Learning Terrain Suitability for Rail using Portugal's Existing Network

Pedro Afonso
PARIS LODRON UNIVERSITY OF SALZBURG

June 30, 2026

Abstract

Railway alignment has its challenges, which is why there has been so much effort put into its automatization. These processes usually include a comprehensive geographic model, with a digital terrain model as one of its many layers.

This paper presents a technique to process the terrain data in a way that takes the directional component of its suitability into account. A Random Forest Model was trained using the existent rail network of Portugal as a mask for positive samples within a Digital Elevation Model and its directional derivatives. This allowed the construction of 8 directional cost grids, which then made it possible to apply a Dijkstra algorithm to test the model with possible paths.

The tests were made in Western Portugal, more specifically the connection between Porto and Lisbon, with the results showing clear signs of the model understanding real terrain suitability.

Keywords: Railway design, Machine Learning, Digital Elevation Models

1 Introduction

With every public construction, there are costs associated with feasibility studies and planning, which often end up representing significant percentages of the total cost of the project. When it comes to railways, these studies include defining the actual path of the railway, which is one of the components that has concentrated the most efforts aiming at its automatization, using so-called railway alignment optimization.

These algorithms come in all shapes and forms, implementing varied techniques, from particle swarm optimization [1] to density-based clustering [2], and usually relying on comprehensive geographic models, where a digital terrain model is simply one of the many layers included. These, however, fail to account for the fact that terrain suitability has a directional nature that cannot be represented with a simple raster (even with its derivatives), i.e the fact that the hillsides of a valley are optimal for lines along the valley, but inadequate for a line that would cross it transversely. For that reason, the author suggests a simple approach (using Portugal's existent rail network) to process the terrain data in a way that takes the directional component into account.

The results of this terrain processing should ideally be used alongside other datasets to create the aforementioned comprehensive geographic model, but they can also be used on their own for identifying rudimentary railway corridors, which is precisely what the author did due to computational (and time) constraints.

2 Datasets

With high-resolution Digital Elevation Models (DEMs) being widely available for Portugal, the arduous task was obtaining the railway lines of the existing network. The author decided to resort to OpenStreetMap (OSM), where a community of volunteers has mapped all the existing and defunct lines of the country, which can be see in (almost[1]) its entirety in Figure 1 on the following page. After careful consideration, the author decided to opt for using solely the main existing, represented in Figure 2, which should be more optimized through modern techniques and for modern trains.

[1] In fact, the map omits lines that have been adapted to the circulation of light rail, as well as sections of lines that have been modernized while keeping most of the path intact.

Furthermore, (for reasons that will become more obvious later) the author decided to omit the sections that were tunnels or bridges. The final dataset, with the lines in a `LineString` format, were obtained using the following Overpass Turbo query:

```
[out:json];

area["ISO3166-1"="PT"]
  [admin_level=2]->.searchArea;

way["railway"="rail"]
  [!"bridge"]
  [!"tunnel"]
  (area.searchArea);

out geom tags;
```

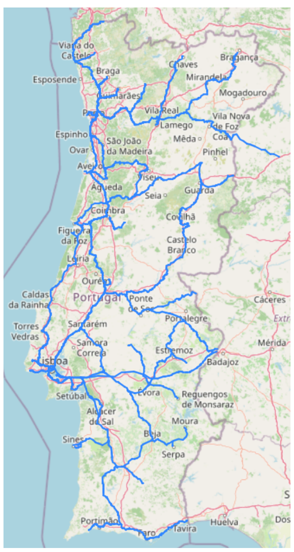

Figure 1: Map of the existing and defunct train lines of Continental Portugal (in blue) overlayed on OSM's standard map layer.

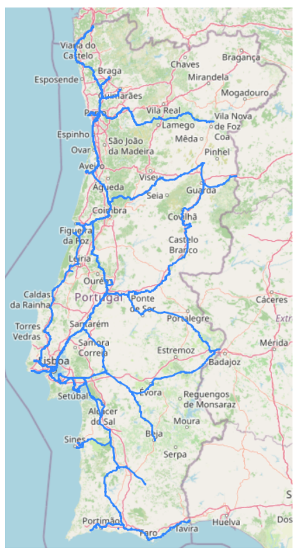

Figure 2: Map of the existing main train lines of Continental Portugal (in blue) overlayed on OSM's standard map layer.

As for the DEM, the author settled with a resolution of 25m, «obtained by the Shuttle Radar Topography Mission by SAR Interferometry» (and kindly made available here), which proved to be good enough for the model to learn the intended behavior, and not so heavy that the author's personal computer could no longer handle its processing without crashing. Figure 3 shows a representation of the DEM used.

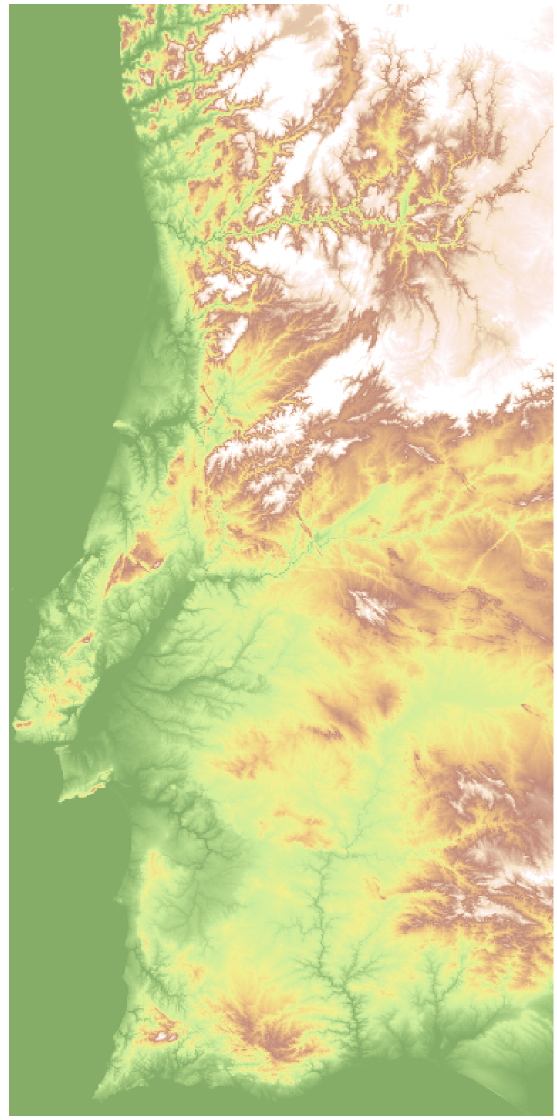

Figure 3: Elevation of Continental Portugal and surrounding areas (Spain and the Atlantic Ocean). The greener the color, the lowest the elevation, with the maximum elevation being ∼2000m.

3 Methodology

3.1 Generating Samples

After loading the railway lines (already missing the tunnel and bridge sections), the author began by creating a 5m buffer of the lines, projecting it to the DEM's Reference System, and then creating a boolean mask of the DEM that was positive for all the cells intersecting the 5m buffer.

Afterwards, the author designed a pipeline that iterated through every cell intersecting a railway line and created 2 positive samples and 6 negative ones, each containing the following features:

- Elevation (directly from the DEM)
- Slope along the railway line (first derivative of the DEM)
- Slope perpendicular to the railway line direction (first derivative of the DEM)
- Curvature along the railway line (second derivative of the DEM)
- Curvature perpendicular to the railway line direction (second derivative of the DEM)
- Cross-term of the second derivative of the DEM
- Slope Magnitude (norm of the first derivative vector)

With each sample basically consisting of a "(location, direction)" tuple, the 2 positive samples corresponded to the intersected cell in both the direction of the railway line and also the inverse direction[2]. The 6 negative samples belonged to 2 different sets: the first set also corresponded to the intersected cell, but, this time, in the 2 directions perpendicular to the railway line; the four remaining negative samples corresponded to 2 cells chosen at random (from within a window of 500m by 500m centered around the intersected cell), with each using the direction of the railway line and its inverse direction.

By now, it should be clear why the tunnels and bridges were omitted: otherwise, the model would not learn to avoid crossing valleys and going through mountains.

3.2 Training the Model

As mentioned in Section 1 on page 1, models are usually trained using comprehensive geographic models, which is also why they need more advanced algorithms, such as particle swarm optimization [1]. The aim of this project, however, was to focus on the processing of the terrain data and, thus, the chosen training method was a simple Random Forest algorithm with 250 estimators and balanced class weights (to account for the negatives samples being three times as much as the positive ones).

Now, if one would start drawing a railway line from a pixel and was able to connect their first pixel to any of its 8 neighboring pixels, there would be 8 possible directions for movement, but the model should assign different costs for each of this directions by considering the features mentioned in Section 3.1.

That is why, after training, the model was used to calculate 8 «Directional Cost Grids» (one for each direction) of the entire DEM raster. This was the most computationally intensive part, having taken more than 1 hour per direction to finish. In the end, the results were saved as both an 8-layer TIF (for visualization purposes) and as a NumPy array, to avoid having to recalculate everything multiple times.

3.3 Finding the Corridor

Unfortunately, with the setup described in Section 3.2, one cannot simply use a simple minimum-cost algorithm, such as the `MCP_Geometric` function from the `skimage.graph`. Instead, the author had to define a Dijkstra algorithm from scratch, taking into account the anisotropic nature of the cost.

Since the start and end point coordinates were defined using latitude and longitude, it was also required for these to be transformed into the Projected coordinate system of the DEM.

4 Results and Discussion

Portugal's first passenger high-speed line will connect Porto to Lisbon, the country's two largest cities. With the project well underway and construction expected to start at the end of the year, the author decided to use it for comparison with the results.

The first attempt was to find the best-fitted path between Porto and Lisbon with no restrictions. The resulting path, as show in Figure 4 on the next page, treads along the coast for most of its length, before going inland through the pass between the Mountains of Candeeiros and Montejunto, and connecting to Lisbon.

This result might appear to be a bit questionable, especially when compared to the existing network, as seen in Figure 5. The current network avoids the coast, with a handful of exceptions to connect to some medium-sized coastal cities, but that does not mean the model learned the wrong behavior; the flat nature of the coast is exactly the type of terrain that is optimal for trains and that the algorithm should detect. The current "avoidance" of the coast is explained by two other factors: low population density and the existence of several protected areas.

[2] It's extremely import to include both directions to force symmetry into the model. It would be sub-optimal if calculating the best path from Porto to Lisbon and from Lisbon to Porto yielded different results.

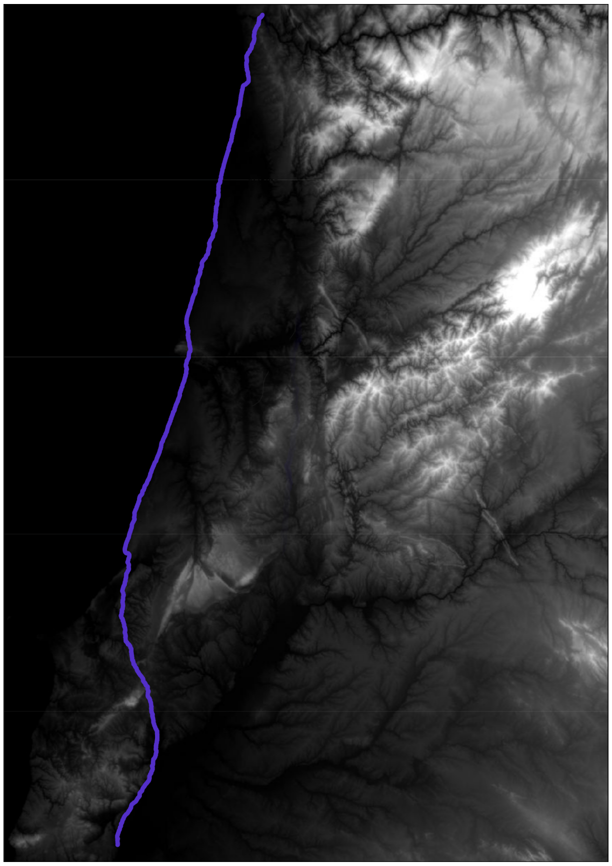

Figure 4: Best path found (in purple) between Porto and Lisbon overlayed on the DEM of the region (the darkest the color, the lowest the elevation).

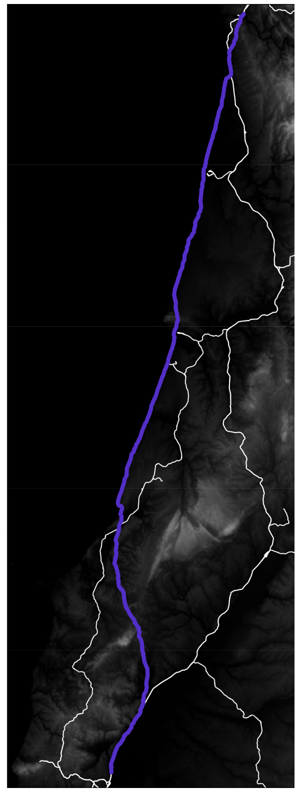

Figure 5: Best path found (in purple) between Porto and Lisbon plotted alongside the existing network used in training (in white).

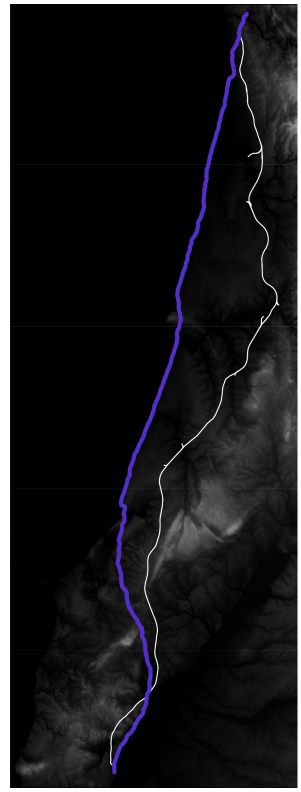

Figure 6: Best path found (in purple) between Porto and Lisbon plotted alongside the future high-speed line (in white).

On the other hand, comparing the results with the high-speed line, as seen in Figure 6, can also provide some insight, namely the fact that both lines go through the pass between the Mountains of Candeeiros and Montejunto, in spite of the high-speed line doing so a bit further north. This further suggests the result is the topographically-optimal path.

Things get more complicated when we force the path to go by Coimbra, home to Portugal's oldest University and a mandatory stop on north-south connection. The coast is now entirely avoided, but, as seen in Figure 7 on the following page, the new path overlays the existing network on long sections of its length (especially in the southern part), hinting at a possible case of overfitting. There is, however, no need to assume the worst. First, and foremost, the "Linha do Norte" (the current line connecting Porto and Lisbon) is the busiest and more important line of the country and, thus has been thoroughly optimized, especially in the southern part, where the hillier terrain forces it through the few available mountain passes. Furthermore, it should be noted that, even the sections where both lines appear identical, there are often differences of the order of magnitude of 500m, which are simply invisible at the scale of the Figure.

5 Conclusion

The trained model yielded some compelling results. As expected, it cannot be used alone to define a railway corridor, hence the need for the comprehensive geographic models (and human supervision), but it appears to have properly learned what attributes make an area suitable for a railway line, and provides an empirical assessment of the suitability.

In the future, it would be interesting to test how this model performs outside the area where it was trained, starting with neighboring Spain, for example, whose similar topography might help with the model transferability. Even more interesting would be to train a model on the entire continent of Europe,

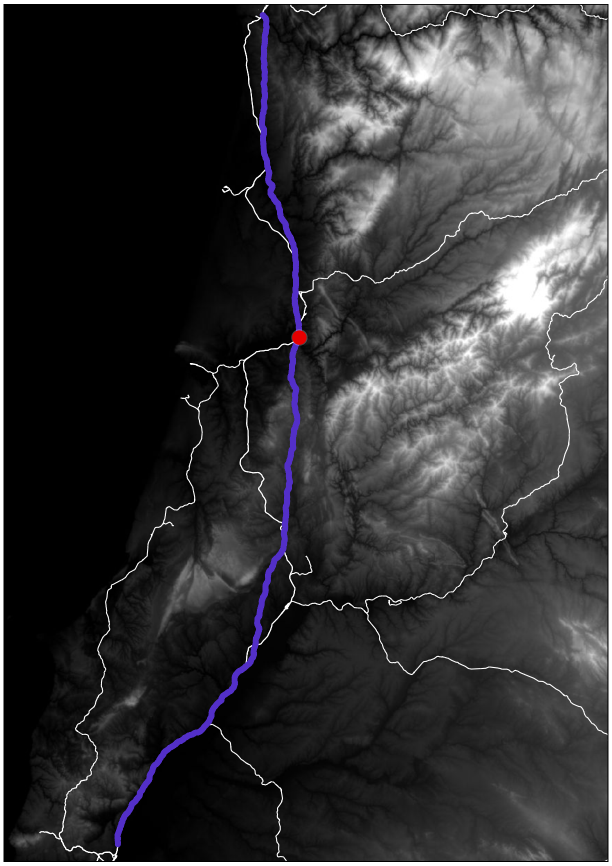

Figure 7: Best path found (in purple) between Porto and Lisbon passing through Coimbra (red dot) overlayed on the DEM of the region (the darkest the color, the lowest the elevation).

whose more varied landscape would make it easier to generalize and apply not only in Europe, but also in other regions of the world.

On the other hand, there is likely merit in training a model with a 10m DEM of the country; this would not only would perfect the results, but also allow to incorporate curvature restrictions, which curb train speed limits, making them especially relevant for high-speed lines.

References

[1] H. Pu, T. Song, P. Schonfeld, W. Li, H. Zhang, J. Hu, X. Peng, and J. Wang, "Mountain railway alignment optimization using stepwise & hybrid particle swarm optimization incorporating genetic operators," *Applied Soft Computing*, vol. 78, pp. 41–57, 2019.

[2] X. Wan, H. Pu, P. Schonfeld, T. Song, W. Li, and L. Peng, "Railway alignment optimization in regions with densely-distributed obstacles based on semantic topological maps," *Integr. Comput. Aided Eng.*, vol. 31, pp. 421–437, July 2024.