Proc. of the European Safety and Reliability Conference (ESREL 2026)
Edited by José C. Matos, Paulo B. Lourenço, Michael Beer, Daniel Oliveira,
Edoardo Patelli, Hélder S. Sousa, Rui A. Silva and Monica Santamaria

doi: 10.3850/ESREL2026061419_esrel26-p26354-cd

A Hybrid Fuzzy-Monte Carlo and DEA Framework for Risk-Based Inspection of Flexible Risers

Anderson B. Custódio
PETROBRAS Research Centre, Rio de Janeiro, Brazil. E-mail: andersonbc@petrobras.com.br

Alyson G. Vaillant, Lucas R. S. Ferreira
PETROBRAS Research Centre, Rio de Janeiro, Brazil.

Marcelo I. Lourenço, Eduardo Meirelles, Caroline Ferraz, Theodoro A. Netto, Henrique Honorato
Subsea Technology Laboratory, COPPE, Federal University of Rio de Janeiro (UFRJ)

Flexible risers are critical subsea assets subjected to numerous failure mechanisms. This work introduces IBR-Flex, a high-resolution Risk-Based Inspection (RBI) framework designed to shift integrity management from traditional calendar routines to a dynamic, risk-based process. The methodology employs heuristic rules to codify expertise into probabilistic models that adjust failure rates based on section attributes. At its core, a dual predictor-corrector algorithm integrates deterministic physical models to forecast evolving risks, updating them dynamically as new evidences, e.g., inspection data, are acquired. To prioritise interventions across four domains - safety, environment, economy, and reputation — Data Envelopment Analysis (DEA) establishes a Relative Risk (RR) efficiency frontier.

Keywords: Flexible pipes, risk-based inspection, RBI, integrity, data science, subsea engineering, fault-tree.

1. Introduction

PETROBRAS operates about 50% of the world's total length of flexible pipes ever manufactured (Loureiro et al., 2024). These pipes, which play a vital role in the company's subsea infrastructure, are subjected to dozens of failure mechanisms. Therefore, they require a high number of inspections to assess their integrity and reduce operational risks. While traditional calendar-based inspections have been a centerpiece on PETROBRAS' flexible pipes integrity management for decades, their ageing and increasing number require a more accurate and risk-based approach. To address this, the IBR-Flex framework (Ferraz et al., 2024) introduces a high-resolution analysis that shifts focus from broad integrity groups to individualised pipe sections.

At IBR-Flex core, a dual predictor-corrector algorithm integrates deterministic physical models to quantify time-dependent failure probabilities, which are dynamically updated whenever new evidence (e.g., inspection data) is acquired. By employing DEA, the methodology establishes a RR efficiency frontier (Honoratto, 2009; Netto et al., 2013; Souza et al., 2023), allowing management to prioritise interventions considering heterogeneous domains ensuring optimal resource allocation as flexible risers get old.

2. Methodology and Algorithm

From its conception, the IBR-Flex framework has been underpinned by the foundational principles illustrated in figure 1.

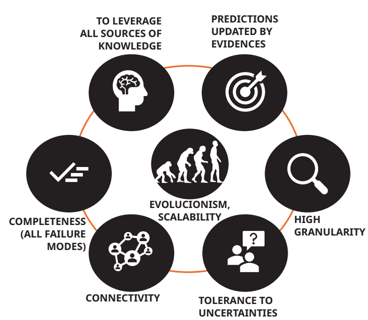

Fig. 1.: Targeting principles of IBR-Flex.

Those principals drive the architectural design of the kernel of the system and its modular ancillary components as presented in Figure 2.

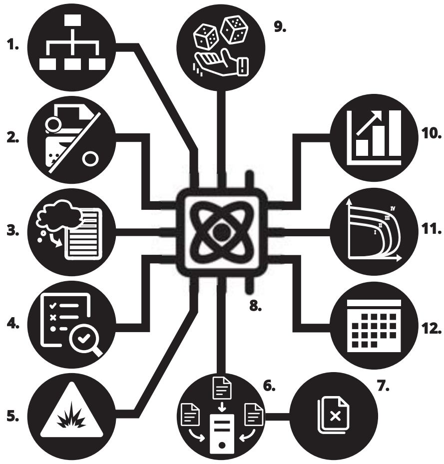

Fig. 2.: IBR-Flex components: 1-Event Trees; 2-Probability calculator; 3-SmartLang interpreter;4-Inspection technique guarantor; 5-Consequences calculator; 6-Easy Data Hub; 7-Missing Data Estimator; 8-The kernel; 9-What-if scenarios composer; 10-Dashboards; 11-Risk frontiers; and 12-Scheduler.

2.1. *Event Trees (ETs)*

ETs are derived from complex Fault Trees (FTs) and serve as the method's starting point, mapping degradation pathways through logical gates. Each ET represents a failure mode (FM), which comprises several failure mechanisms (FMECHs) competing for a certain FM at top. Meanwhile, each FMECH is composed of a large set of events and their relationships, where each event is associated with heuristic rules that process context's attributes, external effects, parent-child dependencies and physical laws to quantify risk levels. To clarify, a FMECH can be represented by a sequence of bottom-up, inspectable or purely predictive events, e.g.: outer sheath damage → annulus flooding → tensile armour corrosion → corrosion-fatigue of tensile wires → tensile wire rupture → section rupture.

At this development, 5 ETs were generated: FM1-hydrostatic collapse; FM2-tensile armour rupture; FM3-pipe burst; FM4-failure of sealing barrier; and FM5-end-fitting leakage.

2.2. *Probability calculator*

For each event, its estimated probability and the related rate over time are calculated through heuristic rules using programmable functions of several parameters (e.g.: water depth; mechanical properties; presence of outer sheath damages; dates; and so on). They operate in a simple sense: from a basic probability, aggressors or weaknesses increase it, and barriers or alleviators do the opposite, and principles of fuzzy logic tackle the inherent unknowns throughout the evaluation. A minimum definition required is:

$$P_{1y} = f(x_1, x_2, x_3, \ldots, x_n) \quad (1)$$

where P_{1y} is the 1-year probability; $f()$ are the heuristics of each single event; and x_n represents the parameters. Some temporal projection models need other parameters (e.g. Weibull shape parameter, etc).

Heuristics were selected as predictors for their ability to codify intuitive expert judgement and riser attributes into numerical rules compatible with fuzzy logic, and because they enable an evolutionary framework, as models are continuously improved through feedback from empirical data. The heuristics are built in an Integrated Development Environment (IDE) using hybrid datasets where virtual instances represent infrequent physical pathologies typically absent in population, as damaged flexible pipes are generally decommissioned.

This module employs large-scale Monte Carlo (MC) simulations during the rule-calibration phase to ensure that heuristic-based predictions are statistically robust across the riser dataset. While a single rule provides a point-estimate for a specific riser section, the MC engine samples from real and virtual pipes to generate a spread of probable outcomes. This process allows experts to evaluate, based on statistical data, the predictor's accuracy and sensitivity.

The IDE results' panel shows: (a) histograms; (b) non-exceedance statistics; (c) correlation plots to measure the influence of each parameter; (d) time series for special individuals (e.g.: holders of extreme probabilities); and (e) sensitivity simulations on the parameters of the probability distribution.

Within the IDE, testing and calibration rely on simplified time-projection functions - such as constant or Weibull models — to extrapolate future risk profiles. While these analytical laws provide a baseline for rule validation, the kernel implements a dynamic temporal mesh to integrate inspection records and environmental updates, enabling the corrector phase to refine probabilities over time.

2.3. *SmartLang interpreter*

Over four decades, PETROBRAS has specified, qualified, and operated flexible pipes, accumulating extensive expertise. Transposing this tacit knowledge into heuristic rules for computational risk prediction became a significant challenge.

SmartLang bridges the gap between natural language used by common people and the code that may be interpreted by computers. It is a scripting language based on a simple and comprehensible commands. Sentences such as ”if the pipe's external sheath has flaws and it was designed before a certain revision of guidelines, then increase probability of fatigue in 40%” can easily be written in SmartLang.

2.4. *Inspection technique guarantor (ITG)*

Originally, the ITG was designed to model inspection assertiveness, using heuristics to adjust probabilities based on technique reliability and riser-specific constraints. This logic is being expanded to incorporate broader concept of new information beyond inspection, ensuring the kernel's predictive mesh remains sensitive to all data that might alter degradation rates, such as the example of sealing barrier mentioned in section 3. Mainly, the ITG defines how to update the event's failure rate and cumulative probability when new data appear.

It can account for the individual's characteristics - e.g., if the pipe's external sheath is too thick to allow ultrasound to penetrate or if it is buried in the seabed.

2.5. *Consequences calculator*

The consequences of each failure mode are assessed across four domains: environmental, safety, economic, and reputation. The consequence application uses rules consistent with the PETROBRAS standards, penalising risks based on factors such as the volume of oil spilled, the function of the riser (e.g., gas injection vs. oil production), and the proximity to personnel. The environmental impacts are calculated by simulating the duration of a potential leak and the sensitivity of the surrounding offshore area. The economic consequences account for both direct repair costs and indirect losses from production downtime. This multi-domain assessment ensures that a section prone to environmental collapse is prioritised alongside a section facing imminent economic loss due to end-fitting leakage.

2.6. *Easy data hub (EDH)*

The EDH is essentially an aggregator and conditioner of data related to materials, design, interconnection, vessel, installation, operation, inspection, and possibly monitoring for each section. It connects with different data sources, tests data validity and completeness, fills gaps where appropriate, and makes conditioned and labelled information available to predictor and corrector algorithms, as well as to other applications. EDH creates replica of datasets under its own management and periodically revisits the original source in order to examine updates.

In early future, the EDH will connect to other expert systems to obtain real-time data from digital twins, fatigue predictors, and other sources, allowing heuristic rules to be fed by these data.

2.7. *Missing Data Estimator (MDE)*

MDE is a set of educated guesses that estimates the value for a single missing parameter whose definition is required by the assessment of probabilities, consequences, or inspection efficiency. It addresses epistemic uncertainty by utilising fuzzy logic principles to infer values for incomplete datasets. Rather than relying on simple averages, the MDE applies heuristic rules that account for linguistic and physical similarities between riser cohorts and environmental contexts. This fuzzy approach ensures that the estimation of missing attributes follows the principle of reasonable conservatism, preventing the underestimation of risk when specific parameters, such as the exact installation date or precise chemical composition of the annulus, are unavailable.

Nevertheless, the ability to estimate some of the missing attributes, there may be situations in which the probability of the failure mode (top event) is very sensitive to the parameter. In these cases, in addition to the suggestion by MDE, the alternative values of this attribute are candidates for the what-if cases to be built.

2.8. *The Kernel*

For every flexible pipe section the kernel performs the following functions:

(1) Temporal management: it establishes a dynamic time mesh with anchor points at relevant operational dates, such as installation, inspection, or environmental shifts.
(2) Probabilistic prediction: Using heuristics and predictive algorithms, it evaluates event-specific probability density functions (PDFs) and numerically integrates them into cumulative failure probability time series for each failure mode.
(3) Correction and feedback: It applies ITG-defined rules to update failure kinetics based on new evidence, while monitoring discrepancies between forecasted and empirical outcomes.
(4) Risk prioritisation: It runs DEA (see details below) to establish a relative risk (RR) ranking across the population.

To address the inherent incommensurability of multi-domain risks, IBR-Flex utilises DEA to establish a RR ranking. Direct comparisons among heterogeneous indicators — such as weighing a 10% probability of a personnel fatality against a 30% likelihood of a subsea oil release or an 80% chance of major production loss — are traditionally subjective and complex. By modelling each riser section as a Decision Making Unit (DMU), the framework establishes a non-parametric efficiency frontier that transforms these diverse consequences into a consistent metric.

2.8.1. *Data Envelopment Analysis (DEA)*

The DEA modelis a non-parametric optimisation technique used to establish a relative risk efficiency frontier by treating each flexible pipe section as a DMU. In this framework, DEA transforms heterogeneous risk indicators across safety, environment, economy, and reputation into a consistent relative ranking.

The algorithm calculates time-dependent weights for each failure mode ($\xi_{FM}(t)$) by integrating failure probabilities with domain-specific consequence multipliers. These weights $\xi_{FM}(t)$ populate the "risk matrix" 1, serving as inputs for the optimisation problem defined by:

$$RR_{DMU(z)} = Maximise(\sum_{k=1}^{z}[\xi_{FM}(t)]_{z,k} * [\varphi_{FM}]_{z,k}) \tag{2}$$

... where RR is the relative risk, $[\varphi_{FM}]_{z,k}$ represents the contribution of each FM weight to the final risk, and z is the total number of DMUs.

Table 1.: DEA Matrix

DMU (z)	FM1 $k=1$	FM2 $k=2$	...	FMk $k=k$
DMU1	$\xi_{FM}(t)_{1,1}$	$\xi_{FM}(t)_{1,2}$	...	$\xi_{FM}(t)_{1,k}$
DMU2	$\xi_{FM}(t)_{2,1}$	$\xi_{FM}(t)_{2,2}$	...	$\xi_{FM}(t)_{2,k}$
...	...	...	...	...
DMUz	$\xi_{FM}(t)_{z,1}$	$\xi_{FM}(t)_{z,2}$	...	$\xi_{FM}(t)_{z,k}$

Mathematically, the DEA algorithm aims to identify the efficiency frontier by maximising the RR for each DMU, thereby quantifying its proximity to the worst-case scenario.

This relative positioning enables integrity managers to prioritise inspections for sections exhibiting the highest risk, even when absolute failure probabilities remain low.

2.9. *What-if scenario composer (WISC)*

The WISC addresses epistemic uncertainty by injecting perturbations, which are prescribed values or assumptions, into relationships and predictions.

Its function enables sensitivity analysis at the expense of added computational effort and also incorporates fuzzy variability principles to address risks that cannot be purely quantified by historical frequency.

By now, few scenario-composition rules have indeed been codified.They try out inspection techniques with different probability of detections or trigger alternative scenarios by setting probabilities of inspectable events to 100% when forecasted risks exceed specific thresholds, allowing the framework to model augmented-probability conditions. These scenarios typically support conservative integrity management by prioritising worst-case outcomes during the decision-making process.

2.10. *Decision support suite*

Once the data are structured by the data management hub (EDH and MDE) and failure probabilities are processed by the kernel, the decision support suite executes a set of post-processing tasks. This suite integrates risk frontier definitions and scheduling guidelines with the augmented cognition provided by dashboards, transforming analytical outputs into actionable integrity management strategies.

The risk frontiers categorise the riser population into response zones by aligning DEA-derived RR with 20-year cumulative failure probability targets (table 2). A 100% RR reference is established by modelling a pipe with 5%-probability of failure and the greatest catastrophic consequence. This architecture segments risers into four zones, ranging from zone 1 (minimal risk, RR $<$ 46%) to zone 4 (imminent hazard, RR $\geq$ 86%), providing a consistent metric for prioritising interventions.

Table 2.: Risk frontiers.

Zone	Relative Risk	20-year cumulative probability
1	0 - 46%	less than 0.2%
2	46 - 66%	0.2 - 2%
3	66 - 86%	2 - 20%
4	$>$ 86%	$>$ 20%

These groupings facilitate automated decision rules: assets maintaining a position within Zone 1 do not require urgent inspection except for collecting life-extension data or addressing epistemic risks. Conversely, pipes projected to enter Zone 3 or 4 within the next year may be flagged for urgent scheduling. Depending on other variables, this framework allows operators to choose between mitigating risk growth through future interventions or decommission assets that approach their safety limits.

To handle the massive data generated by the predictor-corrector algorithm, the dashboards provide a visualisation layer designed for augmented cognition. The module can identify the factors most relevant to risk kinetics for each section. This interface provides access to the inventory of predicted results across all FMs and timeframes, allowing management to compare base-case trajectories against what-if scenarios modified by inspection evidence.

The scheduler optimises both the timing and the methodology of interventions. Since inspection techniques vary in assertiveness and susceptibility to section-specific attributes (e.g., burial or sheath thickness), the scheduler utilises information from base case and what-ifs to simulate how a specific technique might reduce the governing hazard. For example, it can compare the risk reduction provided by a MAPS tool inspection versus torsion monitoring, ensuring that inspection resources are allocated to methods with the highest impact on overall safety.

Ultimately, this suite ensures that the IBR-Flex framework remains scalable and aligned with international risk tolerability standards.

3. The predictor-corrector algorithm

The IBR-Flex framework is powered by a dual predictor-corrector algorithm designed to manage the time-dependent evolution of failure risks across the population of pipes.

Predictor: Heuristic rules codify expertise and physical laws into an executable SmartLang code to estimate the probability of 1-year and its expected variation rate. Suppose that the event to be analised is the PA-11 hydrolysis. The predictor integrates deterministic chemical kinetics with probabilistic reliability models. By deriving the characteristic life and choosing a shape parameter of a Weibull distribution from Arrhenius estimates, the algorithm replaces traditional binary "pass/fail" criteria with a continuous probability density function (PDF), adjusted by section-specific attributes like thermal insulation, temperature, pH, etc.

Corrector: While the probability calculator handles static parameter sets, the kernel manages dynamic updates through a temporal mesh. This mesh utilise "anchor points" (e.g., installation, inspection, or environmental shift dates) to anchor the cumulative risk curve. When new evidence appears, the ITG models how to update the cumulative risk. This prevents the "as-good-as-new" fallacy: imperfect inspection data update the failure kinetics without erroneously resetting the probability to zero, unless the reliability of the technique justifies it.

Beyond inspections, the algorithm treats operational shifts (e.g. changes in pH or fluid content) as triggers to recalculate failure kinetics. Significant mismatches between forecasted and empirical outcomes are recorded in a divergence table, enabling experts to refine heuristics. This feedback loop ensures model evolution and scalability.

Between two dates within which nothing else occurs, the probability numerically integrated by the kernel grows according to the previous rate predicted by the probability calculator (predictor). The effects of inspection records may be ruled by prescribed correction or Bayesian update, depending on the rules defined as per section 2.4. Figure 3 illustrates the influence of inspection efficiency for simple cases of prescribed correction of the cumulative probability.

The graph compares three cases: (a) if no inspection is done, the predicted 2-year probability rises to 81%; (b) if the inspection is done using technique A which has a lower efficiency, the probability instantaneously drops to 34% at year 10 (date of inspection), and it continues to grow afterwards; (c) if the inspection uses technique B which asserts that the event has not occurred by that time, the probability will drop to zero

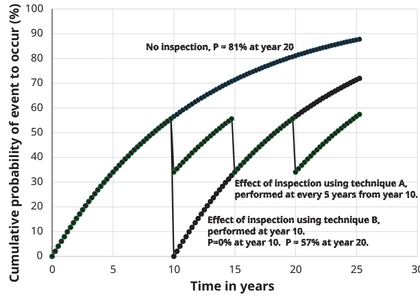

Fig. 3.: Possible effects of inspection technique in the probability evaluated by kernel.

at year 10. In order to achieve confidence levels closer to those of technique B, technique A may be repeated every 5 years.

Moving back to example of hydrolysis of the sealing barrier, no simple inspection (but a section dissection) is expected to be able to significantly update the predicted probabilities and their rate of change. However, changes in fluid contents or in its temperature are new information that triggers the update. In this case, the cumulative probability - resulting from previous aggression history - is not altered by the kernel, but its rate of change will be, based on updated predictions. Nevertheless, whether a section of the flexible pipe is removed, the sealing barrier is dissected and the measured CIV is found below the prediction, the cumulative probability has to be corrected accordingly. These decisions concerning the algorithm's response to new data must be taken by the ITG.

Trials using different inspection techniques are special cases addressed by the WISC.

4. Case study

The IBR-Flex framework was applied to 22 flexible risers across 2 semi-submersible and 2 FPSO units, targeting rupture (FM2), burst (FM3), and leak (FM4) failure modes. The study focused on the high-dynamic top sections, where fatigue and wear are most prevalent. By integrating multi-domain consequences through DEA, the framework generates a RR metric that aligns heterogeneous failure modes onto a single efficiency frontier, correcting qualitative misalignments inherent in calendar-based planning. Four scenarios were analysed to support decision-making: current and 10-year forecasted conditions (std-RR, fut-RR), along with their corresponding what-if variants (whatif-std-RR, whatif-fut-RR) updated by supposedly confirmed inspection evidence.

5. Results and discussion

Looking at the results of present and future operational conditions (std-RR and fut-RR), see Figure 4, which presents the results for environmental and safety domain. In all domains, most of the risers remained below critical risk thresholds, indicating a low urgency for additional inspection, except for considerations of epistemic or methodological uncertainties.

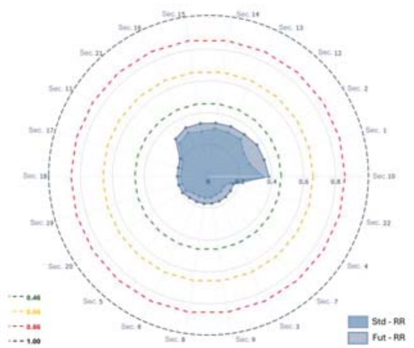

(a) Environmental

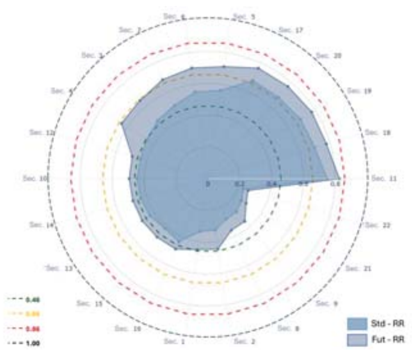

(b) Safety

Fig. 4.: Present and future RR scenarios of the flexible pipe sections / consequence domain

As mentioned in item 2.9, the what-if scenario assume that one inspected event inside a failure mechanism simply happened, i.e., its value jump to 100%, and evaluate its impacts at the top event of the fault tree, or in the failure mode. For those what-if simulations, present (whatif-std-RR) and future operational ones (whatif-fut-RR), as can be seen at Figure 5, which presents the results for environmental and safety domain, again the majority of flexible risers evaluated got positioned below green risk frontier, however, for the safety domain five flexibles crossed the red frontier one of which reached the maximum risk level, 100%.

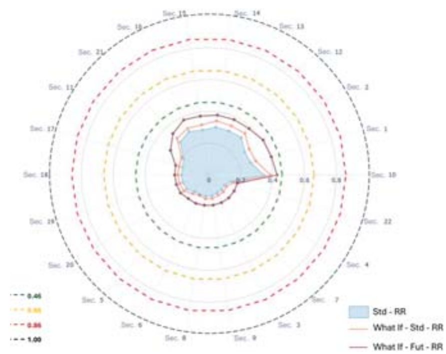

(a) Environmental

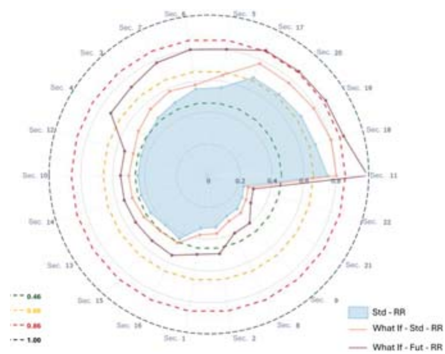

(b) Safety

Fig. 5.: What-if present and future RR scenarios of the flexible pipe sections / consequence domain

Flexible pipe sections positioned on the black risk frontier represent the highest risk intensity within the current population of risers and are flagged for immediate action. However, caution should be exercised in that decision, since it evolves future prediction and sensitivity analysis, therefore it is case by case. For the flexible riser sections mentioned, the present operational condition (std-RR) remains below the red risk frontier, thus, it was needed to apply what was called Scheduler, which is a set of rules built to deal with all risk possibilities. For that flexible riser section, it recommends to shorten the timeframe for the next inspection.

Analysis of current and projected sensitivity result maps shows that IBR-Flex successfully aligns with the risk tolerability matrices of international standards. The results confirm that a dynamic, risk-based approach captures the evolving nature of flexible riser degradation more accurately than fixed interval schedules. Decision taking must account for the fuzzy, scattered nature of the results, which encompasses basic and what-if scenarios.

A minimum of 14% inspection reduction costs was obtained, which could potentially reaches 49%. Therefore, adopting a risk-based approach optimises inspection frequency and techniques, reducing costs and improving earlier detection of critical failures.

6. Conclusions

The IBR-Flex methodology is changing the way PETROBRAS manages its flexible risers, moving from rigid, calendar-based routines to a dynamic, risk-informed process. By codifying expert knowledge into SmartLang heuristics and utilising a predictor-corrector algorithm, the framework addresses the complexities of hundreds of FMECHs.

Tolerance to uncertainty is ensured by fuzzy principles applied in the heuristics and by what-if scenarios generated through WISC. The analysis outcomes may be dispersed across different regions of the risk spectrum, and the decision-support suite incorporates this dispersion to provide more robust recommendations.

The use of DEA for prioritisation enables a more rational and economical allocation of inspection resources, focusing attention on sections with the highest relative risk. Ultimately, this framework enhances environmental safety and operational sustainability, providing a scalable solution for the global offshore industry as assets reach their original design service life.

References

Ferraz, C., T. A. Netto, M. I. L. Souza, H. J. Honoratto, A. G. Vaillant, A. B. Custodio, L. R. S. Ferreira, E. V. Meirelles, and J. V. Moita (2024). Probabilistic data envelopment analysis applied to risk-based inspection of flexible pipes. *Ocean, Offshore and Arctic Engineering* (OMAE2024-133878).

Honoratto, H. J. (2009). Risk ranking for submarine pipeline through flexible application model of data envelopment analysis. M.sc. dissertation, UFRJ.

Loureiro, W. C., A. G. Vaillant, B. N. Borges, C. D. de Lemos, G. C. Campello, M. M. F. Capella, and T. R. Schone (2024). Flexible pipe integrity management strategies applied to lifetime extensionand revitalization of deepwater brownfields. *Offshore Technology Conference* (OTC-35183-MS).

Netto, T. A., H. J. Honorato, and R. Y. Qassim (2013). Prioritization of failure risk in subsea flexible pipes via data envelopment analysis. *Marine Structures 34*, 105–116.

Souza, M. I. L., J. H. Honorato, T. A. Netto, , and C. Ferraz (2023). Dea model applied to risk based inspection for offshore flexible pipelines with virtual relative risk frontier from event trees. In *Materials Technology; Pipelines, Risers, and Subsea Systems*, Volume 3 of *International Conference on Offshore Mechanics and Arctic Engineering*, pp. V003T04A019. ASME.