

What is redundant and what is not? Computational trade-offs in modelling to generate alternatives for energy infrastructure deployment

Lombardi, Francesco; Pickering, Bryn; Pfenninger, Stefan

DOI

[10.1016/j.apenergy.2023.121002](https://doi.org/10.1016/j.apenergy.2023.121002)

Publication date

2023

Document Version

Final published version

Published in

Applied Energy

Citation (APA)

Lombardi, F., Pickering, B., & Pfenninger, S. (2023). What is redundant and what is not? Computational trade-offs in modelling to generate alternatives for energy infrastructure deployment. *Applied Energy*, 339, Article 121002. <https://doi.org/10.1016/j.apenergy.2023.121002>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



What is redundant and what is not? Computational trade-offs in modelling to generate alternatives for energy infrastructure deployment

Francesco Lombardi ^{a,*}, Bryn Pickering ^b, Stefan Pfenninger ^a

^a TU Delft, Faculty of Technology, Policy and Management, Department of Engineering Systems and Services, Delft, Netherlands

^b ETH Zürich, Institute for Environmental Decisions, Department for Environmental Systems Science, Zürich, Switzerland

ARTICLE INFO

Keywords:

MGA
Spatial dissimilarity
SPORES
On-shore wind
Optimization

ABSTRACT

Given the urgent need to devise credible, deep strategies for carbon neutrality, approaches for ‘modelling to generate alternatives’ (MGA) are gaining popularity in the energy sector. Yet, MGA faces limitations when applied to state-of-the-art energy system models: the number of alternatives that can be generated is virtually infinite; no realistic computational effort can discover the complete technology and spatial option space. Here, based on our own SPORES method, a highly customisable and spatially-explicit advancement of MGA, we empirically test different search strategies – including some adapted from other MGA approaches – with the aim of identifying how to minimise redundant computation. With application to a model of the European power system, we show that, for a fixed number of generated alternatives, there is a clear trade-off in making use of the available computational power to unveil technology versus spatial dissimilarity across alternative system configurations. Moreover, we show that focussing on technology dissimilarity may fail to identify system configurations that appeal to real-world stakeholders, such as those in which capacity is more spread out at the local scale. Based on this evidence that no feasible alternative can be deemed redundant a priori, we propose to initially search for options in a way that balances spatial and technology dissimilarity; this can be achieved by combining the strengths of two different strategies. The resulting solution space can then be refined based on the feedback of stakeholders. More generally, we propose the adoption of ad-hoc MGA sensitivity analyses, targeted at testing a study’s central claims, as a computationally inexpensive standard to improve the quality of energy modelling analyses.

1. Introduction

Large-scale energy system optimisation models are increasingly used to support the urgent task of planning for the energy transition [1]. Most typically, they are used to understand how to deploy new energy infrastructure to make energy systems fully carbon-neutral at the country or continental scale while keeping the economic cost for society as low as possible [2]. However, as reported by real-world stakeholders [3], and increasingly acknowledged in the literature [4,5], the provision of a single solution that minimises total economic cost is of little use in practice, for a number of reasons. First, real-world decisions on the energy transition involve a multitude of stakeholders, including local communities, with many other concerns than the total economic cost [6]. Second, modelled costs for future systems are uncertain, for instance due to technology cost projections that need to be best-guessed when optimising for a long-term horizon [7]. It is thus problematic to concentrate on that configuration which ensures the minimum economic cost when some of the apparently more costly options might

end up being just as or even more cost-effective after cost uncertainty is realised in practice. Third and final, the generation and communication of only a single, least-cost solution can create confusion between what is least-cost and what is possible, with dangerous consequences. For instance, it is common to hear claims that a given investment decision, say installing large amounts of bioenergy power supply, is ‘required’ for a country’s energy transition because it is featured in the least-cost solution [8]. In practice, however, strategies without bioenergy may exist within the cost uncertainty range of the least-cost solution.

As a solution to the pitfalls of economic optimisation, some have proposed introducing secondary objectives, such as the minimisation of CO₂ emissions [9,10]. Yet, the alternatives obtained along a multi-objective Pareto front and its near-optimal region cannot ensure that the full range of possibilities gets captured. Other alternatives, driven by unmodelled (or even impossible-to-model) objectives, might also exist and be relevant for real-world discussion [11]. To address the impossibility to model all that might matter in reality, often referred to

* Corresponding author.

E-mail address: f.lombardi@tudelft.nl (F. Lombardi).

<https://doi.org/10.1016/j.apenergy.2023.121002>

Received 21 June 2022; Received in revised form 10 February 2023; Accepted 18 March 2023

Available online 28 March 2023

0306-2619/© 2023 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

as the structural uncertainty of the modelling process [5], approaches known as ‘modelling to generate alternatives’, or MGA, have come to the forefront in recent years [12–15]. Conceptualised by Brill in 1979 [11] and first applied to energy system models a decade ago [12], the basic idea behind MGA is: first, to compute the least-cost solution as a starting point; and second, to change the objective of the problem to search for something as different as possible from the previously-found solution, while enforcing that cost does not increase too much compared to the minimum feasible system cost. The process (also known as the ‘Hop, Skip and Jump’ algorithm) can be repeated indefinitely, each time updating the objective to search for something different from all the previously found feasible solutions.

Yet, MGA displays limitations when applied to sufficiently large models, such as state-of-the-art energy system optimisation models covering the whole of Europe at high temporal and spatial resolution and with many technological options [16,17]. The number of meaningful alternatives that can be generated for such models is *de facto* infinite, and the conventional form of applying MGA often fails to represent the range of available options well, for two reasons.

First, the ‘Hop, Skip and Jump’ algorithm does not span the solution space evenly, leaving even feasible configurations with clear differences in the technology mix unexplored. Previous empirical work showed that configurations with very low or very high shares of a particular technology in the overall capacity mix, located near the ‘corners’, or extremes, of the multi-dimensional decision space, are not found by conventional applications of MGA even when they exist — that is, when they can be eventually found if the model objective is explicitly set to find them [18,19].

Second, to keep the problem computationally tractable, the search for something ‘as different as possible’ is typically applied only to high-level variables of energy system models, such as the total installed capacity for each technology type, even within models of high spatial granularity [19]. This leads to the generation of ‘technologically-distinctive’ feasible solutions, which rely relatively more on technologies that were minimally deployed in the least-cost solution, for instance deploying more solar than wind, at the system scale. However, real-world concerns, such as the social acceptance of new infrastructure, primarily call for alternatives in the way capacity is located at the sub-national scale, even for a fixed, agreed-upon mix of technologies at the national scale [20]. For instance, different ways of spatially distributing wind capacity across sub-national regions, for a roughly fixed amount of total wind capacity to be deployed in the system. While it is true that technologically-distinctive feasible solutions are naturally likely to entail different sub-national distributions of capacity as well [21], this does not necessarily make up for having spatially-distinctive alternatives around a roughly fixed technology mix [20]. On the other hand, an explicit search for such ‘spatially-distinctive’ solutions substantially enlarges the number of potentially interesting alternatives to generate and, thereby, the overall computational burden.

In response to the two limitations above, several recent advancements of MGA with application to energy system modelling have aimed at either spanning the solution space more evenly or explicitly looking for spatially-distinctive solutions; or, in some cases, at achieving both at once. Yet, as mentioned above, the amount of alternatives that may matter for state-of-the-art energy system models of large size is virtually infinite. Even attempts at mapping ‘all’ alternatives [22] or spanning the solution space as evenly as possible [21] can ultimately provide only a finite sample of solutions based on high-level technological dissimilarity. Furthermore, providing a finite set of options leaves open the question of how many alternatives are enough, i.e. which alternatives are redundant and which ones are not.

Acknowledging this, we set out to empirically investigate the computational trade-offs among different possible approaches to the generation of alternatives within high-resolution energy system models. In particular, we test four different ways of generating alternatives

within the SPORES method, an original development of MGA that we have presented in previous work [20]. The use of SPORES does not cover all possible formulations of MGA in the literature; but the high customisability of the SPORES workflow lends itself to tweaking the search towards either spatial or technological dissimilarity explicitly, allowing us to explore the trade-offs between the two. Some of the tweaks of the search strategy that we consider do originate from other recent MGA advancements but are applied here in a spatially-explicit way even when they were originally conceived for application to high-level variables only (see Section 2). With application to a model of the European power system with 97 nodes and 3 h temporal resolution, we show that deciding what is redundant and what is not is far from trivial, and that there is a clear trade-off in the computational efficiency of MGA search strategies in highlighting technological versus spatial dissimilarity of feasible solutions. In particular, we show that focussing on technology dissimilarity may fail to identify system configurations that appeal to real-world stakeholders, such as those in which capacity is more spread out at the local scale. Based on our empirical findings, we propose possible solutions for energy modellers to make their MGA approaches as computationally efficient and practically relevant as possible. We openly release the model and code used for the generation of results on Zenodo [23], to foster transparency, repeatability and further developments [24].

The remainder of the paper is organised as follows. Section 2 presents the latest version of the SPORES workflow, how we tweak its parameters to approximate different search strategies, and the European power system model to which we apply our experiments. Section 3 shows and discusses our findings in terms of the overall discovered decision space for each search strategy, the amount of generated alternatives that are likely relevant for real-world decisions, and differences in the resulting flexibility for spatially moving capacity. We conclude (Section 4) by discussing the implications of our findings for other modellers and policy decision support.

2. Methods

The following subsections present the latest version of the SPORES workflow (Section 2.1), the different search strategies that we use within the workflow for the analysis of computational trade-offs (2.2), and the specific version of the Euro–Calliope model that we use as a case study (2.3).

2.1. SPORES workflow

The SPORES workflow, summarised in Fig. 1 differs from conventional MGA in two key aspects.

First, it searches explicitly for spatially-distinctive options. Conventional MGA searches for a new solution as different as possible to the optimal one by assigning weights (penalties) to aggregate capacity variables proportional to their capacity deployment in the economically optimal solution. SPORES, instead, assigns such weights to spatially-explicit capacity variables. For instance, if wind generation is deployed in the cost-optimal solution, SPORES assigns different penalties to wind generation capacity in each location rather than penalising wind generation overall. This means that the search for something different might result in a configuration that has as much wind generation capacity overall, but distributes it differently in space.

Second, it does so from multiple directions within the feasible, near-optimal solution space. This parallel search from multiple directions arises by anchoring the MGA algorithm to different extremes of the feasible near-optimal space, instead of just using the least-cost solution as the starting point of the search. These additional extremes are identified by explicitly minimising or maximising the system-wide deployment of specific technology-capacity decision variables, a strategy also employed in other recent work [21]. We detail the mathematical steps required to generate alternatives based on such a workflow in the following subsections.

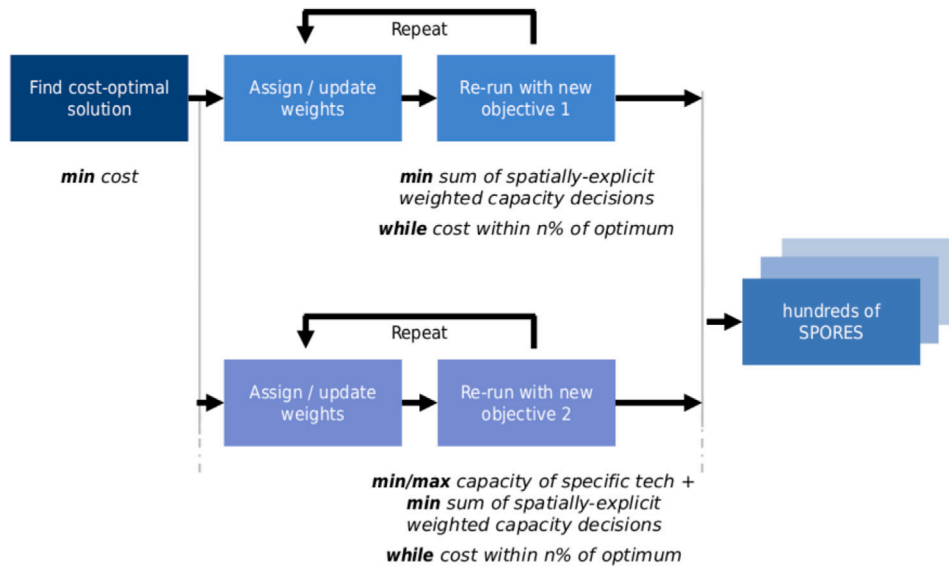


Fig. 1. Schematic representation of the SPORES workflow. The feasible, near-optimal solution space is systematically explored through parallel runs. A base run (top of the figure) searches for alternatives adopting a spatially-explicit version of MGA. Weights are assigned not only to system-wide technology-capacity decision variables but rather to each technology-in-location capacity decision variable. In parallel, several other runs (bottom of the figure) explore the decision space from different directions. These runs add to the search for spatially-explicit alternatives a secondary, technology-explicit objective: the system-wide minimisation (and, alternatively, maximisation) of a specific technology. This ensures that configurations with an extremely high or low deployment of specific technologies do not get lost while also ensuring that several spatially-distinctive configurations of technology deployment are provided around those.

2.1.1. Identification of the cost-optimal solution

We identify the cost-optimal solution, or economic optimum of the problem, by minimising the total annualised system cost, as per Eq. (1).

$$\min : cost = \sum_j \sum_i (c_{fix,ij} x_{ij}^{cap} + \sum_t c_{var,ij} x_{t,ij}^{prod}) \quad (1)$$

s.t. $\mathbf{Ax} \leq \mathbf{b}$
 $\mathbf{x} \geq 0$,

where i and j indicate the i th technology type and the j th location of the model; x_{ij}^{cap} is the decision variable pertaining to the installed capacity of the ij th location–technology pair; $x_{t,ij}^{prod}$ is the decision variable related to the power production of the ij th location–technology pair as a function of time; $c_{fix,ij}$, $c_{var,ij}$ are, respectively, the annualised fixed and variable costs per each location–technology pair; \mathbf{A} , \mathbf{b} , are a matrix and a vector of coefficients that build all the physical constraints in combination with the vector \mathbf{x} of all decision variables.

2.1.2. Assignment of weights

Having identified the mathematical optimum, we assign a strictly positive weight (w_{ij}^n) to the generation capacity decision variables (location–technology pairs) that are non-zero. This weight can be assigned using different approaches, which result in different search strategies. In the first published application of SPORES [20], we assigned weights based on the relative deployment of a technology in a given location compared to the maximum potential for deployment at that location ($x_{ij,max}^{cap}$). The newly-found weight is then summed to the weight obtained in the preceding iteration (w_{ij}^{n-1}), for any iteration other than the first (Eq. (2)). Here, we also test three additional weight assignment methods (see Section 2.2).

$$w_{ij}^n = w_{ij}^{n-1} + \frac{x_{ij}^{cap,n}}{x_{ij,max}^{cap}} \quad (2)$$

2.1.3. Main batch of SPORES

We hence obtain a SPORE by minimising the sum of location-specific weighted capacity decision variables. At the same time, we

constrain the total annualised system cost ($cost_n$) to remain in a neighbourhood of the optimal cost ($cost_0$), as per Eq. (3).

$$\min Y = \sum_j \sum_i w_{ij} x_{ij}^{cap}$$

s.t. $cost_n \leq (1 + s) \cdot cost_0$
 $\mathbf{Ax} \leq \mathbf{b}$
 $\mathbf{x} \geq 0$, (3)

where s is the accepted cost relaxation (also known as cost slack). Such a new, MGA-like objective function formally makes the problem an ϵ -constrained multi-objective optimisation: the minimisation of already-deployed location–technology pairs is the explicit objective and cost is the implicit one. This objective function is similar to the one first applied by DeCarolis [12] to energy system models and also used in more recent work [18], except for the variables and weights being here spatially-explicit.

2.1.4. Parallel batches of SPORES with secondary objectives

The workflow steps outlined in Sections 2.1.1 to 2.1.3 can be as well seen as a spatially-explicit version of the ‘Hop, Skip and Jump’ algorithm. Such an algorithm has been repeatedly shown to struggle with pushing the search for alternatives up to the extreme corners of the multi-dimensional decision space [18,19]. The incorporation of the spatial dimension further exacerbates the problem, as it multiplies the variables at stake and the different feasible system configurations. A spatially-explicit version of the algorithm will focus explicitly on spatial diversity and will further struggle to discover alternatives in which the mix of deployed technologies is radically different. For such reasons, SPORES foresee a further systematic exploration of the decision space from alternative directions in which technology diversity is handled explicitly. They do so by adding to the objective function a second explicit objective: the minimisation (and, alternatively, maximisation) of the capacity of a specific technology.

$$\min (or \max) Y_{2,i} = a \cdot \sum_j x_{ij}^{cap} + b \cdot \sum_j \sum_i w_{ij} x_{ij}^{cap}$$

s.t. $cost_n \leq (1 + s) \cdot cost_0$
 $\mathbf{Ax} \leq \mathbf{b}$
 $\mathbf{x} \geq 0$, (4)

where x_{ij} is the capacity decision variable associated with the technology under minimisation (or maximisation), and a and b are the

weights associated with the different components of the objective function (which can be customised at need, as discussed in 2.2). Other recent work [19,21] has used the minimisation or maximisation of specific technologies as a way to explore the near-optimal solution space, but never in combination with the simultaneous generation of spatially-explicit alternatives, which is a unique feature of SPORES. As anticipated above, the rationale behind the combination of technology-explicit and spatially-explicit objectives is that of maintaining a focus on the discovery of spatially-distinctive options while also making sure that the option space is explored from all the relevant search directions. In other words, ensuring that no technology-distinctive option is left unexplored. If needed, the customisability of the coefficients a and b leaves open the possibility for the modeller to collapse the search into only one of the two objectives, thereby replicating, for instance, the search strategy proposed by Neumann and Brown [19].

This configures the problem as a linearised multi-objective optimisation problem, with two explicit objectives parametrised by the coefficients a and b that add up to the implicit ϵ -constrained cost objective. However, despite the mathematical formulation being multi-objective, the two explicit objectives do not aim to reflect any plausible real-world decision factors, which differentiates the approach from typical applications of multi-objective optimisation [9]. The primary objective is, in fact, just the search for something different from previous iterations, whilst the secondary objective is only needed here as a technical means to operate the spatially-explicit MGA search from alternate extremes of the decision space. Finally, we do not aim to (and do not) find a Pareto-front of optimal solutions by systematically varying the a and b weights of the different components of the objective function, as typically done in multi-objective optimisation. This is for two reasons. First, the resulting Pareto-optimal solutions would not have any real-world meaning since our fictitious objectives and weights do not have any, either. And second, we explicitly want to look at mathematically sub-optimal solutions beyond a fictitious ‘Pareto front’ as long as they are within the defined cost slack.

2.2. Customised search strategies

The SPORES workflow presented throughout Sections 2.1.1 to 2.1.4 lends itself to customisation. Both the weight-assignment method (Eq. (2)) and the relative strength of the two explicit objectives in the parallel batches of SPORES (Eq. (4)) can be modified at need. Different weight-assignment methods may push the search relatively more on spatial versus technology dissimilarity from previous iterations [20]. Similarly, changing the relative strength of the technology-explicit secondary objective in Eq. (4) may allow the search to depart more or less from the identified extreme of the decision space. In this work, we test four different weight-assignment methods and two levels of relative strength for the technology-explicit secondary objective. In both cases, the aim is to test the outcomes, in terms of redundancy of generated solutions, of pushing the search more on technology versus spatial dissimilarity, or vice versa.

2.2.1. Alternative weight-assignment methods

First, we consider the weight-assignment method outlined in Eq. (2), which we hereafter refer to as the *relative-deployment* method. We designed this method with a focus on spatial dissimilarity in previous work [20].

Second, we test the so-called *integer* method originally proposed by Brill [11] and later applied to energy system modelling by DeCarolis [12]. Reported in Eq. (5), it is the simplest method, but has the potential drawback of many variables ending up with the same weight, hampering the search efficiency. We apply it here in a spatially-explicit way, assigning weights to location–technology pairs rather than to system-wide technology variables only.

$$w_{ij}^n = w_{ij}^{n-1} + k_{ij}, \quad \text{with } k_{ij} = \begin{cases} 100, & \text{if } x_{ij} > c \\ 0, & \text{if } x_{ij} \leq c \end{cases} \quad (5)$$

where c is a constant threshold defined to avoid that even very marginal deployments of capacity may receive a weight, which would otherwise entail the risk that almost all location–technology pairs receive a non-zero weight. It is also worth noting that the *integer* weight amounts here to 100 based on the internal unit scaling of our model; a different scaling, say 1 or 10, may make more sense for a model with different units.

Third, we consider a *random* method, in which weights have no rationale and are indeed assigned as random integer numbers (Eq. (6)). This method approximates the random MGA search proposed by other authors [15,25]. Unlike such previous work, though, we do not consider further degrees of randomisation in the objective function, for the sake of consistency with the other analysed methods, and we apply the random weights to spatially-explicit decision variables.

$$w_{ij}^n = w_{ij}^{n-1} + r_{ij}, \quad \text{with } r_{ij} = U(0, 100) \quad (6)$$

where $U(0, 100)$ is a random uniform distribution.

Fourth and final, we propose an original *evolving-average* method (Eq. (7)). The idea of this method is to retain a more explicit memory of past iterations, compared to just having incremental weights. In such a way, one can assign a weight to each location–technology pair based on the distance from the *average* capacity deployed for that pair across all previously found feasible solutions ($\bar{x}_{ij}^{cap,n-1}$), which is kept up to date — in other words, it *evolves*.

$$\left\{ \begin{aligned} w_{ij}^n &= \left| \frac{\bar{x}_{ij}^{cap,n-1} - x_{ij}^{cap,n}}{\bar{x}_{ij}^{cap,n-1}} \right| \\ \bar{x}_{ij}^{cap,n-1} &= \frac{\sum_{n=1}^{n-1} x_{ij}^{cap,n}}{n-1} \end{aligned} \right. \quad (7)$$

2.2.2. Relative strength of spatial- and technology-explicit objectives

The parallel batches of SPORES arising from technology-explicit extremes of the decision spaces can be customised by tweaking the parameters a and b of Eq. (4). By default, both parameters are set to a unitary value, which means they have the same relative strength. Here, we test the outcome of reducing the relative strength (a) to 0.1 of the technology-explicit secondary objective. We want to avoid obtaining, for a batch of SPORES generated around the minimisation of a certain technology, only configurations in which such a technology is fully minimised. A reduction of the relative strength of the technology lever may allow us to obtain configurations in which the given technology is only partially minimised, further improving the technology dissimilarity of the generated alternatives. As for the case of the other customised parameters (see Section 2.1.2), the absolute numbers we adopt here for the a and b parameters make sense within the unit scaling of our model. Other absolute values might be more appropriate for a different unit scaling. However, since our analysis focuses on modifications of the parameters relative to one another, the outcomes remain generally valid.

2.3. Energy system model and MGA setup

We apply our customisations of the SPORES method to a power system model of Europe comprising 34 countries and 97 nodes, or locations, across those. The model is based on the well-established open-source energy system modelling framework Calliope [26] and takes the name of Euro–Calliope. Originally conceived by Tröndle et al. [27], the version we use here is based on a previous work of ours [16] in which we updated the initial (brownfield) electricity grid topology to mirror the most relevant real-world transmission constraints as identified by the e-Highway 2050 project. Besides various hydroelectric technologies whose capacity is assumed to remain constant, the model features seven main technologies for capacity expansion: roof-mounted and open-field solar photovoltaic, on-shore and off-shore

wind, bioenergy power plants, and battery and hydrogen-to-power storage. Each of these is used as a basis for two parallel batches of SPORES (Eq. (4)), one in which the deployment of the technology is maximised and one in which it is minimised, at the system level. Transmission technologies are also allowed to expand, and add up to the parallel batches of SPORES. We generate 10 alternatives for each parallel batch, for a total of 160 alternatives. These add up to the 50 alternatives generated within the main batch, leading to an overall 210 SPORES. For all SPORES, we adopt a slack cost of 10%, in line with previous work [16,19,20]. We refer the readers to the Supplementary Methods for further details about the model and how it has been customised for this work.

Although the Euro-Calliope model version that we use in this work allows for the modelling of all energy sectors, we subset the analysis to the power sector alone, without considering the additional electricity demand arising from the likely electrification of sectors such as building heat and transport. In fact, the aim of the present work is to focus on the computational trade-offs between different approaches to the generation of alternative near-optimal solutions. Subsetting the analysis to the power sector alone allows us to test many approaches while keeping the computational effort feasible. It also allows readers to compare our findings with those of other recent MGA studies applied to energy system models of Europe [19,21,22,28], which are all grounded in a similar model setup. What is more, the analysis of the available decision space for the full decarbonisation of Europe including all energy-consuming sectors has already been undertaken using Euro-Calliope [16] and would therefore be redundant to repeat here.

3. Results

We present our results for the reference case in which the relative strength of the two SPORES objectives is even (see Section 2.2.2), whilst we provide Supplementary Results for the case of a reduced strength of the secondary, technology-explicit objective. We start by analysing the ‘shape’ of the overall discovered decision space across the four search strategies, discussing which decision space is richer in technologically-distinctive options. Hence, we assess how many of the discovered alternatives, in each case, would match a plausible stakeholder interest, such as the limited concentration of onshore wind farms, discussing the efficiency of discovering many spatially-distinctive options across search strategies. Finally, we look for potentially real-world relevant spatial features of feasible configurations that may not be common to all search strategies and discuss how this may impact the practical usefulness of a search strategy.

3.1. Overall discovered decision space

The four considered search strategies showcase substantial differences in the overall discovered decision space, as shown in Fig. 2. These differences can be even more marked if considering only those alternatives generated within the main batch of SPORES (see Fig. S1), but are mitigated when parallel batches with technology-explicit secondary objectives are included, as per the default SPORES workflow presented in Section 2.1. This evidence reinforces the importance of systematically exploring the decision space from multiple directions, in line with recent work [20,21].

Regarding trade-offs between spatial versus technology dissimilarity, the *integer* and *relative-deployment* strategies tend to produce less-sparse solutions, with many alternatives almost overlapping in the tri-dimensional space that considers total renewables, transmission and storage capacity deployed. Conversely, the *random* and the *evolving-average* strategies seem to produce more sparsity and push the alternatives away from each other more, particularly along the transmission-expansion axis, despite both still ending up with a non-negligible number of overlapping solutions. This means that the *random* and

evolving average strategies are more efficient at discovering markedly distinctive options from a technology perspective, whilst the *relative-deployment* and *integer* methods focus more on dissimilarity of spatial deployment around fewer distinctive overall technology mixes. In fact, the alternatives overlapping in the tri-dimensional space outlined in Fig. 2 must not be mistaken for identical, and hence redundant, solutions. On the contrary, they are solutions that, albeit similar in terms of the total deployed capacity of the different technology options, are likely radically different in terms of their spatial configuration of technology deployment. Fig. S3, which expands the cross-search-strategy comparison by looking at the deployment of further disaggregated technologies, confirms the same trend: the *relative-deployment* and *integer* search strategies showcase a higher degree of overlapping solutions, which is particularly apparent for wind generation and transmission capacities, but generally valid for all technologies.

Differences across methods do not change qualitatively for the case in which the secondary objectives aimed at minimising or maximising specific technologies are assigned a relatively weaker weight in the objective function (Fig. S2, complemented by Fig. S4 for technology-disaggregated results). However, as expected, the number of alternatives located at the extremes of the decision space – i.e., those focussing on generating different spatial configurations around an explicit high-level technology feature, say the minimal deployment of bioenergy – decreases slightly overall. To further investigate the trade-off between technology and spatial dissimilarity, we move on to looking at metrics related to spatial aspects and real-world concerns that play out on a more local scale.

3.2. Alternatives that match plausible stakeholder interests

One of the most common real-world concerns when deploying new infrastructure for the energy transition is the social acceptability of on-shore wind farms [6]. It is thus helpful to analyse how many of the generated feasible energy system configurations ensure a reduction in the maximum concentration of on-shore wind farms in any single region relative to the total deployed on-shore wind capacity. Such a metric is of the possible proxies for the richness of spatial deployment options, allowing us to investigate further whether those search strategies that performed less well in terms of technology dissimilarity actually provide richer insights regarding the decision flexibility for moving capacity spatially. If a search strategy produces relatively more alternatives with a high concentration of wind farms, that means it focuses primarily on reducing wind capacity overall or on moving highly-concentrated on-shore wind hubs elsewhere without spreading capacity out. Fig. 3a shows the distribution of the maximum on-shore wind concentration across all the SPORES generated by each search strategy. As hinted at by the results in the previous subsection, search strategies that are less efficient at providing technologically-distinctive solutions (such as the *integer* and *relative-deployment* ones) are those with lower median values for on-shore wind farm concentration in single regions. In other words, they generate many solutions in which wind capacity is sited differently across sub-national regions, eventually leading also to solutions in which capacity is more spread out, which is likely of particular importance when providing practical alternatives to real-world acceptability concerns.

Instead, the search strategies with the wider distribution of values for on-shore wind capacity concentration are the *relative-deployment* and the *evolving-average* ones. However, the median value for the *evolving-average* distribution lies much higher than that of the *relative-deployment* one. This suggests that the *evolving-average* method is highly efficient at spanning the full spectrum of wind concentration options, but – given the finite number of alternatives – does so at the expense of generating further spatial dissimilarity around the found configurations. In other words, there is a trade-off between using the limited computational power and time to efficiently span all the most radically different options from a high-level perspective versus generating many

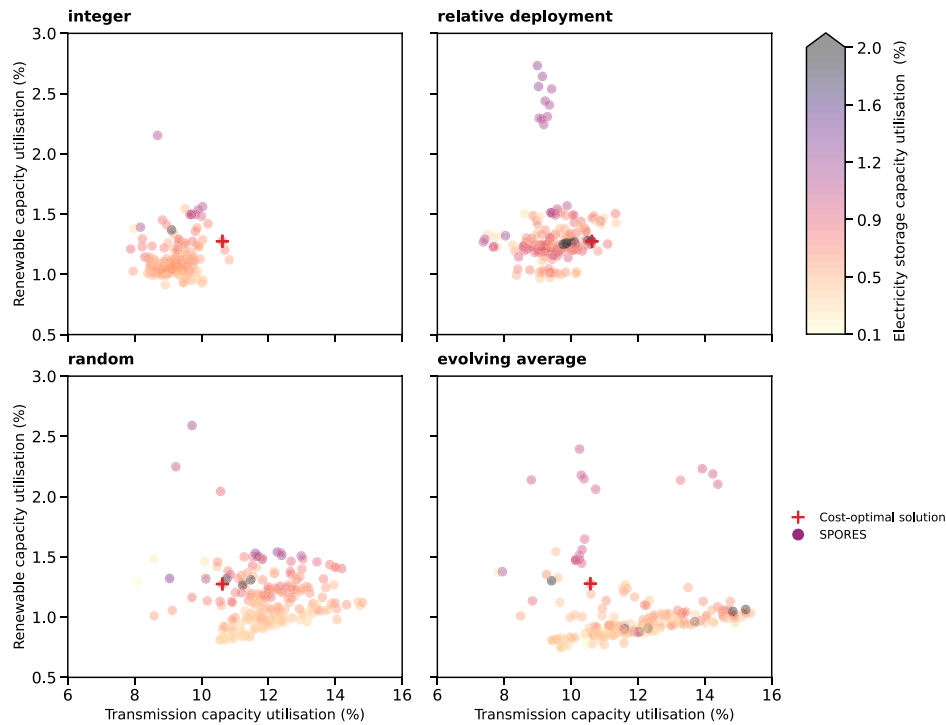


Fig. 2. Generated alternatives in a tri-dimensional space defined by aggregate renewable (Y-axis), transmission (X-axis) and storage (colour coding) capacity deployment compared to maximum deployment potential. Results are provided for all four considered search strategies. The cost-optimal solution is identified in each panel by a red marker.

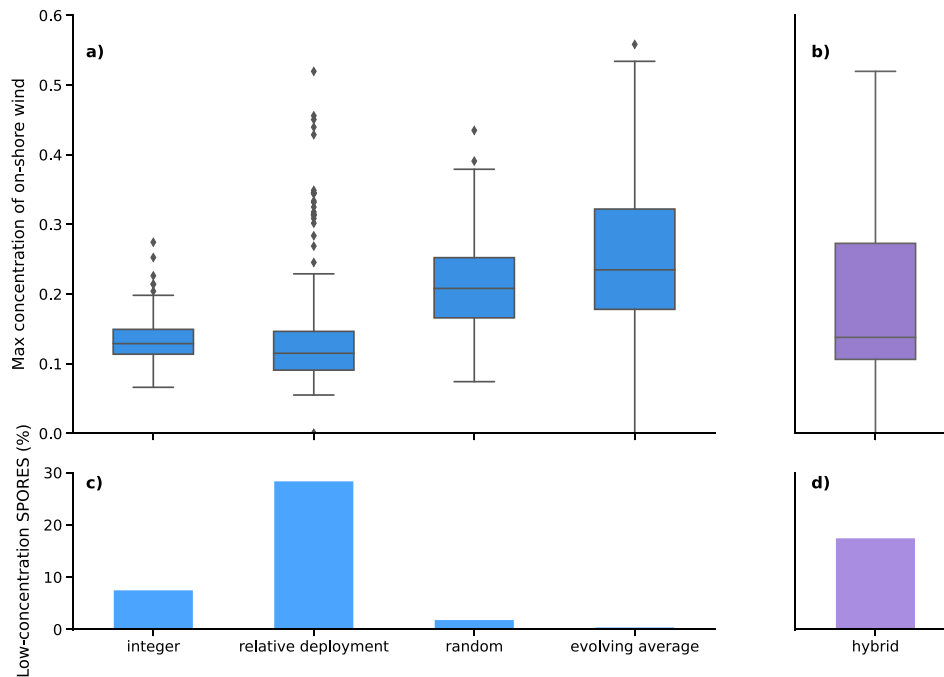


Fig. 3. The top of the Figure shows the data distribution (boxplots) for the maximum concentration of deployed on-shore wind capacity in a single region (across all 97 model locations) compared to total deployed on-shore wind capacity. The boxplots are shown for the four tested weight-assignment methods (panel a) and for a hybrid case that combines ex-post the data from the *relative-deployment* and *evolving-average* methods (panel b); see Section 3.2 for further details about this hybrid case. The bottom of the Figure shows the fraction of generated alternatives with a concentration of on-shore wind capacity at least 20% lower than in the cost-optimal case. Results are again shown for the four tested weight-assignment methods (panel c) and for the hybrid case (d).

feasible spatial configurations in the ‘technology-neighbourhood’ of fewer feasible system configurations.

Another way to quantify this trade-off is to count, for each search strategy, how many of the generated alternatives match the plausible stakeholder interest of limiting on-shore wind capacity. For instance, filtering out alternatives in which on-shore wind capacity

is at least 20% less concentrated than in the least-cost case, as done in previous work [20]. The outcomes of such filtering (Fig. 3c) confirm that the *integer* and, primarily, *relative-deployment* search strategies outperform the others from a spatial-dissimilarity perspective. This is in line with our expectations since the *relative-deployment* weight-assignment method was designed precisely to focus more on spatial

dissimilarity (see Section 2.1.2). The results of the *random* and, particularly, *evolving-average* approaches further strengthen the finding that the more a method is efficient at spanning high-level system configurations, the less it is efficient at producing spatially-distinctive, non-concentrated system configurations.

Based on this apparent trade-off, combining the benefits and limitations of different methods into a *hybrid* search may be helpful. For instance, the main batch of SPORES in which the focus is on spatial dissimilarity may be generated via the *evolving-average* method, which will make sure to span an even range of technology options. The parallel batches, in which technology dissimilarity is handled explicitly, may instead use a *relative-deployment* method to ensure that configurations are also spatially distinctive. The outcomes of such a *hybrid* method are shown in Fig. 3b, d, resulting in a good compromise in terms of stakeholder-appealing, low-wind-concentration solutions.

The results do not change substantially for a weaker anchoring of the search to the extremes of the decision space, except for the *relative-deployment* search strategy that experiences a reduction in the total amount of alternatives with a low concentration of wind farms (Fig. S5). This peculiarity is coherent with our finding so far that the *relative-deployment* method is the one that most efficiently explores the flexibility for moving capacity spatially. The *relative-deployment* search strategy takes advantage of his high spatial focus to generate many spatially-distinctive alternatives around extreme technology features of the decision space within parallel batches of SPORES (see Section 2.1.4). When the anchoring to extreme technology features becomes weaker, the search departs too quickly from the given extreme technology feature for the *relative-deployment* method to eventually generate low-concentration alternatives around it. Other weight-assignment methods are less affected by the same phenomenon because they more naturally tend to depart from the extreme technology feature regardless of the strength of the anchoring. Such a worsening of the performance affects, in turn, also the *hybrid* method. Fine-tuning the relative strength of the two objectives in Eq. (4) to the chosen search strategy appears thus essential to maximise the performance.

3.3. Unique spatial features

We have discussed the concentration of on-shore wind capacity as an example of a plausible criterion for stakeholder discussion and observed a difference across search strategies in the number of alternatives which limit it. This difference can have very concrete implications. In line with previous work [16,20], let us assume that stakeholders may be interested in filtering out the decision space based on multiple criteria at once. For instance, minimising the deployment of bioenergy power plants while simultaneously reducing on-shore wind capacity concentration (with the same threshold as defined in Section 3.2). In this case, the decision spaces discovered by the four search strategies would provide very different outcomes.

The feasible solutions found by the *integer* and *evolving-average* strategies do not feature any option that satisfies the above criteria at once. Minimising bioenergy appears to be only possible if accepting large, concentrated on-shore wind power hubs in key regions. Conversely, both the *relative-deployment* and *random* search strategies include system configurations that satisfy the chosen criteria. Fig. 4 shows, in particular, the feasible system configurations with the lowest concentration of on-shore wind farms that simultaneously allow the reliance on bioenergy to be completely avoided for both the *random* and the *relative-deployment* solution spaces. Albeit both homogeneous in their wind capacity deployment, the configuration found via *random* search (Fig. 4a, c, e, g) has about three times higher on-shore wind deployment overall. This results in several regions still becoming large wind power hubs, although few to no regions carry an unfair share of total capacity deployment. Instead, the solution discovered by the *relative-deployment* method (Fig. 4b, d, f, h) is substantially more balanced in terms of technology mix, with off-shore wind and

solar capacities being deployed alongside a spatially-homogeneous deployment of on-shore wind capacity. Substantial hydrogen storage and further grid reinforcements towards the Iberian peninsula support the balancing of this configuration. In other words, in this arbitrary yet plausible example, a search strategy that seemed not to be particularly efficient at spanning the solution space if considered from a high-level technology perspective, ends up being the only one capable of identifying a solution that meets certain stakeholder preferences when looked at from the perspective of spatial infrastructure deployment.

The situation changes when there is less of a push towards searching the extremes of the decision space. In this case, no weight-assignment method can discover a configuration that simultaneously avoids bioenergy and limits the concentration of on-shore wind capacity. This is consistent with our expectations. As anticipated in Section 2.2.2 and also observed in Section 3.2, the resulting decision space features fewer solutions in which a given technology – say, bioenergy – is fully minimised, and more in which its deployment is only partially reduced. Accordingly, there are no more of those apparently redundant alternatives in which a key part of the technology mix remains the same (i.e., no bioenergy) while capacity is deployed differently from a spatial perspective.

As discussed above, we have selected the example of a desire to simultaneously limit the concentration of wind generation capacity and the reliance on bioenergy based on previous work that identified those as plausible stakeholder interests. However, this is just one illustrative example out of many possible ones, motivated by previous work that identified these criteria as particularly appealing to real-world stakeholders [6,20]. A similar case could be made for other plausible combinations of stakeholder interests, say, the degree and topology of expansion of transmission lines and the type and amount of storage technologies deployed.

4. Discussion and conclusion

Our empirical comparison of different mathematical strategies for the generation of near-optimal alternatives within large-scale, high-resolution power system models aimed at investigating how to make the best use of limited computational power. We wanted to assess which search strategies may generate redundant alternatives and which ones may not. Our findings allow us to conclude the following.

First, the boundaries of the overall decision space do not change much with the selected search strategy, provided that the extremes of such a decision space are systematically explored. This corroborates the findings of other recent studies [16,21].

Second, for an arbitrarily fixed number of generated alternatives (i.e., for a given accepted computational effort), there is a clear trade-off between making use of the computational power to discover alternatives in terms of the high-level technology mix; and using it to generate different spatial configurations of technology deployment in the technology-neighbourhood of fewer high-level technology mixes. The research question should hence guide the choice of the search strategy. If the research focus requires an as-homogeneous-as-possible exploration of the high-level technology mixes, the MGA search can be targeted to such an aim. For instance, relying on the *evolving-average* search strategy proposed in this work or using a brute-force optimisation that targets technology mixes [21]. Yet, such a high-level technology focus is likely – more so, the more limited the computational power available – to miss out on most of the stakeholder appealing, spatially-distinctive system configurations around the found technology mixes. Accordingly, whenever research aims at providing alternatives to support real-world decisions, spatial dissimilarity should be at the core of the search strategy. This corroborates our observation in Section 1, based on the literature, that no single method can capture everything, simply because the set of potentially relevant options is infinite.

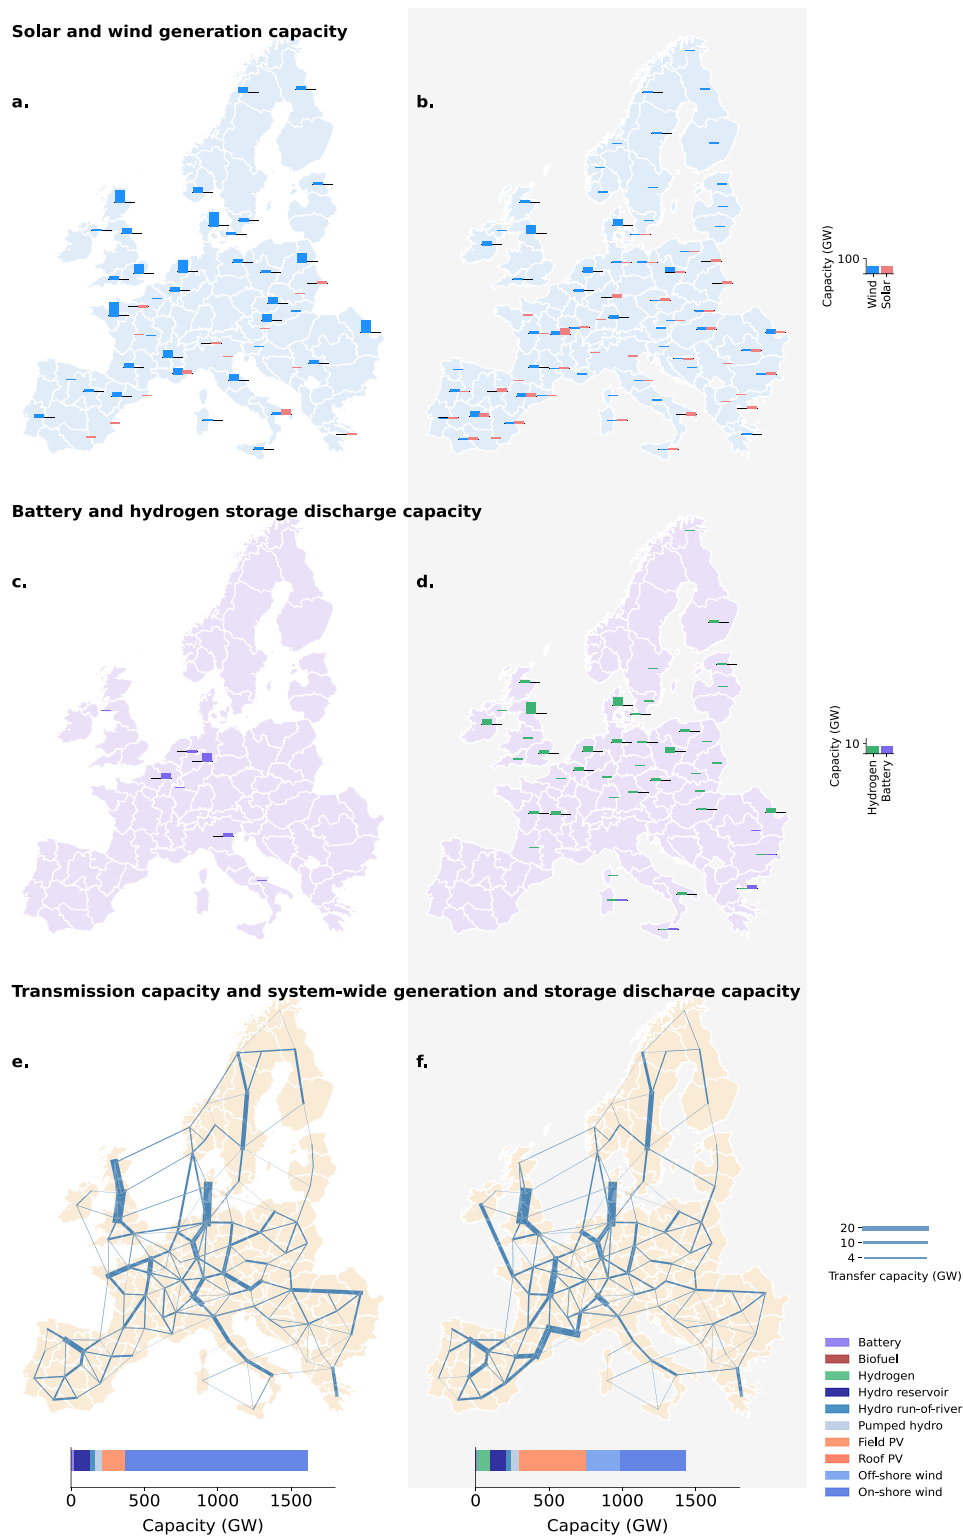


Fig. 4. SPORES with the lowest concentration of on-shore wind capacity and simultaneous minimisation of bioenergy capacity in the *random* (left, grey background) solution spaces. Panels **a–b** show the deployed capacity of solar (rooftop plus open-field PV) and wind (on-shore plus off-shore) in each of the 97 model locations where the locally-deployed capacity of at least 10 GW. Panels **c–d** show the spatial deployment of battery and hydrogen storage (discharge) capacity where the locally-deployed capacity is of at least 1.5 GW. Panels **e–f** show the capacity of transmission lines connecting each location, including both existing and expanded line capacity. Panels **g–h** show the system-wide capacity deployed for all considered generation and storage technologies.

Third, as a consequence of the above findings, and particularly when MGA is used to support practical decisions, the computational workflow should foresee iterations with the relevant stakeholders. In such a way, the initial, inherently inexhaustive decision space can

be refined based on stakeholder feedback, redirecting the available computational power specifically towards their interests and needs. In this framework, it might be ideal to set up the search in a way that initially compromises between spatial and technology dissimilarity,

providing a practically helpful overview of the options. The *hybrid* workflow proposed in Section 3.2 might be the most suited for such an initial exploration. Another option, albeit more computationally burdensome, could be to integrate the abovementioned brute-force methods for the homogeneous exploration of technology mixes with the search for a few spatially-distinctive solutions around each found technology mix. Regardless of the chosen method, stakeholders may then indicate themselves which additional technology combinations they would like to investigate that are not initially available, or for which of the existing technology features they would like to see more options to locate capacity spatially. The conceptualisation of a coherent method for creating such a ‘human–computer feedback loop’ is one further development of our method that we are investigating in the context of an ongoing project, SEEDS.¹

Overall, our results demonstrate that using modelling analyses to outline viable deployment strategies is, in practice, a challenging task. Previous studies have exposed the pitfalls of the common practice of relying on single, cost-optimal results, and have emphasised the potential of MGA to mitigate the provision of misleading insights [15, 18–20]. Our study does confirm that MGA has the potential to provide more meaningful and robust samples of the practically feasible decision space, particularly when carried out from multiple search directions, as done in the most recent literature [16,20,21]. Nonetheless, our results also warn that MGA is not a panacea per se and needs to go along with careful modeller judgement. For state-of-the-art energy system models of large size, even around 200 samples of near-optimal alternatives – let alone the standard provision of a single, optimal solution – may not be enough to robustly assess whether a particular system configuration is not feasible near the economic optimum. Our arbitrary yet highly plausible example of the configuration that minimises bioenergy while limiting on-shore wind capacity concentration is one such case in point. For instance, let us assume we had provided results based on only one of the search strategies tested in this work, say the *integer* one. We might have been tempted to conclude that minimising the use of bioenergy in a fully carbon-neutral European power system is only possible if accepting large, concentrated on-shore wind power hubs. Yet, we have seen that other search strategies lead to a different conclusion. Besides increasing the sample size and search directions as we do in a recent study [16] and adopting more balanced search strategies like the *hybrid* one proposed in this work, a computationally inexpensive solution to improve the robustness of modelling results might be the adoption of ad-hoc MGA sensitivity analyses targeted to specific policy-relevant claims. In the example above, we could have customised the MGA search to explicitly look for a few solutions with minimal bioenergy deployment and constrained concentration of on-shore wind capacity, finding out that those do, in fact, exist. Such a simple MGA-based counterfactual experiment may help corroborate modelling results both in conventional cost-optimisation studies and in cases in which high-performance computing facilities are not available and carrying out a very large MGA analysis is impossible.

A less immediate option to improve the computational viability of MGA analyses for complex energy system design problems could be exploring radically different algorithms, such as heuristic ones. These typically iterate by moving from one sub-optimal solution to an improved one, gradually converging to the optimum. It may be thus sufficient to retain those sub-optimal solutions in memory to obtain both the optimum and the near-optimal alternatives with a single model run. So far, heuristics like particle swarm, genetic algorithms and others have been primarily employed in the context of multi-objective optimisation to find sets of Pareto-optimal energy system design options between two explicit real-world objectives [29]. We have here argued that this type of optimisation faces limits when

applied to complex energy transition questions because of the countless stakeholders and unmodelled objectives involved. However, the same underlying algorithms may be repurposed explicitly towards the generation of alternatives without explicit real-world objectives, as is the goal in MGA. This may allow overcoming the limitations of conventional multi-objective optimisation while retaining the advantage of heuristic algorithms. In addition, depending on the chosen approach, heuristic algorithms may allow for non-linear problem formulations, which cannot be achieved with deterministic optimisation for problems of very large size, such as those we discuss here. Future research should further explore the applicability of these methods to energy planning problems.

Regarding multi-objective optimisation, we have argued that MGA is designed to encompass a broad range of unmodelled objectives, thereby overcoming the limits associated with the finite objectives required by multi-objective methods. Nonetheless, it is worth noting how this is only valid when the applied cost relaxation is large enough to encompass Pareto-optimal solutions located far away from the economic objective. Furthermore, the broader the cost-relaxation space, the higher the number of alternatives required to explore it in a balanced way. For limited computational power and a narrow cost relaxation, multi-objective optimisation might still lead to solutions that lie outside the samples obtained via MGA, for instance, due to being substantially more costly than the selected cost relaxation allows. The research question should guide the choice between MGA and multi-objective optimisation.

Finally, our findings arise from the case study of a highly-resolved European power system, which we select as an ideal example of a ‘state-of-the-art, large-scale energy system model’ that requires advancements in how MGA is applied. However, the conclusions we draw based on such findings are not case-specific. On the contrary, they are valid for any other energy system model of similar (or higher) technical and spatial detail — for instance, the model of another continental power system; or the model of a country’s energy system at a highly granular sub-national resolution. Any such model would face a similar trade-off between spatial and technical dimensions when using limited computational power for generating alternatives.

CRediT authorship contribution statement

Francesco Lombardi: Conceptualization, Formal analysis, Investigation, Methodology, Project administration, Software, Visualization, Writing – original draft. **Bryn Pickering:** Conceptualization, Methodology, Software, Writing – review & editing. **Stefan Pfenninger:** Conceptualization, Methodology, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

All the data and code are made publicly available on Zenodo. The paper provides references and links to those.

Acknowledgements

This work was funded by the SEEDS project. The SEEDS project is supported by the CHIST-ERA grant CHIST-ERA-19-CES-004, the Swiss National Science Foundation grant number 195537, the Fundação para a Ciência e Tecnologia (FCT), Portugal, grant number CHIST-ERA/0005/2019, the Plan Estatal de Investigación Científica y Técnica y de Innovación 2017–2020, Programación Conjunta Internacional 2020 through the Agencia Estatal de Investigación, Spain, and the Estonian Research Council, Estonia, grant number 4-8/20/26. Model runs were performed on the ETH Euler cluster.

¹ <https://seeds-project.org>

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.apenergy.2023.121002>.

References

- [1] Süsner D, Ceglaz A, Gaschnig H, Stavrakas V, Flamos A, Giannakidis G, et al. Model-based policymaking or policy-based modelling? How energy models and energy policy interact. *Energy Res. Soc. Sci.* 2021;75:101984. <http://dx.doi.org/10.1016/j.erss.2021.101984>.
- [2] Deng X, Lv T. Power system planning with increasing variable renewable energy: A review of optimization models. *J Clean Prod* 2020;246:118962. <http://dx.doi.org/10.1016/j.jclepro.2019.118962>.
- [3] Süsner D, Gaschnig H, Ceglaz A, Stavrakas V, Flamos A, Lilliestam J. Better suited or just more complex? On the fit between user needs and modeller-driven improvements of energy system models. *Energy* 2022;239:121909. <http://dx.doi.org/10.1016/j.energy.2021.121909>.
- [4] DeCarolis J, Daly H, Dodds P, Keppo I, Li F, McDowall W, et al. Formalizing best practice for energy system optimization modelling. *Appl Energy* 2017;194:184–98. <http://dx.doi.org/10.1016/j.apenergy.2017.03.001>.
- [5] Yue X, Pye S, DeCarolis J, Li FGN, Rogan F, Gallachóir BÓ. A review of approaches to uncertainty assessment in energy system optimization models. *Energy Strategy Rev* 2018;21:204–17. <http://dx.doi.org/10.1016/j.esr.2018.06.003>.
- [6] Price J, Mainzer K, Petrovic S, Zeyringer M, McKenna R. The implications of landscape visual impact on future highly renewable power systems: A case study for Great Britain. *IEEE Trans Power Syst* 2020;1. <http://dx.doi.org/10.1109/TPWRS.2020.2992061>, Conference Name: IEEE Transactions on Power Systems.
- [7] Ellenbeck S, Lilliestam J. How modelers construct energy costs: Discursive elements in energy system and integrated assessment models. *Energy Res. Soc. Sci.* 2019;47:69–77. <http://dx.doi.org/10.1016/j.erss.2018.08.021>.
- [8] Brown TW, Bischof-Niemz T, Blok K, Breyer C, Lund H, Mathiesen BV. Response to ‘Burden of proof: A comprehensive review of the feasibility of 100% renewable-electricity systems’. *Renew Sustain Energy Rev* 2018;92:834–47. <http://dx.doi.org/10.1016/j.rser.2018.04.113>.
- [9] Prina MG, Casalicchio V, Kaldemeyer C, Manzolini G, Moser D, Wanitschke A, et al. Multi-objective investment optimization for energy system models in high temporal and spatial resolution. *Appl Energy* 2020;264:114728. <http://dx.doi.org/10.1016/j.apenergy.2020.114728>.
- [10] Pickering B, Choudhary R. Quantifying resilience in energy systems with out-of-sample testing. *Appl Energy* 2021;285:116465. <http://dx.doi.org/10.1016/j.apenergy.2021.116465>.
- [11] Brill ED. The use of optimization models in public-sector planning. *Manage Sci* 1979;25(5):413–22. <http://dx.doi.org/10.1287/mnsc.25.5.413>, Publisher: INFORMS.
- [12] DeCarolis JF. Using modeling to generate alternatives (MGA) to expand our thinking on energy futures. *Energy Econ* 2011;33(2):145–52. <http://dx.doi.org/10.1016/j.eneco.2010.05.002>.
- [13] Price J, Keppo I. Modelling to generate alternatives: A technique to explore uncertainty in energy–environment–economy models. *Appl Energy* 2017;195:356–69. <http://dx.doi.org/10.1016/j.apenergy.2017.03.065>.
- [14] Li FGN, Trutnevte E. Investment appraisal of cost-optimal and near-optimal pathways for the UK electricity sector transition to 2050. *Appl Energy* 2017;189:89–109. <http://dx.doi.org/10.1016/j.apenergy.2016.12.047>.
- [15] Berntsen PB, Trutnevte E. Ensuring diversity of national energy scenarios: Bottom-up energy system model with modeling to generate alternatives. *Energy* 2017;126:886–98. <http://dx.doi.org/10.1016/j.energy.2017.03.043>.
- [16] Pickering B, Lombardi F, Pfenninger S. Diversity of options to eliminate fossil fuels and reach carbon neutrality across the entire European energy system. *Joule* 2022;6(6):1253–76. <http://dx.doi.org/10.1016/j.joule.2022.05.009>, Publisher: Elsevier.
- [17] Victoria M, Zeyen E, Brown T. Speed of technological transformations required in Europe to achieve different climate goals. 2022, [arXiv:2109.09563](https://arxiv.org/abs/2109.09563).
- [18] DeCarolis JF, Babae S, Li B, Kanungo S. Modelling to generate alternatives with an energy system optimization model. *Environ Model Softw* 2016;79:300–10. <http://dx.doi.org/10.1016/j.envsoft.2015.11.019>.
- [19] Neumann F, Brown T. The near-optimal feasible space of a renewable power system model. *Electr Power Syst Res* 2021;190:106690. <http://dx.doi.org/10.1016/j.epsr.2020.106690>.
- [20] Lombardi F, Pickering B, Colombo E, Pfenninger S. Policy decision support for renewables deployment through spatially explicit practically optimal alternatives. *Joule* 2020;4(10):2185–207. <http://dx.doi.org/10.1016/j.joule.2020.08.002>, Publisher: Elsevier.
- [21] Neumann F, Brown T. Broad ranges of investment durations for renewable power systems, robust to cost uncertainty and near-optimality, [arXiv:2111.14443](https://arxiv.org/abs/2111.14443) [PhysicsSoc-Ph].
- [22] Pedersen TT, Victoria M, Rasmussen MG, Andresen GB. Modeling all alternative solutions for highly renewable energy systems. *Energy* 2021;234:121294. <http://dx.doi.org/10.1016/j.energy.2021.121294>.
- [23] Lombardi F. Customised pre-built sector-coupled Euro–Calliope model — Focus on the power sector and additional SPORES options. 2022, [http://dx.doi.org/10.5281/zenodo.6655601](https://dx.doi.org/10.5281/zenodo.6655601).
- [24] Pfenninger S, DeCarolis J, Hirth L, Quoilin S, Staffell I. The importance of open data and software: Is energy research lagging behind? *Energy Policy* 2017;101:211–5. <http://dx.doi.org/10.1016/j.enpol.2016.11.046>.
- [25] Sasse J-P, Trutnevte E. Regional impacts of electricity system transition in Central Europe until 2035. *Nature Commun* 2020;11(1):4972. <http://dx.doi.org/10.1038/s41467-020-18812-y>.
- [26] Pfenninger S, Pickering B. Calliope: A multi-scale energy systems modelling framework. *J. Open Source Softw.* 2018;3(29):825. <http://dx.doi.org/10.21105/joss.00825>.
- [27] Tröndle T, Lilliestam J, Marelli S, Pfenninger S. Trade-offs between geographic scale, cost, and infrastructure requirements for fully renewable electricity in Europe. *Joule* 2020;4(9):1929–48. <http://dx.doi.org/10.1016/j.joule.2020.07.018>.
- [28] Grochowicz A, van Greevenbroek K, Benth FE, Zeyringer M. Intersecting near-optimal spaces: European power systems with more resilience to weather variability. 2022, [http://dx.doi.org/10.48550/ARXIV.2206.12242](https://arxiv.org/abs/10.48550/ARXIV.2206.12242).
- [29] Prina MG, Manzolini G, Moser D, Nastasi B, Sparber W. Classification and challenges of bottom-up energy system models — A review. *Renew Sustain Energy Rev* 2020;129:109917. <http://dx.doi.org/10.1016/j.rser.2020.109917>.