Hes·so VALAIS WALLIS

Haute Ecole d'Ingénierie
Hochschule für Ingenieurwissenschaften

BACHELOR THESIS

Dimitri Imfeld

Rolling Stones: Etude de faisabilité d'un modèle de Machine Learning pour l'identification du risque de chute de pierres sur le territoire du Canton du Valais

Informatique et Systèmes de communication (ISC)

ISC-ID-25-7

University of Applied Sciences and Arts Western Switzerland
Faculty of Engineering and Architecture
Haute École d'Ingénierie de Sion

Dimitri Imfeld

Rolling Stones: Etude de faisabilité d'un modèle de Machine Learning pour l'identification du risque de chute de pierres sur le territoire du Canton du Valais

Thesis submitted to the Bachelor degree programme
Informatique et Systèmes de communication (ISC) – Data engineering major
Haute École d'Ingénierie de Sion

Supervisor: **DR. GRÉGORY MERMOUD, PHD**
Co-supervisor: **DR. CÉDRIC TRAVELLETTI, PHD**

Submitted on: July 25, 2025

Abstract

The Service for Natural Hazards of the Canton Valais (SDANA) is responsible for assessing and managing natural hazards, including rockfalls. Rockfalls are a significant natural hazard in mountainous regions, particularly in the Swiss Alps, where steep slopes and geological conditions create an environment favourable to such events.

They wanted, through this project, to explore the potential of machine learning techniques to assist in the creation of rockfall susceptibility maps, with the goal to help identify areas at risk and inform and assist decision-making processes for land use planning and risk management.

The project uses a variety of geospatial data, including digital elevation models (DEMs), geological maps, and hand-made rockfall susceptibility maps used as reference, to train and evaluate different machine learning models.

The results show that machine learning models can effectively learn from the data and produce rockfall susceptibility maps that, while not as precise as expert-made maps, are still accurate enough to identify general areas at risk of rockfalls.

The project demonstrates the potential of using machine learning techniques to assist in rockfall hazard assessment, providing a valuable tool for risk management in the canton of Valais.

Keywords : engineering, data, machine learning, rockfall, natural hazards, rockfall susceptibility mapping

Résumé

Le Service des dangers naturels du Canton du Valais (SDANA) est responsable de l'évaluation et de la gestion des dangers naturels, y compris les chutes de pierres. Les chutes de pierres constituent un danger naturel important dans les régions montagneuses, particulièrement dans les Alpes suisses, où les pentes raides et les conditions géologiques créent un environnement favorable à de tels événements.

Ils souhaitaient avec ce projet explorer le potentiel des techniques d'apprentissage automatique (Machine Learning) pour aider à la création de cartes de susceptibilité de chutes de pierres, dans le but d'identifier les zones à risque et d'informer et d'assister les processus de prise de décision pour l'aménagement du territoire et la gestion des risques.

Le projet utilise une variété de données géospatiales, notamment des modèles numériques d'élévation, des cartes géologiques et des cartes de susceptibilité aux chutes de pierres réalisées manuellement utilisées comme référence, pour entraîner et évaluer différents modèles d'apprentissage automatique.

Les résultats montrent que les modèles d'apprentissage automatique peuvent efficacement apprendre à partir des données et produire des cartes de susceptibilité aux chutes de pierres qui, bien qu'elles ne soient pas aussi précises que les cartes réalisées par des experts, sont néanmoins suffisamment exactes pour identifier les zones générales à risque de chutes de pierres.

Le projet démontre le potentiel de l'utilisation des techniques d'apprentissage automatique pour aider à l'évaluation des dangers de chutes de pierres, fournissant un outil précieux pour la gestion des risques dans le Valais.

Keywords : engineering, data, machine learning, rockfall, natural hazards, rockfall susceptibility mapping

Acknowledgements

I would like to express my thanks to both Dr. Grégory Mermoud and Dr. Cédric Travelletti for their guidance and support throughout this project. Their expertise was invaluable in shaping its direction and ensuring its success.

LLM Usage Notice

This thesis was written with the help of a large language model, which assisted in the writing process for structuring and rewording sections of the text, with the goal of improving clarity. All analysis, reflections and conclusions are mine.

Table of Contents

Chapter 1 – Introduction

Due to its mountainous terrain, the Canton of Valais is particularly prone to rockfalls, which can pose significant risks to people and infrastructure.

These events may also become more frequent due to climate change, especially in alpine areas [1]. This makes it crucial to develop effective methods for assessing and predicting rockfall hazards. Traditional methods of rockfall hazard assessment often rely on expert knowledge and manual mapping, which can be time-consuming.

Currently, these rockfall hazard assessment maps are created manually by engineering experts based on topography, geological features, vegetation, rockfall trajectory simulations and other data, as well as their own knowledge and experience. These maps are commissioned by the cantonal authorities on a case-by-case basis, and are used to inform land use planning and risk management decisions in specific areas. The coverage of the risk maps is therefore limited to a few areas.

The goal of this thesis is not to replace these methods, nor to create a fully automated system; rather, it aims to explore the potential of machine learning techniques to assist in the creation of rockfall susceptibility maps. These maps can help identify areas at risk and inform decision-making processes for land use planning and risk management.
Indeed, one especially interesting aspect of this project is that it could predict rockfall susceptibility on the scale of the whole canton. This would allow for a more comprehensive assessment of rockfall hazards in the region, and could help identify areas that are at risk but have not yet been mapped.

The project uses a variety of geospatial data, including digital elevation models (DEMs), geological maps, as well as the expert hand-made rockfall susceptibility maps as reference, to train and evaluate different machine learning models to predict rockfall susceptibility.

1.1 – Objectives

The main objectives of this project were to:
- Study the state of the art in rockfall hazard assessment and machine learning applications in this field.
- Do a preliminary analysis of the data.
- Develop and evaluate machine learning models to predict rockfall susceptibility.
- Compare the performance of different models and assess their potential for practical applications in rockfall hazard assessment.

1.2 – Rockfall Hazard Assessment

The mapping of rockfall hazards is a crucial aspect of risk management in mountainous regions. It involves identifying areas at risk of rockfalls and assessing the potential impact on people and infrastructure.

To do so, experts typically use a combination of geological, topographical, rockfall trajectory simulations, and meteorological data, along with their own knowledge and experience.

To define the hazard, experts often use a risk assessment matrix (Figure 1), which combines the probability of a rockfall event occurring, with the potential impact intensity that a rockfall event could inflict. The matrix helps to categorize areas into different hazard levels, from low hazard to high hazard, guiding land use planning and risk mitigation strategies.

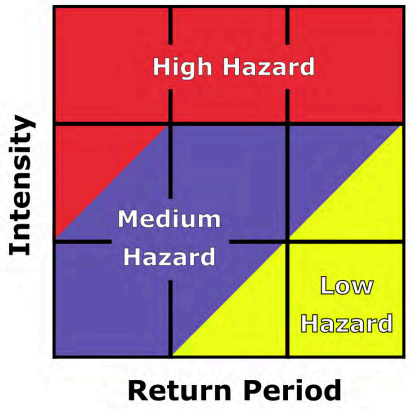

Figure 1 - Standard risk assessment matrix for natural hazards used in Switzerland.

Each study area (Figure 2) possesses one or multiple rockfall source areas, which are the locations where rockfalls originate from. These source areas are typically identified based on geological and topographical features, such as steep slopes, fractured rock formations, and areas with a history of rockfalls. In our study case, these sources are defined solely on slope steepness, as per the data provided by the SDANA.

From these source areas, rockfalls can travel down slopes, potentially impacting people and infrastructure in the surrounding areas. The hazard area is defined as the area at risk of being affected by rockfalls originating from a specific source area. It is typically divided into different zones based on the expected impact intensity, such as high hazard, medium hazard, and low hazard zones.

Also, some areas may be designated as “no hazard”, where the risk of rockfalls is considered negligible.

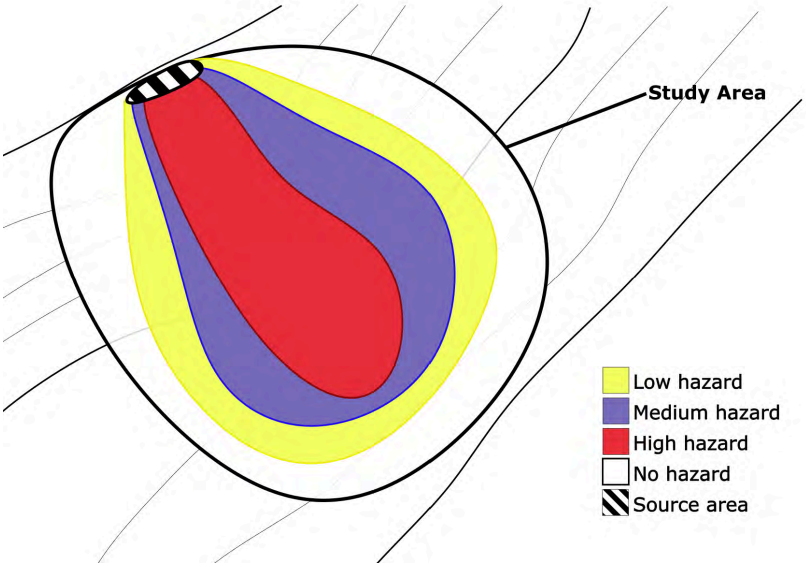

Figure 2 - Illustration of a rockfall hazard area, showing the source area and the different areas of varying hazard.

Chapter 2 – State of the Art

This chapter reviews the current state of research and methodologies in rockfall hazard assessment and the application of machine learning techniques in this field.

A. M. Fanos et al. [2] developped a solution based on both LiDAR and GIS data, using a combination of high-resolution airborne laser scanning (ALS) as well as terrestrial laser scanning (TLS) to capture detailed topographical information. This data is then used to create a digital terrain model (DTM), which is then used to extract features such as slope steepness, aspect, and curvature.

Alongside this, they also a 3D rockfall kynematic model to simulate rockfall trajectories, which gave them a detailed understanding of how rocks would move through the terrain. This also allowed them to extract detailed features from the trajectories, such as bouncing height, velocity, impact location and intensity for each simulated rockfall event.

They then used these features to train multiple machine learning models, including random forests, support vector machines, and k-nearest neighbors (KNN). They found that an hybrid approach combining both random forest and KNN models outperformed the other models, achieving a test accuracy of 0.865 and an AUC of 0.907.

H. Wen et al. [3] studies the use of XGBoost with factor screening and hyperparameter optimization using Grid Search for rockfall probability mapping. Similarly to the previous study, they use a combination of topographical features such as slope, elevation and aspect, but also meteorological data such as precipitation and temperature, as well as geological features such as lythology (rock type).

As a ground truth, they use a sample of 220 historical rockfall events, which they use to train and evaluate their model, alongside a set of 220 randomly selected cells from the study area which are used as negative samples.

Their raster data has a resolution of 90m, which is relatively low. This means that a cell considered as a positive sample may contain the source area, as well as the whole rockfall runout area.

They found that the XGBoost model performed with a test accuracy of 0.6923, improved to 0.7769 after hyperparameter optimization.

Chapter 3 – Methodology

This section, alongside Figure 3, are meant to provide a general overview of the project's pipeline, which includes data loading, model training, to evaluation and prediction. The pipeline is designed to be modular and reusable, allowing for easy experimentation with different models and features.

The following figure illustrates the architecture of the whole pipeline, from data loading, to model training, to evaluation and prediction. It shows how the different components of the pipeline interact with each other, and how data flows through the system.

The models are evaluated using a shared set of metrics, such as precision, recall, F1 and F2 score and AUC. They are also visualized using confusion matrices and ROC curves, which allow for a better understanding of the models' performance.

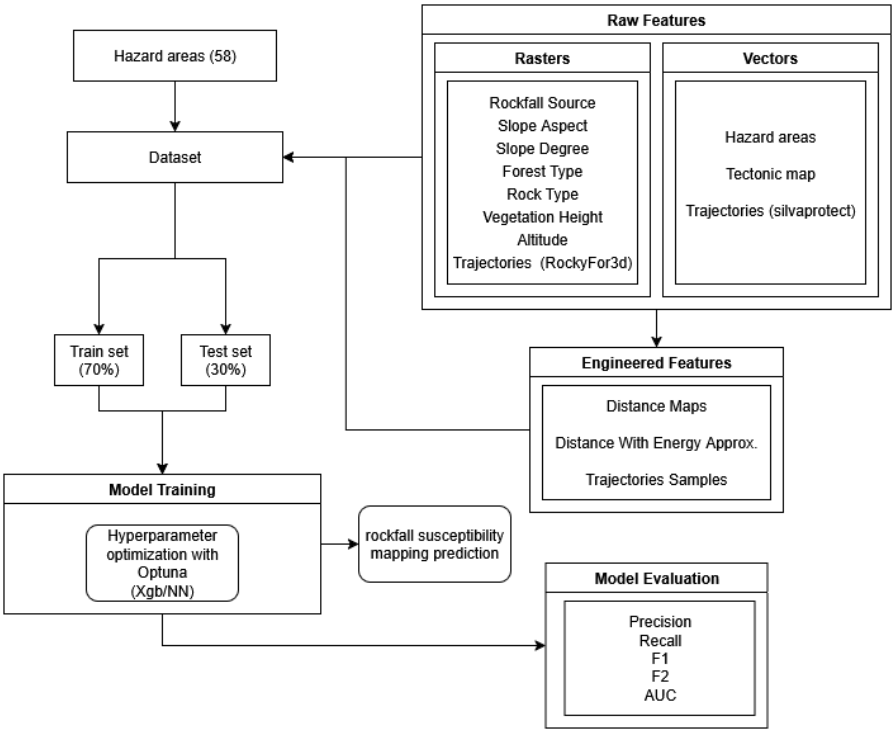

Figure 3 - Architecture of the whole pipeline, from data loading, to model training, to evaluation and prediction. Inspired by similar graph by H. Wen et al. [3]

3.1 – Study Area

The study area for this project is a 16x16 km region located in the Canton of Valais, Switzerland, specifically around the town of Martigny. It's a particularly interesting area for rockfall hazard assessment due to its mountainous terrain, and for the fact that it has been extensively studied in the past, providing a rich dataset of hazard zones for analysis.

This area also contains areas of varying slope, forest coverage, geology, and land use, which allows us to explore a more diverse set of features and their impact on rockfall susceptibility.

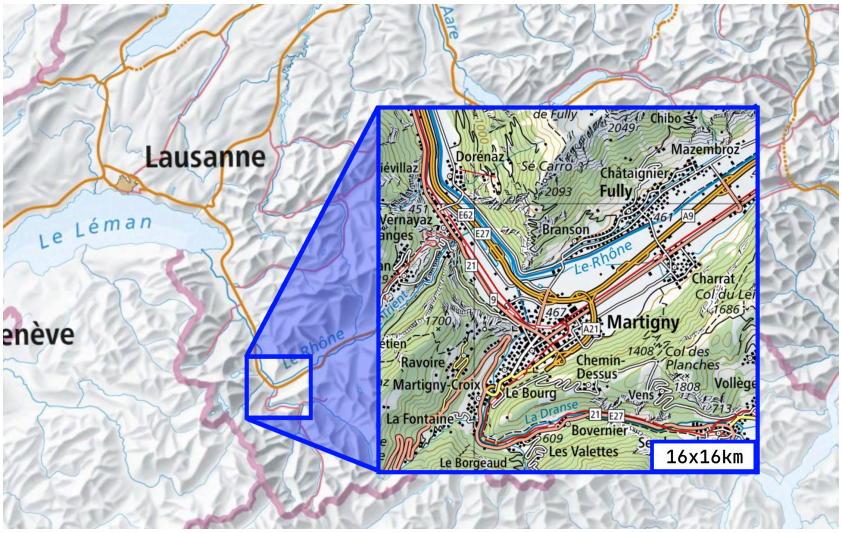

Figure 4 - Map of the study area, the region of Martigny and its surroundings in the Canton of Valais, Switzerland.

3.2 – Data Sources

The data used in this project comes from two main sources: raster data and vector data.

Raster data consists of gridded datasets, where each cell represents a specific geographic location and contains a value for a particular feature. Examples include digital elevation models (DEMs), slope maps, aspect maps, or vegetation coverage. These rasters provide continuous spatial information and are essential for extracting terrain-related features.

All rasters data are in the LV95 coordinate system, which is the Swiss national coordinate system, and have a resolution of 10m. This means that each cell in the raster represents a 10m x 10m area on the ground.

Vector data, on the other hand, represents geographic features as points, lines, or polygons. In this project, vector data includes hazard zones and simulated trajectories from the SilvaProtect project.

The following two tables provide an overview of the vector and raster data used in this project, including their names and descriptions.

Feature Name	Description
Hazard area	Rockfall susceptibility polygons, i.e. what we want to predict. Goes from 0 (no hazard) to 3 (high hazard).
Trajecto SilvaProtect	Simulated rockfall trajectories as vector lines, i.e. the path a rock would take if it were to fall from a certain point, from the Federal Office for the Environment's SilvaProtect-CH project.

Table 1 - Overview of the vector data used in this project.

Feature Name	Description
Aspect	Aspect of the terrain, i.e. the direction the slope is facing, in degrees.
Forest area	Type of forest in that cell.
Rock area	Type of exposed rock in that cell.
Slope	Slope of the terrain in degrees.
Source zone	Source cells of rockfall, i.e. cells that are likely to produce rockfall. Determined by a 55° slope threshold.
SwissAlti3D	Digital elevation model (DEM), which provides the altitude of each cell in meters.
Tectonic map	Map of the different tectonic plates in the region.
Trajecto RockyFor3D	Simulated rockfall trajectories generated using RockyFor3D, a rockfall simulation software. Each pixel contains the number of trajectories that passed through it.
Vegetation height	Vegetation height in meters.

Table 2 - Overview of the raster data used in this project.

The following figures illustrate some of the raster data used in the project, including the slope in the top left, RockyFor3D trajectories in the top right, altitude in the bottom left, and source areas in the bottom right.

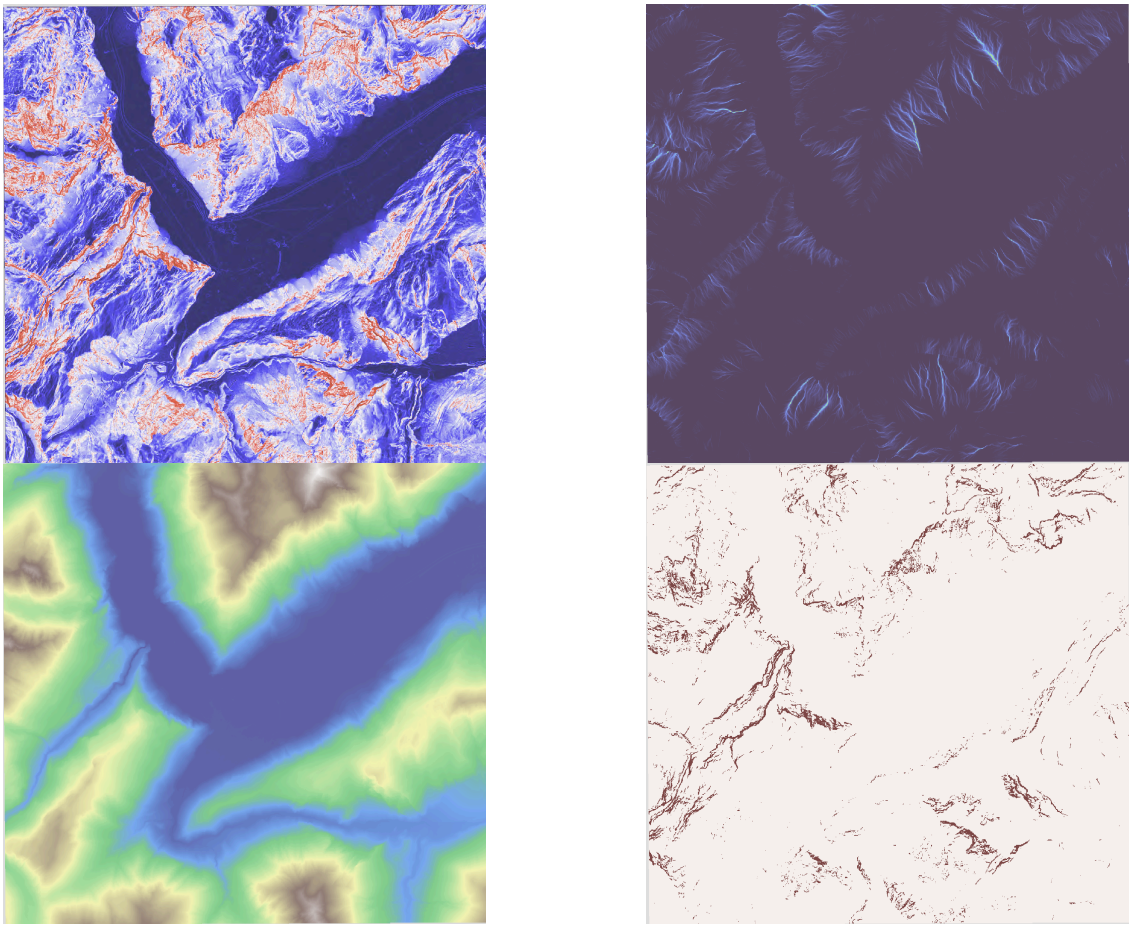

Figure 5 - Example of four raster layers used in the project.

3.3 – Feature Engineering

3.3.1 – Trajectory

Using the SilvaProtect trajectories, we can, for each trajectory, sample features from the raster layers such as the altitude, slope, vegetation coverage and elevation difference along the trajectory. This allows us to capture the context all along the trajectory.

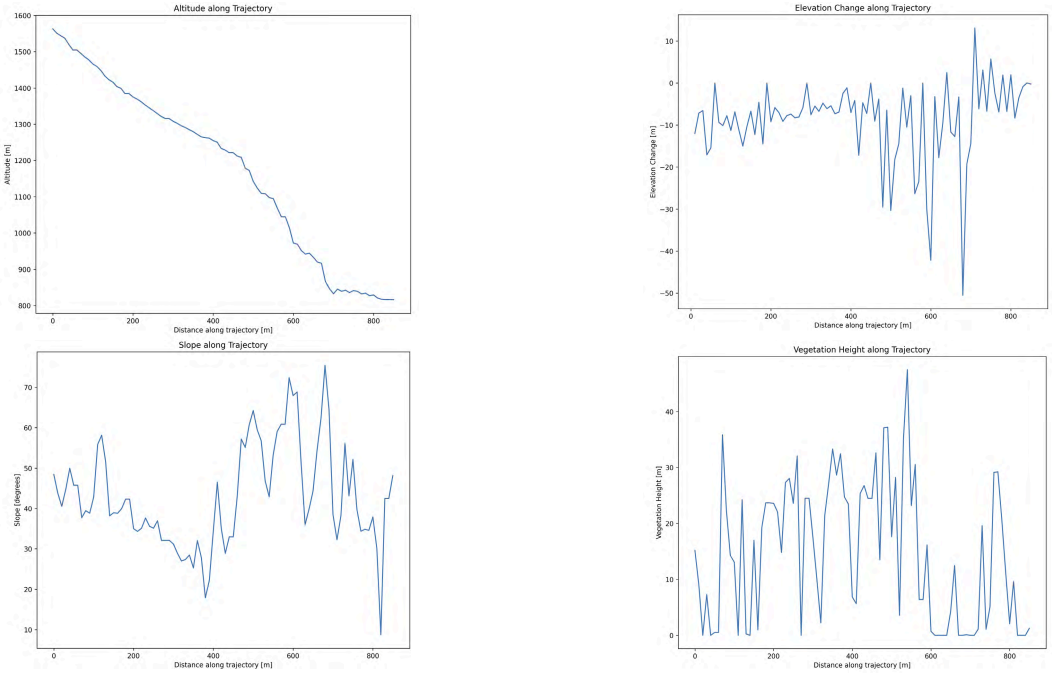

Figure 6 - Example of four feature layers being sampled along a single trajectory, showing the altitude, slope, vegetation coverage and elevation difference along that trajectory.

Some of these features can then be aggregated to create a feature map, which contains the average, maximum, and minimum values of the features extracted from the trajectories that pass through each cell. This allows us to capture the context of the trajectories in a more compact form, while still retaining the information about the trajectories.

This aggregation methodology is interesting, because it allows to give each cell information that is outside of the cell itself, allowing our models to learn patterns from the general context of the trajectories that pass through it. For example, with information such as the maximum drop and rise a rock has experienced while traveling to the cell (i.e. the biggest altitude difference between two cells, respectively negative or positive, it has experienced), a model could potentially learn physical patterns related to rockfall behavior.

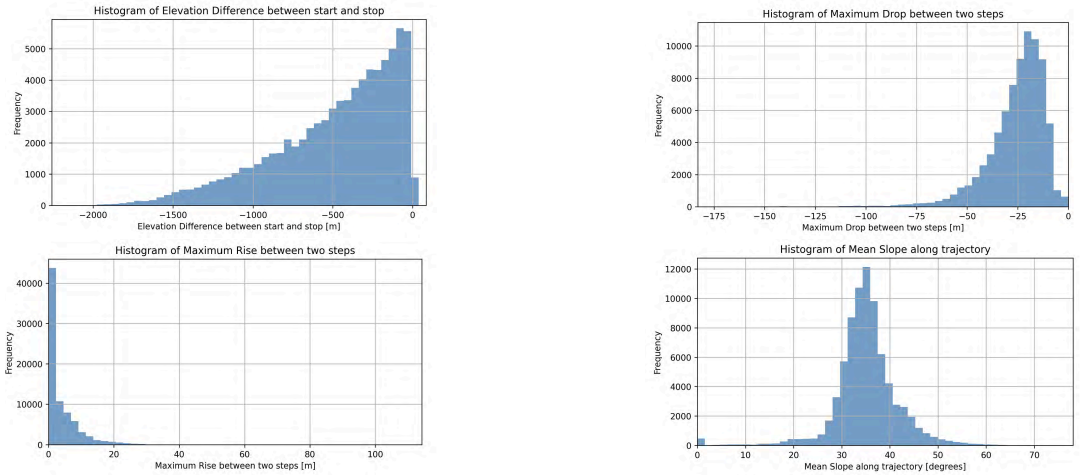

Figure 7 - Example of four features extracted from all trajectories, displayed as histograms. These features are the elevation difference from the start to the end of the trajectory, maximum drop, maximum rise and mean slope along the trajectories.

3.3.2 – Hazard zones

The hazard zones are derived from the “Hazard area” vector dataset, which consists of polygonal geometries representing discrete hazard levels ranging from 0 (no hazard) to 3 (high hazard). While these polygons contain temporal metadata indicating their creation date, they lack explicit zone identification attributes. Additionally, temporal clustering alone proves insufficient due to instances where geographically distinct hazard zones share identical creation dates.

To address this challenge, we implement a two-stage clustering approach: initial grouping by creation date, followed by spatial clustering of geographically close polygons. This gives us our desired hazard zones, each containing one or more polygons of different hazard levels.

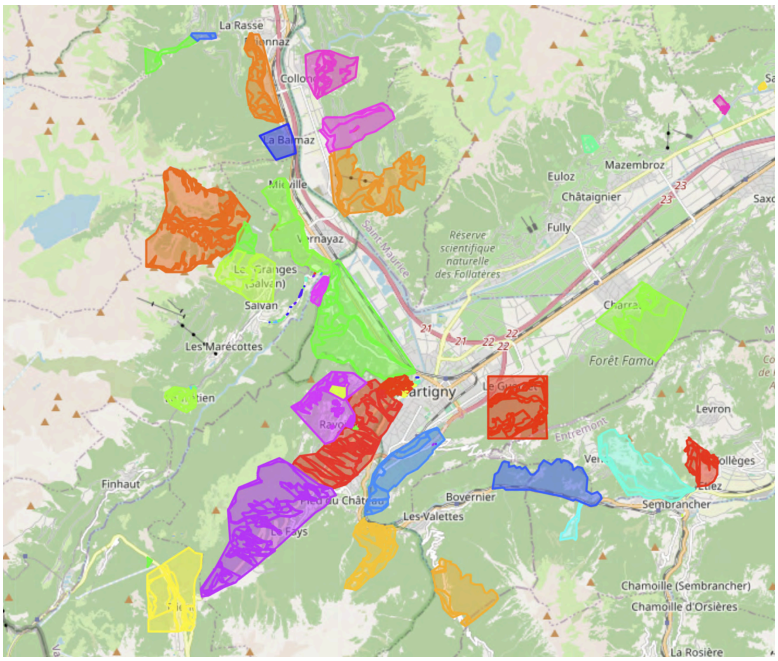

Figure 8 - Grouped hazard zones. Each color represents a different hazard zone, with its different hazard levels.

Since our models work on raster data, we then need to convert these hazard zones into raster data. This is done by rasterizing the polygons, which creates a raster layer where each cell contains the hazard level of the corresponding polygon. The rasterized hazard zone is binarized, with cells of any hazard level (1-3) being classified as positive samples (hazardous), and cells with no hazard (level 0) as negative samples (safe).

Figure 9 - Comparison of the hazard zones as vector polygons (left), and as converted raster cells (right). The rasterized hazard zone is binarized, with cells of any hazard level being considered as positive samples (red), and cells with no hazard being considered as negative samples (white).

After binarization, the hazard zones contain 294′117 cells that are considered as at risk of rockfall, and 97′386 cells that are considered as safe. This means that the dataset is imbalanced, with a ratio of approximately 3:1 between positive and negative samples.

These binarized hazard zones are then used as the ground truth for the models, which will be trained to predict the hazard level of each cell in the training dataset.

3.3.3 – Distance maps

Distance maps are feature layers that represent for each cell the distance to the nearest positive value of a target binary raster layer. Figure 10 shows an example of a distance map, which shows the distance to the nearest rock cell.

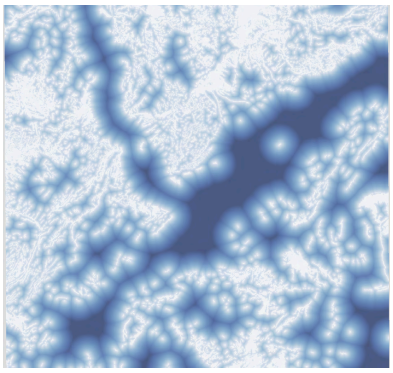

Figure 10 - Example of a distance map, showing the distance to the nearest rock cell.

This same distance map can be computed to compute the nearest source zone. During development, multiple improvements were made to this particular distance map, with the goal of building a feature layer that the baseline model could use.

First, a new type of distance map was created, which only considers target cells that are uphill from the current cell. This was a first step towards creating a distance map that somewhat simulates the behaviour of rockfalls, which typically only travel downhill. This distance map is called the uphill distance map.

Then, from that uphill distance map, a new distance map was created, which scales the distance to the nearest uphill source cell by an energy factor, which is computed as such:

$$\left(\frac{1.0}{1.0 + \text{vertical\_elevation\_diff}^{\text{energy\_dissipation\_exponent}}}\right)$$

That energy factor aims to simulate the behaviour of a falling rock: the steeper the slope, the more energy it has, and the further it can travel. This distance map is called the uphill distance map with energy dissipation.

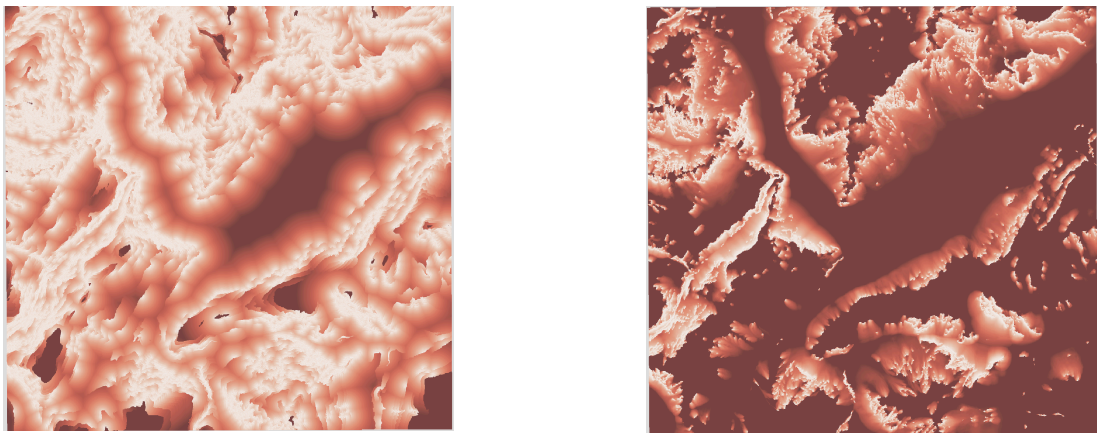

Figure 11 - Comparison of the uphill distance map to the nearest source cell (left), and the uphill distance map with energy dissipation (right). Brighter values indicate a lower distance, which suggests a higher risk of rockfall.

Even though this distance map seems to capture the general patterns of rockfalls, it is important to note that it is still a simplification and in no ways could replace proper simulations.

Indeed, it possesses multiple flaws, such as the fact that identifying the closest uphill source cell does not guarantee rockfall connectivity, as this method cannot detect topographical features (such as ridges or valleys) present in-between that would prevent rocks from actually reaching the target cell.

Also, the energy factor is a purely empirically-determined parameter, which in no sense represents the actual energy dissipation of a rock falling down a slope.

3.4 – Models Architecture

This section describes the different models implemented and evaluated in this study. Each model represents a different approach to rockfall susceptibility prediction, ranging from traditional heuristic methods to modern machine learning techniques.

3.4.1 – Heuristic Model

The heuristic model serves as our baseline, with which we compare the performance of the other models. It uses the uphill distance to source with energy dissipation described above as its sole feature, with the energy dissipation exponent fixed to a value of 1.

Overall, this model simply uses ROC to determine a threshold with which to classify the cells as either hazardous or safe.

3.4.2 – Heuristic + Grid Search Model

This model enhances the basic heuristic approach by optimizing the energy dissipation exponent through grid search.

3.4.3 – XGBoost Model

This model uses the XGBoost algorithm, a popular gradient boosting framework. Its features are selected using the select K best method, which selects the top K features based on their importance scores.

Hyperparameters such as the learning rate, maximum depth of trees, and regularization parameters are optimized using Optuna, a hyperparameter optimization framework.

3.4.4 – XGBoost Dist Model

This model is a variant of the XGBoost model, which only uses the horizontal and vertical distance to the nearest source cell as features. Its goal was to see if XGBoost could learn the same patterns as the heuristic model, but with a more complex model which could potentially discover more complex patterns between the features.

3.4.5 – Neural Network Model

This model implements a multi-layer perceptron (MLP) architecture designed for binary classification by leveraging sklearn's MLPClassifier.

For its architecture, its input layer consists of one neuron per feature, with the number of features depending on the dataset used. The hidden layers consist of two layers, with 32 neurons in the first layer and 64 neurons in the second layer. The output layer has one neuron with a sigmoid activation function, which outputs either 0 or 1, indicating whether the cell is at risk of rockfall or not.

Hyperparameters such as the activation function, solver and learning rate are optimized using Optuna, a hyperparameter optimization framework.

3.5 – Evaluation Metrics

This chapter describes the evaluation metrics used to assess the performance of the machine learning models developed in this project, and the reason why they were chosen in the context of this project.

3.5.1 – Precision and Recall

Precision and recall are fundamental metrics for evaluating classification performance, particularly when the consequences of false positives and false negatives differ.

In our context, a positive prediction indicates that a cell is at risk of rockfall, while a negative prediction indicates that it is not.

Precision evaluates how many of the predicted positive cases are actually correct:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}$$

This metric is crucial when the cost of false positives is high, as it indicates how many of the predicted positive cases are actually correct.

Recall, on the other hand, measures how many of the actual positive cases were correctly identified by the model:

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

Recall is particularly important when the cost of false negatives is high, as it measures how many of the actual positive cases were correctly identified by the model.

These metrics are often used together to provide a more comprehensive evaluation of model performance, especially in imbalanced datasets where one class is much more prevalent than the other.

3.5.2 – F1 and F2 scores

The F1 and F2 scores are widely used to evaluate the performance of classification models, particularly in cases where the data is imbalanced. These metrics provide a combined measure of both precision and recall, offering a more informative assessment than accuracy alone.

The F1 score is the harmonic mean of precision and recall:

$$F_1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

It treats precision and recall as equally important.

The F2 score, on the other hand, emphasizes recall more heavily. This is especially relevant in our context, where minimizing false negatives is crucial (i.e., we want to avoid predicting a cell as not being at risk when it actually is). The general formula for the F2 score is:

$$F_2 = 5 \cdot \frac{\text{Precision} \cdot \text{Recall}}{4 \cdot \text{Precision} + \text{Recall}}$$

These scores allow for a nuanced comparison between models depending on the relative importance of precision and recall.

3.5.3 – AUC

The Area Under the Curve (AUC) is another key metric for binary classification. It measures the model's ability to distinguish between positive and negative classes across all decision thresholds.

An AUC of 0.5 would indicate a model that performs no better than random guessing, while an AUC of 1.0 indicates a perfect model. Likewise, an AUC below 0.5 indicates a model that performs worse than random guessing.

The ROC curve plots the true positive rate (recall) against the false positive rate, and the AUC is the area under this curve.

A higher AUC indicates better discrimination capability between the classes.

Chapter 4 – Results

This chapter presents the findings from the training and evaluation of the five proposed models.

4.1 – Model Performance Comparison

Table 3 presents the comprehensive performance metrics for all models across training and testing datasets.

Key findings:

- The **Heuristic** model serves as a baseline, achieving a test F2 score of 0.755 with a high precision of 0.957 but a lower recall of 0.717, indicating it is conservative in predicting at-risk cells. Overall, it has a good AUC of 0.86, indicating a good classification ability.
- The **Heuristic + Grid Search** improved over the basic heuristic model, improving most metrics, for a slight loss in precision (0.941 vs 0.957) but a significant gain in recall (0.779 vs 0.717), resulting in an F2 score of 0.807, which is a notable improvement over the heuristic model. The AUC remains the same at 0.86, indicating that its ability to classify between safe and hazard cells has not changed significantly.
- The **XGBoost** model achieved a test F2 score of 0.906, with a high recall of 0.937, indicating its strong ability to identify at-risk cells. However, it had a lower precision of 0.799, suggesting it may produce more false positives. Its high training metrics also indicate some overfitting. Early stopping was tested to mitigate this, but it did not significantly improve the results.
- **XGBoost Dist** achieved the highest recall (0.978) but at the cost of precision (0.767). Even though we want to minimize false negatives, this model's low precision indicates it may not be suitable for practical applications. Also, by comparing its metrics with the heuristic model, we can see that it does not learn any new patterns, as it has a much lower AUC (0.77 vs 0.86).
- The **Neural Network** model showed the highest AUC (0.88), demonstrating superior classification ability. It achieved a test F2 score of 0.898, with a good balance between precision (0.851) and recall (0.91). This model's performance indicates that it can effectively learn complex patterns in the data, outperforming the heuristic and XGBoost models in terms of AUC and F2 score.

The grid search optimization improved the basic heuristic model's F2 score from 0.755 to 0.807, demonstrating the value of parameter tuning even for simple models.

Model	Train Precision	Train Recall	Train F2	Test Precision	Test Recall	Test F2	Test AUC
Heuristic	0.941	0.641	0.683	0.957	0.717	0.755	0.86
Heuristic + Grid Search	0.916	0.763	0.789	0.941	0.779	0.807	0.86
XGBoost	0.887	0.99	0.969	0.799	0.937	0.906	0.85
XGBoost Dist	0.794	0.986	0.941	0.767	0.978	0.926	0.77
Neural Net	0.882	0.932	0.922	0.851	0.91	0.898	0.88

Table 3 - Model performance metrics for training and testing datasets.

Overall, it seems that the Neural Network model achieves the best balance between precision and recall, with a high AUC, indicating that it can effectively learn complex patterns in the data.

The XGBoost model also performs well, but it has a lower precision. It is not such a big issue in our case, as false positives are not as critical as false negatives.

The following figure shows the ROC curves for each model, illustrating their classification performance across different thresholds. We can clearly see that the Neural Network model outperforms the others in terms of AUC, followed by the the heuristic models and the XGBoost model, which have similar AUCs.

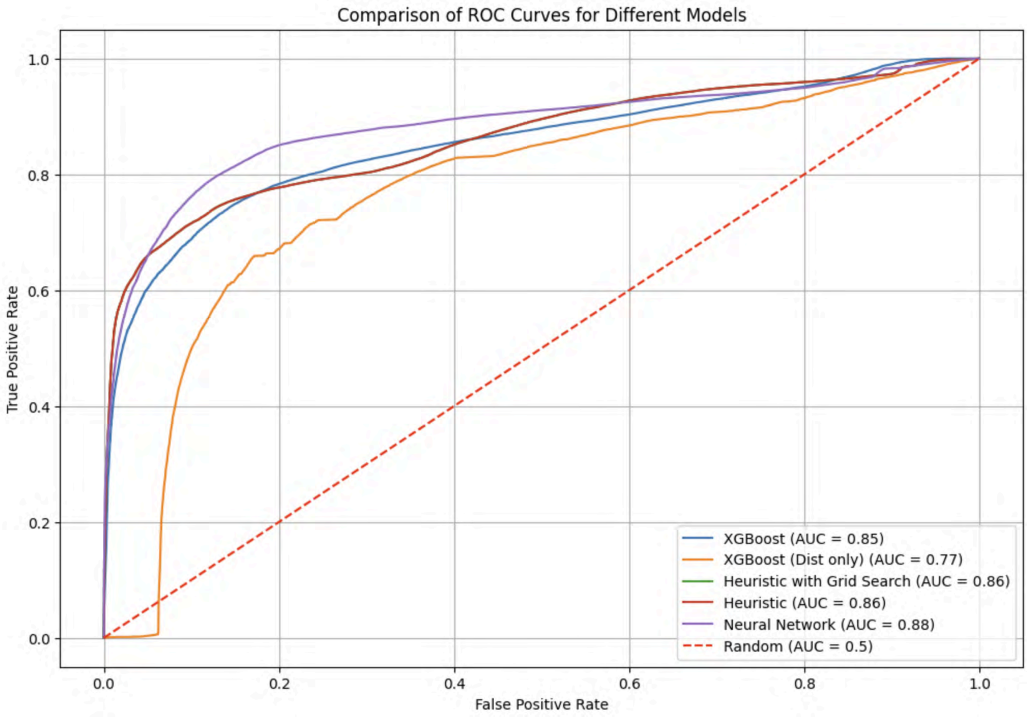

Figure 12 - Comparison of the ROC curves of the different models, showing the AUC for each model. The Neural Network model outperforms the others with an AUC of 0.88.

4.1.1 – Prediction maps

The following figure presents a visual comparison of the predicted hazard maps generated by each model against the ground truth hazard map. The ground truth map is based on the expert-defined hazard zones, while the predicted maps are generated by each model using the features and methods described earlier. Blue pixel represent safe cells (no hazard), while red pixels represent hazardous cells (hazard levels 1-3). White pixels represent cells that are outside of any hazard zones, and are not considered in the evaluation.

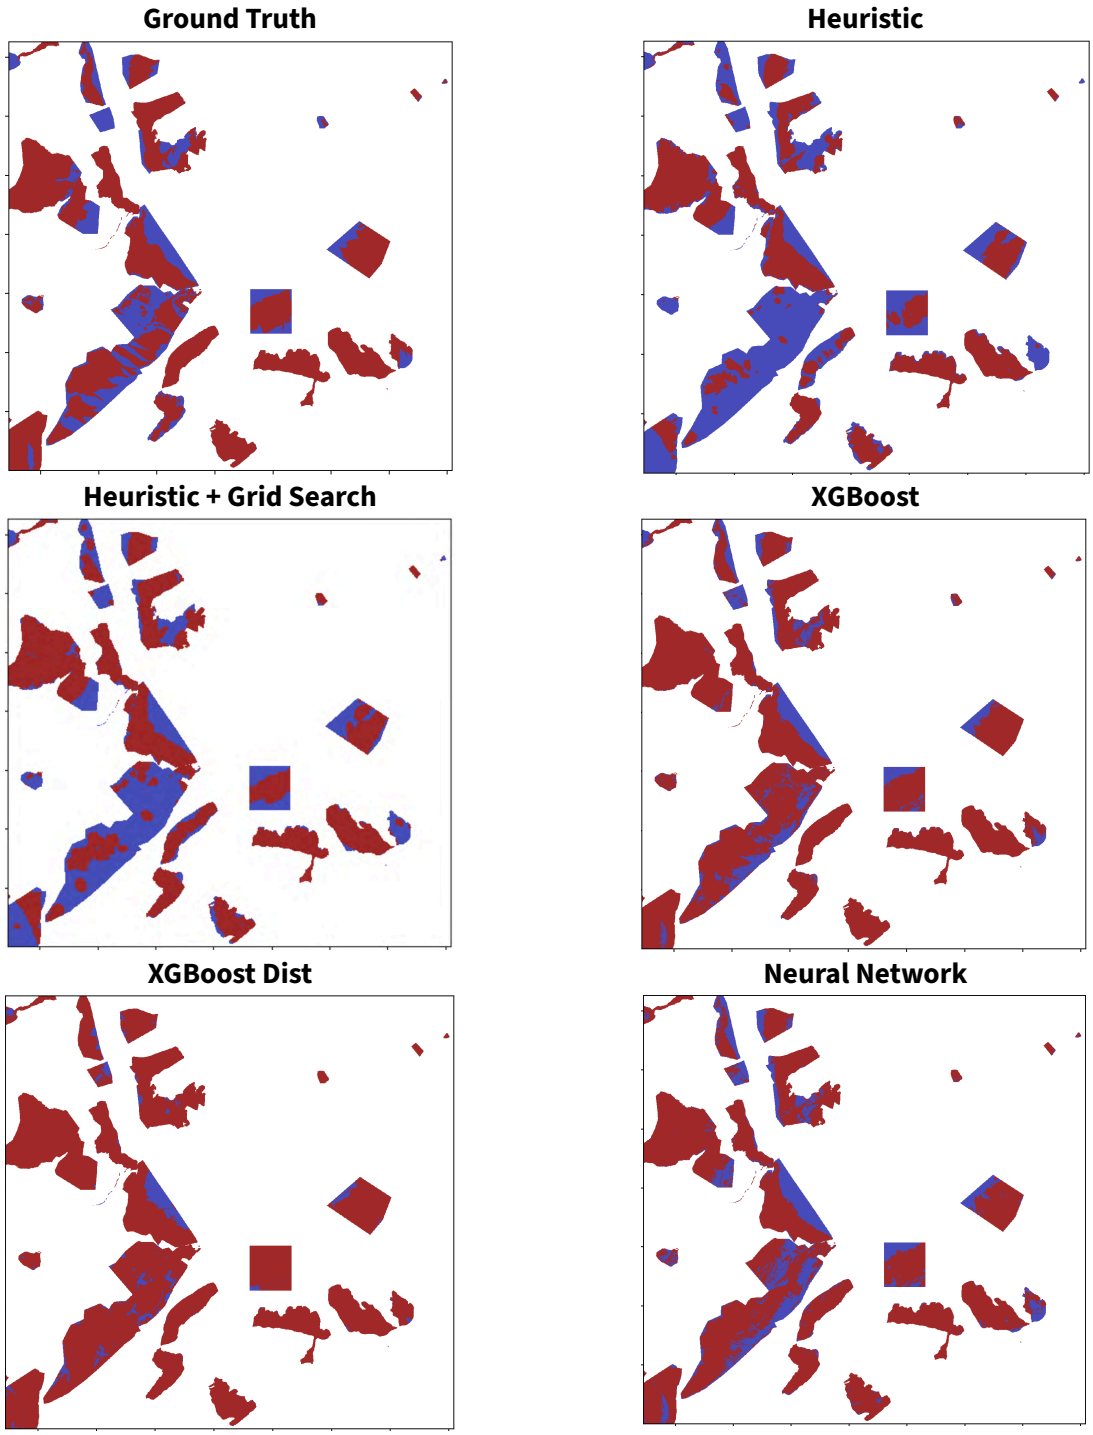

This visualisation allows us to confirm what we observed in the metrics: the NN and XGBoost models are able to capture the general patterns of the hazard zones. The XGBoost Dist model on the other hand captures too broad of hazard zones, which makes its predictions rather useless, as it predicts almost all cells as hazardous. As for the heuristic models, it seems to capture some hazard zones better than others. The reason why will be discussed in Section 5.

4.2 – XGBoost

This section presents the training results of the XGBoost model.

Table 5 presents the hyperparameters that were optimized using Optuna. The selected features were selected using the select K best method, which selects the top K features based on their importance scores.

The selected features are the slope, the RockyFor3D trajectory, the rasterized SilvaProtect trajectory (trajectory_heatmap), the SwissAlti3D altitude and the uphill distance to the nearest source cell with energy dissipation (distance_to_source).

hyperparameter	value
learning_rate	0.8086
n_estimators	96
subsample	0.6568
select_k_best	5
gamma	0.3760
reg_alpha	3.254e-05
reg_lambda	0.4171

Table 5 - Hyperparameters optimized for the XGBoost model using Optuna.

Figure 13 and Figure 14 present the SHAP values and feature importance for the XGBoost model. The RockyFor3D trajectory and the uphill distance to the nearest source cell with energy dissipation are the most important features, with the RockyFor3D trajectory having the most important positive contribution to the model's predictions.

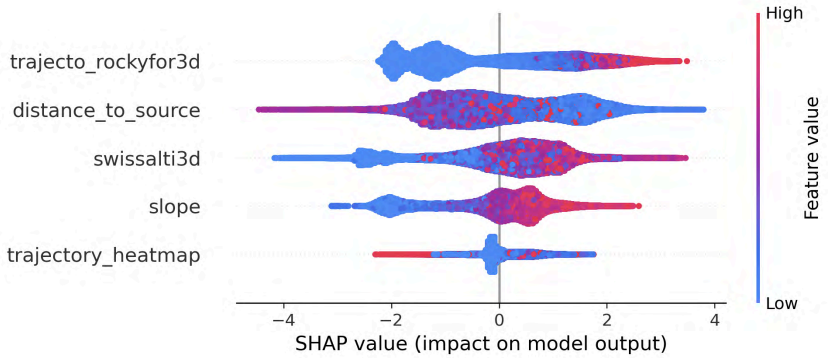

Figure 13 - SHAP values for the XGBoost model. The SHAP values show the contribution of each feature to the model's predictions, with positive values indicating a positive contribution and negative values indicating a negative contribution.

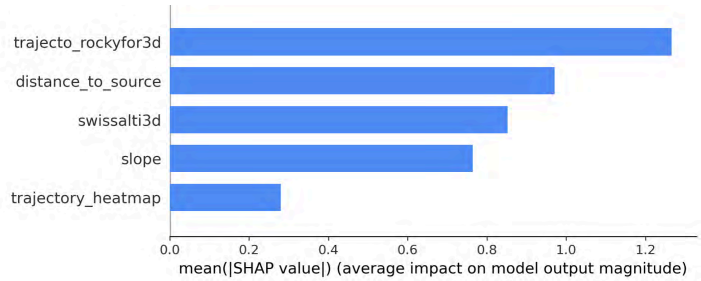

Figure 14 - Mean SHAP values for the XGBoost model, showing the most important features for the model's predictions.

4.3 – Neural Network

Table 6 presents the hyperparameters that were optimized using Optuna for the Neural Network model. Three features were selected using the select K best method, them being the slope, the RockyFor3D trajectory and the altitude from the SwissAlti3D raster.

hyperparameter	**value**
solver	adam
activation	relu
alpha	0.0020
learning_rate_init	0.0003
max_iter	300
select_k_best	3

Table 6 - Hyperparameters optimized for the Neural Network model using Optuna.

Chapter 5 – Discussion

This section tries to interpret the results presented above, and addresses the limitations of this project.

5.1 – Model Performance Analysis

The Neural Network achieved the best overall performance, which can be attributed to its ability to capture complex relationships between terrain features that heuristic models cannot. Its strong classification ability, as indicated by an AUC of 0.88, allows it to outperform both the heuristic and XGBoost models.

In short, with its balanced precision (0.851) and recall (0.91), the NN seems to provide the best trade-off for practical applications.

The XGBoost model also performed well, achieving a high recall (0.937) and F2 score (0.906), but with lower precision (0.799). This suggests that while it is effective at identifying at-risk cells, it may produce more false positives than the NN.

5.2 – Heuristic Models Limitations

In the prediction maps presented above, we may notice a certain area that heuristic models struggle to predict correctly, which is the south-facing slope north of the city of Martigny.

If we were to look more closely at the coverage of the source cells layers in this area, such as in Figure 15, we can notice it is almost devoid any, which explains why the models struggle to predict rockfall susceptibility in this area.

Figure 15 - Problematic area, showing the predicted hazard map for the Heuristic + Grid Search model (left), the source cells (middle) and the RockyFor3D trajectories (right).

We can also notice that the same area is also devoid of RockyFor3D trajectories. Indeed, if we look at the prediction maps for XGBoost and Neural Network models, we can see that they also struggle to predict hazard cells correctly, albeit to a lesser extent than the heuristic models.

5.3 – Feasability of using ML for rockfall susceptibility prediction

The results of the models presented in this thesis show that machine learning techniques can be effectively used to predict rockfall susceptibility, with the Neural Network model achieving the best performance overall.

While not as precise as the expert-made hazard zones, the models are able to capture the general patterns of rockfalls, and could be used to help identify areas at risk of rockfall, especially in areas where expert knowledge is lacking or where the terrain is too complex to be analyzed manually.

However, it is important to note that the models are not perfect, and they should not be used as a replacement for expert knowledge. As shown in the previous section, the quality of the predictions can vary significantly depending on the area, and some areas may be more difficult to predict than others.

They can be used as a complementary tool to help identify areas at risk of rockfall, but they should not be relied upon solely for decision-making.

Chapter 6 – Conclusion

The results of the models presented in this thesis show that machine learning techniques can be effectively used to predict rockfall susceptibility.

One other success of this project is the development of the whole pipeline, from data loading to prediction, in a modular and reusable way, which allows for easy experimentation with different models and features. This pipeline can be used as a starting point for future work in this area, and could even be adapted to other study areas or datasets.

6.1 – Future Work

6.1.1 – Implementing trajectory features

As described in an earlier section, feature engineering procedures have been setup to extract raster data from trajectories. Unfortunately, these features were not used in the models presented in this thesis. However, they could be used in future work, with the hope to improve the models' performance.

To do so, we would need for each cell in the study area to find each trajectories that passes through it, then snip each trajectory so it ends at the cell. Then, we would need to extract the features from these trajectories, such as the maximum and minimum rise a rock has experienced while travelling to the cell, the average slope, the distance travelled, etc.

Finally, the data from all the trajectories would then be aggregated to create a feature map which, for each cell, would contain the average, maximum, and minimum values of the features extracted from the trajectories that pass through it. These feature maps would then be used as input to the models, alongside the other features used in this thesis.

Implementing this feature engineering procedure would require a significant amount of computational resources, as it involves processing a large number of trajectories and extracting features from them. But it could potentially lead to significant improvements in the models' performance, as trajectories are a key aspect of rockfall events, and contain a lot of contextual information that is not captured by the other features used in the models.

6.1.2 – Introduce the concept of slope units

Slope units are terrain segments, delimited by drainage lines. In other words, they area the common area where rainfall water over a certain point flows down to a common point, such as a stream or a river.

As described above, we currently define rockfall hazard areas based on the date provided by the SDANA and the geographic closeness of the different areas. This leads to a huge disparity between the size of the hazard areas, as some are very large, spanning whole side of mountains, while other are very small, only covering a few hundred square meters.

Z. Change et al. [4] successfuly use this concept of slope unit to delineate areas that are prone to landslides. Even though the context is not exactly the same, we could try to apply this concept to rockfalls as well, as it could lead to more consistent hazard areas, and potentially improve the models' performance.

M. Alvioli et al. [5] propose a method to automatically delineate slope units from digital elevation models (DEMs) using a combination of topographic features through their `r.sopleunits` module in GRASS GIS, published as free software under the GNU General Public License. In the same paper, they also propose an optimal set of parameters for the Umbria region in central Italy as well as the methodology to obtain them, which could be used as a starting point for our case.

6.1.3 – Try other ML model architectures

We have only explored a few models in this thesis, and there are many other models that could be used to predict rockfall susceptibility. For instance, we could try using more complex models such as convolutional neural networks (CNNs), which are well-suited for processing spatial data.

CNNs, because of their ability to learn spatial hierarchies through their convolutional layers, could potentially capture more complex patterns in the data, and improve the models' performance.

6.1.4 – Improve rockfall source area definition

As described above, rockfall source areas are defined solely based on slope steepness, as per the data provided by the SDANA. This causes issues, as some areas that have been identified as prone to rockfalls by experts do not possess any rockfall source area uphill, which leads to a total inability to predict rockfalls in these areas based on that feature.

A. M. Fanos et al. [2] propose a solution to this problem, by training a model to predict rockfall source areas based on the intersection of the slope steepness raster and their probability map of rockfall. This model could then be used to more accurately predict rockfall source areas in the study area, and improve the models' performance.

6.1.5 – Improvements of the data loading pipeline

This is a more technical and somewhat minor aspect, but some features need to be generated manually through a command-line script. This causes some unneeded friction in the workflow, as the data loading pipeline is not fully automated. It would be beneficial to integrate feature generation directly into the data loading pipeline, so that each model can request its desired features, and it would automatically generate them when loading the data, or load the cached features if they have already been generated.

6.1.6 – Test the models on the whole Canton of Valais

It could be interesting to see how the models perform on the whole Canton of Valais, as the study area is only a small part of it. This would allow us to see if the models are able to generalize to other areas, and if they can be used to predict rockfall susceptibility in parts of the Canton with different topographical and geological characteristics.

Bibliography

[1] M. Jacquemart *et al.*, “Detecting the impact of climate change on alpine mass movements in observational records from the European Alps,” *Earth-Science Reviews*, vol. 258, p. 104886, Aug. 2024, doi: 10.1016/j.earscirev.2024.104886.

[2] A. M. Fanos, B. Pradhan, A. Alamri, and C.-W. Lee, “Machine Learning-Based and 3D Kinematic Models for Rockfall Hazard Assessment Using LiDAR Data and GIS,” *Remote Sensing*, vol. 12, no. 11, p. 1755, Jan. 2020, doi: 10.3390/rs12111755.

[3] H. Wen, H. Jiwei, Z. Jialan, X. Xuekun, , and M. Liao, “Rockfall susceptibility mapping using XGBoost model by hybrid optimized factor screening and hyperparameter,” *Geocarto International*, vol. 37, no. 27, pp. 16872–16899, Dec. 2022, doi: 10.1080/10106049.2022.2120547.

[4] Z. Chang *et al.*, “Landslide susceptibility prediction using slope unit-based machine learning models considering the heterogeneity of conditioning factors,” *Journal of Rock Mechanics and Geotechnical Engineering*, vol. 15, no. 5, pp. 1127–1143, May 2023, doi: 10.1016/j.jrmge.2022.07.009.

[5] M. Alvioli *et al.*, “Automatic delineation of geomorphological slope units with <tt>r.slopeunits v1.0</tt> and their optimization for landslide susceptibility modeling,” *Geoscientific Model Development*, vol. 9, no. 11, pp. 3975–3991, Nov. 2016, doi: 10.5194/gmd-9-3975-2016.

Appendix

List of Figures