CHALMERS
UNIVERSITY OF TECHNOLOGY

Stochastically Modelling Road Topography

Identifying Road Topography Characteristics and Stochastically Modelling it for Simulations

Master's thesis in systems, control and mechatronics

Jonatan Svensson
Henrik Bertilson

DEPARTMENT OF MECHANICS AND MARITIME SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2025
www.chalmers.se

Master's thesis 2025

Stochastically modeling road topography

Identifying Road Topography Characteristics and Stochastically Modeling it for Simulations

Jonatan Svensson
Henrik Bertilson

Department of Mechanics and Maritime Sciences
Division of Vehicle Engineering and Autonomous Systems
Chalmers University of Technology
Gothenburg, Sweden 2025

Stochastically modeling road topography
Identifying Road Topography Characteristics and Stochastically Modeling it for Simulations
Jonatan Svensson
Henrik Bertilson

Examiner: Fredrik Bruzelius, Chalmers University
Industrial Supervisors: Sebastien Ragot and Fredrik Öijer, Volvo Group
Academic Supervisor: Carl Emvin, Chalmers University

Master's Thesis 2025
Department of Mechanics and Maritime Sciences
Division of Vehicle Engineering and Autonomous Systems Vehicle Dynamics group
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Stochastic Modeling of Road Topography
Identifying Key Characteristics for Modeling and Simulation
Jonatan Svensson
Henrik Bertilson
Department of Mechanics and Maritime Sciences
Chalmers University of Technology

Abstract

Road topography is the 2-dimensional elevation profile of a road and is a critical factor in evaluating vehicle performance and energy consumption in the transport sector. Currently the most used method is to model existing roads. This method requires that a suitable real representative road exists or needs to be constructed in expensive testing tracks. It also runs into the problem of being constrained to one road and can't capture the larger characteristics of an area.
This study aims to find alternatives to this method by modeling road topography stochastically to create an infinite number of roads with varying characteristics. To achieve this, it investigates real world road data and develops various stochastic modeling approaches to generate synthetic road profiles that capture the key topographical characteristics observed in actual roads. The models that were examined were AutoRegressive(1), ARMA(1,5) and Markov chain model. The models were designed to model the slope of the road to capture the characteristics of a set of input roads. The generated models are validated against real world data using multiple comparison metrics, to assess their ability to capture real-world characteristics. This is to evaluate the potential for simulation applications.
To identify which characteristics are essential for road topography modeling, a qualitative study was performed, where experts within vehicle engineering disciplines were interviewed.

Keywords: road, topography, modeling, stochastic processes, simulation, characteristics, AR(1), Markov process

Acknowledgements

We would like to express our sincere gratitude to our supervisors, Carl Emvin, Sébastien Ragot and Fredrik Öijer, and our examiner, Fredrik Bruzelius, for their invaluable guidance, support, and expertise throughout the duration of this project. Your insightful feedback and constructive advice have been instrumental in shaping our work and enhancing our understanding of the subject matter.

Your dedication and encouragement have helped us navigate the challenges of this project, from guidance on methods and ideas, to connecting us to resources and people. We deeply appreciate the time and effort you devoted to our thesis, and we are truly thankful for your mentorship.

Henrik Bertilson
Jonatan Svensson
Gothenburg, Sweden, March 2025

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

AR	AutoRegressive
AR1	First-degree AutoRegressive
ARIMA	Autoregressive Integrated Moving Average
ARMA	Autoregressive Moving Average
BEV	Battery Electric Vehicle
ESS	Effective Sample Size
GTA	Global Transport Application
ICE	Internal Combustion Engine
PELT	Pruned Exact Linear Time (algorithm)
PFLAT	Predominantly Flat (GTA topography classification)
RQ	Research Question
VHILLY	Very Hilly (GTA topography classification)

Nomenclature

Below is the nomenclature of indices, sets, parameters, and variables that have been used throughout this thesis on stochastically modeling road topography.

Parameters

a	Autoregressive coefficient in the AR(1) model
σ_ϵ	Standard deviation of the white noise in the AR(1) model
σ_y	Standard deviation of the the AR(1) process
Δx	Spatial discretization step (m)
L_s	Length of a discretized road segment (m)
L_h	Average hill length (m)
p, d, q	Orders of the ARIMA model
λ	Transition rate in the Markov chain model
P_{ij}	Transition probability from state i to state j (Markov chain)
$p_{y,\min}$	Minimum driving distance percentage for a GTA class
$p_{y,\max}$	Maximum driving distance percentage for a GTA class
e_k	White noise term in the AR(1) model at segment k
y_k	Road gradient value at segment k
h	Altitude (m)
h_0	Initial altitude at the start of a road segment (m)
θ	Road slope angle (degrees or percentage)
$\hat{a}$	Estimated autoregressive coefficient
$\hat{\sigma}_y$	Estimated standard deviation of the AR(1) process
r_0	Sample variance of input data
r_1	One-step sample covariance of input data
λ_i	Transition rate for state i (Markov chain)

G_{ij}	Elements of the generation matrix in a Markov chain process (transition rate from state i to state j)
L	Threshold level in level-crossing analysis
Δt	Sampling interval in level-crossing analysis

Contents

List of Figures

List of Tables

1
Introduction

The study of energy consumption, power transmission, and other vehicle components in trucks and cars has gained significant attention in recent years, driven by the global imperative to reduce CO_2 emissions and optimize vehicle performance for customers. One effective approach to understanding these factors is analyzing real-world data from vehicle operations to identify correlations. For instance, a prior thesis investigated the impact of ambient temperature, traffic conditions, and other variables on vehicle performance[1]. Another notable aspect studied in that project was road grade, where energy consumption was analyzed across varying road inclinations using real-world datasets. The findings of these studies are illustrated in Figure 1.1. Where driving data from a Nissan Leaf model was considered, which is a small family electric vehicle.

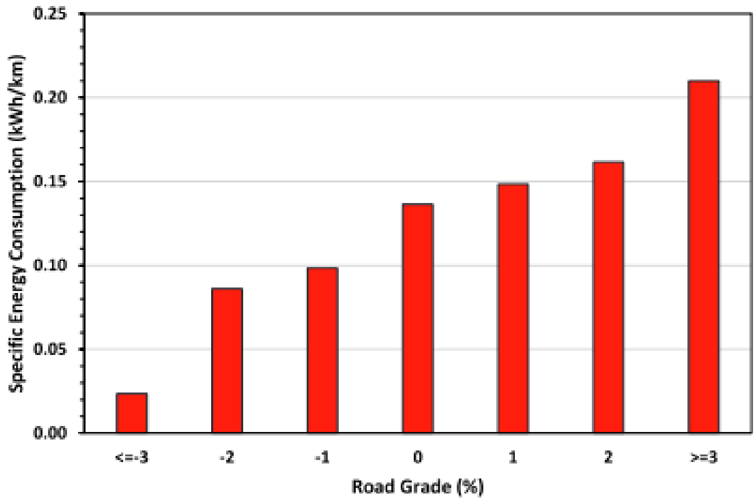

Figure 1.1: Variation of energy consumption with road grade in real drive cycle trips.

While using real-world data provides highly realistic insights into vehicle performance, it is inherently limited to the specific conditions and measurements captured during the data collection. To overcome these limitations, models can be employed to describe the conditions and roads in a systematic way.
Generative models offer several advantages, one of which is controlled variability. This capability allows for the generation of synthetic road profiles with adjustable parameters, enabling the simulation of a wide range of conditions—including scenarios that may not be represented in real-world data. In contrast, real-world data is often

limited to specific geographic regions or road types, thereby restricting its scope and applicability. Generative models help overcome this limitation by generalizing findings and uncovering correlations between key parameters, which may not be immediately evident from observational data alone. Given these benefits, modeling approaches can be a complement to real-world data analyses.. By integrating models into transport operation studies, researchers can simulate diverse conditions, gain broader insights, and develop effective strategies for optimizing vehicle performance and energy efficiency.

The transport operation, which refers to the interaction of vehicles, infrastructure, and external conditions to achieve specific transport goals, involves numerous factors, including road characteristics, weather conditions, traffic flow, and the mission being carried out. Integrating these components with driver behavior and vehicle specifications in simulations is a challenging yet essential task. When conducted effectively, such simulations provide valuable insights into how each element of the transport operation influences overall performance.
To perform such simulations, it is critical to have well-defined models for each component. Laboratory-standard driving cycles often fail to correlate well with real-world driving conditions due to disparities in factors such as geography, traffic and vehicle type[2]. This limitation highlights the need for more precise and adaptable models to improve simulation reliability. This project focuses on developing a robust model for road topography.
Accurate simulation of road topography is fundamental when studying vehicle performance and understanding the interaction between vehicles and their operating environments.
Modern vehicle simulations incorporate detailed input parameters gathered from real-world runs, including road profiles, curvature, speed patterns, weather conditions, and driver behavior. These parameters are combined into operating cycles, which replicate specific driving scenarios to analyze metrics such as energy consumption, brake efficiency, durability, and tire wear.
The idea of representing the road and the operating environment through models is not new and has been explored extensively by Pettersson, and some other researchers from Chalmers.[3], who highlighted its potential to address challenges in vehicle design and operation. Here it was suggested that modeling transport operations—including roads, weather, and traffic conditions—provides a structured approach to understanding and optimizing vehicle performance within their operating environments.
Building upon Pettersson's work, Romano further explored in his PhD thesis how transport missions could be described mathematically, reinforcing the need for structured and precise modeling approaches to capture the complexities of road transport operations[4]. The road topography model used in this previous work, will be one of the models used and implemented for this project as well.

1.1 Problem motivating the project

Although using real-world data is effective due to its high realism, it has significant limitations. Collecting and maintaining real-world data takes time and often is

geographically constrained, limiting its applicability across diverse regions or road types. The data is inherently restricted by the structure of existing roads, limiting its flexibility for studying novel or hypothetical scenarios. To be able to design real-world data, test tracks need to be built, which is expensive.
In comparison, the common methods for generating roads suffer from being too simplified, and lack the connection to the real-world data. This project aims to bridge the gap between these two methods, and find a way to generate road profiles that are grounded in real-world data, and are adaptable and quick to generate. Virtual models can generate synthetic road profiles with tailored characteristics, allowing for exploration beyond the constraints of real-world data. This enables simulations to test a wide range of conditions, paving the way for more comprehensive analyses of vehicle performance under diverse and controlled scenarios, making this a more flexible approach.

1.2 Aim

The objective is to develop and evaluate road topography models that effectively capture the statistical topography properties of real roads and enable the generation of virtual roads for simulation scenarios.
To identify the performance of the models, the underlying characteristics that are most important for road topography need to be identified and examined. This will give a benchmark for which characteristics are important to capture in models, and what is crucial for different disciplines when simulating using road topography.

1.3 Research questions

1. What are the key road topography characteristics?
2. How can road topography be modeled stochastically?
3. How can the generated roads from the topography models be relevant for use in vehicle simulations?

1.4 Limitations

- The project is limited to 20 weeks.
- The data is limited to the availability and quality of data gathered in the Here database.
- The models are limited to relatively simple models.
- No advanced simulations are used for model evaluation.

1.5 Ethical and Environmental Considerations

This research adheres to ethical standards related to data handling, privacy, and academic integrity. All data used in this study, including elevation profiles, geographic coordinates, and road attributes, were sourced from publicly available, open-access

databases or through partnerships with agencies providing anonymized datasets. No personal or identifiable user data was used or collected.
Furthermore, the modeling approach and analysis have been conducted with transparency and reproducibility in mind. Algorithms and methodologies are clearly documented to allow for independent validation. The study also avoids any manipulation of results to fit a desired outcome, ensuring objectivity and scientific rigor.

Modeling road topography contributes to sustainable infrastructure planning and mobility systems. By accurately characterizing road slopes and elevation profiles, the outcomes of this research can support better fuel efficiency assessments, electric vehicle energy consumption modeling, and the identification of road segments with higher environmental impact due to elevation-related factors. This work has a low environmental footprint, as it is computational in nature and does not involve physical field experiments or resource-intensive processes. In the broader context, improved understanding of road topography supports eco-friendly transportation planning, contributing indirectly to environmental conservation efforts and climate change mitigation.

2
Theory

This chapter explains the theory of the mathematics and models used in this thesis. Additionally it introduces definitions used, and explains the theory behind the simulations.

2.1 Road Topography

Road topography in this report refers to a two dimensional elevation profile of the altitude over distance, of a continuous road. Meaning the altitude, measured in meters, on the vertical axis and the distance, mostly measured in kilometers, on the horizontal axis. This can also be described with a slope profile, where the slope is the inclination of the road.
Road topography is often defined as wavelengths of 50 meters and above. Features under 50 meters are considered as road roughness, and are therefore not relevant [5]. This means that when studying real-world roads, only data sampled at intervals of 25 meters or more is considered, since 2 measurements are needed for each wavelength. Similarly, when generating new roads from a model, sampling every 25 meters is deemed sufficient. Using this approach the model has an effect comparable to a low pass filter.

2.1.1 Road Topography Hills

Two major aspects of road topography that go hand in hand are the slope profile and the elevation profile. The slope profile describes the gradient, measured in percentage, of the road on the vertical axis along with the length of the road measured in kilometers on the horizontal axis. While the elevation profile shows the altitude, measured in meters, along the same length.
Examining the characteristics of these two profiles for generated roads compared to real roads can serve as a means to validate a given model to ensure that it captures key features in a realistic way.
Since the point of this project is to model with at least some stochasticity, meaning having some degree of randomness, designing a specific road is not the goal. It is instead to formulate a method for generating roads which capture the most important characteristics of the road, and identifying what these characteristics are.
To describe topography, it is also necessary to set one definition of what a slope is since there are several different ways to define what is meant by a slope. For this thesis it is defined as the path between two points that are both flat (0 degrees

angle), with a continuous positive or negative incline, as shown in Figure 2.1a. When classifying the incline of this hill, one can see that the gradient is not constant, but instead fluctuates from zero, to a positive value, and then back to zero. The slope of the hill can then be described to either have the average gradient of the entire hill see Figure 2.1c, or be split at specific angle intervals see Figure 2.1b. Both definitions are employed in this thesis, with the selected definition specified in each instance to avoid ambiguity.

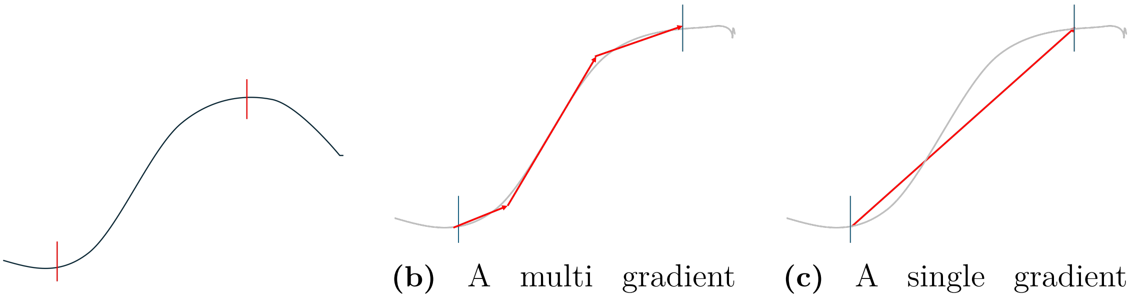

(a) definition of a hill

(b) A multi gradient sampled hill

(c) A single gradient sampled hill

Figure 2.1: Slope splitting of a hill

2.1.2 GTA parameters

Volvo has an internal language to characterize the customer environment and usage called **G**lobal **T**ransport **A**pplications (**GTA**). These are defined to classify a vehicle's usage and environment under its lifetime. The parameters do not only cover road topography but also factors like a vehicle's transport mission, usage and operating environment.

Road topography is currently described in GTA by three parameters: topography, max altitude and max road inclination. Of these three, it is most important to capture the topography class, since it is the characteristics that most affects the vehicle under its lifetime. In GTA topography is defined by the slope of the total driving distance. The classification of the different categories can be seen in Table 2.1. The definitions start at Flat, and if the vehicle exceeds the limitations of the class, it moves up one step until it satisfies the criteria. The classes in order are Flat, Predominantly flat, Hilly, and Very hilly. The classification is defined for the lifetime of a vehicle, but can be used to talk about a shorter time period or a specific road, which will be done in this thesis.

Category	**Topography class**	**Gradient < 3%**	**< 6%**	**< 9%**
Flat	FLAT	>98%	-	-
Predominantly Flat	P-FLAT	-	>98%	-
Hilly	HILLY	-	-	>98%
Very Hilly	V-HILLY	Criteria for Hilly not met		

Table 2.1: Topography classification based on road slope gradients.

To see how this is applied to one specific road, the road profile for the distance

Landvetter-Borås is plotted in Figure 2.2. It shows the altitude profile above, and the grade profile below. The grade is lower than 3% for around 91.5% of the trip, depicted as the blue lines in the top image. This places the road in the Predominantly flat category, since 91.5% is less than the required 98% required to achieve the Flat, but is not placed in the Hilly category since the road has 100% of the distance above 6%.

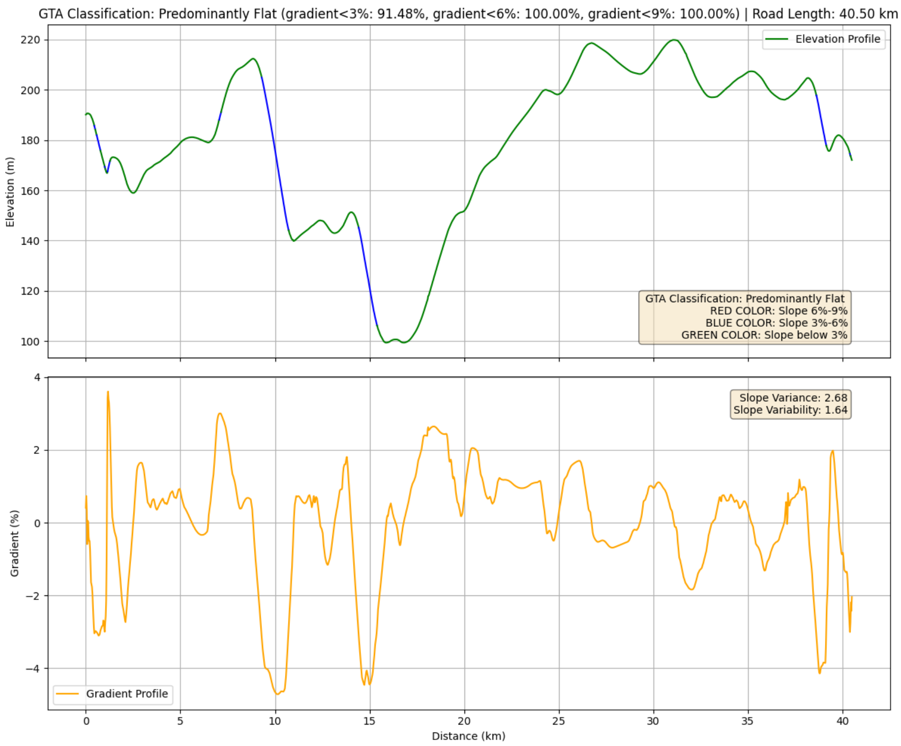

Figure 2.2: Road profile of altitude and slope, with the GTA Topography classification system added. Depicts the road between Landvetter-Borås

2.2 Stochastic Modeling in Transport Applications

In this section three stochastic models will be described. Stochastic is the property of being well defined by a random probability distribution. In modeling this means the output is defined as a probability of how the road is designed, but the specifications of what value is selected is random.

2.2.1 Autoregressive (AR(1)) Model

A first-order AutoRegressive model, or AR(1) model, describes each new value in a time (or space) series as a linear function of its immediate predecessor plus a stochastic term. It has been applied extensively for modeling road topography because it captures both the dependence on the previous slope and the inherent randomness in real terrain. Several studies have used AR(1) processes to generate synthetic road profiles, motivating this project usage of it [5, 6].

In its simplest form [3], the AR(1) model for the road gradient y_k can be written as

$$y_k \;=\; a\, y_{k-1} \;+\; e_k, \quad e_k \sim \mathcal{N}(0, \sigma_e^2), \quad |a| < 1, \tag{2.1}$$

where a is the auto regressive coefficient and e_k is a zero-mean white-noise term with variance σ_e^2. Intuitively, a controls how strongly the new slope y_k depends on the previous slope y_{k-1}. If $|a|$ is close to 1, changes happen slowly, producing long, smooth hills. If $|a|$ is smaller, the gradient fluctuates more rapidly, leading to shorter, more frequent peaks and valleys.

Although a is a convenient mathematical parameter, it can sometimes be more intuitive to express it as a average hill length, L_h. Drawing on autocorrelation and zero-crossing properties, shows that

$$L_h \;=\; \frac{4\pi}{\pi - 2\arcsin(a)}\; L_s, \tag{2.2}$$

where L_s is the physical distance between successive discretized segments (sample rate). This parameter describes the average hill length for the generated roads. When the hill length is defined as zero crossings in the slope profile. This parameter can be useful when changes of the average hill lengths is something considered for simulations. However, this parameter only affects the average hill lengths and not the spread of the hill lengths.

When stationarity in the AR(1) process is achieved, meaning,

$$\mathrm{Var}[y_k] = \mathrm{Var}[y_{k-1}] = \sigma_y^2.$$

Because e_k is independent of y_{k-1}, it follows that

$$\sigma_y^2 \;=\; \mathrm{Var}[y_k] \;=\; \mathrm{Var}[a\, y_{k-1} + e_k] \;=\; a^2\,\sigma_y^2 \;+\; \sigma_e^2. \tag{2.3}$$

Rearranging (2.3) yields

$$\sigma_y^2 \;=\; \frac{\sigma_e^2}{1 - a^2}, \tag{2.4}$$

showing that σ_y^2, the total (long-run) variance of the AR(1) process, is completely determined by the white-noise variance σ_e^2 once the correlation parameter a is specified. In this context, σ_y is often called the topography amplitude, because it measures the overall variability (standard deviation) of the road-gradient distribution. When σ_y is large, steep gradients (both positive and negative) are more probable, producing "hillier" terrain. Conversely, if σ_y is small, large gradients are unlikely and the road profile stays relatively flat most of the time. The term "amplitude" therefore captures how large these vertical fluctuations can be, much like a wave amplitude represents the extent of oscillation above or below a reference level.

As noted by [3], the pair (σ_y, L_h), provides a concise yet physically meaningful way to describe road topography. More generally, there are multiple, equivalent ways to parameterize the same AR(1) model, depending on whether one emphasizes noise variance, total variance, or a physical average crest-to-crest distance (zero crossings in the slope profile).

Although Eqs. (2.1) define the AR(1) process in terms of $\{a, \sigma_e\}$, it can re-express those parameters in either $\{a, \sigma_y\}$ or $\{L_h, \sigma_y\}$ form. In particular:

1. (a, σ_e)**:** Here, a captures the one-step correlation, and σ_e is the amplitude of the white-noise term.
2. (a, σ_y)**:** Here, a still describes the correlation, but $\sigma_y = \sqrt{\text{Var}[y_k]}$ measures the overall standard deviation (or amplitude) of the AR(1) process. From (2.4), one have $\sigma_y^2 = \sigma_e^2/(1 - a^2)$.
3. (L_h, σ_y)**:** Here, L_h is a reparameterization of a that represents an average "hill length" (estimated via zero-crossings or spectral arguments). It remains equivalent to a, just measured in meters rather than as a dimensionless correlation. Meanwhile, σ_y is again the overall topography amplitude.

Crucially, all three formulations describe the same underlying AR(1) process, the choice of parameter set simply depends on which perspective is most convenient. In this project, the focus will primarily be on the (a, σ_y) formulation, since it aligns directly with the parameter-estimation procedures used.

Estimating parameters

One straightforward approach to estimating a and σ_y^2 from real data is the standard correlation method, also known as the Yule–Walker method [6]. Let r_0 denote the sample variance of the input data and r_1 the sample one-step covariance

$$r_1 \;=\; \frac{1}{n-1} \sum_{i=1}^{n-1} (y_{k,i} - \overline{y_k})\,(y_{k,i+1} - \overline{y_k}),$$

where n is the number of observations and $\overline{y_k}$ is the mean. Then the parameter estimates are

$$\hat{a} \;=\; \frac{r_1}{r_0}, \qquad \hat{\sigma}_y^2 \;=\; r_0.$$

Here, $\hat{\sigma}_y^2$ is the estimated total variance of the AR(1) model (which converges to the true variance under stationarity and large sample size), estimated directly as the sample variance of the real road data. Meanwhile, $\hat{a}$ captures the correlation between consecutive observations and is obtained by normalizing the one-step sample covariance r_1 by the total variance r_0. This division yields a correlation coefficient that indicates how strongly each slope depends on its predecessor, regardless of scale.

An alternative estimation approach is the zero-crossing method described in [6], which infers the correlation parameter from the frequency of "hills" (i.e., zero up-crossings in the slope profile) in the measured series. In practice, the zero-crossing estimates for a and σ_y have been shown to closely match those obtained from the Yule–Walker method. Nevertheless, the zero-crossing method might be of interest when the control of hill lengths is crucial. However it only allows such control if a hill is defined as a slope's zero crossing. Which in many cases is not a sufficient definition.

Calculating GTA parameters from variance σ_y

Since a common way to classify roads at Volvo is to use GTA classification. It may be of interest to connect model parameters to GTA classification if possible, since this could be beneficial feature when simulating. A method to connect the GTA

Topography parameters to the $\sigma_{Y|r_i}$ parameter, is discussed by Romano [4], where it is shown that if the road only has one road type, the GTA topography class can be statistically determined by only the variance σ. This can be seen if one begin with the definition of GTA from Table 2.1, and write the thresholds as:

$$p_{y,\min} < P(|Y| \leq y) \leq p_{y,\max}, \tag{2.5}$$

In this context, y represents a threshold for the road grade, while $p_{y,\min}$ and $p_{y,\max}$ define the lower and upper limits for the proportion of the total driving distance that falls within this threshold. Specifically, the GTA classification assigns $y = \{3, 6, 9\}$ and enforces fixed constraints of $p_{y,\min} = 0.98$ and $p_{y,\max} = 1$ on the corresponding driving distance fraction. Then if one define the road grade distribution as following a normal distribution:

$$Y \mid R_t = r_i \sim \mathcal{N}(0, \sigma^2_{Y|r_i}). \tag{2.6}$$

Given that:

$$P(|Y| \leq y \mid R_t = r_i) = 2\Phi\left(\frac{y}{\sigma_{Y|r_i}}\right) - 1, \tag{2.7}$$

where $\Phi(\cdot)$ denotes the Cumulative Distribution Function **(CDF)** of the normal distribution, and R_t is the road type, a method for splitting roads into categories depending on some variable such as speed. If then a given topography class is defined by ensuring that the probability $p_y(y, \sigma_Y, P_R, L_R)$ satisfies the inequality:

$$p_{y,\min} < p_y(y, \sigma_Y, P_R, L_R) \leq p_{y,\max}. \tag{2.8}$$

then this expression establishes a connection between **the topographic parameters** and **road type characteristics** and the **thresholds** y and $p_{y,\min}$. In the special case where only a **single road type** is considered (i.e., $\sigma_{Y|r_1} \equiv \sigma_Y$), the variance alone suffices to fully characterize a GTA class based on topography parameters. The corresponding ranges of σ_Y for different topographic classifications are provided in 2.2

Class	GTA system
FLAT	$\sigma_Y < 1.29$
P-FLAT	$1.29 \leq \sigma_Y < 2.58$
HILLY	$2.58 \leq \sigma_Y < 3.87$
V-HILLY	$3.87 \leq \sigma_Y$

Table 2.2: Topography classification based on the GTA system.

2.2.2 ARIMA and ARMA Models

An ARIMA model[7][8](**A**uto**R**egressive **I**ntegrated **M**oving **A**verage) is a generalized time series approach that combines three components. AR(p), which is the autoregressive part of order p, where each of the previous p observations influences the current

value. I(d), is the differencing part, with differencing of order d, used to transform a nonstationary series into a stationary one by subtracting lagged versions of itself. And the third part, MA(q), which is a moving average part of order q, where each of the previous q forecast errors (residuals) affects the current value.
Hence, an ARIMA(p, d, q) model can be thought of as AR(p) + I (d) + MA(q).
Notably, setting d = 0 converts the model into an ARMA(p,q) specification, as no differencing is applied. Moreover, with d = 0 and q = 0, the model becomes purely autoregressive (AR), and with p = 0 and d = 0, it simplifies to a purely moving average (MA) model.
By allowing for differencing, autoregressive terms, and moving average terms, an ARIMA(p, d, q) model provides a flexible framework to capture both past-value dependence and past-error dependence in potentially nonstationary time series.
Although ARIMA is often introduced as a general time-series model, it has also been applied to characterize and generate road profiles. For example, studies show that an ARIMA model can capture non-stationary features in measured road profile data[9]. Other authors have used ARIMA to forecast road slope or gradient over short horizons. For example, a study applies an ARIMA predictor to estimate the near-future road gradient and vehicle velocity in real time for hybrid-electric vehicle energy management [10]. Although their focus is on optimizing the vehicle's powertrain control rather than creating synthetic road profiles, it demonstrates how ARIMA can successfully handle gradient data in a predictive setting.
One report has also applied the ARIMA approach to shorter segments, with lengths as small as 250 mm, which corresponds to modeling road roughness rather than complete road topography [11]. While this does not rule out its application to longer segments, the study primarily focuses on capturing road roughness. Moreover, it is understood that the elevation profile is predominantly used as input, in contrast to this project where the slope profile will be considered as the input.

2.2.3 Continuous-time Markov Chain Process

A continuous-time Markov process is a type of stochastic process where the system evolves over time through transitions between different states. The key property of such a process is that the future state of the system depends only on the current state and not on the history of previous states [12]. Similar to the AR(1) model in its state-dependency, this process introduces behavioral changes that evolve continuously over time. This makes it well-suited for modeling various dynamic systems, such as road conditions, terrain types, or traffic dynamics, where the likelihood of transitioning between states is dependent only on the current state [13]. Data can be utilized to estimate the transition probabilities between states, allowing the model to accurately capture the underlying system dynamics. The specific parameters derived from the data include the mean duration of time spent in a given state and the transition matrix, which defines the probabilities of transitioning from one state to another.

Mean Length
The mean segment length for each state quantifies how long, on average, a system

remains in a particular state before transitioning to another. This is determined by first identifying all continuous segments where the system stays within the same state. Then, the lengths of these segments are measured, and their average is computed to obtain the mean segment length.

Transition Matrix

Along with the mean length, the transition matrix is also calculated from the data. The transition matrix describes how likely it is for the system to switch from one state to another. To compute this, the number of transitions from one state category to another is recorded throughout the dataset. This is done by counting instances of state transition, considering each transition between consecutive segments. The transition matrix is then derived by normalizing these counts for each row so that each row sums to 1. This ensures that the matrix represents probabilities, with each element P_{ij} in the matrix corresponding to the probability of transitioning from state s_i to state s_j. Mathematically, the transition matrix $\mathbf{P}$ for a system with n_s different states is a square matrix of size $n_s \times n_s$, where n_s is the number of states considered [12].

Generation Matrix Calculation

Once the transition matrix is computed, the generation matrix is derived, which models the continuous-time behavior of the system. The generation matrix $\mathbf{G}$ represents the rates of transitions between different states, and it is used to simulate the time spent in each state before transitioning to another.
The generation matrix is derived from the transition matrix, but with an important distinction: while the transition matrix $\mathbf{P}$ represents probabilities of switching between states, the generation matrix $\mathbf{G}$ represents rates. The off-diagonal elements of the generation matrix correspond to the rate of transitioning from one state to another, while the diagonal elements represent the rate of leaving the current state. The off-diagonal elements G_{ij} of the generation matrix are computed by multiplying the transition probabilities P_{ij} by the rate parameter λ_i for state i. The rate parameter λ_i is associated with the mean segment length $\mathrm{L}_{m,i}$, meaning that:

$$\lambda_i = \frac{1}{\mathrm{L}_{m,i}}.$$

Thus, λ_i quantifies the frequency with which the system transitions from state s_i to another state, with shorter segments corresponding to higher rates of transition.
Mathematically, the off-diagonal elements of the generation matrix G_{ij} are given by:

$$G_{ij} = \lambda_i \cdot P_{ij} \quad \text{for} \quad i \neq j,$$

where λ_i is the rate for state s_i, and P_{ij} is the transition probability from state s_i to state s_j.
The diagonal elements G_{ii} of the generation matrix are computed as the negative sum of all the off-diagonal elements in the corresponding row, ensuring that the total rate of leaving state i is equal to the sum of the rates of transitioning to all other states:

$$G_{ii} = -\sum_{j \neq i} G_{ij} = -\sum_{j \neq i} \lambda_i \cdot P_{ij}.$$

This formulation ensures that the system adheres to the properties of a continuous-time Markov process, where the time spent in each state follows an exponential distribution with the rate λ_i.
The generation matrix can then be used to model transitions between states in the system [4].

2.3 Forward/Backwards simulations

Simulating vehicle performance often involves two main approaches, backward or forward simulations.
In the backward simulation method, the driving cycle provides a speed that the vehicle experience. The required force and engine input power are computed by working backwards from the desired speed through the powertrain. This inverse dynamics approach simplifies the problem by treating the vehicle speed as a non-dynamic parameter, which generally results in computational efficiency. However, it may not capture the transient dynamics arising from environmental factors that prevent the vehicle from reaching the target speed. For example, environmental factors such as traffic signals, stop signs, sharp curves, and steep inclines can prevent the vehicle from reaching the target speed.[14].
In contrast, the forward simulation method treats the vehicle speed as a dynamic state. In this framework, a driver model continuously adjusts control inputs (e.g., accelerator and brake pedal actions) to follow a target speed provided by the environmental factors. The simulation computes the vehicle's response forward in time by solving a system of differential equations and incorporates an environmental model that accounts for factors such as traffic lights, stop signs, road curves, and steep inclines. In this approach, environmental information is sent to the driver model, which then dynamically changes speed and acceleration to maintain the desired speed as closely as possible. This method often yields a more realistic representation of actual driving behavior, as the driver not only strives to follow the target speed but also adapts to events and external constraints.[14].

3

Methodology

This chapter outlines the methodology used to gather, process, and analyze real-world road data for the purpose of simulating and modeling road topography. The process includes data collection from geospatial platforms, analysis of road characteristics, and subsequent modeling of road topographies using statistical approaches. It also contains an interview study which aims to identify the effect road topography has on modeling and simulation, and what is most important to capture. This study is performed to find the answer to RQ1.

3.1 Identifying road characteristics

This section aims to answer RQ1 "What are the key road topography characteristics?", by identify the characteristics of road topography through an interview study. This is to build an understanding for what benchmarks should be used, and what to test in the evaluations of the models.

3.1.1 Interview study of road characteristics

An interview study was conducted to identify the most critical characteristics of road topography for different engineering disciplines. This study was carried out through semi-structured interviews with experts from Volvo Group and Chalmers University. Each interviewee is an engineering expert in their respective field and utilize or have utilized road topography models for simulations. Each discipline provides a unique lens through which road characteristics are interpreted and analyzed. The participant were chosen to get a large range of disciplines represented.

3.1.1.1 Execution

The interviews where semi-structured, which is a mix between the two methods completely structured interviews and unstructured interviews. The semi-structured utilizes predefined questions, but allows the researchers and interviewees to expand these questions to broader subjects [15]. This method is preferred since the interviews can be adapted according to the expertise of the interviewees. The prior limited knowledge of the researchers leads to an adaptability in the questions, where the interview can be changed to capture the most important and interesting areas of the conversation. The disciplines and titles of the 4 interviewees can be seen in table 3.1.

Discipline	Title	Company
Energy consumption	Researcher of Energy consumption in BEVs	Chalmers
Thermal management	System Architect Thermal Management	Volvo
Startability & Durability	Researcher of durability for increased efficiency	Volvo
Drive line	Lead researcher of durability	Volvo

Table 3.1: Disciplines and participants of study

The questions aimed to identify what simulations were used in testing, what kind of simulations were run, and what effect road topography has in these simulations. The questions can be seen below.

1. Which parts of the topography affect your results?
2. How do you use road models?
3. What has the greatest impact on your simulations?
4. Which road characteristic affects your simulations the most?
5. Does elevation need to be controlled to a specific value in your simulations?
6. When using road profiles in previous work, how did you choose the models/roads?
7. How long are the roads you use in simulations?
8. What are your requirement specifications?

3.1.1.2 Discipline analysis from interview study

Energy consumption Energy consumption analysis aims to optimize fuel efficiency and minimize energy losses in both conventional and electrified powertrains. For this interview the area focused solely on Battery Electric Vehicles (BEV). It considers factors such as aerodynamics, rolling resistance, and driving cycles. Here topography is a part of driving cycles, as a part of the road model. Simulations are run to verify the models and calculations are correct against real roads, and to examine how the driving cycle affects the energy consumption. The factors that most impact the energy consumption are total height gain, difference in height from start to end, and traveled distance. This can be explained by the road topography characteristics of slope inclination, length of hills, and elevation.
The interviewed uses simulations to check if their calculations and models are correct. The focus areas are long stretches of road, where the main focus is the inclination and length of hills. BEVs uses regenerative braking, which is the effect of using the engine to brake, which charges the battery. To ensure the regenerative braking is tested in the simulation, it is important that the model generate downhills that are steep enough so braking is needed. However on very steep descents, or when intense braking is needed, the regenerative brakes might be insufficient, and service brakes need to be applied. If this is desired to examine in the simulations, the models must be able to generate very steep downhill slopes. The simulation requires that the max inclination is not too high, since a too high max inclination can lead to the vehicle not being able to cross the hill, and potentially terminate the simulation.
For the entire mission, it is important to acknowledge what the altitude change between start and stop is, since this greatly affects the energy consumption. The effect of this can quickly be calculated with the potential energy equation $E = mgh$, where

the **E**nergy is equal to the **m**ass times **g**ravity constant and **h**eight. Given a truck with a weight of 55 tons, each extra meter in height would increase the consumption as $55 \cdot 9.82 \cdot 1/3600 = 0.151 kWh/m$. Using this and the fact that a 1% increase in inclination gives 10 meters height gain over a km, we get $1.51 kWh/(km \cdot \%)$. If comparisons between different missions are made, it is therefor important that the height difference is the same.

Thermal management The field of thermal management aims to examine the temperatures in vehicle, how they affect performance, and use efficient heat regulation to improve performance, longevity, and safety. Thermal management ensures optimal operating temperatures for critical vehicle components, including the engine, battery, cooling system, and cabin climate control.
When running simulations, roads are selected from a predetermined list depending on which GTA class that is tested. The roads are modeled after real roads, each representing the toughest in its classification. How the toughness of the roads is calculated, and why they were selected is not widely known. More information about the characteristics of what makes both a specific road, and roads in general tough is desired by workers. The roads used in simulations have been the same for a long time. For running simulations, there is a big difference in test between BEVs and Internal Combustion Engines (ICE)s. The combustion vehicles are tested on hills that are infinitely long, only changing the grade. This is since the cost of cooling a ICE is not very high in volume or monetary, so it is worth over dimensioning the cooling system. This means models does not need to be realistic, since they only need the slope model parameter. For BEVs the length of the hill is fixed, since BEVs need more cooling per traveled time which means that length and order of hills are important. The simulations are often short, never longer than a workday, and test specific scenarios. A desire for longer roads is sought after, especially for BEVs since the cooling during both charging and downtime affect the performance.
The current approach to selecting road for simulation has the weakness that what is a tough variant is not clear, and does not have statistical background, but instead is determined by each expert. This makes comparisons hard, since no objective values can be examined, and knowledge why the topography is considered tough can be lost along the way. It is also unsure how these roads are representative of their respective classes.

Startability Startability refers to a vehicle's ability to start and move in steep inclines. Startability specifically tests that a vehicle can start given a specific slope value. For simulations this means that only the max inclination is of concern. These test only need one hill to examine if the vehicle completes its specification, but this slope value must exist for the test to be relevant. This means that if one hill on the road with a specific slope exists, that is enough to verify startability.

Durability Durability engineering ensures that vehicle components can withstand prolonged mechanical stress, environmental factors, and operational wear over time, maintaining performance and safety. Here the focus is on gearbox, damping and brakes. For durability, the relationship between wear and gradient is exponential,

as increased loads on the components such as the gearbox, result in significantly higher wear. This means that a short steep slope will introduce higher stress than a flatter longer slope. This means that it is essential to get a good representation of the variation of hills, both with regards to the length, slope and amount.
For brakes the length of the downhill is of importance, since the brakes get hotter in long hills, which leads to greater wear and tear. This affects BEVs as well as ICEs, since at a point regenerative braking is not enough and conventional braking needs to be applied. An extreme case is if the battery gets fully charged thereby disabling regenerative braking, and then reaching a downhill that is too steep for its conventional braking system. The gearbox deteriorates faster if only a single gear is used for long stretches. To capture this the road model needs long stretches of the same slope.

Driveline The driveline transfers power from the engine or electric motor to the wheels, including components such as the transmission, differential, and axles. Efficient power delivery improve performance, efficiency, and handling.
In powertrain analysis, ensuring the reproducibility of results is of utmost importance. This requires maintaining consistent test conditions while modifying only the variable under investigation. Durability assessments for powertrain require long-term verification, typically conducted over intervals of 12, 24, and 36 months. These tests are performed on standardized road tracks that have remained unchanged for decades, allowing for the accumulation of extensive datasets and providing a high degree of certainty in the results. As a result, durability testing heavily prefer predefined, fixed road routes. This limits the applicability of stochastically generated road models due to their inherent variability. Consequently, artificially generated road profiles are not prioritized for durability assessments. However, in fuel efficiency testing of the driveline, a more diverse set of road conditions is examined to better reflect real-world driving scenarios. In this context, generated road models offer significant advantages, as they naturally encompass a broader range of topographical features, thereby enhancing the representativeness of the simulations.

3.1.2 Essential characteristics for road topography

From the interview study of employees at Volvo and Chalmers, some themes can be gathered for what is important for realistic and relevant simulations. The most important factors were slope and length of hill, which affected all areas. Positive and negative slopes greatly impact energy consumption [16]. This means that the resulting models need to prove that they capture representative information about grade and length, and keep this in the produced roads. To ensure that they achieve this, models will be validated with heatmaps, which is explained in section 3.1.3.
For BEVs, it was especially important to get the length of hills correct, since they are more affected by battery temperature, where the efficiency greatly decreases. They also need to have a good balance of up- and downhills, since they utilize regenerative braking [17], and batteries warm up during regeneration in downhills. For angle of slope, every area is affected, but especially startability and driveline, where it is important to capture the max inclination, a specific angle that the vehicle must

manage to overcome. To ensure this is captured in the model, a model that fulfills this criteria can be chosen, or a specific function can be made to add one slope of the specific inclination and length at a random position in the road.

In the evaluation part of the project comparing the road characteristics for the generated roads from different models with the real roads will be done. But also simulations will be performed. To save time and effort for these simulation certain simpler backward simulations will be used 2.3. When examining the driveline performance, one common approach is to look at the torque generated by the engine, since this affects the whole driveline. Specifically the areas of maximum torque, average torque, and level crossing is of value for comparing results. Level crossing is the method of counting every instance that the output is equal or greater than a value, and plotting it [18] (see Section 3.1.5)

3.1.3 Heatmap of slope lengths

In this chapter it is shown how heatmaps over the relationship between grade, length of hill, and number of instances is created, providing a clear visualization of their distribution across real-world and generated road data. Real-world road segments were divided into segments based on slope, with each segment characterized by its average grade and length. The data is analyzed on the relationship between length and grade of hill, with one result plotted in Figure 3.1. For each instance of a combination of slope and length of hill, the number increases. The values are calculated from rounding to the closest value.

This is used to compare the patterns in the real-world to the output of the generated roads. This analysis highlights the utility of heatmaps in identifying discrepancies between real and generated roads, providing a foundation for refining models and validating against GTA classifications.

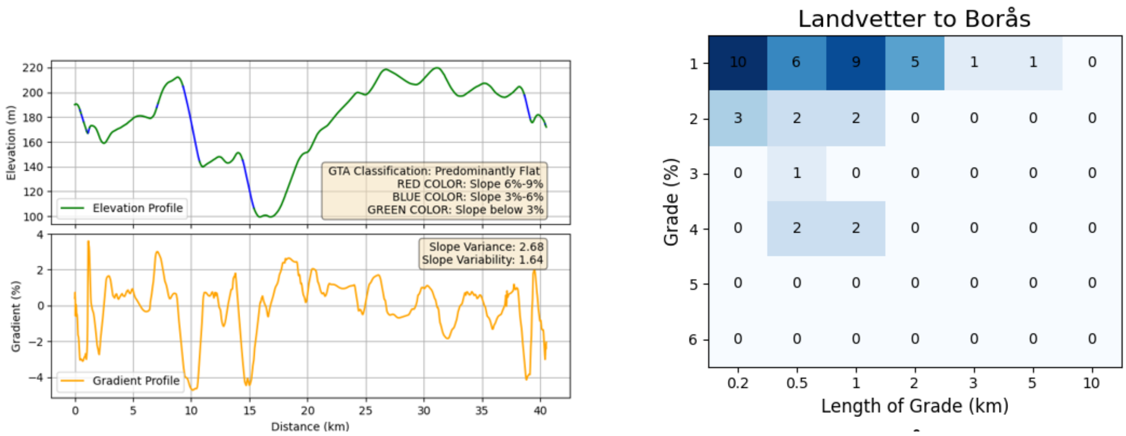

(a) Elevation and slope profile

(b) Heatmap

Figure 3.1: The road between Landvetter to Borås, showing elevation, and slope profile, and the corresponding heatmap

3.1.4 Grade segmented hills

From the definition of a road and its topography explained in section 2.1.1, together with the information from the interview study in section 3.1.1, it is clear to see that the most important characteristics when examining and comparing road topography is length of hills, and their inclination. For this comparison and examination to be valid and interesting, roads need to be compared on the same terms. The problem of classifying the slope and length of a hill, as is done in a heatmap in section 3.1.3, is that if a hill changes gradient in the middle, it is counted as two different hills, even if the slope continue upward the whole time, see Figure 2.1b. If the larger picture instead is of interest, i.e. how long do a section of slopes continue before the road starts to decline, the definition of a hill seen in Figure 2.1c can be used. This definition captures the larger flow of the topography. To achieve this a model that simplifies the road is utilized, where the start and end of a hill is found, and the slope is averaged out between these two points. A realization of this method can be seen in Figure 3.2a.

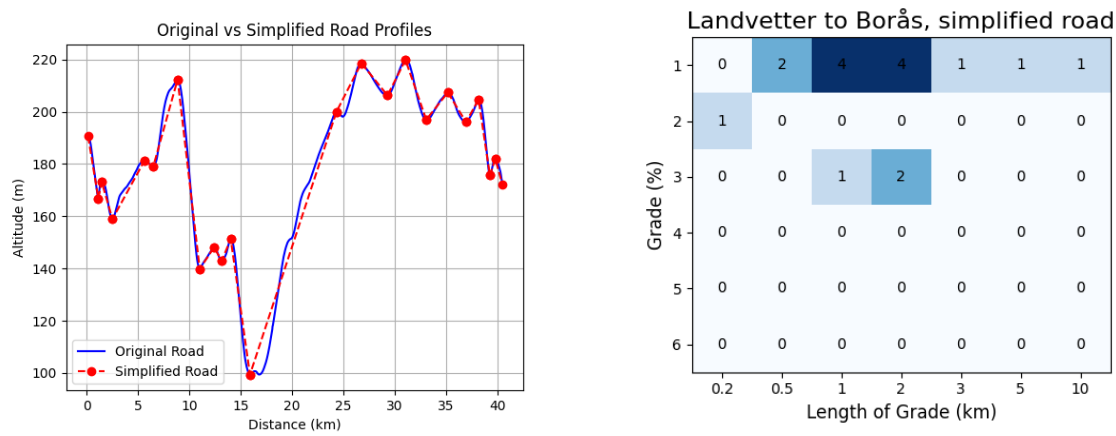

(a) Elevation profile

(b) Heatmap

Figure 3.2: Grade segmented version of road between Landvetter to Borås, showing elevation profile, and the corresponding heatmap

The cutoff for this method is that the slope must change from positive to negative or vice versa, and move at least 2 meters in the other direction. This is to ensure that the change is significant and not just a small bump in the road before returning to the same slope value. The result of this grade segmentation can be seen in the heatmap Figure 3.2b, which when compared to the original heatmaps in Figure 3.1b, the grade segmented road have longer hills, and lower grades. The heatmap has fewer instances, since the average length of hills has gone up. This result is used for smoothing out generated model roads, to be able to examine the topography on a larger scale. This introduces an alternative validation method for identifying how and if the generated roads capture the characteristics of the input data.

3.1.5 Level crossings

Level crossing is the method of counting each occurrence of a parameter exceeding a specific level. Level crossing analysis is a widely used technique in signal processing and time series analysis, applied here to assess variations in road profiles and output from simulations. The method involves counting the instances where a signal, in this case the road elevation, slope profile or torque/power output, crosses a specific threshold level. This level can be either one level, such as the zero crossing that grade segmented roads utilizes in Section 3.1.4, or every level, which is then plotted in a graph. This approach provides insights into the frequency and magnitude of elevation or slope changes, which are crucial for evaluating road classification.

The level crossing method operates by defining a threshold level, denoted as L, and identifying points where the road profile crosses this level. Mathematically, a level crossing occurs at time t when the road profile satisfies:

$$h(t - \Delta t) < L \quad \text{and} \quad h(t) \geq L \tag{3.1}$$

for upward crossings, and

$$h(t - \Delta t) > L \quad \text{and} \quad h(t) \leq L \tag{3.2}$$

for downward crossings, where Δt is the sampling interval [18].

An example of level crossing can be seen in Figure 3.3, where the slope profile 3.3a gets converted to a level crossing plot in Figure 3.3b. The angle of the slope is the y-axis, and the x-axis the number of occurrences of the specific slope grade. The more instances of a crossing, the further the line shifts to the right. . Here we can see that the most common grade is around positive 0.5%, and that the spread is from [-5 : +3.5] %. The advantage of the level crossing is especially apparent when comparing different profiles, since then the intensity and spread can be analyzed.

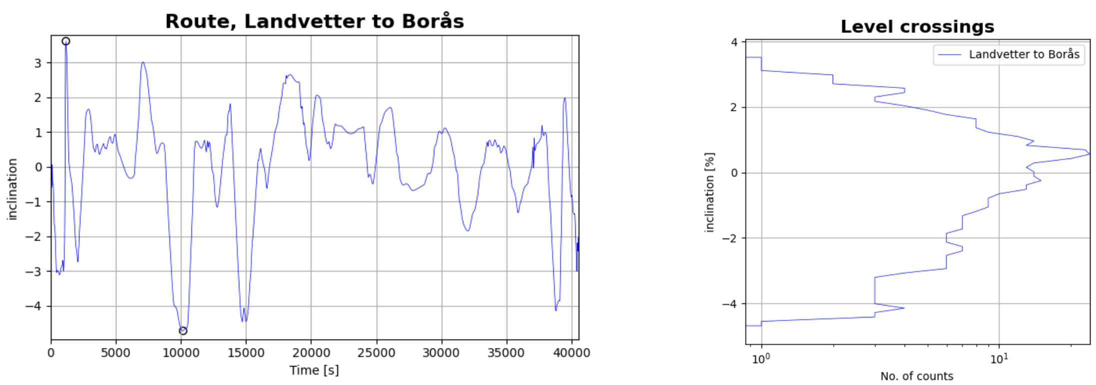

(a) slope profile

(b) Level crossing profile

Figure 3.3: Slope profile for Landvetter-Borås, with slope profile and the level crossing plotted

3.2 Road data processing

To perform the work, road profile data is needed to analyze road characteristics, build model parameters, simulate roads, and verify results. To get accurate and precise data, road profile data was gathered from the HERE maps database [19], a provider of routing datasets. HERE Maps collects data through a combination of satellite imagery, vehicle-based mapping, and crowd-sourced contributions. These diverse data collection methods ensure a broad representation of road conditions and characteristics, although resolution and sample rate may vary between regions.
The process of data acquisition involves identifying and selecting roads from various terrains, including urban, rural, and mountainous areas, to ensure the dataset captures diverse topographical features. Researchers manually curate the roads of interest to ensure relevance to the project objectives. To maintain consistency in the data set, the road profiles is sampled at equal distance intervals. This ensures that the analysis is not biased by varying sample densities along different roads.
The data retrieved from HERE Maps include important attributes such as latitude, longitude, elevation, distance, slopes, speed signs, curvature, headings, and functional class. To test the quality of the data, the elevation was compared to the slope data. The altitude was calculated from the slope data, using the equation

$$h = h_0 + sin(\theta) \cdot L_s \quad (3.3)$$

where h is the altitude, h_0 is the starting altitude, θ is the slope data, and L_s is the sampling distance. The result can be seen in Figure 3.4, which shows that the error over long distances is very small, less than 1 m over a 45 km distance.

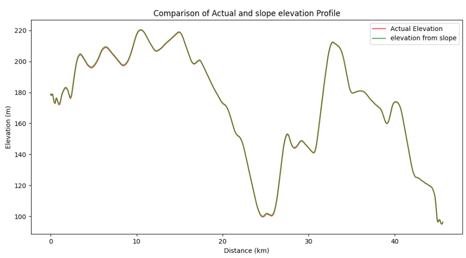

(a) Elevation profile

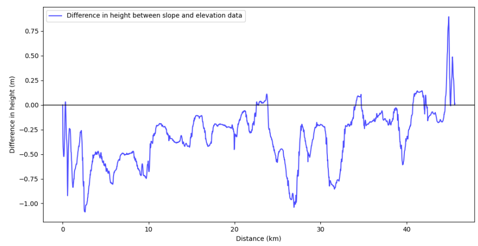

(b) Difference between elevation and calculated road

Figure 3.4: Verification results using Height difference between elevation data and elevation calculated from slopes data of Borås-Landvetter

3.2.1 Analyzing roads

A critical aspect of this research involves analyzing real-life road data to ensure the accuracy of the simulated models. The most important parameters, simulations methods, and characteristics were discussed in the interview study, section 3.1.1. To identify if the models produce generated results that match real-world data, and data that is relevant for studies, an examination of the input data needs to be performed.

To analyze the data from HERE maps and see what differentiates the 4 GTA Topography classes, a heatmap can be seen in Figure 3.5. The road data is normalized so each classification has the same amount of kilometers driven. The data shows that the spread is as expected according to the definition 2.1, with Flat having the lowest slopes, and the slope steadily increasing the higher one goes in the classifications, ending with Very hilly having the highest slope grades. The most common combination for all classes are the smallest slope and shortest distance, which is to be expected since this is most similar to a flat road, and the shorter the measuring distance, the more instances can occur.

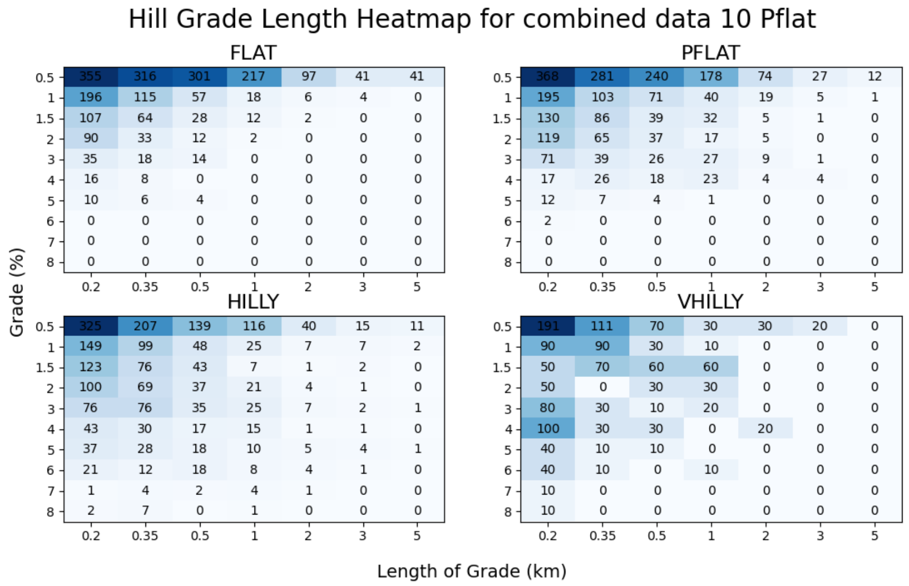

Figure 3.5: Examples of heatmaps of multi grade definition for the 4 GTA classifications

To perform analyses on the models, the data for the PFLAT road will be used as training data. This dataset consists of 9 roads with varying degrees of parameters, but all falling within the PFLAT category. These roads can be seen in table 3.2

Route	**Total Distance (km)**
Arvika-kil	53
Borås-Landvetter	45
Budapest	55
Gothenburg-Kiruna	1,643
Helsingfors-Turku	166
Manchester-Hull	167
Prag-Berlin	347
Prag-Wien	334
Umeå-Vindeln	67
Total	2881

Table 3.2: Total distances for different routes

3.2.2 Functional classes

Another thing that is gathered from Here maps is labels of the roads called functional classes, which categorize roads based on their importance and role in the transportation network. This will later be used when considering splitting up a road into different road types. The definition of functional classes in this project is that of HERE[19], which divides roads into five functional classes as follows:

Functional Class 1: These roads facilitate high-volume, high-speed traffic between major metropolitan areas. They have few speed changes and controlled access, typically highways or freeways.

Functional Class 2: These roads channel traffic to Class 1 roads for efficient city-to-city travel, supporting high-speed, high-volume traffic but with more speed changes than Class 1.

Functional Class 3: Roads that intersect with Class 2 roads, providing high traffic volume but at lower speeds and mobility. They serve as connectors between important regions.

Functional Class 4: These roads handle moderate-speed, high-volume traffic between neighborhoods, connecting with higher-class roads to distribute traffic within urban areas.

Functional Class 5: Roads with lower traffic volume, including neighborhood roads, access roads, and internal connections. This class also includes walkways, truck-only roads, bus-only roads, and U-turn lanes unless part of an intersection.

3.3 Implementing Markov and Autoregressive Models

Markov methods, or Markov models, have previously been proven to be able to reliably stochastically model road modeling in curvature by Maghsood and Johanesson [20]. Previous work has also modeled elevation as a Markov Chain, focusing on transitions between discrete elevation levels to predict chassis loads influenced by road roughness [21]. This prior work provides a foundation for the more general stochastic modeling of road topography pursued in this project. Specifically, two types of Markov-based models will be explored, the AR(1) model and a discrete-state Markov Chain variant. In addition to AR(1), another autoregressive model, ARMA(1,5), will also be examined to capture more complex temporal dependencies in the road topography data. Furthermore, modifications to these models will be considered to address specific challenges encountered during the modeling process.

3.3.1 Implementation of the AR(1) model

The first model implemented and evaluated is the AR(1) model, originally introduced in Section 2.2.1. The parameters of interest in this model are the correlation coefficient a and the total variance σ_y^2, which will be estimated based on the observed data. Based on the estimation method mentioned in the theory section 2.2.1.

3.3.2 ARMA Approach for Road Profile Generation

To evaluate the performance of this approach, a Python library for time series analysis was employed for parameter estimation and model validation, called statsmodels [22]. This library provides robust tools for ARIMA modeling, where parameters are typically estimated using Maximum Likelihood Estimation (MLE). In MLE, the likelihood function is optimized iteratively to obtain the best-fitting values for the autoregressive coefficients, moving average parameters, and the residual variance.

This approach can be tested following a similar procedure done using the AR(1) model. Meaning that model parameters can be estimated using the mentioned library to fit model parameters for real road slope profiles from HEREmaps. And the model could then be used to try to generate similar new slope profiles. However, incorporating the differencing component (d) in the ARIMA model did make the slope profile have a strange behaviour, so it was set to zero, reducing the model to an ARMA formulation.

When the ARMA model was considered, it was found that having a higher order of the AR process or MA order was making the abrupt changes in slope observed with the AR(1) model before significantly moderated. This effect is illustrated in Figure 3.6, where the left image shows a road profile generated with the normal AR(1) process while the right image shows a road generated road profile using the AR(1) process but with the moving average parameter set to 5. Making it a ARMA(1,5) process.

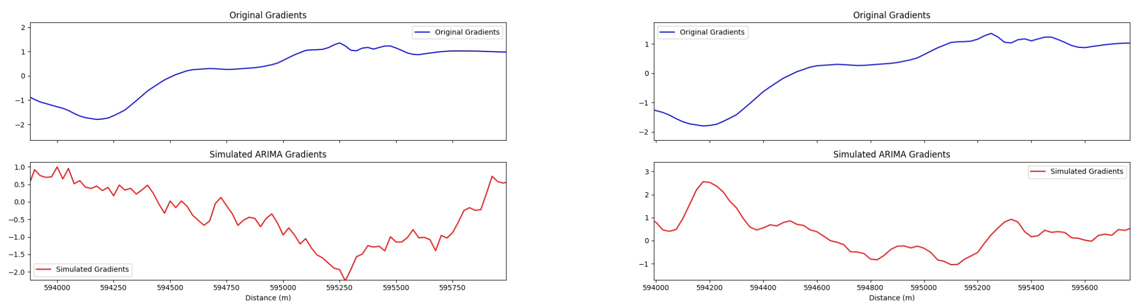

(a) Without MA **(b)** With MA (set to 5)

Figure 3.6: Comparison of road profiles: (a) A generated slope profile from the AR(1) model without a moving average component and (b) A generated slope profile from the ARMA model with the moving average parameter set to 5.

The observed change in slope, approximately ±0.2% over a 100-meter distance, could be interpreted as road roughness. This observation may explain why the ARIMA approach is considered effective for generating road segments over shorter distances. Nonetheless, the ARMA(1,5) model, which extends the AR(1) process by incorporating a moving average component of order 5, will be representing a more complex candidate to be validated alongside other approaches. The ARMA(1,5) model is mathematically expressed as,

$$y_k = a\, y_{k-1} + \theta_1 e_{k-1} + \theta_2 e_{k-2} + \theta_3 e_{k-3} + \theta_4 e_{k-4} + \theta_5 e_{k-5} + e_k, \quad e_k \sim \mathcal{N}(0, \sigma_\epsilon^2). \tag{3.4}$$

It is of course possible to compute the stationary distribution from the ARMA process's parameters, yielding the long-run variance and enabling conversion from specific ARMA parameters to GTA classification. However, this method is considerably more complex than the equivalent calculation for an AR(1) process.

3.3.3 Switching slope values based on a Markov chain process

This approach follows a similar methodology to the one used later in Section 3.4.2, where transitions between different road characteristics, such as speed limits, functional classes, or terrain types, are modeled using a continuous time Markov chain processes 2.2.3. However, instead of switching between these features, the model directly transitions between slope values to construct a slope profile.

At every 25-meter interval, which corresponds to the sampling frequency used throughout this project, a new slope value is generated based solely on the current slope value. Data can be analyzed to derive a transition matrix that represents the probability of transitioning from one slope value to another. Since the model updates the slope at every step without a fixed dwell time in each state, a generation matrix is unnecessary. This is because slope values commonly fluctuate slightly rather than remaining exactly the same. Instead, a transition matrix with self-transitions is sufficient, allowing the model to remain in the same slope state. These self-transition probabilities tend to be higher than the probabilities of switching to a new slope value, as slopes often remain consistent over short distances.

The resolution of the transition matrix, meaning the number of decimal places considered, can significantly impact the model. If all values are rounded to whole numbers, the generated hills become unnaturally straight, potentially reducing realism, as illustrated in Figure 3.7a, where the corresponding transfer matrix is also shown. This matrix is derived from data collected along a HERE Maps route from Gothenburg to Kiruna.

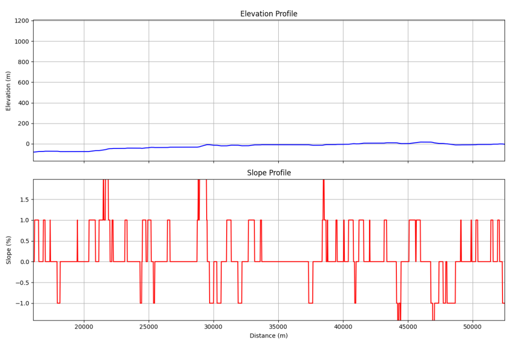

(a) Slope profile when only switching between integer slope values.

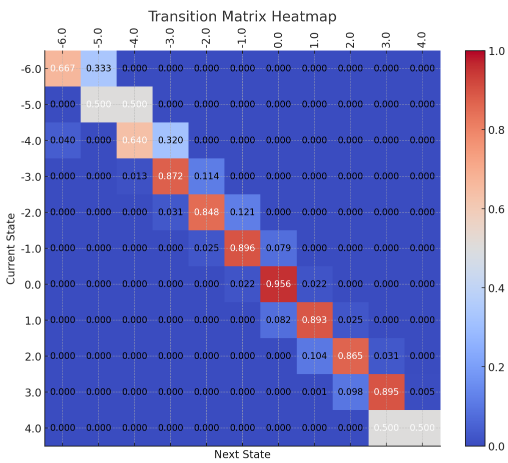

(b) Transition matrix of slope values derived from a route between Gothenburg and Kiruna.

However, this simplification makes the model easier to analyze and interpret. As shown in Figure 3.7b, it is evident that remaining in the same slope value is the most probable transition, while shifts to nearby slope values also have relatively high probabilities.
On the other hand, increasing the resolution results in a much larger transition matrix and requires a more extensive dataset to ensure sufficient representation of all transitions.
It is important to note that this project focuses on modeling road topography rather than road roughness. Therefore, small variations in slope do not need to be captured with high precision.

Statistical properties and control

Similar to the AR(1) process discussed earlier, the Markov Chain process used in this approach also provides useful statistical properties that can be exploited for analysis and control.
One key property is the stationary distribution π, which represents the long-term probability of being in each slope state, see Section A.1. Given a transition matrix P, the stationary distribution satisfies the equation

$$\pi P = \pi \tag{3.5}$$

where π is a row vector whose elements sum to one. This equation implies that, over a long enough generated sequence, the proportion of time spent in each slope state converges to π. The stationary distribution can be obtained by solving the equation

$$\pi(P - I) = 0 \tag{3.6}$$

and normalizing it so that the sum of all elements equals one. This normalization ensures that the stationary distribution represents valid probabilities.
With the stationary distribution, the probability of the generated road profile exceeding a certain slope value can be computed directly. For example, the probability of the slope being greater than a threshold $s_{\text{threshold}}$ is given by:

$$P(S > s_{\text{threshold}}) = \sum_{s > s_{\text{threshold}}} \pi_s \tag{3.7}$$

where π_s is the stationary probability of being in slope state s. Given that the generated road is sufficiently long, this probability approximates the fraction of the total road distance where the slope exceeds the threshold. This allows classification of the generated road into GTA categories in a similar manner to the AR(1) approach, ensuring that the generated road topography adheres to the GTA classification criteria, which can be a useful attribute when generating roads for vehicle simulations.
Another advantage of this Markov chain based approach is that the length of hills with specific slopes can be controlled by modifying the transition matrix. Specifically, increasing the self-transition probability $P(s, s)$ for a slope value s will increase the expected time in that state. The expected number of consecutive segments spent in state s before transitioning away is given by:

$$E[T_s] = \frac{1}{1 - P(s, s)} \tag{3.8}$$

where T_s represents the number of consecutive steps spent in state s before transitioning to a different state. Thus, by increasing $P(s, s)$ while normalizing the row probabilities accordingly, one can extend the average length of hills with that specific slope. Similarly, adjusting transition probabilities to favor certain slope values can increase the frequency of occurrence of those slopes within the generated sequence. These adjustments provide a straightforward way to shape the generated road profile to meet specific statistical or classification requirements. For example, if it is determined that slopes of 3% should have a mean length of 50 meters, this implies that the expected number of consecutive segments in the slope-3 state should be 2, given a sampling interval of 25 meters. To achieve this, one can solve for $P(s, s)$ in Equation 3.6, insert the computed value into the corresponding row of the transition matrix, and then normalize the row to ensure a valid probability distribution.
This level of control, combined with the statistical properties derived from the transition matrix, makes the Markov chain model an effective tool for generating realistic road topographies while still having the ability to affect what roads to be generated, which sometimes is the case when generating road topography for vehicle simulations.

3.4 Modification of the AR(1) model

Each of the three implemented models in the precious section demonstrated certain strengths but also exhibited limitations. In particular, when evaluated using heatmaps of hill lengths, used as a validation metric, their behavior deviated from that observed in real-world data obtained from HERE Maps. The following chapter focuses on the AR(1) model and its modifications, aiming to achieve improved alignment with real world topographic characteristics.

3.4.1 Plotting AR(1) model Parameter Convergence

To better understand the behavior of parameter estimation in real data, how the estimated parameters evolve over different road profiles is explored. Specifically, data from two different slope profiles, one from Zurich to Milano and another from Oslo to Havstad, were used. The estimated parameters were plotted alongside the corresponding elevation profiles to examine convergence time and the effect of terrain changes.
Figure 3.8 presents the results, where the first two plots display the estimated $\hat{\sigma}_y$ parameter, while the third and fourth plots show the estimated auto regressive coefficient,$\hat{a}$. Every plot is plotted over distance and with the elevation profile as reference.
The convergence behavior of $\hat{\sigma}_y$ and $\hat{a}$ strongly depends on the terrain, which is expected. This is particularly evident in Figure 3.8a, where $\hat{\sigma}_y$ and $\hat{a}$ initially stabilize but are later significantly affected by the Alpine mountains. When reaching the peak of the hill, the variance increases as the terrain transitions into a steep descent, introducing larger negative slopes.
In contrast, Figure 3.8b shows a more gradual shift, as the terrain changes less abruptly. However, even in this case, the parameters do not remain within a fixed

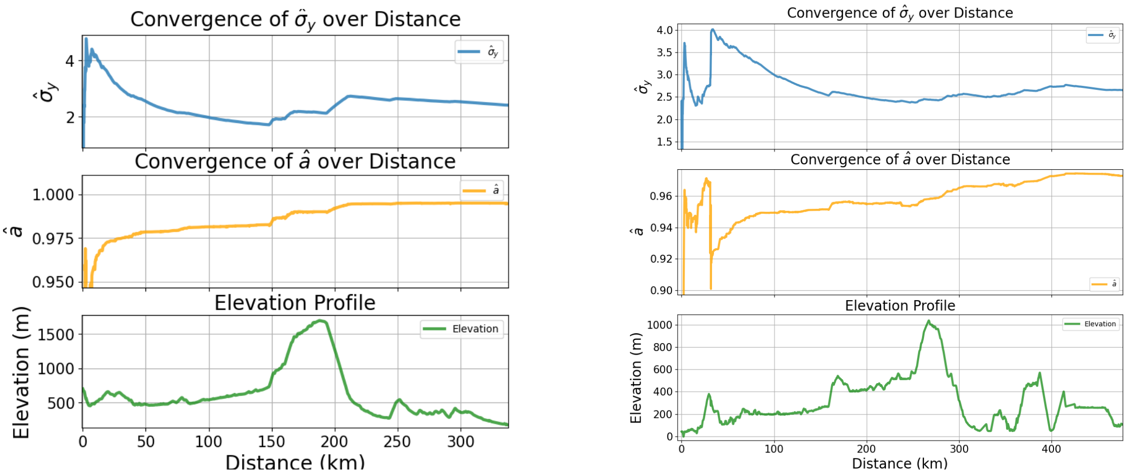

(a) Zurich-Milano $\hat{\sigma}_y$ and $\hat{\sigma}_y$ Convergence

(b) Oslo-Havstad $\hat{\sigma}_y$ and $\hat{\sigma}_y$ Convergence

Figure 3.8: Convergence plots for 2 roads, one from Zurich to Milano and one from Oslo to Havstad. The plots show the estimated $\hat{\sigma}_y$ and $\hat{a}$ parameters up to the given distance. Having them together with the corresponding elevation profile to see how changes affect the parameters.

interval, reflecting continuous terrain variations.
These findings highlight an important consideration for generating road profiles using the AR(1) model. The choice of input data directly affects the estimated parameters, meaning that including only a specific terrain type (e.g., flat or mountainous roads) will tailor the model to that particular topography. Including multiple terrains will produce a more generalized road profile, potentially smoothing out key terrain variations.
This might introduce a challenge, a generalized model may fail to capture the distinct characteristics of terrain-shifting road profiles, which could be more realistic. To address this, one possible approach is to divide the road into segments, estimating specific model parameters for each segment. This method is discussed further in Section 3.4.2.

3.4.2 Combining Ar(1) and Markov chain for road segmentation

Using only an AR model or solely slope-based Markov chains has led to several challenges, especially when generating trajectories over long distances. As discussed previously, the performance of these models depends on the spatial scale considered (e.g., below hundred meters versus below ten thousand meters). Over extended distances, the current models tend to exhibit unrealistic behavior. For example, they do not account for variations in road types. A common scenario is a truck transitioning from a highway to a local road, which alters driving dynamics considerably. Since, highways must adhere to strict construction regulations, whereas smaller roads may have a more varied elevation profile. Additionally, differences in terrain is also something that the single models so far would not be able to handle. Meaning that some areas may possess unique road topography characteristics.

To overcome these limitations, a hybrid approach is proposed. This approach involves using a base model, given its relative performance and ease of implementation, the AR(1) model will be used as the base model for this project. Also, the AR(1) model has been used for this in previous work [4]. ARMA(1,5) or slope-based Markov chain could of course also be chosen. And the augment this base model (the AR(1) model in this case) with a parameter-switching Markov chain that capture more macroscopic effects. Meaning having different AR(1) parameters for different states. Considering this there are two different approaches explored in this project. The first is to have a Markov chain that switches between different road types, meaning having separate parameter values for different road types, and the other way is to switch between terrain types.

Road type dependency

For the first approach, road type dependency, there are two ideas explored here for dividing the road into different road types.

One method to classify these road type segments is based on speed limits, and then separate AR(1) parameters can be estimated for each speed limit. This method was tested following the approach described in [4], where the data is first grouped into road types based on the speed limits (30, 50, 70) km/h. Meaning that road segments are separated into predefined speed categories. The process of transition between different speed limits in a realistic way, when generating a new road, is implemented by using a Markov chain process, see 2.2.3. The speed limits considered in this study are 30, 50, 70, 80, and 110 km/h. Any deviations from these values were rounded to the closest speed limit category.

Using this approach, separate AR(1) parameter values can be assigned to different speed limits, and the transitions between these states can be modeled realistically, based on either an assumption of which roads the driver drives on, or directly taken from the road data from HEREmaps.

Another way of dividing roads into road type segments is done by using functional classes. This classification organizes roads by their role and importance in the transportation network. Just like speed limit, these functional class labels are gotten directly from HERE maps, where every road is label with a number 1-5, see section 3.2.2. Just like with speed limits, the approach taking was using a Markov chain process to switch between functional classes in a realistic way when generating a road, see section 2.2.3.

Using functional classes can sometimes be more useful than relying on speed limits. For example, low-traffic roads in rural areas often have a speed limit of 70 km/h, which is common for roads between towns, but they serve a different function than higher-traffic roads with the same speed limit. It can be suggested, for instance, that more effort was made to flatten road gradients on roads with higher traffic compared to those in rural areas. Also, functional classes from HERE is supposed to be more of a global standard where as speed limits can vary more between countries.

Terrain Type

And a third way to segment roads, is by considering different terrains. One could for instance directly analyze the road characteristics for trucks, by observing how often trucks encounter hills of various lengths. This data could then be used to create a Markov chain. This approach is somewhat similar to the method described in the section on Markov slope chains for road topography modeling, Section 3.3.3, though here the length is not considered only slope changes from HERE maps.
Another approach explored more in this project, is to divide the road into different terrains, by observing when the key parameters fall within a certain interval. When the parameters shift significantly into a new interval, this can be considered a transition to a new terrain type. The road can then be split into these terrain segments, and a Markov chain process can be used to model the transitions between them. For this to work smoothly one can consider only one road type, for instance highway, since this will be the main road type for long drives. It might be noteworthy to say that for a model to be dependent on both terrain and road type, another dimension of the transition matrix would be needed, making the approach more complex and increases the risk of deadlock significantly.
However, since the AR(1), model here is estimated based on the slope variance and the one step covariance. Identifying terrains based on significant changes of these could be a good idea. Since one step covariance seems to be more correlated with the road type and hard to capture as a terrain, only the variance of the slopes will be considered for this project. Meaning that areas where the variance seems to be behaving similarly can be determined using change point detection methods. For this purpose, the Python library Ruptures is highly effective [23]. Ruptures provides several algorithms (such as the Pruned Exact Linear Time (PELT) method which was used for this analysis), that can detect shifts in statistical properties like variance. The penalty parameter used for the PELT algorithm plays a crucial role in this process. Essentially, the penalty controls the trade-off between model complexity and the goodness of fit. A higher penalty requires a more significant statistical change for the algorithm to register a break, resulting in fewer detected change points. Conversely, a lower penalty makes the detection more sensitive, which may lead to over-segmentation. Thus, tuning the penalty value is critical. In this case, a value of 400 was found to provide a balanced segmentation that captures meaningful shifts in variance without introducing excessive noise. A result from using this on a road from Gothenburg to Kiruna can be seen in Figure 3.9.
It is important to note that the scale of analysis plays a crucial role, since estimating the AR(1) parameters typically requires a sufficiently long segment to converge, (on the order of tens of kilometers, depending on the variability of the terrain), otherwise the process wont capture the underlying dynamics reliably.
Another noteworthy insight when studying roads is that in some cases, such as observed in parts of Norway, there are prolonged segments with very persistent slopes. This results in continuous uphill profiles, see Figure (3.10). While these extremes are notable, they are relatively rare and might often be better accounted for by for instance a high auto regressive coefficient, if the AR(1) model is considered. Another method could be to incorporate a Markov distribution that occasionally adds a larger hill, but this idea is not tested in this project.

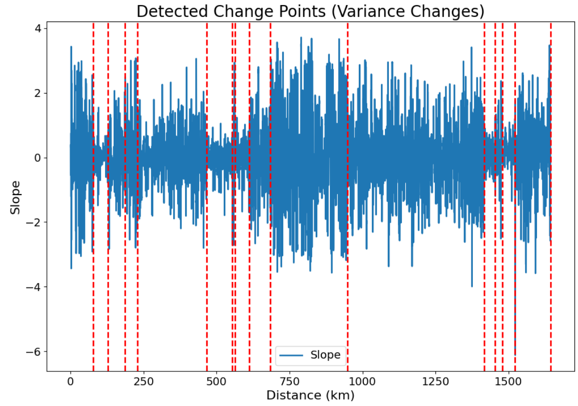

Figure 3.9: Splitting up a road from Gothenburg to Kiruna into different segments based on intervals with similar variance, identified by the PELT algorithm from the Python library Ruptures.

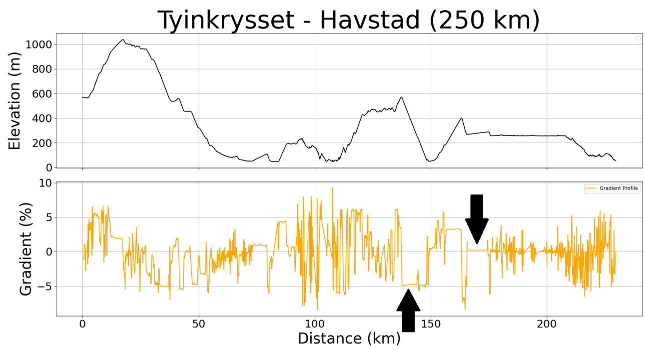

Figure 3.10: A road in Norway showing prolonged segments of equal slope.

Another idea for terrain modeling can be to consider different altitude as different terrains. Where supposedly higher altitude would mean higher slope variance. This correlation might exist to some degree, but as seen in Figure 3.11. Flat plains, low slope variance, can exist at high altitudes. And hilly areas, high slope variance, can exist at sea level.

Due to lack of time, only the road type models (speed limits and functional class) will be evaluated in this project. And the terrain type is left as an exploration of an alternative idea.

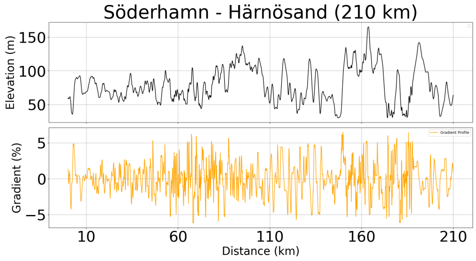

(a) Söderhamn to Härnösand (210 km), Slope variance: 1.61, mean altitude: 78.7m

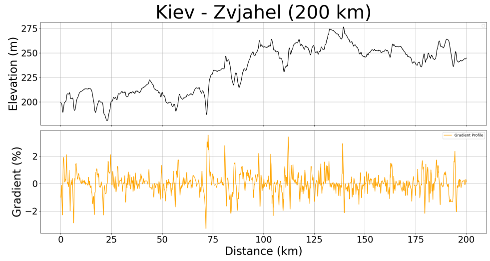

(b) Kiev to Zvjahel (200 km), Slope variance: 0.17, mean altitude: 233.3m

Figure 3.11: Elevation and Slope profiles for 2 mostly highway roads

3.4.3 Altitude consistency and mean reversion

While the combined AR(1) and Markov Chain model offers some benefits, it did not resolve the issue of unrealistic hill lengths compared to those observed in real-world roads, when the heatmap validation metric was considered. One possible explanation is that the models considered so far rely only on the previous slope value, leading to a kind of "wandering" behavior. In this setup, the topography lacks any memory of altitude, which can result in the formation of unnaturally large mountains (the elevation profile wanders with no knowledge of the current altitude). In contrast, real world roads typically exhibit a more balanced pattern, where an uphill is often followed by a downhill. This creates a landscape with frequent, smaller hills, while still occasionally allowing for longer climbs or descents.
One approach to addressing this problem is to introduce mean reversion by applying a weight that "pulls" the generated road back toward a predetermined mean. In this method, the AR(1) model is maintained while an elevation-dependent noise term is added, meaning that the further the current elevation deviates from the specified mean, the greater the noise. The effect of this mean reversion noise is illustrated in Figure 3.12.

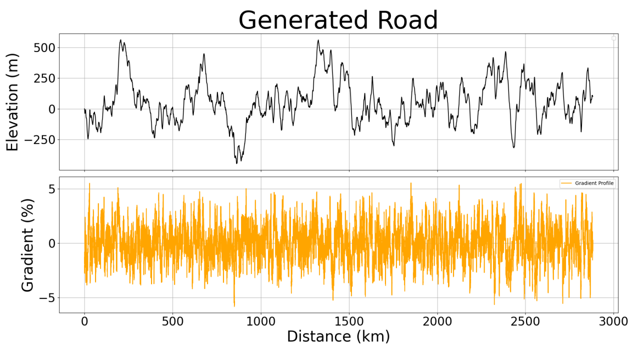

Figure 3.12: Generated road with mean reversion

As shown in Figure 3.12, the generated profiles are not entirely realistic. In real-world scenarios, roads tend to revert faster when falling below a certain elevation, whereas deviations above the mean are more frequent (i.e., hills are more common than valleys below sea level). To address this, it is advisable to use a noise function with a varying amplitude for different elevations. And add this noise to the AR(1) process. One such approach is using the following function,

$$f(E_d) = a_2 \cdot E_d \cdot \left(1 + \frac{c}{1 + e^{-k \cdot E_d}}\right),$$

where the parameters a, c, and k can be estimated, and the input parameter d, is the difference between the current elevation and the specified ground value, for reversion. An example with $a = 1$, $c = -0.9$, and $k = 0.1$ is illustrated in Figure 3.13, which shows both the function plot and the resulting generated road. Where the function noise is supposed to mimic real-world behavior better than just a simple mean reversion, which is the same below and above a specified elevation.

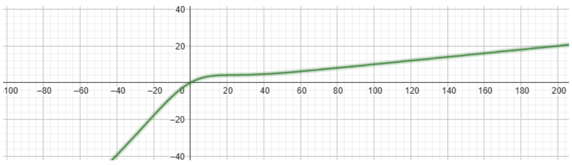

(a) Noise function for realistic noise

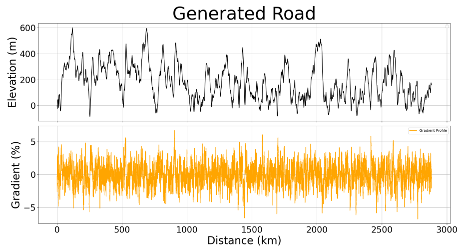

(b) Elevation and slope profile, for a generated road.

Figure 3.13: The noise function used for generated the slope profile with mean reversion, and the corresponding generated slope and elevation profile.

It should be noted that this approach introduces nonlinearity into the process, thereby complicating the attainment of the desired statistical properties. Consequently, fitting the parameters to data and determining the stationary distribution becomes more challenging, although these issues can still be addressed with further evaluation. Mathematically the three different models are described as following,

$$\begin{bmatrix} s_1 \\ s_2 \\ s_3 \\ \vdots \end{bmatrix} = \underbrace{a_1 \cdot \begin{bmatrix} s_0 \\ s_1 \\ s_2 \\ \vdots \end{bmatrix} + \boldsymbol{\varepsilon}}_{\text{AR(1)}} + \underbrace{a_2 \cdot E_d}_{\text{With Simple Mean Reversion}} \cdot \underbrace{\left(1 + \frac{c}{1 + e^{-k \cdot E_d}}\right)}_{\text{With Complex Reversion}},$$

$$E_d = \frac{Sr}{100} \begin{bmatrix} s_0 \\ s_0 + s_1 \\ s_0 + s_1 + s_2 \\ \vdots \end{bmatrix}, \quad \text{with } Sr = 25.$$

where s is the slope values in (%), E_d is the elevation difference vector, Sr is the sample rate, ε is the vector of error terms, a_1 and a_2 are scaling parameters, k and c govern the nonlinear complex reversion part.

The parameters a_1, a_2, c, k, and the properties of the error vector $\boldsymbol{\varepsilon}$ are to be estimated.
Since these models extend beyond a standard AR(1) specification, the Yule-Walker method is not used for parameter estimation. Instead, the parameters of the simple mean reversion model are estimated via an ordinary least squares (OLS) regression using the statsmodels.api library [22]. For the more complex logistic model, the parameters are estimated using the curve_fit function from scipy.optimize [24], which employs a nonlinear least squares approach.

3.5 Adaptation of models for use in simulations

When generating road profiles from models for use in vehicle simulations, it is often necessary to ensure that the generated roads meet specific criteria, such as the number and length of hills, the proportion of slopes exceeding 3%, or other predefined constraints. However, incorporating such constraints directly into a stochastic process can be challenging without compromising its inherent randomness.
For example one limitation of the current approaches, not considering the mean reversion model, is that the generated road elevation profiles tend to exhibit unrealistically high and low altitudes. This arises because there are no constraints to limit these extreme fluctuations. The importance here is to ensure that the road profile conform to the specifications of the area that should be testing. This means that the altitude should not go above or below certain thresholds that are unreasonable, examples could be not going under the sea level, or going up to above 2000 m when driving in Sweden.

Various strategies have been proposed to address this issue. One way to be able to control the altitude is to utilize a method called *detrending* presented by Johannesson. P [5]. This method calculates the average slope offset, and adds it to the slope at every data point. This leads to the end altitude being the same as the starting altitude, since the total slope gain becomes zero. The pros of using this method is that the statistical variance and stochastically of the model is kept, with the disadvantage that the gradient becomes slightly less extreme and that while making sure that the start and end elevation is the same, there can still be unrealistically high/low altitudes in between.

One effective approach to address this challenge while preserving stochastic variability is to generate a large number of road profiles and then select those that satisfy the desired criteria. This method ensures that the resulting roads remain statistically diverse while adhering to the required constraints. An example of this approach is illustrated in Figure 3.14 where only roads with a final elevation within 100 meters of the starting elevation and classified as PFLAT according to GTA criteria have been selected.
Alternatively, multiple road profiles can be generated, and simulations can be executed across the entire set. This approach aims to enhance statistical robustness by reducing the impact of individual anomalies and capturing a broader range of possible road conditions. By analyzing multiple road realizations, it becomes possible to derive more reliable performance metrics and identify trends that might not be evident

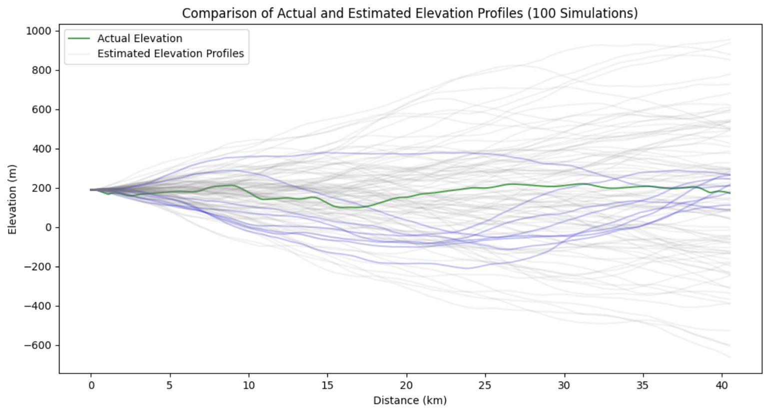

Figure 3.14: Elevation profiles of 100 generated roads from Borås to Landvetter, with the actual road (green) and selected roads (blue).

from a single simulation.
This method is particularly useful in applications where variability in road topography can significantly influence energy consumption, max inclination or altitude. By distributing statistical uncertainty across a diverse set of generated roads, the analysis becomes more representative of real-world conditions. Furthermore, this approach allows for sensitivity analysis, enabling the evaluation of how different terrain features affect key parameters such as fuel efficiency, wear and tear, or traffic flow. Additionally, running simulations on multiple road profiles supports robustness testing in road design and transportation planning. Due to lack of time, this approach is left more as a suggestion, and not evaluated like the other methods.

Another method could simply be to run simulations on longer roads, since the longer a road is, the higher the probability to get a correct estimation of the real gathered data, as seen in figure 3.8b, and described by the Law of large numbers eq A.8.

3.6 Validation and simulation

When it comes to validation and simulation, various aspects will be examined. For driveline and power, the proposed heatmap will be used to compare real roads and generated roads, alongside their corresponding slope and elevation profiles. Additionally, the suitability and limitations of the different models will be discussed. Regarding energy consumption, a forward simulation model developed by Chalmers students in Simulink, is used [25]. This model was developed for the purpose of calculating the total energy consumption for a double combination Battery Electric Vehicle (BEV), with a vehicle weight of 25 tons and a cargo weight of 30 tons. For a round trip from Gothenburg to Viared. This model works for other trips as well and will be applied to both real and generated roads. Using this model the slope profile from the real roads used to validate against. were substituted for the slope profile from the generated road. Furthermore, the speed limits, curvatures, free-flow speeds, and traffic signals were adopted directly from the roads used for the validation, isolating the influence of the slope profile. The weather data used was fixed to a specific date, ensuring consistent conditions across runs. As expected, the plots of power and torque over time at given times, show a strong correlation with the slope profile, as illustrated in Figures 3.15 and 3.16. This reinforces the notion that analyzing slope profiles and slopes of varying lengths, is a robust approach.

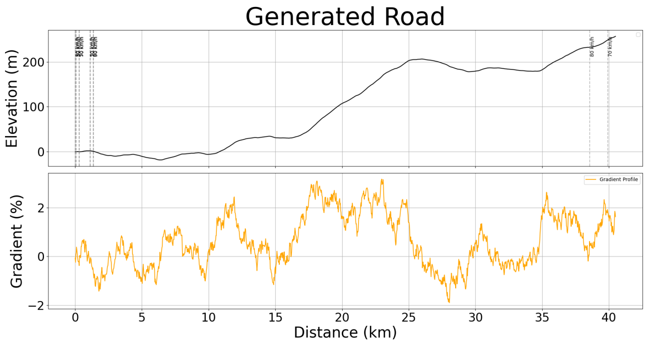

Figure 3.15: The elevation and slope profile of a 40 kilometers long generated road.

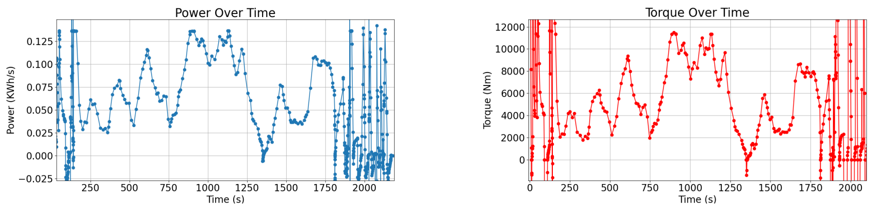

(a) The power at a given time over time using the generated road.

(b) The wheel torque over time using the generated road

Figure 3.16: The torque, taken from the wheel, and power over time profiles looks very similar to the slope profile, highlights its close correlation.

This simulation model will serve as a validation tool by processing real roads from HEREmaps alongside five modified roads in which the original slope profile is replaced with a generated one from one of the project's generative models. To ensure robustness, multiple real roads of different lengths (all predominantly flat) will be tested. For each road, the simulation results of the five generated variants will be averaged and compared with the real road data in a table. Also, histograms showing the different energy consumptions for the runs will be plotted.

4
Results

This chapter presents the results of the energy simulations and characteristics evaluation of the three models: AR(1), ARMA(1,5) and Markov chain. The models were trained on a set of 10 real PFLAT roads for characteristics analysis in heatmaps and level crossing, and the simulations were run each road by themselves. The methods of road segmentation and altitude reversion are evaluated.

4.1 Model performances

This section presents the generated roads, characteristics and results of the models and compares them to the sampled real-world data. For this section only the first three models in this project without any modifications will be considered. This refers to Ar(1) model, ARMA(1,5) model and the Markov slope chain model. Also due to lack of time only these three models will be evaluated using the Simulink energy consumption tool.

4.1.1 Forward simulation of energy consumption

To examine the performance of the models, a PFLAT road was selected, and each of the three models AR(1), ARMA(1,5) and Markov slope chain (see sections 2.2.1, 3.3.2, 3.3.3), were trained on its data according to the descriptions given in the method. To ensure that the result is representative of the model, each model generated five roads to achieve statistical robustness 3.5, and the process was repeated for 4 different roads. These roads were selected to represent 4 different areas of the PFLAT category, one average (Borås), one with a mountain and then a flat plain (Manchester), one long (Prague), and one tough (Arvika). The roads that were generated for the AR(1) model for the distance Landvetter-Borås can be seen in Figure 4.1a, where each generated road has a number.
The roads are then used as input for the topography in the simulink code for vehicle modeling using forward simulation described in Section 3.6. The resulting energy consumption is then plotted in figure 4.1b.
Comparing the two graphs in Figure 4.1, one can see that the elevation has a large effect on the energy consumption, both in total elevation difference and in inclination. The effect of elevation gain can be seen as the four roads that have a higher ending altitude than the original road all have higher energy consumption as well, and vice versa for the last road nr 3. The flat road ends on a higher altitude than the real road but still has a lower consumption, which shows that inclination and topography

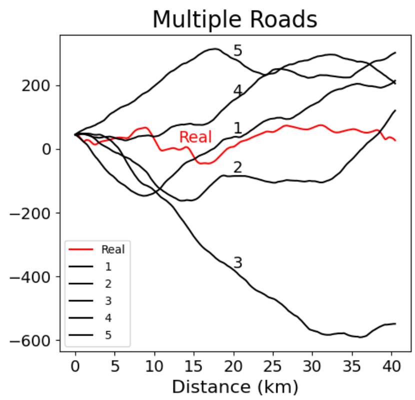

(a) Elevation profile

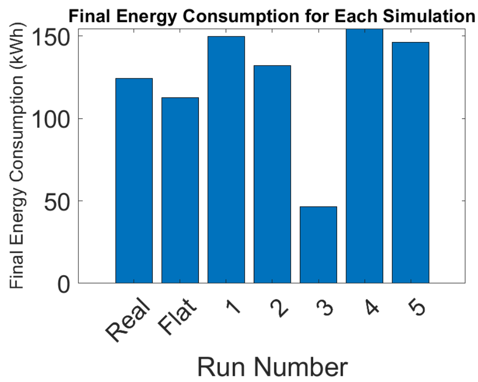

(b) Total energy consumption

Figure 4.1: Elevation profile and Total energy consumption for AR(1) model simulations of Landvetter-Borås. The sets contain the sampled road, a straight road, and 5 generated AR(1) model roads

has an effect on the simulations and concurrently energy consumption.
The performance of the models, together with the result for the original sampled data and a completely flat road is quantified in the table 4.1.

road	**category**	**real**	**flat**	**AR(1)**	**Markov Slopes**	**Arma(1.5)**
Borås-landvetter (40.5 km)	Relative energy	1	0.920	1.017	0.762	0.978
	Δh [m]	-18	0	13.6	-78.1	39.5
Manchester-Hull (167.3 km)	Relative energy	1	0.927	1.051	1.039	0.988
	Δh [m]	-37	0	149.4	-21.4	-56
Prag - Wien (334.8 km)	Relative energy	1	0.824	0.990	1.054	0.930
	Δh [m]	-47	0	357.8	446.9	310.2
Arvika - Kil (53.6 km)	Relative energy	1	0.786	0.974	1.016	0.990
	Δh [m]	46	0	-64.8	-18	80.7

Table 4.1: Table comparing the energy consumption of four real roads with their simulated counterparts. Here, Δh denotes the elevation difference between the start and end points of each road. For each real road, five simulated roads were generated, and the average elevation difference along with the corresponding energy consumption was calculated.

The table depicts two rows for each road. The relative energy, which is the energy of the models compared to the real road data, and the height difference between start and endpoint. It can be observed that the height difference between start and endpoint can vary between models, this is a result of the models not having any input parameter for elevation. This is what causes the large difference in energy consumption between run of the same parameters, discussed earlier in Figure 4.1.

Examining the table, one can see that all results from the generated road models are quite close to the energy consumption of the real road [±7%], except for the Markov slope model for Landvetter-Borås. This has a relative energy consumption difference of 23.8%. This is a huge outlier, and if one examines the elevation profile and energy consumption for this model in Figure 4.2, one can see that the consumption is much lower since almost all roads have a large negative elevation difference. This difference leads to a huge difference in energy consumption. This shows that when using models to stochastically generate roads, it's important to be aware of what can vary in the roads, and either examine them beforehand with multiroad generation, section 3.5, or use a method to remove this variance, such as de-trending, presented in the same section 3.5.

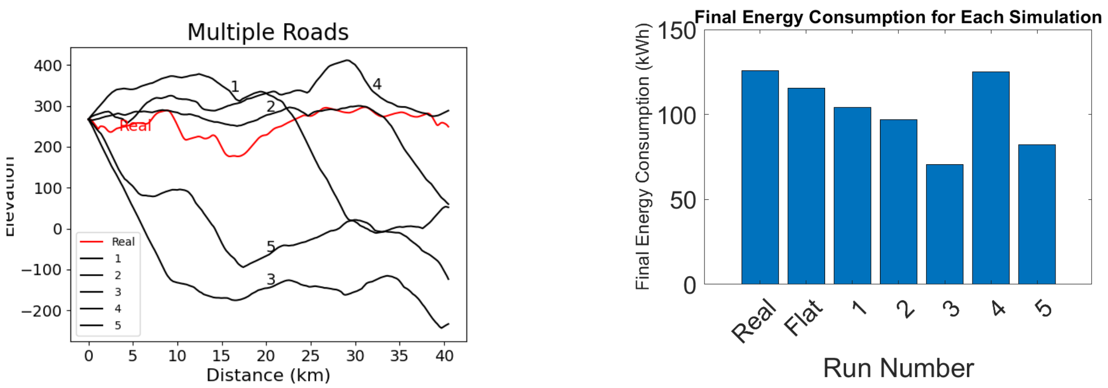

(a) Elevation profile

(b) Total energy consumption

Figure 4.2: Elevation profile and Total energy consumption for Markov chain model simulations of Landvetter-Borås. The sets contain the sampled road, a straight road, and 5 generated AR(1) model roads

4.1.2 Characteristics of models

It is of importance that the generated models are able to capture the true characteristics of the roads, as discussed in chapter 3.1. The study concluded that the most important characteristics were hill grade and length. To examine how well the different models perform in this regard, the aggregated data from 10 PFLAT roads were gathered and AR(1), ARMA(1,5) and Markov models were built on this data. The heatmap for these results can be seen in Figure 4.3.

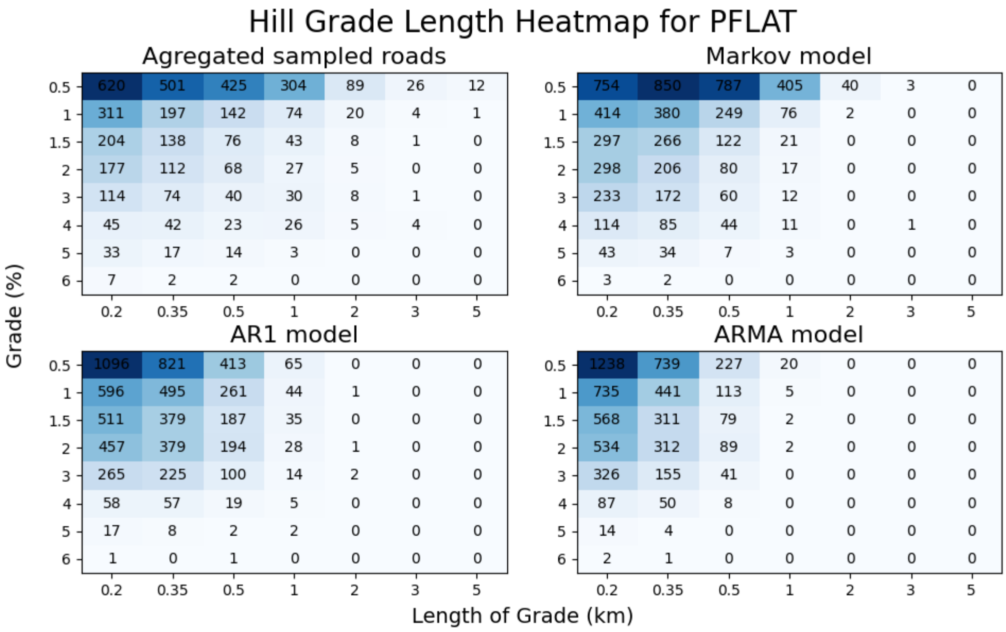

Figure 4.3: Heatmaps for real aggregated roads, and the Markov, AR(1), and ARMA(1,5) models

The results show that as expected the most common inclination is the smallest and shortest (0.5 %, 200 m), since as explained in Section Heatmap 3.1.3, you can fit more short segments than long ones on a given distance, and 0.5% is the closest to flat. When looking at the result one can see that the Markov model is very similar to the sampled data, but with slightly longer hills. This is positive for Markov as it shows it can capture the characteristics of the sampled roads fairly accurately. The two AR models produce similar results to each other, capturing the inclination of the sampled roads correctly up to 4 %, but drops off in inclination afterwards. The models fail to produce longer hills, especially low inclination, long hills are hard to produce.
This can further be examined in the level crossing analysis, see section 3.1.5, seen in Figure 4.4.

Here the Markov slope chain model can be seen to have the same distribution of slope variance and intensity as the sampled roads. This is expected, since the Markov slope chain model is constructed by counting transitions between slope values in a way that closely matches the real sampled data.". The two AR models exert the dame distribution and behavior, namely a standard distribution with a specific variance. This is expected since this is the method that the AR models are built upon.

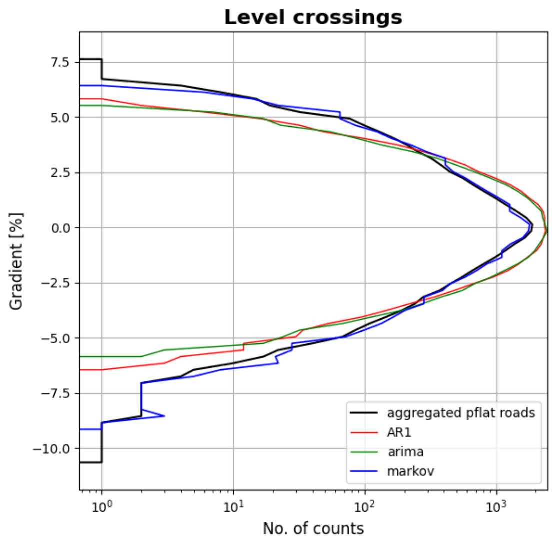

Figure 4.4: The level crossing profile for the grade segmented roads. The sampled roads are in black, and the models in color.

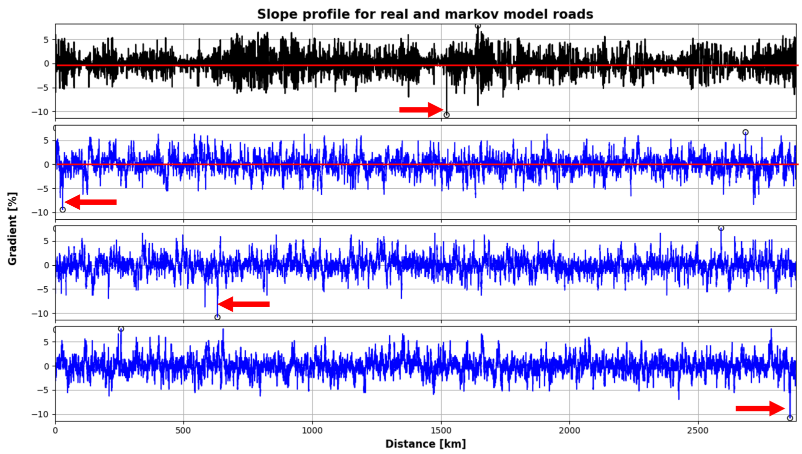

Figure 4.5: Slope profile of gradient, for Markov models (blue) and the sampled roads (black). The arrows point towards the maximum negative slope

To further examine how well the Markov slope chain model captures the dynamics of the sampled roads, one can look at Figure 4.5, where the slope profile of the sampled road, and three generated Markov model roads are plotted.

The red arrows show that both the sampled and the Markov roads all have one instance of a negative slope with gradient of 10 percent. This shows that the slope chain model can capture real instances of intense gradients, and output them at a reasonable frequency.

The red line through the first two profiles is to show that the real roads have a balanced slope profile, with the intensity of positive and negative slopes being the same at most instances. The slope chain model meanwhile have a more fluctuating slope profile, where there can be instances of steep hills in one direction, but small in the other. This occurs since the Markov model is not affected by terrain or remember what have happens longer than one timestep back. This can be seen in the elevation profile in Figure 4.6, where the Markov models reaches higher extremes than the sampled data.

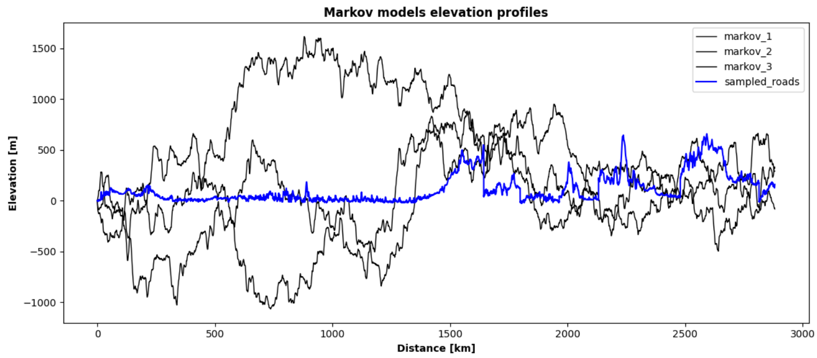

Figure 4.6: Elevation profile for Markov models (blue) and the sampled roads (black)

To examine if the way hills are defined 2.1.1 impact the results, the roads from the PFLAT dataset and corresponding models are transformed using the grade segmented roads method 3.1.4. The resulting can be seen in the heatmap 4.7.

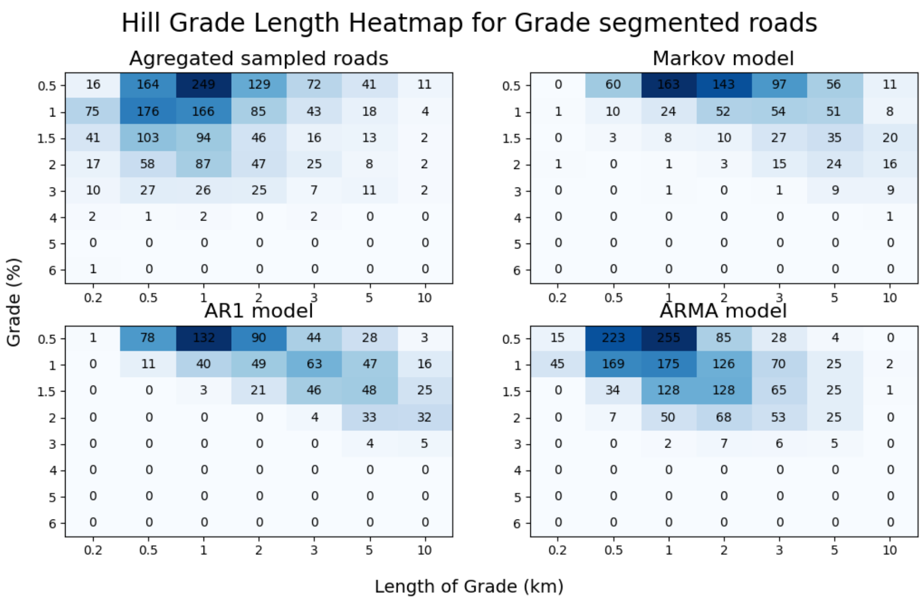

Figure 4.7: Heatmaps for the Grade segmented hills for the real aggregated roads, and the Markov, AR(1), and ARMA(1,5) models

All road segments become longer in this method, and also flatter, as extreme slope values are averaged out. The heatmaps generated for all models exhibit significant changes compared to those based on the previous hill definition (see Figure 4.3), whereas the sampled data has a smaller change. The Models now have a skewed distribution, with a correlation between longer hills and higher inclination, a pattern not observed in the actual road data. This discrepancy may stem from the fact that the models are not constrained by real-world geographical or engineering limitations, but are instead governed solely by the previous slope value. This implies that if a hill has a steep inclination, it will generally be longer—and vice versa—since a slope ends when the inclination returns to 0 degrees, as explained in Section 3.4.3. The reason the sampled real data does not change as much, is because it already is more even in its fluctuations, as roads in the real-world tend to be designed to have a constant gradient, and being connected to the geographical restrictions of the surroundings. From the resulting heatmap, the ARMA(1,5) model is the one that seems to be most similar in appearance to the real roads. This means that the implementation of the previous error values not only smooths out the slope profile, but also has a positive effect on the larger structure of the slope profile. The Grade segmented hills results are further examined in the level crossings Figure 4.8.

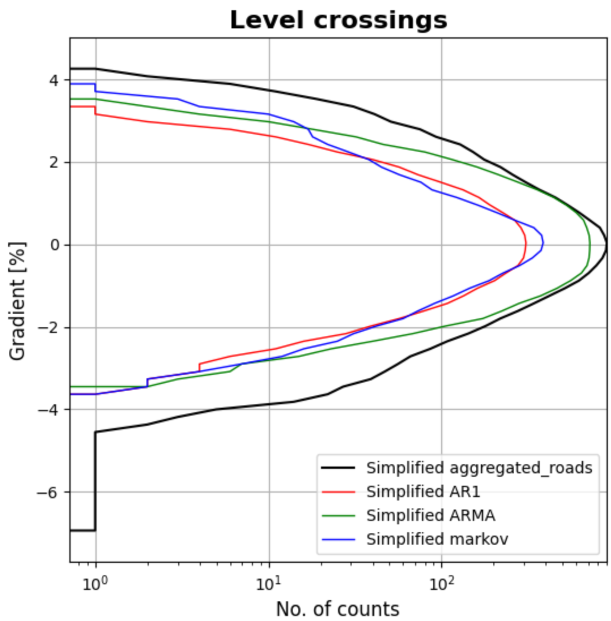

Figure 4.8: The level crossing profile for the grade segmented roads. The sampled roads are in black, and the models in color.

The result show that the Markov model no longer follow the sampled data, and neither does AR(1). The ARMA(1,5) model does perform well for small slopes in the range of ±2%, but worse outside of this spectrum. This align with the result from the heatmap analysis. From this we can gather that there is some characteristics difference between all the models and the real roads when looking at the grade segmented roads, which should mean that the larger topography does not align with reality. How this effect simulation results are not clear, but from the results of the simulations in Table 4.1, the ARMA(1,5) model was the closest to the real roads in energy consumption, which could mean that the characteristics the model capture is important.

4.2 Segmentation of roads using AR(1) and Markov chain

Three road segmentation approaches were explored. Segmentation based on speed limits, functional class and slope variance (terrain type). Due to lack of time only the first two was implemented as state switching Markov chain model, and evaluated against real roads. However to illustrate their application and the segmentation potential, all three methods were tested on a route from Gothenburg to Kiruna, using data provided by HERE Maps.

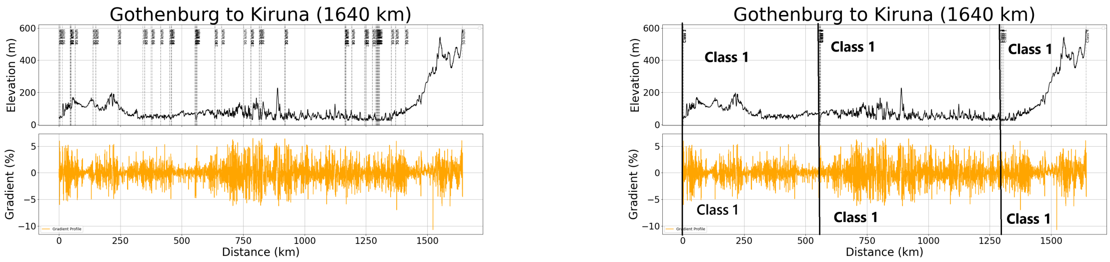

(a) Segmentation based on switches of speed limits.

(b) Segmentation based on switches of road functional classes.

Figure 4.9: Figure a) show segmentation based on switches of speed limits, where the gray lines represent a switch to a new speed limit. The trip was 52% 70 km/h, 27% 90 km/h, 17% 110. and 4% 30, 50 and 80 km/h. While figure b) show segmentation based on switches of road functional classes, where the gray lines represent a switch to a new functional class

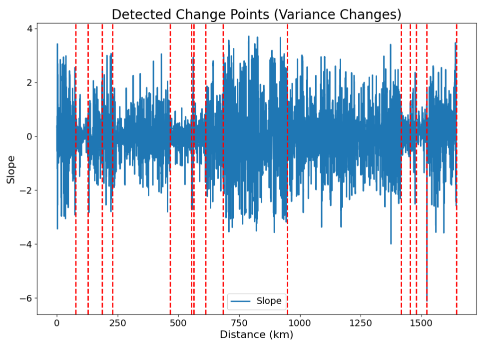

Figure 4.10: Segmentation based on switches of areas with similar slope variance (terrains), where the red lines represent a switch to a new terrain. Based on a route from Gothenburg to Kiruna.

Additionally, an example of a generated road based on functional class data from real trip is presented to showcase potential use cases. The trip used was a trip from Varekil (Orust) to Hålt (A place 10 kilometers into the forest). On this trip three different functional classes occurred while having a speed limit of 70 km/h 98.5% of the trip. The results for this can be seen in figure 4.11

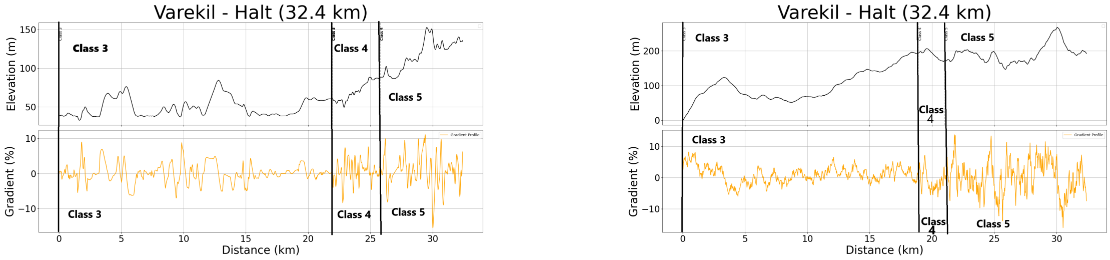

(a) The real trip from Varekil to Hålt. **(b)** A generated road trip.

Figure 4.11: Segmentation based on switches of functional class, where the black lines represent a switch to a new functional class. Where figure a) show the real tip (Varekil - Hålt), and figure b) show the generated road profile using the AR(1) model along with Markov chain to switch functional classes based on the trip from Varekil to Hålt.

Furthermore, the speed limit and functional class models were trained on the same roads used in the earlier model evaluations. These models were then used to generate new road profiles, which were compared to real world data using the grade length heatmap metric. The results of this comparison are shown in Figure 4.13. Similar to the findings from the earlier model evaluations, these two models do not appear to capture hill lengths in a realistic manner either.

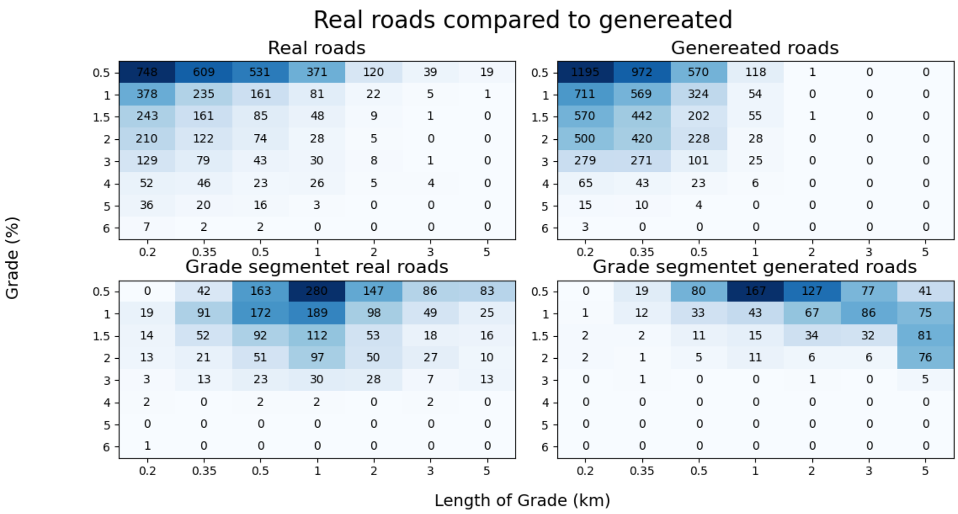

Figure 4.12: Heatmap of grade lengths where the aggregated real roads are compared with the generated roads, which are generated based on the functional class model, trained on the aggregated real roads.

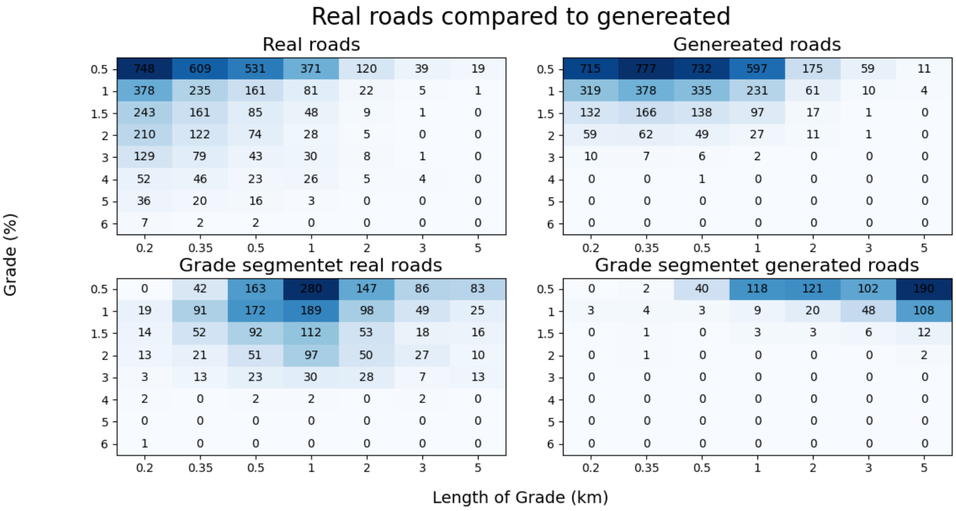

Figure 4.13: Heatmap of grade lengths where the aggregated real roads are compared with the generated roads, which are generated based on the speed limit model, trained on the aggregated real roads.

4.3 Altitude and mean reversion

The last model tried was the altitude reversion model, described in Section 3.4.3. Here the complex reversion model was trained on the same roads as the previous evaluations, and was validated against the real roads using the grade length heatmap metric. The result of this is illustrated in Figure 4.14.

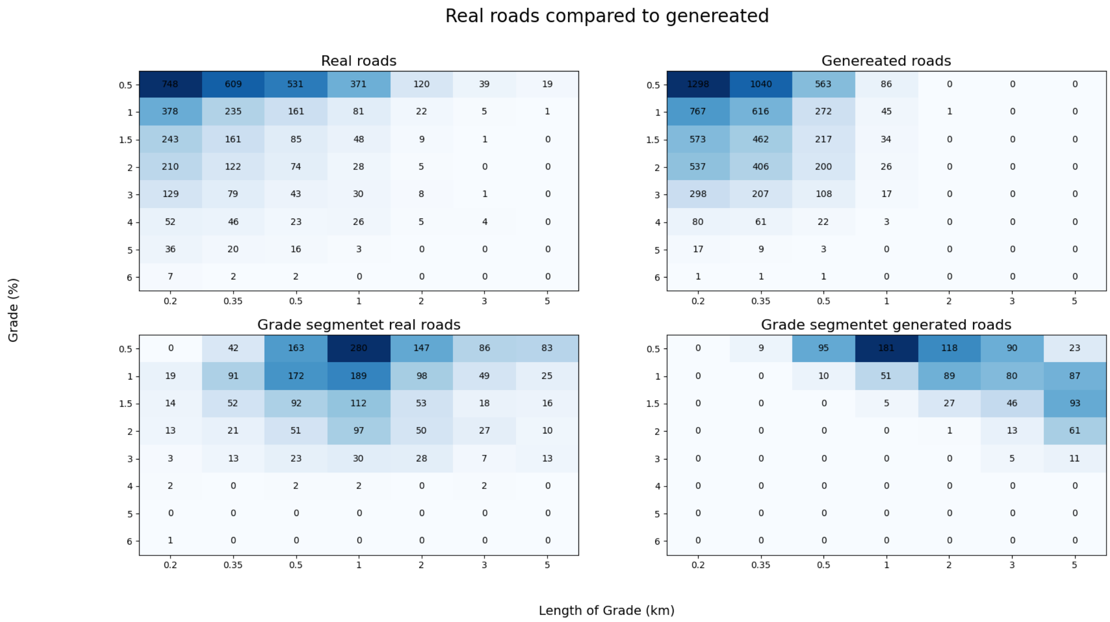

Figure 4.14: Grade length heatmaps for the complex reversion algorithm compared to the real predominate flat roads used for evaluation.

As illustrated by the heatmaps, the generated roads do not achieve particularly better results than other models, and performs similarly to the AR(1) model, which is the basis for the altitude reversion model. It is however more effective at maintaining altitudes closer to the ground level, as seen in 3.4.3, which is to be expected.

5
Discussion

In this chapter, the results and methodology is discussed. The RQs are answered, and the key identifications are examined. Lastly further research is discussed.

5.1 Model evaluation

For this section the AR(1), ARMA(1,5) and the Markov slope chain model evaluation will be discussed. The three models vary a bit and each has its strengths and flaws. The AR(1) model is very simple and has the ability to go from model parameter to GTA classification and vice versa, explained in the AR(1) theory section 2.2.1. Here it also describes a possibility to control the average hill length of the generated roads, it should however be noted that this does not change the overall distribution between longer and shorter hills, only the average hill length. These two aspects can be a benefit when generating roads for vehicle simulations. It is of course also possible to go to stationary distribution for the Markov slope chain model and ARMA(1,5) model, meaning that a given transition matrix or model parameters can give a GTA classification when long enough roads are considered. This is demonstrated for the Markov slope chain model in 3.3.3. For these models however, it is probably not possible to go from GTA classification to transition matrix or model parameters.

Considering simulations for energy consumption it can be seen from the energy consumption table 4.1, that all three models tend to perform quite well, although the ARMA(1,5) model consistently results in slightly lower energy consumption. This difference may simply be due to random variation. In some cases, the generated road profiles show significantly lower energy consumption than real roads. This effect is particularly noticeable for shorter road segments, such as the Markov slope chain model for Borås to Landvetter in Figure 4.2a, where the stochastic generation process can occasionally produce roads that maintain a continuous downward slope, resulting in low energy consumption. To counteract this, it is recommended to use one of the methods described in 3.5, such as detrending, or generate longer roads, that fulfill the criteria of stationary distribution and sample size discussed in A.1. If longer roads are not desirable or feasible, one could instead utilize the method of multi road generation discussed in section 3.5, to generate several roads, and either choose the ones that fulfill the criteria, or do the simulations on several roads and take the average. This method has the disadvantage of taking more resources and more time. Ensuring the generated roads are balanced is something that is extra important when simulating for energy consumption.

A notable advantage of the Markov slope chain model is its ability to capture intense max inclination hills, as illustrated in Figure 4.5. Such spikes are reflective of real-world conditions, where abrupt slope variations occur. Given that the Markov chain model performs well in terms of energy consumption, it could be particularly useful for startability analysis, where the maximum gradient plays a critical role. Startability could also be tested by inserting a slope with the specific grade that should be validated. How this affect the stochasticity is unsure.
When it comes to durability and driveline, which are exponentially affected by grade, we can see from the heatmap 4.3, that the models produce an acceptable spread of slopes, but generally don't produce roads with as long hills as seen from real data, especially the two AR models. If we look at the level crossing of the slope profile 4.4, we can see that the Markov chain method captures perfectly, which is desired. This discrepancy for the AR models poses challenges in powertrain simulations, where long stretches of specific hills are critical for assessing performance under sustained loads.
Though, when defining the slopes using grade segmented hills method 3.1.4, the heatmaps yield markedly different results, with all generative models producing longer slopes than the real-world data. This result shows that the larger structure of the generative models are not realistic, since they will fluctuate more between different grades, but continue in the same slope direction. This finding underscores the importance of precisely defining how hills are characterized when assessing model realism; reliance solely on heatmaps of grade lengths is insufficient.

Another thing that has to be considered when stochastically generating road profile for different disciplines is the symmetry in the slope profile, seen in Figure 4.5, it can be seen that for the real-world roads a spike up, almost always means a spike down of similar intensity. This phenomenon is attributed to the natural topography of road networks, where ascents are typically followed by descents of the same intensity, since the terrain is correlated, and the road usually has a specific ground level with hills on top of it. The lack of this characteristic from the generative models, is considered to be the reason that the grade segmented hills, tend to create longer hills more often. Capturing this characteristic in stochastic models presents a challenge. The altitude and mean reversion model was an effort to try to solve this issue, discussed in section 5.3, however it did not achieve particular good results. One potential solution is to employ a more complex model. An approach is to incorporate elevation as input directly into the modeling process, discussed further in future work 5.4.

5.2 Segmentation of roads

As seen in the results from the road segmentation. The speed limit approach often provides a richer dataset than the functional class method. In such scenarios, the functional class remains nearly constant, whereas speed limits frequently change (see Figure 4.9a). Moreover, speed limits may be more effectively combined with other stochastic models, such as those capturing curviness, traffic lights, stop signs and roundabouts. Because these factors tend to correlate more strongly with speed limits

than with functional classes. And the idea of this has already been explored, seen in [4].
The functional class appears to be more closely related to the variance and one-step covariance used in estimating the parameters for the AR(1) model. This is likely because the functional class reflects the road's importance. For example, high-traffic roads and highways typically receive additional construction efforts to ensure a straighter hill. This construction standardization directly influences the variance and covariance metrics. An example of this effect is illustrated in Figure 4.11a, where it can be seen that the slope profile changes with the functional class, while the speed limit on the road is 70 km/h 98.5% of the distance. In this case it is because the less important road in the forest, differ from the highway by having a high functional class, while still having a speed limit of 70 km/h.
When considering road segmentation based on terrain, i.e. identifying regions where the variance remains relatively constant, the segmentation method becomes inherently correlated with the variance used for parameter estimation. However, it remains challenging to determine the frequency with which drivers traverse these specific areas. Additionally, there is no clear connection between such segmentation and driving patterns for specific missions, for instance driving long-haul or regional trips. An idea could be to consider a specific road type, for instance highway/functional class 1. It may be possible to depict these areas on a map, similar to what have been done regarding a altitude topography map, thereby providing further insight into regions with more challenging terrain. Although related to topographical features, the correlation is not exact, for instance, regions of low altitude can exhibit high slope variance, while some high-altitude areas may display low slope variance, as illustrated in Figure 3.11a.

5.3 Altitude and mean reversion

The weight model for controlling altitude performs better at keeping the elevation close to ground level, however, the generated elevation profiles still appear unrealistic when compared with real-world using heatmaps validation metric. This discrepancy is likely due to the underlying AR(1) process, which tends to wander and struggles to produce the short hills typical of actual landscapes, even though it can generate longer hills. Even the more complex reversion model falls short of capturing the real-world pattern, where numerous short hills are occasionally interspersed with longer ones. This result however, suggest the idea that modeling the behavior of real-world topography might need a very complex model, potentially a machine learning algorithm.

5.4 Future work

The current approach in this project can still be improved in several ways, when for instance more control is a priority, exploring Markov chain models further is suggested. More data could be used to train the possibilities of switching between slope values. One could for instance consider to use a data set consisting of the slope profile of all highways in Europe, making understanding real world slope value switches easier. When it comes to generating based on a Markov chain methods, the possibilities are nearly endless. This approach can perhaps be explored further to better mimic real-world conditions by incorporating additional inputs, such as elevation alongside the slope profile, to enable slope transitions that consider both factors. Consequently, this method could capture the elevation profile more accurately than the current approach. Furthermore, the Markov slope chain could be used to define larger segments of slope transitions, for example, the grade segmented heatmaps in this project might instead be used as input for stochastic generation based on a heatmap of grade lengths. Then different heatmaps could be used for different test departments, where the heatmap is chosen specifically for each test departments need. However, this approach might also need some kind of elevation input to ensure that the hills does not keep going up or down in an unrealistic way, if realistic behavior is critical.

Markov chains could also be combined. For instance, a chain describing the road type, another for slope and elevation and a third for terrain, could be one way of making the road switch more similar to what is seen in the real world. It is though important to have enough data for each Markov chain to avoid a deadlock state.

However, if one wants a realistic behavior, these Markov chains might be very complex and in the end may not yield good results anyway. Real world behavior is complex, as seen when the different model approaches examined in this project does not generate a similar hill distribution when compared with real roads.

Therefore a recommended approach is to use a machine learning approach. Although this idea has not been investigated for this project, machine learning algorithms excel at uncovering complex correlations in large datasets. By training on a lot of data, along with the useful statistics, machine learning algorithms can be better at figuring out correlations. This can then perhaps first be used to identify real roads, that are better representatives for vehicle simulations than the roads used today. After this step is achieved, manually creating virtual roads, skipping the stochasticity, or using a similar but generative AI model, could be used to generate virtual road profiles that are also relevant for vehicle simulations. Since there are some correlations between for example road topography, road type and road curvature. Machine learning and AI models might be a good choice since they can treat a lot of different inputs simultaneously, instead of trying to divide them using for instance Markov chain models. This approach would however make the intuition of model behavior harder to achieve, where simple models is a better choice. But if a relevant, realistic road profile is of high value, a simple model might not be sufficient.

6
Conclusion

The thesis aimed to investigate the use of stochastic models to generate topographical road profiles for use in simulations. The research focused on three areas specified in the RQs.
An interview study found that the crucial characteristics of road topography to capture was grade, length of hill and number of hills.
Using these characteristics, three stochastic models were evaluated and compared to each other and real roads from here maps. It was found that all models and the real roads generally had a similar slope distribution, but depending on how you defined a hill, the generated roads were both shorter (AR(1), ARMA(1,5)) and longer (all models when grade segmented) than the real-world data. The models without modification are not connected to altitude, and could therefore create both large mountains and end up with a huge height distance compared to the start, a problem that was addressed by using mean reversion and altitude reversion. Unfortunately these methods did not seem to solve the discrepancy in characteristics between the generated roads and real road data seen in heatmaps.
When it comes to energy consumption, all three models generate roads that produce acceptable results in the simulations. We therefore cautiously recommend all models, with each having its advantages. AR(1) being the most simple and effective, with the ability to generate large amounts of roads quickly. Markov chain for capturing the characteristics the best. ARMA(1,5) for generating smoother roads than AR(1). For energy consumption where startability is also considered, the Markov slope chain model is suggested. When it comes to power and driveline, there is no clear answer, as behavior of the models seems to differ, and a more thorough research needs to be done on exactly what is considered as something tough for power and driveline. Markov chain models seems to be the best option, with it incorporating the varied gradients of the real-world. The negative is that it still lacks the connection to altitude, and does not create the balanced slope profiles that real-world roads generate. Another possible option is to use ARMA(1,5) model, since its grade segmented heatmap corresponds better than the other models compared to the sampled data in Figure 4.7. This would capture the large scope of the topography of roads.

However, it's not solely about generating realistic topography profiles. All models feature tunable parameters that can influence both the GTA classification and the lengths of the hills. This flexibility can be a significant asset for testing, as it allows tracking performance across various parameter settings.
For road segmentation, dividing the road into functional class is considered the best approach, since this is easiest to implement regarding transport mission and

vehicle description, and can be accessed directly from HERE maps. It does have the disadvantage of not having a rigid classification, but is instead determined according to HERE maps own discretion. However, if the stochastic topography model is supposed to run beside other stochastic models (traffic lights stop signs etc), speed limits is considered better since it correlates more, and can be found in all databases. That is why it has been proposed before [4]. Both the speed limit and the functional class model was compared to real roads using the grade length heat map metric, and performed quite similar to what was seen from the other models.

While the present study does not yield a definitive stochastic road generation model, it does show that modeling road topography as stochastic processes are a viable option for use in simulations. This result lays the groundwork for future research and methodological improvements in the field, where more methods and models can be explored.

Bibliography

[1] Y. Al-Wreikat, C. Serrano, and J. R. Sodré, "Driving behaviour and trip condition effects on the energy consumption of an electric vehicle under real-world driving," *Applied Energy*, vol. 297, p. 117 096, Sep. 2021, ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2021.117096. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261921005444 (visited on 12/04/2024).

[2] X. Zhao, X. Zhao, Q. Yu, Y. Ye, and M. Yu, "Development of a representative urban driving cycle construction methodology for electric vehicles: A case study in Xi'an," *Transportation Research Part D: Transport and Environment*, vol. 81, p. 102 279, Apr. 2020, ISSN: 1361-9209. DOI: 10.1016/j.trd.2020.102279. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920918303274 (visited on 12/04/2024).

[3] P. Pettersson, P. Johannesson, B. Jacobson, F. Bruzelius, L. Fast, and S. Berglund, "A statistical operating cycle description for prediction of road vehicles' energy consumption," *Transportation Research Part D: Transport and Environment*, vol. 73, pp. 205–229, Aug. 2019, ISSN: 1361-9209. DOI: 10.1016/j.trd.2019.07.006. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920919305243 (visited on 12/03/2024).

[4] L. Romano, "The operating cycle representation of road transport missions," en, Ph.D. dissertation, Chalmers University of Technology, 2023. [Online]. Available: https://research.chalmers.se/en/publication/536236 (visited on 12/03/2024).

[5] P. Johannesson, K. Podgórski, I. Rychlik, and N. Shariati, "AR(1) time series with autoregressive gamma variance for road topography modeling," *Probabilistic Engineering Mechanics*, vol. 43, pp. 106–116, Jan. 2016, ISSN: 0266-8920. DOI: 10.1016/j.probengmech.2015.12.006. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0266892015300655 (visited on 12/03/2024).

[6] P. Johannesson, K. Podgorski, and I. Rychlik, "Laplace distribution models for road topography and roughness," *Laplace distribution models for road topography and roughness*, Working Papers in Statistics, 2015.

[7] *SAS/ETS(R) 9.3 User's Guide*. [Online]. Available: https://support.sas.com/documentation/cdl/en/etsug/63939/HTML/default/viewer.htm#etsug_tffordet_sect016.htm (visited on 02/25/2025).

[8] J. V. Kern and D. J. B. Ferris, "DEVELOPING AREMA ROAD PROFILE MODELS WITH A STATIONARY RESIDUAL PROCESS," *IFAC Proceedings Volumes*, 5th IFAC Symposium on Advances in Automotive Control, vol. 40, no. 10, pp. 383–388, Jan. 2007, ISSN: 1474-6670. DOI: 10.3182/20070820-3-US-2918.00052. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1474667015319509 (visited on 02/25/2025).

[9] J. V. Kern, J. B. Ferris, D. Gorsich, and A. A. Reid, "Characterizing 2D road profiles using ARIMA modeling techniques," K. Schum and D. A. Trevisani, Eds., Orlando, Florida, USA, Apr. 2007, p. 65640L. DOI: 10.1117/12.720088. [Online]. Available: http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.720088 (visited on 02/25/2025).

[10] J. Guo, H. He, and C. Sun, "ARIMA-Based Road Gradient and Vehicle Velocity Prediction for Hybrid Electric Vehicle Energy Management," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 6, pp. 5309–5320, Jun. 2019, ISSN: 1939-9359. DOI: 10.1109/TVT.2019.2912893. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/8695855 (visited on 02/25/2025).

[11] J. V. Kern and J. B. Ferris, "Preliminary results for model identification in characterizing 2D topographic road profiles," in *Modeling and Simulation for Military Applications*, K. Schum and A. F. Sisti, Eds., International Society for Optics and Photonics, vol. 6228, SPIE, 2006, p. 622 806. DOI: 10.1117/12.665647. [Online]. Available: https://doi.org/10.1117/12.665647.

[12] P. Hoel, S. Port, C. Stone, and R. Holley, "Introduction to stochastic processes," *Systems, Man and Cybernetics, IEEE Transactions on*, vol. 3, pp. 533–533, Oct. 1973. DOI: 10.1109/TSMC.1973.4309295.

[13] J. Ferris, "Characterising road profiles as markov chains," *International Journal of Vehicle Design - INT J VEH DES*, vol. 36, Jan. 2004. DOI: 10.1504/IJVD.2004.005351.

[14] P. Pettersson, B. Jacobson, F. Bruzelius, P. Johannesson, and L. Fast, "Intrinsic differences between backward and forward vehicle simulation models," *IFAC-PapersOnLine*, 21st IFAC World Congress, vol. 53, no. 2, pp. 14 292–14 299, Jan. 2020, ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2020.12.1368. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2405896320317754 (visited on 03/09/2025).

[15] K. Säfsten and M. Gustavsson, *Research methodology: for engineers and other problem-solvers*. Studentlitteratur, 2020.

[16] L. Guzzella and A. Sciarretta, *Vehicle Propulsion Systems: Introduction to Modeling and Optimization*, en. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, ISBN: 9783642359125 9783642359132. DOI: 10.1007/978-3-642-35913-2. [Online]. Available: https://link.springer.com/10.1007/978-3-642-35913-2 (visited on 02/26/2025).

[17] T. Ghandriz, B. Jacobson, L. Laine, and J. Hellgren, "Impact of automated driving systems on road freight transport and electrified propulsion of heavy vehicles," en, *Transportation Research Part C: Emerging Technologies*, vol. 115, p. 102 610, Jun. 2020, ISSN: 0968090X. DOI: 10.1016/j.trc.2020.102610. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0968090X18311574 (visited on 02/26/2025).

[18] P. H. Brill, "Level crossing methods," in *Encyclopedia of Operations Research and Management Science*, S. I. Gass and M. C. Fu, Eds. Boston, MA: Springer US, 2013, pp. 874–877, ISBN: 978-1-4419-1153-7. DOI: 10.1007/978-1-4419-1153-7_529. [Online]. Available: https://doi.org/10.1007/978-1-4419-1153-7_529.

[19] *HERE Technologies Documentation | HERE Docs*. [Online]. Available: https://www.here.com/docs/bundle/maps-xml-data-specification/page/format-agnostic-data/format_agnostic_data_specification/topics/ra-functional-class.html (visited on 01/29/2025).

[20] R. Maghsood and P. Johannesson, "Detection of the curves based on lateral acceleration using hidden markov models," *Procedia Engineering*, vol. 66, pp. 425–434, 2013, Fatigue Design 2013, International Conference Proceedings, ISSN: 1877-7058. DOI: https://doi.org/10.1016/j.proeng.2013.12.096. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1877705813019292.

[21] J. B. Ferris, "Characterising road profiles as Markov Chains," en, *International Journal of Vehicle Design*, vol. 36, no. 2/3, p. 103, 2004, ISSN: 0143-3369, 1741-5314. DOI: 10.1504/IJVD.2004.005351. [Online]. Available: http://www.inderscience.com/link.php?id=5351 (visited on 02/27/2025).

[22] S. Seabold and J. Perktold, "Statsmodels: Econometric and Statistical Modeling with Python," en, *scipy*, May 2010. DOI: 10.25080/Majora-92bf1922-011. [Online]. Available: https://proceedings.scipy.org/articles/Majora-92bf1922-011 (visited on 03/09/2025).

[23] *Ruptures: Change point detection for signals in Python.* [Online]. Available: https://github.com/deepcharles/ruptures/ (visited on 03/11/2025).

[24] *SciPy documentation — SciPy v1.15.2 Manual.* [Online]. Available: https://docs.scipy.org/doc/scipy/ (visited on 03/23/2025).

[25] G. G. Kamat, H. Ahmadi, M. R. Rajput, P. K. Adiga Nagaraj, W. Yousefi, and L. Perfetti, "Voltage vs. Horsepower: Comparing Heavy Vehicle Performance through Simulation," eng, 2024. [Online]. Available: http://hdl.handle.net/20.500.12380/309141 (visited on 03/09/2025).

[26] P. Hewson, "Bayesian Data Analysis," en, *Journal of the Royal Statistical Society Series A: Statistics in Society*, vol. 178, no. 1, pp. 286–288, Jan. 2015, ISSN: 0964-1998, 1467-985X. DOI: 10.1111/j.1467-985X.2014.12096_1.x. [Online]. Available: https://academic.oup.com/jrsssa/article/178/1/301/7058467 (visited on 03/09/2025).

[27] P. J. Brockwell, R. A. Davis, P. J. Brockwell, and R. A. Davis, “Estimation for arma models,” *Time Series: Theory and Methods*, pp. 238–272, 1991.

[28] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time series analysis: forecasting and control.* John Wiley & Sons, 2015.

[29] G. Grimmett and D. Stirzaker, *Probability and random processes.* Oxford university press, 2020.

A

Appendix 1

A.1 Stationary Distribution and Effective Sample Size (ESS)

Another important aspect to examine is how long the generated roads need to be for the roads to be statistically representative of the input sample. One way to examine this is to examine when the parameters reach an asymptotical state, where the the statistical properties of the process (like mean, variance, and autocorrelation) will not change over time. this is called the stationary distribution π. In road modeling, the concept of a stationary distribution helps describe the long-term behavior of road slopes, meaning if stationary distribution is reached, enough independent data points are in the dataset for the result to to not change over time. The amount needed to achieve this state is called effective sample size (ESS), and measures how much independent information is available in a given dataset.
A **stationary distribution** represents the long-run equilibrium of a stochastic process describing road slopes. If the process generating the slopes follows a probabilistic transition rule, the stationary distribution tells us the probability of encountering different slopes over long distances. Mathematically, for a stochastic process with transition matrix P, the stationary distribution π satisfies:

$$\pi P = \pi, \quad \sum_i \pi_i = 1. \tag{A.1}$$

This means that if one start from the stationary distribution, the slope probabilities remain unchanged over time. The stationary distribution is useful for understanding the expected frequency of various slopes in a given road network.
The **ESS** quantifies how many independent samples are present in a dataset when the data points are correlated. Road slopes tend to exhibit spatial correlation, meaning that consecutive observations are not fully independent. The ESS is given by:

$$n_{\text{eff}} \approx \frac{n}{1 + 2\sum_{t=1}^{T} \rho(t)}, \tag{A.2}$$

where $\rho(t)$ is the autocorrelation at lag t, and T is a cutoff where $\rho(t)$ becomes negligible. A lower ESS means that more data is required to achieve the same statistical precision as in an independent sample [26].
Applying this to the methods used in this thesis one get that in an **AR(1) model**, road slopes evolve according to eq. 2.1 where a is the autoregressive coefficient. The stationary distribution of this process is normally distributed with variance:

$$\sigma^2_{\text{stationary}} = \frac{\sigma^2}{1-a^2}, \tag{A.3}$$

for $|a| < 1$. The ESS for an AR(1) process is given by:

$$n_{\text{eff}} \approx \frac{n(1-a)}{1+a}. \tag{A.4}$$

Larger values of a lead to stronger autocorrelation and a lower ESS, meaning more data is needed for reliable estimates [27][28].
In a **Markov model**, road slopes follow a discrete transition probability matrix P, where each state (slope category) depends only on the previous state. The stationary distribution is found by solving:

$$\pi P = \pi. \tag{A.5}$$

The ESS for Markov chains is computed using the integrated autocorrelation function:

$$n_{\text{eff}} \approx \frac{n}{1 + 2\sum_{t=1}^{T} \rho(t)}, \tag{A.6}$$

where $\rho(t)$ is the lag-t autocorrelation, which for the AR model is $\rho(t) = a^2$. For the Markov model it is given by:

$$\rho(t) = \frac{\gamma(t)}{\gamma(0)} = \frac{E[X_0 X_t] - (E[X])^2}{E[X^2] - (E[X])^2} = \frac{\sum_{i,j} ij \cdot \pi_i (P^t)_{ij} - (\sum_i i\pi_i)^2}{\sum_i i^2 \pi_i - (\sum_i i\pi_i)^2}. \tag{A.7}$$

Slow-mixing Markov chains (where transitions between slope states occur infrequently) result in smaller ESS values, requiring longer simulations or observations to obtain reliable statistics [26]. Using this function, one can calculate certainty that given a specific length of road, the results should be representative of the given inputs to a certain degree. This can be used as a benchmark to evaluate how long the generated roads need to be, to fulfill the requirements of the different areas of simulation and testing.

This is connected to the **law of large numbers**, which states that as the number of observations increases, the sample average converges to the true expected value of the distribution. For a sequence of independent and identically distributed (i.i.d.) random variables $X_1, X_2, \ldots, X_n$ with finite expected value $E[X]$, the LLN states that:

$$\frac{1}{n}\sum_{i=1}^{n} X_i \xrightarrow{n\to\infty} E[X]. \tag{A.8}$$

This means that with a large enough sample size, the sample mean $\bar{X}$ approaches the true mean $E[X]$. However, in Markov models and AR processes, observations are correlated rather than independent. This affects how quickly averages converge, which brings us to stationary distributions and ESS [29].

A.2 Results

A.2.1 Simulations

road	category	real	flat	AR(1)	Markov Slopes	ARMA(1.5)
Borås-landvetter (40.5 km)	Energy [KWh]	129.2	118.8	131.4	98.5	126.3
	Relative energy	1	0.920	1.017	0.762	0.978
	Δh [m]	-18	0	13.6	-78.1	39.5
Manchester-Hull (167.3 km)	Energy [KWh]	472.0	437.7	496.2	490.2	466.1
	Relative energy	1	0.927	1.051	1.039	0.988
	Δh [m]	-37	0	149.4	-21.4	-56
Prag - Wien (334.8 km)	Energy [KWh]	822.2	677.9	814.2	866.8	764.9
	Relative energy	1	0.824	0.990	1.054	0.930
	Δh [m]	-47	0	357.8	446.9	310.2
Arvika - Kil (53.6 km)	Energy [KWh]	171.9	135.1	167.5	174.6	170.1
	Relative energy	1	0.786	0.974	1.016	0.990
	Δh [m]	46	0	-64.8	-18	80.7

Table A.1: Table comparing the energy consumption of four real roads with their simulated counterparts. Here, Δh denotes the elevation difference between the start and end points of each road. For each real road, five simulated roads were generated, and the average elevation difference along with the corresponding energy consumption was calculated.

DEPARTMENT OF SOME SUBJECT OR TECHNOLOGY
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se