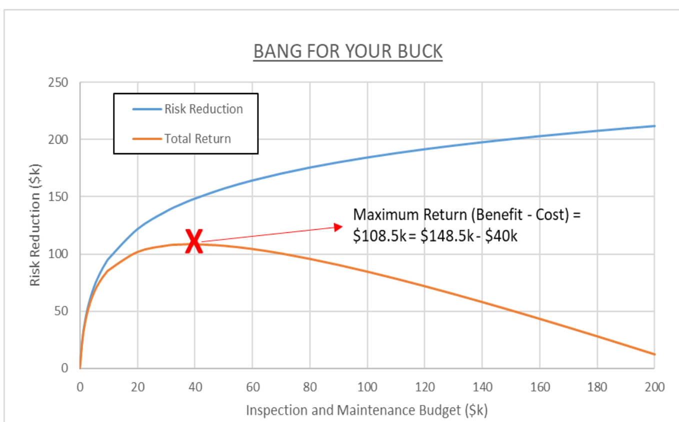

Figure 4 – There is generally an optimal spend on inspection and maintenance that can be found if the risk reduction versus spend is known (the blue curve shown in the above figure). The return on investment (the orange curve above) is the benefit (risk reduction) minus the inspection and maintenance cost. For this example, there is a point of maximal return at the optimal spend of about $40k on inspection and maintenance due to the ever-decreasing slope of the risk reduction curve (less and less risk reduction is obtained for more and more spend on inspection & maintenance). To the right of this point, one is over-inspecting. To the left of the point, one is under-inspecting. The sweet spot is right in the middle, where one gets the most bang for your buck.

Notice that this powerful conclusion is reached by simply knowing the risk reduction versus spend. To find this relationship, one needs to predict the risk of failure both before and after all possible inspection and maintenance activities. This requires an accurate predictive model of failure that is also able to properly quantify the effectiveness of inspection and maintenance at reducing risk. Since failure depends on the accumulation of damage over time from one or more damage mechanisms, this is why predictive models for damage progression are required, along with some specified condition for failure once damage reaches a critical level.

Not only does this approach find the right spend on inspection and maintenance for any given strategy, but it also compares different strategies to find the best one. For example, suppose there are two inspection and/or maintenance strategies, A and B, that use different approaches and methods and may be performed at different times, and a risk reduction versus spend curve is obtained for each. Once the optimal spend is determined for each using the above approach, the best overall strategy is the one having the largest return at its optimal spend. This approach can be expanded to thousands of possible strategies. There is no limit with today's computational resources. Looking at different times for performing some action can also be regarded as different strategies, so this approach not only determines what to do but also at what time it is best to do it.

Thus, making optimal high-level financial decisions for mechanical integrity just boils down to being able to do one thing and one thing well: **determine the risk reduction curve versus spend on inspection and maintenance for every possible strategy under consideration, including both what to do as well as when to do it**. One can use any desired method to do this, but Equity has developed an improved method for doing this that involves using explainable AI models for both prediction and optimal decision making.

4.2 Predicting Failure Time is Key

Key to determining the risk reduction curve is being able to predict the failure time probability distribution function (PDF) for any asset both before and after any planned actions that might be considered or unplanned events (like accidents or natural events) that might occur. The cumulative distribution function (CDF) of this failure time probability distribution is equivalent to the probability of failure (POF) curve, which can be evaluated at any time and multiplied by the consequence of failure to get the risk of failure by operating over that time window. Once the risk is known before and after any action, the risk reduction resulting from performing that action is simply the difference between these two. This all relies on an accurate determination of the failure time distribution, which is why that is so important and central to this whole endeavor.

4.3 Keeping Assets Healthy

Sometimes, how one looks at a problem makes all the difference. Rather than discussing damage all the time (the glass half empty approach), Equity prefers to focus on keeping assets healthy (the glass half full approach). Perfect health is the complete absence of damage. Health degrades as soon as any damage starts accumulating. There are many ways in which health can degrade, by many different damage mechanisms and at many different locations, and in many different components that might comprise the asset as a whole. However, there is only one perfect state of health, and that is what we are always striving to achieve. This is why focusing on keeping assets healthy is a conceptually simpler and more optimistic way of framing this entire endeavor.

To predict the degradation of health from its perfect state, one must understand all possible ways in which damage accumulates, from every possible damage mechanism. This is what Equity is doing by developing improved predictive models for every damage mechanism in API RP 571 and beyond. One way of improving predictions is by using AI, in particular the form of explainable AI (XAI) that Equity has been focusing on for a number of years already. More traditional *black-box* AI methods can be incorporated for special purposes, such as gathering the input data necessary for the XAI models to make more accurate predictions. This data gathering or data digitization process is crucial to making all this work in practice. However, it is a worthy goal to develop these improved predictive models even before such other applications are completed or integrated.

4.4 Optimizing Decisions

Once the failure time probability distribution $f(t_F)$ is known for every asset under consideration both before and after all possible actions that could be performed, one determines the optimal decision strategy that achieve some desired financial objective, such as maximizing the return on investment (ROI), return on capital employed (ROCE), or anything else.

As one simple example, consider a single asset whose failure time probability distribution, $f(t_F)$, is known. The probability of failure (POF) curve is just the CDF of this distribution, $POF = F(t_F)$, where $f(t_F) = dF/dt_F$. The optimal replacement time, t_R, is easily determined for an asset given this failure time distribution if the cost of replacement, C_R, and the cost of failure, C_F (the consequence) are also known. The return on investment, H, is the benefit of operation minus the expected cost of failure up to the replacement time. If one assumes the benefit of operation (revenue) is accrued at a nearly constant rate over time, then the benefit drops out of the lifetime-averaged return defined by $h = H/t_R$. It is normally h that one wants to maximize. That is, **as a business, one normally wants to make as much profit as possible over the shortest period of time**. Assuming the asset will be replaced after its useful life is consumed, this cycle of replacement is assumed to continue indefinitely.

With these basic assumptions it is possible to show that the optimal replacement time of the asset satisfies the following simple, nonlinear equation that maximizes h,

$$1 - F\left(t_R\right) = \frac{C_F}{C_R} \cdot t_R \cdot f\left(t_R\right) \tag{1}$$

An example plot of h versus replacement time in Figure 5 shows that it reaches a maximum value at about 7.4 years when $C_F/C_R = 10$. This is the same optimal replacement time obtained by solving Equation (1).

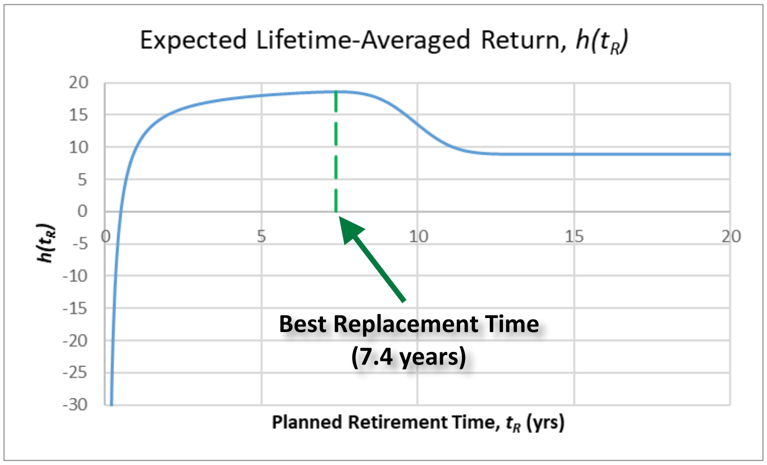

Figure 5 – The optimal replacement time for some aging asset maximizes the total lifetime-averaged return h, which occurs here at about 7.4 years assuming $C_F/C_R = 10$. This agrees with Equation (1). This same principle of maximizing h holds even when very complex decision strategies involving operation, inspection, and maintenance are being evaluated.

This simple example can be extended to much more complicated scenarios where multiple possible decisions could be made at different times, and one wants to know what the best possible combination of decisions is, i.e., the optimal decision strategy. Regardless, the basic principle remains the same, **finding the combination of decisions that results in a maximum lifetime-averaged return**. That would be the recommended optimal decision strategy. For such scenarios, a simple analytical solution may not be possible, but this is where the more sophisticated methods of XAI help, as shown later.

For example, suppose that by whatever method (the above or by using a more sophisticated XAI approach), one is able to find the optimal lifetime-averaged return by comparing five different decision combinations (strategies) as shown in Figure 6. Each of these strategies could consist of many different combinations of decisions related to inspection and maintenance across multiple assets using different inspection methods, repair techniques, replacements, and inspection coverage areas. From a financial perspective, the best strategy is the one that results in a maximum expected lifetime-averaged return h, which would be the 4th strategy as shown here. Any number of different strategies can be compared in this way, representing every possible combination of decisions (could be thousands of different combinations under consideration), but the basic principle for finding the best strategy remains the same – find the combination that maximizes the total lifetime-averaged return h.

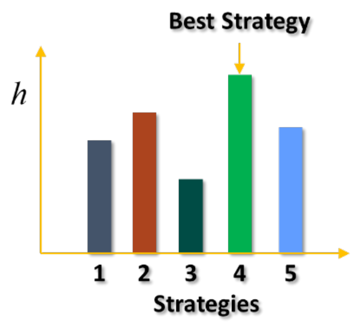

Figure 6 – Comparing different decision strategies to find the best one. Financially, the best strategy is the one that is expected to lead to the greatest lifetime-averaged return on investment, which is the 4th strategy here.

5 EQUITY'S EXPLAINABLE-AI BRAIN

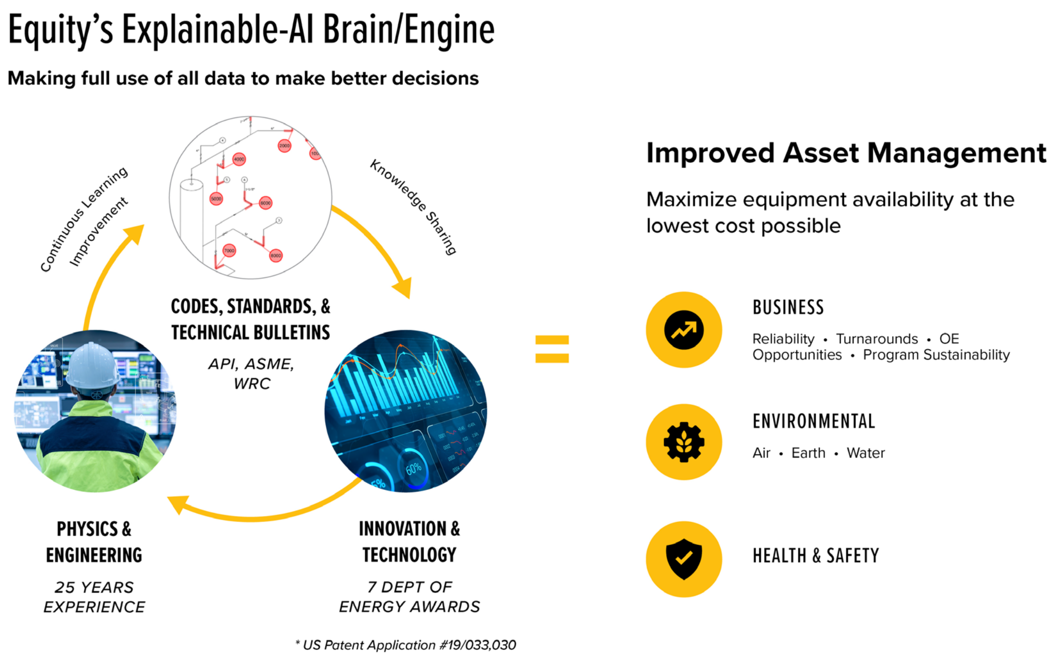

Figure 7 – The Equity BRAIN is the application of Equity's proprietary explainable-AI (XAI) technology (BENGI) to solve the industry's biggest problem of optimal asset integrity once and for all. It is the perfect union between human and machine intelligence that blends and shares knowledge coming from multiple sources to make the best predictions possible regarding damage progression and remaining life so that optimal actions and decisions are taken at precisely the right time. This is all done to maximize equipment availability at the lowest cost possible.

5.1 Solving Reliability Problems Using BENGI – Equity's Explainable AI Engine

In 2018, Equity started applying to Department of Energy (DOE) SBIRs to help build their advanced engineering analysis cloud software platform, The Equity Engineering Cloud (eec). Equity won its first two Phase 1 SBIR awards in 2019 for *High Performance Computing (HPC) for the EEC* and *CAN² - Canister Corrosion Analysis, Assessment and Action Plans): A Predictive Detection and Interpretation Software Platform for Life-Cycle Management of Spent Fuel Canisters*. *CAN²* was the first application of Equity's emerging explainable AI technology to solve a real-world problem (this was before BENGI was born). Phase 2 of *CAN²* was awarded in 2021. Equity has been awarded 7 SBIR grants so far for this and related research, totaling about $3.7 million.

Perhaps the most relevant SBIR award was for Phase 1 of BENGI in 2020. This kicked off Equity's formal development of their unique framework for building and solving industrial-scale, probabilistic, physics-based causal networks, also referred to as Bayesian networks or influence diagrams, that now form the basis of Equity's wide range of XAI solutions. Although these networks can be used to solve just about any problem, Equity's current focus is on using them to solve specific mechanical integrity problems – namely, predicting damage rates and making optimal decisions that achieve some desired financial objective. Phase 2 of BENGI was awarded in 2022. Some of the advantages of XAI versus traditional AI are listed in Figure 8.

Equity was awarded one patent so far in 2022 for the application of Bayesian hierarchical analysis to predict corrosion rates and failure of assets [1] and has applied for a second patent in 2025 for *Asset Lifecycle Optimization Systems and Methods* [2] – a broader application of BENGI technology to solve a wider range of mechanical integrity problems.

Recently, the inclusion of BENGI core technology in software to improve mechanical integrity across the board is being referred to as the *Equity BRAIN – Bayesian Reasoning for Asset Integrity using Networks*.

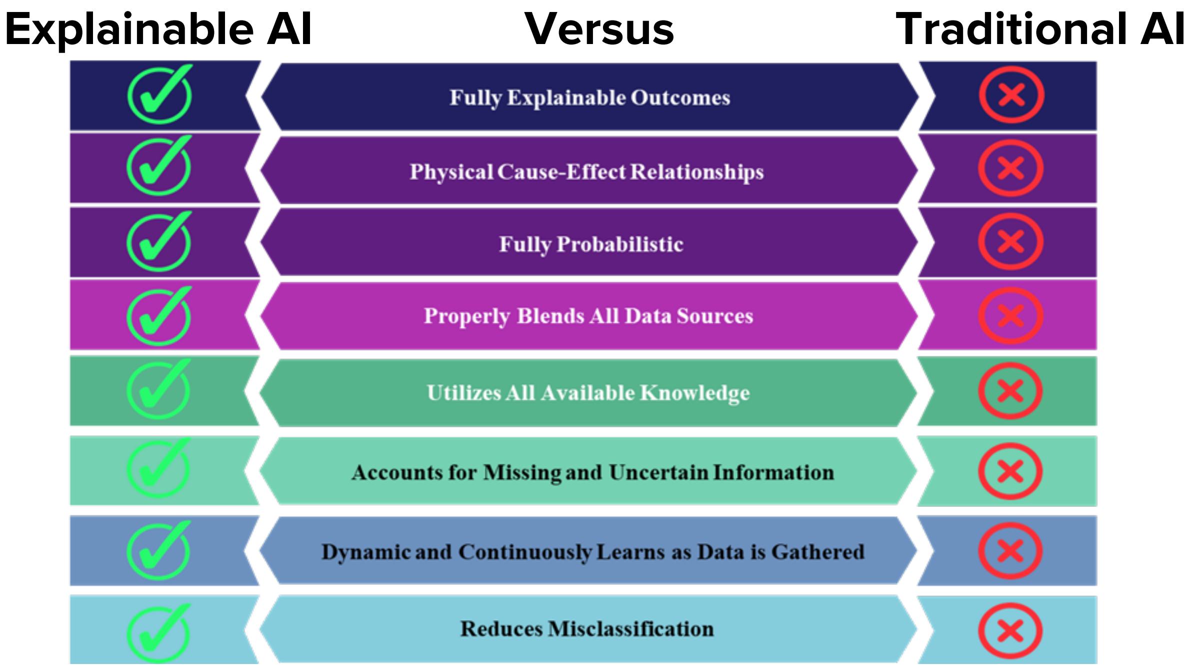

Figure 8 – An overview of the some of the key advantages of explainable AI (XAI) versus traditional AI.

5.2 Predicting Damage Rates and Failure Times

So, what are these probabilistic, physics-based, causal networks and how are they useful? Perhaps one of the simplest examples is shown in Figure 9, where a relatively small network predicts the probabilistic corrosion rate for naphthenic acid corrosion. One of the key features of these networks is their graphical depiction of causal relationships in a very digestible manner. It is very easy to see what factors are being considered as the causes of the corrosion rate – sulfur concentration, TAN, temperature, and velocity – by noting the direction of the arrows pointing from these multiple *causal* nodes towards the single corrosion rate *effect* node. This basic network structure is something typically set up initially by a human expert, but it can be learned automatically as well if there is enough data.

This network is much more than just a visual depiction of causal relationships. Behind the scenes, each node has associated with it a numerical conditional probability table (CPT) that specifies the probability of each possible state of that node given every possible combination of all the states of its parent nodes. This is where the real power of these networks comes to life and where knowledge coming from multiple sources gets encoded into the network. Specifying these CPTs in advance from known principles of science and engineering, and/or learning them from observations and/or data, is the key to these networks being so much more predictive. These networks also learn and get smarter over time as observations in the real world are made, as shown next in Section 5.3. All of this combined is what makes this an **explainable AI approach**. It is easy to understand the causal relationships leading to predictions from their visual nature, and the networks learn and get smarter over time, just like the best human experts do, as nodes are added for observations.

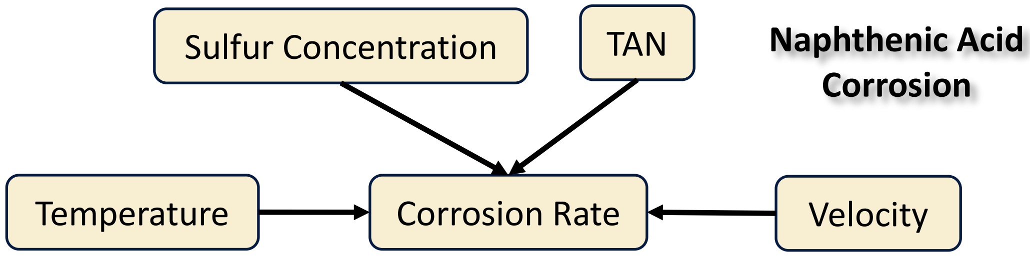

Figure 9 – Predicting the corrosion rate for naphthenic acid corrosion using a probabilistic, physics-based causal network. This is only a simple example that demonstrates the basic causal structure of such a network. Much more complex and expansive networks are used to solve real word problems.

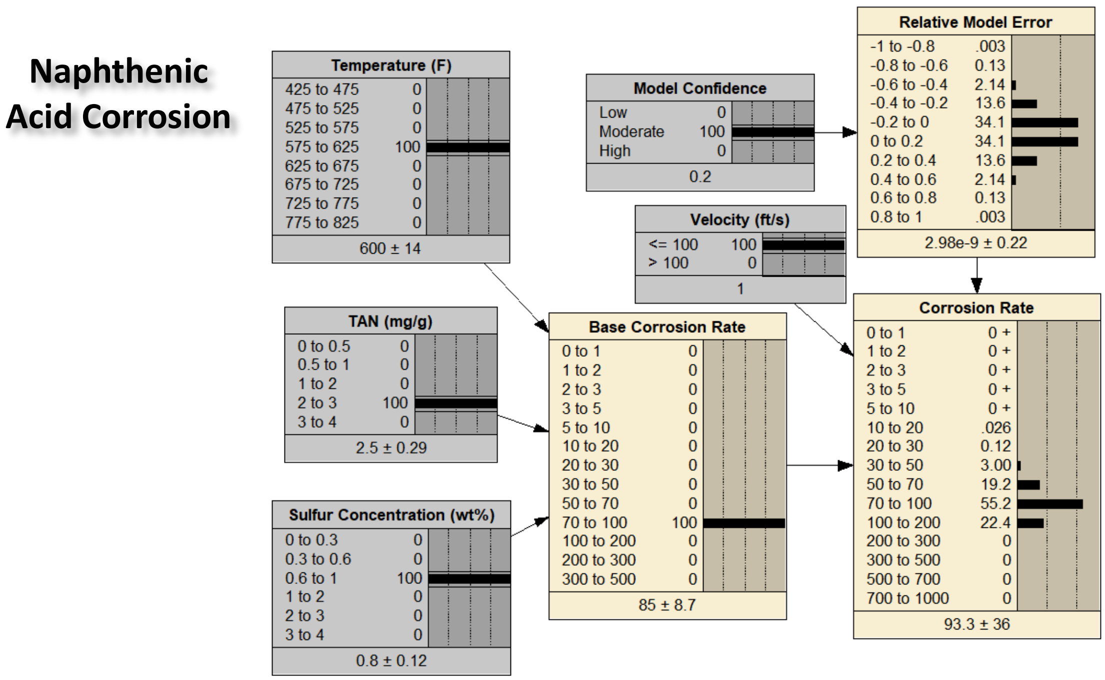

Figure 10 – A more detailed look at a slightly more complex extension of the network in Figure 9 that exposes the discrete states of the nodes and their respective probabilities (indicated by the length of the bars next to each state). The grey nodes have evidence set on them reflecting information known with certainty (the single state selected has 100% probability). The primary result is a predicted probabilistic corrosion rate distribution that also takes into consideration model confidence and error.

Figure 10 provides a more detailed look at a probabilistic, physics-based, causal network for predicting the corrosion rate due to naphthenic acid that now exposes the list of discrete states (possibilities) on each node along with their associated probabilities (as indicated by the length of the bars next to each state and the numerical probability printed to the left of each bar). This network encodes a relatively simple, rule-based method where selecting known values on the causal nodes leads to a precise *base corrosion rate*. Model error is then introduced separately to arrive at the final probabilistic corrosion rate prediction.

Many of these discrete states represent "min to max" ranges of some continuous variable that has been discretized into a finite number of such discrete ranges. The grey nodes have *hard evidence* set on them, meaning only one of their states is selected with 100% probability to specify inputs known with absolute certainty. It is also possible to enter uncertain input information as *soft evidence*, where more than one state is selected with probabilities that add up to one. Most of the CPTs in this network were set based upon predefined expert rules, but they could also be learned by training with large amounts of data, if such data were available. Since such data is not always available, it is the power of these networks that they encode and mix knowledge from multiple, disparate sources, whether it be from human experts, data analysis, physics-based models, or any combination thereof.

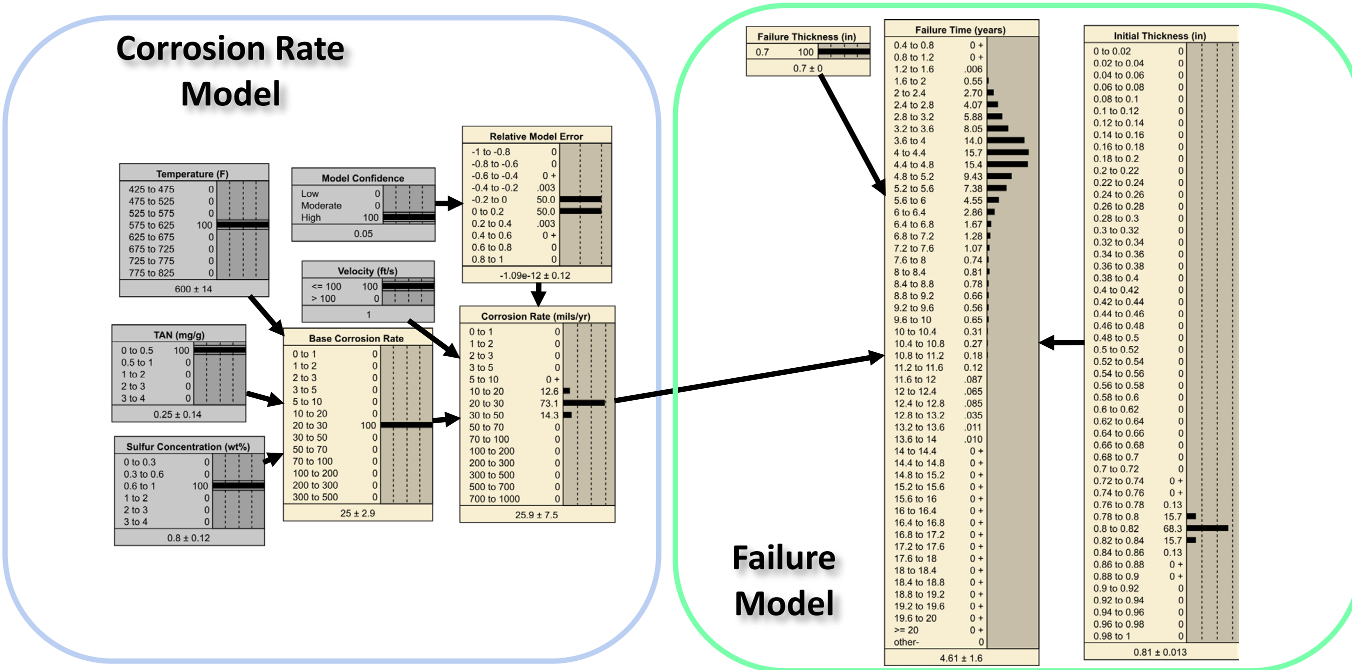

Figure 11 – An extension of the predictive network in Figure 10 that includes additional nodes for the failure time, initial thickness, and failure thickness. This is only one simple illustration of what in reality are much more complex models of damage rate and failure prediction, even though the basic concepts remain the same.

Once there is some predicted corrosion rate, additional nodes are added to the network to predict the probability of failure based on whatever failure criteria is chosen. This is done in a simple manner for demonstration purposes in Figure 11 by assuming there is a failure thickness below which failure will occur (and there is also an uncertain initial thickness). The discrete failure time probability distribution resulting from this is illustrated by the various non-zero bar lengths next to the multiple failure time ranges on the *Failure Time* node. At the bottom of the node, the mean failure time is reported to be 4.61 years with a standard deviation of 1.6 years.

The main purpose of networks like this, no matter how complex they become, is to predict failure time probability distributions like those shown in Figure 12, both before and after any operational, inspection or maintenance related action that could be performed. From these distributions, the life extension (or risk reduction) of these actions is inferred, and optimal decisions are then recommended by comparing the life extension and risk reduction of these activities to their cost.

Such networks demonstrate why this approach is also referred to as **next-generation RBI**. Once the failure time distribution from the improved predictive model is known, other nodes are added for the probability of failure (POF) at any time, which is just the CDF of this failure time distribution evaluated at that time, and if nodes are also added for the cost of failure (the consequence of failure, COF), then risk is easily calculated as the product of this POF and consequence. This allows these networks to be a drop-in replacement for the existing cruder POF models currently in API RP 581, while still allowing for the rest of that procedure to be followed, as is. However, the true power of this new approach is realized only when these improved predictions are combined with enhanced knowledge of how inspection and maintenance activities actually reduce risk, as discussed in Sections 5.3 and 5.4, and the cost optimization algorithms used to find the optimal decision strategy leading to a maximum return on investment given all this information, as discussed in Section 5.5.

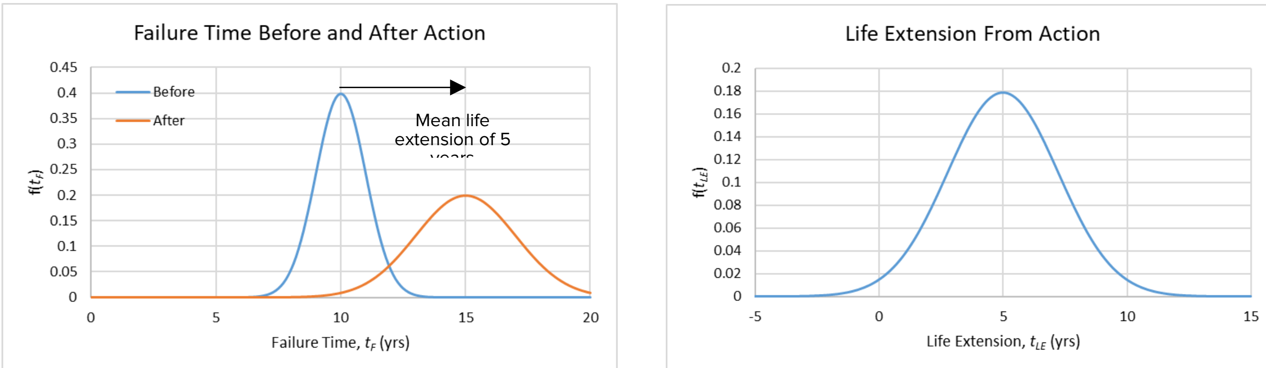

Figure 12 – Illustration of the primary desired result from these predictive networks, the failure time distribution both before and after any potential action is performed. The benefit of the action is then expressed in terms of either the risk reduction it provides over some specified time period, or its life extension as shown here. Beneficial actions such as this come at a cost, and optimal decisions are recommended by performing a subsequent cost-benefit analysis.

There are numerous advantages to using these sorts of networks to develop predictive models such as:

- **The visual nature of the networks makes them explainable** – easier to understand and critically examine model assumptions.
- **They are probabilistic** and properly account for all sources of uncertainty in the inputs as well as the outputs.
- **Inputs can be missing or uncertain** – for missing inputs, the model falls back to using the prior distributions specified by a human expert or from previous data or experience. For uncertain inputs, values are specified using soft evidence (specifying multiple states with probabilities that add up to one).
- **Human expert opinion is accounted for** by setting prior probabilities or by adding additional nodes.
- **An uncertain consequence is accounted for** by entering soft evidence on failure cost nodes.
- **Different sources of knowledge are blended together seamlessly** (predictive models, expert opinion, real-world data and observations) to come up with a single source of truth, as opposed to just picking one or the other.
- **The networks learn from experience** by adding additional nodes for real-world observations (inspection results, sensor readings, etc.) that may differ from prior predictions, making future predictions more accurate (this is also be done by updating prior probabilities with posterior probabilities after new observations are accounted for) – this continual learning is a sign of intelligence.
- **Complex models are built fairly easily** as long as the causes of the problem are fairly well understood.
- **Older rule-based models are given new life by converting them to networks**, since all of the above advantages then apply to these older models and enhance their usefulness while still retaining their encoded historical knowledge.

As an example of converting an older rule-based method into a network, consider the predictive network for CUI likelihood shown in Figure 13 that was built from the simple rules presented in API RP 583 (this network was not in API RP 583 but the rules were laid out in its text). Just by building this visual network, the factors influencing the final result are more clearly exposed and understood (allowing for discussions with human experts to critically evaluate model assumptions and discover possible enhancements). In some sense, all the predictive rules in this entire, lengthy document have been converted into a simple, compact visual representation that fits on less than a single page, exposing the model's rules, causal factors, and assumptions in an easily digestible manner. This breathes new life into older methods, while still retaining their essence.

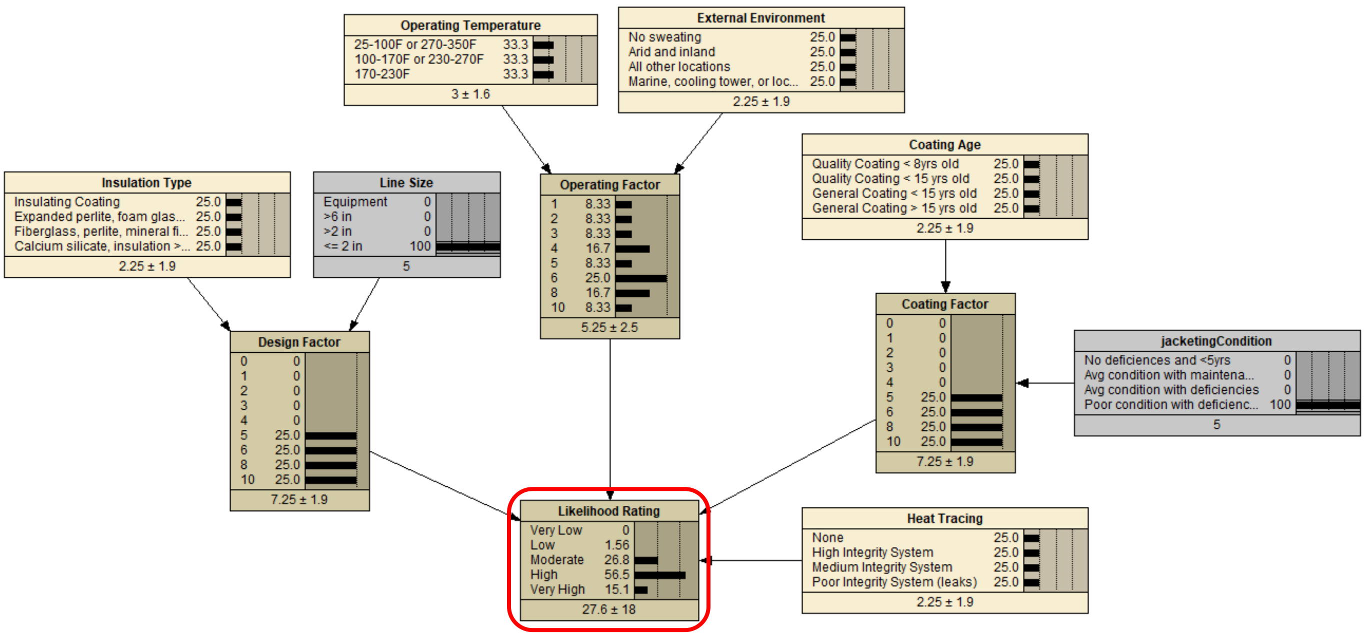

Figure 13 – A probabilistic causal network representing the simple rule-based method in API RP 583 that results in a final likelihood rating for CUI (encircled by the red rectangle).

Another immediate advantage of converting older rule-based methods to probabilistic networks like this is that now missing or uncertain inputs can be entered, and proper probabilistic blending is still used to obtain a final uncertain, but still useful, result. For example, hard evidence is only entered on the *Line Size* and *Jacketing Condition* nodes in Figure 13. The rest of the input nodes are left unspecified, leading to the use of their prior distributions as shown, which may come from previous experience or knowledge but here are mostly set to be uniform distributions (reflecting no specific prior knowledge that any one value is more likely than another). This results in an uncertain likelihood rating with a most-likely value of *High* but with some possibility it could also be *Moderate* or *Very High*, as indicated by their respective probabilities. This uncertain result could still be useful for future decision making.

5.3 Incorporating Inspection Results

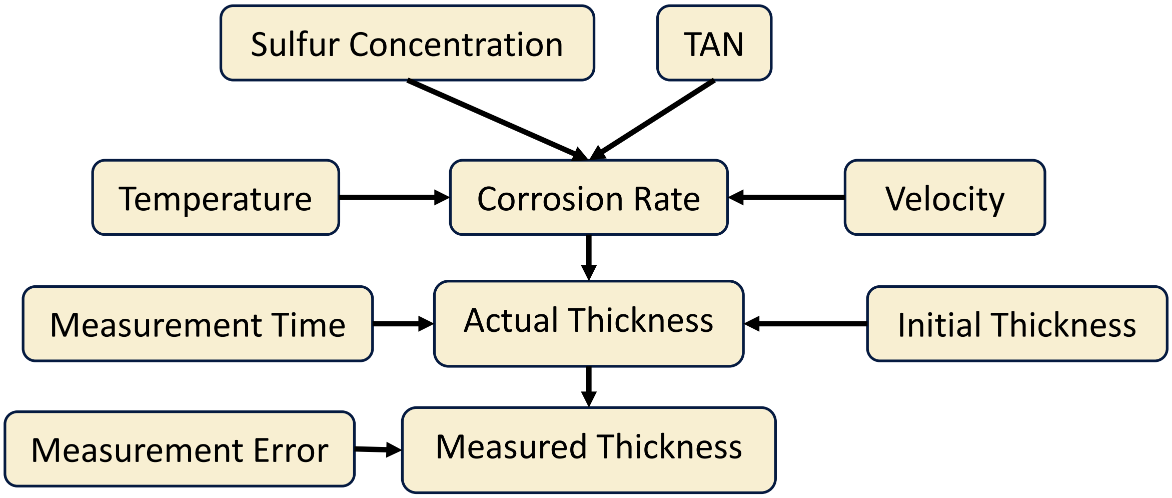

Figure 14 – An extension of the network in Figure 9 to account for a single thickness measured during an inspection. Additional nodes are added to account for the time at which the measurement was taken, the error of the measurement technique, the initial thickness, and the measured thickness value itself. The actual thickness node is what the thickness is predicted to be at the time of the measurement as a result of the predicted corrosion rate. Just by accounting for this single real-world observation in the network, future model predictions will become more accurate. This helps correct any discrepancies there may be between what was predicted before the measurement and what was observed afterwards. This is a key attribute of these networks – they learn and get smarter from on-the-job experience, just like human experts do.

Just by adding additional nodes to the predictive network in Figure 9 to account for a measured thickness taken during a single inspection at some later time, as shown in Figure 14, the prediction of the corrosion rate immediately becomes more accurate and updates the corresponding POF curve that makes use of it, as shown in Figure 15. This is one of the key attributes of these networks that allows them to be classified as AI – **the networks learn and get smarter over time from on-the-job experience, just like human experts do**. Additional nodes are added for every inspection result, using any method or technique, all in the same network, and the model learns and gets smarter (makes more accurate predictions) after incorporating each one.

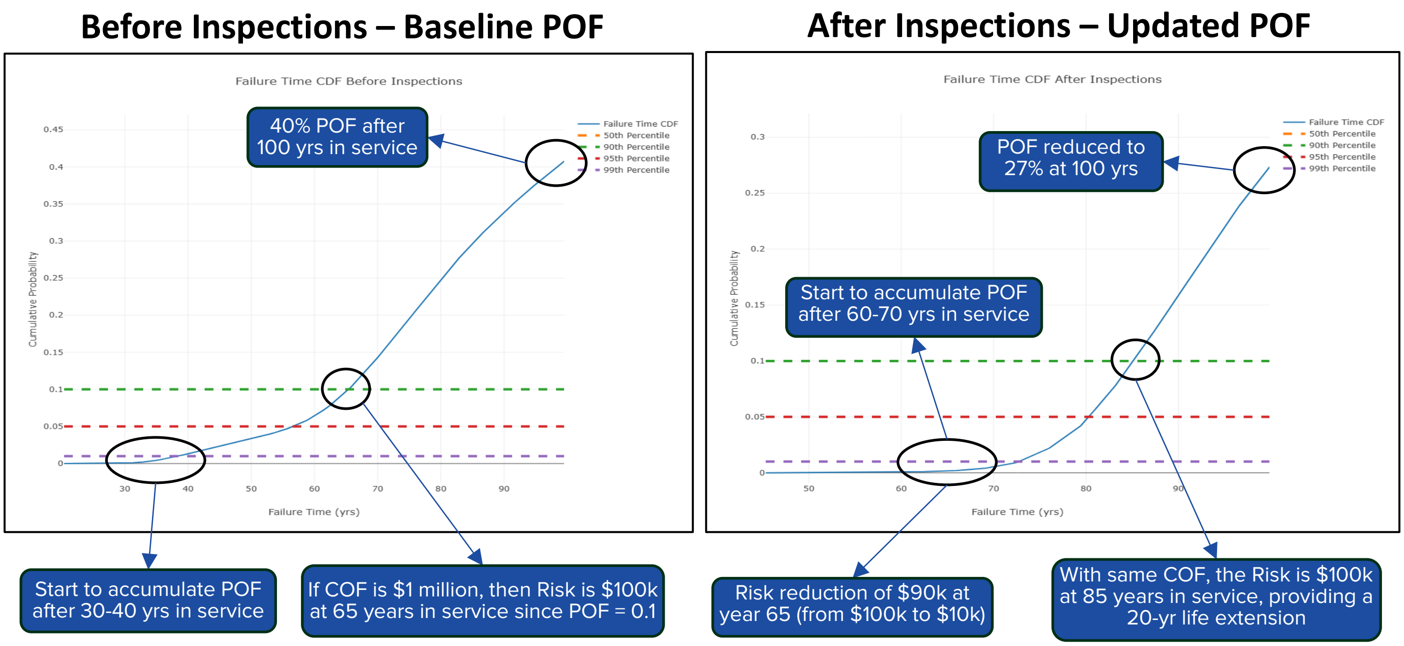

Figure 15 – Accounting for inspection results in an augmented predictive network model like that shown in Figure 14 leads to an updated corrosion rate and, subsequently, an updated POF curve like that shown here. This is for the case where the measured metal loss at the time of inspection is significantly lower than what was predicted, leading to a reduction in the expected corrosion rate that results in a life extension of approximately 20 years (and a risk reduction at year 65 of $90k). This demonstrates the power of these networks to properly account for the beneficial effect of even a single inspection result through proper probabilistic inference. This is an example of making full use of all available data to improve predictions.

What the network in Figure 14 represents is a blending between two sources of knowledge: 1) a predictive physics-based causal model, and 2) inspection data. The single corrosion rate probability distribution (the single source of truth) is updated from both sources of knowledge simultaneously and blended appropriately based on the relative confidence in each source. That is, if one is very confident in the model predictions versus the measurements (due to, perhaps, a large measurement error), then more weight is given to the model predictions, as one would expect. On the other hand, if the uncertainty in the model predictions is much larger than the measurement error, then more weight is given to the measurement. This all happens automatically though the proper laws of probabilistic inference that are encoded in the network structure and the CPTs of each node. This network could be expanded further to blend in a human expert's opinion as well (the third major source of knowledge) by setting prior probabilities or by adding additional nodes appropriately.

5.4 Quantifying Inspection Effectiveness

Another problem solved with these networks is quantifying inspection effectiveness for the purpose of more accurately determining the risk reduction of inspections. One way of doing this was already incorporated into the network of Figure 14, by accounting for measurement error. Inspection methods with larger errors are generally less effective. For inspections focused on *detecting* damage, this is captured by using a similar node for the probability of detection (accounting for both false positive and false negative errors). Going further, when only part of an asset is inspected, the additional question always arises as to what is the risk of damage still being present in the uninspected regions even though none was found within the inspected regions? This question is answered by using networks like that depicted in Figure 16. This network is able to predict the likelihood of having damage in the uninspected regions given that some number of the inspected regions were found to contain damage.

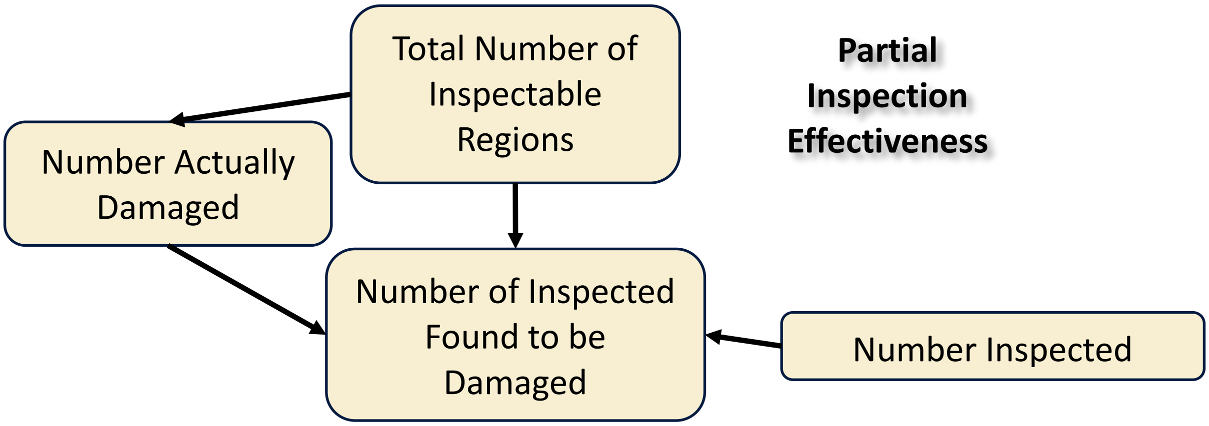

Figure 16 – A simple network used to properly quantify risk reduction from partial inspection coverage. The total inspectable surface area is subdivided up into some smaller number of inspectable regions, and then only some of these are inspected, leaving other regions uninspected. The inspection result is the number of inspected regions that contained damage. The question is then, how likely is there to be damage in the remaining uninspected regions? This is predicted by the probability distribution on the Number Actually Damaged node. From this, the risk reduction due to the partial inspection is inferred from the Number Actually Damaged node, which accounts for both inspected and uninspected regions.

Networks like this are ultimately used to obtain curves of risk reduction versus the number of inspections that are then compared to the cost per number of inspections as shown in Figure 17 to find the optimal number of inspections to perform.

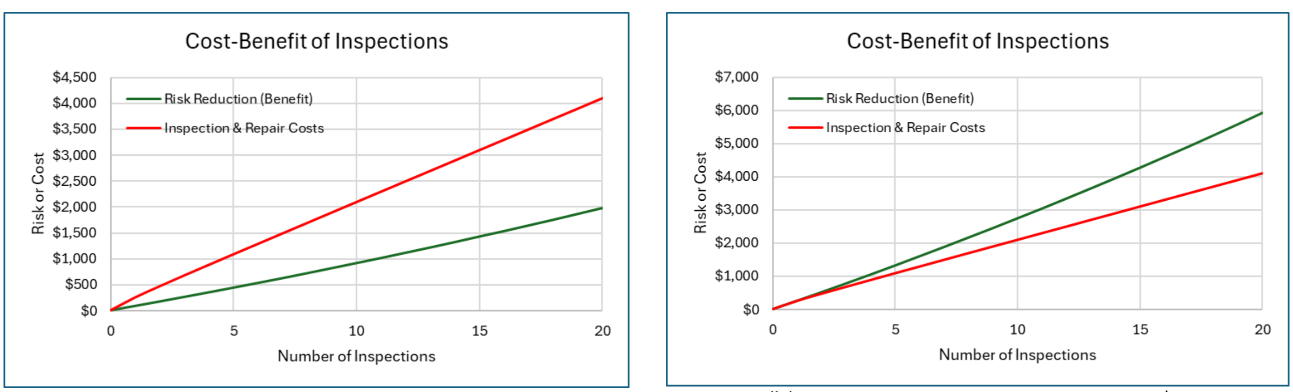

(a) Total Risk Without Inspections is $2,000

(b) Total Risk Without Inspections is $6,000

Figure 17 – Networks like those in Figure 16 are used to obtain curves of risk reduction versus the number of inspections (green curves) that are compared with the cost versus the number of inspections (red curves) to determine the optimal number of inspections to perform. This is for a case where the total inspectable area is subdivided into 20 total smaller inspectable regions. If all 20 regions were inspected, this would be full coverage, and it is assumed that all damage would be found, and the maximum risk reduction achieved. In general, when the total inspection and maintenance costs for full coverage are greater than the total risk (the maximum possible risk reduction), as is the case in (a), then no inspections are worth it at that time, because the inspection cost is always greater than the nearly-linearly increasing risk reduction curve no matter how many inspections there are. However, when the total risk increases later in time as in (b), it now turns out that the maximum return (risk reduction minus cost) is achieved by full inspection coverage (all 20 inspections). Since risk generally increases over time, this means there is some point in time when the optimal decision switches from no inspections to full inspection coverage for this scenario.

5.5 Putting This All Together to Make the Best Decisions

The ultimate use of these networks is to help operators make better decisions regarding inspection and maintenance. This is achieved by adding decision and cost nodes to these networks, as shown in Figure 18. Once this is done, the network automatically finds the best decision strategy as the one that leads to the lowest total cost (maximum utility, maximum return, or maximum profit).

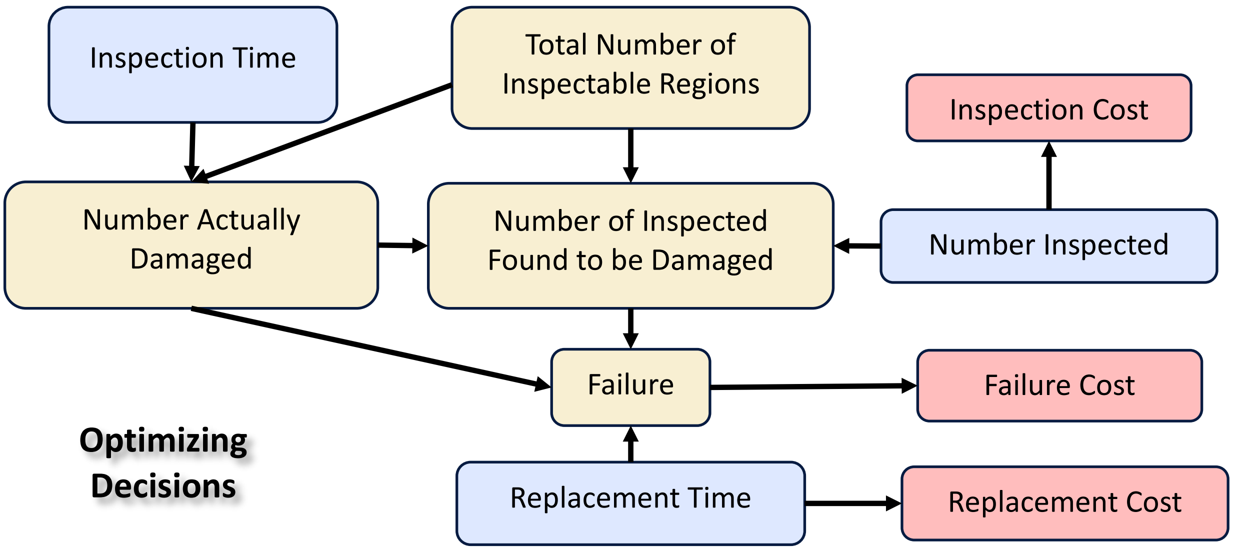

Figure 18 – The predictive network in Figure 16 is augmented to include decision (blue) and cost (red) nodes to find optimal decisions (those that minimize total expected cost in this example). The cost of failure, replacement, and inspection are accounted for (and dependent upon the number of inspections). There are three decisions: When to inspect, how many inspections to perform, and when to replace after the inspections, which depends on what is found during the inspections.

A practical example of using decision networks to find optimal decisions is shown in Figure 19, where several different decisions are evaluated with regards to CUI maintenance using a network similar to the one in Figure 18 but tailored for CUI. The results are presented in a table format so the user sees which is best (the one that minimizes the total expected cost or maximizes the expected return on investment). These calculations are embedded within the CUISight software described later in Sections 6 and 7.

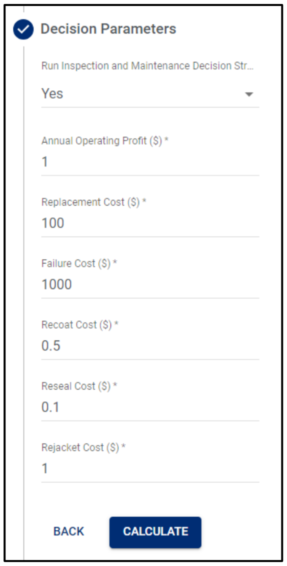

Strategy	Recoat	Recoat Time (yrs)	Rejacket	Rejacket Time (yrs)	Reseal	Reseal Time (yrs)	Replace Time (yrs)	Total Utility ($)
Replace Only	N/A	N/A	N/A	N/A	N/A	N/A	30.000	-2.906
Recoat Only	Yes	15	N/A	N/A	N/A	N/A	40.000	-1.876
Rejacket Only	N/A	N/A	Yes	20	N/A	N/A	30.000	-2.526
Reseal Only	N/A	N/A	N/A	N/A	Yes	10	30.000	-2.681
Recoat and Rejacket	Yes	15	Yes	15	N/A	N/A	40.000	-1.908
Recoat and Reseal	Yes	20	N/A	N/A	Yes	20	40.000	-1.848

This decision provides a 10-year life extension in the most cost-effective way

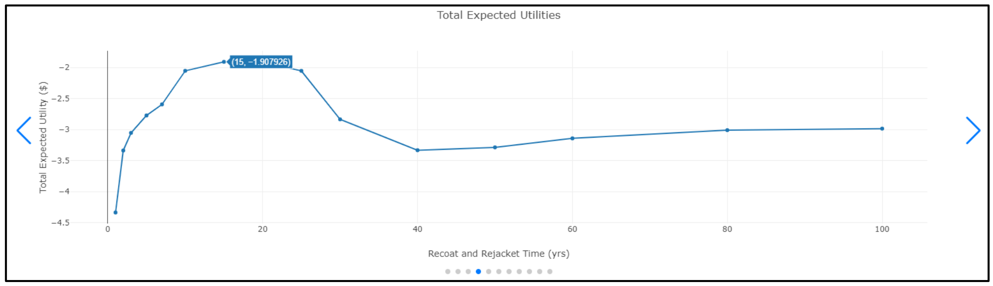

Figure 19 – Optimal decision-making for CUI management consists of evaluating several different maintenance decisions and then finding the best one, which is here to recoat & reseal after 20 years. The optimal time for performing this activity is found by evaluating the total expected utility (benefit minus costs) at each possible time to find the time that maximizes this utility, which occurs for this particular activity after about 20 years as shown in the graph. This maximum utility is higher than the maximum expected utility for every other possible decision, which is why it is the best overall decision to make.

5.6 Process Optimization

The applications of the Equity BRAIN are not only limited to mechanical integrity. Since the overarching goal of asset lifecycle optimization is to maximize the return on investment over the lifetime of the asset, a holistic approach also needs to take into consideration optimization of the process, from which all benefits due to the operation of the asset are derived. This is typically the revenue acquired through the sale of some end product produced by the asset, but it could be anything else. This benefit of operation is accounted for by the green boxes in the life-cycle diagram of Figure 3 and depends upon all design and operational decisions that are made and their effect on mechanical integrity. This interplay between operations and maintenance (O&M) must be accounted for in a holistic way to make accurate predictions of the total expected return on investment over the asset's lifetime.

For example, in the renewable energy field, selecting which renewable feedstock to use is an important problem, and market considerations, government incentives, and availability often drive that choice. Unfortunately, the effect of various feedstocks on the mechanical integrity of equipment is often not well understood or taken into consideration, and this can become especially deleterious due to the impurities present in some feedstocks such as free fatty acids (FFAs), which leads to greatly elevated and hard to predict corrosion rates. The Equity BRAIN is also able to optimize process related decisions like feedstock selection while also addressing mechanical integrity concerns to maximize the entire return on investment.

This is accomplished by simply incorporating other decision nodes into the causal networks like that shown in Figure 19 to account for operational decisions and their costs and benefits along with the costs due to inspection and maintenance. Then, when these networks find the optimal decision strategy, all operational, inspection, and maintenance decisions will be optimized together holistically to maximize the overall expected return on investment. In the end, the optimal overall decision strategy must account for everything, or it will not provide reliable recommendations.

6 SOFTWARE

Equity has integrated all of this technology into its **Equity Engineering Cloud (eec)** environment in the form of several different interconnected software tools that form the building blocks of this new offering as shown in Figure 20. The eec is Equity's first cloud software platform that was launched in 2017. Within the eec, an APM layer that provides advanced capabilities for maintaining asset health has been incorporated known as **HealthSight (HS)** – providing new insights into asset health. HealthSight connects to third-party asset and sensor systems in order to bring data into the system that is required for advanced analysis and decision making. Since damage is what causes health to degrade, within HealthSight is a calculator called **DamageSight (DS)** that predicts the damage rates for at least every damage mechanism in API RP 571 using the explainable AI methods discussed previously. Within DamageSight is a damage mechanism specific module known as **CUISight (CS)** that focuses on one particular costly damage mechanism, Corrosion Under Insulation (CUI).

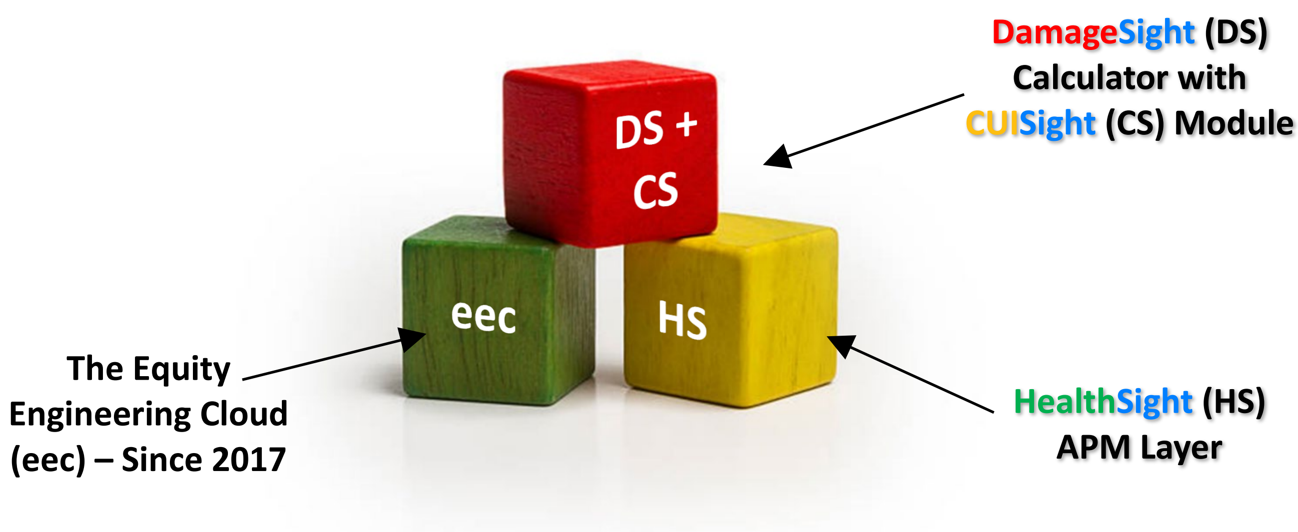

Figure 20 – The building blocks of Equity's new XAI software solution for improved MI are shown here. The Equity Engineering Cloud (eec) has been operational since 2017 and is Equity's cloud software platform for solving a wide variety of engineering related problems. Integrated within the eec is HealthSight (HS), an advanced APM layer for keeping assets healthy. Within HealthSight is DamageSight (DS) – a calculator for predicting damage rates from every possible damage mechanism. A special module within DamageSight called CUISight focuses attention on one particular damage mechanism of great importance, CUI.

6.1 HealthSight

The HealthSight system is designed to *bridge the gap between data and decisions* by taking input from existing asset systems, DCS process sensors, NDE sensors, scheduled inspections, and every other possible source of knowledge and producing output that consists of damage rate predictions, remaining life, and optimal recommended decisions and actions as shown in Figure 21.

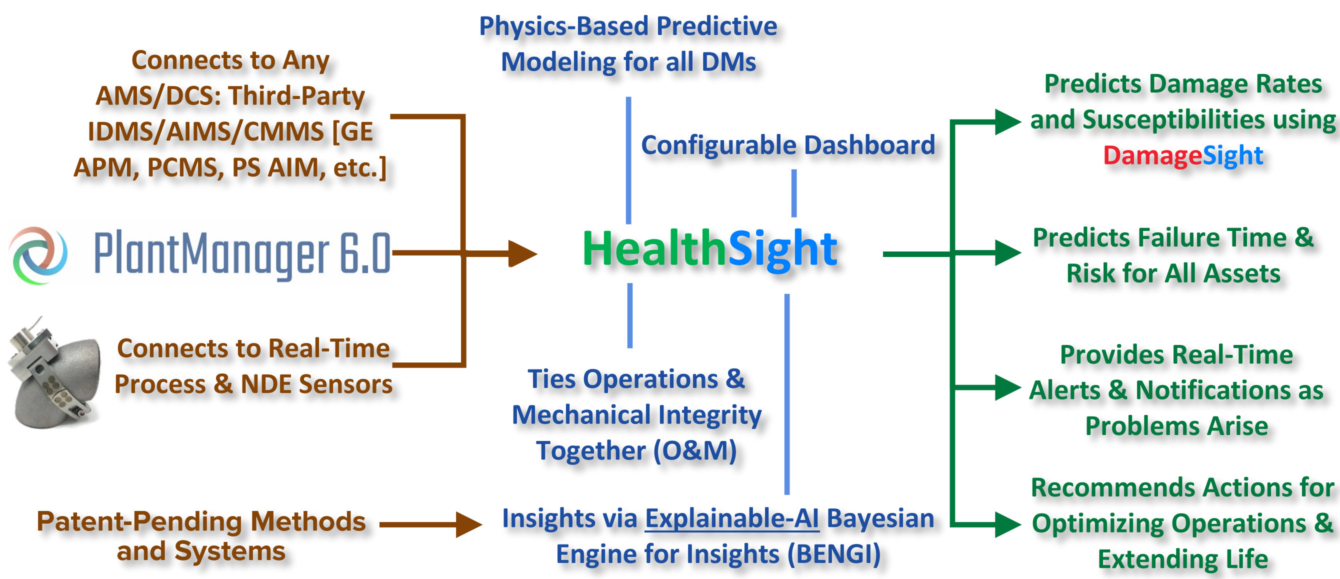

Figure 21 – HealthSight is a smart asset performance management (APM) layer that sits on top of any asset system to provide recommendations for optimal decisions related to design, process, operation, inspection, maintenance, and integrity by monitoring the health of assets in real time using explainable AI methods for improved predictions and to perform cost-benefit analysis.

The particular way in which HealthSight functionality has been seamlessly integrated into the eec user interface menus is shown in Figure 22.

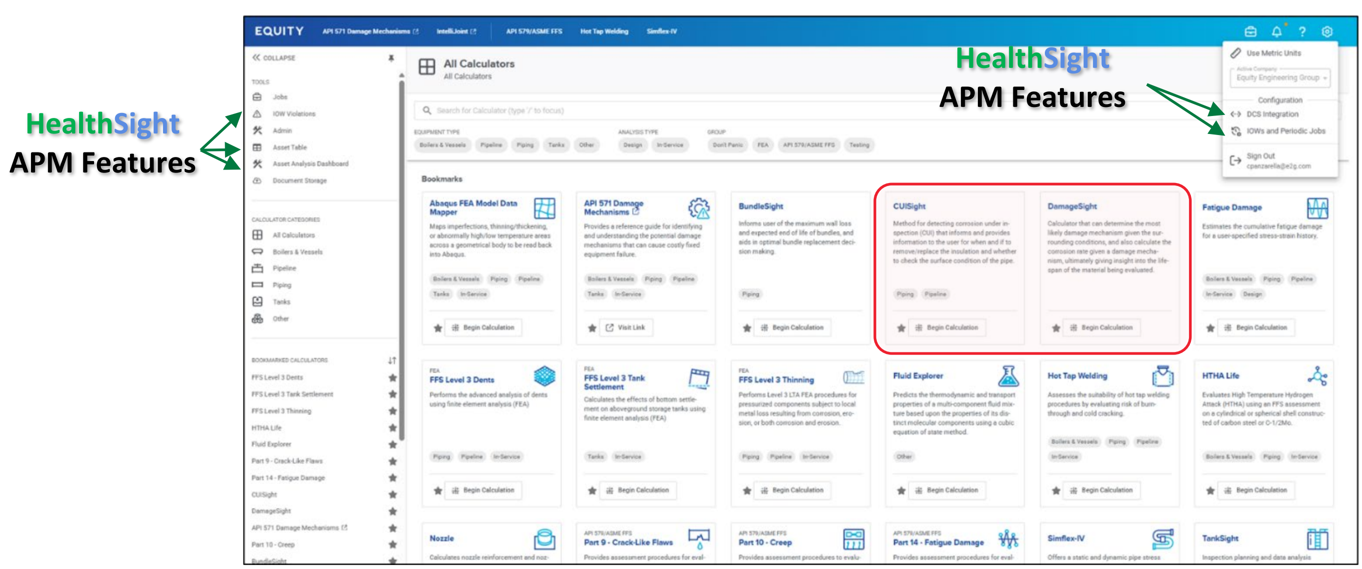

Figure 22 – The HealthSight APM layer integrated seamlessly into the Equity Engineering Cloud (eec) provides additional functionality to connect to third-party asset systems to acquire the necessary data for performing real time calculations on those assets to predict damage rates and remaining life. In order to get the most accurate predictions, these calculations are fed live sensor data from connected DCS systems and embedded NDE sensors. Based upon these continuously updated predictions, any IOW violations (on either the raw input data or on calculated results like accumulated damage or remaining life) are reported in real time, and alerts are raised when situations arise that require human attention. All of this is presented to the user on a user-friendly dashboard that also recommends optimal actions and decisions for achieving maximal return on investment.

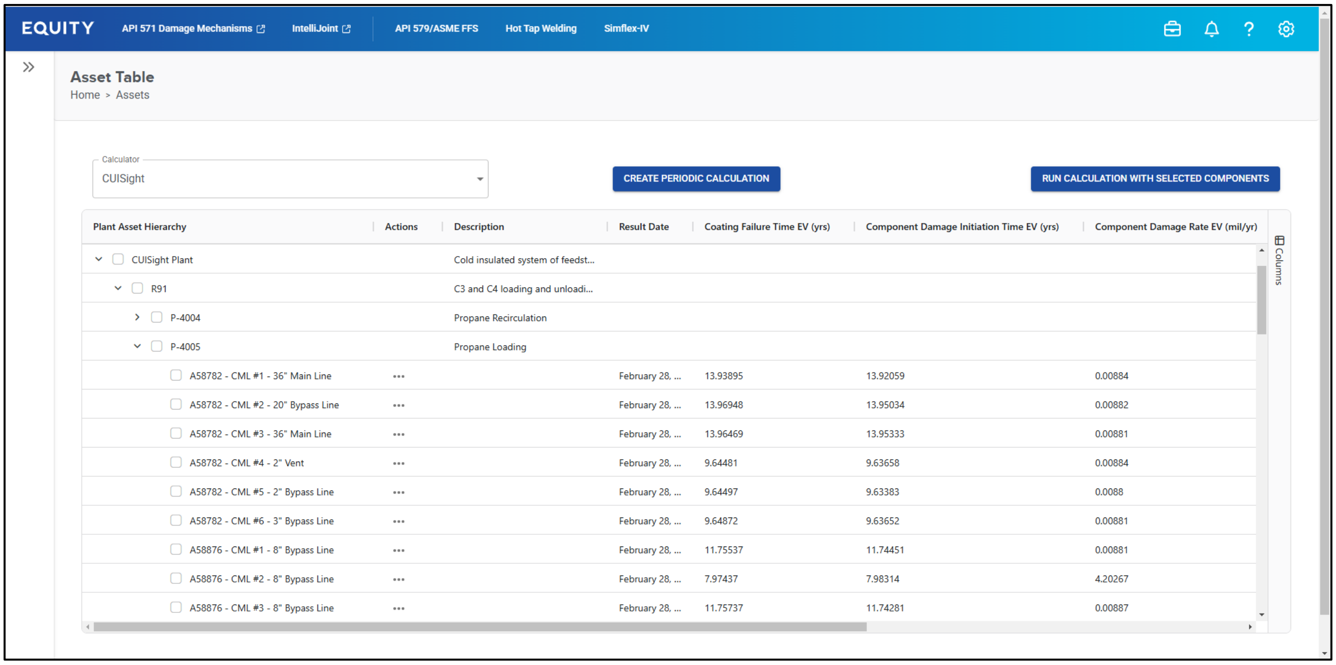

Figure 23 – The asset table built into HealthSight pulls data in from any asset system (Equity's PlantManager ASSET or any other) in order to get the inputs required to perform the desired calculation. This example is configured to run the CUISight corrosion-rate calculator on all these assets automatically in real time (recalculations are triggered as soon as any new input data is received from a sensor or new scheduled inspection to always ensure the most up-to-date predictions).

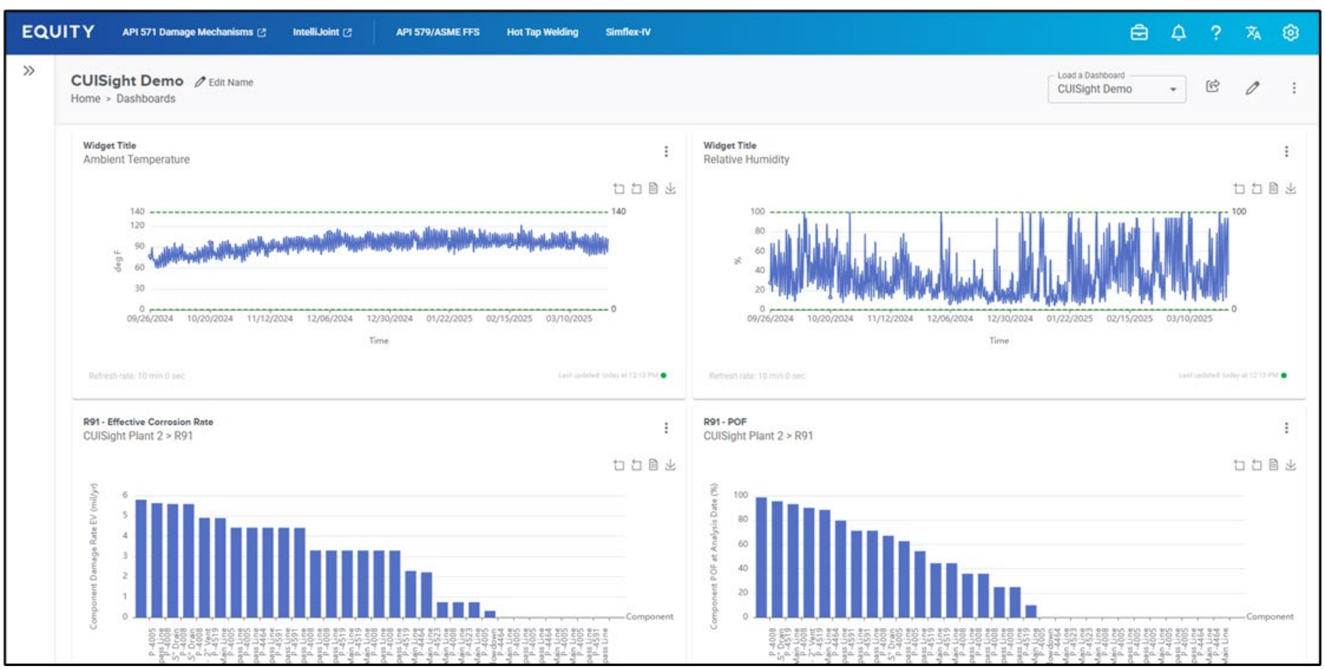

Figure 24 – The configurable HealthSight dashboard is capable of displaying raw sensor data from DCS systems or NDE sensors as well as calculated results such as damage rates, failure times, or the probability of failure (POF). Here, it is shown how calculations can make use of real-time local weather data such as ambient temperature and relative humidity to calculate the effective corrosion rate and POF for CUI on all assets, ranking them according to these calculated results from worst to best.

6.2 DamageSight

The plan for DamageSight is to improve the predictive models for every damage mechanism in API RP 571 and beyond using the patent-pending methods of XAI that Equity has developed. This is depicted visually in Figure 25 using a few simplified probabilistic, physics-based, causal networks (final networks are much more complex).

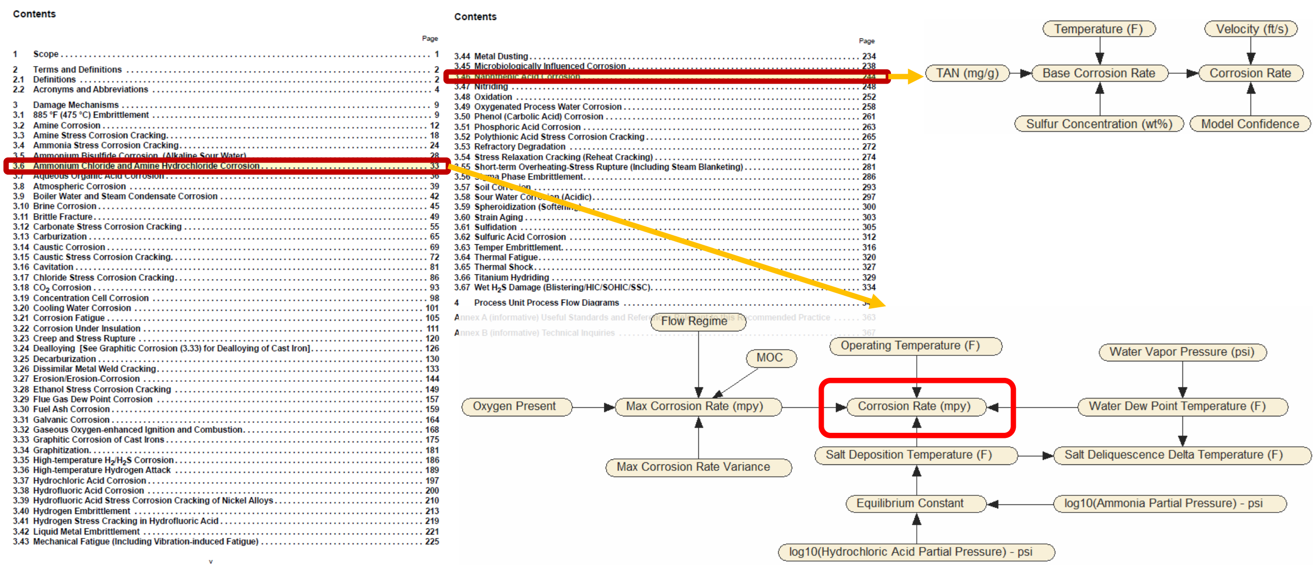

Figure 25 – DamageSight will have improved damage rate calculators for all of the damage mechanisms in API RP 571 and beyond. These mechanisms are being prioritized by importance to Equity's clients. For each mechanism, one or more probabilistic, physics-based, causal networks are being developed. The first damage mechanism to be thoroughly investigated and improved is CUI, resulting in the CUISight module of DamageSight.

6.3 CUISight

The CUISight calculator predicts improved damage rates, metal loss, thickness projections, and probability of failure for assets subject to CUI damage. It is discussed in more detail in Section 7.1.

7 END-TO-END SOLUTIONS

7.1 CUISight Solution as a Service

Since [CUI is a $5 trillion global problem for the industry](), and there was no satisfactory solution until now, it was felt that CUI would be the most-beneficial end-to-end application of all this technology out of the gate, providing the industry with a much better solution right away using the latest emerging technology in the field of XAI. An overview of the improved predictions of this approach and its ultimate value for one of the early adopters of this groundbreaking technology is shown in Figure 26. The bottom line is that Equity's XAI technology leads to more accurate predictions, even with just currently available data (no additional data was needed for these cases beyond what was already readily available), although improved data collection with more prescriptive practices and guidance leads to even better predictions. These more accurate predictions would have led to a total potential savings of $38 million, as indicated, if this system had been adopted by this client from the start (who unfortunately did not have access to this technology before). Going forward, this client will realize similar future savings by adopting this approach at all other facilities and integrating it with their own best practices and standards (which can be further improved by Equity).

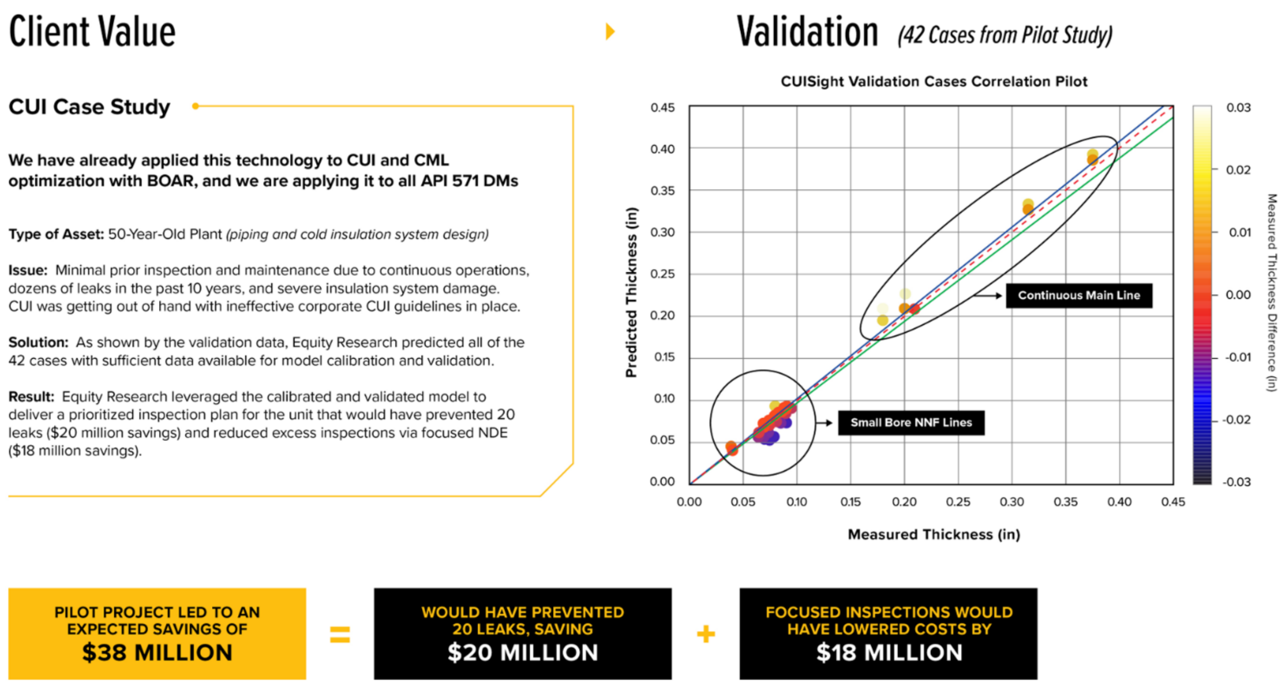

Figure 26 – The application of Equity's new XAI technology to improve predictions of metal loss from CUI led to an expected total savings of $38 million dollars for one of the early adopters of this technology. This savings is a combination of both having fewer failures (savings of over $20 million) and having more focused inspections (savings of over $18 million). In general, it is always true that savings is achieved from either or both of these two buckets (having fewer failures or reducing unnecessary inspections). This same approach is being applied to every other damage mechanism in API RP 571 as well (and beyond).

These improved predictions are the result of using a core probabilistic, physics-based, causal network for CUI like that shown in Figure 27 with many auxiliary networks (not shown here) to account for all available information and knowledge in the proper way so that the model learns and gets smarter over time – **the power of XAI**.

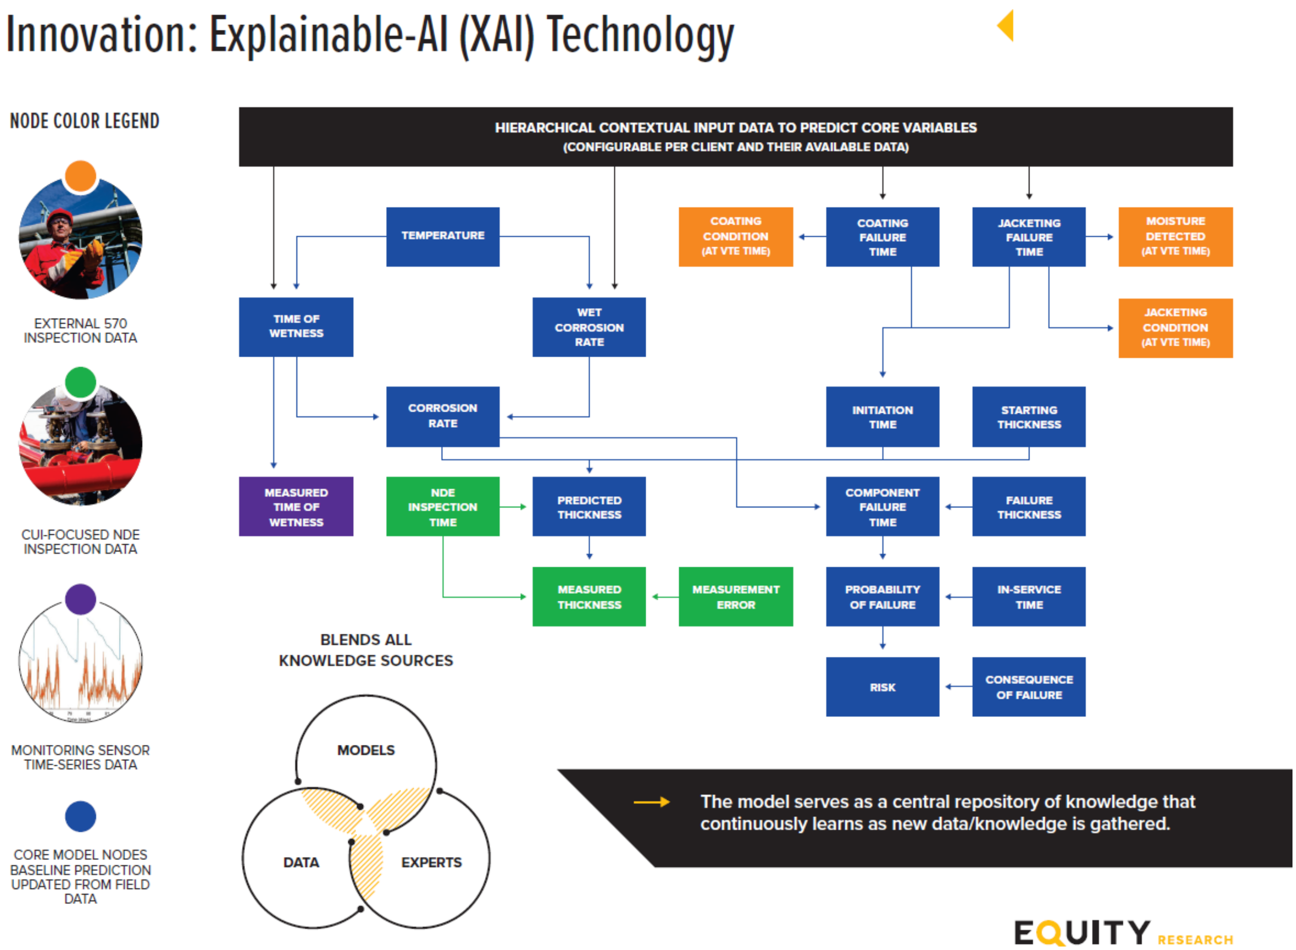

Figure 27 – The primary causes of CUI are moisture getting in under the insulation and coming into direct contact with bare metal once the coating has degraded enough to allow this to happen. For water to get in, the jacketing must first fail. Even once water gets in, the amount of time that water makes direct contact with the metal is not known precisely, but this *time of wetness* is another important factor, along with the temperature-dependent corrosion rate, that directly contributes to future metal loss. Even though these variables are not known precisely, and are treated as random variables, predicting them probabilistically as best as possible allows for better overall predictions of metal loss, and, ultimately, better decisions to be made. Inspection data is incorporated into the network to make it smarter by adding additional green nodes (CUI-focused NDE inspection data), orange nodes (external API 570 inspection data), and purple nodes (live time-series data from monitoring sensors). These nodes remain in the network while making future predictions, allowing these networks to learn over time and get smarter from on-the-job experience just like some human experts do – this is Equity's unique XAI approach to MI in a nutshell. To predict risk, additional nodes are added for the probability of failure and consequence. This allows one to make risk-informed decisions such as optimizing the return on investment if additional cost information is available for all possible inspection and maintenance activities. Alternatively, this improved risk model could be used as a drop-in replacement for what's currently in API RP 581, allowing that standard approach to be used for everything else but with higher quality predictions of risk.

The overall solution provided to Equity's clients is frequently offered as a complete service that not only makes use of the CUISight software for improved predictions but also provides additional expert human guidance to ensure successful implementation. This overall process is outlined at a very high level in Figure 28.

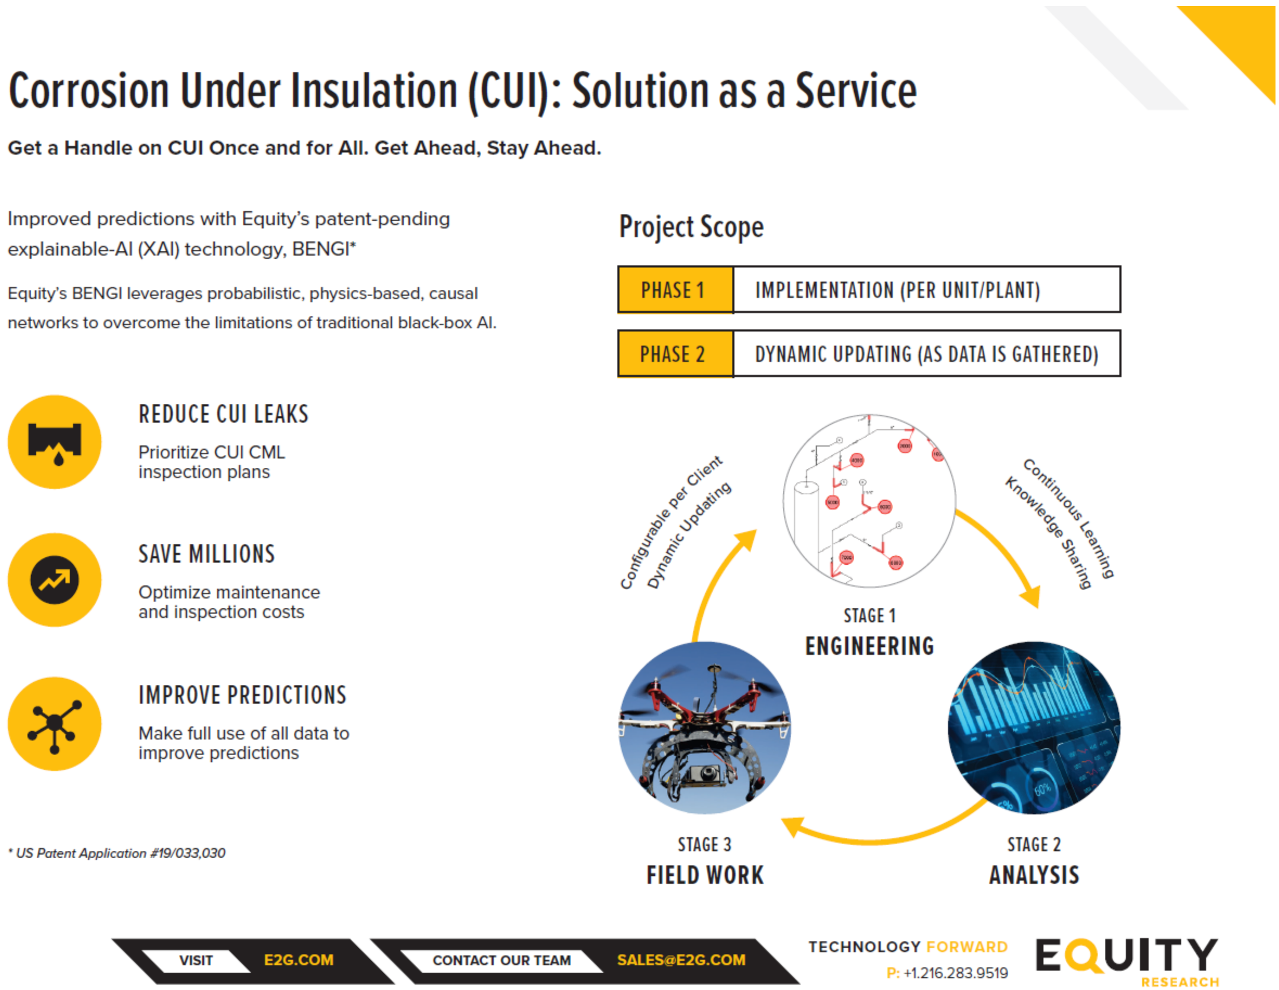

Figure 28 – A high-level depiction of how this new explainable-AI (XAI) solution fits into the overarching CUI Solution as a Service offering by Equity to improve predictions, leading to a reduction in leaks at the lowest cost possible. There are three primary stages of implementation that repeat in a cyclic fashion over the lifetime of the asset. The XAI fits mostly into Stage 2: Analysis in order to plan future inspection and maintenance activities more effectively, after engineering assessments of inspection data (or from design if just starting) are made in Stage 1.

While implementing this procedure as a consulting service, one deliverable is often a prioritized table of *CMLs for CUI* to be inspected, ranked by severity, taking into account the cost of inspection if such data is available (so that money is not wasted on unnecessary or low-risk inspections) as shown in Figure 29. An important quality of this approach is that existing asset systems capable of managing CMLs can still be used, mostly as is, by making use of the *CML for CUI* concept that Equity has introduced. These CMLs for CUI are separate from the CMLs for internal corrosion, so whatever asset system is used must be capable of having separate internal and external circuits at the same time to get the most overall benefit.

Deliverable: Prioritized CUI CML Inspection Plan

CUI CMLs are Dynamic

Line Number	Drawing Number	CML Number	Process Fluid	Inner Diameter (in)	Corrosion Allowance (in)	Failure Time at Target Min POF (yr)	Max Effective Damage Rate EV (mil/yr)	Initiation Time EV (yr)
P-4592	NB-A75583	5	Butane	1	0.089	14.492	5.576398078	7.421801993
P-4592	NB-A75583	4	Butane	1	0.089	14.493	5.575961538	7.420597816
P-4594	NB-A75585	5	Butane	1	0.089	14.495	5.58160674	7.42157892
P-4595	NB-A75586	5	Butane	1	0.089	14.497	5.584541107	7.422263866
P-4594	NB-A75585	6	Butane	1	0.089	14.498	5.5577489903	7.421850298
P-4008	NB-A58933	2	Butane	1.5	0.11	16.375	5.583562354	7.422343178
P-4004	NB-A58857	4	Propane	8	0.162	52.771	1.722932202	11.6412863
P-4455	NB-A75572	4a	Propane	16	0.175	64.442	0.325924489	10.08487774
P-4455	NB-A75572	3a	Propane	16	0.175	64.449	0.325368102	10.08426876
P-4455	NB-A75572	5a	Propane	16	0.175	64.449	0.32581336	10.08694291
P-4455	NB-A56257	6a	Butane	24	0.117	70.73	0.743775301	9.171274224
P-4455	NB-A75573	19a	Butane	24	0.117	70.739	0.744586054	9.171923662

Recommendation: Broad insulation system maintenance of moderate-/high-risk lines in lieu of further NDE, which will give 100% visual inspection opportunity.

Use existing asset systems to manage CUI CMLs.

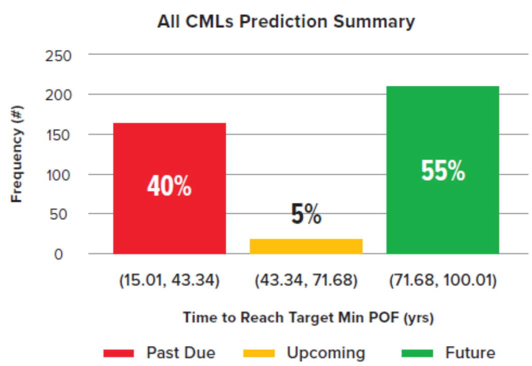

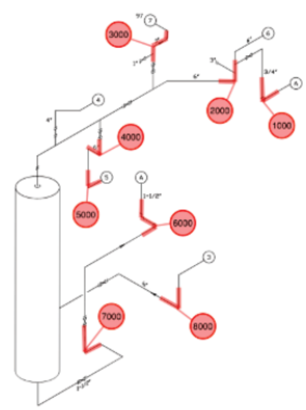

Figure 29 – One of the main deliverables of Equity's CUI Solution as a Service is a prioritized CML inspection plan like this, generated by the CUISight software, where the CMLs are ranked in order of priority, based upon whatever criteria is decided to be most important by the client. No matter what ranking is used, it always depends on the predicted failure time probability distribution in some way, which is why predicting failure time as accurately as possible is key to making the best possible decisions. Here, a simple ranking criterion is used, in terms of the time at which a minimum, specified target probability of failure (POF) is reached (similar to the risk target in standard API RP 581 RBI). The POF is, of course, just the CDF of the predicted failure time distribution, evaluated at this time. The mean damage rate and CUI initiation time are also reported.

8 NOMENCLATURE

C_F = Failure cost.

C_R = Replacement cost.

$f(t_F)$ = Failure time probability distribution function (PDF).

$F(t_F)$ = Failure time cumulative distribution function (CDF).

H = Total return (total benefit minus total costs over the entire lifecycle)

h = Lifetime-averaged return, defined by $h = H/t_R$

POF = Probability of failure at some time t, equivalent to $F(t)$.

t_F = Failure time.

t_R = Replacement time.

9 REFERENCES

[1] E^{2}G CorrSolutions, LLC. (2022). "Failure Prediction and Analysis Techniques." U.S. Patent 11,262,196, filed April 11, 2019, and issued March 1, 2022.

[2] Panzarella, C., Stenta, A., Osage, D. (2025). "*Asset Lifecycle Optimization Systems and Methods*." US Patent Application Number 19/033,030, filed January 21, 2025, pending.

EQUITY RESEARCH

Corporate Headquarters
20600 Chagrin Boulevard, Suite 1200
Shaker Heights, OH 44122

+1.216.283.9519
www.E2G.com