Intersecting near-optimal spaces: European power systems with more resilience to weather variability

Aleksander Grochowicz [a,*,1], Koen van Greevenbroek [b,1], Fred Espen Benth [a], Marianne Zeyringer [c]

[a] *Department of Mathematics, University of Oslo, P.O. Box 1053 Blindern, Oslo, 0316, Norway*
[b] *Department of Computer Science, UiT The Arctic University of Norway, Postboks 6050 Langnes, 9037, Tromsø, Norway*
[c] *Department of Technology Systems, University of Oslo, P.O. Box 70, 2027, Kjeller, Norway*

ARTICLE INFO

ABSTRACT

We suggest a new methodology for designing robust energy systems. For this, we investigate so-called near-optimal solutions to energy system optimisation models; solutions whose objective values deviate only marginally from the optimum. Using a refined method for obtaining explicit geometric descriptions of these near-optimal feasible spaces, we find designs that are as robust as possible to perturbations. This contributes to the ongoing debate on how to define and work with robustness in energy systems modelling.

We apply our methods in an investigation using multiple decades of weather data. For the first time, we run a capacity expansion model of the European power system (one node per country) with a three-hourly temporal resolution and 41 years of weather data. While an optimisation with 41 weather years is at the limits of computational feasibility, we use the near-optimal feasible spaces of single years to gain an understanding of the design space over the full time period. Specifically, we intersect all near-optimal feasible spaces for the individual years in order to get designs that are likely to be feasible over the entire time period. We find significant potential for investment flexibility, and verify the feasibility of these designs by simulating the resulting dispatch problem with four decades of weather data. They are characterised by a shift towards more onshore wind and solar power, while emitting more than 50% less CO_2 than a cost-optimal solution over that period.

Our work builds on recent developments in the field, including techniques such as Modelling to Generate Alternatives (MGA) and Modelling All Alternatives (MAA), and provides new insights into the geometry of near-optimal feasible spaces and the importance of multi-decade weather variability for energy systems design. We also provide an effective way of working with a multi-decade time frame in a highly parallelised manner. Our implementation is open-sourced, adaptable and is based on PyPSA-Eur.

1. Introduction

The climate crisis and tumbling prices for renewable technologies in the last decade are leading to an unprecedented shift to variable renewable and other low-carbon energy sources. The pace and extent at which this transition is projected to take place often corresponds to a complete overhaul of currently existing energy systems within a few decades. This necessitates a renewed understanding of the workings and planning of energy systems, taking into account future unknowns including weather, climate, costs, and politics.

In the domain of energy system modelling, grappling with these unknowns is an open research problem which has prompted many different approaches. In particular, weather variability has been identified

to have a large impact on model solutions: cost-optimal solutions are often fragile in the sense that modelling with one weather year can produce solutions that are infeasible for other weather years. At the same time, recent work involving the relaxation of cost-optimality has revealed the opportunities and insights provided by the near-optimal feasible space of energy system optimisation models. Still, existing methods only map out near-optimal solutions partially or heuristically, and a complete understanding of the theoretical and computational trade-offs has not yet been developed.

In this paper, we introduce a methodology to study robustness of energy systems against uncertainties in inputs and apply it to inter-annual weather variability. To start with, we investigate methods for

---

approximating projections of near-optimal feasible spaces of energy system models. We then intersect these spaces for varying input data to produce solutions which are feasible in all scenarios. Finally, we propose to consider points well in the *interior* of this intersection as candidate robust solutions. This geometric approach lends itself to building increased resilience against uncertainties and perturbation in a systematic way and generating greater flexibility for policymakers.

We are interested in the problem of long-term planning of the European power system with a large share of renewable production capacity in order to meet emission targets. Capacity expansion models are often used for this purpose, in which capacities of generation technologies, storage, and transmission are optimised. The model at the same time ensures feasibility by simultaneously solving the corresponding optimal dispatch problem over a certain time period, "simulating" the operations of the network. Such models are vital exploratory tools used to shape high-level European energy policy.

In this study, we use the open-source, bottom-up energy system optimisation model (ESOM) called PyPSA-Eur (Hörsch et al., 2018) for our implementation. PyPSA-Eur consists of a model-building routine based on the PyPSA (Python for Power System Analysis) framework (Brown et al., 2018), collecting and processing the required input data from various sources. It assembles a faithful representation of the European high-voltage transmission grid and existing generation capacities, and uses Atlite (Hofmann et al., 2021) to compute capacity factor time series for renewable energy sources (PV, wind, hydro), based on historical ERA5 reanalysis weather data. We use PyPSA-Eur with a partial greenfield approach under perfect foresight, including existing transmission (expandable), hydropower and nuclear capacities (both non-extendable), but optimising renewables (onshore & offshore wind power, solar power), gas turbines and storage from zero. For the purpose of demonstrating and validating our methodology, we use a spatial resolution of one node per country, and a 3-hourly temporal resolution (without time aggregation). However, our techniques can be applied also when using a significantly higher resolution and including more technologies and energy sectors.

Usually, these types of energy system models are optimised using one or a few historical weather years (Ringkjøb et al., 2018), or different weather years are used for sensitivity analyses (Lombardi et al., 2020). The issue of weather year variability has been addressed in the literature (Pfenninger and Staffell, 2016; Pfenninger, 2017; Collins et al., 2018; Staffell and Pfenninger, 2018; Hilbers et al., 2020; Craig et al., 2022; Ruhnau and Qvist, 2022) and has been identified as an important factor for ESOM outcomes (Zeyringer et al., 2018). In particular, using only a single weather year as input data for ESOMs can produce design solutions which are over-fitted to that year, and are not feasible in general (Bloomfield et al., 2016; Zeyringer et al., 2018). However, most previous studies running ESOMs with decades of weather data have either focused on a single country such as Germany (Ruhnau and Qvist, 2022), the UK (Pfenninger, 2017; Zeyringer et al., 2018), and the US (Dowling et al., 2020) (a single-node model), or used only a European dispatch model to solve for operations, not capacity expansion (Collins et al., 2018). One study on the impact of different climate scenarios applied a TIMES model for the European power system with a decades-long modelling horizon (Simoes et al., 2021) in which the usage of representative time slices limited the ability to model storage and capture medium- and long-term effects. Pickering, Lombardi and Pfenninger have recently used the sector-coupled Euro-Calliope model in a study of the European energy system at a 96-node two-hourly resolution (Pickering et al., 2022); the main results were generated over a single weather year but validated using 8 additional weather years.

To the authors' knowledge, this is the first paper to run a spatially resolved ESOM for the European power sector with multiple decades of weather data, without aggregating to time slices. For this analysis, we use 41 years of ERA5 reanalysis data (Hersbach et al., 2018) for the European continent (from 1980 until 2020 inclusive). While we

aim to find system designs which are feasible for all weather years under consideration, we base our methods on optimisations with single weather years in order to reduce the computational burden. However, we are still able to optimise our model with all 41 weather years in order to validate our approach.

The robust solutions that we are interested in are near-optimal feasible solutions, meaning that their costs do not go beyond a previously defined threshold. These solutions are "close to cost-optimal" and leave room for alternative objectives and desirable qualities; the additional costs we accept lie below the 9%–23% deviation from cost optimality (due to political, social, or technical reasons) that have been observed in recent years in the UK (Trutnevyte, 2016). Following the works by Neumann and Brown (2021) and Pedersen et al. (2021) we exploit the geometric shape and properties of the near-optimal space defined by the ESOM. Instead of studying the full-dimensional near-optimal feasible space, we study a projection onto 5 relevant dimensions representing total investments in certain technologies.

By varying the weather years as inputs, we then construct one (reduced) near-optimal feasible space for each year. When we intersect these near-optimal feasible spaces, we obtain a space of solutions in which each point represents a set of total investment decisions which are feasible for every year under consideration.

As the most robust candidate in the intersection, we choose the point which lays in the middle, being as far away from being infeasible as possible. This means that changes in total investment decisions (up to a certain point) still leave us in the near-optimal feasible space for every weather year. We then map the total investments back to a full system design, and verify its feasibility by simulating its operations over the entire time period. Note that our form of "robustness" is a geometric concept (laying in the middle of a near-optimal feasible space) and is only loosely connected to robust optimisation.

Apart from contributing to the discussion on energy system robustness, our methods also have implications for ESOM parallelisation. The difficulty in parallelising linear program (LP) solvers has been highlighted as the main barrier preventing ESOMs in taking advantage of increasing computational power (Kotzur et al., 2021). While there are efforts to address this problem at the level of LP solvers (Rehfeldt et al., 2022), we work at the level of model formulation. Finding solutions which are feasible for many weather years by studying the intersection of their respective near-optimal spaces can be an alternative to solving ESOMs with many weather years outright, which is computationally prohibitive. Thus, our methods constitute a way of heuristically replacing one large (difficult to parallelise) optimisation by many optimisations with single weather years.

Finally, we formalise and significantly deepen the understanding of the geometry and approximation of near-optimal feasible spaces of ESOMs. Previously, Pedersen et al. proposed a methodology for approximating near-optimal feasible spaces of ESOMs (Pedersen et al., 2021), and used the results to study the density of certain system design properties under projections of the near-optimal feasible space. Furthermore, Lombardi et al. mapped out the utilisation of renewable capacity, transmission capacity, and storage capacity of chosen near-optimal solutions, depending on different uncertainties, indicating overlaps between these (Lombardi et al., 2020). We detail the dimension reductions involved in working with near-optimal feasible spaces, and how to map back and forth between the different stages. We then propose several variations on a general algorithm for approximating reduced near-optimal feasible spaces, and analyse their convergence characteristics. The application of geometric descriptions of near-optimal spaces to studying different weather years is also novel.

In summary, our paper contributes to the literature on ESOMs in several ways. We formalise a general framework for working with and intersecting near-optimal feasible spaces, which allows us to study uncertainties of different kinds. We apply this framework to a first-of-its-kind study of robustness of highly renewable scenarios for the

European power system to decades of weather data. Beyond robustness, the methods also contribute to parallelisation of energy system optimisation models.

In Section 2 we formalise the methodology and introduce the necessary steps to define "robust" energy system designs. Afterwards, in Section 3, we describe the adapted PyPSA-Eur model we use and our modelling set-up to obtain power systems resilient to 41 years of weather data. In Section 4 we present our main findings on using intersections of near-optimal feasible spaces, features of robust solutions, and performance. We discuss ramifications of our approach in Section 5, before we conclude with Section 6.

2. Methodology and formal definitions

We introduce the methodology used in this paper and describe how we apply this to obtain energy system designs robust to weather variability: in Section 2.1 we revisit the concept of near-optimality. In Section 2.2 we discuss how dimension reduction is necessary to describe the near-optimal feasible space in a computationally tractable manner, and we elaborate on how we approximate near-optimal feasible spaces. Afterwards we introduce robustness as the geometric property of lying in the intersection of different near-optimal feasible spaces (Section 2.3). In Section 2.4 we then justify our choice of the Chebyshev centre as our robust solution of choice, as its location implies maximal stability to perturbations. Finally, in Section 2.5 we suggest different allocations that can translate a point in the low-dimensional intersection of the near-optimal feasible spaces to a spatially resolved (full-dimensional) energy system design.

2.1. Near-optimality

Let a capacity expansion problem be given as the linear program

$$\min c \cdot x \quad \text{such that} \quad Ax \leq b. \tag{1}$$

Here, $x \in \mathbb{R}^N$ is a vector of decision variables, $c \in \mathbb{R}^N$ the coefficients of the objective function, and $A \in \mathbb{R}^M \times \mathbb{R}^N$ and $b \in \mathbb{R}^M$ give the set of M linear constraints. Let $\mathcal{F}$ be the feasible space of the linear program Eq. (1), defined as

$$\mathcal{F} := \{x \mid Ax \leq b\}. \tag{2}$$

Letting $x^* \in \mathcal{F}$ be an optimal solution with objective value $c \cdot x^* = c_{\text{opt}} \in \mathbb{R}$, and $\varepsilon > 0$ a chosen slack level, we define the ε-*near-optimal feasible space* as

$$\mathcal{F}_\varepsilon := \{x \in \mathcal{F} \mid c \cdot x \leq (1 + \varepsilon) \cdot c_{\text{opt}}\}. \tag{3}$$

When ε is clear from the context, we simply refer to $\mathcal{F}_\varepsilon$ as the *near-optimal* space. For general linear programs, $\mathcal{F}_\varepsilon$ is a convex polyhedron, and when x is bounded (as is the case for energy system models), $\mathcal{F}_\varepsilon$ is a convex polytope. To work with $\mathcal{F}_\varepsilon$ geometrically, we can solve the optimisation problem $\min d \cdot x$ s.t. $Ax \leq b$ and $c \cdot x \leq (1 + \varepsilon) \cdot c_{\text{opt}}$ for some objective d in order to find a vertex or boundary point of $\mathcal{F}_\varepsilon$. In the context of ESOMs, this amounts to solving the energy system model once with an alternative objective function.

The definition of a near-optimal space is not new in the context of ESOMs, and previous work has explored the near-optimal space either through uniform sampling as Pedersen et al. (2021), maximally different solutions as DeCarolis (2011) and Price and Keppo (2017), or extreme points of the space (Neumann and Brown, 2021).

2.2. Dimension reduction

The near-optimal space $\mathcal{F}_\varepsilon$ is high-dimensional and complex; large-scale ESOMs typically involve millions of variables (dimensions) and constraints (hyperplanes defining the polytope). In this section we reduce to a much lower-dimensional space in two steps; see Fig. 1 for an overview of the maps and spaces involved.

In the case of capacity expansion models, we are most interested in *investment decision variables* of the linear program, as opposed to all other (operational) decision variables. Specifically, let $x = (x^I, x^O)^T \in \mathbb{R}^N$ be split into investment decision variables $x^I \in \mathbb{R}^{N_{\text{inv}}}$ and operational decision variables $x^O \in \mathbb{R}^{N_{\text{op}}}$, where $N = N_{\text{inv}} + N_{\text{op}}$. Then we define the projection map $\pi : \mathbb{R}^N \to \mathbb{R}^{N_{\text{inv}}}$ by simply forgetting about the operational decision variables. The image

$$\mathcal{F}'_\varepsilon = \pi(\mathcal{F}_\varepsilon) = \{x^I \in \mathbb{R}^{N_{\text{inv}}} \mid x = (x^I, x^O)^T \in \mathcal{F}_\varepsilon\} \tag{4}$$

of $\mathcal{F}_\varepsilon$ under π is the N_{inv}-dimensional ε-near-optimal feasible space of investment variables.

The convex polytope $\mathcal{F}'_\varepsilon$ consists of all points x^I such that an energy system with capacity investments given by x^I is feasible and whose total system cost (including operations) is at most $(1 + \varepsilon) \cdot c_{\text{opt}}$. In short, it is the space of all near-optimal feasible investment decisions. This makes an explicit description of $\mathcal{F}'_\varepsilon$ interesting for decision-makers in order to explore different kinds of near-optimal investments.

However, in a model with a high spatial resolution, the number of investment decision variables N_{inv} is typically still in the hundreds or more (with multiple investment decisions at each node, and transmission expansion). This makes the polytope $\mathcal{F}'_\varepsilon \subseteq \mathbb{R}^{N_{\text{inv}}}$ difficult to work with, visually and mathematically. Specifically, in order to work with $\mathcal{F}'_\varepsilon$ we would want to find a set of points P such that $\mathcal{F}'_\varepsilon$ is the convex hull of P. However, the number of vertices of an N_{inv}-dimensional polytope defined by M hyperplanes is in $O(M^{\lfloor N_{\text{inv}}/2 \rfloor})$ — see Tóth et al. (2017), Chapter 26. This puts a precise description of $\mathcal{F}'_\varepsilon$ in terms of vertices out of reach.

One solution is to map down to a much lower-dimensional space where we group and aggregate investment decision variables. Let $x^I = (x_1, \ldots, x_{N_{\text{inv}}})$ be the individual investment decision variables. Let $T_1, \ldots, T_k$ be a collection of sets of indices with $T_i \subseteq \{1, \ldots, N_{\text{inv}}\}$ and $T_i \cap T_j = \emptyset$. For each index j in one of these sets, we also choose a coefficient/weight c_j. Then we define a linear map $\sigma : \mathcal{F}'_\varepsilon \to \mathbb{R}^k$ as:

$$\sigma(x) = \left(\sum_{j \in T_i} c_j x_j\right)_{i=1}^k. \tag{5}$$

In our case, we take each T_i to be the set of indices identifying decision variables that belong to a specific technology. We weight these decision variables x_j (for $j \in T_i$) by their respective capital costs c_j. Specifically, throughout this paper we consider $k = 5$, with $T_1, \ldots, T_5$ corresponding to transmission expansion, PV expansion, onshore wind expansion, offshore wind expansion and gas turbine expansion respectively. In effect, σ maps a vector x^I of investment decisions to a summary of selected total investment costs.

Let

$$\mathcal{A}_\varepsilon = \sigma(\mathcal{F}'_\varepsilon) = \{\sigma(x^I) \mid x^I \in \mathcal{F}'_\varepsilon\} \subseteq \mathbb{R}^k \tag{6}$$

be the image of $\mathcal{F}'_\varepsilon$ under σ. Then $\mathcal{A}_\varepsilon$ is a k-dimensional convex polytope (since convex polytopes are preserved by linear maps). Note that in our specific choice of $T_1, \ldots, T_5$ we have not included *all* investment decision variables in σ, only those we deemed most important for the particular model instances we work with. Of course, different dimension reductions can be achieved by other choices of aggregation (groups of indices T_i and coefficients c_j). While we have taken the coefficients c_j to be capital costs (making investment in different technologies easier to compare), the coefficients could, for example, also be set to 1 in order to consider only capacities.

The utility of the reduced near-optimal space $\mathcal{A}_\varepsilon$ is as a proxy for system feasibility. If we can describe $\mathcal{A}_\varepsilon$ well, we can quickly assess whether any given set of total investments $y \in \mathcal{A}_\varepsilon$ can result in a feasible system design. However, by aggregating investment decision variables, we lose information on the specific feasible system designs $\sigma^{-1}(y) \subseteq \mathcal{F}'_\varepsilon$ realising the total investments y. The trade-off is that the fewer dimensions k we aggregate to, the easier $\mathcal{A}_\varepsilon$ is to work with, but the less information it gives us. Each point $y \in \mathcal{A}_\varepsilon$ can have a large preimage under σ and π, meaning there may be many near-optimal
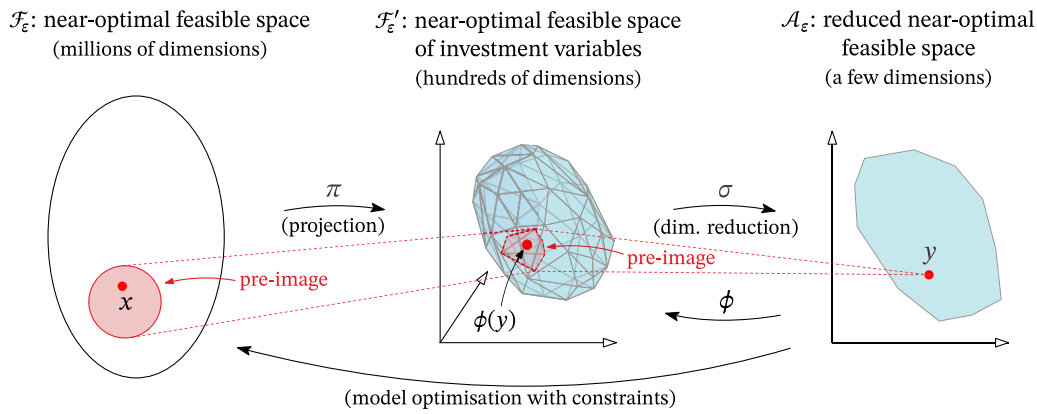
Fig. 1. Illustration of how the spaces $\mathcal{F}_\varepsilon$, $\mathcal{F}'_\varepsilon$ and $\mathcal{A}_\varepsilon$ are connected. The map ϕ can be seen as a composition of a model optimisation "$\min c \cdot x$ s.t. $Ax \leq b$ and $\sigma \circ \pi(x) = y$" composed with the projection π. It gives an explicit model design based on chosen coordinates y in the reduced near-optimal feasible space $\mathcal{A}_\varepsilon$.

feasible solutions $x \in \mathcal{F}_\varepsilon$ with the given total investments y (see Fig. 1). As we discuss in more detail in Section 2.5, to find a specific solution $x \in (\sigma \circ \pi)^{-1}(y) \subseteq \mathcal{F}_\varepsilon$, we must solve a version of the original model Eq. (1) in $\mathbb{R}^N$.

Similarly to Pedersen et al. (2021), we attempt to describe the low-dimensional space $\mathcal{A}_\varepsilon$ explicitly. Again, in order to find an explicit description of the polytope $\mathcal{A}_\varepsilon$, we would like to find (a large subset of) its vertices. We can do this by optimising over $\mathcal{A}_\varepsilon$ with different objective functions or *directions* in $\mathbb{R}^k$. For each direction $d \in \mathbb{R}^k$, the solution to the linear program

$$\max d \cdot y \quad \text{such that} \quad y = \sigma(\pi(x)) \text{ and } Ax \leq b \text{ and } x \in \mathcal{F}_\varepsilon \quad (7)$$

is an extreme point of $\mathcal{A}_\varepsilon$ in the direction d. The above linear program can be solved by solving the original problem in Eq. (1) with the new objective function $-d \cdot \sigma(\pi(x))$ (which is linear), and mapping the solution to $\mathcal{A}_\varepsilon$ by $\sigma \circ \pi$. In effect, for each extreme point of $\mathcal{A}_\varepsilon$ that we want to find, we need to solve the original capacity expansion problem once with an adapted objective function.

While we cannot expect to find all vertices of $\mathcal{A}_\varepsilon$ (quadratic in the number of model constraints when $k = 4, 5$ Tóth et al., 2017), we want to find a set of extreme points P such that their convex hull approximates $\mathcal{A}_\varepsilon$ well. Given a "budget" of n optimisations (and hence n extreme points), the natural question is: how do we choose the directions $d_1, \ldots, d_n$ to optimise in, in order to get a set of points P whose convex hull approximates $\mathcal{A}_\varepsilon$ the best possible?

We use an iterative approach to approximate $\mathcal{A}_\varepsilon$ while filtering on already used (or similar) directions. At each step we optimise in a different direction: depending on which property of the near-optimal feasible space is of interest, there can be many different ways to choose these directions. In our case we are interested in the largest ball within the near-optimal feasible space, the Chebyshev ball (see Section 2.4). Thus, we optimise at each iteration in the normal direction to a facet tangential to the Chebyshev ball of the current approximated polytope. If these have been exhausted, we choose the normal direction to the largest facet by volume. For other approaches to choosing directions, caveats and performance comparisons, see Appendix A. An illustration of one step in this process is shown in Fig. 2. A simplified version of the algorithm (based only on normals to large facets) is given in pseudo-code in Algorithm 1.

Once in the low-dimensional space $\mathbb{R}^k$ (with $k = 5$ in our case), the complexity of the geometric objects of interest is low enough that we can do efficient exact computations on them. In particular, recall that a polytope with n vertices can only have $O(n^2)$ facets for $k = 4, 5$, so computing the convex hull of the vertices, its volume, etc. has a time complexity of $O(n^2)$. We use the qhull software (Barber et al., 1996) for computational geometry related to polytopes. Since the number of vertices n which we can compute (by solving Eq. (7) for each vertex)

is limited, the complexity of $O(n^2)$ is acceptable for our purposes. In particular, note that for our application, the complexity of Algorithm 1 is dominated by the model optimisations, not the computational geometry.

Algorithm 1: Outline of algorithm for approximating $\mathcal{A}_\varepsilon$

$P := \emptyset$;
for $d \in \{e_1, -e_1, e_2, -e_2, \ldots, e_k, -e_k\}$ **do**
    Let y be extreme point on $\mathcal{A}_\varepsilon$ in direction d (optimisation);
    Add y to P;
end
for $i \in \{1, \ldots, n\}$ **do**
    Let H be the convex hull of P;
    Let $F_1, F_2, \ldots$ be the facets of H, sorted by decreasing volume;
    Let d_i be the normal of facet F_i;
    Let d be the first of $d_1, d_2, \ldots$ which is not within a small angle θ of any previously used direction;
    Let y be extreme point on $\mathcal{A}_\varepsilon$ in direction d (optimisation);
    Add y to P;
end
Return convex hull of P;

2.3. Intersections and robust solutions

One of the new ideas we propose is to investigate the intersections of the near-optimal spaces of related capacity expansion problems, or different instances of the same abstract model. Of course, if $\mathcal{F}^{(a)}, \mathcal{F}^{(b)} \subseteq \mathbb{R}^N$ are the feasible spaces of two linear programs A and B, then $\mathcal{F}^{(a)} \cap \mathcal{F}^{(b)}$ is simply the space of all solutions x which are feasible for both problems. More interestingly for capacity expansion problems, consider $\mathcal{F}'^{(a)}_\varepsilon \cap \mathcal{F}'^{(b)}_\varepsilon$: the space of all investment allocations which are both feasible and near-optimal for both A and B.

In our case, we consider the near-optimal spaces for optimisation problems defined with different weather years. Specifically, we use 41 years of reanalysis weather data (1980–2020) in order to compute capacity factors and load time series as input for our model (Section 3.2 and Appendix B). This gives 41 model instances, each defined with the input data from a different weather year. For brevity, let $\mathcal{Y} := \{1980, \ldots, 2020\}$ denote the set of weather years. Then for $i \in \mathcal{Y}$ write

$$\min c^{(i)} \cdot x^{(i)} \quad \text{such that} \quad A^{(i)} x^{(i)} \leq b^{(i)} \quad (8)$$

for the LP in Eq. (1) defined with weather year i. Let $\mathcal{F}^{(i)}$ be the feasible space of the above LP. From these feasible spaces, we want to recover investment allocations which are feasible for each of the weather years.
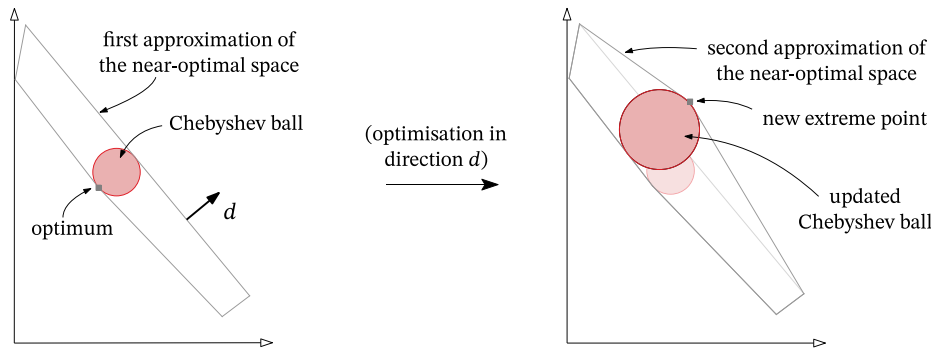
Fig. 2. Illustration of a single step of Algorithm 1. Additionally, the Chebyshev ball (see Section 2.4) is shown at each stage.

Since total system costs vary considerably between cost optimisations with different weather years, we define a uniform system cost bound across all weather years. Specifically, letting $c_{\text{opt}}^{(i)}$ be the optimal objective value (total system cost) for weather year i, we take

$$c^* = \max_{i \in \mathcal{Y}} c_{\text{opt}}^{(i)} \tag{9}$$

to be the highest optimal system cost across all weather years under investigation. In defining near-optimal spaces with the different weather years, we then set the slack relative to c^* instead of relative to each $c_{\text{opt}}^{(i)}$ for each weather year individually:

$$\mathcal{F}_\varepsilon^{(i)} = \{x \in \mathcal{F}^{(i)} \mid c \cdot x \le (1+\varepsilon) \cdot c^*\}. \tag{10}$$

We define $\mathcal{F}_\varepsilon'^{(i)}$ and $\mathcal{A}_\varepsilon^{(i)}$ similarly to $\mathcal{F}'$ and $\mathcal{A}_\varepsilon$ (Eqs. (4) and (6)).

Now, the intersection $\bigcap_{i \in \mathcal{Y}} \mathcal{F}_\varepsilon'^{(i)}$ is the space of investment decisions which are feasible for all weather years under consideration and near-optimal (relative to the most expensive year).[2] We are interested in the intersection of the reduced near-optimal space $\mathcal{A}_\varepsilon^{(i)}$. Repurposing our notation slightly, we write $\mathcal{A}_\varepsilon := \bigcap_i \mathcal{A}_\varepsilon^{(i)}$. Note that $\mathcal{A}_\varepsilon = \bigcap_i \sigma(\mathcal{F}_\varepsilon'^{(i)}) = \sigma\left(\bigcap_i \mathcal{F}_\varepsilon'^{(i)}\right)$. However, while we cannot easily find explicit descriptions of the spaces $\mathcal{F}_\varepsilon'^{(i)}$, we *can* approximate each (low-dimensional) $\mathcal{A}_\varepsilon^{(i)}$ as explained above. This, in turn, enables us to find an explicit (approximate) description of $\mathcal{A}_\varepsilon = \bigcap_i \mathcal{A}_\varepsilon^{(i)}$.

We call a set of total investments $y \in \mathbb{R}^k$ *robust* when $y \in \mathcal{A}_\varepsilon$. This means that for each weather year i, there exists some near-optimal feasible model solution $x^{(i)} \in \mathcal{F}_\varepsilon'^{(i)}$ (including investment and operation decisions) such that $\sigma \circ \pi(x^{(i)}) = y$. Note that the same total investment for each technology, $y \in \mathcal{A}_\varepsilon$, may be spread differently onto the different nodes of the model for each $x^{(i)}, i \in \mathcal{Y}$. Formally speaking, it is plausible that $\bigcap_i \mathcal{F}_\varepsilon'^{(i)} = \emptyset$ even if we find some robust point $y \in \mathcal{A}_\varepsilon$. In Section 2.5, however, we propose different methods for finding robust allocations $x^I \in \bigcap_i \mathcal{F}_\varepsilon'^{(i)}$ such that $\sigma(x^I) = y$ if they exist, and in Section 4.1 we show that this works well in practice.

2.4. Chebyshev centre

Among the total investment decisions in the intersection $\mathcal{A}_\varepsilon = \bigcap_i \mathcal{A}_\varepsilon^{(i)}$, we want to find choices that are not only feasible for all years considered. We want to find the most resilient choice among all the alternatives. We therefore select the point $y_{\text{ch}} \in \mathcal{A}_\varepsilon$ maximally removed in all directions from the boundary of $\mathcal{A}_\varepsilon$, meaning y_{ch} is as far away from being infeasible as possible. This is realised if we pick y_{ch} to be the Chebyshev centre (see Boyd and Vandenberghe (2004), Section 8.5.1), i.e.

$$y_{\text{ch}} = \underset{y \in \mathcal{A}_\varepsilon}{\text{argmax}}(r) \text{ s.t. } B_r(y) \subseteq \mathcal{A}_\varepsilon, \tag{11}$$

where $B_r(y)$ is the ball of radius r around y. Figs. 2 and 3(b) show examples of Chebyshev balls. The point y_{ch} can be found efficiently using a linear program. Specifically, let a_j be the normal vectors of the hyperplanes supporting $\mathcal{A}_\varepsilon$ and b_j the associated offsets, so that each $y \in \mathcal{A}_\varepsilon$ satisfies $a_j \cdot y \le b_j$ for all j. Then y_{ch} is given by

$$\max r \quad \text{such that} \quad a_j \cdot y + r\|a_j\| \le b_j \quad \forall j \text{ and } r \ge 0. \tag{12}$$

2.5. Disaggregating robust solutions

The previous steps leading to the Chebyshev centre, one chosen robust point, have all been performed in the reduced ε-near-optimal feasible spaces and their intersection, $\mathcal{A}_\varepsilon$. The designs that are of ultimate interest to us, however, are elements in $\mathcal{F}'$, including *all* investment decision variables. Thus, we would like to define a function $\phi : \mathcal{A}_\varepsilon \to \mathcal{F}_\varepsilon'$ mapping robust total investments to complete system designs realising those total investments. Being more precise, we need to specify over which weather years the space we map back to is defined — we want a map

$$\phi : \mathcal{A}_\varepsilon \to \mathcal{F}_\varepsilon'^{\mathcal{Y}}, \tag{13}$$

recalling the notation $\mathcal{Y} = \{1980, \dots, 2020\}$. From this map, we want to obtain a robust near-optimal energy system design $x_{\text{rob}} = \phi(y_{\text{ch}})$ at the Chebyshev centre of $\mathcal{A}_\varepsilon$.

We can write our ESOM as in Eq. (8) but instead define it with all 41 weather years as $\min c^{\mathcal{Y}} \cdot x^{\mathcal{Y}}$ such that $A^{\mathcal{Y}} x^{\mathcal{Y}} \le b^{\mathcal{Y}}$; its ε-near-optimal feasible space of investment variables is $\mathcal{F}_\varepsilon'^{\mathcal{Y}}$. Then we define the map ϕ (which we call $\phi^{\mathcal{Y}}$ for clarity) by adding constraints to the above LP to ensure that the solution has total investments given by $y \in \mathcal{A}_\varepsilon$:

$$\phi^{\mathcal{Y}}(y) = \pi(\underset{}{\text{argmin}} \, c^{\mathcal{Y}} \cdot x^{\mathcal{Y}}) \text{ such that } A^{\mathcal{Y}} x^{\mathcal{Y}} \le b^{\mathcal{Y}} \text{ and } \sigma \circ \pi(x^{\mathcal{Y}}) = y. \tag{14}$$

Given that π and σ are linear, the above is still a linear program. Note that $\phi^{\mathcal{Y}}$ may not be well-defined for all y if the corresponding linear program has no solutions; this can happen if some total investments y are realisable for each individual weather year, but not realisable by any one complete system design over all 41 years.

Solving Eq. (14) amounts to solving an ESOM defined with 41 weather years, which is computationally challenging. Indeed, one of the motivations for working with the intersection $\mathcal{A}_\varepsilon$ is that it can be computed on the basis of model optimisations with single weather years.

Thus, we propose two alternatives to the "exact" map $\phi^{\mathcal{Y}}$. We call the exact map and its alternatives *allocations* since they map total investments $y \in \mathcal{A}_\varepsilon$ to a spatially resolved allocation of investments $\phi(y) \in \mathcal{F}_\varepsilon'^{\mathcal{Y}}$. The alternative allocations are heuristics in the sense that they map to designs that are not strictly speaking guaranteed to be feasible. Formally, they map $\mathcal{A}_\varepsilon \to \mathbb{R}^{N_{\text{inv}}}$ but may map some points of $\mathcal{A}_\varepsilon$ outside of $\mathcal{F}_\varepsilon'^{\mathcal{Y}}$ (whereas $\phi^{\mathcal{Y}}(\mathcal{A}_\varepsilon) \subseteq \mathcal{F}_\varepsilon'^{\mathcal{Y}}$). For the first
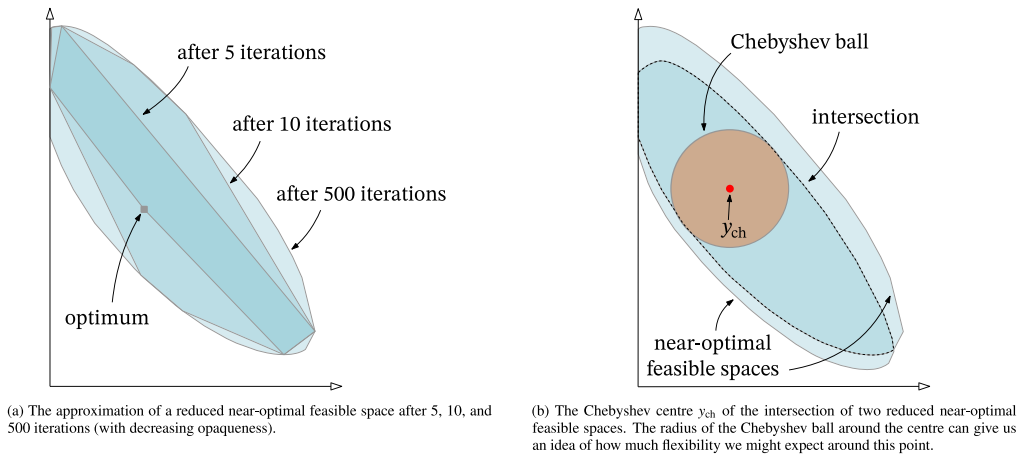
---

[2] In fact, this resembles the near-optimal feasible space of a robust optimisation program defined over the weather years $\mathcal{Y}$. Strictly speaking, however, the different operational variables for different weather years make this a loose generalisation of classical robust optimisation.

(a) The approximation of a reduced near-optimal feasible space after 5, 10, and 500 iterations (with decreasing opaqueness).

(b) The Chebyshev centre y_{ch} of the intersection of two reduced near-optimal feasible spaces. The radius of the Chebyshev ball around the centre can give us an idea of how much flexibility we might expect around this point.

Fig. 3. Near-optimal space approximation, intersection and Chebyshev centre.

Table 1
Overview over three different allocations to obtain "robust" solutions.

Map	Short name	Description	Computation
$\phi^{\mathcal{Y}}$	exact	Fix the investment costs in y_{ch} as additional constraints. Then solve Eq. (1) jointly over all years.	Optimisation over 41 years (134 GB RAM, 35 h)
$\phi^{(i*)}$	conservative	Fix the investment costs in y_{ch} as additional constraints. Then solve Eq. (1) for the year with the highest optimal cost, c^*.	1 single-year optimisations (ca. 3 GB RAM, 0.25 h)
ϕ^{mean}	mean	Fix the investment costs in y_{ch} as additional constraints. Then solve Eq. (1) for all single years and take the mean over all capacities.	41 single-year optimisations (parallelisable, each ca. 3 GB RAM, 0.25 h)

alternative, which we call the "conservative" allocation, let i^* be the most expensive year, for which $c_{\text{opt}}^{(i*)} = c^*$. Then we simply take $\phi^{(i*)}$ as an alternative to $\phi^{\mathcal{Y}}$. That is, for the conservative allocation we take the solution of the ESOM defined with weather year i^*, with the added constraints that $\sigma \circ \pi(x^{(i*)}) = y_{ch}$. Computing $\phi^{(i*)}$ involves only one model optimisation with a single weather year.

For the second alternative, which we call the "mean" allocation, we follow the idea of the conservative allocation $\phi^{(i*)}$, but involve the other weather years more. Indeed, we define ϕ^{mean} as

$$\phi^{\text{mean}} = \frac{1}{|\mathcal{Y}|} \sum_{i \in \mathcal{Y}} \phi^{(i)}. \tag{15}$$

Computing ϕ^{mean} involves $|\mathcal{Y}|$ model optimisations with single weather years. In our case, $|\mathcal{Y}| = 41$. The exact and alternative allocations are summarised in Table 1.

For comparison, we also define a "baseline" point in $\mathbb{R}^{N_{\text{inv}}}$ which is obtained by scaling up investments uniformly in the optimal solution for the most expensive weather year. This is explained in more detail in Section 4.3.

3. Implementation

3.1. Modelling set-up

We base our implementation on the PyPSA-Eur 0.4 model (Hörsch et al., 2018), which is itself based on the general PyPSA framework

(version 0.18) (Brown et al., 2018). While we have modified PyPSA-Eur for our purposes (as described in Appendix B), especially in order to support multiple weather years, we have kept the model set-up relatively close to the defaults as described in Hörsch et al. (2018). We use the model in a partial greenfield configuration, where existing transmission, nuclear[3] and hydro capacities at current (2020) capacities are included in the model from the start, but all other technologies start at zero capacity.[4] The extendable technologies included in the model are transmission (both AC and relevant DC connections), battery and hydrogen storage, onshore and offshore wind power, solar power and open-cycle gas turbines. The model is run with a single investment period and perfect foresight. We limit annual CO_2 emissions to 95% of 1990 levels.[5] A one-node-per-country[6] spatial resolution and a 3-hourly temporal resolution is chosen. Note, however, that the spatial and temporal resolution can be increased easily (as in PyPSA-Eur); the resolution is limited in this paper in order to allow extensive validation of our methods. We model the year 2030 with 41 distinct historical weather years (1980–2020) driving renewable capacity factors and electricity demand based on the default PyPSA-Eur cost assumptions[7] (given in 2013 EUR) for the year 2030. Thus all weather years are viewed as different potential realisations of 2030, and we can compare investment and operational costs of multi-year optimisations to single-year optimisations by taking annual averages.

Using enough weather data to accurately represent long- and short-term dynamics and extreme events is difficult in ESOMs, considering the resulting model size and increased solving complexity. However limited the lessons of historical weather data are on future weather (van der Wiel et al., 2019), further driven more and more by climate change, the extreme events and variability represented here will still likely offer insights for future designs. We capture historical climate change implicitly here, whereas incoming trends and changes through climate change are hard to predict and an active field of research in itself (Wohland et al., 2017; Schlott et al., 2018; Kozarcanin et al., 2019; Bloomfield et al., 2021).

---

[3] Nuclear power in Germany is removed from the model.

[4] PyPSA-Eur can also include existing biomass capacities in the model. Due to their limited capacities, they do not lead to significant deviations in results compared to the scenario we considered, so we choose to omit biomass for simplicity of the setup.

[5] We present additional results with a 100% emission reduction in Appendix C.

[6] Except for countries in multiple synchronous zones (Denmark, Spain, Italy, UK), which are represented through two nodes.

[7] https://github.com/PyPSA/pypsa-eur/blob/v0.4.0/data/costs.csv (accessed 06/10/2022).

3.2. Data

The data we use as input for our model lean heavily on the data and sources used in PyPSA-Eur. For instance, we use cost data based on the default costs considered in PyPSA-Eur, which are estimates for 2030[7] (in 2013 EUR). However, time series input data have been extended to 41 weather years.

We use reanalysis data to generate capacity factors for the renewable energy sources; our source is the hourly ERA5 dataset (Hersbach et al., 2018) for the time period from 1980 up to and including 2020. The open-source tool Atlite (Hofmann et al., 2021) translates weather data to hourly capacity factors for solar PV, on- and offshore wind. The inflow profiles for hydropower are generated similarly, however they are corrected to fit production data given by Administration (2022). We describe in detail how we process the data in Appendix B.

The load input data are generated using a regression model on 1980–2020 ERA5 temperature data. The regression is based on hourly country-level ENTSO-E load data from 2010 to 2014 (ENTSO-E, 2022), as well as the temperature data. We conduct a two-staged regression with a similar approach to that used in Bloomfield et al. (2020). This allows us to generate 41 years of country-level temperature-dependent load data matching the weather data we use. Finally, we scale the demand by a factor of 1.13 according to load projections for 2030 by the European Commission.[8] More details on the generated load data can be found in Appendix B.

3.3. Modeller's decisions

In this section, we discuss various details and choices regarding the implementation of the methods described in Section 2. For the exact code, including configuration options and installation and running instructions, we refer to the GitHub repository.[9]

One of the first decisions we have to make is choosing a suitable slack level ε. For small ε the intersection $\mathcal{A}_\varepsilon$ may be empty; a priori it is not clear how large ε has to be for $\mathcal{A}_\varepsilon$ to be nonempty. In our case, we choose $\varepsilon = 5\%$ on top of the most expensive weather year, but found that $\mathcal{A}_\varepsilon$ is even nonempty with $\varepsilon = 2.5\%$; this may change with a different modelling set-up. For comparison, Trutnevyte found in Trutnevyte (2016) that the transition in the UK energy system from 1990 to 2014 deviated between 9 and 23% from the cost optimum.

Given that the dimension reduction map σ is our main tool in working with near-optimal spaces (see Section 2.2), we need to define it carefully. Both the number of dimensions k that σ maps to and which investment decision variables are mapped to each dimension must be considered. We have investigated the convergence of Algorithm 1 with $k = 2, 3, \ldots, 7$ (Appendix A) and we find that it is tractable to work with this number of dimensions. The decision variables that are mapped to each dimension (the sets $T_1, \ldots, T_k$ in the notation of Section 2.2, here abbreviated to "dimensions") should be chosen meaningfully. On one hand, including as a dimension a technology which is not utilised in cost-optimal solutions may be detrimental, as much computational effort will be expended on potentially irrelevant solutions including this technology. Moreover, the Chebyshev centre of the resulting space must include at least a Chebyshev radius worth of that technology, which may be sub-optimal. On the other hand, *not* including as a dimension a technology which plays a significant role in any near-optimal solution can limit the usefulness of the results. In our case, we choose to map investment decision variables for transmission expansion, solar, onshore wind, offshore wind and gas to 5 respective dimensions.

Lastly, we choose the coefficients c_j in the definition of σ to be the investment cost associated with the investment decision variable

x_j; this has the advantage of mapping to the single unit of EUR in every dimension, making different dimensions easily comparable. In contrast, for example, comparing renewable capacity expansion in MW and storage capacity expansion in MWh directly is more difficult and not as useful for the purpose of working with a Chebyshev ball. For some applications one might consider scaling the weights c_j for certain dimensions, for example, when robustness to changes in investment in one technology are more important than for other technologies.

A key driver of the computational demands of our approach is the desired quality of approximation of near-optimal spaces. How many iterations of Algorithm 1 are needed, depends on the number of dimensions k, the mode of finding new directions to explore (see Fig. A.1) and the intended use-case. A fixed number of iterations can be chosen, or the algorithm can be ended once some convergence criteria is satisfied. Either way, we refer to Appendix A for a discussion of the various trade-offs involved.

The other factor influencing the computational effort is how much time and computing power every single optimisation takes. This is typically driven by the number of technologies, inter-temporal relations between different variables (e.g. through storage), temporal resolution (Hoffmann et al., 2020) and spatial resolution (Tröndle et al., 2020; Frysztacki et al., 2021). Deciding on the model complexity must be done in light of the research question at hand. For our application with weather years, we do assume that each individual model is defined over a time period of one calendar year.

4. Results

We first present the main results pertaining to system design with 41 weather years, and the use of near-optimal spaces in this context (Section 4.1). This is followed by a more detailed description of the characteristics of our proposed robust solutions (Section 4.2). The subsequent subsection focuses on validation results and a comparison of the different robust design allocations (Section 4.3). Finally, we touch on computational results and parallelisability (Section 4.4).

4.1. Weather years and intersection

First of all, optimising our model with each of the 41 considered weather years individually shows large discrepancies in the respective optimal solutions, re-affirming the importance of considering a large set of different weather years. Optimal total system costs range from 121 billion EUR for 2020 to 152 billion EUR in 1985.[10] The composition of investment by technology also differs significantly between weather years, with especially the onshore- and offshore wind investment varying by up to around 20 and 25 billion EUR between years (corresponding to variations up to 243 GW for onshore wind and 82 GW for offshore wind), respectively. The inter-year variability in optimal investments is illustrated in Fig. 5, where the investments are compared to the robust solution x_{rob}.

Meanwhile, for this study we also conduct the first capacity expansion optimisations of a spatially resolved model for the European power system with 41 weather years directly (one node per country, 3-hourly resolution). These optimisations took in the order of 1–2 days (using two threads) and up to 134 GB of memory. The optimal annualised total system cost for the 41-year model is 137 billion EUR; only slightly higher than the average of the total system costs of optimisations with single weather years at 134 billion EUR. See also Fig. 7 for a breakdown into investment (for extendable technologies)- and variable costs and comparison with the optimisations with single weather years. Apart

---

[10] These and all the following total system costs are annualised and include investment in new capacities as well as variable costs (in 2013 EUR), but not existing capacities. See Section 3.1 for details on which existing technologies are included in the model.
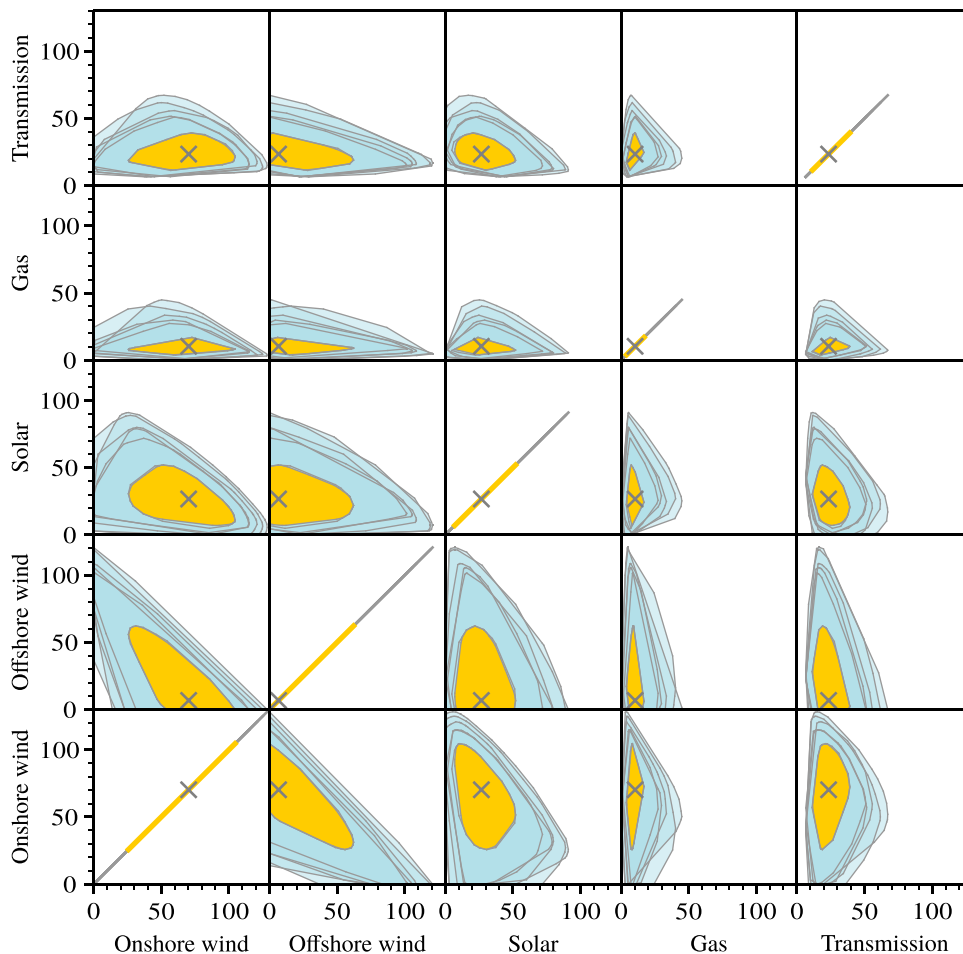
Fig. 4. Projections of the near-optimal spaces for different weather years and their intersection. All values are annualised total investment costs per technology. For illustrative purposes, we only plot the near-optimal spaces for 6 out of 41 weather years (in different hues of blue). The intersection of all 41 near-optimal spaces is filled in yellow and the Chebyshev centre is marked with a cross.

from providing a basis of comparison, optimising the model with 41 weather years also allows us to compute the exact robust map $\phi^{\mathcal{Y}}$ in order to validate the alternative mean and conservative allocations (Section 4.3).

We implement (as in Section 3) the methodology of intersecting near-optimal spaces (laid out in Section 2) using PyPSA-Eur. With a slack level of $\varepsilon = 0.05$ as in Eq. (9), we obtain a nonempty intersection $\mathcal{A}_\varepsilon$ and a robust allocation x_{rob} that is fully feasible using all the weather data of 41 years. Thus we show that there are robust solutions with less than 5% additional costs (on top of the most expensive year). We even find robust solutions which are less than 5% more expensive than a system optimised with the entire period of 41 years (see Fig. 7). Furthermore, the space $\mathcal{A}_\varepsilon$ offers significant flexibility for policymakers beyond our point of reference x_{rob}, and beyond what flexibility is shown by previous MGA approaches.

Fig. 4 shows projections of a selection of near-optimal spaces $\mathcal{A}_\varepsilon^{(i)}$ for weather years $i \in \{1985, 1989, 1996, 2006, 2014, 2020\}$ in addition to the intersection $\mathcal{A}_\varepsilon = \bigcap_{i \in \mathcal{Y}} \mathcal{A}_\varepsilon^{(i)}$ over all weather years. While the spaces are 5-dimensional (with the dimensions being investment in onshore wind, offshore wind, solar, gas and transmission expansion), they have been

projected down to all possible pairs of technologies considered. The Chebyshev centre, marked by a cross, is located within the intersection which consists of the robust solutions. The figure reveals that there is significant flexibility in these dimensions; while a certain amount of investment in renewables is needed, the investment can be shifted between different technologies while staying feasible and near-optimal. Note also that the near-optimal spaces for different years resemble each other in shape and location in space, but mainly differ in size. This indicates that the effect of "difficult" weather years on modelling is mainly that they restrict the size of the feasible design space.

We find that the Chebyshev radius of $\mathcal{A}_\varepsilon$ is 3.43 billion EUR, coming near to the theoretical maximum possible radius of 3.80 billion EUR given by the chosen slack level. Indeed, note that the distance between cost-optimal solutions and the near-optimal cost constraint is $c_{\mathrm{opt}}^* \cdot \varepsilon$, meaning that any near-optimal space can have a Chebyshev radius of at most $c_{\mathrm{opt}}^* \cdot \varepsilon/2 \approx 3.80$ billion EUR. The result means that the total investments in technologies that make up the dimensions of $\mathcal{A}_\varepsilon$ can change by up to 3.43 billion EUR (corresponding to 2.35% of the total cost of the robust system) in any direction, starting at the robust point y_{ch}. The resulting (potentially reduced) total investments can still result in a feasible design for every weather year.
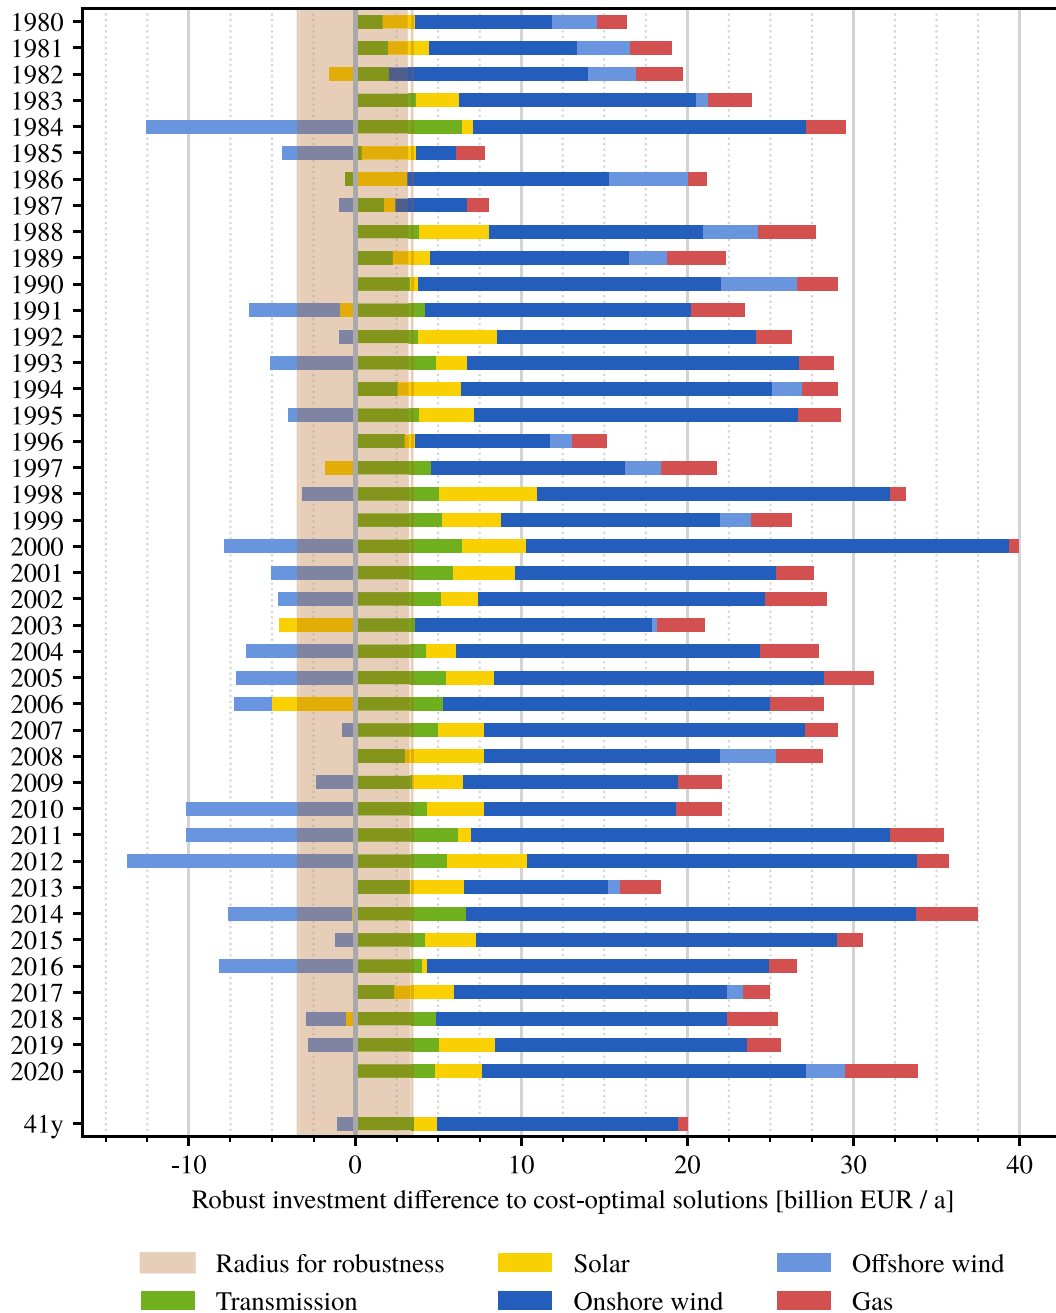
Fig. 5. Comparison of total investments for selected technologies in the optimal solutions for each weather year and the optimal solution with all weather years ("41y") to the robust point y_{ch}. Positive values mean greater investment in the given technology by the robust solution.

Given that we work with an absolute objective bound (1.05 times the cost of the most expensive year) for all near-optimal spaces, we find that the smallest and largest near-optimal spaces differ in volume by a factor of about 79. This means that some weather years by themselves allow for many more different near-optimal feasible solutions than others. Put differently, some weather years restrict the system design much more than others. The smallest and largest near-optimal spaces come from the weather years 1985 and 2020 respectively, and the difference in volume corresponds to an average scaling factor of $79^{1/5} \approx 2.4$ in every dimension. Meanwhile, the intersection of the near-optimal spaces has a volume that is 79% and 36% of the volumes of the near-optimal space for years 1985 and 1987, respectively, and is between

1% and 10% of the volume of all other near-optimal spaces. This means that except for 1985 and 1987, 10% or less of near-optimal solutions for any particular weather year are feasible for all other weather years under consideration.

In fact, we find that the optimal solutions with the years 1985 and 1987 actually have total investments that lie within the intersection $\mathcal{A}_\varepsilon$. When operated over the entire weather year dataset, these designs are practically feasible, with negligible load shedding. These results can inform the choice of weather year to model with — if only a single (or few) years can be chosen.

We note, however, that while there are weather years to which our model does not "over-fit" in a single-year optimisation (1985 being
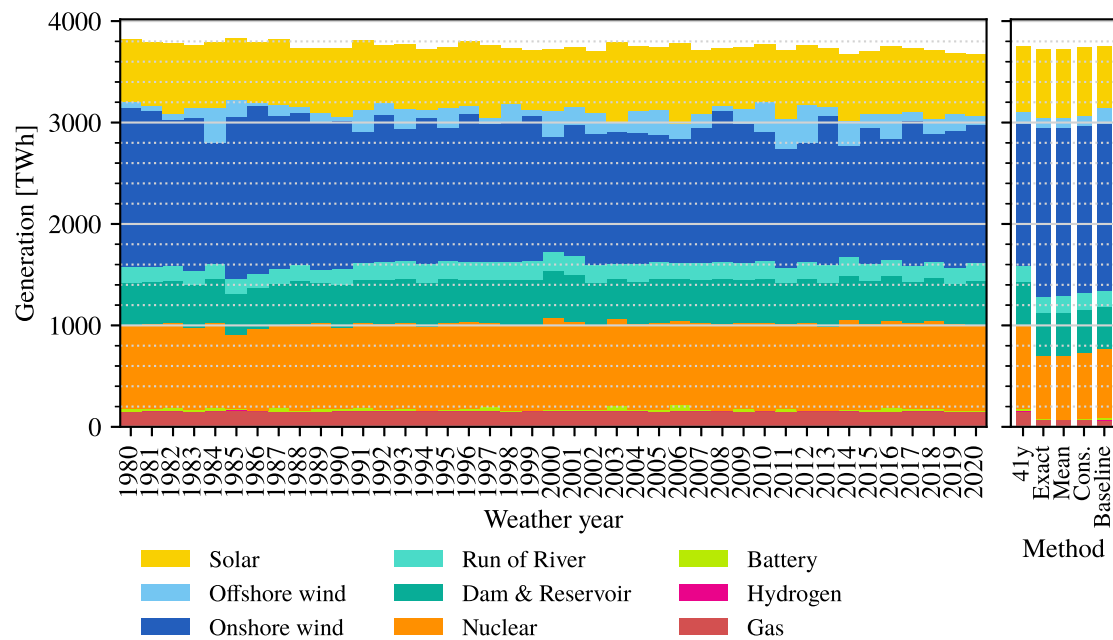
Fig. 6. Annual average generation mixes for the optimal solutions for each single weather year, the optimal solution with all weather years ("41y") and the robust allocations as well as the baseline design. The baseline design is used for comparison in validation (Section 4.3).

such a year resulting in a generally applicable design), this result could be particular to our modelling set-up.

4.2. Robust design characteristics

The robust solutions we compute, coming from the Chebyshev centre of the intersection $\mathcal{A}_\varepsilon$, have several interesting properties. First, we compare the investment composition of the robust point y_{ch} to that of optimisations using just a single year of weather data. We then analyse total investment and operational costs. Finally, we present results related to the CO_2 limit.

Fig. 5 shows the differences in investment per technology between the robust point and each of the optimal designs for individual years. The robust point is characterised by more investment in onshore wind, solar, transmission, and gas (sorted by decreasing additional investments). Meanwhile, most optima from single weather years over-invest significantly in offshore wind power compared to the robust allocation due to higher relative costs. In this particular set-up where gas can smooth the electricity production, the cost benefits of additional onshore wind capacities outweigh the potential of offshore wind power in the most favourable years. We conclude that here onshore wind power contributes more than other technologies to robustness, followed by solar power and transmission capacity.[11] This holds as well when one compares the investments in the robust ("exact") allocation to the optimal solution using all weather years (see the "41y" row in Fig. 5).

Fig. 6 shows the annual (for the optimisations with more weather years, average) total electricity generation per technology. This shows the increased importance of onshore wind and solar for the robust system. Meanwhile, the figure also shows that although solutions for individual years typically under-invest in gas turbines relative to the robust ones, the robust solutions actually generate less power with gas (and nuclear) in total. This reflects the fact that while additional gas capacity is needed for robustness, the additional investment in renewables leads to a reduced dependence on gas for "day-to-day" operations.

On that topic, Fig. 7 shows the variable and total system costs of systems optimised with single weather years, as well as the robust system and a system optimised with the full dataset of 41 weather years. It illustrates that the robust design has a higher total investment cost than designs for individual years, while the (average) operating costs are lower due to the reduced use of gas and nuclear as mentioned above. Recall that while we set the slack ε to 5%, we see that the investment costs in fact lie only 0.4% above what would have to be invested based on the most expensive year. With the (average) variable costs of the robust system being lower due to a strengthening of renewables, the total (annualised) system cost of the robust solution is about 146 billion EUR and actually lower than the system costs for some optimal solutions with single weather years. This is because the total system cost for the robust solution is averaged over all 41 weather years, with some being more expensive than others.

We also see that robust system designs emit less CO_2 in our tests compared to the single-year optimisations, and use 48% of the given CO_2 limit.[12] This is again because robust designs direct more of the total system cost into capital investment of renewables and less into variable costs including gas, the only source of emissions in our model. However, we should note that when we operate the design obtained from e.g. a system optimisation with the single weather year 1985 over the entire time period (over which it is practically feasible), it also does not use up the whole CO_2 limit. Although capacity expansion optimisations with single weather years always use up the CO_2 limit, the designs which are adapted to difficult years such as 1985 have enough renewable capacities that they do not use up the CO_2 limit in a typical year.

Finally, Fig. 8 shows the geographical differences in investment between the exact robust allocation and the cost-optimal solution with 41 weather years. We see that the majority of additional onshore wind power in the robust system is allocated to the UK, followed by Poland, France, Estonia and Sweden. The transmission capacity to the UK also receives more investment in the robust allocation. Meanwhile, although

---

[11] The conclusions can change with different assumptions — see Appendix C.

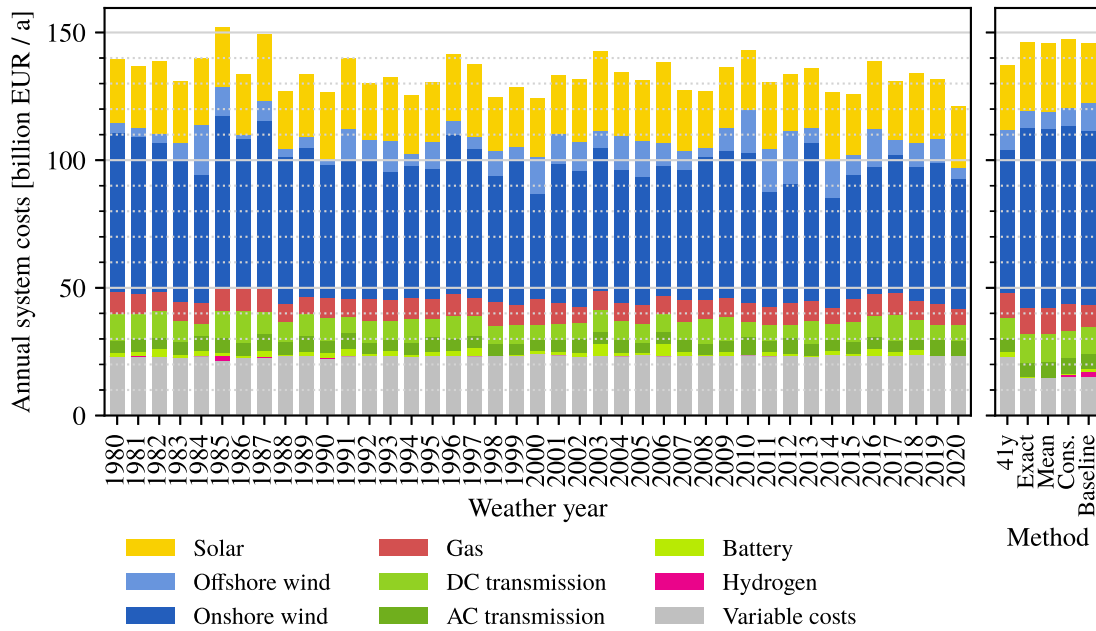[12] The optimal solution with all 41 weather years uses the whole CO_2 limit as well.

Fig. 7. Comparison of cost-minimal designs based on optimisations over individual years to the annualised costs of robust designs (exact, mean and conservative, respectively), an optimisation with all 41 years ("41y") and the baseline design.
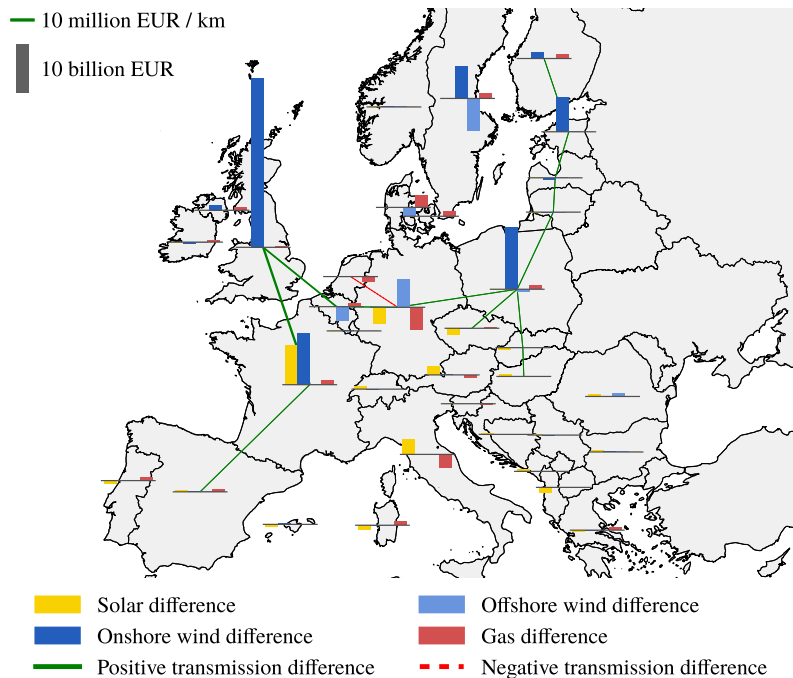

Fig. 8. Difference in investment between the exact robust allocation and the optimum solution with 41 weather years. Bars above the baselines mean that there is more investment for that technology in the robust solution. For transmission there are only positive differences (negative differences are too small to appear on this scale), and AC and DC transmission have been combined. For both bars and transmission lines, the same area means the same investment difference in EUR. All costs are annualised.

the system invests less in offshore wind in total, more offshore wind investment is shifted to Denmark.

4.3. Validation of robust allocations

Recall from Section 2 that we can use a (large) number of optimisations with single weather years to find a robust point of *total* capacity allocations $y_{ch} \in \mathcal{A}_\varepsilon$ — the Chebyshev centre in the intersection of near-optimal spaces of model instances using different weather years.

However, in order to find *specific* (per-node) investment allocations fitting the given robust totals and being feasible for every weather year, we have to map y_{ch} back to $\mathcal{F}'^{\mathcal{Y}}_\varepsilon$. The exact allocation $\phi^{\mathcal{Y}}$ does this by solving the original ESOM (with the additional constraint that $\sigma \circ \pi(x) = y_{ch}$) with 41 weather years. Section 2.5 gives two alternatives for allocating the robust capacities to individual nodes: the "mean" and "conservative" allocations, both based only on a number of optimisations with one single weather year at a time (see Table 1). Recall that both the mean and conservative allocations have the same coordinates

Table 2

Performance of the different allocation methods and the baseline for comparison. Note that total load shedding (over 41 years) here includes technical infeasibility and over-budget operations as explained earlier. The relative load shedding is with respect to the total load of all 41 weather years.

Allocation method	Total load shedding [TWh]	Relative load shedding [%]
Exact	0.0	0.000
Conservative	46.7	0.032
Mean	119.0	0.081
Baseline	132.3	0.090

in $\mathcal{A}_\varepsilon$ as the exact allocation, namely y_{ch}, but are not guaranteed to be feasible for all 41 weather years.

In this section, we investigate the quality of these heuristic allocations. The basis for comparison is two-fold: the exact robust allocation on the one hand, and a "naïve" baseline allocation based solely on the most expensive weather year on the other hand. Specifically, in order to make the designs more comparable, the baseline allocation is obtained by taking the allocation from the single most expensive weather year (1985 in our case) and uniformly scaling all expandable capacities up by such a factor that the total capital cost of the whole network equals that of the exact robust design. This way, the only difference between the exact, mean, conservative, and baseline designs is how investment is allocated (by technology and spatially), not the total investment volume.

The validation of the robust allocations consists of a "stress test" where we operate the systems over the entire dataset of weather years (i.e. only optimise the dispatch and not the investments). This dispatch optimisation includes the 95% CO_2 reduction constraint as in the capacity optimisation. Similarly to how we ensure that the total investment cost is equal between the four different allocations, we also ensure that all designs keep to the same operational budget; namely the total operational costs of the exact robust design. By allowing costly load shedding (with variable costs of 7 300 EUR/MWh, as in Price and Zeyringer (2022)), we make sure that the models are solvable. The quality of the allocations is then measured in the amount of load shedding; 0 load shedding means complete feasibility. In the presence of a hard operational budget constraint, it should be understood that total unmet demand combines actual unmet demand due to technical infeasibility and over-budget operations. Without the operations budget constraint, a number of designs – including all robust allocations and also some optimal solutions for single years – are practically feasible with all weather years (see Section 4.1), but simply have higher operational costs. In general, however, optimal solutions from single years fail to serve the entire time period reliably.

Table 2 shows that the exact allocation has 0 load shedding as expected, and is followed by the conservative allocation in quality. The mean and baseline designs perform slightly worse, but still only up to about 0.1% of the total load is shed or produced over budget. This shows that the conservative and mean heuristics produce results that are practically feasible.

4.4. Performance and computational effort

Our work demonstrates methods for working with decades of weather and demand data in a parallelised setting. Namely, the mean and conservative allocations produce robust system designs with decades of weather data while only requiring model optimisations with single weather years which can be parallelised effectively. This is done by using the centre point of the intersection of near-optimal spaces for different weather years to obtain a robust set of total capacities per technology. From these we then compute specific per-node capacity allocations adhering to the robust totals for each individual weather year, and average the resulting capacities.

This is particularly relevant because solving ESOMs is usually memory-constrained: a single large model may take more memory to

solve than is available on common systems. While effort is being spent on splitting up and parallelising the solving process (Rehfeldt et al., 2022), this is not yet practical. With our methods we avoid solving models defined over a period of decades, and instead obtain results based on many smaller runs — a process which is easily parallelised. The workflow is now constrained by the number of available processors, which is better suited to current computational developments which are in the direction of more, not faster processing cores.

While computational time often varies significantly based on a multitude of factors (exact modelling set-up, model size, nature of objective and constraints, numerical issues among others) and is difficult to predict, the memory requirements are easier to derive directly from model size (number of technologies, temporal and spatial resolution). With our modelling set-up (see Section 3.1), a capacity expansion optimisation with a single weather year takes 3.5–4 GB of memory using the commercial optimisation software Gurobi (Gurobi Optimization, 2022). Thus, our methods put optimal energy system design with decades of weather data within reach of typical desktop computers for the first time. For comparison, optimising with 41 years of weather data in one model takes approximately 41 times as much memory, around 134 GB. For models with a higher spatial and temporal resolution, our methods may currently be the only viable method of designing a system with decades of weather data; we estimate that optimising a typical PyPSA-Eur model with an hourly resolution and 181 nodes (as suggested in Frysztacki et al. (2021)) with 41 weather years could take around 1.7TB of memory, while this can be split into single-year optimisations taking around 45 GB of memory each using our methods.

With our modelling set-up, a model optimisation with one weather year takes 23 min of CPU time[13] on average using Gurobi with 2 threads. Spending 10 initial optimisations along unit vectors in positive and negative direction, 150 for the algorithm to approximate the near-optimal space for each year and finally 1 optimisation per year to compute per-node allocation of robust allocations, the mean allocation takes approximately $41 \cdot (10 + 150 + 1) \cdot 23 = 151\,823$ min ≈ 2530 h of CPU time to compute. On a cloud computing platform with 64 CPU cores (and around $4 \cdot 64 = 256$ GB of memory), this equates to around 40 h of wall time.[13] This is actually comparable to the wall time it typically takes to solve a single capacity expansion optimisation with 41 weather years of the same model on the same platform, which is in the range of 1–2 days.

We also investigate different methods for approximating near-optimal spaces of ESOMs, focusing on how optimisation directions are chosen, and how many optimisations are needed for a good approximation (for a different number of dimensions or projection variables). The results presented in Appendix A support the configurations we choose as default options.

5. Discussion

The urgent transition to energy systems based on intermittent renewable generation comes at a time with increasing computational power and availability of extensive climate data. However, while using ever larger models is helping us to understand the detailed functioning of future energy systems, it does not necessarily improve the understanding of uncertainties and resilience. In this paper, we propose a framework for producing energy system designs that are robust to uncertainty, taking advantage of the geometry of near-optimal spaces. We apply the framework to a study involving 41 weather years; the importance of using as much weather data as possible has been shown in the literature and is confirmed by our findings. Our framework can

---

[13] We make a distinction between the *CPU time* of a process, which is the sum of time spent on the process by all CPUs, and the *wall time* of a process, which is the elapsed real time between process start and finish. Thus CPU time may be larger than wall time for a parallelised process.

help policymakers to investigate different alternatives (similarly to the predecessors of this methodology, MGA and MAA) and overcome the fallacy of single solutions that can become infeasible through marginal perturbations.

Our choice of using the Chebyshev centre as a robust point is motivated by it being maximally tolerant to changes in investment in any direction. This is one of the distinctions to robust optimisation, where the optimum solution (with respect to the worst-case scenario) lies at the boundary of the intersection of feasible spaces, thus becoming infeasible with just marginal perturbations. At the same time, we see that even a naïve over-investment starting at the cost-optimal solution with a difficult weather year results in a system design which is practically feasible for all considered weather years. The choice of the Chebyshev centre can be seen as an attempt to find the most advantageous over-investment for robustness. In other words, while all additional investment contributes to robustness to some extent, our methods are aimed at finding the most efficient additional investments for robustness.

The first runs of a spatially resolved European power system optimisation model with four decades of weather data build on advances in open data accessibility and computational resources. Our results show that robust solutions stand out from standard cost-optimal solutions by a relatively higher investment in onshore wind and solar power. Furthermore, the increased investment in renewables reduces the usage of gas and nuclear, which lowers the CO_2 emissions of the system by more than 50% (in comparison to optimising the system with 41 weather years). By adjusting the slack level ε, our methods allow for a trade-off between robustness and additional costs.

Beyond looking at how weather years influence total system cost or investment composition, we see the effects of the choice of the weather year on the entire near-optimal space. Whereas the shapes and locations are fairly uniform, we see that near-optimal spaces for different weather years vary significantly in size. We still find a significant amount of flexibility within their intersection, mainly limited by a small number of most difficult weather years.

In previous studies utilising MGA techniques, dimension reduction has been used somewhat implicitly in order to make a systematic exploration of reduced near-optimal spaces feasible. Mapping down to a lower-dimensional space was seen as a way of summarising a few key properties of model solutions. However, the *geometry* of reduced near-optimal spaces has been largely unexploited, a gap we aimed to fill with this paper. Moreover, we investigate and propose practical methods for going back to spatially explicit investment decisions from points in the reduced near-optimal space. Mapping the Chebyshev centre of an intersection of near-optimal spaces back to a robust system design is just one application of this idea.

At the same time, approximating and intersecting near-optimal spaces as a way of designing systems which are feasible for many weather years is a reasonable alternative to optimising with many weather years directly from a computational perspective. Having split the problem into many smaller optimisations with single weather years, our methods are easily parallelised. This represents a new approach to optimisation problems which previously have been considered intractable or at least very impractical. In particular, an application of our methods to larger models (including sectoral coupling, higher spatial or temporal resolution) would be interesting and computationally feasible. While we expect the heuristic allocations to work at higher spatial resolutions (see also Frysztacki et al., 2022 for related work), their quality should be investigated under such conditions.

We highlight that we study just one out of many definitions of robustness (Moret et al., 2016; Maggioni et al., 2017); note also that our techniques do not directly follow the concept of *robust optimisation*. In all generality, robustness is a relative concept and is directed towards some uncertainty. Although these uncertainties can be well understood, some of them may be hard to quantify (e.g. political, societal changes), other could be epistemic (e.g. extreme weather events, changes in

costs, or misspecifications of the model), or even aleatoric (as the actual future weather conditions). For instance, Stirling (2010) locates robustness in the overlap of problematic levels of knowledge about possibilities and probabilities. All in all, our approach contributes to a wide academic debate about how to deal with uncertainty and robustness in energy systems modelling.

6. Conclusion

In this article we find that studying the near-optimal feasible space is a helpful tool to achieve more robust solutions against uncertainties for energy systems. When investigating weather variability, this enables us to quantify the variations and to find alternative designs that allow some flexibility for policymakers.

Since we utilised historical climate reanalysis data, we have not incorporated the consequences of climate change that an energy system in transition will face, nor the possibility of unseen extreme events. It would be interesting to understand which developments will limit flexibility and what events are defining for the design of a future (climate-)robust energy system. Moreover, we see that a small number of weather years including 1985 and 1987 constrain our system design the most, having relatively small near-optimal feasible spaces. These results call for a deeper understanding of which meteorological properties of these weather years, including extreme and compound events, are determining for energy system design.

While we have applied a notion of robustness in the reduced near-optimal space $\mathcal{A}_\varepsilon$, we have not considered robustness at lower levels of the network. So although the point $y_{\text{ch}} \in \mathcal{A}_\varepsilon$ is robust to a shift in investments of 3.4 billion EUR, the system $x_{\text{rob}} \in \mathcal{F}'_\varepsilon$ is not robust to such a shift at any one particular node in the network. There may also be many system designs with very different per-node capacities mapping to the same point y_{ch}. Another interplay with the spatial dimension is whether different regions contribute differently to the robustness of the whole system, as hinted at by Fig. 8. Thus, spatial aspects of robustness form an interesting avenue for future research.

More generally speaking, our methods leave a lot of room for different choices of dimension reduction: choosing which variables to aggregate and how. The choice leading to the most suitable reduced near-optimal space should be considered application-specific, depending on the technologies of interest, their relative importance and use-case for the reduced near-optimal space. Mapping to a space of investment costs as we did is a neutral choice, but which reduction is the most efficient for more specific purposes is still an open question.

The results we present here may depend on our particular modelling set-up. To which extent specific technologies contribute to robustness and flexibility may change if other technologies are included (e.g. through sector coupling), additional restrictions are introduced (e.g. on transmission) or the spatial and temporal resolutions of the model are improved. And whereas our model has a single investment period, it would be interesting to apply our methods with a model considering transition pathways through multiple investment periods. As our implementation is open-source and customisable, it should be adaptable to this setting.

Last but not least, our methodology can also be used to investigate other uncertainties besides weather variability. As near-optimal spaces strongly depend on cost assumptions, future applications of the present framework can contribute to an improved understanding of robustness in the face of uncertain costs.

CRediT authorship contribution statement

Aleksander Grochowicz: Conceptualisation, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualisation, Original draft & editing. **Koen van Greevenbroek:** Conceptualisation, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualisation, Original draft & editing. **Fred Espen Benth:** Conceptualisation, Methodology, Supervision, Review & editing. **Marianne Zeyringer:** Conceptualisation, Methodology, Supervision, 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.

Acknowledgements

Code & data availability

The code for our approach and its documentation can be found at https://github.com/aleks-g/intersecting-near-opt-spaces/tree/v1.0.1, and is made available under the GPL 3.0 license. Some of the data included in the above repository (load and hydropower capacity factors) were generated using the workflow available at https://github.com/aleks-g/multidecade-data/tree/v1.0 (code licensed GPL 3.0, data licensed CC-BY-4.0).

The reanalysis data we used for this study (for renewable capacity factors and temperature-dependent load data) were downloaded from the Copernicus Climate Change Service (C3S) Climate Data Store (Hersbach et al., 2018) using the Atlite software (Hofmann et al., 2021) as documented in the main repository. We have additionally made these weather data more easily available at https://doi.org/10.11582/2022.00034[14]; they are shared under the "License to use Copernicus Products",[15] which is comparable to the CC-BY license. Neither the European Commission nor the European Centre for Medium-Range Weather Forecasts is responsible for any use that may be made of the Copernicus information or data it contains.

Given the substantial size of the weather data needed for this study, we generated a number of PyPSA networks which can be used to reproduce our results without having to download and process the ERA5 data. They are available at https://doi.org/10.5281/zenodo.6683829 under the CC-BY-4.0 license.

All other data used in our model is directly inherited from PyPSA-Eur and also openly available as described in Hörsch et al. (2018).

Appendix A. Direction generation

We give a more detailed overview on how to explore the near-optimal space $\mathcal{A}_\varepsilon$ (here in the original sense as a reduced near-optimal space and not the intersection), using the notation from Section 2. Recall that we can compute an approximation of $\mathcal{A}_\varepsilon$ by finding a number of its vertices (or rather, extreme points), and each such point is obtained by solving the linear program in Eq. (7) with some different objective (*direction*) d. As described in Algorithm 1, we first optimise over $\mathcal{A}_\varepsilon \subseteq \mathbb{R}^k$ in each of the cardinal directions (positive and negative) in order to obtain a first full k-dimensional approximation. Each of these optimisations can be performed in parallel. Thereafter, in which direction d we choose to optimise over $\mathcal{A}_\varepsilon$ has a significant effect on how well $\mathcal{A}_\varepsilon$ can be approximated in a limited number of optimisations.

We investigate and compare three methods that generate directions in which to optimise over $\mathcal{A}_\varepsilon$. The first method is the simplest and

consists of choosing directions uniformly at random. For the remaining two methods, the idea is to compute the convex hull of the points obtained so far after every optimisation, say H_i, and use the geometric properties of H_i in some way to generate the "next" direction d_{i+1}. The convex hull is computed using the program qhull (Barber et al., 1996). The three methods are as follows:

1. "random-uniform": Choose a random vector from the uniform distribution on the sphere in $\mathbb{R}^k$.
2. "facets": Choose the normal vector to the facet of H_i with the largest volume.
3. "maximal-centre-then-facets": Compute the Chebyshev centre y_{ch} of H_i and the ball of maximal radius around it by solving Eq. (12). Of all facets of H_i tangential to this ball, choose the one with the largest associated dual variable in Eq. (12) and take its normal vector. If already used, take the normal vector to the facet of H_i with the largest volume.

With all these methods, we also use a filtering procedure: we discard vectors that have already been used. For each method, it is clear how to generate another direction if the first was discarded: for example for the "facets" method we choose the normal of the facet with the second-largest volume if the first direction was discarded. For the "maximal-centre-then-facets" methods we fall back on the "facets" method when all normals to facets tangential to the Chebyshev ball have already been used.

In the filtering procedure, we employ an angle threshold θ such that a potential direction d is discarded if it is within θ degrees of any previously used direction. If the filter discards all possible directions, we reduce the angle threshold θ by 20%. This is repeated every time a method "runs out of directions", until θ falls below a pre-defined minimum angle θ_{min}, at which point the whole algorithm is terminated.

Note that the approximation of $\mathcal{A}_\varepsilon$ can be parallelised effectively for any of the three direction generation methods, in the sense that multiple optimisations in different directions can be run in parallel. For the latter two methods, this means that the convex hull H_{i-P} must be used in the calculation of the ith direction when there are P parallel optimisations. When P is large, this means some of the generated direction could be slightly inferior (being generated with an older hull H_{i-P}).

We compare the performance of different methods in Fig. A.1. The plots show that the three different direction generation methods have different characteristics, but also that their performance varies substantially between different spaces (different weather years). Generally speaking, we see that the "random-uniform" method attains the largest volume in the long term, while the "facets" and "maximal-centre-then-facets" methods attain similar volumes and have a stronger performance initially. In terms of radius, the "random-uniform" method performs worse, while the "maximal-centre-then-facets' method converges the quickest initially.

We see that a large number of iterations is needed to converge in terms of volume, with the "random-uniform" attaining the highest volume, but still not converging after 2000 iterations. Meanwhile, the other two methods based on facet normals make large strides initially, but display a false convergence below the actual volume after about 1000 iterations. Convergence in terms of the radius is better (especially for the "maximal-centre-then-facets" method), but can be more erratic than the convergence of volume. 150 iterations with the "maximal-centre-then-facest" method were chosen as a compromise between accuracy and computational demand for this paper.

In Fig. A.2 we compare convergence of volume between different numbers of dimensions k of the spaces $\mathcal{A}_\varepsilon \subseteq \mathbb{R}^k$. For this plot, we use the "maximal-centre-then-facets" direction generation method and approximated the near-optimal spaces for the weather year 2020, but use different dimension reduction maps σ. The final approximated volumes are all normalised to 1. We see that 3- and especially 2-dimensional spaces are quickly approximated, while the convergence
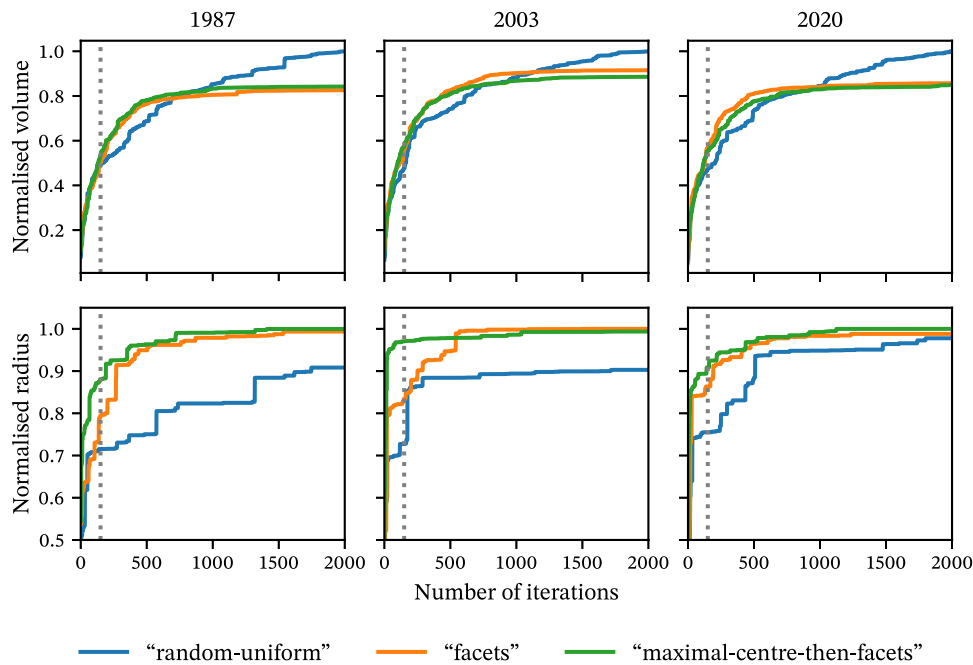

Fig. A.1. Performance of the different direction generation methods in terms of volume and Chebyshev radius convergence. We selected three years (among those, the years with the highest and lowest optimal costs) and approximated their near-optimal spaces with 2000 iterations. For each of the years the plotted volumes and radii have been normalised by the largest volume and radius obtained by any method for that year. The dotted lines mark 150 iterations.
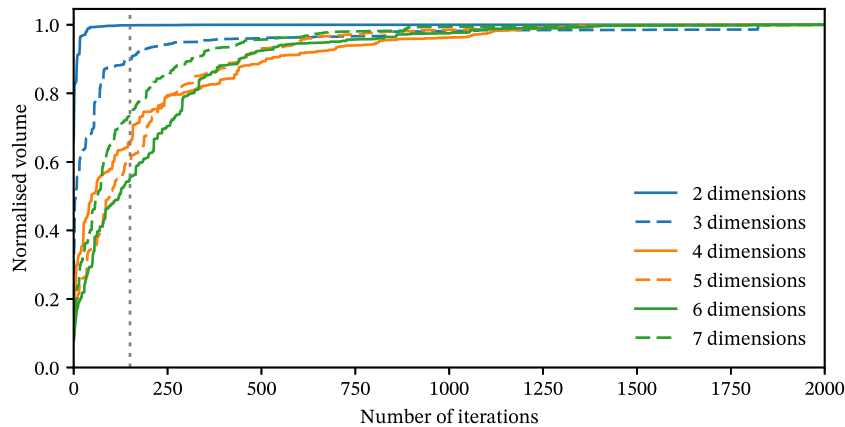

Fig. A.2. The convergence of volume in approximating the near-optimal space for the weather year 2020, reduced to different numbers of dimensions. The direction generation method is "maximal-centre-then-facets". The vertical dotted lines marks 150 iterations. The volume is normalised for each of the dimensions individually.

is slower for higher dimensional spaces. However, we do not find a significant difference in convergence between 4, 5, 6 or 7 dimensions.

The results in this section can be used to inform a termination criterion for Algorithm 1. The simplest option is to terminate the algorithm after a fixed number of iterations. Alternatively, the algorithm may be terminated after the volume or Chebyshev radius of H_i have not changed more than δ percent between successive iterations for the last N_{conv} iterations. In this case, we advise that N_{conv} be chosen as large as possible, since we can see from Fig. A.1 that the convergence on volume and especially radius is often somewhat erratic.

Appendix B. Data

B.1. Load data

In this article we use load data based on two regressions that were trained on hourly country-level ENTSO-E data from 2010 to 2014

from ENTSO-E (2022).[16] The aim is to infer country-level synthetic load data for each weather year between 1980 and 2020, whose profiles relate to weather patterns but are otherwise directly comparable. In other words, we disregard long-term changes in demand due to demographic and technological developments.

First we infer weekly load profiles (at an hourly resolution) for each country; for the purpose of this regression we treat holidays for each country as Sundays (using the Python package `python-holidays` Montel, 2022). Specifically, for each country c, we first divide the hourly demand values by the daily average value. On these normalised values, we conduct a regression based on the following model formulation:

$$D_c^{\text{norm}}(t) = \alpha_c \cdot t + \alpha_{c,\, t \bmod 168}, \tag{B.1}$$

---

[16] Due to inconsistencies for Swiss ENTSO-E data, we additionally used data from the Swiss transmission operator (Swissgrid, 2022).
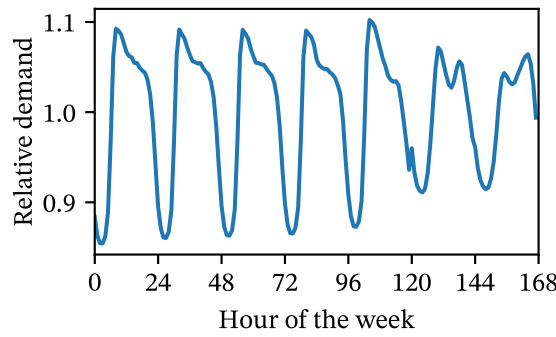
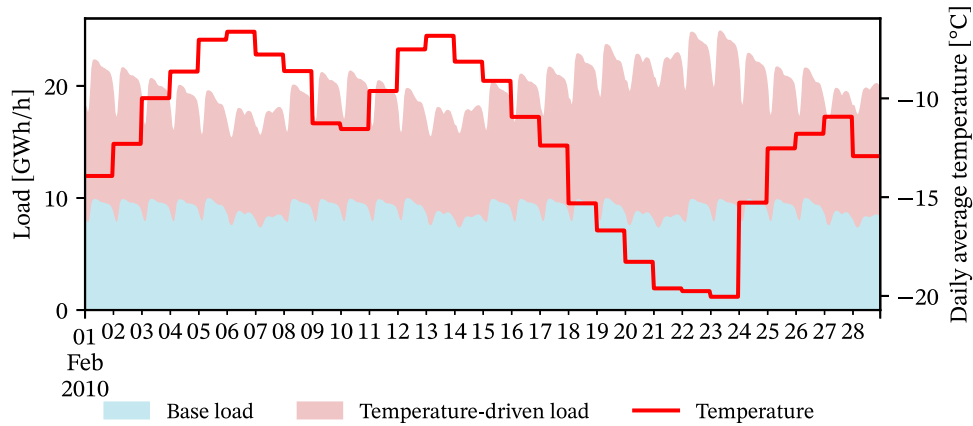**Fig. B.1.** Weekly load profile for Norway, based on the regression described in Eq. (B.1).

Fig. B.2. Load data (here for Norway in February 2010) split into base load and temperature-driven load. The underlying temperature data are daily, country-wide averages from ERA5 reanalysis. Base load and temperature-driven load are derived from the regression described in Appendix B.
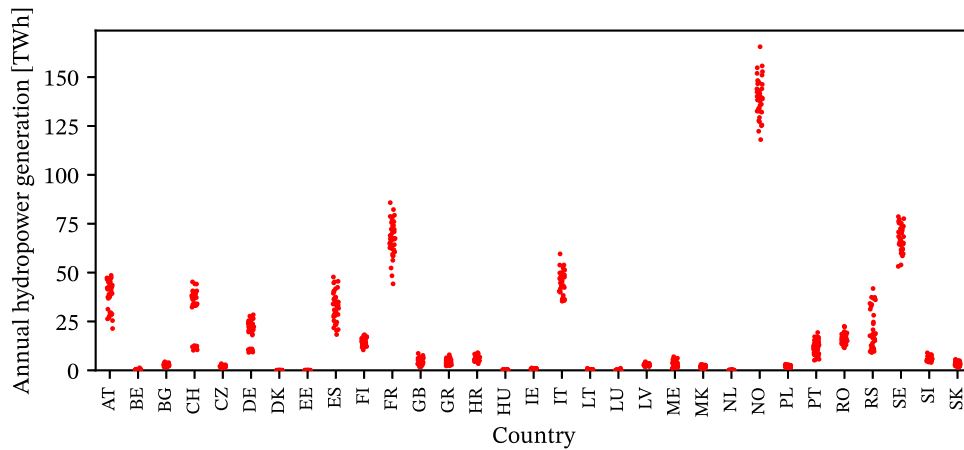

Fig. B.3. Variability of annual hydropower generation (1980–2020) based on EIA data (Administration, 2022), normalised by the reported hydropower capacities from 2020.

where α_c is an annual linear trend component, and the parameters $\alpha_{c,t \bmod 168}$ describe the weekly profile (see an example of this in Fig. B.1).

Afterwards we use the concept of heating and cooling degree days (HDD and CDD resp.) as in Benth et al. (2008) to find temperature-independent daily demand and demand driven by heating or cooling demand. For simplicity we only use one threshold for both HDD and CDD, which we set to be 15.5 °C as (Spinoni et al., 2015) use for HDD:

$$\text{CDD}_c(d) = \max\{T_c(d) - 15.5, 0\},$$

$$\text{HDD}_c(d) = \max\{15.5 - T_c(d), 0\},$$

where $T_c(d)$ is daily average temperature in country c during day d. Note that by using the daily averages we represent smoothing effects

of thermal inertia on heating and cooling demand (compare Fig. B.2). The country-wide temperatures are computed from ERA5 reanalysis data (Hersbach et al., 2018) via the open-source tool Atlite (Hofmann et al., 2021). We now conduct a regression on the daily average load, $D_c(d)$, with dummy variables for each weekday (where national holidays are classified as Sundays), and exogenous variables given by heating degrees and cooling degrees:

$$D_c(d) = \beta_{c,\text{weekday}(d)} + \beta_c^{\text{cooling}} \cdot \text{CDD}_c(d) + \beta_c^{\text{heating}} \cdot \text{HDD}_c(d). \quad (B.2)$$

For both regressions, we test the parameters for statistical significance; in some cases we thus set the trend parameter α_c and the cooling parameter β_c^{cooling} to 0.
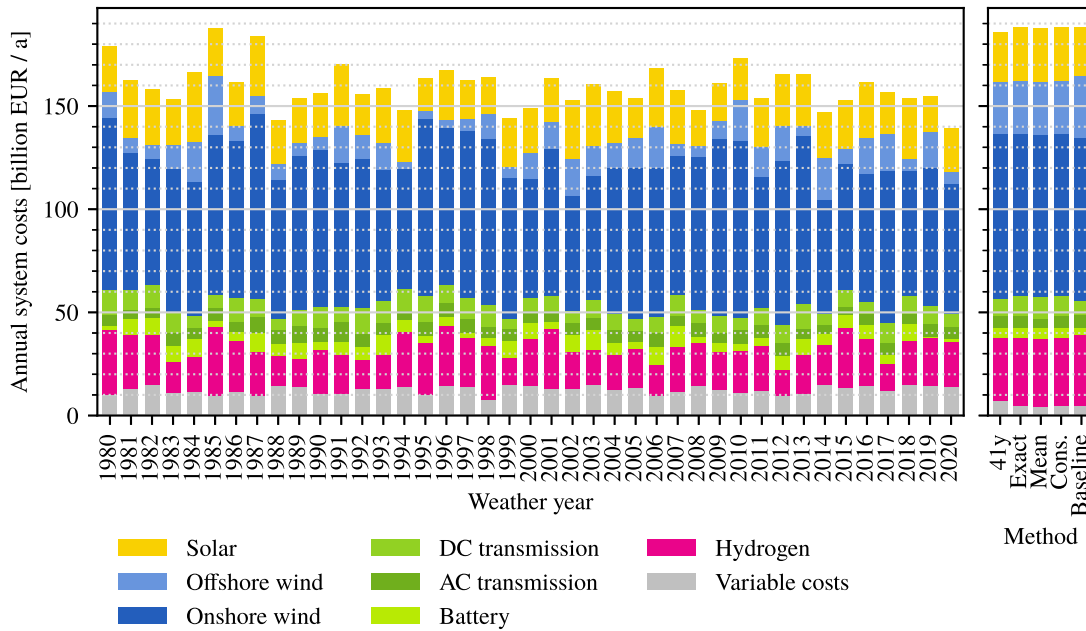
Fig. C.1. Comparison of cost-minimal designs under a 100% emission reduction based on optimisations over individual years compared to the annualised costs of robust designs (exact, mean, conservative and baseline, respectively) and an optimisation with all 41 years ("41y").
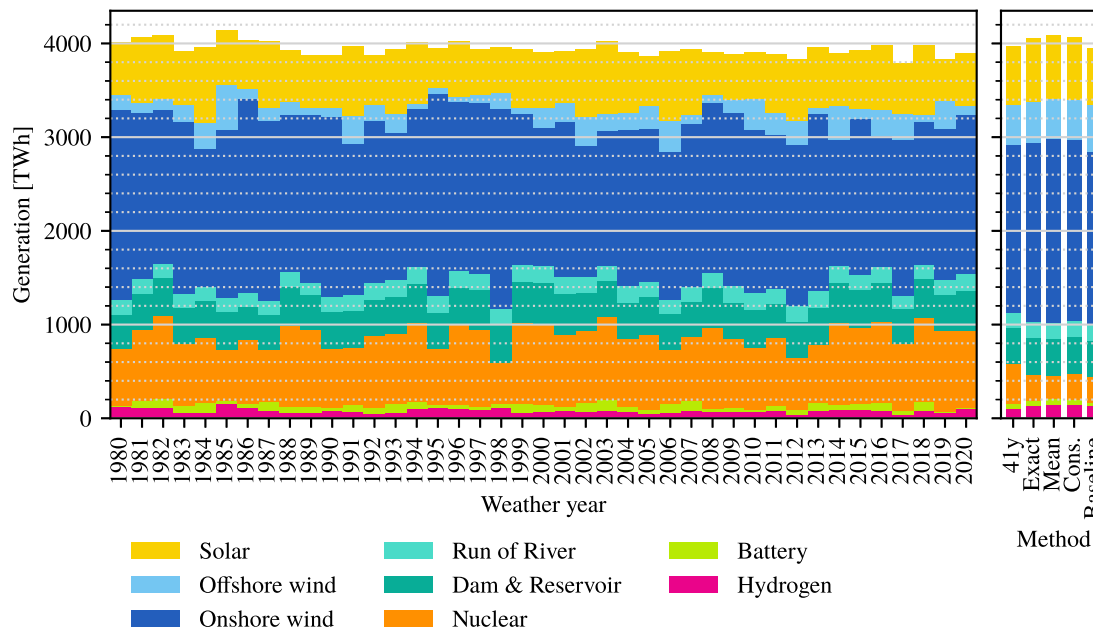

Fig. C.2. For 100% emission reduction, annual average generation mixes for the optimal solutions for each single weather year, the optimal solution with all weather years ("41y") and the robust allocations.

With these regressions and the temperatures for 1980–2020 from ERA5, we can compute the artificial load for each hour as follows:

$$\widetilde{D}_c(t) = \alpha_{c,t \bmod 168} \cdot \left[\alpha_c \cdot (t \bmod 8760) + \beta_{c,\text{weekday}(t)} + \beta_c^{\text{cooling}} \cdot \text{CDD}_c(t) \right.$$
$$\left. + \beta_c^{\text{heating}} \cdot \text{HDD}_c(t) \right], \tag{B.3}$$

where the index c is over countries, t is the time in hours, $\alpha_{c,i}$ the regression parameter for the ith hour of the week, α_c is the annual trend component, $\beta_{c,j}$ the regression parameter for weekday j, and β_c^{cooling}, β_c^{heating} are the regression parameters for one degree of cooling/heating

demand for country c. We abuse notation slightly by writing $\text{CDD}_c(t)$ to mean $\text{CDD}_c(d)$ where d is the day containing t (and likewise for HDD). For an example, see Fig. B.2.

We validate the regression on hourly ENTSO-E load data on the country level for 2015, and show it to be a good fit — more information about this can be found in the GitHub repository.[17]

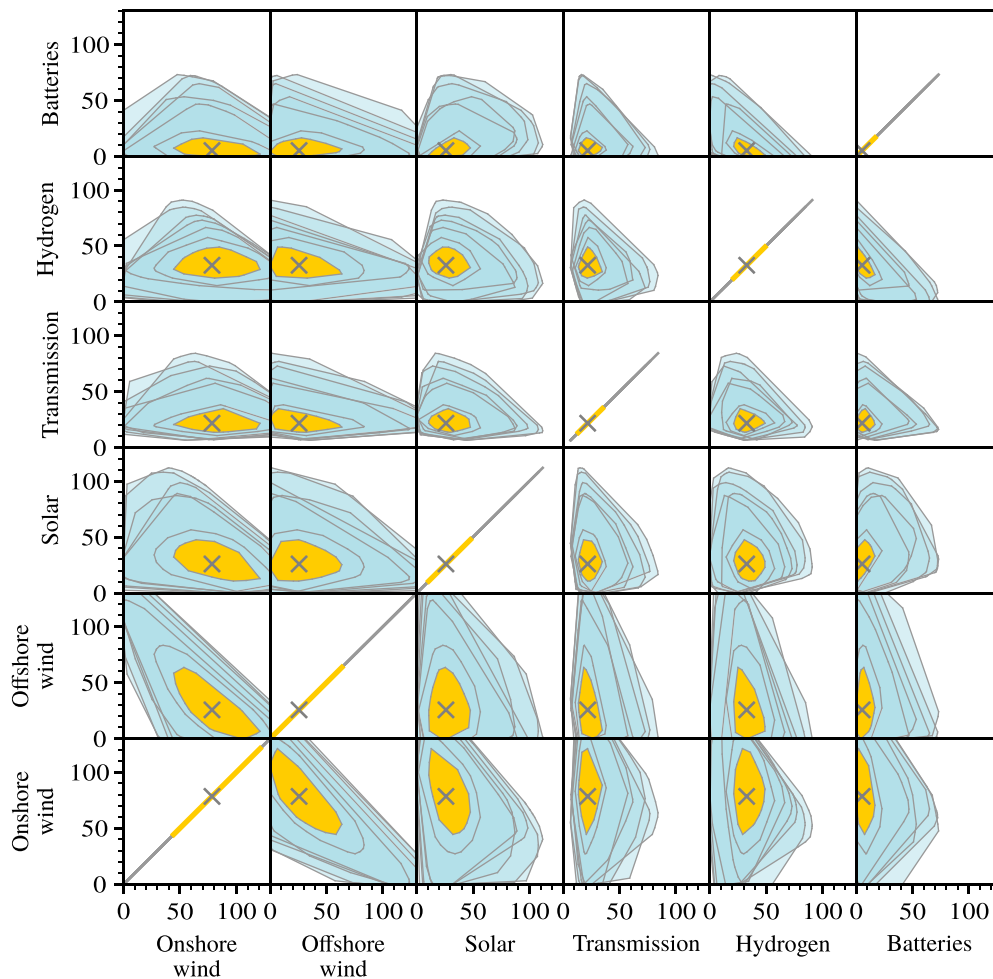[17] https://github.com/aleks-g/multidecade-data/tree/v1.0.

Fig. C.3. Projections of the near-optimal spaces for different weather years and their intersection under 100% emission reduction. All values are annualised total investment costs per technology. For illustrative purposes, we only plot the near-optimal spaces for 6 out of 41 weather years (in different hues of blue). The intersection of all 41 near-optimal spaces is filled in yellow and the Chebyshev centre is marked with a cross.

In accordance with load projections for 2030 by the European commission[18] we increased the demand in each country by 13%.

B.2. Hydropower data

For hydropower data we follow the approach in PyPSA-Eur which uses ERA5 reanalysis data to generate inflow profiles (using Atlite) that are then scaled by historical country-level hydro generation data from the US Energy Information Administration (EIA) (Administration, 2022). We have extended the default dataset in PyPSA-Eur to cover the entire period of 1980 to 2020, and we have also normalised EIA's production data to EIA's capacity levels of 2020[17]:

$$\text{gen}_c(y) = \text{nom\_gen}(y) \cdot \frac{\text{cap}_c(2020)}{\text{cap}_c(y)},$$

where $\text{gen}_c(y)$ and $\text{nom\_gen}_c(y)$ are the normalised and reported hydropower generation for country c in year y respectively, and $\text{cap}_c(y)$ is the reported hydropower capacity for country c in year y.

This allows a comparison throughout different years without any trends in infrastructure development. To avoid anachronisms, we have distributed generation and capacities of former countries onto the current states (based on the first year of current borders, e.g. 1993 for Czechia and Slovakia, or the sum of West and East Germany)[17].

The historical generation data ensure that the inflow profiles are scaled to reasonable values (see the general approach in Schlachtberger et al. (2017) and in particular Fig. 4 in Liu et al. (2019)); we are interested in variability and not trends, therefore we want fixed capacities to have comparable weather years (see Fig. B.3).

Appendix C. Additional results of a 100% emission reduction

We showcase the impact of a 100% emission reduction (as opposed to the 95% reduction studied in the main text) on the results we present in this paper. As the other assumptions remain unchanged and gas turbines are the only sources of CO_2 emissions in our model, this

---

[18] https://ec.europa.eu/clima/document/download/ec1acac9-10fe-4eeb-915f-cad388990e0f_en, Fig. 44 (accessed 23/06/2022).
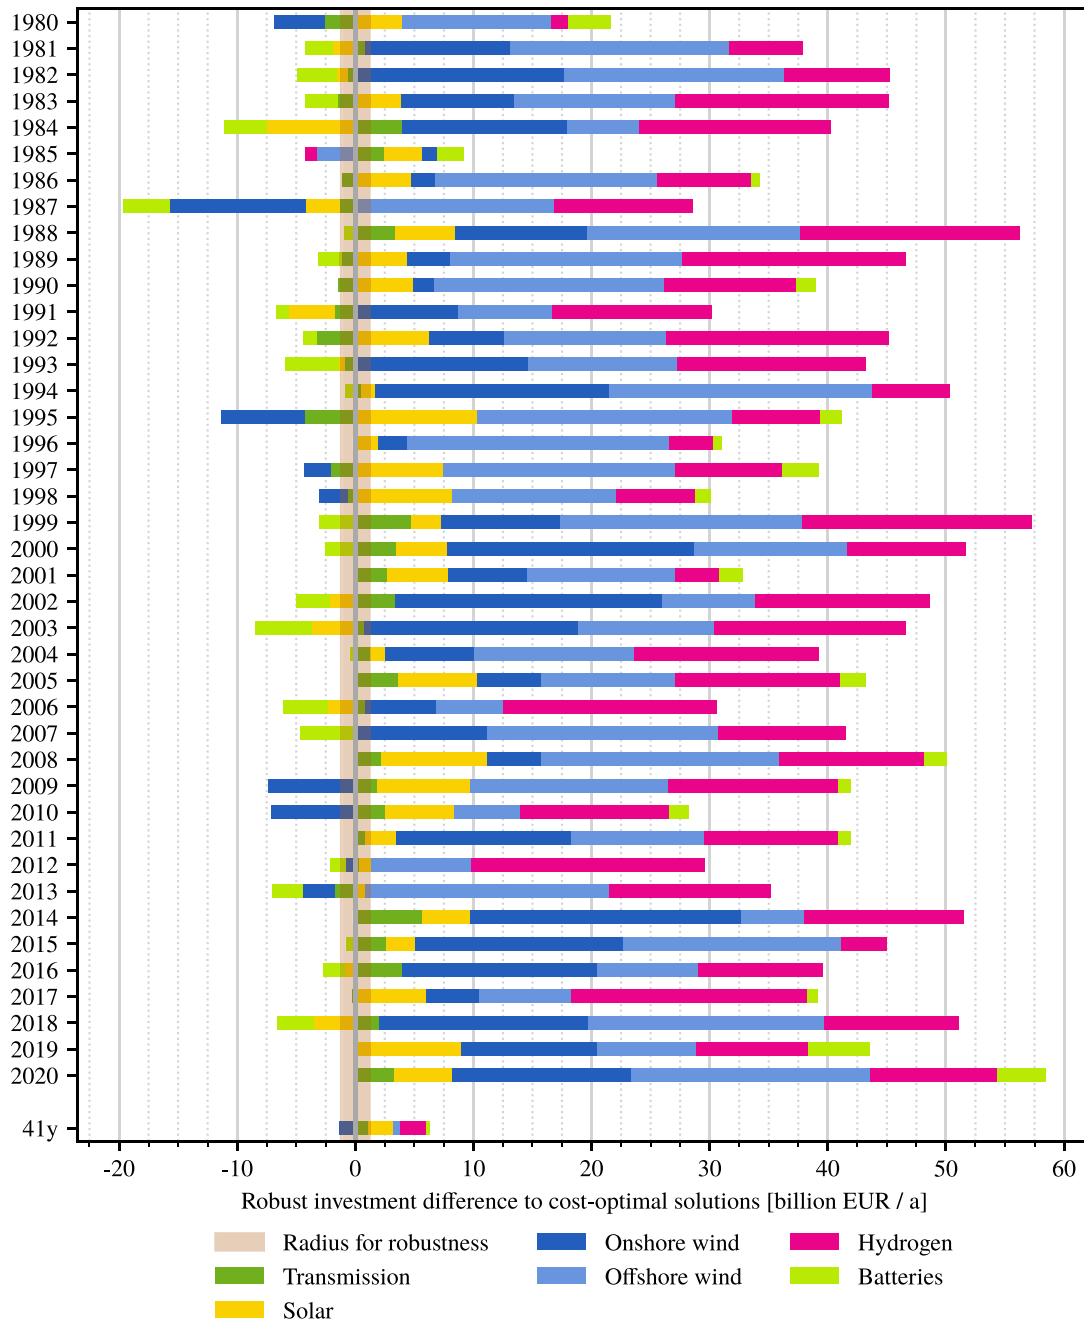
Fig. C.4. For 100% emission reduction, a comparison of total investments for selected technologies in the optimal solutions for each weather year and the optimal solution with all weather years ("41y") to the robust point y_{ch}. Positive values mean greater investment in the given technology by the robust solution.

corresponds to replacing gas as a generation technology with carbon-neutral alternatives. The results sharpen the uncertainty that weather variability introduces to renewable power systems. We investigate a different reduced near-optimal feasible space than before (replacing the gas investment dimension by two new dimensions, investment in battery and hydrogen storage). This means that the Chebyshev centre under this new reduction is additionally robust to changes in investment in battery and hydrogen storage, which is not the case for the results in the main text.

Without gas as a dispatchable generation technology, the costs of the now carbon-neutral power systems increase by 15 to 40 billion EUR/a, depending on the weather year. The last percentage points of emission reductions are thus overproportionally costly (as also shown in Neumann and Brown (2021)), in particular in "more difficult" years. Wind power in combination with storage technologies see an increase

in investment (see Fig. C.1), most pronounced in the 1985, the year with the highest optimal costs.

Fig. C.2 shows that the different optimal power systems are also mostly driven by (onshore) wind power, as in the 95% reduction case. Additionally, it depicts a strengthening of battery and hydrogen storage, which previously did not appear in the optimal solutions, although their shares of generation vary throughout the years. Similarly, the optimal share of nuclear generation becomes more volatile and decreases not only in the robust allocations (which are characterised by higher renewable investment by design), but also in the 41-year optimisation. As nuclear power has higher variable costs than renewable generators, the existing capacities are not fully used, indicating the competitiveness of newly installed renewable capacities.

As with the 95% emission reduction, significant investment in both onshore wind power and solar power is necessary for a power system

that can withstand weather variability. Figs. C.1 and C.3 show additionally the need for investment in hydrogen storage and additional transmission capacities. In comparison to optimal solutions for different weather years, additional investment in wind power and hydrogen strengthens robustness against weather variability (see Fig. C.4). The full decarbonisation increases the value of offshore wind power which did not feature prominently in the robust solution under 95% emission reduction (see Fig. 5). Finally, it should be noted that there is a significant amount of investment flexibility for onshore wind, offshore wind and solar among the robust solutions for all 41 weather years.

References

Administration, U.E.I., 2022. International - U.S. energy information administration (EIA). https://www.eia.gov/international/data/world/electricity.

Barber, B.C., Dobkin, D.P., Huhdanpaa, H., 1996. The quickhull algorithm for convex hulls. ACM Trans. Math. Software 22 (4), 469–483. http://dx.doi.org/10.1145/235815.235821.

Benth, F.E., Benth, J.Š., Koekebakker, S., 2008. Stochastic Modeling of Electricity and Related Markets. In: Advanced Series on Statistical Science & Applied Probability, 11, World Scientific, http://dx.doi.org/10.1142/6811.

Bloomfield, H., Brayshaw, D., Charlton-Perez, A., 2020. ERA5 derived time series of European country-aggregate electricity demand, wind power generation and solar power generation: Hourly data from 1979–2019. http://dx.doi.org/10.17864/1947.272.

Bloomfield, H.C., Brayshaw, D.J., Shaffrey, L.C., Coker, P.J., Thornton, H.E., 2016. Quantifying the increasing sensitivity of power systems to climate variability. Environ. Res. Lett. 11 (12), 124025. http://dx.doi.org/10.1088/1748-9326/11/12/124025.

Bloomfield, H.C., Brayshaw, D.J., Troccoli, A., Goodess, C.M., De Felice, M., Dubus, L., Bett, P.E., Saint-Drenan, Y.M., 2021. Quantifying the sensitivity of European power systems to energy scenarios and climate change projections. Renew. Energy 164, 1062–1075. http://dx.doi.org/10.1016/j.renene.2020.09.125.

Boyd, S.P., Vandenberghe, L., 2004. Convex Optimization. Cambridge University Press, Cambridge, UK.

Brown, T., Hörsch, J., Schlachtberger, D., 2018. PyPSA: Python for power system analysis. J. Open Res. Softw. 6 (1), 4. http://dx.doi.org/10.5334/jors.188.

Collins, S., Deane, P., Ó Gallachóir, B., Pfenninger, S., Staffell, I., 2018. Impacts of inter-annual wind and solar variations on the European power system. Joule 2 (10), 2076–2090. http://dx.doi.org/10.1016/j.joule.2018.06.020.

Craig, M.T., Wohland, J., Stoop, L.P., Kies, A., Pickering, B., Bloomfield, H.C., Browell, J., De Felice, M., Dent, C.J., Deroubaix, A., Frischmuth, F., Gonzalez, P.L.M., Grochowicz, A., Gruber, K., Härtel, P., Kittel, M., Kotzur, L., Labuhn, I., Lundquist, J.K., Pflugradt, N., van der Wiel, K., Zeyringer, M., Brayshaw, D.J., 2022. Overcoming the disconnect between energy system and climate modeling. Joule http://dx.doi.org/10.1016/j.joule.2022.05.010.

DeCarolis, J.F., 2011. Using modeling to generate alternatives (MGA) to expand our thinking on energy futures. Energy Econ. 33 (2), 145–152. http://dx.doi.org/10.1016/j.eneco.2010.05.002.

Dowling, J.A., Rinaldi, K.Z., Ruggles, T.H., Davis, S.J., Yuan, M., Tong, F., Lewis, N.S., Caldeira, K., 2020. Role of long-duration energy storage in variable renewable electricity systems. Joule 4 (9), 1907–1928. http://dx.doi.org/10.1016/j.joule.2020.07.007.

ENTSO-E, 2022. Power statistics. https://www.entsoe.eu/data/power-stats/.

Frysztacki, M.M., Hagenmeyer, V., Brown, T., 2022. Inverse methods: How feasible are spatially low-resolved capacity expansion modeling results when dis-aggregated at high resolution?. http://dx.doi.org/10.48550/arXiv.2209.02364, arXiv:2209.02364, URL http://arxiv.org/abs/2209.02364.

Frysztacki, M.M., Hörsch, J., Hagenmeyer, V., Brown, T., 2021. The strong effect of network resolution on electricity system models with high shares of wind and solar. Appl. Energy 291, 116726. http://dx.doi.org/10.1016/j.apenergy.2021.116726.

Gurobi Optimization, L., 2022. Gurobi optimizer reference manual.

Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Rozum, I., Schepers, D., Simmons, A., Soci, C., Dee, D., Thépaut, J.-N., 2018. ERA5 hourly data on single levels from 1959 to present. http://dx.doi.org/10.24381/cds.adbb2d47.

Hilbers, A.P., Brayshaw, D.J., Gandy, A., 2020. Importance subsampling for power system planning under multi-year demand and weather uncertainty. In: 2020 International Conference on Probabilistic Methods Applied To Power Systems (PMAPS). IEEE, Liege, Belgium, pp. 1–6. http://dx.doi.org/10.1109/PMAPS47429.2020.9183591.

Hoffmann, M., Kotzur, L., Stolten, D., Robinius, M., 2020. A review on time series aggregation methods for energy system models. Energies 13 (3), 641. http://dx.doi.org/10.3390/en13030641.

Hofmann, F., Hampp, J., Neumann, F., Brown, T., Hörsch, J., 2021. Atlite: A lightweight python package for calculating renewable power potentials and time series. J. Open Source Softw. 6 (62), 3294. http://dx.doi.org/10.21105/joss.03294.

Hörsch, J., Hofmann, F., Schlachtberger, D., Brown, T., 2018. PyPSA-Eur: An open optimisation model of the European transmission system. Energy Strategy Rev. 22, 207–215. http://dx.doi.org/10.1016/j.esr.2018.08.012.

Kotzur, L., Nolting, L., Hoffmann, M., Groß, T., Smolenko, A., Priesmann, J., Büsing, H., Beer, R., Kullmann, F., Singh, B., Praktiknjo, A., Stolten, D., Robinius, M., 2021. A modeler's guide to handle complexity in energy systems optimization. Adv. Appl. Energy 4, 100063. http://dx.doi.org/10.1016/j.adapen.2021.100063.

Kozarcanin, S., Liu, H., Andresen, G.B., 2019. 21St century climate change impacts on key properties of a large-scale renewable-based electricity system. Joule 3 (4), 992–1005. http://dx.doi.org/10.1016/j.joule.2019.02.001.

Liu, H., Andresen, G.B., Brown, T., Greiner, M., 2019. A high-resolution hydro power time-series model for energy systems analysis: Validated with Chinese hydro reservoirs. MethodsX 6, 1370–1378. http://dx.doi.org/10.1016/j.mex.2019.05.024.

Lombardi, F., Pickering, B., Colombo, E., Pfenninger, S., 2020. Policy decision support for renewables deployment through spatially explicit practically optimal alternatives. Joule 4 (10), 2185–2207. http://dx.doi.org/10.1016/j.joule.2020.08.002.

Maggioni, F., Potra, F.A., Bertocchi, M., 2017. A scenario-based framework for supply planning under uncertainty: Stochastic programming versus robust optimization approaches. Comput. Manag. Sci. 14 (1), 5–44. http://dx.doi.org/10.1007/s10287-016-0272-3.

Montel, M., 2022. Python-Holidays 0.12. https://github.com/dr-prodigy/python-holidays.

Moret, S., Bierlaire, M., Maréchal, F., 2016. Robust optimization for strategic energy planning. Informatica 27 (3), 625–648. http://dx.doi.org/10.15388/Informatica.2016.103.

Neumann, F., Brown, T., 2021. The near-optimal feasible space of a renewable power system model. Electr. Power Syst. Res. 190, 106690. http://dx.doi.org/10.1016/j.epsr.2020.106690.

Pedersen, T.T., Victoria, M., Rasmussen, M.G., Andresen, G.B., 2021. Modeling all alternative solutions for highly renewable energy systems. Energy 234, 121294. http://dx.doi.org/10.1016/j.energy.2021.121294.

Pfenninger, S., 2017. Dealing with multiple decades of hourly wind and PV time series in energy models: A comparison of methods to reduce time resolution and the planning implications of inter-annual variability. Appl. Energy 197, 1–13. http://dx.doi.org/10.1016/j.apenergy.2017.03.051.

Pfenninger, S., Staffell, I., 2016. Long-term patterns of European PV output using 30 years of validated hourly reanalysis and satellite data. Energy 114, 1251–1265. http://dx.doi.org/10.1016/j.energy.2016.08.060.

Pickering, B., Lombardi, F., Pfenninger, S., 2022. Diversity of options to eliminate fossil fuels and reach carbon neutrality across the entire European energy system. Joule 6 (6), 1253–1276. http://dx.doi.org/10.1016/j.joule.2022.05.009.

Price, J., Keppo, I., 2017. Modelling to generate alternatives: A technique to explore uncertainty in energy-environment-economy models. Appl. Energy 195, 356–369. http://dx.doi.org/10.1016/j.apenergy.2017.03.065.

Price, J., Zeyringer, M., 2022. highRES-Europe: The high spatial and temporal resolution electricity system model for Europe. SoftwareX 17, 101003. http://dx.doi.org/10.1016/j.softx.2022.101003.

Rehfeldt, D., Hobbie, H., Schönheit, D., Koch, T., Möst, D., Gleixner, A., 2022. A massively parallel interior-point solver for LPs with generalized arrowhead structure, and applications to energy system models. European J. Oper. Res. 296 (1), 60–71. http://dx.doi.org/10.1016/j.ejor.2021.06.063.

Ringkjøb, H.-K., Haugan, P.M., Solbrekke, I.M., 2018. A review of modelling tools for energy and electricity systems with large shares of variable renewables. Renew. Sustain. Energy Rev. 96, 440–459. http://dx.doi.org/10.1016/j.rser.2018.08.002.

Ruhnau, O., Qvist, S., 2022. Storage requirements in a 100% renewable electricity system: Extreme events and inter-annual variability. Environ. Res. Lett. http://dx.doi.org/10.1088/1748-9326/ac4dc8.

Schlachtberger, D.P., Brown, T., Schramm, S., Greiner, M., 2017. The benefits of cooperation in a highly renewable European electricity network. Energy 134, 469–481. http://dx.doi.org/10.1016/j.energy.2017.06.004.

Schlott, M., Kies, A., Brown, T., Schramm, S., Greiner, M., 2018. The impact of climate change on a cost-optimal highly renewable European electricity network. Appl. Energy 230, 1645–1659. http://dx.doi.org/10.1016/j.apenergy.2018.09.084.

Simoes, S.G., Amorim, F., Siggini, G., Sessa, V., Saint-Drenan, Y.-M., Carvalho, S., Mraihi, H., Assoumou, E., 2021. Climate proofing the renewable electricity deployment in Europe - Introducing climate variability in large energy systems models. Energy Strategy Rev. 35, 100657. http://dx.doi.org/10.1016/j.esr.2021.100657.

Spinoni, J., Vogt, J., Barbosa, P., 2015. European degree-day climatologies and trends for the period 1951–2011. Int. J. Climatol. 35 (1), 25–36. http://dx.doi.org/10.1002/joc.3959.

Staffell, I., Pfenninger, S., 2018. The increasing impact of weather on electricity supply and demand. Energy 145, 65–78. http://dx.doi.org/10.1016/j.energy.2017.12.051.

Stirling, A., 2010. Keep it complex. Nature 468 (7327), 1029–1031. http://dx.doi.org/10.1038/4681029a.

Swissgrid, 2022. Production and consumption. https://www.swissgrid.ch/en/home/operation/grid-data/generation.html.

Tóth, C.D., Goodman, J.E., O'Rourke, J., 2017. Handbook of Discrete and Computational Geometry. CRC Press, http://dx.doi.org/10.1201/9781315119601.

Tröndle, T., Lilliestam, J., Marelli, S., Pfenninger, S., 2020. Trade-offs between geographic scale, cost, and infrastructure requirements for fully renewable electricity in Europe. Joule 4 (9), 1929–1948. http://dx.doi.org/10.1016/j.joule.2020.07.018.

Trutnevyte, E., 2016. Does cost optimization approximate the real-world energy transition? Energy 106, 182–193. http://dx.doi.org/10.1016/j.energy.2016.03.038.

van der Wiel, K., Stoop, L.P., van Zuijlen, B.R.H., Blackport, R., van den Broek, M.A., Selten, F.M., 2019. Meteorological conditions leading to extreme low variable renewable energy production and extreme high energy shortfall. Renew. Sustain. Energy Rev. 111, 261–275. http://dx.doi.org/10.1016/j.rser.2019.04.065.

Wohland, J., Reyers, M., Weber, J., Witthaut, D., 2017. More homogeneous wind conditions under strong climate change decrease the potential for inter-state balancing of electricity in Europe. Earth Syst. Dyn. 8 (4), 1047–1060. http://dx.doi.org/10.5194/esd-8-1047-2017.

Zeyringer, M., Price, J., Fais, B., Li, P.-H., Sharp, E., 2018. Designing low-carbon power systems for great britain in 2050 that are robust to the spatiotemporal and inter-annual variability of weather. Nat. Energy 3 (5), 395–403. http://dx.doi.org/10.1038/s41560-018-0128-x.