

THE ROLE OF MACHINE LEARNING FOR FLEXIBILITY AND REAL OPTIONS ANALYSIS IN ENGINEERING SYSTEMS DESIGN

Caputo, Cesare;
Cardin, Michel-Alexandre

Imperial College London, Dyson School of Design Engineering

ABSTRACT

Flexibility analysis helps improve the expected value of engineering systems under uncertainty (economic and/or social). Designing for flexibility, however, can be challenging as a large number of design variables, parameters, uncertainty drivers, decision making possibilities and metrics must be considered. Many available techniques either rely on assumptions that are not suitable for an engineering setting, or may be limited due to computational intractability. This paper makes the case for an increased integration of Machine Learning into flexibility and real options analysis in engineering systems design to complement existing design methods. Several synergies are found and discussed critically between the fields in order to explore better solutions that may exist by analyzing the data, which may not be intuitive to domain experts. Reinforcement Learning is particularly promising as a result of the theoretical common grounds with latest methodological developments e.g. decision-rule based real options analysis. Relevance to the field of computational creativity is examined, and potential avenues for further research are identified. The proposed concepts are illustrated through the design of an example infrastructure system.

Keywords: Flexibility in Design, Large-scale engineering systems, Machine learning, Real Options, Computational design methods

Contact:

Cardin, Michel-Alexandre
Imperial College London
Dyson School of Design Engineering
United Kingdom
m.cardin@imperial.ac.uk

Cite this article: Caputo, C., Cardin, M.-A. (2021) 'The Role of Machine Learning for Flexibility and Real Options Analysis in Engineering Systems Design', in *Proceedings of the International Conference on Engineering Design (ICED21)*, Gothenburg, Sweden, 16-20 August 2021. DOI:10.1017/pds.2021.573

1 INTRODUCTION

Climate change impacts, pandemics, evolving economies, volatile markets and innovation are just a few examples of the sources of uncertainty driving the need for adaptability as a means to improve lifetime performance of engineering systems, making their design increasingly challenging. These systems are characterized by a high degree of complexity, multidisciplinary interactions, and large irreversible capital investments with a long operating life. Uncertainty in the context of an engineering design problem may then be defined as making up the difference between the knowledge available to decision makers and that necessary to optimally plan the system (de Neufville et al., 2019). Given system design configurations are often determined in early decision-making phases, current design frameworks tend to simplify considerations of uncertainty, partly to maintain tractability and partly because of limited knowledge, leading to designs that can be highly vulnerable to change (de Neufville and Scholtes, 2011). Flexibility has emerged as a potential and underexploited paradigm to help systematically improve system performance under uncertainty, by allowing it to capture more upside opportunities and reduce downside exposure through the potential for reconfiguration and changeability. The definition of flexibility used here, inspired from real options theory, is the “right, but not the obligation, to change a system in the face of uncertainty” (Trigeorgis, 1996). The higher the associated level of uncertainty, generally the more enhanced the value of embedding flexibility into the system for its expected performance. Flexibility in design, however, is a paradigm which is not fully exploited in the engineering community as decision makers often do not fully understand the added value it may bring to a project, while the costs are readily evident. There are tools available already which are well suited to designing for flexibility, both qualitative and quantitative based, helping to identify possible strategies and estimating their value. These can be very important as they help bridge the gap between the perceived and actual value of inclusion of flexibility measures in the system. In fact, previous studies have shown improvements routinely ranging from 10-30% in economic terms when using flexible design methods compared to standard ones, with applications across energy (Kuznetsova et al., 2019), infrastructure (Cardin et al., 2015) and telecommunications (de Weck et al., 2004), among others. Flexibility, therefore, changes the objective from optimizing for a “design point” to finding solutions that maintain strong performance across the entire distribution of outcomes (Cardin, 2014). Nonetheless, current flexibility design tools present significant limitations on scope, applicability and integration into holistic systems engineering methods. The complexity of the design problem can be very significant, as a large number of design variables, parameters, uncertainties, decision making possibilities and metrics must be considered in the context of the system, with a great computational expense. The majority of flexibility design tools, partly to maintain computational feasibility, are therefore focused on recognizing the value of flexibility measures rather than generating designs or suggestions for how to implement them. There has been little work done on stimulating creativity for alternative solutions by exploration of the design space, with some qualitative based frameworks or generic flexibility strategies normally employed. These generic strategies, however, are not well suited to easily identify system specific flexibility value enhancing measures outside of a few applications.

The complexity of flexible design, therefore, tends to necessitate taking assumptions which may be unrealistic in an engineering setting and/or a significant limitation of the design space from that found in the true problem. These methodological limitations, combined with the emergence of newly available datasets, such as patents or open access design databases, highlight the potential role for Machine Learning (ML) in this field. By extracting useful patterns out of available datasets and increasing the ability to tackle high dimensional problems with neural networks, the role of ML within engineering design offers great promise and has been increasingly investigated in recent years (Rolnick et al., 2019). This paper makes the case for further integration of ML techniques into the design process for flexible engineering systems, shifting the focus from previous works by specifically focusing on ML as an enabler of flexibility. Starting from a short review of the literature on flexibility in design tools and ML, the argument is made that there may be a number of optimal solutions embedded within the design space not being considered with standard methods. Several synergies are identified between ML and flexibility in design which point to the potential for greater combination of these two fields and a preliminary framework integrating some of the methods identified as most promising is presented, based on an adaptation of some previous work, illustrated through a simple design case study. Some general recommendations and directions for further research to complement available techniques are then made.

2 BACKGROUND AND RELATED WORK

2.1 Flexibility in Design

A flexible design provides the potential for changeability of a system under different and unpredictable conditions through inclusion of some potential strategy into the system configuration. Typical project evaluation approaches for engineering systems tend to be based on likely scenario predictions for major uncertainty drivers and Discounted Cash Flow (DCF) analysis, which is not able to account for the value of decision making over the project lifetime (de Neufville et al., 2019). A review of flexibility in engineering design from (Cardin, 2014) builds upon the definition of a flexible design concept as providing the ability for an engineering system to "adapt, change and be reconfigured, if needed, in light of uncertainty realizations". This definition, therefore, extends past the technical domain to consider managerial decision-making as an enabler of valuable flexibility during operation, also known as a "real option". A real option "in" the system is composed of flexibility within the design, also referred as technical enablers, which give decision makers the choice to strategically decide when and how to exercise this flexibility based on uncertainty realizations. Real option "on" the system involves managerial flexibility, and do not require as much technical and engineering expertise to be enabled e.g., abandonment, investment deferral (Bowman and Moskowitz, 2001; Trigeorgis, 1996).

Previous works have generally attempted at structuring designing for flexibility through a sequential process workflow in order to provide standardized and reproducible procedures. In the framework and taxonomy presented by Cardin (2014), five main phases are recognized for a flexible design process: 1) Baseline Design; 2) Uncertainty Recognition; 3) Concept Generation; 4) Design Space Exploration; 5) Process Management. Baseline design, conducted with traditional methods, is necessary to provide a point of comparison and understand relative performance of outputs. Procedures in phase 1 can help to generate new design solutions, enrich current information on available solutions or help navigate the entire design process. Uncertainty recognition in phase 2 is crucial and can be a defining factor in the design of a project subject to stochasticity, as estimates from this phase may skew results towards largely unrealistic conditions or ill-prepared designs for realized scenarios. Procedures in this phase are largely based on probabilistic or statistical approaches to help designers identify, quantify, and characterize the major uncertainty drivers (de Weck et al., 2007). Practical approaches to model them include stochastic simulation-based diffusion models, binomial lattice, decision trees and scenario planning. The selection of input parameters, assumptions and mathematical approaches have a significant effect on predictive accuracy, with different techniques better suited for different types of uncertainties or design problems - although all presenting some limitations. The most relevant phases for this paper are Concept Generation and Design Space Exploration, since they benefit largely from data-driven techniques.

2.1.1 Concept Generation

The tools available for concept generation, with a particular focus on enhancing system flexibility, can be quite varied, ranging from qualitative to computational approaches, and can be further split into strategy generation and enabler identification (Cardin, 2014) - or real options "on" and "in" systems, respectively. These two parts are complementary as they both must be present to allow exercising system flexibility under uncertainty, during operations. During the strategy generation phase, generic real options strategies can be used to guide creativity systematically e.g., (1) defer capital investment until favourable market conditions arise, (2) stage or phase asset deployment strategically over time instead of deploying all capacity at once, (3) alter operating scale by expanding or contracting output production capacity, (4) abandon a project doomed to fail and resell assets at salvage value, (5) switch production inputs and/or outputs to accommodate different markets or missions, (6) invest in R&D to capitalize on future technology and cash flows if the initial investment is successful, and/or (7) combine the above (Trigeorgis, 1996). Creativity then stems from the need to understand how to integrate these approaches into the specific system at hand. The Integrated Real Options Framework formalized this latter step by mapping potential strategies to physical enablers in the design. The logical multiple domain matrix is an extension of this work to generate concepts for flexibility via a more holistic analysis of interactions among the system domains (Mikaelian et al., 2011).

In terms of enabler identification, a variety of approaches are also present. An initial starting point may be product or system safeguards, connecting modularity to degree of uncertainty as shown by Gil (2007). Another commonly used tool in this phase is the Design Structure Matrix (DSM). The DSM lists system

components and domains in the rows and columns of a square matrix in order to map out dependencies and relationships among them. It is often used in combination with Change Propagation Analysis (CPA) to understand the response of different design variables to changes in design or functional requirements and identify candidate areas for value enhancing flexibility (Giffin et al., 2009). The method has been previously successfully extended to include non-technical domains (e.g., Engineering Systems Matrix from Bartolomei et al. (2012)) as well as Bayesian network analysis (Hu and Cardin, 2015), however the effort and level of detail required to build this representation can be limiting. A similar analytical approach is followed when using the Theory of Inventive Problem Solving (TRIZ) in the context of designing for change to identify candidate areas for modularization (Chechurin and Collan, 2019). The Filtered outdegree method assigns a value to the perceived degree of changeability of a system, with possible end states filtered based on a transition cost function and CPA (Shah et al. 2008). While this method is most useful for measuring flexibility in a system, it can also help identify enablers through embedding of real options when combined with CPA or other techniques, essentially serving to decrease the cost of transition in the outdegree (Shah et al. 2008). Explicit design variable evaluation is another approach, consisting of identifying the main design variables and quantitatively evaluating their impact on lifetime performance, and thus identifying most impactful changes within the design (Cardin, 2014).

2.1.2 Design Space Exploration

Following the generation of flexible design concepts in phase 3, potential strategies must be evaluated on a comparative basis to understand what the best performing system solutions may be. Real options theory stems from the financial literature theory, and thus gives an array of quantitative methods to quantify the value of flexibility under different scenarios, allowing better comparison of alternative design concepts (de Neufville and Scholtes, 2011). Dynamic (DP) and stochastic programming (SP) can be used for flexibility analysis, seeking to compute an optimal policy in form of maximization of the recursive Bellman reward equation, normally by dividing a complex problem into a sequence of simpler subproblems, starting at the final stage (Sutton and Barto, 2018). This folding back process continues to project start, at which point expected lifetime performance can be computed and compared to inflexible baselines. DP is better suited to handle infinite horizon problems or problems with large number of stages compared to SP. Both approaches are however limited by data inefficiency, memory storage requirements and dimensionality (Martínez-Costa et al., 2014). Binomial lattice analysis is similar conceptually to DP, other than the fact that uncertainty either goes up or down following each stage. The value of flexibility in these methods is normally estimated as the difference in expected performance between the baseline and flexible design solutions. These tools, however, present a number of assumptions on path independence, knowledge of precise probability distributions and market conditions such as arbitrage enforced pricing which make it usually unsuitable for implementation in an engineering setting, as decisions made during early design phases may influence the techno-economic feasibility of various lifetime measures (Cardin, 2014). Furthermore, results are of limited relevance if realised uncertainties deviate from predicted scenarios (Martínez-Costa et al., 2014).

A number of approaches have been proposed in recent years to help address these limitations such as screening methods (Bartolomei et al., 2006) or design catalogues (Cardin and Geltner, 2016). Drawbacks of these techniques, however, may include model resolution loss, limited guarantee of solution global optimality, and difficulties with estimating aggregate preferences to calculate expected utility in a design context (Cardin, 2014). Decision rules (DR) is another method, defined as a triggering mechanism or signal to exercise a particular flexibility on the system (Caunhye and Cardin, 2017). Conditional-go, linear and constant DR are some of the generic options for problem formulation. The DR approach has been previously shown to estimate very similar value for flexibility measures compared to standard ROA or DP, with better computational efficiency while providing readily usable guidelines (Cardin et al., 2017). The benefit of this approach is that it can model flexibility in a more intuitive manner for decision makers by allowing straightforward combination of physical design variables with managerial aspects affecting operation, giving insights on optimal exercise time. This allows for both to be simultaneously optimized stochastically with a more transparent approach than standard DP. The DR method is, however, fairly static in nature in the sense that the stochastically optimal DR are normally determined at project start and not adjusted thereafter. Thus, while it allows some degree of flexibility as it permits system response to realized scenarios and trends, it is incomplete in that it does not allow more adaptive decision making and optimization of performance over time, only making limited use of available information.

2.2 Machine Learning

ML is the general term describing the field of study of using algorithms and computational power to extract usable models and patterns out of empirical data. Generally, it can be categorized into supervised and unsupervised learning and there is an abundance of potential algorithms available. Supervised learning provides labelled datasets to predict their relationship and is primarily used in classification and regression. Unsupervised learning instead seeks to find patterns and distributions from unlabelled data and can be used in clustering or feature extraction applications, among others. Recent advances can be largely attributed to the integration of "Deep Learning". This layer-by-layer learning structure within a neural network allows for much more complex feature extraction than was possible before. The input to each layer after the first is normally composed of a weighted sum from the previous one, transformed via a generally nonlinear activation function to link the information between different layers and perform the calculations. Throughout training, these weights between layers are updated to optimize the loss function defined initially, using gradient information across the layers (Schmidhuber, 2015). Depending on how the different layers are connected, they may be referred to as multilayer perceptron (MLP), convolutional neural network (CNN), recurrent neural network (RNN) or long short-term memory network (LSTM), among others, with each architecture appropriate for specific tasks.

The potential of ML and data driven methodologies to enhance the systems design process is of great interest, and examples of potential applications within the engineering design space are abundant (Panchal et al., 2019). There has been great success working with and extracting patterns out of time series data to increase prediction and simulation accuracy (Rolnick et al., 2019). These models can handle much larger data sets than would be feasible with conventional approaches to reproduce the behaviour of complex non-linear systems, often with much better computational efficiency than analytical solutions. In the design of renewable energy systems, for example, neural networks and support vector regression are being increasingly used to predict operating conditions and better characterize the technical requirements to meet project objectives throughout its lifetime (Voyant et al. 2017). In the design space exploration and optimization phase, classification algorithms can be used to identify appropriate regions of feasible design space among a number of candidate solutions. Mapping the boundaries of the feasible design space can help guide engineers on more efficient exploration of the space while limiting the need for computationally expensive simulations at every stage. Support Vector Machines, Random forests, Gaussian Naive Bayes and ANNs are some of the prevalent techniques for this purpose that were comparatively evaluated by (Sharpe et al. 2019) on several different design problems. The findings from their study suggest that while no approach is universally dominant, these techniques hold significant potential in helping to facilitate and speed up design space exploration across various disciplines. The success stories of these applications have, however, been heavily dependent on the quality and quantity of data fed to the network, which tends to require significant time-intensive data pre-processing to avoid issues such as bias or overfitting, among others. The majority of applications so far have focused on increased accuracy of prediction, scenarios, surrogate modelling, and classification as inputs into design but have yet to be integrated into an end-to-end design process to stimulate creativity about alternative flexibilities or strategies.

In the review by Toivonen and Gross (2015) on the intersection between ML and computational creativity, the authors argue that while the most obvious use of ML would be as a means of increasing exploratory creativity, there are a number of other potential applications. For instance, classification or regression could be used to learn the evaluative or objective function to be used in the design optimization procedure, thus developing a new way to analyse the system at hand. Another application could be using predictive models to generate new designs or solutions based on incomplete ones (combinational creativity), which could be extended to learning the generative model directly (exploratory creativity). Data mining is also argued to be useful for discovering nonobvious links among system components and the resulting implications for design.

2.3 Reinforcement Learning

Reinforcement learning (RL) does not fall within either of the traditional ML paradigms of supervised or unsupervised learning. It involves an agent interacting with an environment over time as part of a sequential decision-making problem with the objective to maximize a reward signal. The past few years have witnessed a strong resurgence of RL as a result of its combination with deep neural networks as well as advances in computational capacity and algorithmic techniques, leading to great achievements via Deep RL (DRL). DRL is much better suited to handle time-dependent, control tasks

and uncertainty planning under conflicting signals than traditional ML (Sutton and Barto, 2018). This is essential as it is at the core of the motivations for flexibility in design, although with less limiting assumptions than in standard methods. Generally, each time step (t) the agent finds itself at a state (s_t) and selects an action (a_t) following a policy $\pi(s_t / a_t)$. As a result of taking the action, the agent transitions to the successive state (s_{t+1}), receiving a corresponding scalar reward (R_t), according to the environment's dynamics or model for reward function, and state transition probability (Sutton and Barto, 2018). This repeated interaction is captured in Figure 1, including an illustration of the role of neural networks to approximate the value function within the overall process.




Figure 1: Graphical overview of DRL (Mao et al. 2016)

$$G_t = \sum_{t=0}^{\infty} \gamma^t R_t \quad (1)$$

The reward signals the agent seeks to maximize may also be referred to as the return (G_t), or the discounted accumulated reward per episode (R_t), as described by equation 1, where γ refers to the discount rate used in financial modelling to account for the time value of money. This is comparable to the objective function in a traditional systems engineering optimization problem and is used by the agent to formulate a decision-making policy to navigate the environment. The majority of RL implementations in engineering have centred on more classical control problems, which can more easily fit into its formalization and stimulate agent learning, such as in optimal microgrid operations (Perera et al. 2020). As far as the authors know, there is still no use of DRL in the context of flexibility in systems design.

3 ML AUGMENTED FLEXIBLE DESIGN FRAMEWORK

The proposed ways in which ML could be integrated into a flexibility design framework are summarized in Figure 2, with further details in section 4. The motivation and potential will be illustrated through a simple case study of the design of a multi-story parking garage subject to demand and price uncertainty, inspired from the example by Cardin et al. (2015). The inflexible baseline would be to construct everything at project start and benefit from economies of scale, while the flexible alternative starts with a smaller capacity but with the ability to expand later on. There are a number of possible design alternatives that deliver the same flexible expansion capability, such as inclusion of larger foundations, stronger material, or fully modular implementation (i.e., stacking prefabricated levels on top of each other). No specific ML integration is included for Phase 1 (Baseline Design), however any ML method applicable to general engineering design could be used during this stage to define the optimal inflexible baseline. During phase 2, ML could be used to build more accurate models to understand how price and demand uncertainty will evolve over time and help make better informed design decisions on how flexibility could be most impactful. This could lead designers to try to focus on either generic capacity expansion, or for example target expansion for pricing past a certain threshold only. Surrogate modelling may be implemented at this stage to determine expansion cost function from a dataset of previous expansion decisions in other projects more accurately and easily than via analytical means. ML classification or clustering algorithms could then be used to identify the most flexible regions of the design space, rather than simply the feasible ones, and understand the relevant trade-offs during concept generation. The clustering criteria could be varied for a more holistic evaluation, in order to assess for example regions of the system design space that are well suited to respond to price but not demand uncertainty, and vice versa and the resulting design implications. Based on the results from phase 2, decision makers could choose to focus on designs allowing adaptability to demand uncertainty, for example, and look at their common value adding features.

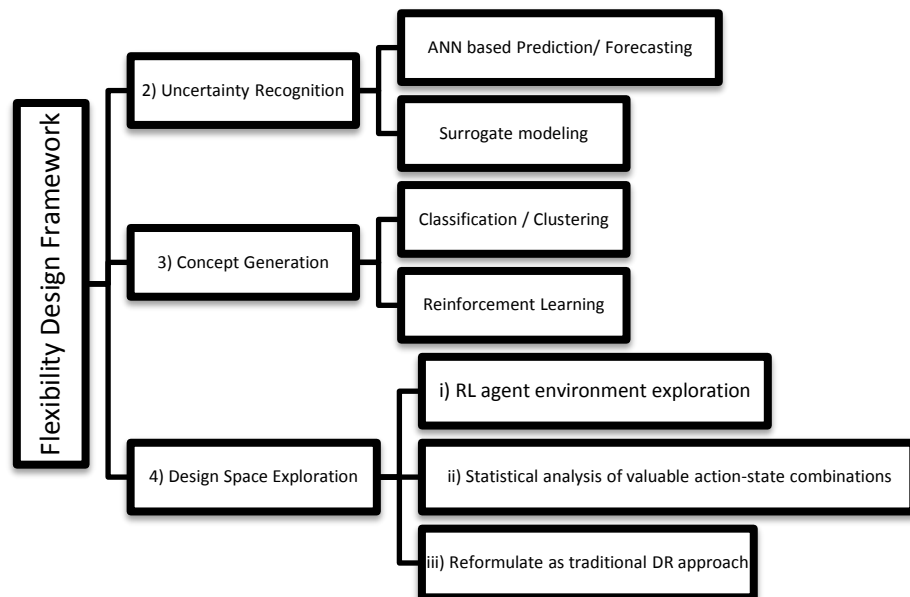


Figure 2: Proposed integration of ML into design framework from (Cardin, 2014)

In terms of RL applications, there is immediate potential for supporting the decision rule method to flexibility analysis, particularly during phase 4, as shown in Figure 2. A decision rule approach for the multi-story garage design could be formulated as (IF demand for the last 2 years exceeds current capacity THEN expand capacity by 1 floor, up to max allowable capacity, ELSE do nothing). Once flexible design concepts and strategies are generated in phase 3 and all relevant actions recognized within the RL environment, the training of the agent would be launched. In the garage case, the step 0 decision space would include the decision on starting levels of floors and flexibility enabler (i.e., stronger columns or prefabricated containers). This would help integrate the early design decision with operational insights for responding to uncertainties. It could also be better suited to produce combinations of decision rules by more systematically exploring the design space for value enhancing measures. For example, this RL based exploration may suggest building 1 floor only once capacity reaches a certain level to avoid investment risk, while building 2 or 3 floors up to that level in order to ramp up demand more quickly. Another purely theoretical example in the garage case may be that the agent finds building an extra half floor on some expansion years as part of its optimal policy, which would likely not have been considered in a domain expert defined flexibility analysis. This could allow for the garage to ramp up capacity more quickly if demand greatly exceeds expectations on some years, reducing lost revenue during the building period and reducing liquidity risks. Advanced statistical analysis of optimal action-state combinations could yield a more "explainable" process. These value enhancing actions could be re-formulated as a traditional DR approach, providing readily usable guidelines for decision makers.

4 DISCUSSION

The review of the relevant literature suggests there is potential for the development of novel tools to analyse flexibility in systems design exploiting recent advances in machine learning. Current methods available to designers and engineers are well suited for the enhancement of system flexibility compared to a traditional or baseline design process, but present a number of analytical and practical limitations. They tend to rely on either domain expertise or generic approaches and thus are limited in the number and type of solutions they are able to generate in phase 3 of the design process, based on what has previously been done or is known to decision makers. There is likely solutions and configurations embedded within the increasingly available data, from sources such as patents and publications, which can increase delivered expected value that may not be fully considered by engineers and decision makers. The argument here is not for ML to be used for the standalone design of flexible engineering systems or to generate mathematically more optimal solutions, but rather as an important tool that can help practitioners complement standard methods in order to achieve better performing solutions. Below is a discussion of how ML techniques can be used to assist the different phases of the design process for flexibility.

4.1 Enhanced Uncertainty Recognition via ANNs

Increasing accuracy and understanding of relevant uncertainties is essential in any engineering design problem. We argue, however, that they are particularly important when designing for flexibility as the objective is to maximize adaptability to realised uncertainties. In the simplest of cases, ML methods may be used to generate more representative input scenarios and parameters to be used in scenario planning, decision trees or other relevant techniques from 2.1. When accounting for highly complex physics-based problems, surrogate modelling may also be helpful to simulate scenarios helping to understand potential random shocks from external factors at a greatly reduced computational expense as shown by (Sarkar et al. 2019). This may also allow a much higher number of scenarios to be evaluated within the design process. Historical data on trends for the uncertain variable would, however, likely be necessary for improved accuracy, which may become a limiting factor in a lot of applications.

4.2 Classification and Clustering for Flexible Concept Generation

Appropriate clustering and classification implementations during phase 3 could greatly accelerate the design process for flexibility. They may be used to provide a list of common features for previous system designs well suited to deal with a particular uncertainty identified from phase 2, or to identify how a specific design compares to others in the field. The advantage of ML is that it may better tackle computational challenges as compared to existing techniques. Furthermore, it is reasonable to assume that flexible designs for a particular system may share common features (i.e., stronger columns for garage or oversized inverter for solar PV system) not found in their inflexible counterparts. These methods would allow for much faster identification of the features and for designers to choose to prioritise them if needed from very early stages. It may then be possible to build a predictive model that is able to generate a flexible design solution based on an inflexible baseline or a few basic inputs, for example only specifying max initial investment and local material costs, based on combinations previously found within the dataset of previous multi-story garage designs. The integration of all these applications, however, would necessarily require a large data set of system designs, both flexible and inflexible. This is not normally available for most design problems; however, it is something that could be developed over time if designs were made more accessible or even built artificially using techniques such as Generative Adversarial Networks (GANs) (Shu et al. 2019). There are also a number of potentially interesting applications more closely related to the field of computational creativity. For instance, ML may be used to develop novel metrics measuring flexibility in system design by looking at a large number of candidate designs. These novel metrics or evaluation function could then be used to estimate the flexibility of different solutions considered later for further evaluation. These ways of estimating flexibility may be unintuitive to practitioners but could uncover potential value enhancing methods or configurations that do not appear flexible at first, or vice versa. For example, using a stronger material may seem advantageous at first but can lead to vulnerability to future prices when expansion decisions are made, thus lowering the actual value of flexibility in the system compared to using a larger amount of a weaker material upfront.

4.3 RL Based Design Space Exploration

RL is likely to be more impactful and effective than any other ML technique for flexibility in design. The sequential decision-making is well suited to examining impact of design and subsequent operational decisions affecting system performance. As such, the impact of exercising different types of flexibility over time can be assessed quantitatively over a large number of scenarios. It should be noted that sequential decision-making nature refers to the RL agent exploration of the environment, rather than the design process itself. It is thus expected that the design process will remain more iterative using the proposed framework, where insights gained from RL exploration over time can be collated to inform design decisions, rather than modelling the agent as the designer directly.

We make the case that RL holds great potential to solve combinatorial optimization problems such as the ones found in flexibility decision analysis, as the reward stochasticity found in decision making and the reduced reliance on human intuition inputs allows it to explore more solutions than possible with comparable methods. For example, it could be used in combination with decision rules to create more dynamic and adaptable systems under uncertainty. Looking at a standard implementation, decision rule variables and their potential feasible ranges are normally identified at start by domain expertise and their optimal values determined through simulation. In an RL based approach, these would not need to be

identified at start but rather just all possible actions recognized within the environment. Over a much larger number of scenarios than possible with standard methods, the RL agent could then develop their own decision rules to maximize utility in an environment which are not limited to the form chosen by the designer at start (i.e., IF-THEN). Using neural networks, they can build a much more complex value function representation and understand when and how to implement different and new types of flexibilities. Furthermore, an RL implementation is better able to handle conflicting objectives and multiple sources of uncertainty than typical decision analysis methods, thus is more representative of the engineering design setting (Perera et al. 2020). This methodology could also enhance exploratory creativity, as the input values for design variables are much less restricted than in standard optimisation procedures given the advantages of parallel computation and resulting algorithms. The exploration of the design space could also be easily extended outside the economic realm in order to enhance system sustainability and resilience by integration directly into the reward function. Furthermore, the theoretical common ground shared with standard methods in decision analysis allows for a more straightforward solution procedure to be created out of it, such as the decision rule one presented above.

5 CONCLUSIONS

The ideas proposed here, while seeming promising, remain unproven outside of a theoretical basis. There are a number of potentially interesting ways in which ML could be used to complement flexibility in design methods, presenting several avenues for further research. The RL based approach seems particularly promising as a result of methodological similarities with standard methods but with assumptions that can be more realistic in a flexibility in design setting, as well as a better ability to handle multiple objectives. Potential for integrating other ML methods remains strong and more closely aligned with that found in a general engineering design setting. Particularly interesting could be applications in energy systems, where the multiple sources of flexibility, increasing uncertainty, conflicting objectives and large action spaces create challenging design problems that may be addressed with the ML techniques. There remains significant room for further research, however, starting from a mathematical implementation, case study and benchmarking against standard methods such as SP on a complex design task. The acceptability by decision makers of design configurations using this method may also remain limited due to its “black-box” style computational properties. In terms of decision analysis, ROA, DP and SP are better understood and more likely to be implemented in practice. Reproducibility has also been found to be inconsistent in RL implementations, particularly when implemented in a new setting such as this. Another issue may be the ease of communication to decision makers when comparing to a standard decision rule approach. Looking at a more thorough statistical analysis of the agent’s optimal policy and actions under different conditions may help yield some valuable insights and be more acceptable to decision makers. Data sets necessary for classical ML applications may remain a limiting factor or require significant domain expertise manipulation for particular applications. Nonetheless, there is certainly potential for increased integration of ML into design for flexibility methods given some of these limitations can be addressed.

REFERENCES

- Bartolomei, J.E. et al., (2006). 9.1.3 Screening for Real Options “In” an Engineering System: A Step Towards Flexible System Development. *INCOSE International Symposium*, 16(1), pp.1241–1257. <http://dx.doi.org/10.1002/j.2334-5837.2006.tb02809.x>.
- Bartolomei, J.E. et al., (2011). Engineering Systems Multiple-Domain Matrix: An organizing framework for modeling large-scale complex systems. *Systems Engineering*, 15(1), pp.41–61. <http://dx.doi.org/10.1002/sys.20193>.
- Boden, M.A., (2009). Computer Models of Creativity. *AI Magazine*, 30(3), p.23. <http://dx.doi.org/10.1609/aimag.v30i3.2254>.
- Bowman, E.H. & Moskowitz, G.T., (2001). Real Options Analysis and Strategic Decision Making. *Organization Science*, 12(6), pp.772–777. Available at: <http://dx.doi.org/10.1287/orsc.12.6.772.10080>.
- Cardin, M.-A., (2014). Enabling Flexibility in Engineering Systems: A Taxonomy of Procedures and a Design Framework. *Journal of Mechanical Design*, 136(1). <http://dx.doi.org/10.1115/1.4025>
- Cardin, M.-A., de Neufville, R. & Geltner, D.M., (2015). Design Catalogs: A Systematic Approach to Design and Value Flexibility in Engineering Systems. *Systems Engineering*, 18(5), pp.453–471.
- Cardin, M.-A. et al., (2013). Empirical evaluation of procedures to generate flexibility in engineering systems and improve lifecycle performance. *Research in Engineering Design*, 24(3), pp.277–295.
- Cardin, M.-A., Xie, Q., Ng, T. S., Wang, S. and Hu, J. (2017) An approach for analyzing and managing flexibility in engineering systems design based on decision rules and multistage stochastic programming. *IISE Transactions*, Vol. 49, No. 1, pp. 1-12.

- Caunhye, A.M. & Cardin, M.-A., (2017). An approach based on robust optimization and decision rules for analyzing real options in engineering systems design. *IISE Transactions*, 49(8), pp.753–767.
- Chechurin, L. & Collan, M. eds., (2019). *Advances in Systematic Creativity*. <http://dx.doi.org/10.1007/978-3-319-78075-7>.
- de Neufville, R. & Scholtes, S., (2011). *Flexibility in Engineering Design*.: <http://dx.doi.org/10.7551/mitpress/8292.001.0001>.
- de Neufville, R. & Smet, K., (2019). Engineering Options Analysis (EOA). *Decision Making under Deep Uncertainty*, pp.117–132.: http://dx.doi.org/10.1007/978-3-030-05252-2_6.
- de Weck, O.L., de Neufville, R. & Chaize, M., (2004). Staged Deployment of Communications Satellite Constellations in Low Earth Orbit. *Journal of Aerospace Computing, Information, and Communication*, 1(3), pp.119–136. <http://dx.doi.org/10.2514/1.6346>.
- Giffin, M., de Weck, O., Bounova, G., Keller, R., Eckert, C. (2009) 'Change propagation analysis in complex technical systems', *Journal of Mechanical Design*, 131(8). <http://dx.doi.org/10.1115/1.3149847>.
- Gil, N., 2007. On the value of project safeguards: Embedding real options in complex products and systems. *Research Policy*, 36(7), pp.980–999. <http://dx.doi.org/10.1016/j.respol.2007.03.004>
- Hatchuel, A. & Weil, B., (2009). C-K design theory: an advanced formulation. *Research in Engineering Design*, 19(4), pp.181–192. <http://dx.doi.org/10.1007/s00163-008-0043-4>.
- Hu, J. & Cardin, M.-A., (2015). Generating flexibility in the design of engineering systems to enable better sustainability and lifecycle performance. *Research in Engineering Design*, 26(2), pp.121–143.
- Kuznetsova, E., Cardin, M.-A., Diao, M. and Zhang, S. (2019) 'Integrated decision-support methodology for combined centralized-decentralized waste-to-energy management systems design', *Renewable and Sustainable Energy Reviews*, 103, 477-500. <http://dx.doi.org/10.1016/j.rser.2018.12.020>.
- Li, Y. (2017) 'Deep reinforcement learning: An overview', *arXiv preprint arXiv:1701.07274*.
- Mao, H., Alizadeh, M., Menache, I. and Kandula, S. (2016) 'Resource management with deep reinforcement learning', in *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*, 50-56. <http://dx.doi.org/10.1145/3005745.3005750>
- Martínez-Costa, C., Mas-Machuca, M., Benedito, E. and Corominas, A. (2014) 'A review of mathematical programming models for strategic capacity planning in manufacturing', *International Journal of Production Economics*, 153, 66-85. <http://dx.doi.org/10.1016/j.ijpe.2014.03.01>
- Mikaelian, T., Rhodes, D.H., Nightingale, D.J. and Hastings, D.E. (2011) 'A logical approach to real options identification with application to UAV systems', *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 42(1), 32-47. <http://dx.doi.org/10.1109/tsmca.2011.2157133>.
- Panchal, J.H., Fuge, M., Liu, Y., Missoum, S. and Tucker, C. (2019) 'Machine Learning for Engineering Design', *Journal of Mechanical Design*, 141(11).
- Perera, A.T.D., Wickramasinghe, P., Nik, V.M. and Scartezzini, J.-L. (2020) 'Introducing reinforcement learning to the energy system design process', *Applied Energy*, 262, 114580.
- Rolnick, D., Donti, P.L., Kaack, L.H., Kochanski, K., Lacoste, A., Sankaran, K., N., Jaques, N. and Waldman-Brown, A. (2019) 'Tackling climate change with machine learning', *arXiv preprint arXiv:1906.05433*.
- Sarkar, S., Mondal, S., Joly, M., Lynch, M.E., Bopardikar, S.D., Acharya, R. and Perdikaris, P. (2019) 'Multifidelity and Multiscale Bayesian Framework for High-Dimensional Engineering Design and Calibration', *Journal of Mechanical Design*, 141(12), available: <http://dx.doi.org/10.1115/1.4044598>.
- Schmidhuber, J. (2015) 'Deep learning in neural networks: An overview', *Neural networks*, 61, 85-117.
- Shah, N.B., Viscito, L., Wilds, J., Ross, A.M. and Hastings, D.E. (2008) 'Quantifying flexibility for architecting changeable systems', in *6th Conference on Systems Engineering Research*, Los Angeles, CA, Citeseer
- Sharpe, C., Wiest, T., Wang, P. and Seepersad, C.C. (2019) 'A comparative evaluation of supervised machine learning classification techniques for design applications', *Journal of Mechanical Design*, 141(12).
- Sutton, R.S. and Barto, A.G. (2018) *Reinforcement learning: An introduction*, MIT press.
- Shu, D., Cunningham, J., Stump, G., Miller, S.W., Yukish, M.A., Simpson, T.W. and Tucker, C.S. (2019) '3D Design Using Generative Adversarial Networks and Physics-Based Validation', *Journal of Mechanical Design*, 142(7), available: <http://dx.doi.org/10.1115/1.4045419>.
- Toivonen, H. and Gross, O. (2015) 'Data mining and machine learning in computational creativity', *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 5(6), 265-275.
- Trigeorgis, L. (1996) *Real options: Managerial flexibility and strategy in resource allocation*, MIT press
- Voyant, C., Notton, G., Kalogirou, S., Nivet, M.-L., Paoli, C., Motte, F. and Fouilloy, A. (2017) 'Machine learning methods for solar radiation forecasting: A review', *Renewable Energy*, 105, 569-582.