Università degli Studi di Padova

Department of Information Engineering

Tesi di Laurea Triennale in Ingegneria dell'Informazione

Study and Evaluation of Correlation Techniques for Automotive Sensor Data

Supervisor
Marco Giordani
Università di Padova

Co-supervisor
Paolo Testolina
Università di Padova

Candidate
Matteo Pozzebon

To my parents, whose endless sacrifices have been my bedrock. To my friends, for the laughter and companionship amidst the challenges. And to the boundless world of discovery that fuels my journey.

Abstract

Fully autonomous driving systems are equipped with sensors able to guarantee fast detection and recognition of sensitive objects in the environment. However, the resulting huge volumes of data generated from those sensors may be challenging to handle for standard communication technologies.

Along these lines, in this thesis we test and validate different methods to evaluate correlation of automotive data on a synthetic dataset, to decide the most convenient way of transmitting data.

The study investigates correlation methods such as Pearson's correlation coefficient, the Chamfer Distance, the Iterative Closest Point (ICP), and the Normal Distribution Transform (NDT) algorithms. We also develop a new score function combining Chamfer Distance and correlation to detect significant changes in the automotive data.

The research utilizes a synthetic dataset and examines both static and dynamic contexts. Our findings highlight the interplay between correlation and Chamfer Distance in dynamic environments, and demonstrate that automotive data are highly correlated, even though redundancy is needed to guarantee accuracy, especially in critical areas where safety requirements are particularly critical.

I sistemi di guida completamente autonomi sono dotati di avanzati sensori in grado di rilevare e riconoscere rapidamente oggetti nell'ambiente circostante. Tuttavia, la generazione di enormi volumi di dati da parte di questi sensori può rappresentare una sfida considerevole per le tecnologie di comunicazione standard.

In questa tesi, ci proponiamo di esaminare e convalidare diversi metodi per valutare la correlazione dei dati automobilistici utilizzando un dataset sintetico, al fine di determinare il modo più efficace per trasmettere tali dati. La nostra ricerca esplora vari metodi di correlazione, tra cui il coefficiente di correlazione di Pearson, la Chamfer Distance, l'algoritmo Iterative Closest Point (ICP) e la Normal Distribution Transform (NDT). Inoltre, sviluppiamo una nuova funzione di punteggio che combina la Chamfer Distance e la correlazione al fine di individuare cambiamenti significativi nei dati automobilistici.

La nostra indagine si basa su un dataset sintetico e considera contesti sia statici che dinamici. I nostri risultati mettono in evidenza l'importante interazione tra correlazione e Chamfer Distance in ambienti dinamici e dimostrano che i dati automobilistici sono intrinsecamente correlati, pur richiedendo una certa ridondanza per garantire l'accuratezza, specialmente nelle aree critiche in cui sono essenziali i requisiti di sicurezza.

Contents

Listing of figures

1
Introduction

In the ever-evolving landscape of automotive technology, the integration of sensors and data-driven systems has become paramount in ensuring safety, efficiency, and overall performance. The increasing complexity and volume of data generated by these sensors [1] present both opportunities and challenges for the automotive industry. However, this surge in data production has posed significant challenges in terms of data transmission and processing.

The use of data received from external sources is essential for several reasons:

- Broad Coverage: It provides extensive geographic coverage and global information that a single vehicle's sensors cannot collect.

- Safety: It enables real-time alerts about accidents, road conditions, and emergency vehicles, enhancing road safety.

- Efficiency: It optimizes routes to avoid congestion and reduce travel times, contributing to efficiency and fuel savings.

The need to optimize data transmission from vehicles to cloud-based processing machines has become a pressing issue, driven by the desire to reduce bandwidth consumption, latency, and processing costs.

Key processing tasks include object detection, semantic segmentation, and lane tracking for safe navigation. Sensor fusion combines data from various sensors to enhance perception [2]. Anomaly detection identifies unusual patterns, while mapping and localization enable precise positioning. Additional tasks encompass traffic analysis and environmental monitoring for comprehensive automotive intelligence.

Transmitting all the raw sensor data to the cloud is not only resource-intensive but also often unnecessary, as a substantial portion of this data remains uninformative due to static environmental conditions.

To address these challenges, in this thesis, we analyze various correlation techniques for automotive sensor data, focusing on the study and evaluation of methods to reduce data transmission to a cloud-based processing machine.

This research is rooted in the "value of information" paradigm, as examined in [3, 4]. Our primary objective is to evaluate the appropriateness of data transmission. We achieve this by assessing various correlation techniques and devising a novel method that amalgamates diverse metrics. This approach is aimed at identifying instances when the surrounding environment has undergone significant changes, justifying the need for data transmission.

By achieving this, we aim to find the balance between the need for data accuracy and the efficient utilization of resources.

However, it's crucial to strike a balance in data transmission timing. Waiting too long to transmit may lead to outdated or non-representative data stored in the cloud server. This, in turn, can degrade the accuracy of autonomous driving operations. Therefore, finding the optimal moment for data transmission is paramount to ensure the effectiveness and reliability of autonomous driving systems.

While the importance of data from external sources is recognized, this research focuses on evaluating data transmission appropriateness using correlation techniques. Unlike previous studies, this thesis aims to strike a balance by identifying instances when the environment undergoes significant changes, justifying the need for data transmission. It also proposes a novel scoring function that combines Chamfer Distance and Pearson's Correlation Coefficient to optimize data transmission timing, considering both metrics. This approach contributes to the efficiency and reliability of autonomous driving systems, making it a novel contribution to the field.

Previous work [5, 6] has primarily focused on point cloud registration and change detection utilizing segmentation and data fusion with external datasets. In other work, such as [7] a neural network has been introduced to work in registered point clouds to detect environmental changes. In contrast, this thesis concentrates on the specific domain of automotive sensor data, with a focus on LiDAR sensor data. It explores various correlation algorithms, including Pearson's correlation coefficient, Chamfer Distance, Iterative Closest Point (ICP) registration algorithm, and Normal Distribution Transform (NDT) registration algorithm, and their applicability in identifying meaningful changes in the point cloud data. The development and evaluation of a new scoring function that combines multiple metrics to optimize data transmission timing represent a unique aspect of this research.

This thesis aims to investigate and evaluate correlation techniques that can be applied to automotive sensor data, specifically data obtained from top-mounted LiDAR sensors. The primary research objectives are as follows:

- To explore and understand the theoretical foundations and context of the study in the automotive sensor data domain.
- Study various correlation algorithms and their applicability in both static and dynamic contexts, focusing on Pearson's correlation coefficient, Chamfer Distance, Iterative Closest Point (ICP) registration algorithm, and Normal Distribution Transform (NDT) registration algorithm.
- Investigate the evolution of LiDAR point clouds and evaluate the effectiveness of the selected correlation techniques in identifying meaningful changes in the environment. Moreover, our goal is to develop and assess a new score function that combines Chamfer Distance and Pearson's Correlation Coefficient not only to identify significant changes in the point cloud data but also to optimize the decision regarding when it is appropriate to transmit data, considering both metrics.

This thesis is organized into five chapters. Each chapter builds upon the results and considerations of the previous one.

- Chapter 2 provides the theoretical framework and context for the study.
- Chapter 3 delves into the detailed analysis of sensor data filtering and the study of the correlation algorithms in different scenarios.
- Chapter 4 constitutes the core of the research, where the selected algorithms and tools are applied to an ordered collection of environmental data points (sequence) to detect changes in the surrounding environment and identify the optimal moment for data transmission.
- The final chapter provides a concise summary of our investigation and offers insights into potential future research directions.

This thesis concentrated on enhancing data transmission efficiency within automotive sensor systems through the application of correlation techniques. The research revealed that fine-tuning parameters, implementing downsampling, and configuring field of view settings have substantial effects on correlation outcomes. Additionally, a novel change score function was devised to facilitate the identification of noteworthy changes in the data. Throughout this study, we will discuss the methodologies, results, and insights gained from our investigations.

Autonomy will come a long way in the next few years.
There will not be a steering wheel in 20 years.

Elon Musk

2 Theory and Context of Study

This thesis exclusively relies on SELMA [8], a synthetic dataset tailored for autonomous driving research. SELMA distinguishes itself through its extensive multi-sensor data, featuring 24 different sensors, a vast and diverse dataset comprising over 20 million samples that span various weather and daytime scenarios. Additionally, SELMA aligns seamlessly with benchmark datasets and maintains an open-access policy. Notably, the dataset offers numerous scenes across varying traffic densities and locations, with an impressive sampling rate of 30 frames per second, significantly surpassing the typical 10 frames per second. This higher framerate proves invaluable for registration algorithms and correlation calculations, facilitating enhanced data processing and precision.

The ego vehicle incorporates three LiDARs, one is positioned atop the vehicle, while the remaining two are situated close to the front-right and front-left headlights.

Our primary focus will center on the data obtained from the LiDAR sensor positioned atop the vehicle, as it offers a comprehensive 360-degree field of view. In contrast, the other sensors are largely redundant for our study and do not contribute substantial information relevant to our specific applications.

Each LiDAR sensor produces a point cloud consisting of three-dimensional vectors. Since SELMA is a synthetic dataset, we have access to additional valuable information: We can associate each point with the specific object or actor it belongs to and we have precise information about the position of every object.

2.1 LiDAR point cloud transformation

The point clouds will undergo preprocessing steps such as Voxelization and Downsampling. Additionally, we will register the evolving point cloud, aligning it to the previous ones to account for the ego vehicle's movement during data acquisition.

Voxelization [9] is the process of converting a three-dimensional object or space into a grid of volumetric pixels, known as voxels, which can be thought of as the 3D equivalent of 2D pixels. Each voxel represents a small cubic element within the 3D space and can store various properties or attributes, such as color, density, or material information.

Downsampling [9] a point cloud via voxelization is a process that reduces the density of point data while preserving important structural information. In this study, we performed the following steps:

- We divided the 3D space containing the point cloud into a grid of equally sized cubic voxels.

- We placed each point from the original point cloud into the voxel it falls within, finding the voxel's coordinates that correspond to the point's location.

- If more than one point falls in the same voxel, an averaging aggregation operation is performed. This aggregates the attributes of the points within each voxel into a single value.

The resulting point cloud will have fewer points than the original, but it retains the overall structure and important features of the original data. Downsampling via voxelization has mainly been useful for reducing the computational load, and also for simplifying our data for visualization.

Registration refers to the process of aligning two or more 3D point clouds in a common coordinate system. This alignment is crucial when dealing with data obtained from sensors like LiDARs, as it ensures that the information collected from different viewpoints or time instances can be effectively combined and analyzed.

The Normal Distribution Transformation (NDT) and The Iterative Closest Point (ICP) algorithms are fundamental techniques for LiDAR point cloud registration.

ICP [10] operates iteratively to find the optimal transformation (translation and rotation) that minimizes the distances between corresponding points in two point clouds.

The NDT algorithm [11] models the distribution of points in the 3D space using Gaussian distributions. It is particularly suitable for aligning point clouds with complex structures or in environments with limited features.

Pearson's correlation coefficient [12, 13] is a statistical measure used in Lidar data analysis to quantify the degree of linear dependence between two sets of 3D point clouds. Understanding the correlation between different parts of a scene can aid in detecting objects and tracking their movement.

This coefficient is especially valuable for the change detection analysis we will perform. It will be used in conjunction with Chamfer Distance to determine the optimal transmission timing by utilizing correlation as one of the key factors.

It is obtained as follows:

$$C(a_i, b_i) = \frac{\sum_{j=1}^{N}(a_j - \bar{a})(b_j - \bar{b})}{\sqrt{\sum_{j=1}^{N}(a_j - \bar{a})^2 \sum_{j=1}^{N}(b_j - \bar{b})^2}}$$

Where:

- $C(a_i, b_i)$ is the normalized cross-correlation between points a_i and b_i.

- N is the number of corresponding points.

- $\bar{a}$ and $\bar{b}$ are the centroids of point clouds A and B, respectively.

- a_j and b_j. are the coordinates of the j-th corresponding point.

The Chamfer Distance metric addresses the problem of measuring the dissimilarity between two shapes by quantifying the cost or distance required to transform one shape into another. This cost is typically determined by the Euclidean distance between corresponding points, but variations of the algorithm may use different distance metrics based on specific application requirements.

In our study, the Chamfer Distance is calculated by summing the Euclidean distance between each point of one point cloud and the closest neighbor in the other point cloud divided by the total number of points in the starting point cloud:

$$\text{Chamfer}(A, B) = \frac{\sum_{a \in A} \min_{b \in B} \text{distance}(a, b)}{size(A)}$$

$$\text{Chamfer}(B, A) = \frac{\sum_{b \in B} \min_{a \in A} \text{distance}(b, a)}{size(B)}$$

The total chamfer distance between the point clouds is the resulting sum:

$$\text{Chamfer} = \text{Chamfer}(A, B) + \text{Chamfer}(B, A)$$

We performed all calculations within the MATLAB environment, utilizing the Computer Vision Toolbox [14].

Specifically, we employed the following algorithms which were already present in the toolbox: pcregistercorr for correlation, pcregisterndt for NDT registration, and pcregistericp for ICP registration.

Our goal is to make it safe and easy for people to move around. This technology can save thousands of lives, give people greater mobility, and free us from a lot of the things we find frustrating about driving today.

Chris Urmson

3

Study of Algorithms in static and dynamic contexts

In this chapter, we aim to fine-tune the voxelization parameters in terms of autocorrelation and cross-correlation, maximize the performance of registration algorithms, and conduct a comparison between correlation and the evolution of the machine's position.

3.1 Study of Correlation in a Static Environment

Our investigation commenced with an examination of the correlation algorithm considering a single data point. Initially, we delved into the autocorrelation analysis of the top-mounted LiDAR with itself, followed by an exploration of the cross-correlation between the top-mounted LiDAR and the combined data from the pair of front-mounted LiDARs. In both instances, our objective was to attain correlation results as close to 1 as possible.

The pcregistercorr algorithm offers flexibility through two key parameters: Grid Size and Grid Step. Grid Size denotes the dimensions of the square occupancy grid, expressed as a scalar value in meters. Conversely, Grid Step represents the size of individual grid cells, also expressed as a scalar value in meters. This iterative algorithm aims to determine the optimal transformation between the input pair of point clouds. In each step, it computes the correlation between the point clouds. Upon completion, it provides the best match as the output: the transformation with the highest correlation (correlation peak).

In order to enhance our results and optimize performance, we implemented some filtering steps. Points in close proximity to the vehicle(< 3m) have been eliminated to mitigate "blind spots" between sensors, which can lead to reduced correlation between the two point clouds. Such points serve no specific purpose and do not contribute meaningful information.

Similarly, points located at a significant distance from the vehicle(> 30m) have been removed for similar reasons. Given that sensors have varying perspectives, the captured views of objects in the environment can exhibit substantial differences. Consequently, this leads to the emergence of distinct point clusters at considerable distances from the vehicle, further diminishing the correlation between the two point clouds. While such information would certainly be beneficial for detecting distant vehicles and pedestrians, it is deemed unnecessary for the current analysis.

We also employed downsampling by selecting a voxel size of 0.5 meters for each voxel. This downsampling approach proved valuable in reducing the data density while preserving essential information, contributing to more efficient processing and analysis.

Additionally, points associated with the road have been removed due to their potential interference with the applied algorithms. We employed a straightforward selection criterion based on point height, removing points below 0.35 meters. Floor points, which represent the ground surface, tend to be stationary in a point cloud dataset regardless of the vehicle's motion. As the vehicle moves, these floor points remain in relatively fixed positions within the point cloud.

As a result, when analyzing the point cloud data, a high concentration of these stationary floor points can create an imbalance in the distribution of points. This imbalance can affect the calculation of correspondences or similarities between points from different scans or viewpoints. The point of maximum correspondence, which is the point in the point cloud that matches most closely with another point, tends to be biased towards these stationary floor points due to their abundance.

Consequently, by removing the floor points during downsampling or preprocessing, the point cloud is left with a more balanced distribution of points, reducing the influence of stationary features and allowing for more accurate correspondences to be established, especially with objects or structures that are not part of the ground surface.

Furthermore, without removing the floor points, we would inevitably observe a higher correlation between independent samples, which is an undesirable outcome.

To investigate the influence of parameter adjustments on correlation peaks, we executed the algorithm with numerous parameter combinations.

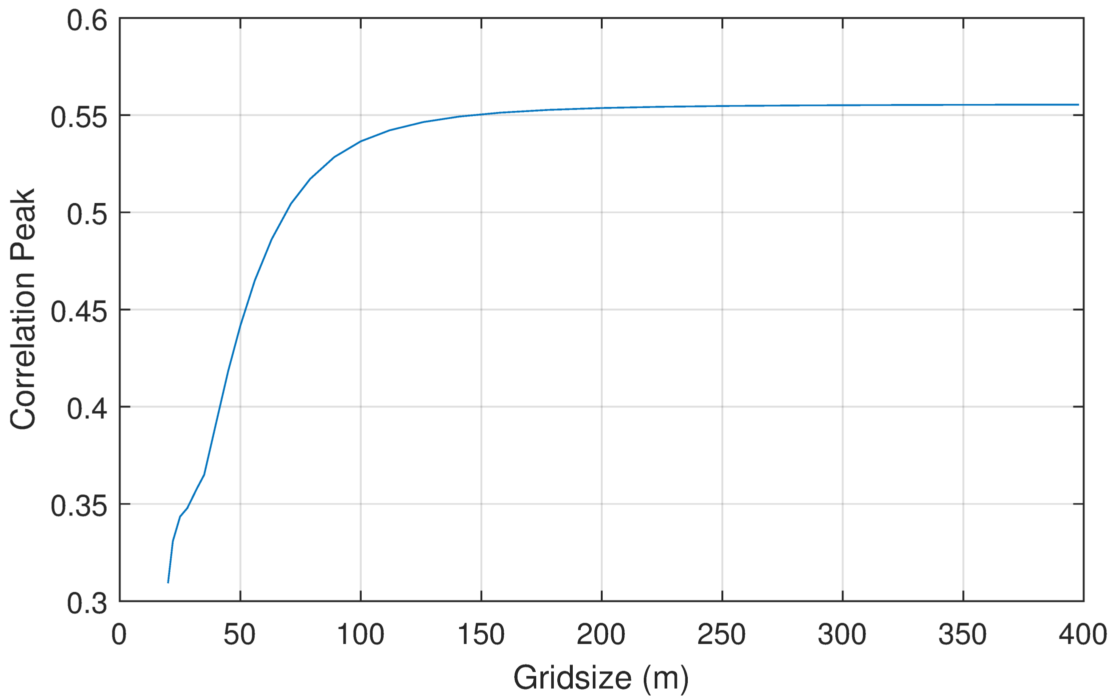

Figure 3.1: Top-Mounted LiDAR Auto Correlation Peak Results varying Grid Size

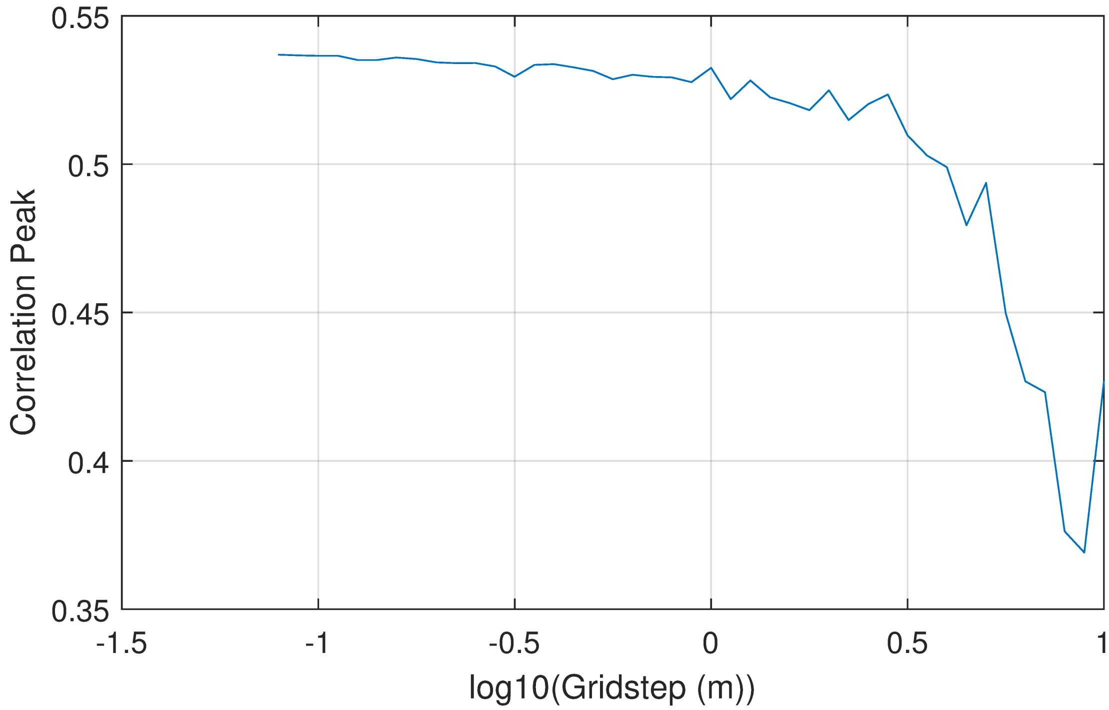

Figure 3.2: Top-Mounted LiDAR Auto Correlation Peak Results varying Grid Step

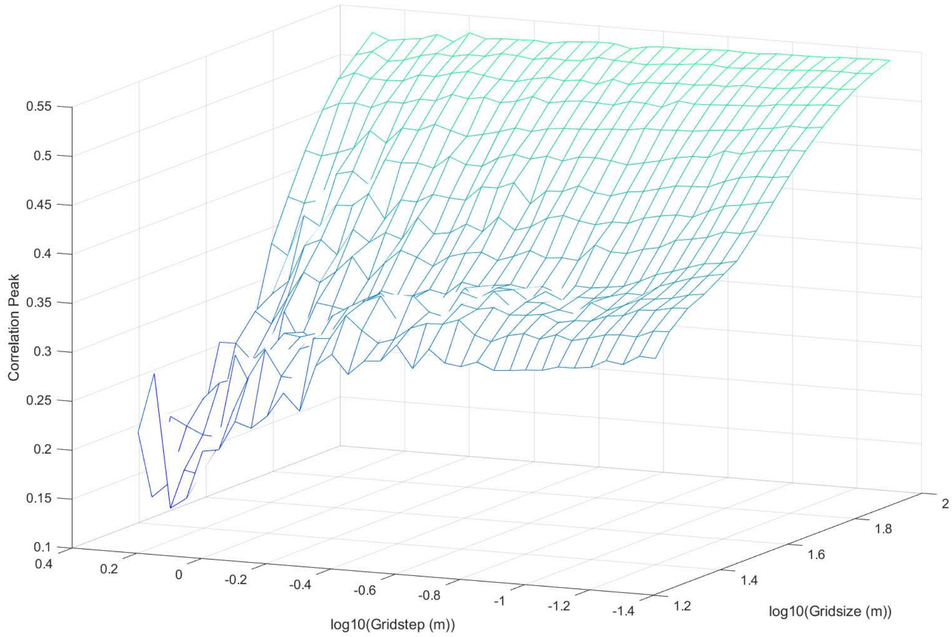

Figure 3.3: Top-Mounted LiDAR Auto Correlation Peak Results varying both parameters

In figure 3.1 we can see that increasing the grid size enhances the match, whereas in figure 3.2 increasing the grid step yields the opposite effect.

An interesting aspect to consider is when the correlation value surpasses 0.5 as the grid size expands to approximately 4-5 times the maximum distance. Beyond this threshold, the correlation appears to plateau at around 0.56 for larger values. In contrast, maintaining the grid step at acceptable levels (below 0.5 meters) has a minimal impact on the final outcome. However, it is worth noting that exceeding a certain value in the grid step poses a challenge to maintaining meaningful results, as it disrupts the coherence of the point cloud data analysis.

This divergence between the expected correlation values of 1 and our results can be attributed to the voxelization process, where the transformation of point cloud data into volumetric elements affects the alignment and matching of points. Figure 3.3 is a good summary of our results.

To complete this initial study, we assessed the cross-correlation between the union of the front LiDARs and the top-mounted LiDAR.

In the cross-correlation results, a noteworthy distinction becomes evident: reducing the

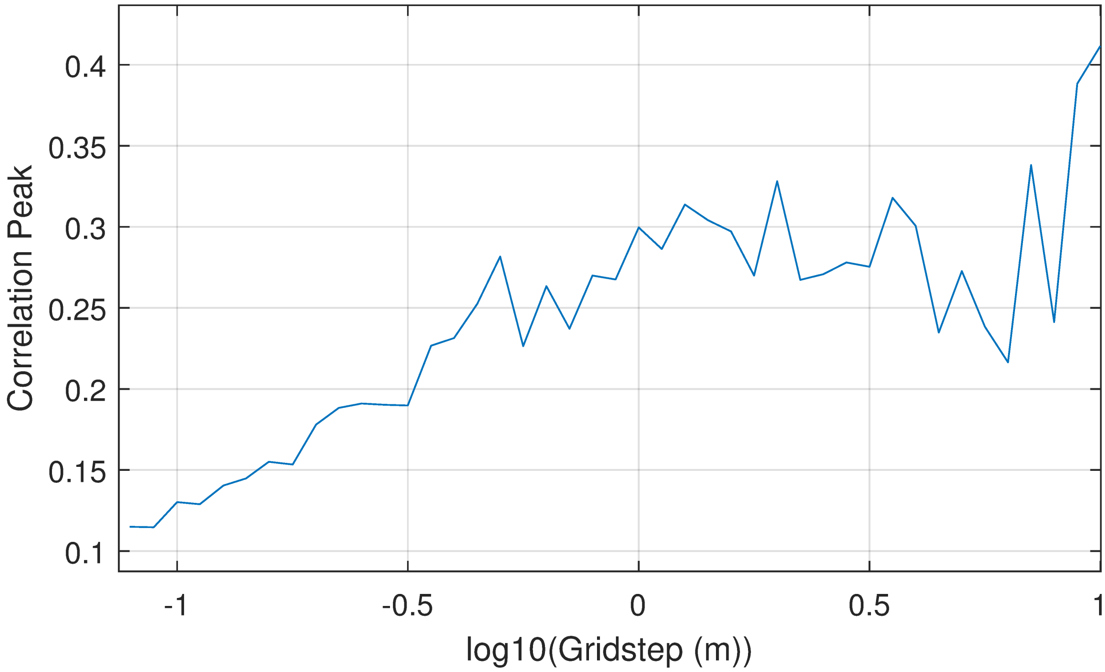

Figure 3.4: Cross-Correlation Peak Results varying Grid Step

grid step leads to a reduction in the resulting correlation peak (figure 3.4). This can be attributed to the fact that a higher grid step value effectively conceals the subtle variations and small changes that exist between the two point clouds, improving the match. Meanwhile, varying Grid Size doesn't yield any significant result in this scenario (figure 3.5). Figure 3.6 displays both variations simultaneously.

We can attest the best peak at 0.14. This observation holds particular significance when compared to the inherent auto-correlation of a point cloud, which consistently reaches a maximum correlation coefficient of 0.56.

Additionally, it is essential to emphasize that the translation value remains almost negligible in this context due to the precise alignment of all sensors during the pre-processing steps, rendering their offsets inconsequential.

When the grid step is raised significantly, we observe peaks in the Translation Vector and RMSE (figure 3.8 and 3.7), and the results start to lose coherence. Therefore, it can be concluded that you can increase the grid step to around 0.5m to "reduce noise" resulting from the different sensor perspectives while maintaining a significant correlation value, provided that the RMSE and translation vector modulus remain acceptable.

If the step size is raised beyond 10 meters, it is possible to achieve a cross-correlation of 1,

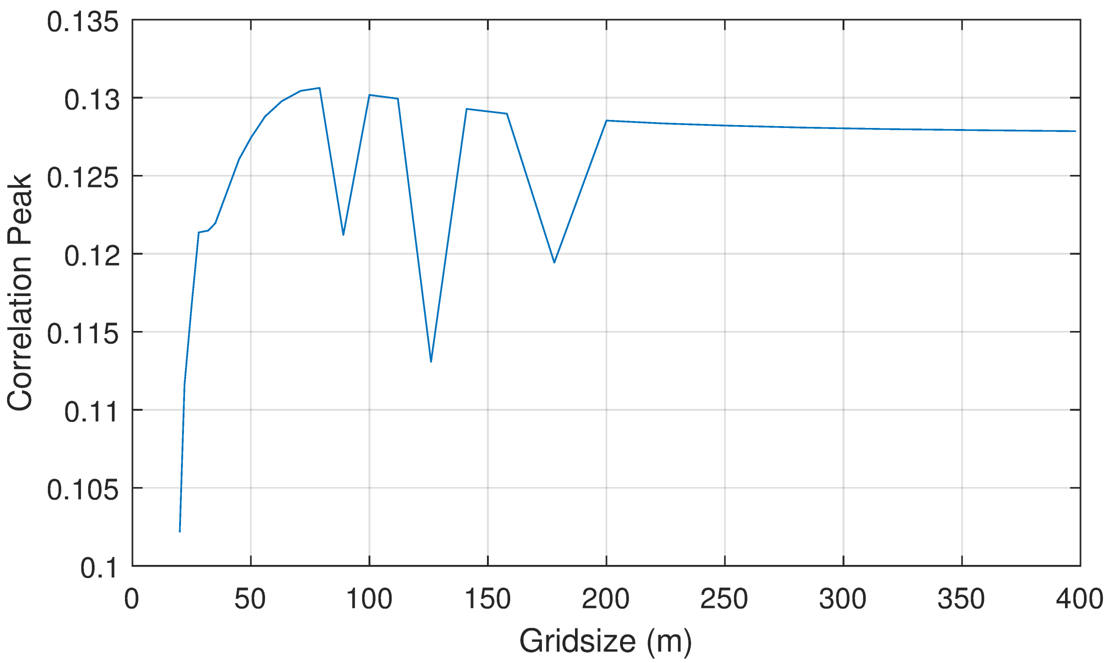

Figure 3.5: Cross-Correlation Peak Results varying Grid Size

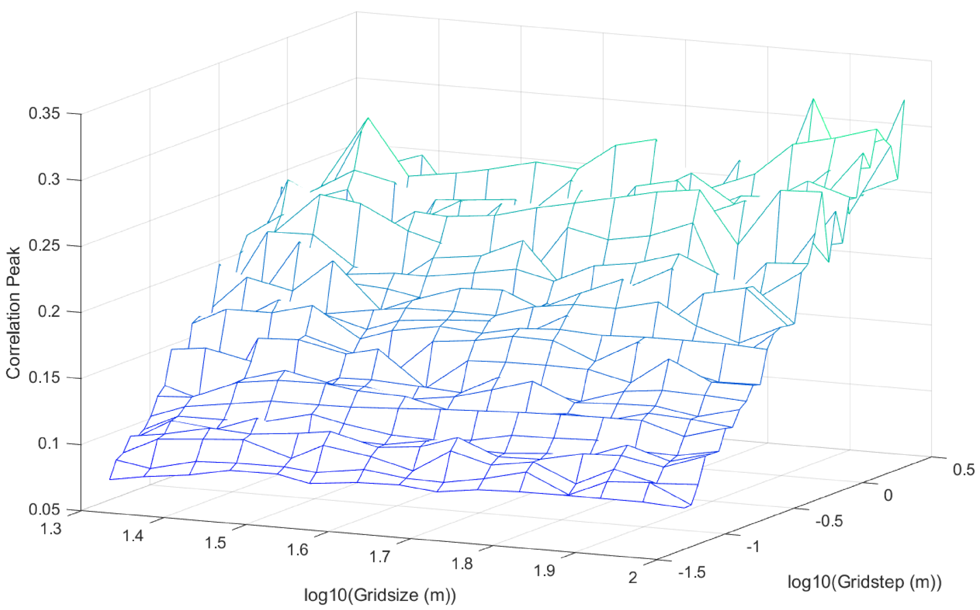

Figure 3.6: Cross-Correlation Peak Results varying both parameters

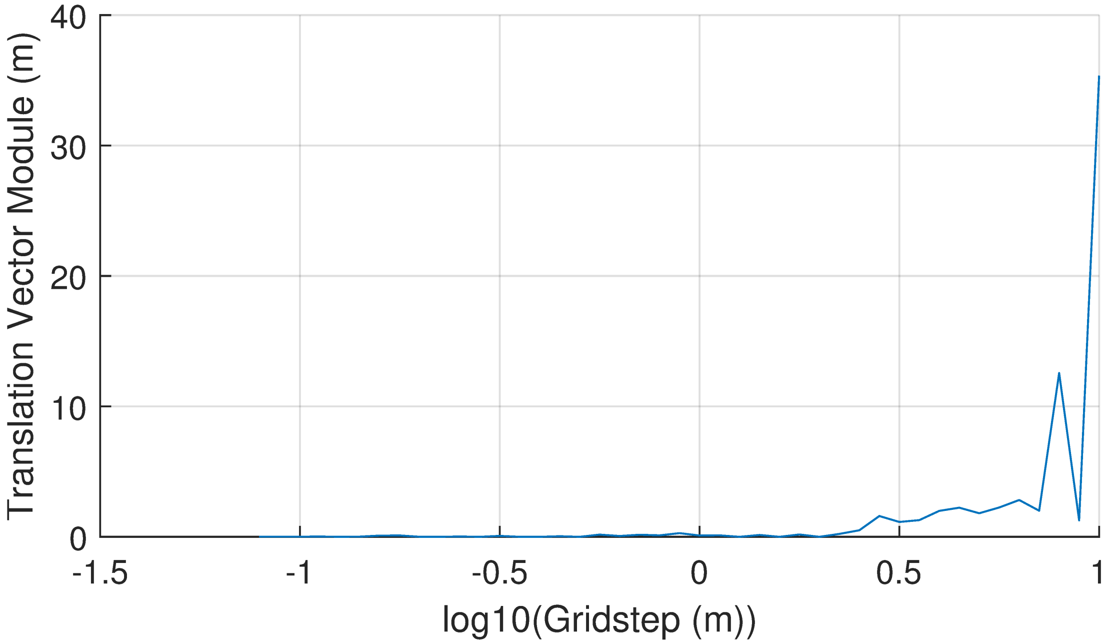

Figure 3.7: Translation Vector Module varying Grid Step

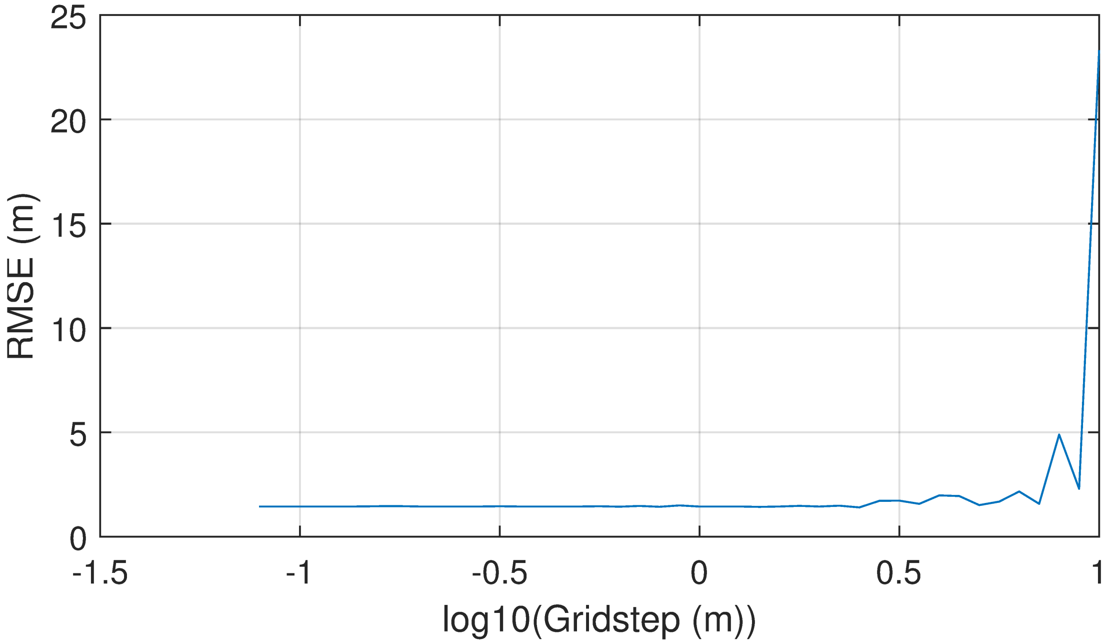

Figure 3.8: RMSE varying Grid Step

however, the steps are so large that the results lose significance as the predicted transformation drifts out and becomes incorrect. We have also experimented with other scenarios, such as a turn at a busy intersection with many nearby vehicles, and these observations hold true even in critical conditions (though, naturally, the correlation value is lower in such cases).

3.2 Study of Correlation in a Dynamic Environment

The next step is the analysis of a dynamic scenario. Here we tested the performance of the 3 registration algorithms in our study.

Based on the results of the previous section, it has been determined that reasonable parameters for the study of dynamic scenes can be set as Grid Step $= 0.5m$ and Grid Size $=$ FOV$(m) * 3.5$.

The accuracy of registration through pcregistercorr and pcregisterndt is intimately tied to the judicious selection of voxelization parameters. In the previous section, we conducted a study in a controlled environment to identify acceptable parameters that yield meaningful results (with low Root Mean Square Error and translation vector) without imposing excessive computational demands.

While ICP does not require any specific parameters, NDT involves voxelizing the space using a grid step value similar to the one employed in the correlation algorithm. Both algorithms follow a coarse-to-fine approach, implying that it is advisable to supply the algorithm with an initial predicted transformation. This initial transformation serves as the starting point for the iterative process. Consequently, we will propagate the transformation acquired from the previous time step into the subsequent one.

The accuracy of registration algorithms can be enhanced by expanding the area under consideration. By doing so, the sensor's field of view encompasses more static points from the background, which aids in the registration process.

Expanding the field of view enhances registration accuracy and fosters stability, particularly in instances far away from the initial reference point. Conversely, reducing the field of view reduces the computational complexity but at the cost of diminished accuracy.

The accuracy of the transformation operation can be measured against the true data (i.e., our ground truth) available from SELMA. Figure 3.9 presents the results obtained for a field of view of 35m and figure 3.10 for 70m.

We also computed the Root Mean Square Error (RMSE) between the actual translation magnitude and the predicted value generated by each algorithm. On average, in the first

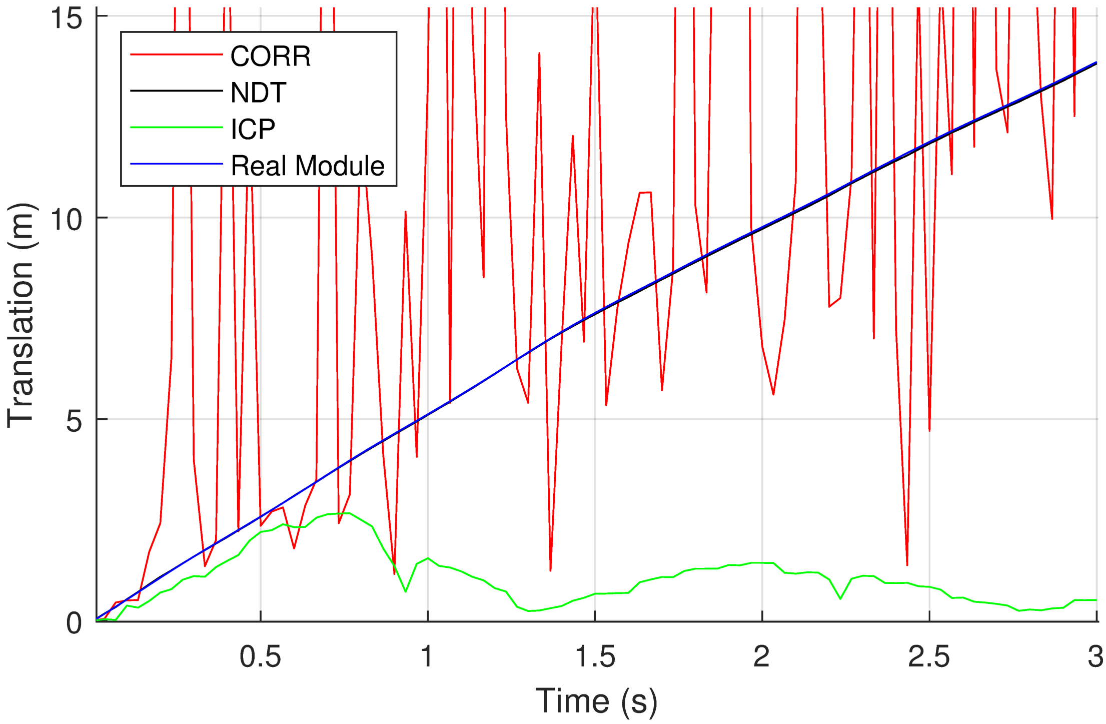

Figure 3.9: Registration Algorithms comparison with 35m FOV

three seconds of a dynamic sequence with a FOV of 35m, NDT has an RMSE of 0.05m, ICP of 7.5m, and CORR of 13m. For a larger FOV of 70m, the RMSE drops, as more static points from the background are considered. NDT has an RMSE of 0.04m, ICP of 0.9m, and CORR of 0,15m.

NDT consistently exhibits the highest accuracy, regardless of whether the field of view is wide or narrow.

In contrast, while ICP demonstrates remarkable speed and reliability within a limited time frame, its accuracy diminishes rapidly as the point cloud evolves.

The correlation algorithm is the slowest and most complex one as it also performs the correlation calculation.

3.3 Comparison between Car Movement and Metrics Change

The final analysis step, preceding our exploration of the core objective, involves examining the evolution of the metrics: correlation and Chamfer Distance.

To normalize the correlation peaks within the expected [0,1] range, we can consider the

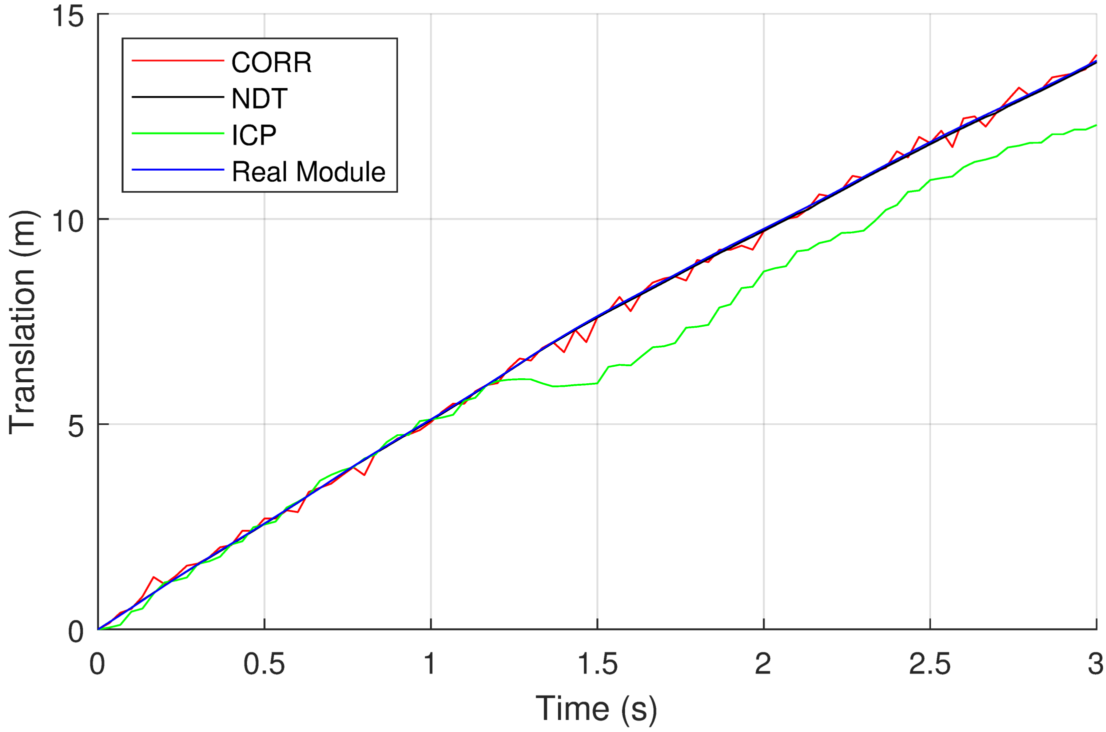

Figure 3.10: Registration Algorithms comparison with 70m FOV

autocorrelation of the initial sample as a perfect match and then divide all subsequent correlation peaks by this reference value.

In a dynamic scenario, where the car is moving, correlation decreases rapidly after a few frames, and registration errors become more frequent, primarily due to the entry and exit of point clusters. As shown in Figures 3.11 and 3.12, correlation typically experiences a drop when objects or actors enter/exit the sensor's field of view. There is not much distinction between whether these objects are mobile or stationary; the primary factor influencing correlation is the movement of the ego vehicle. In contrast, chamfer distance is more sensitive to the movements of actors within the sensor's field of view. This difference in sensitivity highlights the distinct nature of these two metrics in capturing changes within the environment, with correlation primarily reflecting vehicle motion and Chamfer Distance responding to the movement of objects or actors in the vicinity. Chamfer Distance typically exhibits a slower rate of increase in response to these changes.

Moreover, by reducing the maximum distance, the entry and exit of point clusters become more frequent, leading to a reduced percentage of background points in the point cloud. This adjustment influences correlation by causing an earlier drop. It also results in chamfer

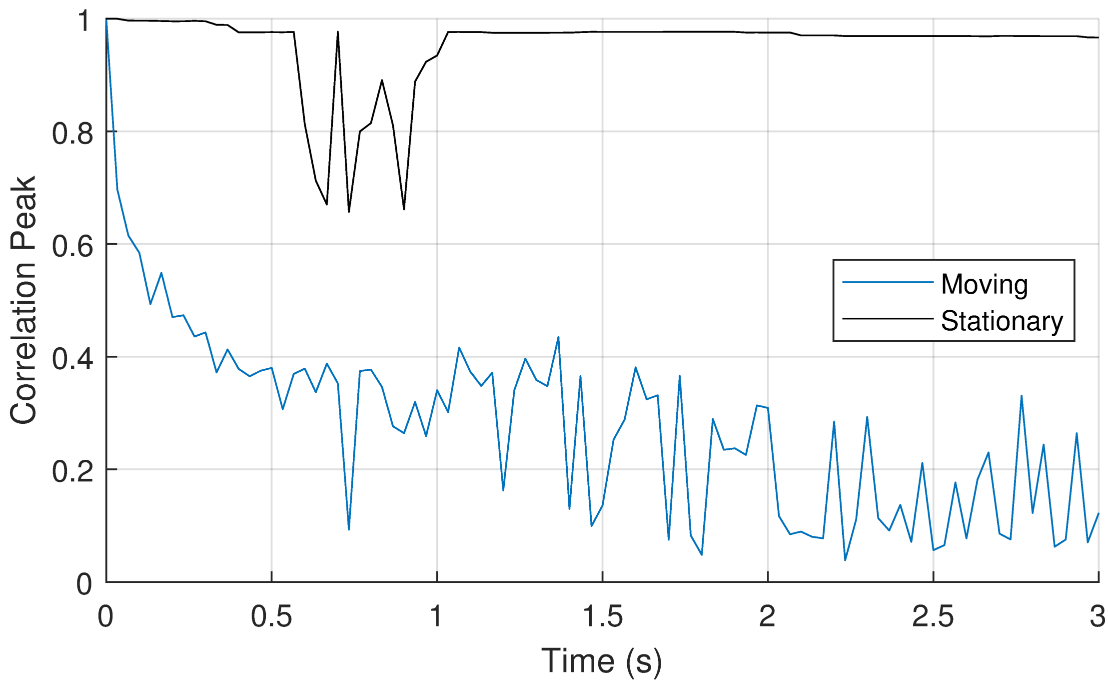

Figure 3.11: Correlation Peak evolution comparison with 50m FOV

distance increasing more rapidly. This effect occurs because the algorithm will place greater emphasis on nearby clusters, which are typically moving vehicles.

Prior registration of the point cloud significantly assists the Chamfer Distance algorithm. An accurate prior registration enables the algorithm to compute the Euclidean distance primarily on the moving objects, optimizing its performance and focus.

On the contrary, when the vehicle remains stationary, the correlation remains highly stable, and the Chamfer Distance exhibits slower growth.

An intriguing phenomenon occurs when the field of view is narrow(< 35m). We can observe an oscillatory behavior in the chamfer distance without previous registration (figure 3.13), which is likely induced by nearby vehicles in the opposite lane. As these vehicles progress, the one behind gradually assumes the position of the one in front, resulting in a reduction in the distance between the clusters of the starting and current point clouds. This highlights the importance of point cloud registration as it helps mitigating this phenomenon.

In summary, these observations offer valuable insights into the behavior of correlation and chamfer distance metrics in dynamic environments. They emphasize the importance of considering both the specific context and the nuances of the metrics themselves when interpreting the results of point cloud analysis, ultimately contributing to a more comprehensive

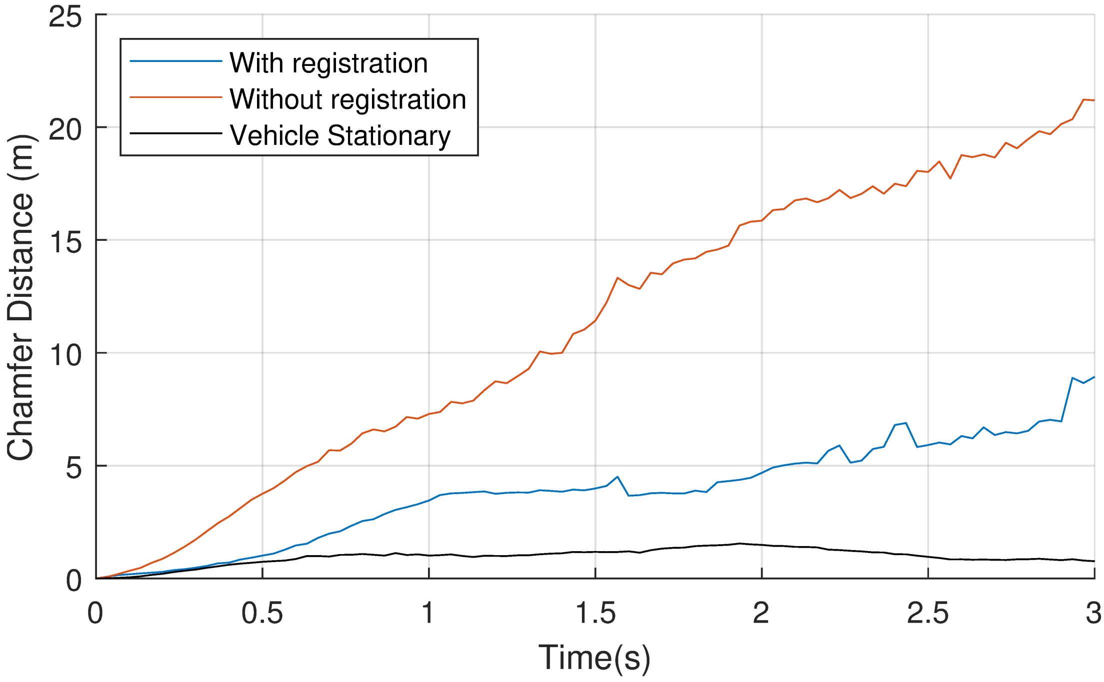

Figure 3.12: Chamfer Distance evolution comparison with 50m FOV

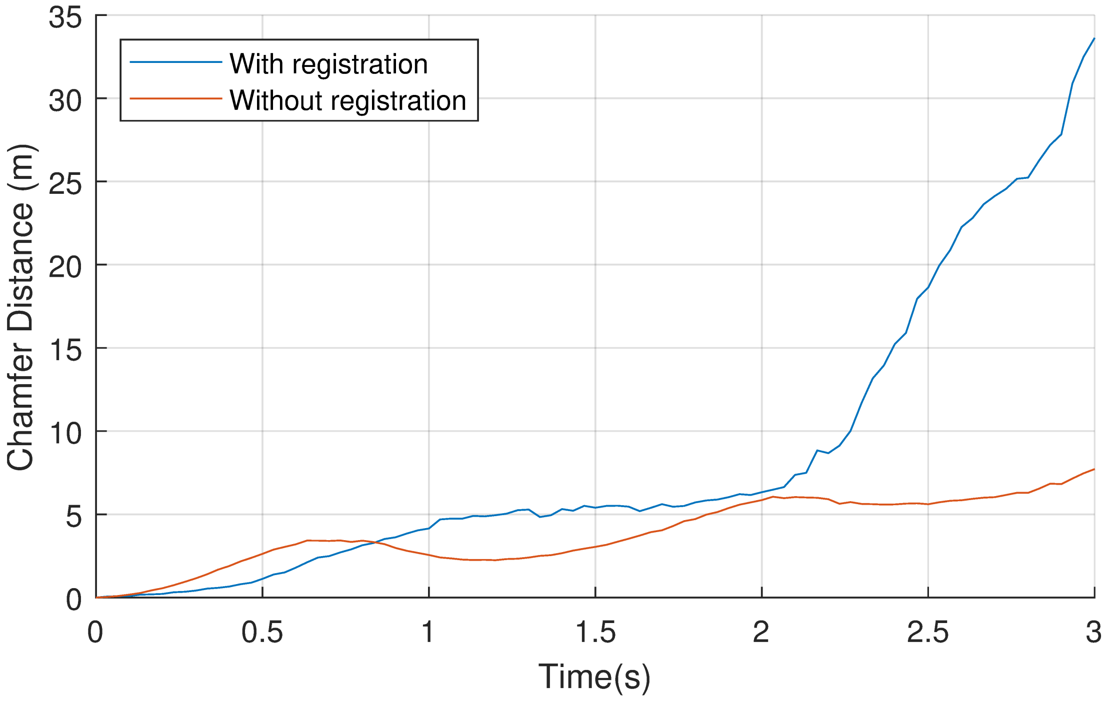

Figure 3.13: Chamfer Distance oscillating behavior with narrow FOV 30m

understanding of the environment and its dynamics.

Fully self-driving cars are here. It's not a promise; it's not a research project. It's not a limited pilot, a demo, or a beta test. It's here.

John Krafcik

4 Study of change detection in LiDAR point clouds

Change Detection is the process of identifying and quantifying differences or alterations in the environment as observed through automotive sensor data. In this chapter, we delve into the evaluation and optimization of change detection in automotive sensor data using correlation and Chamfer Distance metrics. We will discuss the methodologies employed to determine the parameters, their implications in various scenarios, and the potential for parameter optimization.

4.1 A Novel Change Detection Score Function

Our approach involves calculating the correlation and Chamfer Distance peaks (after point cloud registration) to assess changes in the environment. A simple change score can be introduced: it is defined with a percentage coefficient, α, and consequently ($1 - \alpha$), allowing us to weigh the importance of correlation and Chamfer Distance differently. Our change score is formulated as:

$$\text{score} = \alpha \left(\frac{1}{\text{correlation}} - 1 \right) + (1 - \alpha) \cdot \text{Chamfer Distance} \tag{4.1}$$

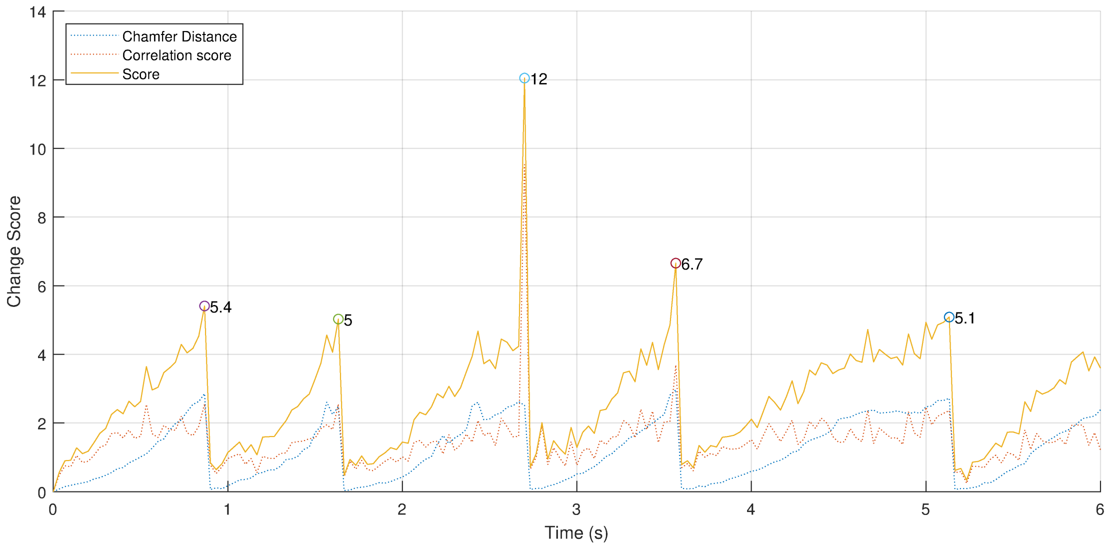

Figure 4.1: Score Function Evolution for a dynamic environment with FOV=50m

In our experiments, we carefully selected values for α, the threshold, and the FOV distance to strike a balance between the two metrics and optimize change detection.

In this chapter, we set the desired average actor(vehicles and pedestrians) movement to 3 meters and we fine-tuned the parameters by referencing the ground truth data from SELMA, enabling us to calculate the average movement of nearby actors effectively.

To accomplish this, the score function's threshold was set at 5, and data transmission was triggered when this threshold was exceeded. Figure 4.1 illustrates the evolution of the metrics. Upon reaching the threshold, data transmission commenced. We then designated the current data point as the new starting point and tracked the metrics' evolution from that moment onward.

With our parameters, in scenarios with the ego vehicle in motion, the retransmission point typically falls between 0.5 and 1.5 seconds, depending on the aforementioned conditions.

Figure 4.2 displays the average movement of nearby actors measured from the initial data point, indicating the value at the time of data transmission.

In addition, we considered initiating transmission if two consecutive instances surpassed the score threshold. This precaution aimed to mitigate false positives in cases where one of the score functions erroneously identified a peak. However, such occurrences were infrequent, and we opted for a streamlined approach to avoid unnecessary complexity.

Correlation typically experiences a drop when objects or actors enter or exit the sensor's field of view (figure 3.11), which is set at 50 meters in our setup. There is not much difference

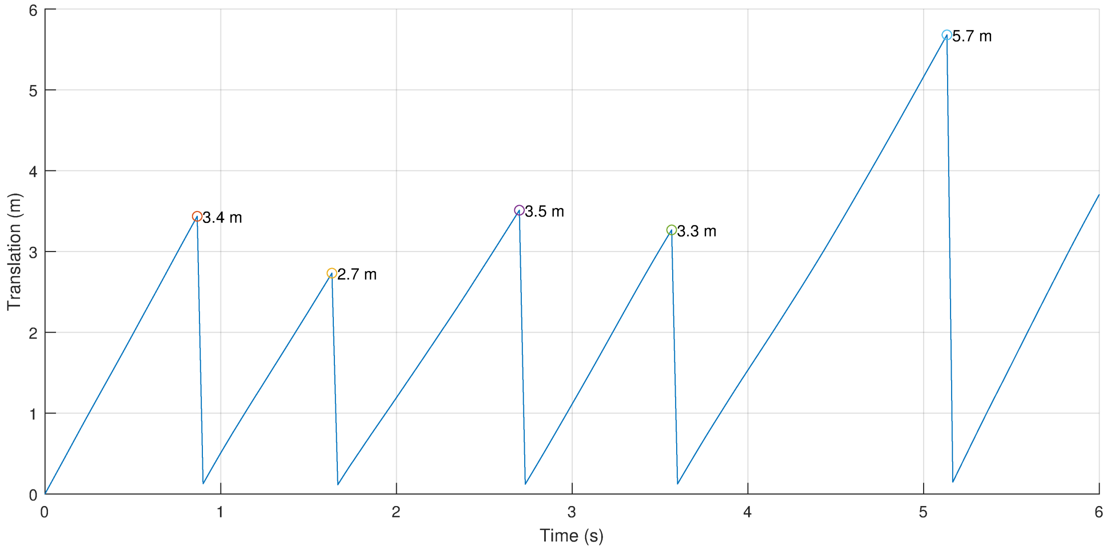

Figure 4.2: Average nearby Actors Movement for a dynamic environment with FOV=50m

between whether these objects are mobile or stationary; the primary factor influencing correlation is the movement of the ego vehicle. Chamfer Distance, on the other hand, is more sensitive to the movements of actors within the sensor's field of view.

These parameter values were chosen after experimentation and proved effective across various scenarios. Ideally, training a neural network could potentially yield even more optimal parameter values. However, it is worth noting that our currently selected values for α, the threshold and the FOV distance already achieve a good balance.

In this setup, both metrics complement each other well. In cases of a radical change in the scenario, we can adjust the weight given to each metric based on the specific change considered more important.

To understand the impact of different traffic intensities and actor speeds (including the ego vehicle), we conducted various experiments. Notably, we observed that the inter-transmission interval tended to shorten when there was an increase in the number of vehicles involved and their speed. Conversely, transmission was delayed in conditions with fewer vehicles or slower speeds.

We also experimented with varying the maximum distance parameter and observed consistent and interesting results. Reducing the maximum distance from 50 meters to 35 meters limits the inclusion of static background points in the processed point cloud. Consequently, this adjustment challenged the correlation algorithm, making it more difficult to achieve an accurate match and leading to more frequent score peaks. We increased the contribution of

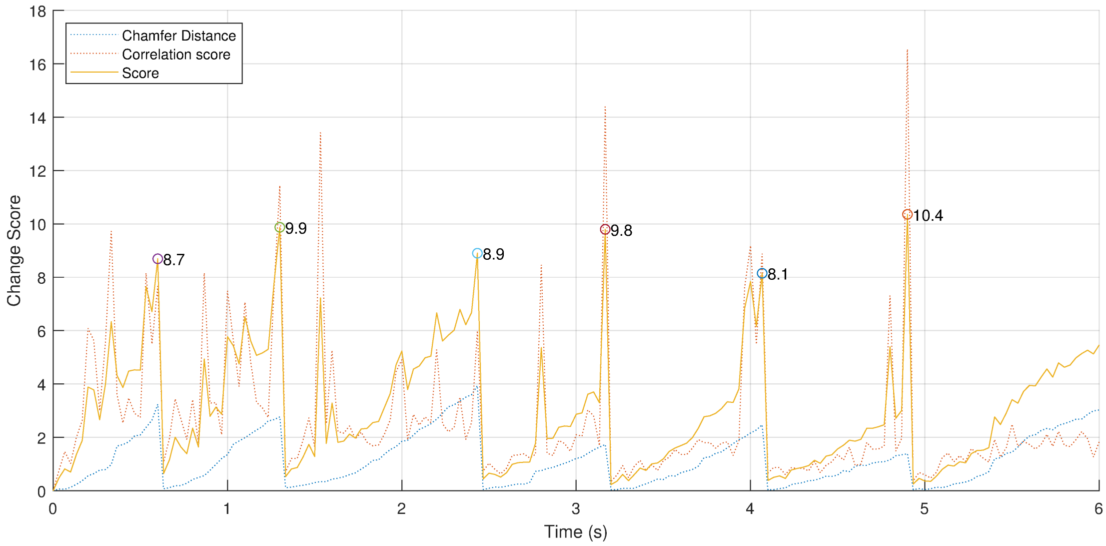

Figure 4.3: Score Function Evolution for a dynamic environment with FOV=35m

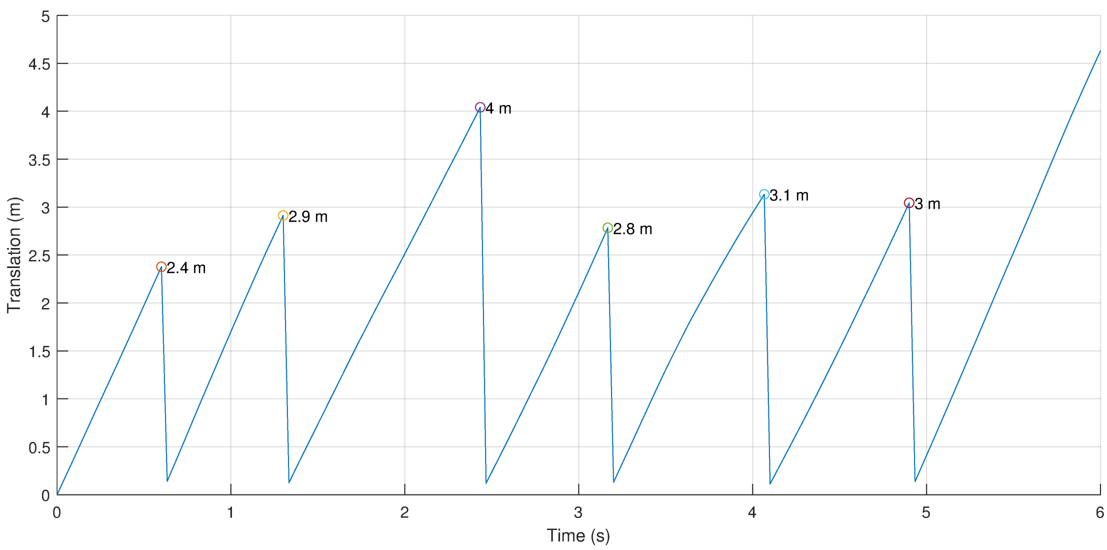

Figure 4.4: Average nearby Actors Movement for a dynamic environment with FOV=35m

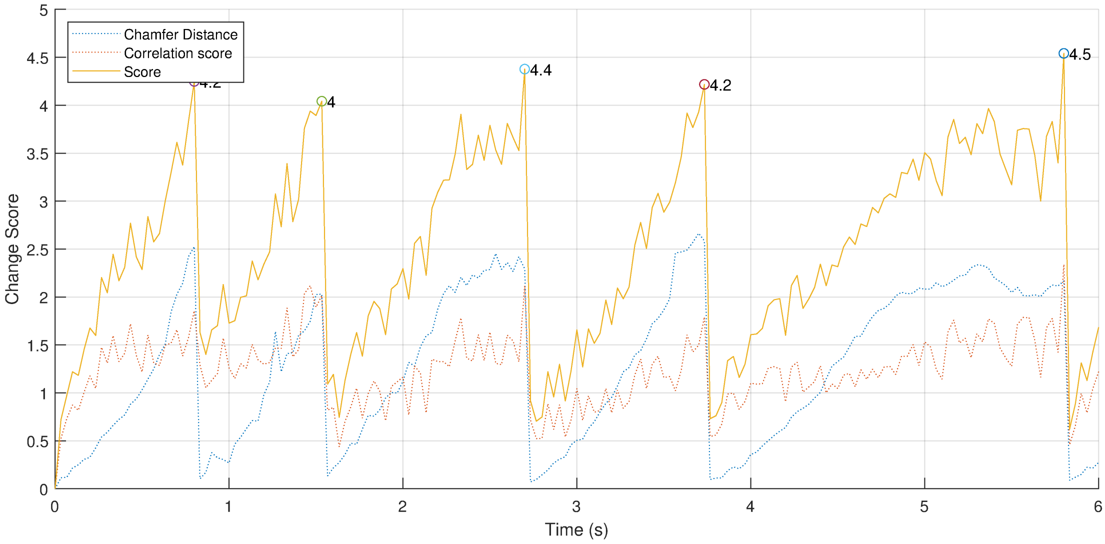

Figure 4.5: Score Function Evolution for a dynamic environment with FOV=70m

Chamfer Distance to the score function (to 75%) while decreasing the weight given to correlation (to 25%) to maintain the desired average actor movement. We also raised the score function threshold from 5 to 8 to increase stability. (Figure 4.3 and 4.4)

Conversely, increasing the maximum distance from 50 meters to 70 meters resulted in slower growth and fewer peaks in both correlation and Chamfer Distance. We adjusted the weight distribution to 60% correlation and 40% Chamfer Distance, with a threshold set to 4. (Figure 4.5 and 4.6)

Adjusting the maximum distance parameter can be beneficial depending on the scenario, with wider sensor fields of view being useful for rural areas and narrower fields of view for urban environments.

In static scenarios (e.g. when the ego vehicle is stationary at a traffic light) transmission may not be triggered due to a lack of significant environmental changes. To address this issue, we considered the option of forcing retransmission after a few seconds to account for potential changes that might not be detected otherwise.

A different approach is to once again optimize the parameters: we can set the threshold to 2, give more weight to the Chamfer Distance (75%), and reduce the FOV to 35m. The results are displayed in figure 4.7 and 4.8.

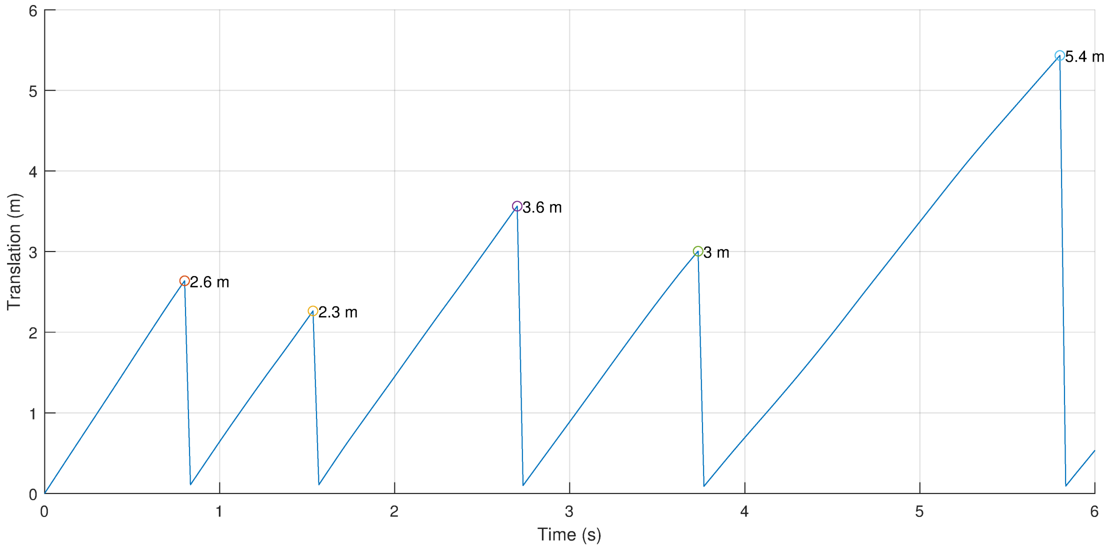

Figure 4.6: Average nearby Actors Movement for a dynamic environment with FOV=70m

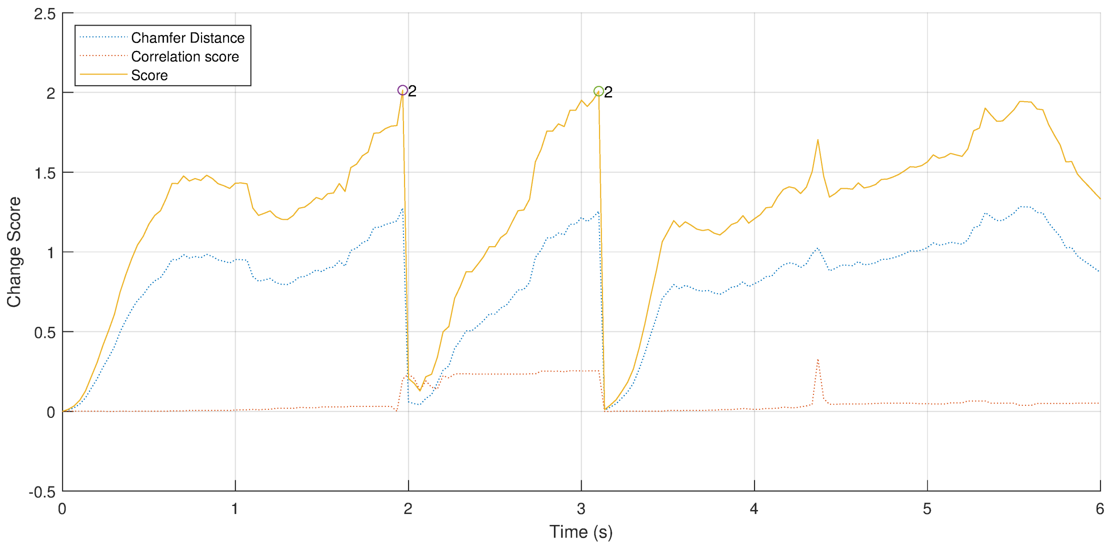

Figure 4.7: Score Function Evolution for a static environment with FOV=35m

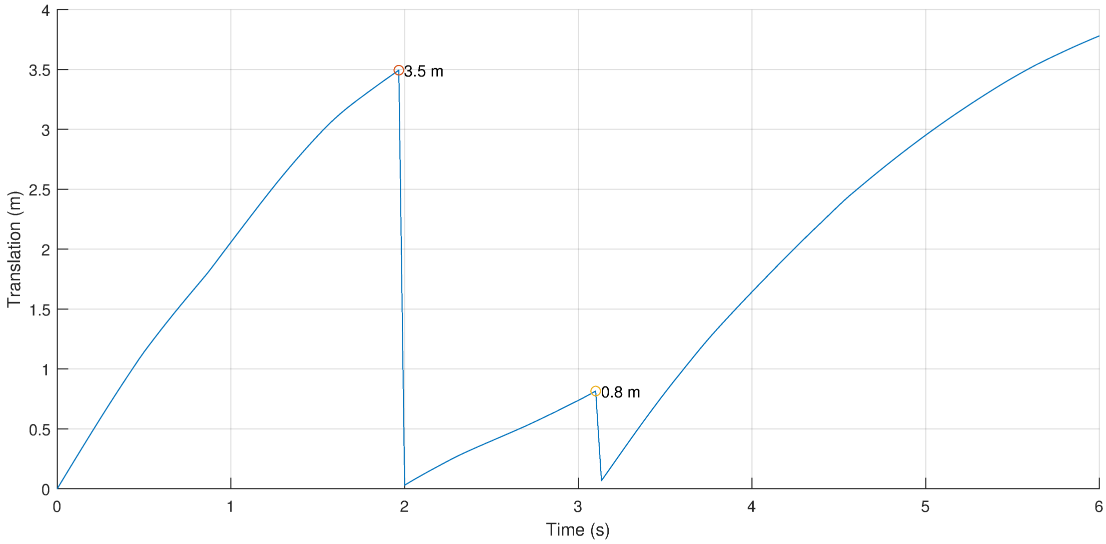

Figure 4.8: Average nearby Actors Movement for a static environment with FOV=35m

4.2 Change Detection Conclusions

In this chapter, we have examined and fine-tuned change detection. We introduced a new change score function that combines correlation and Chamfer Distance metrics. We carefully selected parameters to strike an optimal balance between these metrics, ensuring effective change detection. Our experiments revealed that the retransmission moment is influenced by factors such as traffic intensity and vehicle speed, with more vehicles and higher speeds triggering earlier retransmission. We also explored parameter adjustments for various scenarios, demonstrating adaptability between urban and rural environments.

5 Conclusion and Future Works

This thesis has delved into the intricate world of automotive sensor data, with a specific focus on correlation techniques to optimize data transmission and resource utilization. Through an in-depth study of algorithms in static and dynamic contexts, including Pearson's correlation coefficient, Chamfer distance, ICP, and NDT, we have gained valuable insights into their applicability and performance.

In static scenarios, we observed the importance of parameter adjustments, such as grid size and grid step, in optimizing correlation results. Downsampling and the removal of irrelevant points, such as those near the vehicle or on the road, were shown to enhance correlation outcomes. Furthermore, the analysis of dynamic scenes emphasized the significance of field of view settings in registration algorithms, with NDT displaying superior accuracy.

The study of change detection in LiDAR point clouds revealed the nuanced behavior of correlation and Chamfer distance metrics in response to environmental changes. By devising a change score function, we provided a practical means of identifying significant changes. Our experiments highlighted the influence of traffic intensity, actor speed, and sensor field of view on the retransmission point, offering valuable guidance for parameter optimization.

While our chosen parameter values strike a balance for various scenarios, there is potential for further optimization, such as through machine learning techniques. Future research could explore automated parameter selection and adaptive change detection strategies to enhance the robustness of automotive sensor systems.

References

[1] C. Cao, M. Preda, and T. Zaharia, "3d point cloud compression: A survey," in *The 24th International Conference on 3D Web Technology*, ser. Web3D '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 1–9. [Online]. Available: https://doi.org/10.1145/3329714.3338130

[2] V. Rossi, P. Testolina, M. Giordani, and M. Zorzi, "On the role of sensor fusion for object detection in future vehicular networks," in *2021 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, 2021, pp. 247–252.

[3] T. Higuchi, M. Giordani, A. Zanella, M. Zorzi, and O. Altintas, "Value-anticipating v2v communications for cooperative perception," in *2019 IEEE Intelligent Vehicles Symposium (IV)*, 2019, pp. 1947–1952.

[4] M. Giordani, A. Zanella, T. Higuchi, O. Altintas, and M. Zorzi, "Investigating value of information in future vehicular communications," in *2019 IEEE 2nd Connected and Automated Vehicles Symposium (CAVS)*, 2019, pp. 1–5.

[5] Örkény Zováthi, B. Nagy, and C. Benedek, "Point cloud registration and change detection in urban environment using an onboard lidar sensor and mls reference data," *International Journal of Applied Earth Observation and Geoinformation*, vol. 110, p. 102767, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0303243422000939

[6] B. Gálai and C. Benedek, "Change detection in urban streets by a real time lidar scanner and mls reference data," in *Image Analysis and Recognition*, F. Karray, A. Campilho, and F. Cheriet, Eds. Cham: Springer International Publishing, 2017, pp. 210–220.

[7] B. Nagy, L. Kovács, and C. Benedek, "Changegan: A deep network for change detection in coarsely registered point clouds," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 8277–8284, 2021.

[8] P. Testolina, F. Barbato, U. Michieli, M. Giordani, P. Zanuttigh, and M. Zorzi, "Selma: Semantic large-scale multimodal acquisitions in variable weather, daytime and viewpoints," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 7, pp. 7012–7024, 2023.

[9] F. Pomerleau, F. Colas, R. Siegwart, and S. Magnenat, "Comparing icp variants on real-world data sets," *Autonomous Robots*, pp. 133–148, 34 2013.

[10] P. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992.

[11] P. Biber and W. Strasser, "The normal distributions transform: a new approach to laser scan matching," in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003) (Cat. No.03CH37453)*, vol. 3, 2003, pp. 2743–2748 vol.3.

[12] Y. Huang and F. Da, "Registration algorithm for point cloud based on normalized cross-correlation," *IEEE Access*, vol. 7, pp. 137 136–137 146, 2019.

[13] P. Holmes, "Correlation: From picture to formula," *Teaching Statistics*, vol. 23, no. 3, pp. 67–71, 2001. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-9639.00058

[14] "Matlab: Computer vision toolbox," https://it.mathworks.com/products/computer-vision.html.