

Multi-Objective Coordination Graphs for the Expected Scalarised Returns with Generative Flow Models

Conor F. Hayes^{1,*}, Timothy Verstraeten², Diederik M Roijers^{2,3}, Enda Howley¹, Patrick Mannion¹

*c.hayes13@nuigalway.ie

¹National University of Ireland, Ireland

²Vrije Universiteit Brussel, Belgium

³HU University of Applied Sciences Utrecht, the Netherlands

Abstract

Many real-world problems contain both multiple objectives and agents, where a trade-off exists between objectives. Key to solving such problems is to exploit sparse dependency structures that exist between agents. For example in wind farm control a trade-off exists between maximising power and minimising stress on the systems components. Dependencies between turbines arise due to the wake effect. We model such sparse dependencies between agents as a multi-objective coordination graph (MO-CoG). In multi-objective reinforcement learning (MORL) a utility function is typically used to model a user's preferences over objectives. However, a utility function may be unknown a priori. In such settings a set of possibly optimal policies must be computed. Which policies are optimal depends on which optimality criterion applies. If the utility function of a user is derived from multiple executions of a policy, the scalarised expected returns (SER) must be optimised. If the utility of a user is derived from the single execution of a policy, the expected scalarised returns (ESR) criterion must be optimised. For example, wind farms are subjected to constraints and regulations that must be adhered to at all times, therefore the ESR criterion must be optimised. For MO-CoGs, the state-of-the-art algorithms can only compute a set of optimal policies for the SER criterion, leaving the ESR criterion understudied. To compute a set of optimal policies under the ESR criterion, also known as the *ESR set*, distributions over the returns must be maintained. Therefore, to compute a set of optimal policies under the ESR criterion for MO-CoGs, we present a novel *distributional multi-objective variable elimination* (DMOVE) algorithm. We evaluate DMOVE in realistic wind farm simulations. Given the returns in real-world wind farm settings are continuous, such as generated power, we utilise a generative flow model known as *real-NVP* to learn the continuous return distributions to calculate the *ESR set*.

Keywords: multi-objective, multi-agent, coordination graphs, (nonserial) dynamic programming, expected scalarised returns, generative flow models

1. Introduction

In many real-world problems a trade-off must be made between multiple, often conflicting objectives (Vamplew et al., 2021; Reymond et al., 2022). For example, learning a wind farm control policy is a multi-objective problem (Padulaparthi et al., 2022). Generally, a wind turbine aims to maximise its power production; however, this can lead to higher stress being induced on the turbine's components. Therefore, there is a trade-off between power production and induced damage. However, real-world problems can also have multiple agents (Rădulescu et al., 2020; Hayes et al., 2020). In this work, we focus on multi-objective multi-agent settings.

Many multi-agent settings have a sparse dependency structure between the different agents, which can be exploited. For example, a phenomenon known as the wake effect occurs when upstream turbines extract energy from the wind, which increases the turbulence intensity of the wind flow and reduces the available energy for downstream turbines (González-Longatt et al., 2012). In order to compensate for this effect, upstream turbines may aim to change their yaw angle, i.e., the direction that they are facing w.r.t. the incoming wind vector. This leads to a deflection of the wake, which reduces the power production of the misaligned turbines, but increases the energy production of the downstream turbines, which may lead to a higher energy production over the entire wind farm (van Dijk et al., 2016). Therefore, it is crucial to consider the actions taken by different agents jointly to guarantee optimality. However, deriving a joint action control policy by directly considering the actions jointly is intractable given the number of possible solutions scales exponentially with the number of agents in the system. Fortunately, the dependency structure between agents is often sparse. In the case of the wind farm control setting, we can exploit the wake effect and the resulting dependencies,

given that the decisions made by upstream turbines only affect neighboring downstream turbines (Verstraeten et al., 2020). To this end, we aim to model a wind farm as a multi-objective coordination graph (MO-CoG) (Rollón and Larrosa, 2006; Marinescu et al., 2012; Delle Fave et al., 2011b).

In the multi-objective reinforcement learning (MORL) literature a utility function is typically used to model a user’s preferences over objectives (Hayes et al., 2022c). In the real-world a user’s utility function may be unknown. In this case, a set of optimal policies must be computed (Roijers et al., 2013). In MORL multiple optimality criteria exist. Which criteria to optimise is determined by how the utility of a user is derived. If the utility of a user is derived from multiple executions of a policy (Roijers et al., 2013), the scalarised expected returns (SER) criterion must be used to learn a set of optimal policies (for example the Pareto front (Wang and Sebag, 2012)). However, if the utility of a user is derived from a single execution of a policy (Roijers et al., 2018; Hayes et al., 2021a), then the expected scalarised returns (ESR) criterion should be optimised. The set of optimal policies under the ESR criterion is known as the *ESR set* which must be computed using return distributions (Hayes et al., 2021c, 2022d). It is important that the correct optimality criteria is selected, given the policies, and set of optimal policies, computed under the SER criterion and ESR criterion can be different (Rădulescu et al., 2019; Hayes et al., 2022b,a). For example, a wind farm controller is subjected to various restrictions, such as contractual constraints or grid regulations. It is important to adhere to these at all times. Therefore, the utility should be defined on the basis of a single policy execution, and thus the ESR criterion should be optimised.

Given return distributions must be utilised to compute the ESR set, a distributional approach to MORL must be taken. Recently, a number of distributional MORL algorithms, which utilise categorical distributions, have been deployed to learn policies (Hayes et al., 2022a; Reymond et al., 2021; Hayes et al., 2021a). In real-world problems the returns are often continuous. For example, the power and turbulence values for a wind turbine are continuous. While categorical distributions can be used, they cannot accurately and efficiently represent return distributions with continuous returns. Therefore, methods which can compute return distributions when the returns are continuous must be utilised in real-world settings like wind farm control.

The state-of-the-art MO-CoG algorithms cannot compute sets of optimal policies for the ESR criterion (Hayes et al., 2022a,b), because they use expected value vectors (Rollón, 2008; Delle Fave et al., 2011a; Marinescu et al., 2012; Roijers et al., 2015). Therefore, we propose a new nonserial dynamic programming (Rosenthal, 1977) algorithm for MO-CoGs that we call *distributional multi-objective variable elimination* (DMOVE). DMOVE maintains set of return distributions to compute a set of optimal policies under the ESR criterion. Given the returns for wind farm control are continuous we utilise a generative flow model (Durkan et al., 2019; Ho et al., 2019) known as *real-NVP* (Dinh et al., 2016) to learn return distributions for each policy. We evaluate DMOVE using the state-of-the-art FLORIS wake simulator (NREL, 2021) with different realistic wind farm configurations of varying sizes and dependencies. Specifically, we first learn a model of the vector-valued local reward functions in the MO-CoG from interacting with FLORIS, and then employ DMOVE to extract the *ESR set*.

2. Background

In this section we introduce the relevant background information on multi-objective coordination graphs, dominance criteria in multi-objective reinforcement learning, and generative flow models.

2.1 Multi-Objective Coordination Graphs

In the context of multi-objective reinforcement learning (MORL) (Hayes et al., 2022c; Roijers et al., 2013), it is possible to utilise multi-objective coordination graphs (MO-CoG) (Rollón, 2008) to compute a set of optimal policies. A MO-CoG is a multi-objective extension to coordination graphs and is a tuple $(\mathcal{D}, \mathcal{A}, \mathcal{U})$. $\mathcal{D} = \{1, \dots, n\}$ is a set of n agents. \mathcal{D} is factorised into p , possibly overlapping, groups of agents \mathcal{D}^e . $\mathcal{A} = \mathcal{A}_1, \dots, \mathcal{A}_n$ is the set of joint actions. \mathcal{A}^e denotes the set of local joint actions for the group \mathcal{D}^e . $\mathcal{U} = \mathbf{u}^1, \dots, \mathbf{u}^p$ is a set of p , d -dimensional local payoff functions. The joint payoff for all agents is the sum of local payoff functions: $\mathbf{u}(\mathbf{a}) = \sum_{e=1}^p \mathbf{u}^e(\mathbf{a}^e)$. The dependencies between the local reward functions and agents can be represented by a bipartite graph with a set of nodes \mathcal{D} and a set of edges, \mathcal{E} . In this setting the nodes, \mathcal{D} , are agents and components of a factored payoff function, and an edge $(i, \mathbf{u}^e) \in \mathcal{E}$ exists if and only if agent i influences component \mathbf{u}^e . The set of all possible joint action value vectors is denoted by the set \mathcal{V} .

Multi-objective variable elimination (MOVE) extends traditional variable elimination (Koller and Friedman, 2009) to problems with multiple-objectives. MOVE methods have been used to compute sets of optimal policies for MO-CoGs. However, the present MOVE methods compute solution sets for expected (or deterministic) reward vectors (Roijers et al., 2015; Rollón and Larrosa, 2006) and do not take the underlying distributions over the rewards into account. Therefore, they are suitable for the scalarised expected returns (SER) criterion, but not for the expected scalarised returns (ESR) criterion.

2.2 Dominance Criteria in Multi-Objective Reinforcement Learning

In multi-objective reinforcement learning (MORL) (Hayes et al., 2022c), a user’s preferences over objectives are represented by a utility function. When applying a user’s utility function, the MORL literature generally focuses on two optimality criteria (Roijers et al., 2013). Calculating the expected value of the return of a policy before applying the utility function leads to the scalarised expected returns (SER) optimisation criterion:

$$V_u^\pi = u \left(\mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \mathbf{r}_t \mid \pi, \mu_0 \right] \right). \quad (1)$$

The SER criterion is the preferred optimality criteria in scenarios where the user aims to maximise their utility over multiple policies executions (Hayes et al., 2022c). The majority of MORL literature focuses on the SER criterion (Wang and Sebag, 2012; White, 1982). Applying the utility function to the returns and then calculating the expected value leads to the ESR optimisation criterion:

$$V_u^\pi = \mathbb{E} \left[u \left(\sum_{t=0}^{\infty} \gamma^t \mathbf{r}_t \right) \mid \pi, \mu_0 \right]. \quad (2)$$

In certain settings a user’s utility may be derived from a single execution of a policy, therefore the expected scalarised returns criterion should be optimised (Hayes et al., 2022d, 2021a,b; Roijers et al., 2018; Malerba and Mannion, 2021).

In many scenarios the utility function of a user is unknown a priori, and set of optimal policies must be computed (Roijers et al., 2013). To compute a set of optimal policies under the ESR criterion, a distribution over the returns must be maintained (Hayes et al., 2021c, 2022d,b). By utilising the distribution over the returns a dominance criterion known as ESR dominance can be used to determine a partial ordering over policies (Hayes et al., 2021c, 2022d). To calculate ESR dominance, the cumulative distribution function (CDF) of the each return distribution under consideration must be calculated. For a return distribution \mathbf{z}^π , the CDF of \mathbf{z}^π is denoted by $F_{\mathbf{z}^\pi}$. A return distribution \mathbf{z}^π ESR dominates a return distribution $\mathbf{z}^{\pi'}$ if the following is true:

$$\mathbf{z}^\pi >_{ESR} \mathbf{z}^{\pi'} \Leftrightarrow \forall \mathbf{v} : F_{\mathbf{z}^\pi}(\mathbf{v}) \leq F_{\mathbf{z}^{\pi'}}(\mathbf{v}) \wedge \exists \mathbf{v} : F_{\mathbf{z}^\pi}(\mathbf{v}) < F_{\mathbf{z}^{\pi'}}(\mathbf{v}). \quad (3)$$

2.3 Generative Flow Models

To model return distributions where the returns are continuous we utilise generative flow models (Kobyzev et al., 2020; Papamakarios et al., 2021) which are a class of probabilistic generative models (Tomczak, 2022) that learn non-linear models in continuous spaces through maximum likelihood. Generative flow models utilise bijective functions which enable exact and tractable density evaluation and inference on continuous data using the change of variables formula (Kingma and Dhariwal, 2018).

Consider an observed data point x generated from sampling a random variable X , a latent variable Z and a bijective function $f : X \rightarrow Z$ (and thus $f^{-1} : Z \rightarrow X$ exists). Therefore, the change of variables formula can be used to define a model distribution on the random variable X as follows:

$$p_X(x) = p_Z(f(x)) |\det(\nabla f(x))|, \quad (4)$$

where $\nabla f(x)$ is the Jacobian of f at x . To generate samples from the original sample space, X , a random sample, z , is drawn from the latent space, Z , according to p_Z , and the inverse, $x = f^{-1}(z)$ is computed. In this work we use a model known as *real-valued non-volume preserving transformations* (real-NVP) (Dinh et al., 2016) which utilises the fact the determinant of a triangular matrix can be efficiently computed as the product of its diagonal terms to learn a tractable and flexible model.

Dinh et al. (2016) stack a sequence of simple bijections, where each simple bijection is updated using only part of the input vector with a function which is simple to invert. Dinh et al. (2016) refer to each simple bijection in the sequence as an *affine coupling layer*. Each *affine coupling layer* takes a D dimensional input x and d , where $d < D$ and outputs y . The output y can be defined as follows:

$$y_{1:d} = x_{1:d} \quad (5)$$

$$y_{(d+1):D} = x_{(d+1):D} \odot \exp(s(x_{1:d})) + t(x_{1:d}), \quad (6)$$

where s is a scale function, t transformation function and are functions from $\mathcal{R}^d \rightarrow \mathcal{R}^{D-d}$ and \odot is the element-wise product. Dinh et al. (2016) show that the Jacobian of the transformation is triangular and therefore it is efficient to compute the determinant. A limitation of utilising *affine coupling layers* is during the forward transformation some components are left unchanged. Dinh et al. (2016) overcome this issue by combining *affine coupling layers* in an alternating pattern, therefore elements which are left unchanged in one transformation are modified in the next (see supplementary information (SI) for more details).

3. Multi-Objective Coordination Graphs for the Expected Scalarised Returns

The current literature on multi-objective coordination graphs (MO-CoGs) focuses exclusively on the SER criterion (Rollón and Larrosa, 2006; Roijers et al., 2015). As shown by Hayes et al. (2022b), methods that compute solution sets for the SER criterion cannot be used under the ESR criterion. Before we outline our new method to compute solution sets for the ESR criterion, we must define a set of optimal policies under the ESR criterion for MO-CoGs.

To determine a partial ordering over policies under the ESR criterion, a dominance relation known as ESR dominance can be used (Hayes et al., 2022d). To calculate ESR dominance a return distribution for each local payoff function must be calculated. For MO-CoGs we define a return distribution, \mathbf{z} , as the distribution over the returns of a local payoff function. Therefore, we define \mathcal{Z} where $\mathcal{Z} = \mathbf{z}^1, \dots, \mathbf{z}^p$ is a set of p , d -dimensional return distributions of local payoff functions. The joint payoff for all agents is the sum of local payoff return distributions: $\mathbf{z}(\mathbf{a}) = \sum_{e=1}^p \mathbf{z}^e(\mathbf{a}^e)$. The set of all possible joint action return distributions is denoted by \mathcal{V} .

By utilising ESR dominance (Equation 3) we can define a set of optimal policies under the ESR criterion for a MO-CoG as a set of ESR non-dominated global joint actions \mathbf{a} and associated return distributions of local payoff functions, $\mathbf{z}(\mathbf{a})$, known as the *ESR set*:

Definition 1 *The ESR set (ESR) of a MO-CoG, is the set of all joint actions and associated payoff return distributions that are ESR non-dominated,*

$$ESR(\mathcal{V}) = \{\mathbf{z}(\mathbf{a}) \in \mathcal{V} \mid \nexists \mathbf{z}(\mathbf{a}') \in \mathcal{V} : \mathbf{z}(\mathbf{a}') >_{ESR} \mathbf{z}(\mathbf{a})\}. \quad (7)$$

4. Distributional Multi-Objective Variable Elimination

To solve multi-objective coordination graphs (MO-CoGs) for the expected scalarised returns (ESR) criterion we define a novel *distributional multi-objective variable elimination* (DMOVE) algorithm. DMOVE utilises return distributions and ESR dominance to compute the *ESR set*.

Generally multi-objective variable elimination (MOVE) methods translate the problem to a set of value set factors (Roijers et al., 2013, 2015). Given under the ESR criterion return distributions must be utilised, the problem must first be translated to a set of return distribution set factors (RSFs), f , where each RSF f^e is a function mapping local joint actions to set of payoff return distributions. On initialisation each RSF is a singleton set containing a local actions payoff return distribution and is defined as follows:

$$f^e(\mathbf{a}^e) = \{\mathbf{z}^e(\mathbf{a}^e)\} \quad (8)$$

It is possible to describe the coordination graph as a bipartite graph whose nodes, \mathcal{D} , are both agents and components of a factored RSF, and an edge $(i, f^e) \in \mathcal{E}$ if and only if agent i influences component f^e . It is important to note an agent node is joined by an edge to a factored RSF component if the agent influences the RSF. Therefore, the dependencies

can be described by setting $\mathcal{E} = \{(i, f^e) \mid i \in \mathcal{D}^e\}$. To compute an *ESR set*, DMOVE treats a MO-CoG as a series of local sub-problems. DMOVE manipulates the set of RSFs by computing local *ESR sets* (LESRs) when eliminating agents. To calculate LESRs we first must define a set of neighbouring RSFs, f_i .

Definition 2 *The set of neighbouring RSFs f_i of agent i is the subset of all RSFs which agent i influences.*

Each agent i has a set of neighbour agents, n_i , where each agent in n_i influences one or more RSFs in f_i (see SI for a visualisation). To compute a LESR, all return distributions for the sub-problem must first be considered, \mathcal{V}_i , by calculating the following:

$$\mathcal{V}_i(f_i, \mathbf{a}^{n_i}) = \bigcup_{a^i} \bigoplus_{f^e \in f_i} f^e(\mathbf{a}^e), \quad (9)$$

where \bigoplus is the cross sum of sets of return distributions. In Equation 9 all actions in \mathbf{a}^{n_i} are fixed however, not a^i , and \mathbf{a}^e is formed from a^i and the appropriate part of \mathbf{a}^{n_i} . Once $\mathcal{V}_i(f_i, \mathbf{a}^{n_i})$ has been computed for agent i , a LESR can be calculated by applying an ESR pruning operator. Therefore, a LESR is defined as follows:

Definition 3 *A local ESR set, a LESR, is the ESR non-dominated subset of $\mathcal{V}_i(f_i, \mathbf{a}^{n_i})$:*

$$LESR_i(f_i, \mathbf{a}^{n_i}) = ESR(\mathcal{V}_i(f_i, \mathbf{a}^{n_i})), \quad (10)$$

When calculating a LESR a new RSF, f_{new} , is generated which is conditioned on the actions of the agents in n_i :

$$\forall \mathbf{a}^{n_i} f_{new}(\mathbf{a}^{n_i}) = LESR_i(f_i, \mathbf{a}^{n_i}). \quad (11)$$

The set of RSFs, f , must then be updated with f_{new} . Therefore, we remove the RSFs in f_i from f and update f with f_{new} . To do so, we define a new operator, known as the `eliminate` operator:

$$\text{eliminate}(f, i) = (f \setminus f_i) \cup \{f_{new}(\mathbf{a}^{n_i})\}. \quad (12)$$

Computing Equation 11 and Equation 12 removes agent i from the coordination graph. Therefore, the nodes and edges of the coordination graph are updated, where the edges for each agent in n_i are now connected to the new RSF, f_{new} .

Utilising the steps outlined above, we can now define the DMOVE algorithm (Algorithm 1). DMOVE first translates the problem into a set of return distribution set factors (RSFs) and removes agents in a predefined order, \mathbf{q} . DMOVE calls the algorithm `eliminate` which computes local *ESR sets* by pruning, and updates the set of RSFs. Once all agents have been eliminated the final factor from the set of RSFs, f , is retrieved, pruned and returned. The resulting set, \mathcal{S} , contains ESR non-dominated return distributions, known as the *ESR set*, and the associated joint actions. DMOVE only executes a forward pass, and calculates joint actions using a tagging scheme (Roijers et al., 2015).

Algorithm 1: DMOVE (\mathcal{U} , `prune1`, `prune2`, `prune3`, \mathbf{q})

```

1 Input:  $\mathcal{U} \leftarrow$  A set of local payoff functions;  $\mathbf{q} \leftarrow$  an elimination order
2  $f \rightarrow$  create one RSF for every local payoff function in  $\mathcal{U}$ 
3 while  $\mathbf{a}^{n_i} \in \mathcal{A}^{n_i}$  do
4   |  $i \leftarrow \mathbf{q}.\text{dequeue}()$ 
5   |  $f \leftarrow \text{eliminate}(f, i, \text{prune1}, \text{prune2})$ 
6 end
7  $f \leftarrow$  retrieve final factor from  $f$ 
8  $\mathcal{S} \leftarrow \text{prune3}(f(a_\emptyset))$ 
9 Return  $\mathcal{S}$ 
    
```

The `eliminate` algorithm (Algorithm 2) calculates the LESRs and updates the set of RSFs, f . To calculate the LESRs the function `CalculateLESR` is utilised. We define `CalculateLESRi` as follows:

$$\text{CalculateLESR}_i(f_i, \mathbf{a}^{n_i}, \text{prune1}, \text{prune2}) = \text{prune2} \left(\bigcup_{a^i} \bigoplus^*_{f^e \in f_i} f^e(\mathbf{a}^e) \right), \quad (13)$$

Algorithm 2: `eliminate($f, i, \text{prune1}, \text{prune2}$)`

```
1 Input:  $f \leftarrow$  A set of RSFs;  $i \leftarrow$  an agent
2  $n_i \leftarrow$  the set of neighbouring agents of  $i$ ;  $f_i \leftarrow$  the subset of  $f$  that  $i$  influences;  $f_{new}(a^{n_i}) \leftarrow$  a new RSF
3 for  $a^{n_i} \in \mathcal{A}^{n_i}$  do
4   |  $f_{new}(a^{n_i}) \leftarrow \text{CalculateLESR}_i(f_i, a^{n_i}, \text{prune1}, \text{prune2})$ 
5 end
6  $f \leftarrow f \setminus f_i \cup \{f_{new}\}$ 
7 Return  $f$ 
```

where \oplus^* is the prune and cross-sum operator defined by Roijers et al. (2015). The `CalculateLESRi` function prunes at two different stages; `prune1` is applied after the cross-sum has been computed and `prune2` is applied after the union over all a^i , this results in incremental pruning (Cassandra et al., 1997). Both DMOVE (Algorithm 1) and `eliminate` (Algorithm 2) are parameterised by pruning operators. To compute the *ESR set*, the pruning algorithm `ESRPrune` (see SI, Hayes et al. (2022b,a)) is used as each pruning operator.

The state-of-the-art MORL algorithms that compute sets of optimal solutions for the ESR criterion focus on utilising categorical distributions to represent return distributions (Hayes et al., 2022b,d). However, categorical distributions are discretised and therefore cannot fully represent the distribution of observations for continuous data. Therefore, we utilise a generative flow model (Tomczak, 2022) known as *real non-volume preserving transformations (real-NVP, Dinh et al. (2016), see Section 2.3)*, to model return distributions with continuous returns.

As already mentioned to compute DMOVE, the problem must first be decomposed into a set of RSFs, f . To calculate f a set of *real-NVP* models, \mathcal{X} is used, where a *real-NVP* model exists per payoff function, $x^e \in \mathcal{X}$. To compute a return distribution for each local action, \mathbf{a}^e , for each payoff function each *real-NVP* model is conditioned on the appropriate local action \mathbf{a}^e during training. The local action, \mathbf{a}^e , is then used to condition the scale, s , and transformation, t , functions for each individual *affine coupling layer* as follows:

$$y_{1:d} = x_{1:d} \quad (14)$$

$$y_{(d+1):D} = x_{(d+1):D} \odot \exp(s(x_{1:d}, \mathbf{a}^e)) + t(x_{1:d}, \mathbf{a}^e), \quad (15)$$

where both s and t are represented by a neural network. Figure 1 displays a computational graph of the forward propagation of a conditioned *real-NVP* instance for a single *affine coupling layer*. As previously highlighted in Section 2.3, during the forward transformation some components are left unchanged, therefore each *real-NVP* model must have multiple *affine coupling layers* to ensure each component is modified during the training phase. To train a *real-NVP*

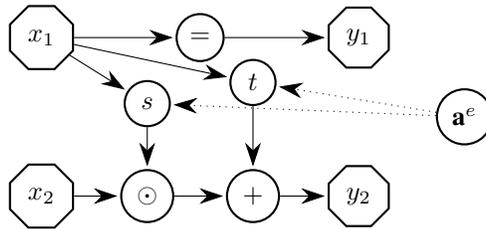


Figure 1: Computational graph for the forward propagation of a *real-NVP* instance, conditioned on a local action, \mathbf{a}^e .

instance, samples from the underlying return distributions for each payoff function must be gathered. To do so, we call an algorithm `learn` which randomly executes a global joint action, \mathbf{a} , in the environment. Once \mathbf{a} has been executed a set of vectorial returns \mathbf{r} , is received and is decomposed into local returns, \mathbf{r}^e , for each group. Each local return, \mathbf{r}^e , and joint local action, \mathbf{a}^e are then pushed to a replay buffer, where each *real-NVP* instance has its own replay buffer. Using the data stored in the replay buffer, it is possible to train each *real-NVP* at different increments, t_{inc} , using a number of random samples from the replay buffer.

Once each *real-NVP* model has been trained the set of RSFs f can be calculated. To generate the required return distributions for each local action, \mathbf{a}^e , for each payoff function, \mathbf{u}^e , \mathbf{a}^e is passed as input to the corresponding *real-NVP* instance, x^e . A random sample is taken from the latent space of x^e and the inverse is then computed which

generates an observation from the observed data space. To generate a return distribution for \mathbf{a}^e this process is repeated n times. The constructed return distribution for each local action, \mathbf{a}^e , can be used to construct the singleton set outlined in Equation 8 for each RSF in f .

The following algorithm `learn` (Algorithm 3) is defined which outlines how each *real-NVP* instance is trained and calls `DMOVE`, which returns a set of optimal return distributions and corresponding global joint actions.

Algorithm 3: `learn()`

```

1 Input:  $steps \leftarrow$  a number of random global joint actions to execute;  $\mathcal{U} \leftarrow$  a set of payoff functions
2  $\mathcal{X} \leftarrow \emptyset$ ;  $\mathcal{F} \leftarrow \emptyset$ ;  $t_{inc} \leftarrow$  increments for real-NVP training
3 for  $u \in \mathcal{U}$  do
4   |  $\mathcal{X} .append(real-NVP())$ 
5 end
6 for  $step \in steps$  do
7   |  $\mathbf{a} \leftarrow$  random global joint action
8   |  $\mathbf{r} = env.execute(\mathbf{a})$ 
9   | for  $x^e \in \mathcal{X}$  do
10  |   |  $x^e.push(\mathbf{r}^e, \mathbf{a}^e)$ 
11  |   | if  $step = t_{inc}$  then
12  |   |   |  $x^e.train()$ 
13  |   | end
14  | end
15 end
16  $\mathcal{S} \leftarrow DMOVE()$ 
    
```

5. Experiments

Generally, offshore wind farms are designed in a symmetric grid-like structure, given this approach is advantageous for planning and construction of the wind farms (Tao et al., 2020). Utilising a grid-like structure can cause wake the turbines in the farm are in close proximity (Tao et al., 2020). To construct the dependencies required to model the problem as a MO-CoG, we utilise a geometric approach to analyse the wake field for a given wind direction as proposed by Verstraeten et al. (2021). To do so, a downstream turbine, t , is considered to be dependent on a reference turbine, t_{ref} if and only if t is geographically positioned within a specified radius and is within a specified angle from the incoming wind vector. We consider that downstream turbines located at an angle of 22.5° and a radius of 1km apart from a given reference turbine lie outside the wake field of that reference turbine (Verstraeten et al., 2021). Using the identified dependencies it is possible to model the problem as a MO-CoG. We deploy `DMOVE` to compute the *ESR set* using several wind farms with different numbers of turbines and different grid-like structures.

Each wind farm has specific wind conditions. For each wind farm we investigate the incoming wind vector at 30° with a wind speed of 11 m/s (Verstraeten et al., 2021). Under the conditions presented the wake effect is strong and the dependencies for each farm created by the wake effect are shown in Figure 2. For each wind farm, each turbine is placed 500m apart along the x-axis and 400m apart on the y-axis (Verstraeten et al., 2021).

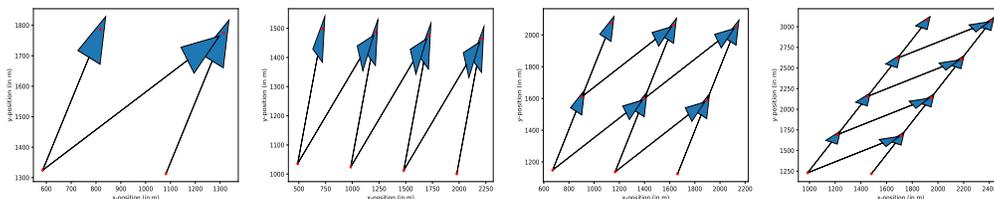


Figure 2: The dependency graphs for wind farms with 4, 8, 9 and 10 turbines.

To simulate each wind farm at the specified conditions, a state-of-the-art simulator known as FLORIS is used (NREL, 2021). A reference turbine model for the LW-8MW turbine is also used to model the turbines dynamics (Desmond et al., 2016). These models can accurately represent the scale and conditions of offshore wind farms.

We aim to use DMOVE to compute a set of optimal policies under the ESR criterion where the turbulence intensity and power for each turbine is maximised by optimising the yaw for each turbine. We choose to minimize turbulence intensity, as it is a direct contributor to fatigue loads, and thus to the overall lifetime of the turbine (Frandsen and Thøgersen, 1999). DMOVE considers each turbine as an agent. Each turbine has the following possible yaw angles, y :

$$y = [-10, -5, 0, 5, 10] \quad (16)$$

Therefore, the possible yaw angles can be considered as the action space for each agent (Verstraeten et al., 2020). For each joint action, \mathbf{a} , a reward \mathbf{r} is received where the reward can be decomposed into reward vectors for each group, \mathbf{r}^e , as follows:

$$\mathbf{r}^e = [-\text{turbulence}, \text{power}]. \quad (17)$$

Therefore, DMOVE aims to find a set of optimal return distributions over the turbulence and power returns, and the associated global yaw angles. All required parameters for DMOVE are presented in SI.

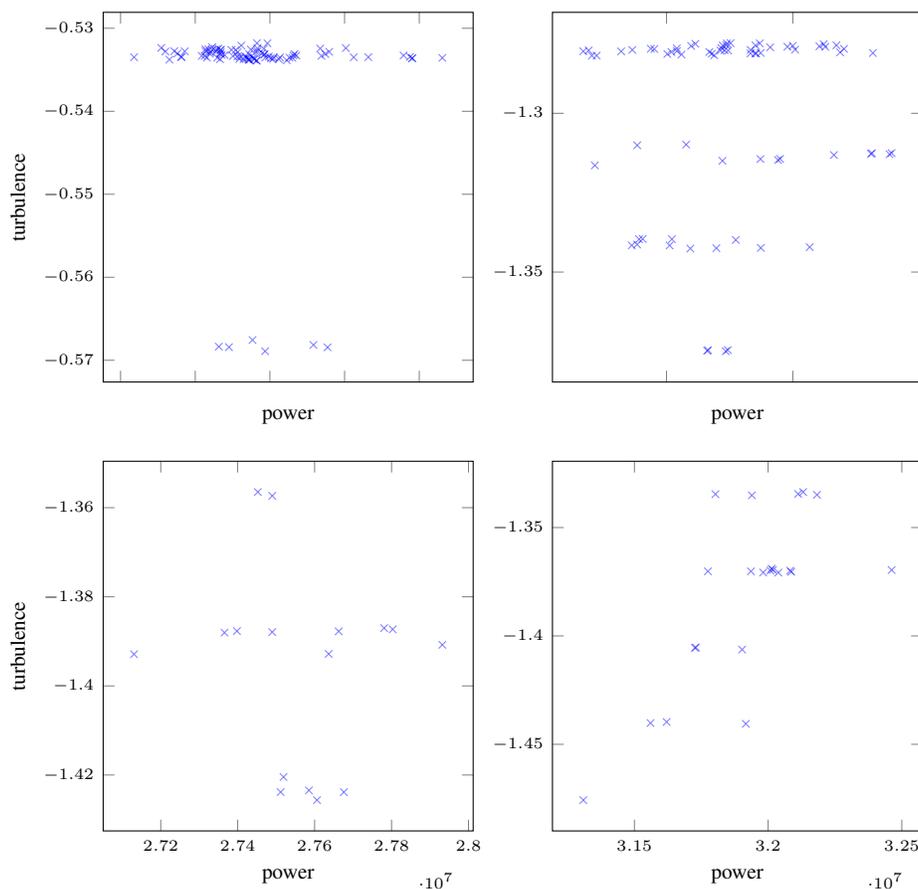


Figure 3: The expected value vector for each policy in *ESR set* for wind farms with 4, 8, 9 and 10 turbines.

Figure 3 presents the expected value vectors of ESR non-dominated return distributions of the *ESR sets* computed by DMOVE for wind farms with 4, 8, 9 and 10 turbines. Given the large number of policies in the *ESR set* it is not possible to present the distributions for each individual policy. However, Figure 4 presents the return distributions for a policy in the *ESR set* computed by DMOVE for each group in a 4 turbine wind farm. Figure 4 contrasts the learned

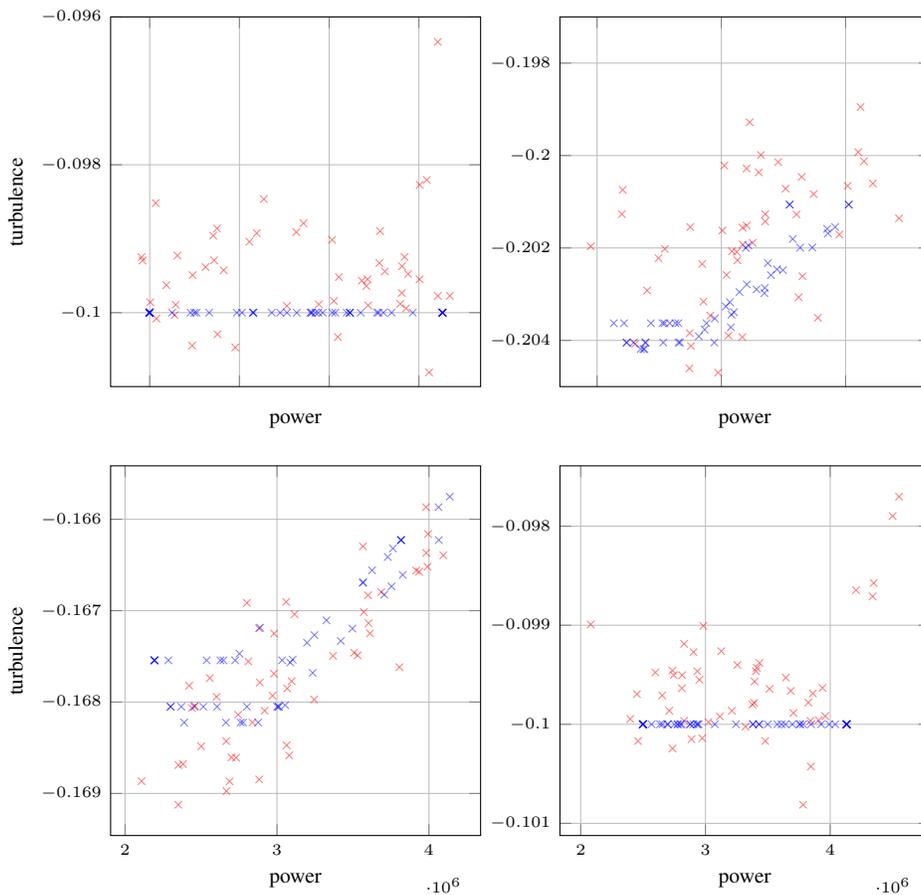


Figure 4: The learned distributions (learned by real-NVP using DMOVE, red) and the underlying distributions (blue) of power and turbulence for each group in a 4 turbine wind farm.

distributions (red) with the underlying return distributions (blue). It is clear from Figure 4 the *real-NVP* models utilised by DMOVE can accurately learn the underlying return distributions for each specific turbine.

The learned distributions presented in Figure 4 highlight the effectiveness *real-NVP* models have in learning continuous return distributions. To date only categorical distributions have been used to compute policies for the ESR criterion. As highlighted by Hayes et al. (2022b), categorical distributions have limitations when dealing with continuous data. Therefore the return distributions learned when using multi-objective categorical methods cannot accurately represent return distributions when the returns are continuous. The results presented in Figure 4 show that *real-NVP* models can be used to replace categorical distributions and learn good representations of return distributions in real-world settings with continuous returns.

Given that wind farm control is subjected to strict regulations and constraints, it is crucial that an operator has sufficient information about the potential outcomes that a policy execution may have and the likelihood of each outcome. Therefore, using the distributions presented in Figure 4 an operator can perform a more informed analysis of potential policies. Given that in wind farm control regulations and constraints need to be adhered to at any moment, and thus each policy execution is important, having the outlined distributions available can help the operator avoid selecting policies which may have some likelihood of violating these regulations or constraints.

Given the rise in deployment of artificial intelligence systems in the real-world, it is important such systems are explainable, trustworthy and safe (Mannion et al., 2021). Distributional methods, like DMOVE, can be used to provide explainability for multi-objective decision making processes. In the context of wind farm control, risk assessment is

crucial when developing operation & maintenance strategies. Having a distribution over the returns available can aid operators in understanding how and why an adverse event occurred as such an event will be present in the distribution of returns. Additionally, distributional methods can also bring outlier events to the attention of the operator before a decision is made, which can help the operator to take decisions to avoid unsafe events. Therefore, it may be possible to utilise distributional multi-objective approaches for expanding the field of trustworthy and safe artificial intelligence (Mannion et al., 2021). As already highlighted, the state-of-the-art MORL approaches focuses on computing policies using expected value vectors (Roijers et al., 2015; Bryce et al., 2007). Making decisions based on expected value vectors reduces the explainability and safety of such algorithms given the distribution over the potential outcomes is lost when the expectation is computed. For example, risks in wind farm control depend on short-term events that may occur (e.g., unexpected downtime or grid curtailment). An expected value vector may not effectively capture uncommon events given their probability of occurring may be small. Expected value vectors can also be difficult to interpret especially in environments where outcomes can be stochastic, like wind farm control. In contrast, return distributions can easily be interpreted by a system expert. Therefore, distributional methods could be utilised as a decision making aid in complex real world problem domains.

6. Related Work

Multi-objective variable elimination focuses on computing sets of optimal policies for the SER criterion (Rollón and Larrosa, 2006; Roijers et al., 2015; Marinescu et al., 2012). For example, Pareto multi-objective variable elimination (PMOVE) (Rollón and Larrosa, 2006), and convex multi-objective variable elimination (CMOVE) (Roijers et al., 2015) are used to compute the Pareto front and convex coverage set respectively. The outlined methods all utilise expected value vectors to determine a partial ordering over policies and therefore cannot be used to learn optimal policies for the ESR criterion (Hayes et al., 2022b,a).

The majority of research for the expected scalarised returns criterion largely focuses on utilising categorical distributions to compute return distributions. Hayes et al. (2022b) compute the *ESR set* using multi-variate categorical distributions in sequential decision making settings when the utility function is unknown. Reymond et al. (2021) utilise a multi-variate categorical distribution to learn a return distribution in MORL settings where the utility function of a user is known. However, in this work we utilise generative flow models, which is a class of generative modelling.

Generative modelling (Tomczak, 2022) has become a very popular area of research where models such as general-adversarial networks (Goodfellow et al., 2014) (GANs), variational auto-encoders (Kingma and Welling, 2013) (VAEs) and generative flow models (Kingma and Dhariwal, 2018; Dinh et al., 2016) have made significant advancements in artificial intelligence specifically in the area of supervised learning. Recently, generative flow models have been used in sequential decision making settings to model complex distributions. Ma et al. (2020) utilise a *real-NVP* model conditioned on a state in reinforcement learning settings to model the complex distribution of a policy.

7. Conclusion & Future Work

In this work, we present a novel *distributional multi-objective variable elimination* (DMOVE) algorithm that can compute the *ESR set* for real-world problem domains where the returns are continuous. We show that DMOVE can compute the *ESR set* in real-world wind farm settings and also highlight how the *real-NVP* models can accurately learn the continuous return distributions by interacting with the environment.

In a future publication we aim to perform a computational analysis on DMOVE to calculate the computational implications of maintaining return distributions to calculate policies. We also aim to utilise generative flow models in multi-objective sequential decision making settings, like multi-objective Markov decision processes (MOMDPs), to compute the *ESR set* in MOMDPs with continuous state-action space and continuous returns.

Acknowledgements

Conor F. Hayes is funded by the National University of Ireland Hardiman Scholarship. This research was supported by funding from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” program.

References

- Daniel Bryce, William Cushing, and Subbarao Kambhampati. Probabilistic planning is multi-objective. *Arizona State University, Tech. Rep. ASU-CSE-07-006*, 2007.
- Anthony Cassandra, Michael L. Littman, and Nevin L. Zhang. Incremental pruning: A simple, fast, exact method for partially observable markov decision processes. In *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, UAI’97, page 54–61, San Francisco, CA, USA, 1997. Morgan Kaufmann Publishers Inc. ISBN 1558604855.
- Francesco Delle Fave, Ruben Stranders, Alex Rogers, and Nicholas R Jennings. Bounded decentralised coordination over multiple objectives. In *Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS’11)*, pages 371–378, 2011a.
- Francesco M. Delle Fave, Ruben Stranders, Alex Rogers, and Nicholas R. Jennings. Bounded decentralised coordination over multiple objectives. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS ’11, page 371–378, Richland, SC, 2011b. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 0982657153.
- Cian Desmond, Jimmy Murphy, Lindert Blonk, and Wouter Haans. Description of an 8 mw reference wind turbine. In *Journal of Physics: Conference Series*, volume 753, page 092013. IOP Publishing, 2016.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/7ac71d433f282034e088473244df8c02-Paper.pdf>.
- Sten Frandsen and Morten L Thøgersen. Integrated fatigue loading for wind turbines in wind farms by combining ambient turbulence and wakes. *Wind Engineering*, pages 327–339, 1999.
- Francisco González-Longatt, P Wall, and V Terzija. Wake effect in wind farm performance: Steady-state and dynamic behavior. *Renewable Energy*, 39(1):329–338, 2012.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- Conor F. Hayes, Enda Howley, and Patrick Mannion. Dynamic thresholded lexicographic ordering. In *Adaptive and Learning Agents Workshop (at AAMAS 2020)*, May 2020.
- Conor F Hayes, Mathieu Reymond, Diederik M Roijers, Enda Howley, and Patrick Mannion. Distributional monte carlo tree search for risk-aware and multi-objective reinforcement learning. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1530–1532, 2021a.
- Conor F. Hayes, Mathieu Reymond, Diederik M. Roijers, Enda Howley, and Patrick Mannion. Risk-aware and multi-objective decision making with distributional monte carlo tree search. In: *Proceedings of the Adaptive and Learning Agents workshop at AAMAS 2021*, 2021b.

- Conor F. Hayes, Timothy Verstraeten, Diederik M. Roijers, Enda Howley, and Patrick Mannion. Dominance criteria and solution sets for the expected scalarised returns. In *Proceedings of the Adaptive and Learning Agents workshop at AAMAS 2021*, 2021c.
- Conor F Hayes, Diederik M Roijers, Enda Howley, and Patrick Mannion. Decision-theoretic planning for the expected scalarised returns. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 1621–1623, 2022a.
- Conor F. Hayes, Diederik M. Roijers, Enda Howley, and Mannion Patrick. Distributional multi-objective value iteration. In *Adaptive and Learning Agents Workshop (AAMAS 2022)*, 2022b.
- Conor F. Hayes, Roxana Rădulescu, Eugenio Bargiacchi, Johan Källström, Matthew Macfarlane, Mathieu Reymond, Timothy Verstraeten, Luisa M. Zintgraf, Richard Dazeley, Fredrik Heintz, Enda Howley, Athirai A. Irisappane, Patrick Mannion, Ann Nowé, Gabriel Ramos, Marcello Restelli, Peter Vamplew, and Diederik M. Roijers. A practical guide to multi-objective reinforcement learning and planning. *Autonomous Agents and Multi-Agent Systems*, 36(1):26, 2022c. ISSN 1573-7454. doi: 10.1007/s10458-022-09552-y. URL <https://doi.org/10.1007/s10458-022-09552-y>.
- Conor F. Hayes, Timothy Verstraeten, Diederik M. Roijers, Enda Howley, and Patrick Mannion. Expected scalarised returns dominance: A new solution concept for multi-objective decision making. *Neural Computing & Applications*, 2022d.
- Jonathan Ho, Xi Chen, Aravind Srinivas, Yan Duan, and Pieter Abbeel. Flow++: Improving flow-based generative models with variational dequantization and architecture design. In *International Conference on Machine Learning*, pages 2722–2730. PMLR, 2019.
- Diederik P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Ivan Kobyzev, Simon JD Prince, and Marcus A Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):3964–3979, 2020.
- Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- Xiaobai Ma, Jayesh K Gupta, and Mykel J Kochenderfer. Normalizing flow policies for multi-agent systems. In *International Conference on Decision and Game Theory for Security*, pages 277–296. Springer, 2020.
- Federico Malerba and Patrick Mannion. Evaluating tunable agents with non-linear utility functions under expected scalarised returns. In *Multi-Objective Decision Making Workshop (MODeM 2021)*, 2021.
- Patrick Mannion, Frederik Heinz, Thommen George Karimpanal, and Peter Vamplew. Multi-objective decision making for trustworthy ai. *Multi-Objective Decision Making Workshop (MODeM 2021)*, 2021.
- Radu Marinescu, Abdul Razak, and Nic Wilson. Multi-objective influence diagrams. *arXiv preprint arXiv:1210.4911*, 2012.
- NREL. Floris. version 2.4, 2021. URL <https://github.com/NREL/floris>.
- Venkata Ramakrishna Padullaparthi, Srinarayana Nagarathinam, Arunchandar Vasan, Vishnu Menon, and Depak Sudarsanam. Falcon-farm level control for wind turbines using multi-agent deep reinforcement learning. *Renewable Energy*, 181:445–456, 2022.
- George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021.
- Roxana Rădulescu, Patrick Mannion, Diederik M. Roijers, and Ann Nowé. Multi-objective multi-agent decision making: a utility-based analysis and survey. *Autonomous Agents and Multi-Agent Systems*, 34(10), 2020.

- Mathieu Reymond, Conor F. Hayes, Diederik M. Roijers, Denis Steckelmacher, and Ann Nowé. Actor-critic multi-objective reinforcement learning for non-linear utility functions. *Multi-Objective Decision Making Workshop (MO-DeM 2021)*, 2021.
- Mathieu Reymond, Conor F Hayes, Lander Willem, Roxana Rădulescu, Steven Abrams, Diederik M Roijers, Enda Howley, Patrick Mannion, Niel Hens, Ann Nowé, et al. Exploring the pareto front of multi-objective covid-19 mitigation policies using reinforcement learning. *arXiv preprint arXiv:2204.05027*, 2022.
- Diederik M. Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113, 2013.
- Diederik M. Roijers, Shimon Whiteson, and Frans A. Oliehoek. Computing convex coverage sets for faster multi-objective coordination. *Journal of Artificial Intelligence Research*, 52:399–443, 2015.
- Diederik M. Roijers, Denis Steckelmacher, and Ann Nowé. Multi-objective reinforcement learning for the expected utility of the return. In *Proceedings of the Adaptive and Learning Agents workshop at FAIM 2018*, 2018.
- Emma Rollón. *Multi-objective optimization in graphical models*. PhD thesis, Universitat Politècnica de Catalunya, 2008.
- Emma Rollón and Javier Larrosa. Bucket elimination for multiobjective optimization problems. *Journal of Heuristics*, 12(4):307–328, 2006.
- Arnie Rosenthal. Nonserial dynamic programming is optimal. In *Proceedings of the ninth annual ACM symposium on Theory of computing*, pages 98–105, 1977.
- Roxana Rădulescu, Patrick Mannion, Diederik Roijers, and Ann Nowé. Equilibria in multi-objective games: A utility-based perspective. In *Adaptive and Learning Agents workshop (at AAMAS 2019)*, May 2019.
- Siyu Tao, Andrés Feijóo, Jiemin Zhou, and Gang Zheng. Topology design of an offshore wind farm with multiple types of wind turbines in a circular layout. *Energies*, 13(3):556, 2020.
- Jakub M Tomczak. *Deep Generative Modeling*. Springer Nature, 2022.
- Peter Vamplew, Benjamin J Smith, Johan Kallstrom, Gabriel Ramos, Roxana Radulescu, Diederik M Roijers, Conor F Hayes, Fredrik Heintz, Patrick Mannion, Pieter JK Libin, et al. Scalar reward is not enough: A response to silver, singh, precup and sutton (2021). *arXiv preprint arXiv:2112.15422*, 2021.
- Mike T. van Dijk, Jan-Willem Wingerden, Turaj Ashuri, Yaoyu Li, and Mario A. Rotea. Yaw-misalignment and its impact on wind turbine loads and wind farm power output. *Journal of Physics: Conference Series*, 753(6), 2016.
- Timothy Verstraeten, Eugenio Bargiacchi, Pieter JK Libin, Jan Helsen, Diederik M Roijers, and Ann Nowé. Multi-agent thompson sampling for bandit applications with sparse neighbourhood structures. *Scientific reports*, 10(1): 1–13, 2020.
- Timothy Verstraeten, Pieter-Jan Daems, Eugenio Bargiacchi, Diederik M Roijers, Pieter JK Libin, and Jan Helsen. Scalable optimization for wind farm control using coordination graphs. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1362–1370, 2021.
- Weijia Wang and Michèle Sebag. Multi-objective Monte-Carlo tree search. volume 25 of *Proceedings of Machine Learning Research*, pages 507–522, Singapore Management University, Singapore, 04–06 Nov 2012. PMLR.
- DJ White. Multi-objective infinite-horizon discounted markov decision processes. *Journal of mathematical analysis and applications*, 89(2):639–647, 1982.