SEEDS project - Deliverable 1.1

Empirically derived methods to increase the computational efficiency of modelling to generate alternatives in energy system analysis (TurboSPORES)

ETH Zurich / TU Delft

F. Lombardi, f.lombardi@tudelft.nl

October 25th, 2021

"TurboSPORES" is the outcome of the research activities carried out within task 1.1 of the SEEDS project. As the name suggests, TurboSPORES is a computational improvement of the existing SPORES approach developed by ETH Zurich to generate, out of energy system optimisation models, a wide range of spatially-explicit, equally-feasible alternatives for renewable capacity deployment. The reason why SPORES need a computatinal upgrade is for the method to be applicable to models of any size, such as for instance the fully-sector-coupled Portuguese energy system model that we plan to use in the SEEDS project.

An expanded version of the concepts and results summarised here is planned to be turned into a peer-reviewed Journal publication. The code upgrades developed to generate "TurboSPORES" are already publicly available on GitHub as a development branch of the open-source energy modelling framework Calliope[1], and will be fully integrated in the new release of the latter.

[1] https://github.com/FLomb/calliope/tree/0.6.6_spores_customisation

1. Minimal background

There's a good agreement about the fact that, to achieve a rapid and deep decarbonisation of energy systems, we need to deploy a lot more new renewable capacity. Yet, how much capacity of each technology we should deploy, and where, remain open questions. To shed light on such questions, we typically use energy system optimisation models, which allow to identify the one system configuration that would make decarbonisation possible at the lowest cost for society.

Unfortunately, cost-optimality has little real-life meaning. Many factors other than cost are at stake in a real-life decision process about infrastructure deployment, and most of them cannot be parametrised into a mathematical problem. Not only, cost itself, in a model, is affected by the uncertainty around all the assumptions about future technology cost. As such, it is silly to fixate on what is the minimum-cost configuration whereas some of the apparently-more-costly options might be just as cheap, or even cheaper, after such an uncertainty is realised in practice[2].

Based on such considerations, it seems clear that it would be much better to explore all the feasible different alternative system configurations in an economic neighbourhood of the cost-optimal one; and methods to do so are known since the 80s in the field of Operations Research under the name of "Modelling to Generate Alternatives", or MGA. The problem is that, as of yet, applications of MGA to energy system models have focussed primarily on generating alternatives in terms of technology mixes. What is instead additionally critical to know is whether there are, for a given technology mix, spatially-distinctive ways of locating capacity - for instance wind farms - at the sub-national scale. This is what the SPORES method has been conceived to do[3].

1.1. The SPORES algorithmic workflow

The algorithmic workflow behind the SPORES method is the following. First, we identify the optimal solution by cost-minimisation, just as a starting point. Then, we assign weights, or scores, to each "location-technology combination"

---

[2] A. Krumm, D. Süsser, P. Blechinger, Modelling social aspects of the energy transition: What is the current representation of social factors in energy models?, Energy, 239, 2022, https://doi.org/10.1016/j.energy.2021.121706.

[3] F. Lombardi, B. Pickering, E. Colombo, S. Pfenninger, Policy Decision Support for Renewables Deployment through Spatially Explicit Practically Optimal Alternatives, Joule, 4, 2020, pp. 2185-2207, https://doi.org/10.1016/j.joule.2020.08.002.

proportional to their spatially-explicit degree of technology deployment. For instance, let's imagine the cost-optimal solution for a simple energy system with 3 regions is represented by the sketch in Figure 1. We would assign a high weight to Wind in Region 1, because it's deployed a lot, and a low weight to Wind in Region 2, where it's little deployed.

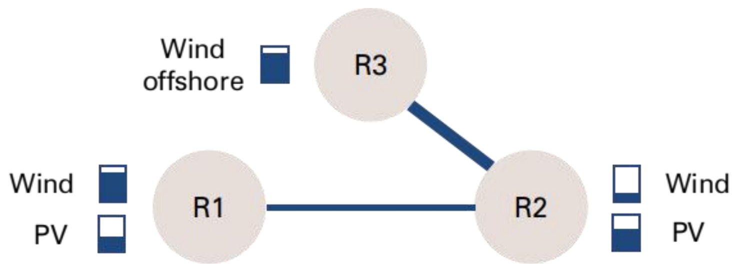

Figure 1. Fictitious cost-optimal solution for a simple energy system.

By assigning weights in such a way, one can reformulate the problem with a new objective function, which now seeks to minimise the appearance of those location-technology capacity variables which have been already explored in the previous iteration, i.e., those that have been assigned a high weight. At the same time, as in conventional MGA, one should impose that the cost cannot exceed by more than a certain amount the identified minimum possible cost. The difference with conventional MGA, however, is that, while the latter just focussed on technological diversity, SPORES make explicit the search for both technologically- and spatially-distinctive solutions at once. Such a search for distinctive solutions can be repeated as many times as wanted, each time updating the weight of location-technology capacity variables, ultimately generating hundreds of SPORES

In addition, the SPORES workflow foresees parallel runs in which a secondary objective is added to the search for spatially distinctive solutions (Figure 2). Such a secondary objective is the overall minimisation (or, alternatively, maximisation) of the capacity of a specific technology. This is needed to make sure that alternatives are originated along different search directions, from several vertexes of the decision space, rater than just along the direction arising from the cost-optimal vertex. The results section shows a clear example of what this means in practice.

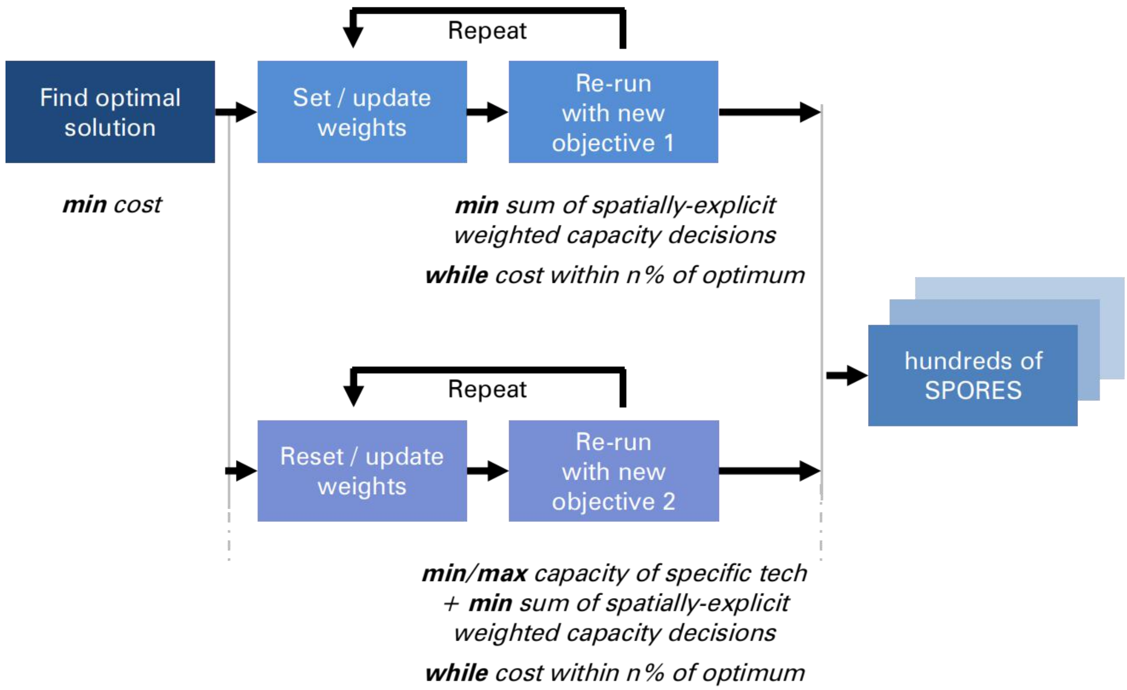

Figure 2. Schematic representation of the SPORES algorithmic workflow.

2. Alternative SPORES formulations for computational improvement

While the SPORES method in the aforementioned form has been successfully applied to a 20-region, power-sector-only model of the Italian energy system, its application to models of larger size requires making the computation as efficient as possible, for instance reducing the generation of "redundant" alternatives - similar-enough to each other to be skipped altogether. To this end, we test here 4 alternative methods to assign a weight to each location-technology capacity variable. Some of them come from existing MGA studies in which they were applied at the technology level only, but here we apply all of them to spatially-explicit capacity variables.

1. The integer weight-assignment method, by which we assign an incremental score of 1 to a given location-technology capacity variable whenever this has been non-zero in the previous iteration.
2. Therelative-deployment method, which is what we used for the Italian study. In this case, the score is proportional to the amount of capacity of technology deployed with respect to the maximum potential for deployment of that technology in the given location.
3. The random method, by which a random number between 0 and 1 is set as the weight, without any particular rationale.

4. The evolving-average method, introduced here for the first time. This is the only method that explicitly keeps a memory of what has happened in previous iterations. In fact, it keeps track and continuously updates the average capacity deployed for each technology in each region, and, the closer to such an average is the capacity deployed by the subsequent iteration, the higher the assigned weight.

3. Results

We can now move forward to assessing how each of these methods performs. Figure 3 shows the first 50 SPORES arising from the basic run of the workflow across all the considered weight-assignment methods. In fact, focussing on just the first run of SPORES, excluding parallel runs, allows to see clearly the effect of each weight-assignment method, without the noise brought about by the more complex full workflow.

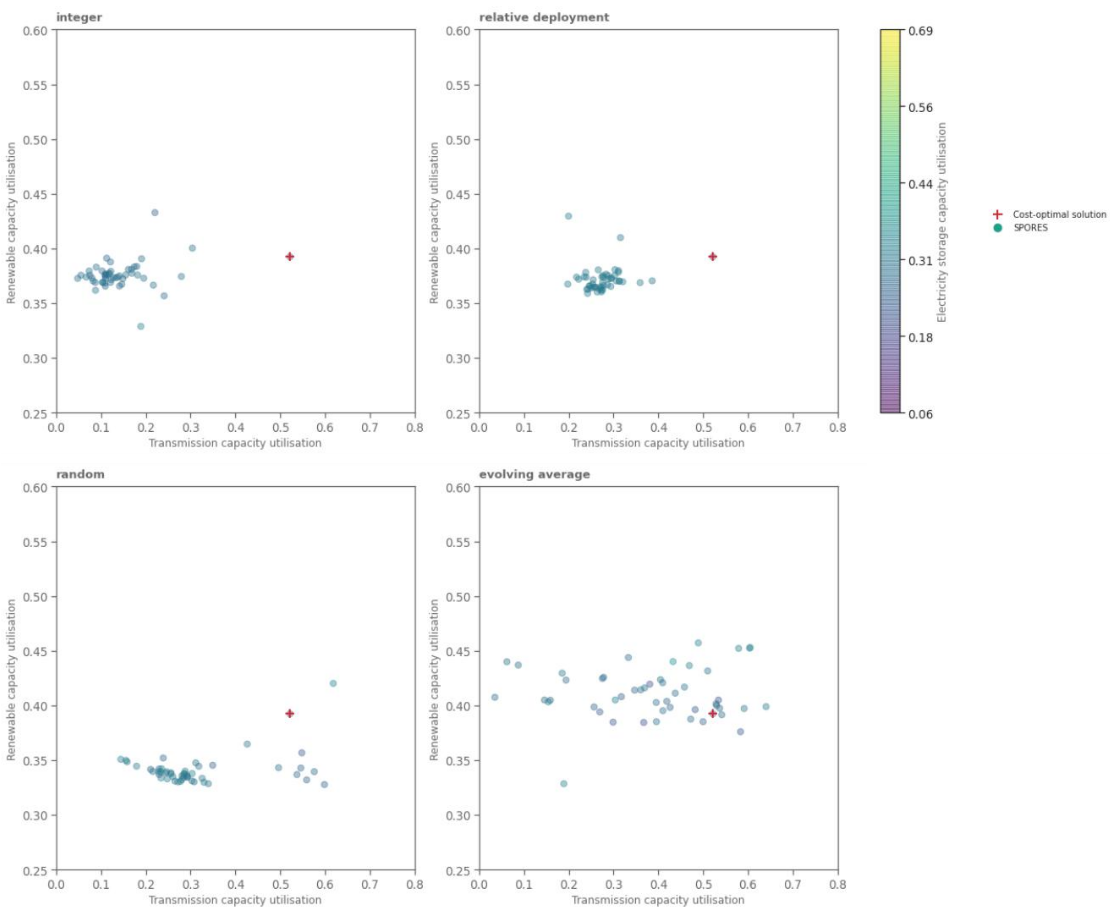

Figure 3. First 50 SPORES arising from the basic run of the workflow, across weight-assignment metods.

This confirms, on the one hand, how important it is to have these paralell runs to explore the solution space from multiple directions and capture the boundaries correctly. On the other hand, it seems that some methods are still better than others in terms of avoiding redundancy. Integer and relative-deployment methods, in particular, seem to produce more redundant solutions, i.e., solutions that almost overlap due to having similar technology mixes. Nonetheless, these almost-overlapping solutions are precisely what SPORES seek, i.e., solutions with similar overall technology mix, but with spatially-distinctive ways of locating capacity at the sub-national space. Therefore, it is not so straightforward to decide what is redundant and what is not. What the results clearly highlight is a trade-off in the computational efficiency of different weight-assignment methods: the better a method is at finding technologically-diverse solutions (such as the evolving-average one), the worse it is at generating spatial diversity around a roughly fixed technology mix, and vice versa.

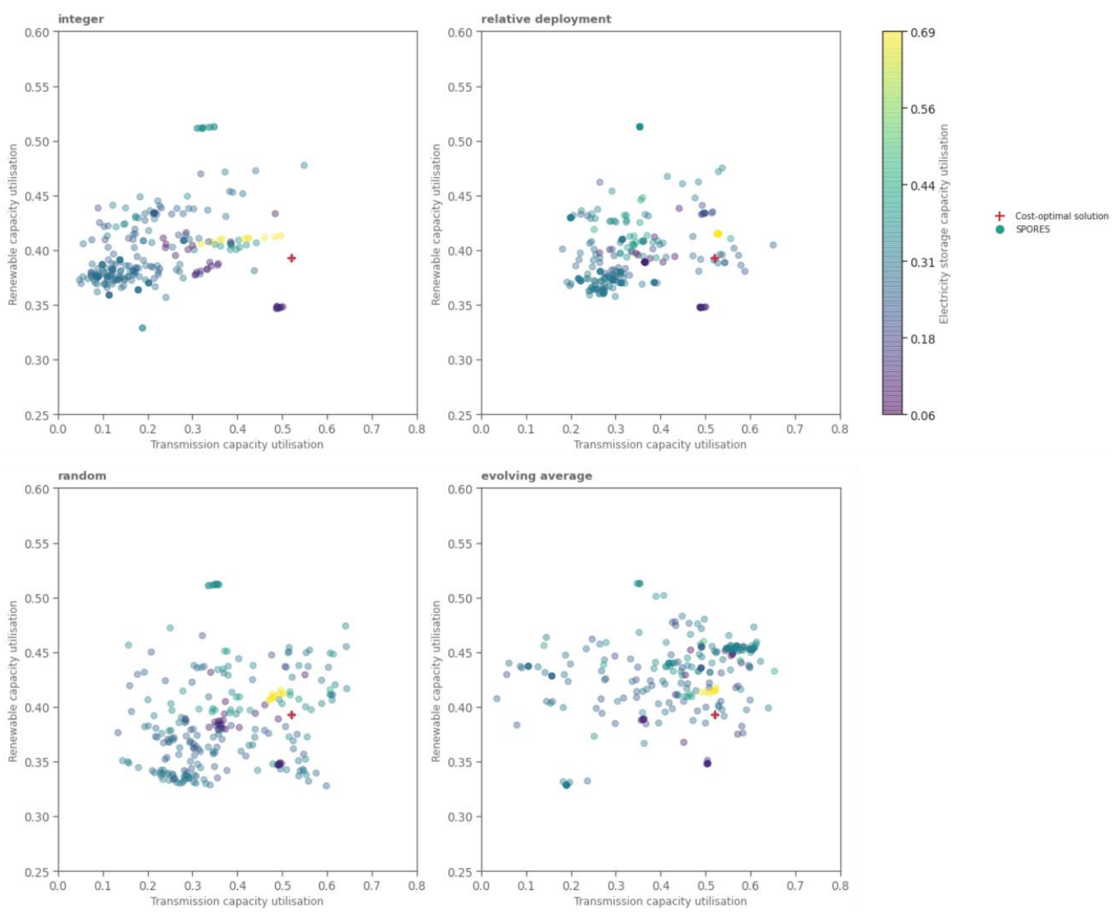

Figure 4. Full set of SPORES arising from a complete run of the workflow, including parallel runs, across weight-assignment metods.

More precisely, if we count how many SPORES, for each method, satisfy some of the criteria that we had identified in the Italian study as potentially appealing to different real-life stakeholder groups (such as the low concentration of wind farms or the high utilisation of newly-built transmission lines), we see that integer and relative-deployment methods substantially outperform the others (Figure 5). And this is precisely thanks to a higher presence of apparently-redundant solutions, with similar technology mix, but distinctive spatial configurations. A good compromise could be thus combining the pros and cons of different methods into a hybrid one. For instance, one that uses the evolving-average method for the initial, basic run of SPORES to explore quickly many relevant technology mixes, and the integer method to generate spatially-distinctive alternatives from multiple search directions through parallel runs - in which the the aspect of technological diversity is already handled explicitly through the secondary objective.

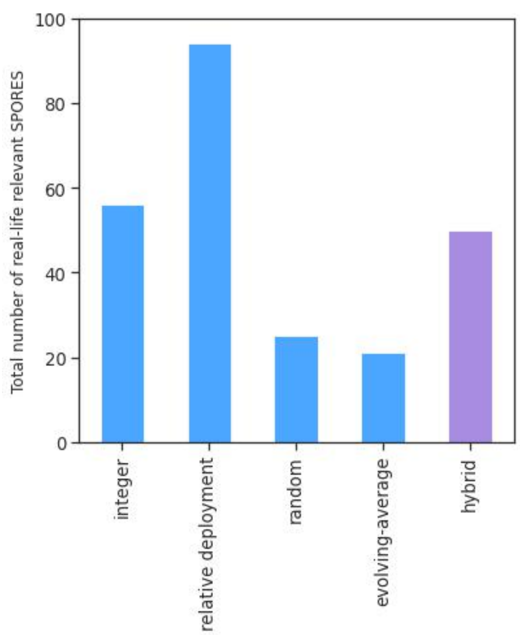

Figure 5. Number of real-life relevant SPORES across weight-assignment methods. The hybrid method combines the evolving-average one for the basic run of the workflow (see Figure 2) and the integer one for parallel runs with a technology-explicit secondary objective.

Conclusions

We can summarise a few key takeaways from this computational experiment:

- First, there is a trade-off between diversity of technology-mix and diversity of spatial configuration in exploring the solution space.

- Second, this also means that no alternative is really redundant. Even a brute-force exploration of all technology mixes, as proposed by a recent pre-print, would likely miss most of the real-life appealing, spatially-distinctive options around the found technology mixes. In short: no method can capture everything, simply because the set of potentially relevant options is not a finite one.
- Third, all this considered, the ideal strategy would be to have an initial as-exhaustive-as-possible exploration of the decision space, followed by iteration with stakeholders who could indicate themselves those technology mixes around which they would like to see more spatially-distinctive options.
- Fourth and final, a hybrid workflow might be the most suited for such an initial exploration. The evolving-average method alone could still be an acceptable first-guess option when computational resources are limited and heavy parallelisation is not possible.

Coming back to the end goal of making SPORES applicable to models of large size (such as a fully-sector coupled Portugal), we can say that, with a hybrid workflow and the heavy parallelisation ensured by our high-performance computing facilities, we can now easily achieve such a goal in a matter of hours. But iterations with stakeholders - as foreseen by the SEEDS project - will be essential to refine the set of alternatives in a way that better covers real-life trade-offs.