International Journal of Business Quantitative Economics and Applied Management Research
Volume-7, Issue-9, 2023 ISSN No: 2349-5677

EXPLAINABLE PORTFOLIO INTELLIGENCE FOR MULTI-PROJECT ENGINEERING INVESTMENT SELECTION UNDER DYNAMIC UNCERTAINTY

Yogesh Gawade
Industrial & Systems Engineering
Oakland University, Rochester, Michigan, USA.
Yogeshgawade@oakland.edu

Abstract

Engineering companies regularly invest limited capital in a portfolio of related projects that typically have uncertain costs, durations and returns as conditions are subject to change, and the tools that are best able to manage these uncertainties are often the least visible. A systematic review is provided and aggregated 24 peer-reviewed papers on project portfolio selection, data-driven project portfolio optimization, uncertainty quantification and explainable AI. The corpus is split into three thematic streams and the main model families are contrasted with regards to uncertainty treatment, dynamic capability, computational cost and interpretability. The synthesis demonstrates that non-dominated sorting can be done very efficiently, and that the out-of-sample R-squared in a 60-year equity study can be improved from 0.16 percent to 0.40 percent by neural predictors, the area under the curve of credit models improved from 0.81 to 0.93 with gradient-boosted credit models and Shapley explanations. An example of a 10 project portfolio over 495 portfolio options identifies a risk adjusted selection of projects that involves a net present value loss of 1.4 percent, a diversity gain of 13.1 percent, and a transparent distribution of the diversity gain across the selected projects using Shapley attribution. In this work, an integrated five-layer architecture is proposed, which combines aleatoric uncertainty separation, epistemic uncertainty learned valuation, multiobjective and two-stage optimization, and layered explanation under governance review. Results suggest that explainability should not be an add-on to the end of optimization, but rather be included as a layer of the decision-making process, and that models that come first should be interpretable when there are high demands for accountability.

Index Terms—Project portfolio selection, explainable artificial intelligence, multiobjective optimization, uncertainty quantification, Shapley additive explanations, deep reinforcement learning, engineering investment, dynamic decision-making, systematic review.

I. INTRODUCTION

Typically, engineering and construction companies are rarely faced with just one investment decision, but rather a series of potential projects compete for limited capital, labor, and management resources. Project Portfolio Selection (PPS): This process involves selecting a subset of candidates to achieve organizational objectives, while keeping stakeholders' resource constraints in mind. The explicit use of risk as a selection criterion has been found to influence the composition of the preferred portfolio for contractor firms as projects, which would be

selected for profit, may be rejected due to their risk level when evaluated as part of the portfolios [1]. This is made harder by the dynamic uncertainty as costs are changed, demands are adjusted, and technical feasibility is impacted as projects develop and external factors vary. Data-driven estimators are becoming more commonplace as tools for prediction of these quantities today, but no information is contained in a point forecast about its own reliability. To fill this void, uncertainty quantification (UQ) is a field of study that aims to add calibrated uncertainty to predictions, and Bayesian approximation as well as ensemble learning are two main families of UQ methods for deep neural networks with their own specific cost-benefit analysis [2]. The second challenge has to do with trust. Recommendations, even those that are correct, are not likely to be accepted by the engineering board, nor will they pass regulatory review and audit. Explainable artificial intelligence (XAI) is the research field focused on making the behaviour of learned models intelligible to human stakeholders, and an initial systematic survey of the field, based on 381 papers, showed that there is a high level of research activity in this area with no consensus on what explanation is [3]. The current review review is therefore focused on how explainability and uncertainty and optimization can be synergized in a coherent portfolio intelligence capability.

II. REVIEW PROTOCOL AND ANALYTICAL FRAMEWORK

The review took the form of a structured synthesis of 24 peer-reviewed sources. Works were limited to those that contained a model, algorithm, empirical benchmark, or conceptual taxonomy that involved the selection and/or allocation of capital between competing alternatives in an environment of uncertainty. Financial portfolio studies were kept when the mechanism carries over to engineering investment – e.g., the mean–variance and conditional value-at-risk portfolios were found to be not robust to the estimation error that occurs when the moments of the distributions of the performance metrics are learned from limited data, a problem also occurring when the distributions of project return are estimated from small historical samples [4]. Thematic streams represent the thematic theme of each source and comparison attributes represent the attributes used as common denominators across sources.Thematic streams summarize each source's thematic theme and comparison attributes summarize attributes used as common denominators across sources. Explainability attributes were defined with a taxonomy, distinguishing between transparent models, which are intelligible by design, due to the ability of simulation, decomposability or algorithmic transparency, and post-hoc techniques, which explain an opaque model after training, through feature relevance, local approximation, visualization or examples; the taxonomy also identified explainability as part of a wider agenda on responsible artificial intelligence, including fairness, accountability, and privacy [5].

International Journal of Business Quantitative Economics and Applied Management Research
Volume-7, Issue-9, 2023 ISSN No: 2349-5677

TABLE I. MAPPING OF REVIEWED LITERATURE TO THEMATIC STREAMS AND FOCAL CONCEPTS

Thematic stream	Focal concept	References
Portfolio selection and optimization	Risk-aware and cancellation-aware selection	[1], [7]
	Multiobjective and metaheuristic search	[8], [9]
	Stochastic, fuzzy and dynamic programming	[10], [17], [23]
	Reviews of portfolio selection research	[18], [19]
Learning-based valuation and allocation	Data-driven portfolio estimation and return prediction	[4], [11]
	Deep reinforcement learning allocation	[14], [22], [24]
Uncertainty and explainability	Bayesian, ensemble, aleatoric and epistemic uncertainty	[2], [13]
	Explainability taxonomies and surveys	[3], [5], [12]
	Feature and outcome attribution	[15], [20]
	Human-centered and interpretable-first design	[16], [21]
	Applied explainability in finance	[6]

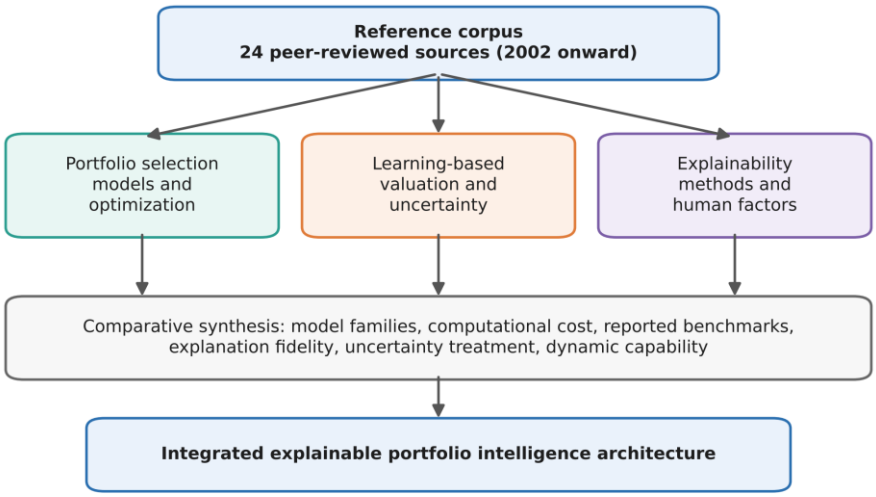


Fig. 1. Review Framework Linking the Reference Corpus, Three Thematic Streams, Comparative Synthesis and the Integrated Architecture

Where engineering-specific studies were not available, cross-domain evidence was weighted. The most similar comparison is credit risk management, where lending, just like capital budgeting, is a sunk cost and subject to regulatory authorization. An extreme gradient boosting (XGBoost) classifier was built on a dataset of 15,045 small and medium enterprises with an area under the receiver operating characteristic curve (AUC) of 0.93 against the logistic regression

International Journal of Business Quantitative Economics and Applied Management Research
Volume-7, Issue-9, 2023 ISSN No: 2349-5677

with an AUC of 0.81, and the Shapley values were correlated to group borrowers with similar explanatory profiles for auditing purposes [6].

III. ENGINEERING PORTFOLIO SELECTION MODELS

A. Static and Stochastic Formulations

The canonical PPS formulation is a binary knapsack problem that selects projects to maximize the total value of projects within a given budget B (see Eq. 1), with vi and ci representing the value and cost of project i, respectively, and xi an indicator variable that takes the value of 1 if project i is selected and 0 otherwise. The dynamic twist on public research and development (R&D) portfolios is that projects can be dropped at any point up to the completion of the project, allowing the budget to be freed up but the value to be lost. Direct modeling of cancellation has been shown to result in differing funding decisions from deterministic knapsack models due to the shifting value of the released budgets from cancelled projects [7].

$$\max_x \sum_{i=1}^{n} v_i x_i \,, \text{ s.t. } \sum_{i=1}^{n} c_i x_i \le B, \; x_i \in \{0,1\} \tag{1}$$

B. Multiobjective and Risk-Aware Search

A typical engineering portfolio isn't typically trying to accomplish one goal; expected return and risk are balanced together alongside strategic alignment and resource balance. The non-dominated sorting genetic algorithm (NSGA-II) is still the benchmark for such problems, which are called elitist multiobjective optimization problems. It sorts non-dominating solutions quickly, with the per generation sorting cost reduced from O(MN3) to O(MN2) for M objectives and a population size of N, the elitism keeps the best solutions at each generation, and the crowding distance (Eq. 2) keeps diversity on the frontiers without requiring the user to tune a sharing parameter as in [8]. The resulting cost difference is summarized in Table 2 and Fig. 2, and is more than 3 orders of magnitude in favor of N = 1600.

$$d_i = \sum_{m=1}^{M} \frac{f_m(i+1) - f_m(i-1)}{f_m^{\max} - f_m^{\min}} \tag{2}$$

TABLE II. DOMINANCE COMPARISONS PER GENERATION FOR EARLIER AND FAST NON-DOMINATED SORTING

Population size (N)	Objectives (M)	O(MN³) comparisons	O(MN²) comparisons	Reduction factor
100	2	2.00×10^6	2.00×10^4	100
100	3	3.00×10^6	3.00×10^4	100
400	3	1.92×10^8	4.80×10^5	400
800	3	1.54×10^9	1.92×10^6	800
1600	3	1.23×10^{10}	7.68×10^6	1600

a. Counts computed from the asymptotic complexity expressions; constant factors are omitted.

International Journal of Business Quantitative Economics and Applied Management Research
Volume-7, Issue-9, 2023 ISSN No: 2349-5677

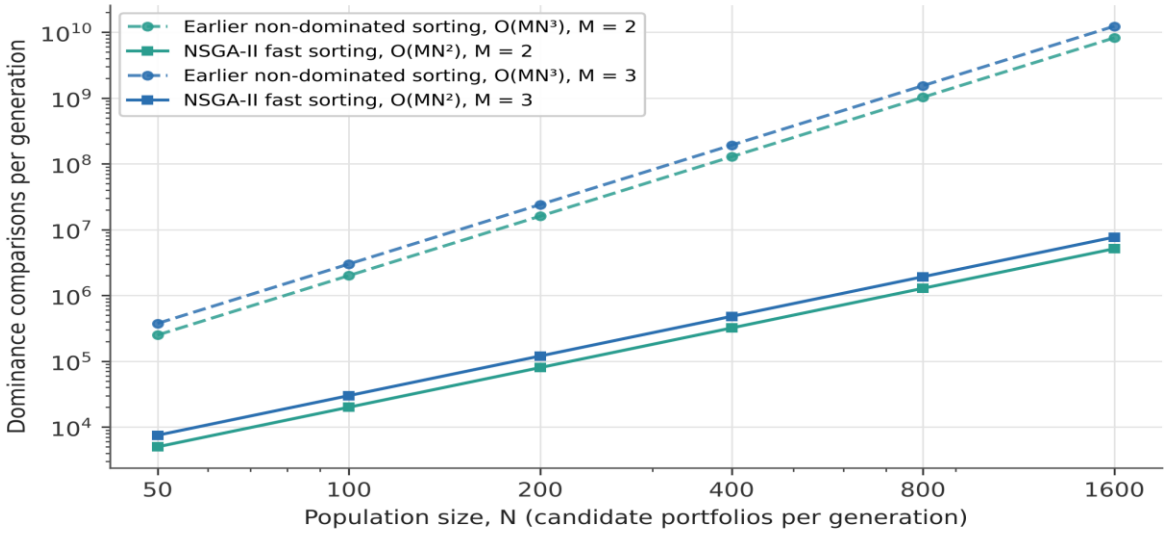


Fig. 2. Growth of Dominance Comparisons with Population Size for Two and Three Objectives on Logarithmic Axes

The most direct way that risk is introduced is in the form of modern portfolio theory (MPT), where the variance of the portfolio value is calculated from the covariances among the projects, as listed in Eq. 3. The resulting problem is a nonlinear integer program in which metaheuristics are applied, and harmony search, which mimics the process of improvisation among musicians, has been incorporated with MPT in order to select project portfolios that optimize the return and variance of the projects [9]. To show this trade-off, a ten-project engineering example was built up for this review (Table 3) with a 0.3 uniform pairwise correlation and with a budget of 100 million USD. An exhaustive search of all possible budget feasible portfolios resulted in 495 portfolios, of which 47 were Pareto-efficient (Fig. 3). The highest expected net present value (NPV) portfolio was 36.2 million USD with 12.87 million USD standard deviation, while the risk adjusted optimum (Eq. 4) with $\lambda = 0.5$ had a value of 35.7 million USD and standard deviation of 11.18 million USD, which led to a 1.4 percent loss of expected value and a 13.1 percent reduction of dispersion.

$$\sigma_P{}^2 = \sum_{i=1}^{n}\sum_{j=1}^{n} x_i x_j \rho_{ij} \sigma_i \sigma_j \quad (3)$$

$$\max_x [E(V_P) - \lambda \sigma_P], \ \lambda \geq 0 \quad (4)$$

International Journal of Business Quantitative Economics and Applied Management Research
Volume-7, Issue-9, 2023 ISSN No: 2349-5677

TABLE III. ILLUSTRATIVE CANDIDATE PROJECTS USED FOR THE WORKED PORTFOLIO EXAMPLE

Project	Cost (million USD)	Expected NPV (million USD)	NPV standard deviation (million USD)	Selected at risk-adjusted optimum
Substation upgrade	18	6.2	2.1	Yes
Bridge retrofit	25	7.5	3.4	No
Water treatment expansion	22	6.8	2.0	No
Rail signalling	15	5.1	1.6	Yes
Solar farm	30	11.4	6.2	Yes
Data center cooling	12	4.3	1.5	Yes
Port automation	28	9.9	5.4	No
Hospital wing	20	5.6	1.8	No
Road resurfacing	10	2.9	0.8	No
Wind repowering	24	8.7	4.6	Yes

a. Values constructed for illustration; budget of 100 million USD, pairwise correlation of 0.3 and risk-aversion coefficient of 0.5.

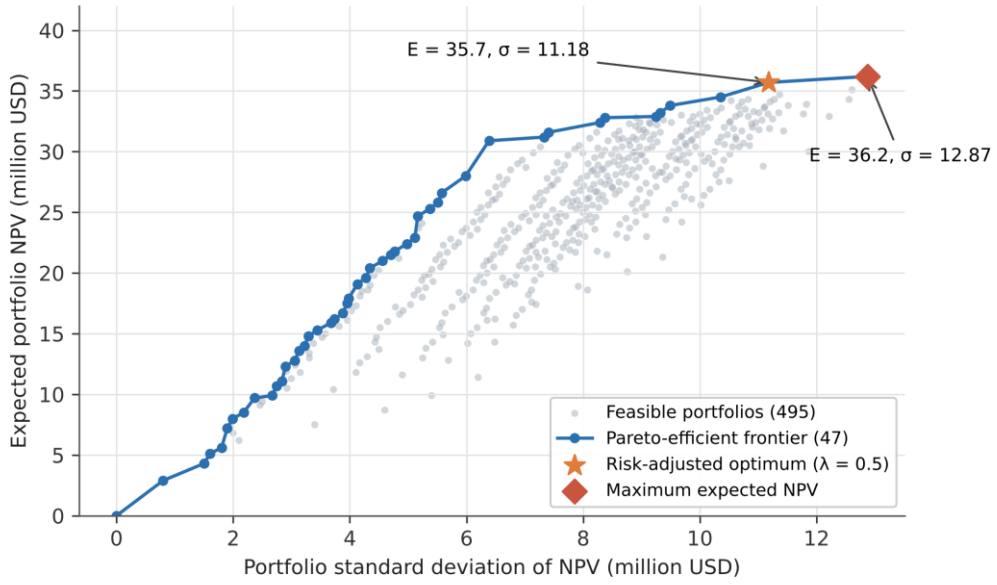


Fig. 3. Risk–Return Space of 495 Budget-Feasible Portfolios Showing the 47-Point Pareto Frontier and Two Reference Optima

C. Resource-Coupled Dynamics

In addition to this, there are dynamics that arise from resource interdependence. If several projects share the same pool of multi-skilled staff and if the stochastic durations of work packages are drawn from a common distribution, selection, scheduling and staffing become inextricably intertwined, and the feasible project portfolio in terms of the expected values for the stochastic durations can become infeasible due to the realized durations falling off the common distribution. This stochasticity has been explicitly represented by planning approaches that were reported to get better results than planning based on expected values alone [10]. Table 4 shows a comparison of the main model families with the attributes coded in Section II and illustrates that none of the engineering-oriented families offer a natural explanation for its recommendations other than the transparency of its own formulation.

TABLE IV. COMPARISON OF PORTFOLIO MODEL FAMILIES AGAINST REVIEW ATTRIBUTES

Model family	Uncertainty treatment	Dynamic capability	Native explainability	Scalability
Deterministic binary knapsack	None; point estimates	Static	High; linear objective	High
Mean–variance integer program	Covariance of project values	Static	Moderate	Moderate; nonlinear
Multiobjective evolutionary search	Objective trade-offs; scenario averages	Re-run each period	Moderate; frontier plots	$O(MN^2)$ per generation
Fuzzy group multicriteria ranking	Linguistic vagueness	Static	High; visible weights	High
Stochastic and two-stage programming	Probabilistic scenarios, cancellations	Multi-period recourse	Low to moderate	Limited by scenario count
Deep reinforcement learning	Learned from trajectories	Sequential and adaptive	Low; post-hoc required	High after training

IV. LEARNING-BASED VALUATION AND DYNAMIC UNCERTAINTY

A. Learned Valuation

The parameters used in the valuation models of any portfolio are themselves forecast, and the loss of accuracy in forecasting is a significant problem in the closest area of data-richness, namely asset pricing, where machine learning has significantly enhanced forecasting accuracy. Nearly 30,000 equities were compared over 60 years and more than 900 candidate predictors showed that tree ensembles and neural networks are superior to linear benchmarks, with the value-weighted long–short neural portfolio having an annualized Sharpe ratio of 1.35 [11].

These gains are achieved using models that are opaque to their internal accounts and workings, which may lead to a conflict of engineering governance. A popular survey of methods for explaining black-box models identifies four problem areas: modelling explanation, explanation of a particular instance of the model, perturbations of the model, and the design of a model

from the beginning that is easy to explain [12]. Outcome explanation relates to the justifications behind why a certain project was chosen in the context of portfolio intelligence, while model explanation relates to characterizing the basis of the valuation logic used in the entire candidate pool.

B. Uncertainty Decomposition and Sequential Allocation

Project portfolios do not exhibit the same dynamic uncertainty. Aleatoric uncertainty is uncertainty due to randomness in the system, e.g. weather delays, commodity price fluctuations, which cannot be alleviated with the collection of additional data, while epistemic uncertainty is due to lack of knowledge, e.g. an untested technology, and can be reduced through the accumulation of additional data. The total predictive entropy TPE of a probabilistic predictor with learned parameters θ from some set of data D can be decomposed additively into an expected aleatoric term and an epistemic mutual-information term, as demonstrated in Eq. 5 [13]. This separation is decision-relevant, as it indicates where staged investment or further feasibility study is useful, in view of epistemic uncertainty.

$$H[p(y|x,D)] = E_{p(\theta|D)}\big[H[p(y|x,\theta)]\big] + I(y;\theta|x,D) \tag{5}$$

If decisions about allocations are repeated, sequential decision-making under partial observability can be expressed and solved using deep reinforcement learning (DRL). A partially observable Markov decision process (POMDP) formulation of equity trading was used to teach an agent a twin-delayed deep deterministic policy gradient (TD3) with an annualized Sharpe ratio of 2.68 using the unseen test data [14]. This review uses the Sharpe ratio of Eq. 6 as the standard risk-adjusted performance metric and the benchmarks shown are summarized in Table 5 and presented graphically in Fig. 4. These findings are based on data-rich financial markets; engineering portfolios lack many of the observations and the transferred performance should be considered a maximum.

$$SR = \frac{E(R_P) - r_f}{\sigma_P} \tag{6}$$

TABLE V. REPORTED BENCHMARKS OF LEARNING-BASED MODELS RELEVANT TO PORTFOLIO VALUATION

Application domain	Comparison	Metric	Baseline	Learned model	Relative change
Equity return prediction, about 30,000 stocks over 60 years	OLS-3 versus NN3	Monthly out-of-sample R^2 (%)	0.16	0.40	+150%
Equity return prediction	OLS-3 versus random forest / GBRT	Monthly out-of-sample R^2 (%)	0.16	0.33 / 0.34	+106% / +113%

International Journal of Business Quantitative Economics and Applied Management Research
Volume-7, Issue-9, 2023 ISSN No: 2349-5677

Application domain	Comparison	Metric	Baseline	Learned model	Relative change
Credit default, 15,045 enterprises	Logistic versus XGBoost	AUC	0.81	0.93	+14.8%
Equity long–short allocation	NN4 value-weighted portfolio	Annualized Sharpe ratio	–	1.35	–
Equity trading, partially observable	TD3 agent on unseen data	Annualized Sharpe ratio	–	2.68	–
Futures trading, 50 contracts, 2011–2019	DRL versus time-series momentum	Risk-adjusted return	Momentum	DRL	Outperformed

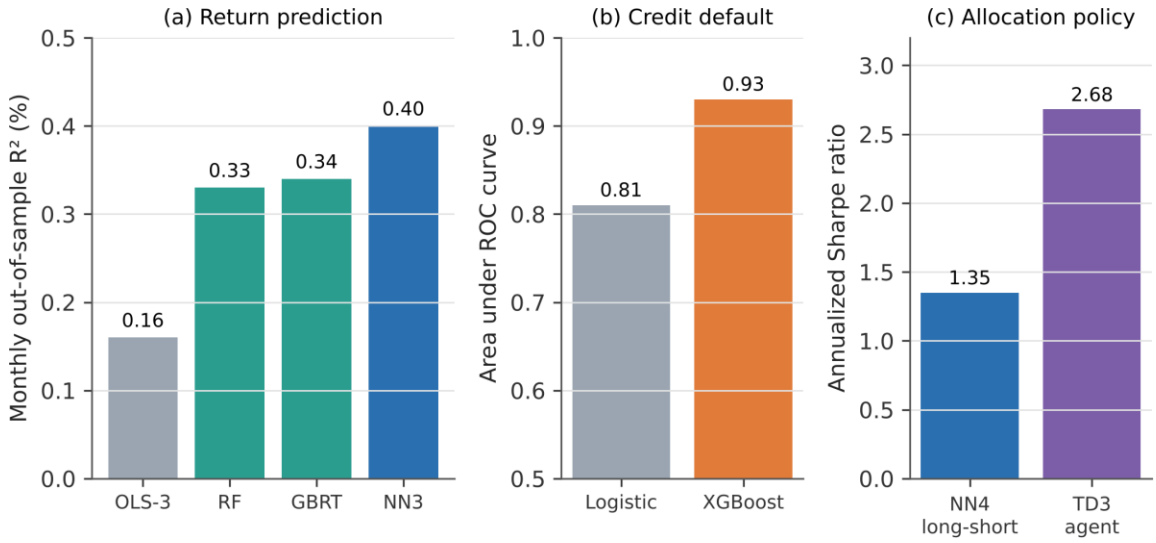


Fig. 4. Reported Benchmarks for Return Prediction, Credit Default Classification and Allocation Policies

V. EXPLAINABILITY MECHANISMS FOR PORTFOLIO DECISIONS

Shapley additive explanations (SHAP) stand out among the post-hoc techniques by having an axiomatic foundation: the attribution φi in Eq. 7 is the only allocation that satisfies local accuracy, missingness and consistency. Exact computation takes exponential time in number of players, but the TreeExplainer algorithm runs in O(TLD2) time for T trees with depth D and L leaves, and synthesizes local explanations to global summaries of interaction effects [15]. In a portfolio context, projects can be viewed as players and risk adjusted portfolio value as the characteristic function. Applied to the five-project optimum of Section III, the 9.00 million USD cost of the solar farm is distributed to the projects, and every project receives more than its stand-alone value due to the diversification gain of 2.41 million USD shared among them according to their marginal contributions (Fig. 5).

International Journal of Business Quantitative Economics and Applied Management Research
Volume-7, Issue-9, 2023 ISSN No: 2349-5677

$$\varphi_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!(|F|-|S|-1)!}{|F|!} [v(S \cup \{i\}) - v(S)] \quad (7)$$

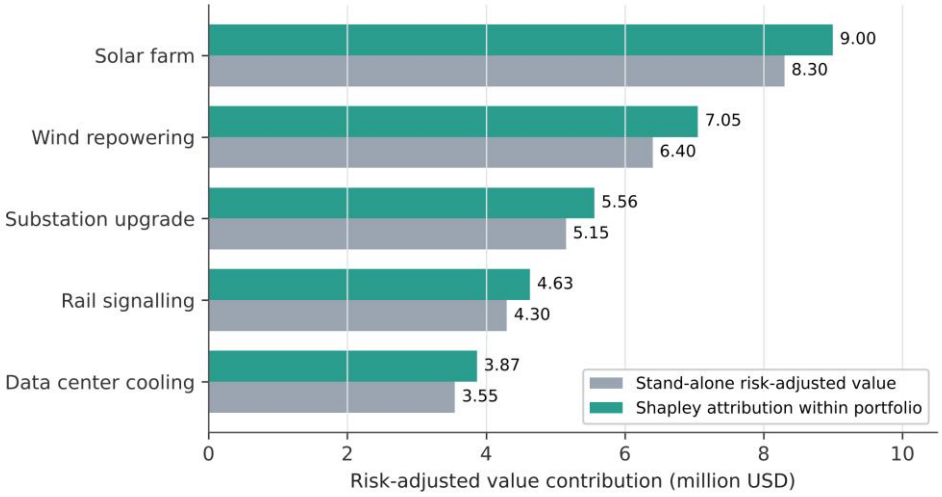


Fig. 5. Stand-Alone Risk-Adjusted Value Versus Exact Shapley Attribution for the Five Projects in the Risk-Adjusted Optimum

An attribution is required but not sufficient to provide an explanation. Philosophy and psychology and cognitive science research suggests that human explanations are contrastive (in response to a why question they address why one outcome is the case, rather than another); limited (they contain only a few causes); causal (they rely on causal relations, not simply statistical ones); and social (they are designed around the beliefs of the listener) [16]. The key question for a portfolio board is usually a contrastive one, like why invest in rail signalling instead of the bridge retrofit? An explanation layer should thus be able to provide contrasts between neighbouring frontier portfolios but it should not simply be an attributions report. A complementary approach are transparent multicriteria methods. In [17] a group decision process for selecting the sustainable infrastructure projects was developed which used interval type-2 fuzzy sets as models of the vagueness of the expert linguistic judgments and a new ranking method called MORAS that allowed multiple decision makers to rank infrastructure alternatives according to sustainability criteria. These are always interpretable since the weights of the criteria and the steps in the aggregation process are visible, but they do not learn from data, and have to be recalibrated manually as conditions change.An attribution is required but not sufficient to provide an explanation. Philosophy and psychology and cognitive science research suggests that human explanations are contrastive (in response to a why question they address why one outcome is the case, rather than another); limited (they contain only a few causes); causal (they rely on causal relations, not simply statistical ones); and social (they are designed around the beliefs of the listener) [16]. The key question for a portfolio board is usually a contrastive one, like why invest in rail signalling instead of the bridge retrofit? An

International Journal of Business Quantitative Economics and Applied Management Research
Volume-7, Issue-9, 2023 ISSN No: 2349-5677

explanation layer should thus be able to provide contrasts between neighbouring frontier portfolios but it should not simply be an attributions report.

A complementary approach are transparent multicriteria methods. In [17] a group decision process for selecting the sustainable infrastructure projects was developed which used interval type-2 fuzzy sets as models of the vagueness of the expert linguistic judgments and a new ranking method called MORAS that allowed multiple decision makers to rank infrastructure alternatives according to sustainability criteria. These are always interpretable since the weights of the criteria and the steps in the aggregation process are visible, but they do not learn from data, and have to be recalibrated manually as conditions change.

VI. COMPARATIVE SYNTHESIS AND DISCUSSION

A literature review of the PPS finds that uncertainty is the predominant modelling concern. Based on the mathematical structure, the approach to uncertainty, and the solution method and case study, fuzzy, stochastic and robust formulations were revealed as the main uncertainty treatments, with exact and metaheuristic solution techniques both represented [18]. No classification dimensions appeared in the literature with the term explainability. This means that the term interpretability has been used as an implicit by-product of transparent formulations and not as a deliberate capability.

A complementary review over the 20 years 1999-2019 reflected a continuation of the growth of PPS publications and a gradual evolution towards the inclusion of uncertainty and multiple objectives [19]. To compare the reviewed sources, they are plotted in a common timeline in Fig. 6, which reveals that explainability methods and learning-based allocation have grown after 2016, whereas engineering PPS studies focused on fuzzy, stochastic and metaheuristic formulations. Their integration is thus a relatively new and ongoing process.

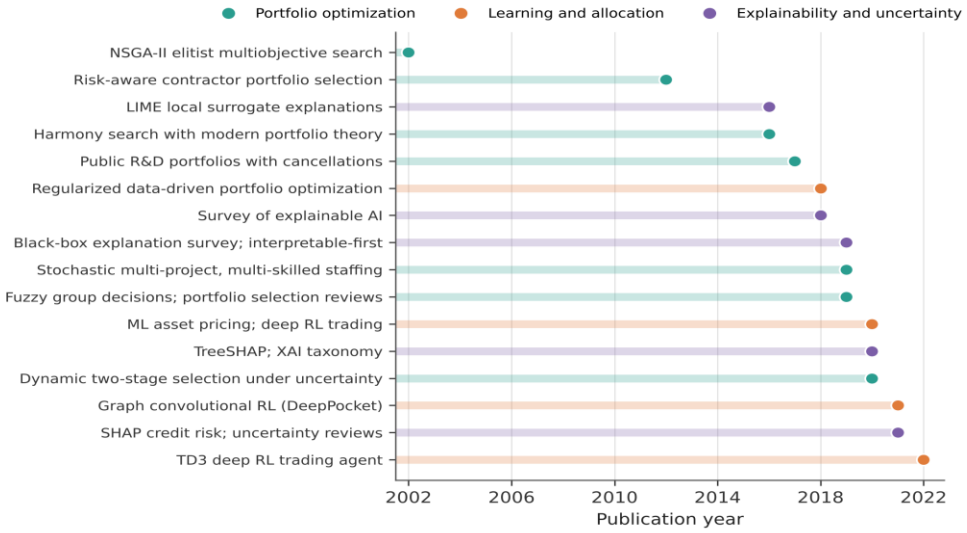


Fig. 6. Chronological Development of the Reviewed Contributions Across the Three Thematic Streams

We present a first comparison between explanation fidelity and cost. Local interpretable model-agnostic explanations (LIME) [20] fit an interpretable surrogate g in a neighborhood πx of the instance x being explained by solving Eq. 8 for a tradeoff between local infidelity L and model complexity Ω, then a submodular pick procedure chooses a minimal, non-redundant set of instances that summarizes the overall behavior. LIME is also model agnostic, low cost, and can explain the model, while TreeSHAP can give exact and consistent attributions only for tree-based valuation models. These trade-offs are summarized in Table 6.

$$\xi(x) = \arg\min_{g \in G} L(f, g, \pi_x) + \Omega(g) \qquad (8)$$

A second comparative analysis is whether or not there is a role for post-hoc explanation in high stakes decisions. There has been some work on explaining black-box models, which must be approximate and may even be unfaithful to the underlying computation, and work has also been done on interpretable models that often are similarly accurate on structured data, in which case interpretable-first design should be favored if there are requirements for accountability [21]. Engineering capital allocation is a criterion that is met by the irreversible commitment of public or shareholder funds in the form of engineering. The synthesis then prioritizes transparent valuation models or decomposable valuation models where data is tabular, and of modest size, and opaque learners are used for data-compressed subproblems with faithful attribution and documented validation.

TABLE VI. COMPARISON OF EXPLANATION APPROACHES FOR PORTFOLIO DECISIONS

Approach	Scope	Model dependence	Computational cost	Fidelity	Portfolio use
Transparent fuzzy multicriteria ranking	Global	Intrinsic	Low	Exact by construction	Criteria-based ranking
Interpretable-first models	Global and local	Intrinsic	Low to moderate	Exact	Valuation of tabular project data
LIME	Local; global via submodular pick	Model-agnostic	Moderate; sampling	Approximate; neighborhood-dependent	Query on a single project
Exact Shapley values	Local to global	Model-agnostic	2^M coalitions; M = 30 gives 1.07×10^9	Exact and axiomatic	Project-level attribution of portfolio value
TreeSHAP	Local to global	Tree ensembles	$O(TLD^2)$; T = 500, L = 64, D = 6 gives 1.15×10^6	Exact for trees	Explanation of learned valuations

International Journal of Business Quantitative Economics and Applied Management Research
Volume-7, Issue-9, 2023 ISSN No: 2349-5677

Approach	Scope	Model dependence	Computational cost	Fidelity	Portfolio use
Contrastive explanation	Local	Model-agnostic	Low once frontier exists	Depends on frontier quality	Why portfolio A rather than B

VII. INTEGRATED ARCHITECTURE AND OPEN CHALLENGES

Graph-structured learning is a way to model the interdependencies of engineering portfolios. DeepPocket used graph nodes to represent financial assets, and edges to represent the relationships between them, extracting features with a restricted stacked autoencoder, and learned allocation policies with actor–critic networks of graph convolutions, including in a period of extreme market turbulence [22]. Immediate: If there are shared resources, technical requirements, common risk contributors, etc., they can be represented as edges, and the valuation layer can flow between projects that are related.

Multistage formulations with recourse are best suited to describe dynamics. A two-stage mathematical programming model of project evaluation and selection under uncertainty showed that the linkage of evaluation and selection can be adjusted throughout multiple planning periods, as the uncertain parameters are realized, and that the two-stage model can provide superior results [23]. The first stage selection x is then determined to maximize Q(x, ξ) in the generic form in Eq. 9 over the set of uncertain scenarios ξ. Fig. 7 combines these into a five-layer explainable portfolio intelligence architecture that initiates re-optimization periodically as information flows in through a governance loop.

$$\max_{\mathbf{x}} \mathbf{c}^T\mathbf{x} + E_{\xi}[Q(\mathbf{x},\xi)], \text{ s.t. } \mathbf{Ax} \leq \mathbf{b} \quad (9)$$

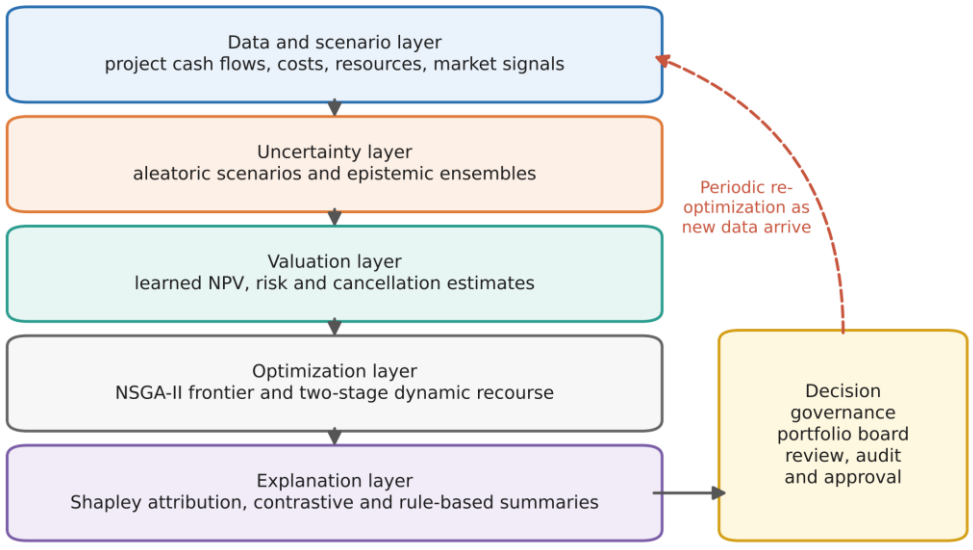


Fig. 7. Five-Layer Explainable Portfolio Intelligence Architecture with Governance Feedback and Periodic Re-Optimization

There are several open challenges that still exist, as listed in Table 7. Incorporating volatility scaling into DRL policies trained on 50 liquid futures contracts between 2011 and 2019 outperformed classical time-series momentum strategies, but in the case of engineering portfolios, there are orders of magnitude fewer decision episodes, hence the need for sample efficiency, simulation fidelity and validation. Other challenges include costs associated with exact attribution as the portfolios expand, the risk of explanations being used to justify decisions instead of to question them, and the requirement for audit trails that meet the needs of procurement and public investment regulation. There is a need for assessment of the quality of the explanation provided with decision makers, not just using the algorithmic metrics, in order to ensure ethical deployment.

TABLE VII. OPEN CHALLENGES AND CORRESPONDING OPTIMIZATION DIRECTIONS

Challenge	Manifestation in engineering portfolios	Optimization direction
Data scarcity	Few comparable historical projects for training	Regularized estimation, transfer from analogous domains, simulation
Epistemic uncertainty	Novel technologies and untested delivery models	Ensemble uncertainty estimates and staged investment triggers
Attribution cost	Coalitions grow as 2^n with portfolio size	Tree-based exact methods, sampling approximations, project grouping
Explanation fidelity	Surrogates drift from the valuation model	Interpretable-first models and reported fidelity measures
Dynamic re-planning	Cancellations, delays and resource conflicts	Two-stage recourse and rolling-horizon re-optimization
Governance and ethics	Public and shareholder accountability	Contrastive reports, audit logs and mandatory human review
Regulatory compliance	Procurement and investment approval rules	Documented model validation and traceable decision records

VIII. CONCLUSION AND FUTURE DIRECTIONS

A total of 24 sources were reviewed and synthesized from engineering portfolio selection, learning-based valuation, uncertainty quantification and explainable artificial intelligence. Its contributions include a three-stream classification of the literature, a comparison of model families, computational cost, reported benchmarks and explanation approaches, a worked 10-project example that quantifies the risk–return trade-off and demonstrates how to perform exact Shapley attribution of the value of a portfolio, and a five-layer architecture that combines uncertainty separation, learned valuation, multiobjective optimization and two-stage optimization, as well as governed explanation. The key takeaway is that explaining is a key component of portfolio intelligence, not something added on when the time is right; and that interpretable-first valuation should be encouraged, wherever capital is irrevocable and accountable to the public. Future work should confirm the architecture on actual multi-project

engineering datasets, assess the contrastive explanations using portfolio boards and create scalable methods for attribution that are still exact when the number of candidates increases.

REFERENCES

1. H. Abbasianjahromi and H. Rajaie, "Developing a project portfolio selection model for contractor firms considering the risk factor," Journal of Civil Engineering and Management, vol. 18, no. 6, pp. 879–889, 2012, doi: 10.3846/13923730.2012.734856.
2. M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U. R. Acharya, V. Makarenkov, and S. Nahavandi, "A review of uncertainty quantification in deep learning: Techniques, applications and challenges," Information Fusion, vol. 76, pp. 243–297, 2021, doi: 10.1016/j.inffus.2021.05.008.
3. Adadi and M. Berrada, "Peeking inside the black-box: A survey on explainable artificial intelligence (XAI)," IEEE Access, vol. 6, pp. 52138–52160, 2018, doi: 10.1109/ACCESS.2018.2870052.
4. G.-Y. Ban, N. El Karoui, and A. E. B. Lim, "Machine learning and portfolio optimization," Management Science, vol. 64, no. 3, pp. 1136–1154, 2018, doi: 10.1287/mnsc.2016.2644.
5. A. Barredo Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. García, S. Gil-López, D. Molina, R. Benjamins, R. Chatila, and F. Herrera, "Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," Information Fusion, vol. 58, pp. 82–115, 2020, doi: 10.1016/j.inffus.2019.12.012.
6. N. Bussmann, P. Giudici, D. Marinelli, and J. Papenbrock, "Explainable machine learning in credit risk management," Computational Economics, vol. 57, no. 1, pp. 203–216, 2021, doi: 10.1007/s10614-020-10042-0.
7. M. Çağlar and S. Gürel, "Public R&D project portfolio selection problem with cancellations," OR Spectrum, vol. 39, no. 3, pp. 659–687, 2017, doi: 10.1007/s00291-016-0468-5.
8. K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," IEEE Transactions on Evolutionary Computation, vol. 6, no. 2, pp. 182–197, 2002, doi: 10.1109/4235.996017.
9. H. N. Esfahani, M. H. Sobhiyah, and V. R. Yousefi, "Project portfolio selection via harmony search algorithm and modern portfolio theory," Procedia - Social and Behavioral Sciences, vol. 226, pp. 51–58, 2016, doi: 10.1016/j.sbspro.2016.06.161.
10. T. Felberbauer, W. J. Gutjahr, and K. F. Doerner, "Stochastic project management: Multiple projects with multi-skilled human resources," Journal of Scheduling, vol. 22, no. 3, pp. 271–288, 2019, doi: 10.1007/s10951-018-0592-y.
11. S. Gu, B. Kelly, and D. Xiu, "Empirical asset pricing via machine learning," The Review of Financial Studies, vol. 33, no. 5, pp. 2223–2273, 2020, doi: 10.1093/rfs/hhaa009.
12. R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi, "A survey of methods for explaining black box models," ACM Computing Surveys, vol. 51, no. 5, Art. no. 93, 2019, doi: 10.1145/3236009.

13. Hüllermeier and W. Waegeman, "Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods," Machine Learning, vol. 110, no. 3, pp. 457–506, 2021, doi: 10.1007/s10994-021-05946-3.
14. T. Kabbani and E. Duman, "Deep reinforcement learning approach for trading automation in the stock market," IEEE Access, vol. 10, pp. 93564–93574, 2022, doi: 10.1109/ACCESS.2022.3203697.
15. S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee, "From local explanations to global understanding with explainable AI for trees," Nature Machine Intelligence, vol. 2, no. 1, pp. 56–67, 2020, doi: 10.1038/s42256-019-0138-9.
16. T. Miller, "Explanation in artificial intelligence: Insights from the social sciences," Artificial Intelligence, vol. 267, pp. 1–38, 2019, doi: 10.1016/j.artint.2018.07.007.
17. V. Mohagheghi, S. M. Mousavi, J. Antuchevičienė, and Y. Dorfeshan, "Sustainable infrastructure project selection by a new group decision-making framework introducing MORAS method in an interval type 2 fuzzy environment," International Journal of Strategic Property Management, vol. 23, no. 6, pp. 390–404, 2019, doi: 10.3846/ijspm.2019.10536.
18. V. Mohagheghi, S. M. Mousavi, J. Antuchevičienė, and M. Mojtahedi, "Project portfolio selection problems: A review of models, uncertainty approaches, solution techniques, and case studies," Technological and Economic Development of Economy, vol. 25, no. 6, pp. 1380–1412, 2019, doi: 10.3846/tede.2019.11410.
19. V. Mohagheghi, S. M. Mousavi, and M. Mojtahedi, "Project portfolio selection problems: Two decades review from 1999 to 2019," Journal of Intelligent & Fuzzy Systems, vol. 38, no. 2, pp. 1675–1689, 2020, doi: 10.3233/JIFS-182847.
20. M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why should I trust you?': Explaining the predictions of any classifier," in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Association for Computing Machinery, 2016, pp. 1135–1144, doi: 10.1145/2939672.2939778.
21. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," Nature Machine Intelligence, vol. 1, no. 5, pp. 206–215, 2019, doi: 10.1038/s42256-019-0048-x.
22. Soleymani and E. Paquet, "Deep graph convolutional reinforcement learning for financial portfolio management – DeepPocket," Expert Systems with Applications, vol. 182, Art. no. 115127, 2021, doi: 10.1016/j.eswa.2021.115127.
23. M. Tavana, G. Khosrojerdi, H. Mina, and A. Rahman, "A new dynamic two-stage mathematical programming model under uncertainty for project evaluation and selection," Computers & Industrial Engineering, vol. 149, Art. no. 106795, 2020, doi: 10.1016/j.cie.2020.106795.
24. Z. Zhang, S. Zohren, and S. Roberts, "Deep reinforcement learning for trading," The Journal of Financial Data Science, vol. 2, no. 2, pp. 25–40, 2020, doi: 10.3905/jfds.2020.1.030.