
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Maanpää, Jyri; Pesonen, Julius; Hyyti, Heikki; Melekhov, Iaroslav; Kannala, Juho; Manninen, Petri; Kukko, Antero; Hyyppä, Juha

Dense Road Surface Grip Map Prediction from Multimodal Image Data

Published in:
Pattern Recognition - 27th International Conference, ICPR 2024, Proceedings

DOI:
[10.1007/978-3-031-78447-7_26](https://doi.org/10.1007/978-3-031-78447-7_26)

Published: 03/12/2024

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Maanpää, J., Pesonen, J., Hyyti, H., Melekhov, I., Kannala, J., Manninen, P., Kukko, A., & Hyyppä, J. (2024). Dense Road Surface Grip Map Prediction from Multimodal Image Data. In A. Antonacopoulos, S. Chaudhuri, R. Chellappa, C.-L. Liu, S. Bhattacharya, & U. Pal (Eds.), *Pattern Recognition - 27th International Conference, ICPR 2024, Proceedings* (pp. 387–404). (Lecture Notes in Computer Science; Vol. 15317). Springer.
https://doi.org/10.1007/978-3-031-78447-7_26

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.



Dense Road Surface Grip Map Prediction from Multimodal Image Data

Jyri Maanpää^{1,2}(✉) , Julius Pesonen¹ , Heikki Hyyti¹ ,
Iaroslav Melekhov² , Juho Kannala^{2,3} , Petri Manninen¹ ,
Antero Kukko¹ , and Juha Hyypä¹

¹ Finnish Geospatial Research Institute FGI, National Land Survey of Finland,
02150 Espoo, Finland

{jyri.maanpaa,julius.pesonen}@nls.fi

² Department of Computer Science, Aalto University, 02150 Espoo, Finland

³ University of Oulu, 90570 Oulu, Finland

Abstract. Slippery road weather conditions are prevalent in many regions and cause a regular risk for traffic. Still, there has been less research on how autonomous vehicles could detect slippery driving conditions on the road to drive safely. In this work, we propose a method to predict a dense grip map from the area in front of the car, based on postprocessed multimodal sensor data. We trained a convolutional neural network to predict pixelwise grip values from fused RGB camera, thermal camera, and LiDAR reflectance images, based on weakly supervised ground truth from an optical road weather sensor.

The experiments show that it is possible to predict dense grip values with good accuracy from the used data modalities as the produced grip map follows both ground truth measurements and local weather conditions, such as snowy areas on the road. The model using only the RGB camera or LiDAR reflectance modality provided good baseline results for grip prediction accuracy while using models fusing the RGB camera, thermal camera, and LiDAR modalities improved the grip predictions significantly.

Keywords: Grip prediction · Autonomous driving · Convolutional neural networks

1 Introduction

Harsh winter conditions pose unique challenges to autonomous driving. According to the Road Weather Management Program by the U.S. Department of

J. Maanpää and J. Pesonen—Shared an equal contribution to this work as first authors.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-78447-7_26.

© The Author(s) 2025

A. Antonacopoulos et al. (Eds.): ICPR 2024, LNCS 15317, pp. 387–404, 2025.

https://doi.org/10.1007/978-3-031-78447-7_26

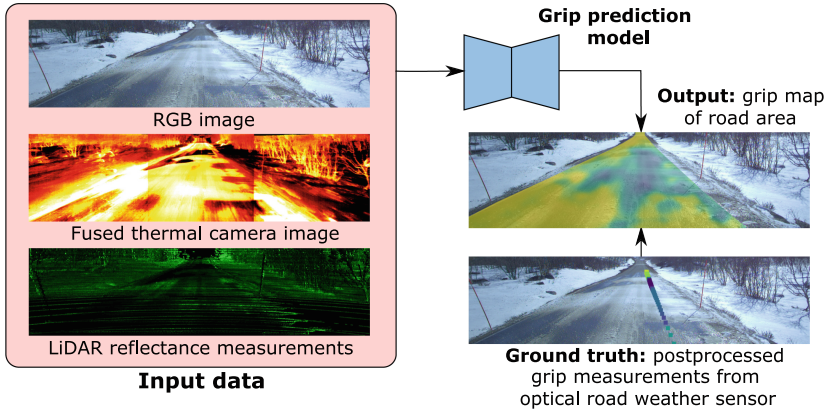


Fig. 1. Our work presents a grip prediction model, which operates on pixelwise fused RGB camera, thermal camera, and LiDAR reflectance measurements and predicts a dense grip map of the road area. The ground truth for training is obtained with an optical road weather sensor that provides road grip measurements which are postprocessed with GNSS trajectories and external calibrations to match the input data.

Transportation, 24 % of weather-related vehicle crashes in the U.S. occur on snowy, slushy, or icy pavement, and 15 % happen during snowfall or sleet each year [20]. Besides low visibility, significant challenges posed by winter conditions include changes in road surface slipperiness. Snowy and icy road surfaces, in particular, can drastically reduce the friction between vehicle wheels and the road compared to dry and wet roads. Therefore, autonomous driving systems must be capable of distinguishing such scenarios, requiring specialized sensing solutions.

Several approaches exist for estimating the grip on the road. However, the greatest shortcoming of most of these methods in the sense of autonomous driving has been the lack of ability to sense the road ahead of the vehicle, thus only allowing the vehicle to react in situations where the slippery conditions have already affected the driving. To enable sensing of the road further ahead, cameras or other forward-facing, longer-range sensors must be deployed.

In addition, the grip can often vary between different sections of the road depending on local snow, ice, and water layer thicknesses. For example, snowy roads with frequent traffic usually have clear tire tracks, which human drivers follow to avoid the snowy areas on the road. Human drivers can also distinguish sudden icy or wet areas on the road allowing them to either avoid these areas or decrease driving speeds accordingly. Therefore producing dense grip predictions would enable autonomous vehicles to react to the sensed conditions in a manner that human drivers can achieve.

Besides traditional vision by RGB cameras, even more accurate grip predictions could be achieved by combining measurements from other long-range sensors. For example, many LiDARs measure the return intensity of the sent infrared laser pulse, which could be used to differentiate ice and water on the

road due to their different optical properties. In addition, a thermal camera might be able to differentiate some road surface layer types, such as snowy and clear areas on the road. These sensors are commonly used in autonomous vehicles, so their positive effect on grip prediction could be easily adopted by the industry.

In this paper, we introduce a pixelwise road surface grip prediction model based on a convolutional neural network (CNN) to generate dense grip map predictions based on fused images from front-facing RGB and thermal cameras as well as LiDAR reflectance measurements. The ground truth grip values were provided by an optical road weather sensor, the data of which were projected on the other sensor data during postprocessing using 3D transformations between different sensors and the postprocessed trajectory of the data collection. The main idea and different sensor modalities are presented in Fig. 1. We collected a 37 h (1538 km) dataset with our autonomous research platform ARVO (Fig. 2) within different adverse weather conditions and preprocessed it for the aims of this study.

This work extends the previous work by Pesonen [19] and provides a new ablation study of the grip prediction accuracy between different input data modalities. The previous approach was also improved with a more consistent training and validation setup and extended testing. The capability of the dense grip map prediction was measured using quantitative error measurements and qualitative analysis for road areas where ground truth measurements could not be obtained. The study shows that the dense grip predictions are improved with the fused RGB, thermal camera, and LiDAR inputs, while the model relying on the sole RGB inputs, already, greatly improves the resolution of any prior camera-based grip prediction methods.

Our contributions to the state of the art are: 1) We developed a novel method to collect and process a dataset with pixelwise matching of multimodal images and sparse road grip measurements. 2) We proposed a model to predict a dense grip map of the road area in diverse weather conditions. 3) We compared the grip prediction accuracies of models using RGB images, thermal images, and LiDAR reflectance measurements as model input modalities both separately and with every combination using multi-encoder-fusion.

We shared a demo of our models in a Gitlab repository to allow readers to test our methods.¹

2 Background

While dense road surface grip map prediction has only been proposed in our earlier work [19], methods for grip prediction have been proposed before using various sensor setups. Road surface grip measurement methods can be roughly divided into non-contact and contact-based measurements, which have been addressed in surveys by Ma et al. [12] and Acosta et al. [1] respectively. Even

¹ <https://gitlab.com/fgi.nls/public/grip-prediction>.

though contact-based grip sensors and models relying on vehicle information, such as wheel rotation speeds, are incapable of producing the required predictions for grip in front of the car, their use has been essential for evaluating the later-developed non-contact methods. According to the survey by Ma et al. [12] the most prominent non-contact-based methods rely on infrared spectroscopy, computer vision, optical polarisation, or radar detection.

Most of the proposed camera-based road surface grip prediction methods have relied on classification of different road surface conditions without providing scalar estimates of the surface grip [18, 22] or by using a two-part process where the classification result is further used to generate a scalar estimate of the road surface grip [4, 10, 24, 28]. Few models have also been suggested for directly generating scalar grip estimates [2, 5, 16]. However, as a common limitation, all of the models rely on either generating a single prediction for the whole input image or for small regions of interest in predefined shapes. In some of the studies the ground truth labeling was generated by expert annotators [22, 28], in one using a portable pendulum tester [4], in one using friction wheel trailer measurements [10], in one with vehicle response [2], and in one with an optical sensor [16].

Models generating pixelwise outputs have become popular in many tasks, such as semantic segmentation and monocular depth estimation. In semantic segmentation, models are trained to classify each pixel of the input image. Solutions proposed for the task, some of which have also found use in many other problems, include U-net [21], FPN [11] and DeepLabV3+ [3].

Monocular depth estimation is a task more similar to the one presented in this paper, as the labels are scalar distance values instead of discrete classes as in the case of segmentation. In addition, the training labels could originate from sparse measurements such as LiDAR readings. Such weak supervision has been applied to monocular depth estimation with sparse labels by Guizilini et al. [6] showing similarity to our grip prediction task due to the comparable sparsity of the ground truth labels.

As both optical road weather sensors and LiDARs use lasers to measure the return intensity of the measured object, the use of LiDARs for road surface condition prediction shows potential. Ruiz-Llata et al. [23] and Shin et al. [26] showed that different road surface conditions can be detected using LiDAR measurements. Sebastian et al. [25] proposed the use of a LiDAR-based CNN for simultaneous road condition and weather classification. While the use of 3D LiDARs was proposed in the studies, their use for dense road surface grip prediction was not investigated in depth.

As noted, the prior literature is concentrated on low-resolution grip predictions using individual sensor inputs. This study aims to fill the gap by introducing both data and methods for generating dense predictions using multimodal input data.

3 Data

In this section, we describe the collection and preprocessing of the data used to train and evaluate the proposed grip prediction methods.

3.1 Dataset Collection

We collected a 37 h or 1538 km dataset of different driving conditions with the sensor setup in our autonomous driving research vehicle ARVO, of which an older version is presented in our earlier study [13]. The dataset includes various driving conditions, such as daytime and nighttime, snowfall, snow-covered roads, slushy conditions, rain, and wet roads. It also contains data from several road types, such as highways, urban roads, and paved and unpaved local roads. The dataset was collected mostly in the capital region of Finland during fall and winter, with a smaller part collected in Western Lapland during spring. The dataset was postprocessed to contain samples at a frequency of 2 fps and after automatic filtering of low-quality data, the dataset had 237 067 samples.

For this project, we used data from a forward-facing RGB camera (Basler MED ace 2.3 MP 164 color), three forward-facing thermal cameras (FLIR ADK, 24° FOV), a roof-mounted 128-beam rotating LiDAR (Velodyne Alpha Prime VLS-128), a GNSS Inertial Navigation System (INS) (Novatel PwrPak7-E1), and a mobile road weather sensor (Vaisala Mobile Detector MD30). An image of the car with highlighted sensor locations is shown in Fig. 2. The left and right thermal cameras were horizontally tilted approximately 23° outwards from the center camera to achieve a combined field of view covering a larger horizontal angle. All sensors were synchronized with GNSS INS triggering signals except the road weather sensor, which was synchronized manually during postprocessing.

We chose these long-range sensor modalities for this research due to their common use in autonomous driving development. We also noted in our preliminary studies that different types of snow can have type-specific features in thermal cameras. LiDAR reflectance measurements (laser pulse return intensity amplitude which is normalized with distance internally by LiDAR sensor) could also provide single-band spectral information on the surface material as the LiDAR sensor used in this study uses 903 nm wavelength lasers, which has



Fig. 2. The research vehicle ARVO used for data collection. The long-range sensors shown in box A are 1. LiDAR, 2. RGB camera in a weatherproof housing and 3. thermal cameras. The road weather sensor is shown in box B.




Fig. 3. The distribution of grip values and road surface states provided by the road weather sensor measurements in the complete unprocessed dataset. We observe that most of the data is collected within dry, wet, or snowy conditions.

different extinction coefficients between water and ice according to the data by Palmer and Williams [17] and the data by Warren and Brandt [27].

The road weather sensor Vaisala Mobile Detector MD30 is an optical sensor that estimates water, ice, and snow layer thicknesses using three laser intensity measurements in different wavelengths. The operating principle of the sensor is not publicly available, but an earlier sensor prototype is presented in a master’s thesis [9]. The sensor uses an internal model for calculating the grip estimate of the road, most likely based on the three surface layer thickness values. Based on earlier studies on optical sensors [14, 15], we assume that an optimistic upper limit of the sensor grip estimate accuracy is 0.1. However, the grip estimate can be more accurate within clear conditions with constant grip. The surface layer thickness and grip estimates are measured with a 40 fps sampling rate. In addition, the sensor provides road surface condition class, road and air temperatures, and other meteorological measurements. In our analysis, we have assumed that the sensor’s grip measurements are sufficiently accurate that it is reasonable to imitate the sensor measurements for a dense grip map, even though the sensor grip values are likely to contain some inaccuracies as the grip between the tires and the road surface is a complex physical phenomenon.

The grip and road surface condition distribution in our dataset is visualized in Fig. 3. We observe that the two most prominent road states are dry and snowy conditions. These conditions have two distinct grip coefficients, which are 0.82 for dry road and 0.35 for snowy road. Additionally, we note that wet roads usually have a grip of less than 0.8, and the smallest grip of 0.1 is observed on icy roads or roads with very thick layers of water.

3.2 Datasplit

The dataset includes data collections from 18 days, many of which shared the same data collection locations. To ensure that the training, validation, and test sets were collected from different locations while maintaining similar weather

condition distributions, we used a geofencing-based approach to choose the validation and test sets from the full dataset. We chose circular areas within approximately one-kilometer intervals from the data collection area so that all samples collected within these areas are included either in the validation or the test set. The rest of the data is included in the training set, except for any positions less than 55 m from any border of the chosen circular areas. This 55-m gap assures that no observations are shared between validation, testing, or training. With this data split, we achieved a qualitatively similar distribution of weather conditions between training, validation, and test sets. Additional qualitative data filtering was also done at this stage. In the end, the training set has 159 801 samples (79.1%), the validation set has 15 343 samples (7.6%) and the test set has 26 783 samples (13.3%).

There is a possibility that some conditions of the input data, such as illumination, would allow the model to fit to these conditions and learn the general grip conditions on specific data collection dates. Therefore, we used three separate data collections, with 16 139 samples in total, as additional test drives to demonstrate the accuracy of the model regardless of this effect.

3.3 Pixelwise Matching of Modalities

To obtain pixelwise pairs of image data and ground truth road weather measurements, we used the following preprocessing approach. We calibrated all cameras intrinsically and extrinsically and measured the 3D locations and orientations of each sensor. Due to the hardware-based synchronization, we also know the time correspondences between each of the sensors. The GNSS trajectory was postprocessed using base-station data to increase its accuracy.

We chose the RGB camera image as the reference frame of the data as it has the highest resolution regarding the front area of the car. The road weather sensor measurements were overlaid on the RGB images with the following procedure: first, we used the postprocessed trajectory and the external transformation between the INS reference frame and the road weather sensor measurement locations to project the road weather sensor measurement positions to a 3D trajectory. This trajectory of the measurements is then transformed to the RGB camera coordinates and projected to the RGB camera image plane. Therefore, we obtained RGB camera images where the road weather measurement points, which were recorded soon after the RGB camera capture time, are overlaid. To improve the data quality, we only included road weather measurement points within 50 m of the cameras and excluded the measurement points behind any obstacles.

The LiDAR point clouds were motion-corrected with the postprocessed trajectory and projected to the RGB camera pixel coordinates. We also accumulated more LiDAR points from the lower part of the three previous scans to include more reflectance measurements from the nearby road area.

The thermal cameras required a more complex pixelwise matching with the RGB camera. Initially, we generated approximate range images from the LiDAR point clouds projected onto the RGB camera. For each RGB camera pixel, we

identified a corresponding 3D point from the range image, projected that point onto a single thermal camera frame, and determined the corresponding thermal value from the thermal camera image. As a result, we obtained thermal camera images projected onto each RGB camera frame. We normalized the pixel values of the left and right thermal cameras to match the scale of the center thermal camera, ensuring a value distribution close to the shared image border.

As the thermal pixel values correspond to the thermal flux in the pixel with a varying scale due to online calibration, the raw thermal values were not considered suitable for this work. Therefore, we normalized the thermal camera pixel values within each frame so that a sample area from the road has a consistent distribution with zero mean and unit variance. In addition, the borders of the thermal camera images had lower values within cold conditions due to the operation of the thermal camera sensor. We alleviated this effect by determining the systematic error distribution for each thermal camera image and subtracting that error to obtain an image with a more homogenous value distribution. The data preprocessing is described in more detail in the preliminary results of our work [19].

An example of the pixelwise matched sensor data can be seen in Fig. 1. The pixelwise matching quality varies between frames and occasionally distant areas or tall objects closer to the camera might appear unaligned between different sensors. We considered this effect negligible for this work, as the road surface is mostly well aligned and the road surface is usually large and homogenous, alleviating any problems that could be caused by the slight unalignment.

4 Methods

In this section, we present our model for the grip prediction, the training setup, and the performance evaluation methods.

4.1 Model

To generate dense predictions of the road surface grip using the multimodal input data, we propose using a convolutional neural network trained with the sparse pixelwise matched road weather measurements as the ground truth labels. Our models are based on Feature Pyramid Network (FPN) [11] which is adapted to predict pixelwise scalar values for regression. The FPN model was chosen as it was shown efficient for the task in our preliminary studies.

We trained our models with every combination of the collected input modalities to measure their effect on grip prediction accuracy. The models utilizing a single input modality are based on the standard FPN implementation which takes an image tensor as the input. However, the multimodal models include separate encoders for each input modality, and their features are concatenated channel-wise within each feature scale before being forwarded to the decoder. We implemented this feature-level fusion approach due to finding occasional lower-quality samples in some input modalities, meaning it was useful for the model to

learn to discard these features in the corresponding situations. For each of the model encoders, we used ResNet-18 [7].

The outputs of the model are the predicted grip and the predicted water, ice, and snow layer thicknesses for each pixel. The grip prediction is the primary task of the model and the prediction of different surface layer thicknesses is used as an auxiliary task to support the learning, as it has been shown to improve the prediction accuracy of the obtained model in our prior experiments [19]. The model architecture and the training scheme are illustrated in Fig. 4.

In most frames, more road weather measurement points were visible further away from the car. These distant points also contain less information as the resolution of the RGB camera and other sensors concerning the road surface was smaller. We alleviated the effect of these distant points by weighting the road weather measurements within each image based on their y -coordinate in the RGB image plane: the weight of each measurement point decreases linearly from the bottom of the image to the estimated horizon level. With this approach, we could approximately balance the prediction accuracy over the whole road area. For validation and testing the pixelwise weights were normalized within each frame so that their mean is one. For the training, the normalization was performed on the unfiltered road weather measurements, which included some overlapping positions, leading to slightly larger weights on average.

The predicted grip and surface layer thickness values are compared to the sparse ground truth values from the postprocessed road weather sensor data

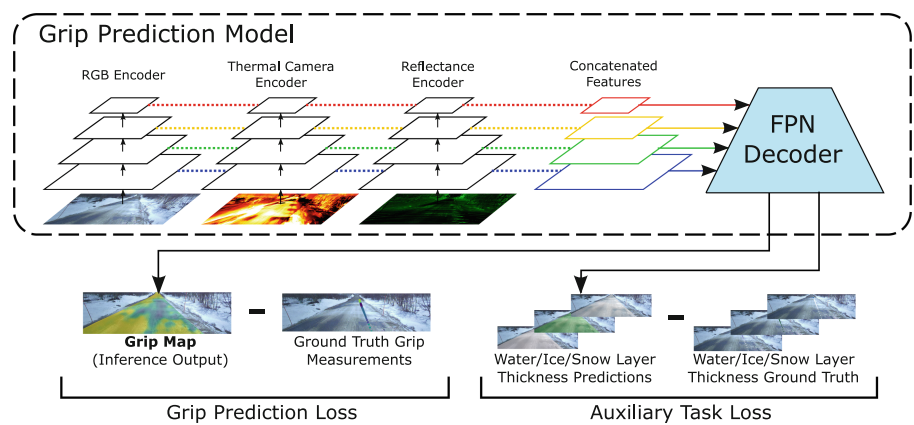


Fig. 4. The model architecture and training scheme for the model using all data modalities. Each input data modality has a separate encoder and their features are concatenated within each feature scale before the FPN decoder. The loss is evaluated both for the grip and the surface layer thickness prediction tasks.

with the following loss function:

$$\mathcal{L}(x, y_p, y_a, w|\theta) = \frac{1}{N} \sum_{i=1}^N w^i (y_p^i - f_p^i(x|\theta))^2 + \frac{\lambda}{3N} \sum_{l=1}^3 \sum_{i=1}^N w^i (y_a^{l,i} - f_a^{l,i}(x|\theta))^2, \quad (1)$$

where x is the input image tensor, N is the number of pixels containing ground truth grip values in the sample, w^i is the weight for pixel i , y_p^i , and $f_p^i(x|\theta)$ are the ground truth and model output for grip value at pixel i and $y_a^{l,i}$ and $f_a^{l,i}(x|\theta)$ are the ground truth and model output for surface layer l value at pixel i . The first term denotes the weighted mean square error for the grip prediction task and the second term denotes the similarly weighted mean square error for the prediction of surface layer thicknesses. The parameter λ is used to adjust the effect of the supportive auxiliary task and in our experiments, it was set to 1.0.

4.2 Training Setup

Each model was trained using the Adam optimizer [8] for 38 epochs with a batch size of 32 and a learning rate of $1e-3$. The FPN model used a dropout rate of 20% in its last layer. The models were compared using the instances which achieved the best validation loss during the training.

As our method was designed to predict the grip and the layer thicknesses based on the surface appearance, we avoided excessive augmentation to maintain accurate predictions. Some augmentations were still used to ensure appropriate generalization. Therefore, it was chosen to apply small random scale and rotation augmentation with a 30% probability, horizontal flip with a 50% probability, small random blur with a 30% probability, and random color jitter to the RGB images with a 30% probability.

4.3 Performance Evaluation

The model performance was evaluated with a root mean square error (RMSE) between the predicted and the ground truth grip values. A similar weighting as in the training loss (1) was applied during validation and test error evaluation as we wanted to measure the grip prediction accuracy balanced over the road area. Due to this weighting, the mean square error was evaluated for each frame separately, and these sample-wise square errors were averaged before evaluating the square root. The test accuracy is reported both for the test set from the main data collection and the three extra test drives with no correspondence to the main dataset.

In addition to the main evaluation metric RMSE, we assessed the test set performance using mean absolute error (MAE), pixelwise unweighted RMSE, and grip weighted RMSE. The pixelwise unweighted RMSE is similar to the main RMSE error, except it omits the weighting in the equation (1). To calculate the grip weighted RMSE, the grip range is divided into ten bins. The weight for each

bin is calculated from the inverse of the number of ground truth measurements in that bin, with all weights scaled to a mean of one. The framewise evaluated mean squared error within each bin is then weighted with the corresponding bin weight, and the square root of the mean of these weighted errors is calculated to obtain the grip weighted RMSE. This process imitates the error for a dataset with a uniform distribution of grip measurements. The grip weighted RMSE is also weighted in a pixelwise manner, similar to the training loss.

However, the error evaluation alone could not show if the grip predictions are valid over the road areas which rarely contain ground truth measurements. Therefore we also performed qualitative analysis on the model output to estimate how well the grip map follows the slipperiness expected by human drivers.

5 Results

In this section, we first analyze the quantitative errors from our validation and test sets and then inspect the qualitative performance of different models.

5.1 Validation and Test Set Errors

We performed the experiments by training the model with different sensor modalities as the input to observe the effect of each sensor on the grip prediction accuracy. The validation set, test set, and separate test drive dataset information and RMSEs achieved with each model are found in Table 1. The additional metrics for the test set are listed in Table 2.

For the validation and test set all obtained errors are significantly smaller than the standard deviation of the dataset, which insists that the models could learn to predict useful grip values. In most experiments, the best or second-to-best results are achieved with the model that uses all data modalities. Using RGB images provides the best accuracy when compared to other data modalities, but the model using only the LiDAR reflectance achieves comparable results with the RGB model. While the combination of RGB and thermal images does not improve performance over using RGB data alone, combining thermal and reflectance data provides similar improvements as the combination of RGB and reflectance. This indicates that the RGB and thermal information may overlap significantly, but also provide information unavailable from LiDAR reflectance alone. Therefore, almost all of the best or second-best results in Table 1 are achieved using some combination of LiDAR reflectance and a higher-resolution image input.

The separate test drive results confirm that the use of several data modalities improves the accuracy and the models have not noticeably overfit to the training data. Even though the standard deviation in each test drive is close to the model errors, it should be noted that the conditions in a single test drive are mostly constant and the models have predicted at least the general conditions in the test drive. However, there is variance and inconsistency in the separate test drive results as the amount of data is small, and adverse effects in one modality could

decrease the performance of a single model. Some differences between the results could also be explained by the specific driving conditions, as the dark conditions in Test drive 1 might benefit the performance of the models using reflectance.

The results of the additional test set metrics in the Table 2 follow the expected previous results with RMSE. As the grip weighted RMSE has larger values than the previous RMSE, we conclude that less prevalent grip values are more difficult to predict than the prevalent values. However, the grip weighted RMSE is below 0.1 which is relatively accurate as grip prediction performance.

We also verify the FPN model choice by comparing the performance of U-net and DeepLabV3+ using only RGB inputs. The results are shown in Table 3. While DeepLabV3+ performed best on the test set and two of the separate test drives, the comparable results confirm that our validation loss-based model choice should not hinder our results and further model optimization can be left as future work along other hyperparameters.

In addition, a scatter plot of the grip and different surface layer thickness predictions in the test set is shown in Fig. 5. The surface layer thickness predictions mostly follow the ground truth values within a relatively small error range while the predicted grip values have a larger error distribution, partly due to misinterpretation of snowy conditions.

5.2 Qualitative Performance

Besides the error evaluation based on the ground truth road weather sensor measurements, we evaluated the grip map prediction over the complete road area qualitatively. Several example scenarios and grip map predictions from the final proposed model in different road weather conditions are shown in Fig. 6. Additionally, examples from the other introduced models and a comparison of

Table 1. Dataset information and grip prediction RMSE for different models on the validation set, test set, and separate test drives. Different data modalities are abbreviated where RGB denotes RGB camera, T denotes thermal camera and R denotes LiDAR reflectance measurements. The best-achieved error in each set is in bold and the second-to-best is underlined.

| | Validation set | Test set | Test drive 1 | Test drive 2 | Test drive 3 |
|--------------------------|----------------|---------------|----------------------|---------------|---------------|
| Weather condition | Varying | Varying | Snowy,snowfall, dark | Snowy | Wet, slushy |
| Grip mean | 0.6474 | 0.659 | 0.399 | 0.557 | 0.649 |
| Grip SD | 0.2037 | 0.201 | 0.104 | 0.140 | 0.159 |
| # samples | 15 343 | 26 783 | 5 746 | 2 042 | 8 351 |
| Modalities | RMSE | RMSE | RMSE | RMSE | RMSE |
| RGB | 0.0657 | 0.0589 | 0.1041 | 0.1497 | 0.1062 |
| T | 0.0794 | 0.0772 | 0.1248 | 0.1670 | 0.1361 |
| R | 0.0677 | 0.0591 | <u>0.0992</u> | 0.1262 | 0.0944 |
| RGB + T | 0.0655 | 0.0605 | 0.1024 | 0.1416 | 0.1069 |
| RGB + R | <u>0.0638</u> | 0.0565 | 0.1038 | 0.1418 | <u>0.0917</u> |
| T + R | 0.0664 | 0.0586 | 0.1056 | 0.1038 | 0.0906 |
| RGB + T + R | 0.0632 | <u>0.0575</u> | 0.0974 | <u>0.1118</u> | 0.0994 |

Table 2. Additional test set metrics of the models using different modalities.

| Modalities | MAE | Pixelwise unweighted RMSE | Grip weighted RMSE |
|------------|---------------|---------------------------|--------------------|
| RGB | 0.0242 | 0.0698 | 0.0912 |
| T | 0.0323 | 0.0875 | 0.0955 |
| R | <u>0.0236</u> | 0.0734 | 0.0947 |
| RGB+T | 0.0248 | 0.0720 | <u>0.0909</u> |
| RGB+R | 0.0248 | <u>0.0680</u> | 0.0939 |
| T+R | 0.0231 | 0.0720 | 0.0974 |
| RGB+T+R | 0.0238 | 0.0674 | 0.0908 |

Table 3. Comparison of different model architectures with only RGB input. Errors in RMSE.

| Model | Validation set | Test set | Test drive 1 | Test drive 2 | Test drive 3 |
|----------------|----------------|---------------|---------------|---------------|---------------|
| FPN (proposed) | 0.0657 | <u>0.0589</u> | 0.1041 | <u>0.1497</u> | <u>0.1062</u> |
| U-net | 0.0692 | 0.0604 | 0.1065 | 0.1466 | 0.1127 |
| DeepLabV3+ | <u>0.0673</u> | 0.0583 | <u>0.1061</u> | 0.1462 | 0.1026 |

the impact of different modalities on the qualitative results are shown in Fig. 7 and in the supplementary material. In all figures in this work, the road area is segmented manually as the model does not differentiate the road area from the input data.

In general, we observe that the model output is smooth and is often constant when there are no variations in road weather conditions, such as when the road is completely dry, completely wet, or completely covered in snow. The model predictions could also mostly follow the boundaries between snowy and clear areas as seen in scenarios presented in Fig. 6 where clear tire tracks can be seen on otherwise snowy roads. Some conditions are still difficult to detect, such as the second scenario on the right column, in which the model could not detect the low grip of an area covered with deep water.

In addition, the model performance is unclear in some conditions that are further from the usual ground truth data locations, such as on the adjacent lane. The model output also could not follow sharp changes in grip values as the model seems to average the grip on relatively large prediction areas. This is likely due to the sparsity and varying data quality of the ground truth road weather measurements.

In Fig. 7 we show performance differences between models using different data modalities as inputs. In some examples, the thermal and reflectance-based single modality models misclassify the grip conditions of the whole scene as the data modality can not differentiate the current condition correctly. However, the model using each data modality seems to combine the correct predictions from the single modalities into a consistent representation of the grip map.

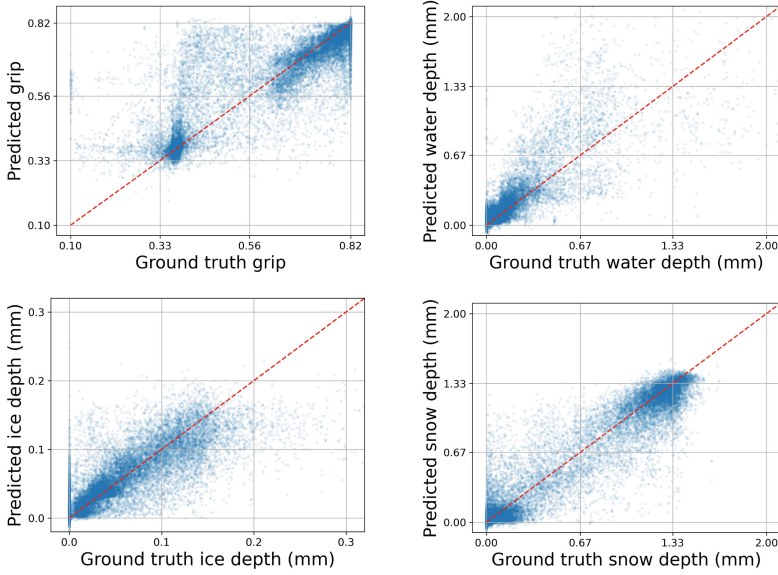


Fig. 5. Scatter plots of predicted grips and layer thicknesses produced by the best, proposed model (RGB+T+R). The x-axis represents the ground truth values and the y-axis the predictions. The plots were generated using 50 000 random measurements and corresponding predictions from the test set. The red dashed line represents the position of correct predictions. (Color figure online)




Fig. 6. Visualisations of the qualitative performance of the final model (RGB+T+R). The ground truth labels are shown using 14-by-14-pixel colored squares drawn on the RGB input image.

6 Discussion

The results support the original hypothesis on the accuracy of the dense grip map and the benefit of additional data modalities besides the RGB camera. Even




Fig. 7. Example grip prediction visualizations for each single-modality model and the RGB+T+R model. Some samples were chosen to highlight the performance differences between models.

though a large portion of the accuracy is obtained by separating the dry and snowy conditions from each other, it seems the model can perform in other road weather conditions as well.

However, one has to consider several error sources, as the optical road weather sensor is not designed to measure all the complex phenomena that could affect the grip between the tires and the road. In addition, the correct synchronization and alignment of the road weather sensor data was challenging. There is also a risk that the data split into training, validation, and test sets could cause some samples in different sets to have too many similarities meaning it's possible that some overfitting could not be observed from the validation and test set results. However, the results on the data from the three extra test drives defend the validation and test set results. In general, one would need an even larger representation of different weather conditions in the dataset to obtain a model with less bias and higher accuracy in several real-life road weather conditions. Despite these limitations, our results show evidence of the performance of our method.

7 Conclusions

This study presents a novel method to predict a dense grip map of the road area from multimodal image data with a convolutional neural network. The models using RGB or 3D LiDAR reflectance measurements provide the best baseline results, whereas the highest accuracy predictions are achieved with sensor fusion using modality-wise encoders. The use of thermal camera images also

shows potential, while their contribution is smaller than that of the RGB and reflectance measurements. The results follow those of earlier studies in proving that the RGB camera is a powerful tool for detecting road surface conditions while also providing major steps in using 3D LiDAR reflectance measurements for dense grip prediction both alone and alongside RGB cameras.

The best model configuration using a combination of all three input modalities achieves an RMSE of 0.0632 and an RMSE of 0.0575 on the diverse validation and test sets respectively. The results from separate test drives also prove the system's usability in unseen conditions. In addition, the qualitative results show the model recognizing various shapes of snow, ice, and water layer distributions affecting the grip prediction. These qualitative results were also improved with the model that uses multimodal inputs with the fusion of encoder features.

To achieve a reliable implementation of this method for autonomous driving, one should collect a large dataset with improved sensor data quality and an even more diverse and balanced set of road and weather conditions. It could also be investigated if one could improve the prediction accuracy by switching the reference image plane from the presented RGB camera frame to another plane, such as the bird's-eye view of the road area or even the 3D frame of the LiDAR. Finally, one should develop methods to predict the uncertainty of the grip prediction output to fuse the output from this method reliably with autonomous driving systems.

Acknowledgements. Funded by the European Union (grant no. 101069576). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency (CINEA). Neither the European Union nor the granting authority can be held responsible for them.

The contribution of the authors is as follows: Maanpää and Pesonen planned the study, performed experiments, and wrote the manuscript. Maanpää also preprocessed the dataset and Pesonen performed preliminary experiments before this study. Manninen, Maanpää and Hyyti developed the research vehicle for data collection with the research group and Maanpää collected the dataset. Hyyti, Melekhov, and Kannala advised in the planning of the study. Kukko and Hyyppä supervised the project.

In addition, we would like to thank Eugeniu Vezeteu for his help in data collection and sensor calibration and Paula Litkey for participating in the vehicle development. We would also like to thank Eero Ahokas for GNSS trajectory processing and Josef Taher for their advice during this work.

References

1. Acosta, M., Kanarachos, S., Blundell, M.: Road friction virtual sensing: a review of estimation techniques with emphasis on low excitation approaches. *Appl. Sci.* **7**(12), 1230 (2017)
2. Cech, J., Hanis, T., Kononisky, A., Rurtle, T., Svancar, J., Twardzik, T.: Self-supervised learning of camera-based drivable surface roughness. In: 2021 IEEE Intelligent Vehicles Symposium (IV), pp. 1319–1325. IEEE (2021)

3. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Proceedings of the European Conference on Computer Vision (ECCV), pp. 833–851. Springer (2018)
4. Du, Y., Liu, C., Song, Y., Li, Y., Shen, Y.: Rapid estimation of road friction for anti-skid autonomous driving. *IEEE Trans. Intell. Transp. Syst.* **21**(6), 2461–2470 (2019)
5. Du, Z., Skar, A., Pettinari, M., Zhu, X.: Pavement friction evaluation based on vehicle dynamics and vision data using a multi-feature fusion network. *Transp. Res. Rec. J. Transp. Res. Board* **2677**(11), 219–236 (2023)
6. Guizilini, V., Li, J., Ambrus, R., Pillai, S., Gaidon, A.: Robust semi-supervised monocular depth estimation with reprojected distances. In: Conference on Robot Learning, vol. 100, pp. 503–512. PMLR (2020)
7. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778. IEEE (2016)
8. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. In: 3rd International Conference on Learning Representations, (ICLR) (2015)
9. Kivi, J.: Algorithmic road state modelling. Master's thesis, Aalto University School of Science (2019)
10. Langstrand, J.P., Randem, H.O., Thunem, H., Hoffmann, M.: Using deep learning to classify road surface conditions and to estimate the coefficient of friction. In: 2023 IEEE Intelligent Vehicles Symposium (IV), pp. 1–8. IEEE (2023)
11. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2117–2125. IEEE (2017)
12. Ma, Y., Wang, M., Feng, Q., He, Z., Tian, M.: Current non-contact road surface condition detection schemes and technical challenges. *Sensors* **22**(24), 9583 (2022)
13. Maanpää, J., Taher, J., Manninen, P., Pakola, L., Melekhov, I., Hyypä, J.: Multimodal end-to-end learning for autonomous steering in adverse road and weather conditions. In: 2020 25th International Conference on Pattern Recognition (ICPR), pp. 699–706. IEEE (2021)
14. Malmivuo, M.: Comparison study of mobile optical friction and temperature meters 2013. Publications by Finnish Transport Agency (2013)
15. Malmivuo, M.: Test of optical MD30 sensors within Mäntsälä contract – winter season 2021–22 (in Finnish). Publications by Finnish Transport Infrastructure Agency (2023)
16. Ojala, R., Seppänen, A.: Enhanced winter road surface condition monitoring with computer vision. arXiv preprint [arXiv:2310.00923](https://arxiv.org/abs/2310.00923) (2023)
17. Palmer, K.F., Williams, D.: Optical properties of water in the near infrared. *J. Opt. Soc. Am.* **64**(8), 1107–1110 (1974)
18. Panhuber, C., Liu, B., Scheickl, O., Wies, R., Isert, C.: Recognition of road surface condition through an on-vehicle camera using multiple classifiers. In: Proceedings of SAE-China Congress 2015: Selected Papers, pp. 267–279. Springer (2016)
19. Pesonen, J.: Pixelwise Road Surface Slipperiness Estimation for Autonomous Driving with Weakly Supervised Learning. Master's thesis, Aalto University School of Science (2023)
20. Road Weather Management Program of U.S. Department of Transportation: Snow and ice (2023). https://ops.fhwa.dot.gov/weather/weather_events/snow_ice.htm. Accessed 8 Feb 2024

21. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 234–241. Springer (2015)
22. Roychowdhury, S., Zhao, M., Wallin, A., Ohlsson, N., Jonasson, M.: Machine learning models for road surface and friction estimation using front-camera images. In: *2018 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE (2018)
23. Ruiz-Llata, M., Rodríguez-Cortina, M., Martín-Mateos, P., Bonilla-Manrique, O.E., López-Fernández, J.R.: Lidar design for road condition measurement ahead of a moving vehicle. In: *2017 IEEE Sensors*, pp. 1–3 (2017)
24. Šabanovič, E., Žuraulis, V., Prentkovskis, O., Skrickij, V.: Identification of road-surface type using deep neural networks for friction coefficient estimation. *Sensors* **20**(3), 612 (2020)
25. Sebastian, G., Vatttem, T., Lukic, L., Bürgy, C., Schumann, T.: Rangeweathernet for lidar-only weather and road condition classification. In: *2021 IEEE Intelligent Vehicles Symposium (IV)*, pp. 777–784. IEEE (2021)
26. Shin, J., Park, H., Kim, T., et al.: Characteristics of laser backscattering intensity to detect frozen and wet surfaces on roads. *J. Sens.* 8973248 (2019)
27. Warren, S.G., Brandt, R.E.: Optical constants of ice from the ultraviolet to the microwave: a revised compilation. *J. Geophys. Res. Atmos.* **113**(D14) (2008)
28. Zhao, T., Guo, P., Wei, Y.: Road friction estimation based on vision for safe autonomous driving. *Mech. Syst. Signal Process.* **208**, 111019 (2024)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

