



Veröffentlichungen der DGK

Ausschuss Geodäsie der Bayerischen Akademie der Wissenschaften

Reihe C

Dissertationen

Heft Nr. 814

Daniel Fitzner

**Estimation of Spatio-Temporal Moving Fields at High
Resolution**

München 2018

Verlag der Bayerischen Akademie der Wissenschaften

ISSN 0065-5325

ISBN 978-3-7696-5226-0

Diese Arbeit ist gleichzeitig veröffentlicht in:

Wissenschaftliche Arbeiten der Fachrichtung Geodäsie und Geoinformatik der Universität Hannover

ISSN 0174-1454, Nr. 338, Hannover 2018



Veröffentlichungen der DGK

Ausschuss Geodäsie der Bayerischen Akademie der Wissenschaften

Reihe C

Dissertationen

Heft Nr. 814

Estimation of Spatio-Temporal Moving Fields at High Resolution

Von der Fakultät für Bauingenieurwesen und Geodäsie

der Gottfried Wilhelm Leibniz Universität Hannover

zur Erlangung des Grades

Doktor-Ingenieur (Dr.-Ing.)

genehmigte Dissertation

Vorgelegt von

Dipl.-Geoinf. Daniel Fitzner

Geboren am 02.09.1980 in Osnabrück

München 2018

Verlag der Bayerischen Akademie der Wissenschaften

ISSN 0065-5325

ISBN 978-3-7696-5226-0

Diese Arbeit ist gleichzeitig veröffentlicht in:

Wissenschaftliche Arbeiten der Fachrichtung Geodäsie und Geoinformatik der Universität Hannover

ISSN 0174-1454, Nr. 338, Hannover 2018

Adresse der DGK:



Ausschuss Geodäsie der Bayerischen Akademie der Wissenschaften (DGK)

Alfons-Goppel-Straße 11 • D – 80 539 München
Telefon +49 – 331 – 288 1685 • Telefax +49 – 331 – 288 1759
E-Mail post@dgk.badw.de • <http://www.dgk.badw.de>

Prüfungskommission:

Vorsitzender: Prof. Dr.-Ing. Ingo Neumann

Referentin: Prof. Dr.-Ing. habil. Monika Sester

Korreferenten: Prof. Dr. Edzer Pebesma (WWU Münster)
Prof. Dr.-Ing. Steffen Schön

Kommissionsmitglied: Prof. Dr.-Ing. Uwe Haberlandt

Tag der mündlichen Prüfung: 16.06.2017

Abstract

Rainfall, hurricanes, tsunamis, eruptions of volcanoes or earthquakes are examples of dynamic environmental phenomena that are monitored by either remote sensors like imaging satellites or weather radar, or in-situ sensors, such as rain gauges, seismic sensors or weather buoys. Often, the goal is the prediction of future states of the phenomena in order to issue early warnings to citizens. With the increasing sensing capabilities of modern sensors, the processing capabilities of modern computers and the advances in communication technology, the data can be collected at high spatial, temporal and data resolutions. In addition, it is often easily and automatically accessible, sometimes in real-time, giving rise to more automated applications that make use of it. However, often, the number of sensors is not sufficient, and the high-resolution data is not available at locations where it is needed. Then, estimation and interpolation methods are required or the installation of new sensors and sensor types, possibly with a different resolution and measurement quality. Further, additional knowledge of the phenomena is sometimes required and beneficial, such as information on atmospheric motion used for precipitation forecasting. These considerations lead to the three topics that are investigated in this thesis.

The first part of the thesis is concerned with the **quantitative estimation of precipitation** from weather radar and rain gauges at 1-min temporal resolution at locations with no rain gauge. Different interpolation methods for estimating 1-min precipitation intensities are compared in *cross-validation* and a specific focus is on spatio-temporal methods that use information on the field motion, estimated from weather radar using optical flow, and on methods for spatio-temporal radar bias correction. The results show that both, motion-based and bias correction methods, deliver superior estimation performance when compared to solely spatial rain gauge interpolation methods or those based on uncorrected weather radar only.

The number and distribution of rain gauges is often not sufficient for certain use cases, for example, the estimation of *areal rainfall* in hydrology. As introducing new rain gauges everywhere they are required is an expensive process, **cars** are investigated **as a potential new data source for precipitation information**. A set of test cars has been equipped with sensors and the sensor readings have been compared with the assumed 'ground truth' rainfall derived from weather radar and rain gauges by the method that has shown best estimation performance previously. As expected, the results show a clear dependency between car speed and the sensor readings. In addition, a positive dependency between the speed-corrected sensor readings and the assumed 'ground truth' could be proven.

The third part deals with the topic of the **estimation of motion of dynamic fields from in-situ sensor data**, such as rain gauges or cars, for example, for use cases where no images of the phenomenon are available. The proposed algorithm is based on a well-known, image-based optical flow algorithm, which is adapted to the specifics of the envisioned use cases, mainly to the irregularity of data and potential distributed processing capabilities of the sensors.

Keywords: Quantitative precipitation estimation, car sensors, motion field estimation

Zusammenfassung

Starkregen, Fluten, Wirbelstürme, Tsunamis oder Vulkanausbrüche sind Beispiele dynamischer Umwelt Ereignisse die durch Fernerkundungssensoren wie Satelliten oder Wetterradar oder in-situ Sensoren wie Regenmessstationen, Seismographen oder Wetterbojen überwacht werden. Die erhobenen Daten speisen Informationssysteme, mit dem Ziel, Nutzern Informationen oder sogar Vorhersagen der Phänomene zu liefern und zeitnah Warnungen, z.B. Unwetterwarnungen, zu publizieren. Mit den zunehmend genaueren Sensoren, den Entwicklungen moderner, immer leistungsfähigerer Computer und den Vorteilen moderner Kommunikationstechnologien werden die Daten immer hochaufgelöster in Raum und Zeit und immer einfacher zugreifbar. Doch nicht immer sind diese Daten auch dort verfügbar wo sie benötigt werden, sodass entweder räumliche oder raum-zeitliche Interpolation vorhandener Daten notwendig wird oder die Installation weiterer, ggf. neuartiger Sensoren. Zudem werden häufig weiterführende Informationen zum Phänomen benötigt, die die Vorhersage erleichtern oder erst ermöglichen, beispielsweise Informationen zur Bewegung der Atmosphäre zur Kurzfristvorhersage von Regenereignissen. Diese Betrachtungen führen zu den drei Themenbereichen der vorliegenden Arbeit.

Der erste Teil befasst sich mit der **quantitativen Regenschätzung** aus Wetterradardaten und Messungen stationärer Regenstationen, an Orten ohne Regenstation in hoher zeitlicher Auflösung. Verschiedene Interpolationsmethoden zur Regenschätzung werden eingeführt und mittels Kreuzvalidierung verglichen. Im Fokus stehen dabei Methoden, die Informationen zur Bewegung des Regenfeldes, hergeleitet mit Methoden der 'Computer Vision' aus Wetterradarbildern, in den Schätzprozess integrieren. Zudem wird eine neuartige Methodik zur Korrektur der raum-zeitlichen Abweichung zwischen Wetterradar- und Stationsmessungen eingeführt. Der Vergleich der Methoden zeigt, dass die Integration von Bewegungsinformationen sowie die Abweichungskorrektur vorteilhaft für die Schätzung der 1-minütlichen Regenintensität an Orten ohne Regenstation ist.

Die aktuelle Verteilung der Regenmessstationen reicht für bestimmte hydrologische Anwendungen nicht aus, beispielsweise die Schätzung des räumlichen Regens. Da die Aufstellung neuer Messstationen ggf. teuer und nicht immer möglich ist, soll in diesem Teil der Arbeit die Nutzung von **Autos als potentielle Regenmessstationen** untersucht werden. Dazu wurden Testfahrzeuge mit Sensoren ausgestattet und die so erhobenen Daten wurden mit den Regenschätzungen derjenigen Methoden als Referenz verglichen, die zuvor die besten Schätzergebnisse zeigten. Die Ergebnisse zeigen einen deutlichen Zusammenhang zwischen Autogeschwindigkeit und Automessungen. Zudem konnte gezeigt werden, dass ein Zusammenhang zwischen geschwindigkeitskorrigierten Sensorwerten und den Referenzen besteht.

Der dritte Teil der vorliegenden Arbeit beschäftigt sich mit der Schätzung der Bewegung raum-zeitlicher Felder, wenn keine Bilder des Phänomens, also beispielsweise aus Wetterradardaten generierte Bilder, vorliegen, und daher die Methoden der 'Computer Vision' nicht ohne Weiteres angewendet werden können. Es wird ein **Algorithmus zur Bewegungsschätzung aus Messungen von in-situ Sensoren** vorgestellt, der einen aus der 'Computer Vision' bekannten Algorithmus des 'optischen Flusses' den Anforderungen irregulär verteilter, dezentral agierender Sensoren anpasst.

Schlagnworte: Quantitative Niederschlagsschätzung, Autosensoren, Bewegungsschätzung

Publications

This thesis is based on the work carried during my PhD studies as part of the RainCars¹ project of the German Research Foundation (DFG), whose support is gratefully acknowledged. It contains material such as text and images previously published in the following publications:

Journal Articles

Fitzner, D.; Sester, M.; Haberlandt, U.; Rabiei, E. Rainfall estimation with a geosensor network of cars - theoretical considerations and first results. *Photogrammetrie, Fernerkundung, Geoinformation*, 2013, 93-103.

Fitzner, D.; Sester, M. Estimation of precipitation fields from 1-minute rain gauge time series - comparison of spatial and spatio-temporal interpolation methods. *International Journal of Geographical Information Science*, 2015, 29, 1668-1693.

Fitzner, D.; Sester, M. Field Motion Estimation with a Geosensor Network. *ISPRS International Journal of Geo-Information*, 2016, 5, 175.

Refereed articles in conference proceedings

Fitzner, D.; Sester, M. Decentralized gradient-based field motion estimation with a wireless sensor network. In *Proceedings of the 5th International Conference on Sensor Networks*, Rome, Italy, 19-21 February 2016; pp. 13-24. (Recipient of the '**Best Student Paper Award**')

¹DFG, SE645/8-2

Contents

1	Introduction	11
1.1	Motivation, research questions and overview of the approach	11
1.1.1	Rainfall estimation at high spatial and temporal resolution	11
1.1.2	Precipitation estimation with cars	14
1.1.3	Motion estimation from in-situ sensor data	14
1.2	Outline of this thesis	15
2	Basics	17
2.1	Precipitation	17
2.1.1	Resolution, accuracy and precision of precipitation measurements	18
2.1.2	In-situ point measurements of precipitation by rain gauges	19
2.1.3	Weather radar	20
2.2	Wireless Sensor Networks	23
2.2.1	Modeling sensor networks	23
2.2.2	Sensor network algorithms and protocols	24
2.3	Statistics	26
2.3.1	Basics and notation	26
2.3.2	Regression	27
2.3.3	Stochastic processes	30
2.3.4	Stochastic filtering and the Kalman filter	31
2.3.5	Geostatistics	34
2.4	Interpolation methods	40
2.4.1	Inverse-Distance-Weighted	40
2.4.2	Ordinary kriging	42
2.4.3	Regression kriging	42
2.4.4	Cross-validation for performance assessment	43
2.5	Optical flow	43
2.5.1	Optical flow intensity conservation	43
2.5.2	Gradient-based optical flow	44
2.5.3	Probabilistic optical flow	45
3	Related Work	47
3.1	Quantitative precipitation estimation from rain gauges, weather radar and other data sources	47
3.1.1	Precipitation estimation with weather radar	47
3.1.2	Precipitation estimation by interpolation of rain gauges measurements	48
3.1.3	Geostatistical merging of radar and rain gauge data	50
3.1.4	Motion-based methods used in nowcasting	50
3.1.5	New data sources for precipitation estimation	51

3.2	Decentralized estimation with geosensor networks	52
3.2.1	Estimation of spatio-temporal field properties with GSN	52
3.2.2	Object-tracking with GSN	53
4	Methodology for precipitation intensity estimation at 1-min resolution from radar and rain gauges	55
4.1	Time-window approach for estimation	56
4.1.1	Estimation of field motion	56
4.1.2	Weather radar upsampling	57
4.1.3	Variogram estimation	58
4.2	Estimation methods	59
4.2.1	Spatial rain gauge interpolation methods	60
4.2.2	Space-time symmetric rain gauge interpolation method	60
4.2.3	Space-time asymmetric rain gauge interpolation methods	60
4.2.4	Radar-rain gauge merging methods	61
4.2.5	Estimation methods solely based on radar	65
4.3	Summary	65
5	Methodology for precipitation intensity estimation with car sensors	67
5.1	Car sensors	67
5.1.1	Wiper Frequency Sensor	67
5.1.2	Xanonex optical sensor	68
5.1.3	Other sensors investigated	69
5.1.4	Experimental setup and preprocessing	69
5.2	Theoretical considerations for the calibration of the W-R relationship in the field .	70
5.3	Dependency between car speed, windscreen angle and sensor readings	71
5.3.1	Manually-operated windscreen wipers	72
5.3.2	Automatically-operated windscreen wipers	72
5.3.3	Xanonex optical sensor	73
5.4	Summary	73
6	Methodology for motion field estimation with a geosensor network	75
6.1	Algorithm overview	76
6.2	Network and field model	78
6.3	Gradient constraint estimation in the network	78
6.3.1	Gradient constraint estimation from irregular data	78
6.3.2	Requirements on node stationarity and sampling synchronicity	80
6.3.3	Estimation of partial derivative error	81
6.3.4	Gradient constraint selection and derivation of gradient constraint error . .	82
6.4	Temporal coherence: Kalman filter for recursive motion estimation	85
6.4.1	Estimation of process noise Q	86
6.4.2	Estimation of measurement noise R	86
6.4.3	Difference to common Kalman filtering problems	87
6.5	Algorithm protocol	87

6.6	Algorithm complexity	89
6.6.1	Communication complexity	89
6.6.2	Load balance	90
6.6.3	Computational complexity of partial derivative estimation	90
6.6.4	Computational complexity of motion estimation	91
6.7	Summary	92
7	Results	93
7.1	Precipitation intensity estimation at 1-min resolution from radar and rain gauges .	93
7.1.1	Study area and data basis	93
7.1.2	Performance assessment via cross-validation	94
7.1.3	Exploratory and visual data analysis	95
7.1.4	Radar estimation and rain gauge cross-validation results	98
7.1.5	Summary	105
7.2	Precipitation intensity estimation with cars	107
7.2.1	Study area and data basis	107
7.2.2	Selection of the reference method	107
7.2.3	Manually-operated windscreen wipers	108
7.2.4	Automatically-operated windscreen wipers	110
7.2.5	Xanonex optical rain sensor	113
7.2.6	Results of experiments on the VW rain track	116
7.2.7	Summary	117
7.3	Field motion estimation with a geosensor network	119
7.3.1	Study Area, sensor network and deployment strategies	119
7.3.2	Error measures	120
7.3.3	Setting the filter parameters	121
7.3.4	Results - simulated field	122
7.3.5	Results - radar field	127
7.3.6	Summary	128
8	Summary and discussion of the research hypotheses	131
8.1	Discussion of research hypotheses 1 and 2: 1-min precipitation intensity estimation	131
8.2	Discussion of research hypothesis 3: precipitation estimation with cars	133
8.3	Discussion of research hypothesis 4: decentralized motion estimation	134
8.4	Outlook	135
9	Appendix	137
9.1	Discussion on the 'frozen field' distance function	137
9.2	Executable Kalman filter equations for the motion estimation algorithm	140
9.3	Controllability and Observability of the Kalman filter for motion estimation	141
	List of Figures	145
	List of Tables	147
	References	149

1 Introduction

The main goal of the work presented in this thesis is the accurate estimation of spatio-temporal moving fields at high temporal and spatial resolutions. This generic topic is investigated in three different parts. While the first two parts deal with precipitation fields, the third part is domain independent and can be applied to other use cases and spatio-temporal fields.

In the first part, the goal is the accurate estimation of precipitation intensities at locations without a rain gauge at 1-min temporal resolution. This is done by combining the available data sources of rain gauges and weather radar using new methods of data analysis and data integration. In the second part, cars are introduced as new sensor types for precipitation. The third part provides a new, decentralized algorithm for the estimation of the motion of a dynamic field by the nodes of a geosensor network.

This chapter provides a discussion on the motivation for the research in each of these parts and an overview of the remaining chapters of this thesis.

1.1 Motivation, research questions and overview of the approach

In this section, motivating use cases for the research are provided for the three topics covered in this thesis as well as the derived research questions that have been guiding the work. Further, an overview of the work carried out is given.

1.1.1 Rainfall estimation at high spatial and temporal resolution

Heavy rainfall and flooding can cause severe damage to humans, the environment and the infrastructure. Rainfall data at high spatial and temporal resolution is important to reduce the impact of flooding events. For example, rainfall data is one of the most important inputs to hydrological models for flood prediction. However, still, it is difficult to obtain high quality rainfall data at the resolutions required for the applications. Traditional rain measurement devices collect rain in can-like devices to be read out by human operators, e.g. once a day. Nowadays, many rain gauges are digitally accessible and provide data at high temporal resolutions. Examples are those operated by the German Weather Service (DWD) that measure automatically and provide a sampling rate of 10-minutes (sometimes 1-min) at a high data resolution of up to 1/100 mm. In Figure 1.1, left, the number of DWD rain gauges starting their operation is displayed from 1990 to 2016 (the end of operation of a particular rain gauge is disregarded).¹ With the new millennium, the number

¹Data provided by the DWD on <https://www.dwd.de/DE/leistungen/klimadatendeutschland/stationsliste.html>, last access 26.10.2016

of rain gauges with automatic, high resolution measurements significantly increased (type MN), leading to approx. 1300-1400 stations in Germany.




Figure 1.1: Left: DWD Sensor types starting operation per year since 1990. RR: Gauges with daily measurements. SY: Gauges with hourly, automatic measurements. MN: Gauges with automatic measurements with (at least) 10-min resolution. Multiple values might correspond to a single rain gauge that has been extended with new sensor types. **Right:** Distribution of automatically-recording rain gauges in lower saxony (111 gauges of type MN). Black dots: Hanover area rain gauges used in this work, providing 1-min aggregations of precipitation with a data resolution of 1/100mm automatically. Black circle: Hanover radar coverage

The rain gauges are complemented by weather radar devices that provide less accurate rainfall estimates but a higher spatial coverage. The coverage of the Hanover weather radar device located at Hanover airport is displayed in Figure 1.1, right. Quantitative estimation of precipitation from both data sources, radar and rain gauges is an active research field. The goal is to combine the higher accuracy of stationary rain gauges with the high spatial coverage of the more erroneous radar estimates. Most of the approaches for merging radar and stationary rain gauges such as the ones from an hydrological background like Haberlandt (2007) or Verworn & Haberlandt (2011) or the RADOLAN product provided by the DWD (Winterrath et al., 2012) consider temporal resolutions of hours, e.g. by calculating hourly sums of precipitation values from combined data provided by both types of devices. Only few methods consider higher temporal resolutions, e.g. 10-minute resolution in the recent study of Berndt et al. (2014). While hydrological models might also benefit from higher temporal resolutions, this work requires 1-min temporal resolution for the derivation of 'ground truth' precipitation intensities for the car data collected in the field. Then, 10-min or even hourly resolution would make any fine grained car data analysis impossible, as the car moves, the precipitation field moves and precipitation rates can significantly change within such periods. When moving from hourly to 10-min or even 1-min aggregations, accurate precipitation rate estimation becomes increasingly difficult (Section 3.1). It is likely that, at 1-min resolution, the integration of the highly accurate but sparsely distributed rain gauges and the more erroneous weather radar becomes even more important. Thus, the first research hypothesis to be investigated in this thesis is:

Research Hypothesis 1: The integration of weather radar and rain gauges is beneficial for accurate 1-min point estimates of precipitation at locations with no rain gauge compared to estimates provided by radar or rain gauges only.

In addition, at high temporal resolutions such as 1-min resolution, the horizontal motion of the precipitation field, called *advection*, gains importance. An example of motion vectors derived from weather radar is provided in Figure 1.2, where a time series of weather radar images is displayed, separated by 20 minutes (left to right) with motion vectors estimated using an image-based Optical Flow (OF) algorithm.




Figure 1.2: Time series of 5-min rain intensities derived from weather radar images (original size 230×230 pixels with a pixel resolution of 1 km^2) separated by 20 minutes from left to right. Motion vectors estimated with an optical flow algorithm.

The visual exploration of animated radar time series gives the impression that the motion information might be very valuable for improving precipitation estimation at locations with no rain gauge. Thus, the second research hypothesis to be investigated is:

Research Hypothesis 2: The integration of motion information improves the quality of point estimates of precipitation at 1-min resolution compared to solely spatial or spatio-temporal interpolation methods that do not include motion information.

For the investigation of these two research hypotheses, the performance of different standard and motion-based deterministic and geostatistical interpolation methods are compared in *cross-validation* and with the estimation performance of radar-rain gauge merging methods. Most of these methods have not been applied to precipitation intensity estimation at 1-min temporal resolution before. Thus, this part of work is mainly empirical as it applies existing methods to new data. In addition, three new estimation methods will be introduced and compared in the experiments. First, the well-known 'frozen field' model and distance function is used as a parameter in an inverse-distance-weighted interpolation, resulting in a new method termed 'frozen field' or FF-IDW (Section 2.4), that has not been described previously in a similar way. Another focus is on the newly introduced method RADAR-DBC for dynamic radar-rain gauge bias correction

(Section 4.2.4). Finally, this method is combined with the 'frozen field' model in a new method termed 'frozen field' regression kriging (FF-RK).

1.1.2 Precipitation estimation with cars

While the previous hypotheses concern problems in integrating weather radar and rain gauges for the derivation of highly accurate point- or areal precipitation estimates, another option is to increase the density of measurement devices, e.g. by using cars as mobile rain gauges. The general idea is that a high number of low-accuracy sensors could beat a low number of high-accuracy sensors (rain gauges) or weather radar for the estimation of areal rainfall, for example, the estimation of the mean precipitation intensity over a certain area to be used in hydrological models. The idea has first been introduced by Haberlandt & Sester (2010) and subsequently been investigated in a research project funded by the German Research Foundation called RainCars.² In the project, the windscreen wiping activities of the cars and the readings of the optical sensors controlling the wipers were investigated as indicators for rainfall. In the laboratory, a positive dependency between the wiper frequency (e.g. windscreen wipes per minute) or optical sensor readings and rainfall measured using standard devices could be established (Rabiei et al., 2013). This work investigates this dependency in the field, which leads to the following research hypothesis:

Research Hypothesis 3: There is a statistical dependence between the wiper frequency and other precipitation-related sensor readings of cars in the field and precipitation intensities provided by standard devices like weather radar or rain gauges.

For investigating this research hypothesis, cars were equipped with sensors for recording wiper activity and the car's precipitation sensor readings, in order to derive a base relationship to rainfall measured by rain gauges or estimated by weather radar, termed **Wiper-Rainfall (W-R) relationship**. In addition, experiments on the rain track of Volkswagen were executed. Technical infrastructures and requirements for car data access e.g. technologies for 'Floating Car' Data are outside the scope of this work. While the present work is concerned with experiments using cars in the field, the benefit of cars for areal rainfall estimation and hydrological applications has been investigated in Rabiei et al. (2016) in computer experiments.

1.1.3 Motion estimation from in-situ sensor data

Information on the motion behavior of environmental phenomena that can be modeled as spatio-temporal fields, such as precipitation or oceanographic fields, is important in a number of applications, for example precipitation forecasting or the analysis of ocean dynamics. Often, motion information is derived from images of the field, such as weather radar images or oceanographic satellite images, e.g. infrared images of the sea surface temperature. However, it is not always

²DFG, SE645/8-2

possible to detect the phenomena by remote sensing techniques. Therefore, images are not always available but only point measurements collected by in-situ sensors such as rain gauges or weather buoys. These considerations lead to the fourth research hypothesis:

Research Hypothesis 4: The motion of dynamic spatio-temporal fields can be estimated from sensor data provided by irregularly distributed in-situ sensors measuring the field.

For the investigation of this hypothesis, a proof-of-concept is provided in the form of an algorithm that is capable of estimating the motion of spatio-temporal dynamic fields by irregularly distributed in-situ sensors. The sensors are assumed to be attached to computing and communication facilities and are therefore considered as the communicating nodes of a geosensor network (GSN). A well-known optical flow algorithm is used as a basis and adjusted to the specifics of GSNs and spatio-temporal fields, such as the irregularity of samples, the strong constraints on communication and computation and the assumed motion constancy over sampling periods.

1.2 Outline of this thesis

Chapter 2 introduces the necessary basics such as principles of precipitation measurements with radar and rain gauges, basic statistical and geostatistical concepts, linear filtering and optical flow as well as core concepts of wireless (resp. geo-) sensor networks. In **Chapter 3**, a compilation of related work is provided, covering the related research areas of quantitative precipitation estimation, precipitation estimation with new data sources and the decentralized estimation of spatio-temporal field properties with GSN.

Chapter 4 describes the methodology for quantitative precipitation estimation at 1-min resolution. **Chapter 5** is concerned with rainfall estimation using cars. The algorithm for the decentralized estimation of field motion is provided in **Chapter 6**.

In **Chapter 7**, experiments are conducted to evaluate the methods developed and described in the three previous chapters. First, the cross-validation results for the estimation of precipitation intensity at 1-min resolution and the results of the car data analysis are provided. Then, simulations illustrate the performance of the motion estimation algorithm introduced in Chapter 6.

A concluding discussion of the research hypotheses, a summary of the main findings and an outlook are provided in **Chapter 8**.

2 Basics

This chapter introduces the basic methodological, statistical and technological concepts underlying the presented work, starting with the basics of precipitation and precipitation measurements in Section 2.1. Then, background knowledge on Wireless Sensor Networks is provided in Section 2.2, which is relevant for the approach of field motion estimation with a geosensor network provided in Section 6. Section 2.3 introduces basic statistical concepts that are relevant at different places throughout the work, such as the geostatistical methods of Section 4 or the linear regression methods for W-R function calibration provided in Section 7.2. In Section 2.4, those interpolation methods that are used in Section 4 are introduced. Finally, Section 2.5 provides an introduction to the gradient-based optical flow method that underlies the work of motion estimation provided in Section 6.

2.1 Precipitation

Precipitation can be distinguished into falling precipitation such as rain or snow, and deposited precipitation, such as fog. Falling precipitation reaching the ground is usually characterized by precipitation depth measured in *mm* over a horizontal area for a certain accumulation interval. The volume of water, e.g. for hydrological applications, can then be derived by considering the size of the area. The most common unit for precipitation measurements is the precipitation depth per time, which is called *precipitation intensity* provided in *mm/time*, mostly in *mm/h*. In this work, *mm/h* is used as the unit of measure as it is most common in hydrological applications and most readers with hydrological background will associate rain events with particular intensities given in that unit. In this work, 1-min aggregations are considered, i.e. the precipitation depth for an accumulation interval of 1 minute. Therefore, the given *precipitation intensity* in *mm/h* is hypothetical in that it gives an indication of the intensity that would occur if the 1-min precipitation depth would last 1 hour.

Liquid precipitation is usually characterized by the two tightly coupled properties of *precipitation intensity* and *drop size*. Table 2.1 displays an intensity classification provided by the German Weather Service.¹

¹http://www.dwd.de/DE/service/lexikon/lexikon_node.html, last access: 09.12.2016

Table 2.1: Classification of liquid precipitation provided by the DWD.

Precipitation type	Characterization and drop size	Intensity (mm/h)
Drizzle	Fine, dense and liquid precipitation with a drop diameter ≥ 0.1 to ≤ 0.5 mm.	light: ≤ 0.2 moderate: > 0.2 and ≤ 0.5 heavy: > 0.5
Rain	Liquid precipitation with drop diameter > 0.5 to ≤ 5 mm with an average diameter of 1 to 2 mm.	light: < 2.5 moderate: ≥ 2.5 and < 10.0 heavy: ≥ 10.0 very heavy: ≥ 50.0

Precipitation events can be roughly classified into *convective* and *stratiform* events (Houze, 1981). Convective events emerge from the vertical movement of warm air masses into colder regions above and are characterized by heavy rainfall intensities, a high variability in space and a short duration in time. Stratiform events usually have a longer duration, a large spatial extend of homogeneous rainfall and show lighter rainfall intensities. They are characterized by low vertical air motion and usually appear in conjunction with areas of low-air pressure.

2.1.1 Resolution, accuracy and precision of precipitation measurements

The term resolution can be applied to different dimensions. First, it refers to either the *temporal resolution*, which denotes the sampling rate at which precipitation data is provided by a measurement device. For rain gauges, the temporal resolution denotes the rate at which the rain gauge reports data, e.g. every minute or every hour. Then, the temporal resolution is usually equal to the aggregation period: if a rain gauge reports data every minute, the data represents aggregations over that particular minute and not some other time interval, e.g. hours (even if the unit is mm/h , see previous section). In case of weather radar, the temporal resolution denotes the rate at which new radar data is generated. The data underlying this work has the temporal resolution of 5 min. In contrast to rain gauges, the radar data is snapshot data and not aggregated over the whole 5 min time period but over shorter periods (Section 2.1.3).

Spatial resolution refers to gridded data and denotes the size in geographic coordinates of a single grid cell, e.g. $1 km^2$ for the weather radar images used in this work. The term *spatial coverage* refers to the spatial extent or area that is covered by the measurement device or measurement device network. In case of weather radar, the spatial coverage is the area under the radar umbrella. In case of rain gauges, the term is used here with a rather fuzzy meaning to denote the area that is covered by the rain gauge network.

The term *measurement resolution* denotes the 'smallest change in a physical variable that causes a variation in the response of a measuring system' (World Meteorological Organization, 2008). For weather radar and rain gauges, the measurement resolution denotes the smallest change in precipitation depth that triggers a change in the response of the rain gauge, e.g. $0.01 mm$.

Accuracy denotes 'the extent to which a measurement agrees with the true value' (World Meteorological Organization, 2008). In contrast, the term *precision* denotes how close repeated mea-

surements of a single quantity are, e.g. repeated measurements of a certain fixed precipitation depth.

2.1.2 In-situ point measurements of precipitation by rain gauges

Most of the point measurement devices for precipitation measurements consist of open cylinders collecting falling precipitation that is led through a funnel to the underlying measurement device. The cylinder and funnel are similar for most of the devices, which then differ in (a) their method of determining precipitation intensity from the collected water, (b) whether they record the data (analog or digital), and (c) their temporal sampling or accumulation rate. The most simple type of rainfall measurement device accumulates incoming rainfall in a can-like container over a certain time period, often one day (called non-recording gauge of type Hellmann), and has to be read out and emptied by a human operator. Another type that has been widely used are *floating gauges* where the water enters a float chamber and the rising water level in the chamber is then transformed into the motion of a pen on a chart. *Tipping bucket* rain gauges use a light metal container that is balanced above a midpoint (fulcrum) and collects water in two equally sized compartments. The seesaw-like construct resides in one of the two positions and water flowing through the funnel then fills the upper compartment. When the compartment reaches a certain weight, corresponding to a certain known volume of rain, the container tips, generates a digital signal, the water flows out of the compartment and the process repeats with the other compartment. Finally, *weighting-recording gauges* estimate precipitation intensity by weighting the incoming amount of water.

Sources of error in gauge-based quantitative estimation of ground precipitation

Rain gauge measurements of precipitation often underestimate the true precipitation due to wind, splashing from the collector or surface wetting and evaporation, with wind being the most influencing factor (Legates, 2000). Due to the dynamic nature of wind, this underestimation is also dynamic and varies with the event. In addition, rain gauges provide point measurements such that interpolation is required to derive precipitation rates at locations with no rain gauge. As precipitation events and especially strong events are often small-scale in nature, interpolation is problematic (Legates, 2000), especially for short aggregation periods. For example, rain clouds with strong rainfall in a convective event are often much smaller than the spacing between neighboring rain gauges. Then, all rain gauges might report zero precipitation intensities although there are strong intensities in open space in between the gauges.

The PLUVIO rain sensor

The PLUVIO Rain Sensor developed by the company OTT is a weighting, digitally recording rain gauge that is often used by the DWD in its official network of rain gauges.² The PLUVIO has a collecting area with the size of 200 cm^2 and provides intensities with a measurement resolution of 0.01 mm at a temporal resolution (accumulation rate) of 1 min with timestamps in the unit MEZ. The maximum reading is the value 99 in the unit 0.01 mm/min , corresponding to $\approx 60\text{ mm/h}$.

²It was confirmed in a personal communication with Ernst Walter from DWD on 14.09.2016 that all rain gauges in the Hanover area, and thus, all rain gauges used in this work, are PLUVIO devices.

The height of the collector above ground is determined by the height of the terrain above sea level (Kuner, 2015). The terrain height is below 500 *m* at each of the gauge locations within the study area. In this case, the height of the collecting opening of the PLUVIO devices is exactly 1 *m* above terrain.

2.1.3 Weather radar

Electromagnetic waves are scattered and partially reflected by particles of condensed water in the atmosphere (called *hydrometeors*). Weather radar makes use of this property by estimating hydrometeors and precipitation rate from reflectivity strength of radar signals emitted and sensed by a terrestrial radar station. Properties of weather radar devices are the wavelength that is used, the beam width, the range and the scan strategy that determines the temporal resolution and horizontal angle of the radar beam. In order to receive high spatial resolutions, small beam widths are beneficial. With increasing distance from the station, the beam width increases and since the error also increases (see next section), the range of radar stations is usually restricted to a maximum of 200 *km* (World Meteorological Organization, 2008). Radar reflectivity is expressed by the radar reflectivity factor Z in the unit mm^6/m^3 that is derived from the incoming signals by the radar equation (Weigl, 2015). The reflectivity factor Z is derived from the reflectivity in the unit dBZ by $Z = 10^{dBZ/10}$ (Bartels et al., 2004). The power-law relationship between the radar reflectivity factor Z and precipitation intensity was initially introduced by Marshall & Palmer (1948) and is termed Z-R relationship where R is the precipitation intensity, usually expressed in mm/h

$$Z = a \times R^b \quad (2.1)$$

where the coefficients a and b determine the relationship and have to be calibrated. Throughout Germany, the DWD operates a network of 17 weather radar stations providing both low-elevation precipitation scans as well as atmospheric scans for 3-dimensional hydrometeors. The DWD radar stations operate in the C-Band radar with a frequency of $\approx 5.6GHz$ at a wavelength of 5*m* (Bartels et al., 2004). The DWD distinguishes between qualitative products that deliver the raw reflectivity values and quantitative products such as RADOLAN (Bartels et al., 2004) that provide precipitation intensities and are derived by post-processing, e.g. by combining radar information with rain gauge data. The precipitation scan that provides the basis for most of the products is run once (one sweep) every 5 minutes at a low elevation following the terrain, a beam width of 1° at a rotation speed of 2*rpm* (elevation of 0.8° with rotation speed of $12^\circ/s$, (Seltmann et al., 2013; Weigl, 2015) and a range of 128*km*. The elevation angle is determined by the orography of the terrain surrounding the radar station and is in between 0.5° to 1.8° (Yen et al., 2005). The DX product used throughout this work delivers data in 255 discrete classes of reflectivity values where every class captures reflectivity values in a range of 0.5*dBZ* with a total range of $-32.5dBZ$ to $+95dBZ$ being recorded (Bartels et al., 2004). The standard DWD parameter set for Equation 2.1, derived from years of drop size spectrum measurements (Weigl, 2015), is $a = 256$ and $b = 1.42$. Often, the Z-R relationship is applied together with a coordinate transformation from the delivered

polar coordinates to a regular grid. Figure 2.1 displays the standard DWD Z-R relationship that is applied in this work.




Figure 2.1: DWD standard Z-R relationship applied in this work.

The radar data used throughout this work has been preprocessed by the Institute of Hydrology and Water Resources Management of the Leibniz University of Hanover. This preprocessing delivers precipitation intensities on a regular grid by applying the standard DWD parameter set and a polar-to-cartesian coordinate transformation. This pre-processing is further described in Section 4.1.2 and also in Berndt et al. (2014).

Potential errors in radar-based quantitative estimation of ground precipitation

Errors in the quantitative estimation of ground precipitation using radar can be classified into those that stem from (a) errors in reflectivity measurement, (b) errors introduced by the Z-R relationship and (c) errors that are associated with effects of estimating ground precipitation from radar reflectivity above ground in the atmosphere (Legates, 2000). All of these factors might result in poor estimates of ground precipitation and are further described in the next paragraphs and displayed in Figure 2.2.




Figure 2.2: Weather radar principle. A weather radar device (red box) emits radar beams (red line) and measures the response reflected by water in the atmosphere. Problem of radar reflection on arbitrary objects (a), reduced radar echoes beyond strong precipitation due to absorption and scatter (b), temporal bias between radar estimates and ground precipitation due to time that the rainfall requires to reach the ground and potential spatial bias between radar estimates and ground precipitation intensities due to advection (c), radar overshoot (d) and increased beam elevation and beam broadening with increasing spatial distance from the radar (e).

Problems due to reflectivity errors: Radar beams are not only reflected by liquid or solid precipitation (ice), but also by other objects in the atmosphere or on the ground, resulting in precipitation estimates where there is no precipitation in reality (Figure 2.2(a)). Further, radar beams can be absorbed by strong precipitation, resulting in reduced radar echoes beyond (Figure 2.2(b)). Further, with increasing distance from the radar, the beam broadens and an equal amount of precipitation might lead to lower reflectivity values far away from the radar (World Meteorological Organization (2008), Figure 2.2(e)). In addition, especially with the increased elevation of the beam above ground with increasing distance from the radar, the possibility of the existence of precipitation below the beam increases (overshooting), Figure 2.2(d), resulting in a general underestimation of precipitation for areas far away from the radar station (Legates, 2000). In addition, the melting process of solid to liquid precipitation heavily increases weather radar reflectivity values, an effect which is known as 'bright band'.

Errors in Z-R relationship: The Z-R relationship depends on drop-size distribution and is also influenced by hail and snow (Legates, 2000; Bartels et al., 2004). Bartels et al. (2004) provide the example of a reflectivity value of 45 *dBZ* due to 'bright band', which leads to a precipitation intensity of 30 *mm/h*, when the standard DWD parameter set is applied (Figure 2.2). Without the effect of 'bright band', the *dBZ* value might have been 35 *dBZ*, leading to a precipitation intensity of 6 *mm/h*. Further, even without the presence of 'bright band', hail or snow, the standard Z-R relationships that are applied by the DWD are usually average relationships derived from years of drop size spectrum measurements (Weigl, 2015). However, such average relationships often underestimate heavy rainfall occurring during convective events and overestimate light rainfall (Legates, 2000). Therefore, in order to receive reliable precipitation estimates, the Z-R relationship has to be adjusted to the current type of rain (see e.g. Alfieri et al. (2010) and also Section 3.1). Further, due to the problem of beam broadening described previously, the Z-R relationship also depends on spatial configuration: the Z-R relationship far away from the radar station might be different from the one close by.

Spatio-temporal bias between radar estimates and ground precipitation: Although precipitation scans are usually performed at the lowest angle possible and more or less follow the terrain (Seltmann et al., 2013), precipitation is still estimated in the atmosphere while the precipitation intensity at the ground is usually of interest. Since precipitation in the atmosphere takes time to reach the ground, the existence of a temporal lag (or bias) between radar and ground precipitation is likely (Figure 2.2(c)). For example, the radar beam elevation angle of 0.8° for the precipitation scan provided by the DWD could result in a beam height of ≥ 1 *km* at the range of 100 *km* from the radar station.³ With a falling speed of 5 *m/s*, it takes more than 3 minutes for rain structures detected by radar in the atmosphere to reach the ground. For a falling speed of 2 *m/s* for small rain drops, this number increases to more than 8 minutes. Especially for short time periods, such as 1-min periods, the temporal lag is important to consider when estimating ground precipitation from weather radar. However, this is also important and considered in the provision of hourly precipitation rates such as those provided by the DWD RADOLAN product. In the

³This is a conservative estimate made based on the model provided by Doviak & Zrnic (1993), displayed in Fig. 2.8, page 23. A precise estimation of beam heights as a function of distance could be made but is not considered highly relevant for this work.

presence of strong horizontal wind or advection, there is the additional problem that precipitation might reach the ground at a different location in space and time of where it has been estimated by radar, resulting in an unsystematic not only temporal, but also spatial bias (Legates (2000), Figure 2.2(c)). It is difficult to provide quantitative estimates of the spatial bias but it can be expected to be in the range of a few kilometers. For example, the DWD RADOLAN product uses the 9 radar pixels in polar coordinates surrounding the rain gauge location to determine the pixel whose hourly sum is closest to the hourly sum of the rain gauge measurements.

2.2 Wireless Sensor Networks

In recent years, **Wireless Sensor Networks** (WSNs), characterized as a set of computing nodes equipped with sensors and radio communication facilities, emerged as a new research field with applications in numerous application areas (Lewis et al., 2004), posing new challenges to algorithm design (Boukerche, 2008). Mostly, WSNs are considered to be *ad-hoc* networks that are self-configuring and do not require an a-priori established communication infrastructure. Vehicular ad-hoc networks (VANETs) are a special case of mobile ad-hoc networks where the mobile nodes are vehicles, e.g. cars (Yousefi et al., 2006). VANETs play an important role in traffic research, for example, to reduce traffic jams and accidents by issuing warning messages to nearby vehicles.

WSN and ad-hoc network research has mainly been concerned with developments in hardware and software, time synchronization, energy efficiency, communication protocols and algorithmic foundations like information routing. The driving force behind WSN research is often the limited *node energy* (Guy, 2006): algorithms are designed in such a way that node energy is saved in order to maximize node operation time. The related field of **geosensor networks** (GSNs) deals with those applications of WSNs where nodes are distributed in geographic space and sense spatial or spatio-temporal phenomena. Up to now, most of the research on WSN and GSN remains largely theoretical and has often not reached the implementation stage (see e.g. Duckham et al. (2015) for a detailed elaboration on reasons for the gap between theory and practice in GSN research). WSN and GSN research strongly relies on sensor network models. The most common basic modeling assumptions for WSNs and GSNs are described in the next section.

2.2.1 Modeling sensor networks

The most common model for a WSN is to represent sensor nodes as vertices of a graph $G(V, E)$, where $V \subset \mathbb{R}^2$ is the set of nodes in the plane and E is the set of edges representing bi-directional communication links (Duckham, 2012; Schulz, 2007; Schmid & Wattenhofer, 2008). In case of volatile and mobile networks, the graph is modeled as a function of time $G(t) = (V(t), E(t))$, where $V(t)$ is the volatile set of sensor nodes and $E(t)$ is the time-dependent set of communication links (Duckham, 2012). A common graph model for G is the unit disk graph (UDG), which assumes radio communication in between nodes without any disruptions caused by obstacles but only exhibiting path loss, which is defined as the decrease in signal strength with increasing spatial distance. In a UDG, a communication link in between two nodes n_i and n_j exists if and only if the Euclidean distance in the plane is less than the maximum communication range, which

is usually denoted with r and often assumed to be the same for all nodes. Nodes can either communicate directly or via intermediate nodes, which is called *multi-hop* communication. The maximum transmission range r that can be assumed in a WSN model depends on the technology to be used in a real deployment, which usually allows communication ranges in between 10^1 to 10^3 meters (Sohraby et al., 2007). Technologies and models for long-range communication in WSNs have been investigated, e.g. up to 10 km in Willis & Kikkert (2005). Figure 2.3 displays visual examples of two different WSN (random and VANET) with different communication distances.




Figure 2.3: Visual examples of wireless sensor network nodes (white dots) and communication connections (lines). Nodes deployed in a random way in the study area (left) and nodes (cars) of a VANET (right).

In reality, sensors can be deployed in space either deterministically or randomly. Often, a completely deterministic placement of the nodes is not possible but nonetheless, a certain control over the random placement can be assumed, e.g. concerning the density (Younis & Akkaya, 2008). Schmid & Wattenhofer (2008) argue that when modeling WSNs, worst-case node distributions are preferable when it comes to evaluating algorithms.

2.2.2 Sensor network algorithms and protocols

Sensor network algorithms either operate *globally* with a central unit having access to the whole graph and collected data. Or, in a *distributed* way where every node implements the algorithm, knows only its own state and data and exchanges messages with neighboring nodes. A *decentralized* system can be characterized as "a special case of a distributed system where no single component knows the entire system state" (Duckham, 2012; Lynch, 1996).

Analyzing node energy demand of algorithms in WSN models is usually focused on *data transmission*, which is assumed to be more costly than computation (Duckham, 2012). Transmission energy consumption analysis is most often done by *message counting* (Duckham, 2012; Schulz, 2007). However, energy consumption by idle listening is a major factor (Sohraby et al., 2007) but is usually taken care of by lower level wireless communication protocols such as IEEE 802.15.4., e.g. by scheduling the communication. Therefore, it is mostly neglected in higher-level algorithm specifications and analyses such as the ones provided by Duckham (2012). Other measures of resource consumption are *load balance*, which answers the question on the distribution of energy consumption among the nodes, and *computational complexity / algorithmic efficiency* in terms of run-time and memory space, as in classical algorithmic analyses (Schmid & Wattenhofer, 2008),

e.g. the number of micro controller instructions and the amount of memory required. For the specification of sensor network algorithms, especially distributed or decentralized algorithms, there is no agreed upon pseudo-code formalism or protocol structure. In this thesis, the protocol structure of Duckham (2012) is used. Protocol 1 shows an example, following the introductory example in Duckham (2012).

Protocol Example	
Protocol header	$\left\{ \begin{array}{ll} \textit{State. trans. Sys.:} & \{ \text{STATE 1, STATE 2} \} \\ \textit{Initialization:} & \text{All nodes in state STATE 1} \\ \textit{Restrictions:} & \text{Graph } G = (V, E) \text{ with nodes } n_i \in V \\ & \text{and communication links } E. \\ \textit{Data:} & \text{Data to be stored and processed by the node} \\ \textit{Parameters:} & \text{Parameters of the algorithm} \end{array} \right.$
Node states and actions	$\left\{ \begin{array}{l} \text{STATE 1} \\ \quad \textbf{broadcast } \langle MESSAGE \rangle \\ \quad \dots \\ \quad \textbf{become STATE 2} \\ \text{STATE 2} \\ \quad \textbf{whenever} \\ \quad \dots \\ \quad \textbf{receiving } \langle MESSAGE \rangle \\ \quad \dots \\ \dots \end{array} \right.$

Protocol 1: Protocol Example after (Duckham, 2012) (modified).

The protocol structure encompasses *restrictions*, *events*, *actions* and *states*. In short, *restrictions* define the "assumptions that are made about the environment in which the algorithm will be operating" (Duckham, 2012). Following Duckham (2012), it mainly encompasses network restrictions on the assumed communication graph. *Events* refer to the external or internal events such as an incoming message or the availability of a new sensor measurement at a node. Events can be considered similar to interrupts in microcontroller code. *Actions* are operations that a node can execute such as sending a message or changing the state. They are atomic in the sense that they can not be interrupted by events. Actions are usually written in bold letters in the protocol, such as **broadcast**. Finally, the definition of *states* allows a node to react differently to similar incoming messages. In the protocol, they are written in Courier font, such as STATE 2.

The protocol header gives generic information on the algorithm such as the state at node initialization and the allowed state transitions. Further, the restrictions in which the algorithm operates and information on the algorithm parameters and stored data is provided. The section on *node states and actions* provides the instructions for each state. There, certain predefined keywords denoting actions together with standard control flow constructs are used. The keywords and associated meanings for actions used throughout this thesis and derived from Duckham (2012) are: **broadcast**, denoting a message transmission action to each neighboring node; **become**, denoting

a state transition action; **whenever**; denoting an event such as a new measurement of the sensors attached to the node and **receiving**, denoting a message receive event. More details on the protocol structure and syntax can be found in Duckham (2012).

2.3 Statistics

For this section, some basic statistical knowledge on random variables, expectations and probabilities is assumed. After an introduction to notational conventions used throughout this work, those statistical concepts are introduced that are most relevant for this work and whose understanding is considered of vital importance for the remaining chapters. The concepts are described mainly from a theoretical point of view concerning the *population statistics*. The formulas for calculating *sample statistics*, which are used later in this work, can be found in any statistical or geostatistical text book.

2.3.1 Basics and notation

Random variables are denoted with capital letters throughout the thesis, e.g. X or Y . If not indicated otherwise, random variables are real-valued, i.e. take values $x \in \mathbb{R}$ and are therefore continuous random variables with *probability density functions* (*p.d.f.*), which are written as $f_X(x)$. Collections of random variables are explicitly referred to as *random vectors* and denoted with arrows, e.g., \vec{X} . The expectation of a random variable is defined as the value that the random variable takes on average and is denoted with $E(X)$ or μ_X . The variance of a random variable X is denoted with σ_X^2 (sometimes also with $var(X)$) and is defined as the expectation of the squared deviation from the mean. A random variable X with Gaussian p.d.f. with mean μ_X and variance σ_X^2 is written as $X \sim N(\mu_X, \sigma_X^2)$. Given two random variables X and Y , the expected value of X given a particular value $y \in Y$ is denoted with $E(X|Y = y)$ or simply $E(X|y)$. Similarly, the conditional p.d.f. of a random variable X given a particular value of random variable Y is written as $f_X(x|Y = y)$ or $f_X(x|y)$.

Given two random variables X and Y , the covariance is a measure of their dependency or joint variability.

$$cov(X, Y) \equiv E[(X - E(X))(Y - E(Y))] \quad (2.2)$$

The correlation is then the covariance normalized by the product of standard deviations.

$$\rho(X, Y) \equiv \frac{cov(X, Y)}{\sigma_X \sigma_Y} \quad (2.3)$$

$\rho(X, Y)$ is also called Pearson's correlation coefficient and is a measure of linear dependency in the interval $[-1, 1]$ where -1 indicates negative correlation, 0 indicates complete linear independence and 1 positive correlation. The calculation of the *sample covariance* or *sample correlation* from realizations $x \in X$ and $y \in Y$ requires that the number of realizations is equal to both random variables and that they can be paired into 2-tuples, e.g. (x_i, y_i) . Often, the realizations are the result of measurements and are paired, if both have been collected at the same time. In this work,

Pearson's correlation coefficient is one of the accuracy measures for the precipitation intensity estimation methods, i.e., correlations between estimates and measurements are computed.

2.3.2 Regression

While $p(X,Y)$ gives an indication for the linear dependency between two random variables, it does not deliver the functional dependency itself and therefore, does not allow prediction such as predictions of values of a random variable Y from realizations of X . The purpose of linear regression is to estimate the coefficients of the assumed linear dependency between a dependent random variable Y and an assumed independent, measured variable or vector X . The assumption underlying linear regression models is that the conditional expectation of Y given a particular (real- or vector-) value $x \in X$ is a linear function of x , a relationship which is called the *population regression line*.

$$E(Y|X = x) = \alpha + \beta x \quad (2.4)$$

When X is a scalar variable ($x \in \mathbb{R}$), this is called a simple linear regression, otherwise, a multiple linear regression. As the relationship of Equation 2.4 is a relationship on the conditional mean, it can not be expected to hold exactly for each pair of realizations of the variables X and Y . Thus, a slightly different relationship has to be assumed including an error term ϵ_i that describes the difference between the conditional mean and the response variable:

$$y_i = \alpha + \beta x_i + \epsilon_i \quad (2.5)$$

The model of Equation 2.5 assumes an error in the dependent variable Y only (it assumes that X is not a random variable). The error can be understood as a representation of quantities that have an effect on the response variable but are not represented as independent variables in the model. Approaches for estimation under the assumption that X is random as well are termed *deming regression* or *error-in-variables* models.

For estimating the coefficients α and β , samples in the form of tuples of predictor and response values are required, e.g. pairs (y_i, x_i) for a simple linear regression with a scalar predictor variable X . The most common way of estimating the coefficients α and β from a set of n tuples is known as ordinary least squares (OLS). It provides estimates a of α and b of β such that the sum of squares of all ϵ_i ($1 \leq i \leq n$, where n is the number of tuples) is minimized. If the errors ϵ_i have expectation zero ($E(\epsilon_i) = 0$), are independent and have equal variance (a property often denoted with $\epsilon_i \sim NID(0, \sigma^2)$, where *NID* means "normally and independently distributed"), the OLS estimator is the "best linear unbiased estimator" or short BLUE (this is known as the *Gauss-Markov theorem*).

Definition 1 (BLUE). A linear estimator $\hat{\theta}$ of an unknown quantity θ is called the "best linear unbiased estimator" (or short **BLUE**), iff $E(\hat{\theta}) = \theta$ ('unbiasedness') and the variance of the error $var(\hat{\theta} - \theta)$ is minimized ('best', sometimes also called 'efficient').

For a linear regression estimator, unbiasedness translates to $E(a) = \alpha$ and $E(b) = \beta$. 'Best' means minimum variance of $a - \alpha$ and $b - \beta$. The formula for calculating an OLS estimate of the regression coefficients of a simple linear regression can be found in any textbook on adjustment theory such

as Niemeier (2002) and is implemented in almost any statistical software package (such as Excel or MatLab) and is therefore not repeated here. The deviations of the measured values y_i from the values predicted by the model (i.e., $\hat{y}_i = a + bx_i$) are then called *residuals* ($r_i = y_i - \hat{y}_i$).

Algorithms for solving non-linear least squares problems are known and described in the literature. Whenever a non-linear function is to be estimated (e.g. for variogram model fitting as described in Section 2.3.5), the Levenberg-Marquardt algorithm (Moré, 1978) is used.

Measures of fit for linear regression models

Different measures for estimating the quality of the regression model and predictions exist. The basic measure that is often used in practice to assess the error of predictions or estimates is the *standard error of the estimate*, here denoted with SE, which can be calculated from sample data using Equation 2.6.

$$SE \equiv \sqrt{\frac{\sum_{i=1}^n r_i^2}{n - m}} \quad (2.6)$$

where n is the number of tuples and m is the number of parameters of the regression model ($m = 2$ for a simple regression). Standard errors can also be calculated for estimates of the regression coefficients. While the SE for the estimates gives a proper measure of the prediction error in the units of the response variable, it does not give any information on the 'goodness' of the model. The *coefficient of determination*, usually denoted with R^2 , is the most common measure for this. For a simple linear regression, the square of Pearson's correlation coefficient (Equation 2.3) is equal to R^2 .

Violations of linear regression assumptions, omitted variables and (multi-) co-linearity

The basic assumption of linear regression models, namely $\epsilon_i \sim NID(0, \sigma^2)$, can be violated with different implications. When the mean of errors is different from zero, there will be a bias in the estimated intercept a , i.e. $E(a) \neq \alpha$. Figure 2.4, (a), displays an example of a true (unknown) *population regression line* (red line) and the estimated regression line (black line) resulting from OLS fitting on the sample data with errors that do not have zero mean.

The unequal variance of the error terms is called **heteroscedasticity** (Figure 2.4, (b)). Heteroscedasticity does not affect the unbiasedness of the estimated coefficients but leads to (a) wrong error measures, such as underestimates of the standard error of the coefficients and (b), estimators that are not longer best or efficient, i.e. there might be other estimators providing smaller variances of the estimated parameters (Sykes, 1993). Therefore, the OLS estimator is no longer BLUE under heteroscedasticity. An extension of OLS under heteroscedasticity in the data is weighted least squares (WLS).




Figure 2.4: **(a) Non-zero mean:** An example displaying the true population regression line (red) and the OLS estimated regression line (black) under errors with non-zero mean. **(b) Heteroscedasticity:** Linear dependency between two variables showing heteroscedasticity, **(c) Omitted variables:** Linear dependency for two sets of data collected e.g. from two different measurement time series and influenced by factors unaccounted for in the model. Colored lines represent the true population regression line for each set, a black line the OLS fit to both sets. **(d) Non-linear dependency** Non-linear dependency between two variables and a wrong linear model resulting in correlation of regression residuals with predictor variable.

Often, the samples (tuples of predictor and response variables) for regression calculation are collected from time series or spatial data. If there are influencing quantities that are not represented in the model as predictor variables ('omitted variables'), and if these quantities show temporal or spatial autocorrelation, the errors and thus, the residuals after model fitting will be temporally or spatially correlated. Figure 2.4, (c) shows the extreme case of two different subsets or groups of data, e.g. collected from two different time series or at two different spatial locations, where the intercept α in the true population regression line differs for each of the two subsets. After OLS fitting, the residuals appear as autocorrelated residuals in the analyses. For example, for detecting such cases when the scatter plot is not as obvious as in the example, scatter plots of residuals versus time or residuals versus locations might help.

Another source of correlated errors and thus, correlated residuals after model fitting is a wrong functional form of the model, i.e. a violation of the basic linear regression assumption of Equation 2.4. Then, the residuals are correlated with the predictor variables, for example, for a certain range of predictor variables, the residuals might be positive while for other they are negative (Figure 2.4, (d), displays an example). As shown in the figure, the scatter plot gives a good indication of this wrong model assumption.

In general, correlated errors do not result in biased estimates but in underestimates of the standard error of the estimated coefficients and in estimators that are no longer efficient or best. An extension of OLS to cope with correlated errors is the method of generalized least squares (GLS) that is still BLUE under the assumptions of heteroscedasticity and correlated errors. However, methods such as GLS should not be applied without investigations of the reasons of correlation.

The occurrence of correlation among two or more predictor variables in a multiple linear regression is called (multi-) co-linearity. This does not affect the optimality of the results but its reliability, when reliability can be characterized as: "Results are reliable when estimates obtained are not sensitive to slight variations in the data" (Watson, 1983). There are several ways and tests to detect (multi-) co-linearity described in the literature. Often, scatter plots between pairs of predictor variables are useful for detecting linear or non-linear dependencies between predictor variables, for

example, for the detection of the dependency between car speed and wiper frequency explained later in Chapter 5.

2.3.3 Stochastic processes

A *time series* is a sequence of observations ordered in time, e.g. a sequence of vehicle positions or rain gauge measurements. Often, each single value of a time series at a particular time period t is considered a realization of a random variable. In this work, discrete time periods are considered and therefore, t is a time index: $t \in T \subseteq \mathbb{N}$. Such time series of random variables are often called *random* or *stochastic processes*:

Definition 2 (Stochastic process). A time-discrete **stochastic process** is a collection or set of random variables X_t or random vectors \vec{X}_t where $t \in T$ is a time index such as integer-values of the set $T = 1, \dots, 356$ representing days of the year.

Particular realizations of the random variable X_t , denoted as $x_t \in X_t$, are often called *states* of the stochastic process. In the following, real-valued random variables are considered. However, the definitions extend to random vectors. An important property of stochastic processes is the Markov property:

Definition 3 (Markov process). A stochastic process is called a **Markov process**, if the *p.d.f.* of each random variable X_t of the process for all $t \in T$ conditioned on all previous realizations of it is equal to the *p.d.f.* conditioned on the previous realization only: $f_X(x_t|x_0, \dots, x_{t-1}) = f_X(x_t|x_{t-1})$.

In other words, the process is a *Markov process*, if the current state depends on the previous state only (all previous states of the process do not give more information on the current state than the previous state). Markov processes play an important role in stochastic filtering theory (Section 2.3.4).

Given a single stochastic process, the term *auto-covariance* refers to the covariance of two random variables of that single process. Then, the covariance between any two random variables of the stochastic process X , say X_t and X_s ($cov_X(X_t, X_s)$) is often written as a function of the two time periods: $cov_X(t, s)$, with $s, t \in T$. In practice and with only single realizations of X_t or X_s available, empirical estimates of $cov_X(t, s)$ can not be computed and certain simplifications are necessary. One of these simplifications that might hold or not hold in practice is the property of *weak stationarity* of the stochastic process:

Definition 4 (Weak Stationarity). A stochastic process is said to exhibit **weak stationarity** if (a) the mean of the random variables is independent of time, i.e. $E(X_t) = \mu_X$, $\forall t \in T$, and (b) the covariance depends on the time difference only, i.e. $cov_X(t, s) = cov_X(t + \tau, s + \tau)$, $\forall t, s, \tau \in T$ and (c) it has finite variance.

Under weak stationarity, the *auto-covariance* becomes solely a function of time lag and is defined as:

$$cov_X(\tau) \equiv E[(X_t - \mu_x)(X_{t+\tau} - \mu_x)] \quad (2.7)$$

Autocorrelation is derived from Equation 2.7 similar to the derivation of correlation from covariance as introduced in Equation 2.3. *Cross-covariance* refers to the covariance of two stochastic processes X_t and Y_t . If the assumption of weak stationarity applies to both, the cross-covariance is defined as:

$$cov_{XY}(\tau) \equiv E[(X_t - \mu_x)(Y_{t+\tau} - \mu_Y)] \quad (2.8)$$

The *empirical cross-correlation* for a particular time lag τ , which is often used in this work, mainly for the calculation of cross-correlations between radar and rain gauge time series, can then be calculated as displayed in Equation 2.9.

$$corr_{XY}(\tau) = \frac{1}{T-1} \sum_{t=1}^T \frac{(x_t - \mu_X)(y_{t+\tau} - \mu_Y)}{\sigma_X \sigma_Y} \quad (2.9)$$

where T is the number of time steps of the time series, x_t (y_t) is the single realization of random variable X_t (Y_t) at time step t and μ_X (μ_Y) and σ_X (σ_Y) are mean and (sample) standard deviation, resp., calculated in the usual way.

2.3.4 Stochastic filtering and the Kalman filter

Stochastic filtering and smoothing addresses the problem of estimating a quantity of interest at a particular time, e.g. room temperature or vehicle position, from potentially inaccurate measurements of the quantity or measurements of related properties. Mathematically, the quantity of interest is considered a realization of an unobservable stochastic process X_t that is to be estimated from measurements that are considered realizations of another stochastic process Z_t that is observable (the realizations z_0, z_1, \dots, z_n are known for a particular time period), and which is somehow related to X_t , e.g. via a *conditional probability density function*. The unknown quantities x_0, x_1, \dots, x_n to be estimated, e.g. the object position, are often called *hidden states*.

In filtering approaches, the goal is to estimate the state x_t at time t and its p.d.f. from measurements up to and including time step t : z_0, z_1, \dots, z_t . This dependency is often written as the p.d.f. of x_t conditioned on the current and all previous measurements: $f_X(x_t|z_0, z_1, \dots, z_t)$ (Maybeck (1979) describes this in detail with an easy-to-follow example).

Instead of computing the conditional p.d.f. each time a new measurement arrives, filtering theory assumes a so-called **dynamic model** where the p.d.f. of the current state of the process x_t is conditioned on the previous state only: $f_X(x_t|x_{t-1})$ (for more details and justifications, see Särkkä (2013), p. 30ff). Therefore, filtering assumes that the stochastic process of interest is a Markov process (Definition 3). Filtering usually requires the specification of an initial p.d.f. for x_0 at time step 0, i.e. a **prior probability distribution**. In addition, filtering theory introduces a **measurement model** in the form of a p.d.f. of the measurement at time step t , conditioned on the current state: $f_Z(z_t|x_t)$ (introductions to the probabilistic and Bayesian origins of filtering theory can be found in Welch & Bishop (2006); Särkkä (2013) or Maybeck (1979)). Together with

an initial distribution of x_0 at the initial time step, this forms a so-called **probabilistic state space model** (Särkkä, 2013).

$$\begin{aligned} x_0 &\sim f_X(x_0) \\ x_t &\sim f_X(x_t|x_{t-1}) \\ z_t &\sim f_Z(z_t|x_t) \end{aligned} \tag{2.10}$$

The different filtering approaches then differ in their approach to computing the conditional probabilities. The most widely used filter implementation is the Kalman filter.

Kalman Filter

The Kalman filter first introduced by Kalman (1960) is a special case of a stochastic filter. It is time-discrete, assumes a linear dynamic model, a linear measurement model and Gaussian white noise for both, prediction by the dynamic model and measurement. When the Kalman filter assumptions are met by the application, the filter equations minimizes the error variance for each parameter or hidden state variable.

Specifying the initial p.d.f. of the hidden state at time zero is done by setting the state to a specific value or vector, e.g. the assumed position of the moving object at the time, the filter is initialized. In addition, an initial distribution has to be provided in the form of a prior variance or covariance matrix of the state, denoted with P_0 . Typically, the Kalman filter then executes two alternating steps: the *prediction* of the state to the next time step by the dynamic model, and the *measurement update* by the measurement model.

Dynamic model: In a time discrete Kalman filter, the current state of the process x_t at time t is assumed to evolve linearly from the previous state at time $t - 1$, involving mean-zero Gaussian noise w_t with a certain covariance Q_t , i.e. $w_t \in N(0, Q_t)$, that accounts for influencing factors not modeled by the dynamic linear model (Equation 2.11). The covariance Q_t is a parameter of the Kalman filter algorithm.

$$x_t = F_{t-1}x_{t-1} + w_{t-1} \tag{2.11}$$

where F_{t-1} is the linear function (called transition matrix) that transforms state x_{t-1} to current state x_t . The linearity of the dynamic model and the assumption of Gaussian noise ensures that the filter propagates the mean and variance of the process to the next time step (Maybeck, 1979), i.e. if x_{t-1} has a Gaussian p.d.f., x_t will have a Gaussian p.d.f. This is derived from the property that a linear transformation of a Gaussian random variable is again a Gaussian random variable. Similarly, the sum of two Gaussian random variables is again a Gaussian random variable. In most filter applications, the transition matrix F_{t-1} and the noise are time-independent and the time parameter is skipped. The Kalman filter equations also propagate the covariance matrix of the previous state, i.e. P_{t-1} , to the current state. The noise co-variance Q_t is a parameter in Kalman filtering and has a strong influence on filter behaviour and performance.

Measurement model The Kalman measurement model describes, how the state and the measurements are related. Again, it is assumed that the relationship is linear and that there is some measurement error involved:

$$z_t = H_t x_t + v_t \tag{2.12}$$

where H_k is the so-called *measurement matrix* that relates the current state to the observation (measurement) and v_k is zero-mean Gaussian noise $v \in N(0, R_k)$ with covariance matrix R_k , which is a parameter of the Kalman filter. Again, the measurement matrix and the noise are often time-independent and then, the time parameter is skipped.

Equations 2.11 and 2.12 describe the Kalman filter assumptions of state evolution in time and the relationship between hidden states and measurements. The Kalman filter algorithm then minimizes the trace of the state covariance matrix P , which corresponds to the sum of the variances of each state element. This means, states are computed in a way that the sum of variances of all state elements is minimized. The equations are described in any textbook or tutorial on filtering theory, e.g. Welch & Bishop (2006), and are therefore not repeated here.

Kalman filter tuning

The behaviour of the Kalman filter is strongly dependent on the relationship between the two parameters prediction noise Q and measurement noise R . For the extreme case that the prediction noise is zero or the zero-matrix, the state of the filter converges over time to a constant value or vector that is no longer influenced by any measurements. Vice versa, if the measurement noise is zero, the measurements are assumed to be highly accurate and have a strong influence on the state, which then follows the measurements closely (Maybeck, 1979).

Kalman filter for linear regression models

The Kalman filter can also be used to estimate the coefficients of a linear regression model (Section 2.3.2), when the calibration tuples are derived from time series (see Watson (1983) or Särkkä (2013) for details). Then, the coefficients α and β are represented by the state of the process, again denoted with x_t , which is a vector of dimension n , where n represents the number of parameters of the model (for simple linear regression including an intercept, $n = 2$). When the parameters are assumed to be constant in time, the transition matrix F is the identity matrix I and the process noise is assumed to be zero. The Kalman filter dynamic model becomes:

$$x_t = x_{t-1} \quad (2.13)$$

Then, the linear regression model of Equation 2.5 has a direct expression in the Kalman measurement equation. For example, for a simple linear regression from a scalar real-valued predictor variable X to a scalar real-valued response variable Y with scalar parameters α and β and calibration tuples (y_t, x_t) , the Kalman filter measurement equation becomes:

$$y_t = \begin{bmatrix} 1 & x_t \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} + \epsilon_t \quad (2.14)$$

where ϵ_t is the measurement error. It is immediately obvious that the measurement matrix $H_t = \begin{bmatrix} 1 & x_t \end{bmatrix}$ is time dependent as it contains the predictor variables, which are assumed to be non-random, deterministic variables. The relationship between the computation of the regression coefficients with the Kalman filter equations and the computation using OLS is extensively discussed in Watson (1983) or Särkkä (2013).

2.3.5 Geostatistics

In geostatistics, another dimension is added as a parameter to the random variables that are now collections of random variables in space or space-time and usually called *stochastic fields*. In the remainder of this section, the common geostatistical notation is used representing a spatial or spatio-temporal location not as a subscript but a parameter.

Definition 5 (Stochastic field). A **stochastic spatial field** is a function $Z(s)$ from the domain of spatial coordinates $s = (x, y) \in \mathbb{R}^2$ to the domain of random variables. A time-discrete **stochastic spatio-temporal field** is a function $Z(s, t)$ from the domain of spatio-temporal coordinates to the domain of random variables, where $t \in \mathbb{N}$ is a temporal 'coordinate' (time index).

Individual realizations of a stochastic spatial or spatio-temporal field are written as $z(s)$ or $z(s, t)$, in the following. Sometimes, spatio-temporal locations (s, t) are abbreviated as $u = (s, t)$. Often, the goal is the estimation of realizations of the random field at locations and time steps where no realization is available. For estimation, information on the covariance of the random variables is employed, in the following denoted with $cov(Z(s_i, t_i), Z(s_j, t_j))$ for a pair of spatio-temporal random variables $Z(s_i, t_i)$ and $Z(s_j, t_j)$ at locations $s_i = (x_i, y_i)$ and $s_j = (x_j, y_j)$ and times t_i and t_j , resp., or short $cov(Z(u_i), Z(u_j))$.

In practical applications it is often difficult to estimate the empirical covariance between arbitrary random variables and therefore, certain simplifications are necessary that make estimation possible. Such simplifications concern the dependency of the covariances in relation to the spatial and temporal difference in locations. In the following, some of these simplifications and characterizations of spatial and spatio-temporal covariances are provided, following the book of Cressie & Wikle (2011), pp. 34, 307, ff., and the implications for precipitation fields are discussed.

Properties of spatial and spatio-temporal covariances

If it is assumed that the covariance between two spatial random variables $cov(Z(s_i), Z(s_j))$ does not depend on locations but only on the difference vector, it is called *stationary* and can be written as a function of spatial differences as displayed in Equation 2.15.

$$cov(Z(s_i), Z(s_j)) = C_s(s_i - s_j) \quad (2.15)$$

where C_s is a function on \mathbb{R}^2 . If, in addition, the covariance is assumed to be direction-independent, it is called *isotropic* and is usually written as a function of Euclidean distance, as displayed in Equation 2.16.

$$C_s(s_i - s_j) = C_s(\|s_i - s_j\|) = C_s(\|h_{ij}\|) \quad (2.16)$$

where $\|h_{ij}\| = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}$ is the Euclidean distance between the two locations. Analogous to the spatial case, a stationary spatio-temporal covariance function is defined as a two-parameter function of spatial distance and temporal difference as displayed in Equation 2.17.

$$cov(Z(s_i, t_i), Z(s_j, t_j)) = C_{st}(s_i - s_j, t_i - t_j) = C_{st}(h_{ij}, \Delta t_{ij}) \quad (2.17)$$

with $s_i = (x_i, y_i)$, $s_j = (x_j, y_j)$ and $h_{ij} = (x_i - x_j, y_i - y_j)$ are as above and $\Delta t_{ij} = t_i - t_j$ and thus, C_{st} is a function on $\mathbb{R}^2 \times \mathbb{N}$. Analogous to the spatial case, spatial isotropy in a spatio-temporal setting can be defined (Cressie & Wikle, 2011). Finally, the assumption of a *fully symmetric* spatio-temporal covariance is defined as displayed in Equation 2.18.

$$\text{cov}(Z(s_i, t_i), Z(s_j, t_j)) = \text{cov}(Z(s_i, t_j), Z(s_j, t_i)) \quad (2.18)$$

This symmetry obviously leads to the loss of any spatio-temporal dynamics as, e.g., induced by wind. Fully symmetric spatio-temporal covariances can be established in different ways: one option is to estimate spatial and temporal covariances separately and then combine both using either a product, product-sum or sum model, as done by Cesare et al. (2001). Then, the estimation of the spatio-temporal model is straightforward.

$$C_{st}(h, t) = C_s(h) \times C_t(t) \quad (2.19)$$

It is known that combinations (sums, products, etc.) of positive definite (conditional negative definiteness) are positive definite (conditional negative definite) (Gneiting et al., 2006). Thus, the model Equation 2.19 of is a valid covariance model, if $C_s(h)$ and $C_t(t)$ are valid models (Gneiting et al., 2006).

In order to establish non-symmetric spatio-temporal covariance functions, a relationship known as the *Taylor hypothesis*, originating from Taylor (1938), is useful (Equation 2.20, following the definitions of Gneiting et al. (2006)).

$$C_{st}(0, \Delta t) = C_{st}(V_f \Delta t, 0) \quad (2.20)$$

where C_{st} is a stationary spatio-temporal covariance function and V_f is the field motion vector expressed in the units of spatial and temporal separation. Equation 2.20 relates spatial and temporal covariance using known field motion. A particular covariance model satisfying Equation 2.20 is the so-called '*frozen field*' model described by Gneiting et al. (2006) or Cressie & Wikle (2011), shown in Equation 2.21.

$$C_{ff}(u_i, u_j) \equiv C_{st}(h_{ij}, \Delta t_{ij}) = C_s(h_{ij} - V_f \Delta t_{ij}) \quad (2.21)$$

where C_s is a stationary spatial covariance function. The model satisfies the Taylor hypothesis since $C_{st}(0, \Delta t_{ij}) = C_s(-V_f \Delta t_{ij}) = C_s(V_f \Delta t_{ij}) = C_{st}(V_f \Delta t_{ij}, 0)$. Equation 2.21 is a straightforward way of reducing a non-symmetric spatio-temporal covariance function to a stationary spatial covariance in the case that field motion information is available. The '*frozen field*' relationship can be best explained by the means of a simple example of samples at a fixed position of a completely constant field that moves forward along a constant motion vector in one dimensional space (Figure 2.5).

The image shows that a sample taken at a certain spatial location (black dots) at a previous time step is relocated (blue dots) to another spatial location at the current time step by the field motion. Thus, the temporal covariance for a certain temporal lag Δt at a fixed location is equal to the spatial covariance for spatial distance $V_f \Delta t$ for zero temporal lag. A particular advantage of Equation 2.21 is the ease of estimation and model fitting: when the motion vector is available, only




Figure 2.5: Example of a 'frozen field' satisfying Taylor's hypothesis in one-dimensional space. A 'frozen field' (red curves) moves along a constant motion vector (here: scalar) in one-dimensional space from left to right over a fixed spatial location sampling the field (black dots). Past samples at the location relocate to the current time step according to the motion vector.

the estimation of a spatial covariance function is required. Equation 2.21 is called the 'frozen field' model since it assumes that the field does not develop but a change in field values solely comes from constant motion. This means, the covariance in between any two spatio-temporal locations does not change (decrease) with increasing temporal difference. When isotropy is assumed for the employed spatial covariance, then Equation 2.22 also holds.

$$C_{ff}(u_i, u_j) = C_s(\|h_{ij} - V_f \Delta t_{ij}\|) \quad (2.22)$$

In order to exemplify the behavior of the 'frozen field' distance function $d_{ff}(u_i, u_j) = \|h_{ij} - V_f \Delta t_{ij}\|$, it can be reformulated as follows:

$$\begin{aligned} d_{ff}(u_i, u_j) &= \left\| \begin{bmatrix} \Delta x_{ij} \\ \Delta y_{ij} \end{bmatrix} - V_f \Delta t_{ij} \right\| = \left\| \begin{bmatrix} x_i - x_j \\ y_i - y_j \end{bmatrix} - \begin{bmatrix} v_X \Delta t_{ij} \\ v_Y \Delta t_{ij} \end{bmatrix} \right\| = \left\| \begin{bmatrix} x_i - x_j - v_X \Delta t_{ij} \\ y_i - y_j - v_Y \Delta t_{ij} \end{bmatrix} \right\| \quad (2.23) \\ &= \sqrt{(x_i - x_j - v_X \Delta t_{ij})^2 + (y_i - y_j - v_Y \Delta t_{ij})^2} \\ &= \sqrt{(x_i - (x_j + v_X \Delta t_{ij}))^2 + (y_i - (y_j + v_Y \Delta t_{ij}))^2} \\ &= \left\| \begin{bmatrix} x_i - (x_j + v_X \Delta t_{ij}) \\ y_i - (y_j + v_Y \Delta t_{ij}) \end{bmatrix} \right\| = \|s_i - (s_j + V_f \Delta t_{ij})\| = d_{eucl}(s_i, s_j + V_f \Delta t_{ij}) \end{aligned}$$

where $V_f = [v_X, v_Y]^T$ is the field motion vector in spatial directions X and Y and d_{eucl} is the Euclidean distance. Thus, the 'frozen field' distance function calculates the distance between any two spatio-temporal locations u_i and u_j as the distance between the spatial location s_i of u_i and the spatial-location s_j of u_j moved to the temporal 'snapshot' t_i by the motion vector.

d_{ff} and thus, any covariance function building on it is symmetric in the sense that $d_{ff}(u_i, u_j) = d_{ff}(u_j, u_i)$. However, it is space-time asymmetric since there is indeed a difference between '30 km east, 30 minutes earlier' and '30 km east, 30 minutes later'. For one case, the product $V_f \Delta t_{ij}$ is

subtracted, for the other it is added to the spatial difference vector $[\Delta x_{ij}, \Delta y_{ij}]^T$. There are now two cases where $d_{ff}(u_i, u_j) = 0$:

1. for $u_i = u_j$ and
2. for $\begin{bmatrix} \Delta x_{ij} \\ \Delta y_{ij} \end{bmatrix} = V_f \Delta t_{ij}$. This is the case, when estimation and sample location are exactly separated by the field motion vector V_f .

It is obvious that even when V_f is constant, there can be certain spatio-temporal configurations where both cases can occur at the same time for some i, j, k with $u_i = u_j$ and $u_i \neq u_k$. This is the case when there is a sample at spatio-temporal estimation location ($u_i = u_j$ for some $1 \leq j \leq n$) and there is another sample u_k with $u_i \neq u_k$ exactly separated by u_i by the field motion vector V_f .

Equation 2.22 assumes a constant motion vector. When C_{ff} is used for estimation, its validity is of vital importance and varying motion vectors (i.e. motion vector fields) might result in invalid covariances, since, depending on the context, i.e. the pair of spatio-temporal locations u_i and u_j , a single sample is moved differently. In that case, the kriging matrices (Section 2.4.2) might no longer be valid. Curriero (2006) shows that valid covariance functions might render invalid when used with non-Euclidean distances. He shows that the conditions of a metric are *necessary* conditions for a distance function to be used as a parameter in (spatial) covariance functions. In Appendix 9.1, a mathematical discussion is provided (building on the findings of Curriero (2006)), which shows that the 'frozen field' distance function is no longer a metric when used with non-constant motion vectors, e.g. motion vector fields. Thus, whenever the covariance function C_{ff} is used for estimation in this work, it is used with a single motion vector V_f . This means, the kriging covariance or semi-variance matrix (Section 2.4.2) is built with a single motion vector V_f (the application of the model to precipitation field interpolation is further described in Section 4.2.3).

Spatial and spatio-temporal empirical and theoretical variograms

Often, the geostatistical approach to optimal estimation (kriging, see Section 2.4.2) is formulated using semi-variances rather than covariances. Therefore, this section discusses certain properties of the stochastic field Z using semi-variances and relates them to the properties of covariances. A more complete discussion on the relationship between co- and semi-variances can be found in Gneiting et al. (2001), who also provide the basis for this section.

Definition 6 (Spatial Intrinsic Hypothesis). A spatial random field $Z(s)$ is said to satisfy the **spatial intrinsic hypothesis** (or to exhibit *intrinsic stationarity*) iff (if and only if) the expectation is constant, i.e. $E(Z(s)) = m$ for all spatial locations and the variance of $Z(s) - Z(s + h)$ only depends on the spatial distance vector $h \in \mathbb{R}^2$, i.e. $\text{var}(Z(s) - Z(s + h)) = 2\gamma(h)$.

Then, $\gamma(h)$ is called the spatial (semi-) variogram. The definition of intrinsic stationarity extends to spatio-temporal fields:

Definition 7 (Spatio-Temporal Intrinsic Hypothesis). A spatio-temporal random field $Z(u)$ is said to satisfy the **spatio-temporal intrinsic hypothesis** iff the expectation is constant, i.e. $E(Z(u)) = m$ for all spatio-temporal locations $u = (s, t)$ and the variance of $Z(s, t) - Z(s +$

$h, t + \Delta t$) only depends on the spatial distance vector $h \in \mathbb{R}^2$ and temporal difference Δt , i.e. $\text{var}(Z(s, t) - Z(s + h, t + \Delta t)) = 2\gamma(h, \Delta t)$.

$\gamma(h, \Delta t)$ is called the spatio-temporal (semi-) variogram. Isotropy of spatial variograms is similar to isotropy for covariances:

Definition 8 (Isotropic Variogram). A spatial variogram $\gamma(h)$ is called **isotropic**, iff it can be written as a function of spatial distance only, ignoring directions: $\gamma(h) = \gamma(\|h\|)$.

Isotropic spatial variograms are the most widely used characterizations of dependence in geostatistics, since they can be easily estimated from data using Equation 2.24.

$$\gamma_s^*(h) = \frac{1}{2|N(h)|} \sum_{(u_j, u_i) \in N(h)} [z(u_j) - z(u_i)]^2 \quad (2.24)$$

where $N(h) = \{(s_j, s_i) : t_i = t_j \wedge d(s_j, s_i) \approx h\}$ is the set of tuples of spatio-temporal locations of the same time step, separated by spatial distance class h (classifying spatial distances into distance classes is called *binning*). d is Euclidean distance and $z(u_j)$ and $z(u_i)$ are field samples. $|\cdot|$ is set cardinality. $\gamma_s^*(h)$ is called the *empirical spatial (semi-)variogram*. In a similar way, pure empirical temporal variograms, denoted with $\gamma(t)^*$, can be estimated using:

$$\gamma_t^*(t) = \frac{1}{2|N(t)|} \sum_{(u_j, u_i) \in N(t)} [z(u_j) - z(u_i)]^2 \quad (2.25)$$

where $N(t) = \{(u_j, u_i) : s_i = s_j \wedge |t_j - t_i| \approx t\}$ is the set of tuples of spatio-temporal locations that coincide in space but differ in time by temporal difference class t . Often, t is a time index, i.e. $t \in \mathbb{N}$ and thus, no binning of values is required.

Covariances and variograms not only deliver stochastic information about the random fields but can be used for prediction (see Section 2.4.2). Then, so-called *theoretical* covariance functions or variograms need to be fitted to the empirical covariances or variograms. However, not every function is a valid theoretical variogram, but only those that are *conditionally negative definite* (see Gneiting et al. (2001) or Cressie (1993) for a definition). Thus, certain function types are common in geostatistical interpolation, which are known to satisfy this property. One particular theoretical variogram model is the exponential model of Equation 2.26.

$$\gamma(h) = c_o + c * (1 - \text{EXP}(-|h|/a)) \quad (2.26)$$

where c_o is the variogram value at spatial or temporal separation zero, known as the nugget. c is the partial sill (i.e. $\text{sill} = c_o + c$) and a is the range parameter. For exponential models, the separation $3 \times a$ is known as the *effective range*, denoted with a_{eff} , which is the spatial distance where the variogram value reaches 95% of the sill. As random spatial variables separated by a distance larger than the a_{eff} can be considered uncorrelated, it is an indicator for the spatial window over which samples should be integrated for interpolation. Figure 2.6 displays an example.




Figure 2.6: Example of empirical (dots) and theoretical exponential variogram (line) and the different parameters of the exponential model.

For spatio-temporal fields, the estimation of spatio-temporal variograms is theoretically possible. However, in practice, it is unlikely that the data contains enough samples for every combination of time and space lags. Therefore, a common approach in spatio-temporal geostatistics is to estimate spatial and temporal empirical variograms separately, e.g. using Equations 2.24 and 2.25, and fit theoretical variogram models separately and then combine both models.

A common (symmetric) variogram model is derived from the product model for covariances of Equation 2.19. When the stationarity assumption of Equation 2.17 holds for the field under consideration and if the variograms are upper-bounded, the relationship between variograms and covariances can be expressed using the covariance-semivariance relationship exemplified e.g. in Gneiting et al. (2001): $C_s(h) = C_s(0) - \gamma_s(h)$ and $C_t(h) = C_t(0) - \gamma_t(h)$. Then, for the separable covariance model of Equation 2.19, the spatio-temporal variogram is given by (Cesare et al. (2001)):

$$\gamma_{st}(h,t) = C_s(0) \times \gamma_t(t) + C_t(0) \times \gamma_s(h) - \gamma_s(h) \times \gamma_t(t) \quad (2.27)$$

where $\gamma_{st}(h,t)$ is the spatio-temporal variogram, $C_s(0)$ is the spatial covariance for space-lag zero and $C_t(0)$ is the temporal covariance for time-lag zero. Both, $C_s(0)$ and $C_t(0)$ can be estimated from the data but are also available from the theoretical variogram models: $C_s(0)$ corresponds to the *partial sill* of the spatial variogram model, i.e. parameter c of the spatial model. Similarly, $C_t(0)$ corresponds to the *partial sill* of the temporal variogram model. This is the approach followed in this work for symmetric spatio-temporal interpolation and thus, the symmetric model becomes: $\gamma_{st}(h,t) = c_s \times \gamma_t(t) + c_t \times \gamma_s(h) - \gamma_s(h) \times \gamma_t(t)$, where c_s is the partial sill of the spatial variogram model, c_t is the partial sill of the temporal variogram model. In this work, when this model is used for geostatistical interpolation, the method is called **ST-OK**.

The non-symmetric spatio-temporal covariance of Equation 2.22, i.e. the 'frozen field' model, translates to variograms by replacing the Euclidean distance parameter of the spatial variogram by the 'frozen field' distance function d_{ff} . In this work, when the distance function d_{ff} is used for geostatistical interpolation, the method is called **FF-OK**.

2.4 Interpolation methods

The following sections introduce the different interpolation methods that are used.

2.4.1 Inverse-Distance-Weighted

The standard spatial inverse-distance-weighted (IDW) formula for calculating an estimate $\hat{z}(s_0)$ at a spatial location s_0 where no sample (measurement) of the field exists from n available samples $z(s_i), 1 \leq i \leq n$ is displayed in Equation 2.28.

$$\hat{z}(s_0) = \begin{cases} \frac{\sum_{i=1}^n \frac{w_i(s_0)}{\sum_{k=1}^n w_k(s_0)} z(s_i)}{z(s_i)} & \text{if } d(s_0, s_i) > 0 \forall i, 1 \leq i \leq n, 0 \neq i \\ z(s_i) & \text{if } \exists z(s_i) \text{ with } d(s_0, s_i) = 0 \end{cases} \quad (2.28)$$

where $d(\cdot)$ is some metric on \mathbb{R}^2 , usually the Euclidean distance, and $w_i(s_0)$ is the weight for the field sample $z(s_i)$ at estimation location s_0 . In IDW, weights are the inverse Euclidean distance between estimation and sample location, as shown in Equation 2.29.

$$w_i(s_0) = d(s_0, s_i)^{(-p)} \quad (2.29)$$

where the power parameter p controls the decrease in weight with increasing distance.

Spatio-temporal IDW

In order to extend the standard IDW with time, the spatial distance $d(s_0, s_i)$ can be replaced with a spatio-temporal distance $d_{st}(u_0, u_i)$. Again, a particular weight can be calculated as the inverse of d_{st} , similar to Equation 2.29. For symmetric spatio-temporal interpolation, an intuitive candidate for $d_{st}(\cdot)$ is the Euclidean metric on $\mathbb{R}^2 \times \mathbb{R}$, displayed in Equation 2.30.

$$d_{st,eucl}(u_0, u_i) \equiv \sqrt{\Delta x_{0i}^2 + \Delta y_{0i}^2 + (c_T \Delta t_{0i})^2} \quad (2.30)$$

where $\Delta x_{0i} = x_0 - x_i$, $\Delta y_{0i} = y_0 - y_i$ and $\Delta t_{0i} = t_0 - t_i$. c_T relates spatial and temporal scale (e.g. kilometres to hours) and is often referred to as a spatio-temporal anisotropy factor. See e.g. Gräler et al. (2012) on how to estimate c_T in a geostatistical setting. Equation 2.30. is obviously space-time symmetric in the sense that it e.g. does not differ between samples '30 km east, 30 minutes earlier' and '30 km east, 30 minutes later' in the weighting. Thus, it is the IDW counterpart of fully symmetric spatio-temporal covariance functions. For extending IDW with non-symmetric

spatio-temporal functionality, d_{st} could be replaced by the 'frozen field' distance function d_{ff} as introduced previously and again displayed in Equation 2.31.

$$d_{ff}(u_0, u_i) \equiv \left\| \begin{bmatrix} \Delta x_{0i} \\ \Delta y_{0i} \end{bmatrix} - V_f \Delta t_{0i} \right\| \quad (2.31)$$

where $\| \cdot \|$ is the Euclidean metric on \mathbb{R}^2 and V_f is a known motion vector of the field to be interpolated. Equation 2.31 reduces the spatio-temporal to a spatial distance under consideration of the known motion vector and is commonly termed the 'frozen field' (FF) model (Cressie & Wikle (2011), see also Section 2.3.5). Surprisingly, the distance function has not been used in an IDW setting in a similar way before, although it is a straightforward way of reducing a spatio-temporal non-symmetric interpolation problem to a spatial one without the need for variogram estimation and model fitting. Further, the model provides promising results and even outperforms its geostatistical counterpart (Section 7.1). This IDW interpolation using d_{ff} as the distance function is called the **FF-IDW** model in the following.

Example: FF-IDW

In Figure 2.7, a regular grid of 4×4 positions $[x_i, y_i, t_i]^T$ with unit spacing $1 \leq i \leq 4$ and a discrete time index of $1 \leq t \leq 2$ is shown. Each position holds a field value, indicated by the different colors, and the whole field is assumed to move forward along the motion vector $V_f = [v_X, v_Y]^T = [1, 1]^T$, i.e. the field travels 1 space unit per 1 time unit in each of the two directions.




Figure 2.7: 'Frozen field' example on a regular grid.

Therefore, e.g. the field at $u_1 = [x_1, y_1, t_1]^T = [3, 1, 1]^T$ has moved to $u_2 = [x_2, y_2, t_2]^T = [4, 2, 2]^T$. Then, when calculating an estimate of the field at u_2 , i.e. $\hat{z}(u_2)$, the measurement $z(u_1)$ should receive maximum weight. This is ensured by the 'frozen field' distance function of Equation 2.31:

$$\begin{aligned} d_{ff}(u_2, u_1) &= \left\| \begin{bmatrix} x_2 - x_1 \\ y_2 - y_1 \end{bmatrix} - \begin{bmatrix} v_X(t_2 - t_1) \\ v_Y(t_2 - t_1) \end{bmatrix} \right\| = \left\| \begin{bmatrix} 4 - 3 \\ 2 - 1 \end{bmatrix} - \begin{bmatrix} 1(2 - 1) \\ 1(2 - 1) \end{bmatrix} \right\| \\ &= \left\| \begin{bmatrix} 1 \\ 1 \end{bmatrix} - \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\| = \left\| \begin{bmatrix} 0 \\ 0 \end{bmatrix} \right\| = 0 \end{aligned}$$

Then, with Equation 2.28, it follows that $\hat{z}(u_2) = z(u_1)$.

2.4.2 Ordinary kriging

Ordinary kriging (OK) is a geostatistical method for interpolation, initially described by Matheron (1971), based on work of Danie G. Krige. In contrast to IDW, the variables involved are considered realizations $z(u)$ of a stochastic field $Z(u)$, which can be a pure spatial ($u \in D \subseteq \mathbb{R}^2$), temporal ($u \in D \subseteq \mathbb{R}$) or spatio-temporal field ($u \in D \subseteq \mathbb{R}^2 \times \mathbb{R}$). Similar to IDW, the field at an unsampled location u_o is estimated using a weighted linear combination of measurements, as displayed in Equation 2.32.

$$\hat{z}(u_o) = \sum_{i=1}^n w_i(u_o) z(u_i) \quad (2.32)$$

where $z(u_o)$ is the estimate at unsampled location u_o , n is the number of available measurements $z(u_i)$ and $w_i(u_o)$ is the kriging coefficient (or weight) of observation $z(u_i)$ at location u_o . However, in kriging, the coefficients are not calculated from spatial distance but from the covariances (resp. semi-variances) of the random variables involved and optimized in order to provide the best linear unbiased estimator (or short BLUE, see Definition 1) for estimates $\hat{z}(u_o)$ of $Z(u_o)$. Optimizing for the two criteria of BLUE estimators leads to the OK system of Equation 2.33 (derivations of the formula can be found in geostatistical textbooks such as Cressie (1993), p. 120ff).

$$\begin{bmatrix} W \\ m \end{bmatrix} = \begin{bmatrix} V & 1 \\ 1 & 0 \end{bmatrix}^{-1} \begin{bmatrix} d \\ 1 \end{bmatrix} \quad (2.33)$$

where $W = [w_1(u_o), \dots, w_n(u_o)]^T$ is the column vector of weights to be calculated, V is the $n \times n$ matrix of variogram values, e.g. derived from a spatio-temporal stationary variogram $\gamma(u_i - u_j, t_i - t_j)$ for all pairs of measurements u_i, u_j . d is the column vector of variogram values for derived from pairing estimation location u_o with all measurements u_i , e.g. $\gamma(u_o - u_j, t_o - t_j)$. m is called the Lagrange parameter that ensures unbiasedness, i.e. $\sum_{i=1}^n w_i(u_o) = 1$. Kriging comes with a measure of interpolation quality in the form of a *mean-squared prediction error*, known as the *kriging variance*.

$$\sigma_k^2(u_o) = 2 \sum_{i=1}^n w_i(u_o) \gamma(u_o - u_j, t_o - t_j) - \sum_{i=1}^n \sum_{j=1}^n w_i(u_o) w_j(u_o) \gamma(u_i - u_j, t_i - t_j) \quad (2.34)$$

This formulation of the kriging equations assumes *intrinsic stationarity*, e.g. a constant mean $E(Z(u)) = \mu_Z$ for all locations u in the domain D . However, as noted by Cressie (1993) (p.123), this is not an OK requirement but due to the variogram.

2.4.3 Regression kriging

Regression kriging (**RK**), which is sometimes also called residual kriging, no longer requires the assumption of a constant mean. RK assumes that the mean of the stochastic field Z is linearly related to another random variable or random vector of interest, say Y , via the linear regression equation:

$$E(Z(u)|Y(u) = y(u)) = \alpha + \beta y(u) \quad (2.35)$$

where α is the scalar intercept and β is the scalar or vector valued 'slope' parameter. Thus, RK requires that the predictor variable is available at every location u in the domain of interest. Then, RK estimates the regression coefficients a of α and b of β using tuples of realizations of predictor and response variable. For doing this, the method of Generalized Least-Squares is applied (GLS, see Section 2.3.2) to account for the assumed non-zero spatial correlation among the errors. Afterwards, at an unsampled location u_0 , the regression model is used for estimating the mean $\mu_Z(u_0) = a + by(u_0)$ of $Z(u_0)$. Then, the regression residuals of all samples (which are assumed to be spatially correlated) are interpolated using OK to provide an estimate of the residual $\hat{R}(u_0)$, which is added to the mean. Thus, the final RK estimate is provided by:

$$\hat{z}(u_o) = a + by(u_0) + \hat{R}(u_0) \quad (2.36)$$

In principle, the variogram for the interpolation of the residuals has to be estimated from the residuals. However, in practice, the variogram is often estimated from the original values. Further, as mentioned by Haberlandt (2011), the combination of regression with residual interpolation assumes both contradicting requirements, independent (for regression) and dependent (for interpolation) residuals. However, as mentioned by Haberlandt (2011), the method is often successfully applied in practical applications.

In this work, the method is applied using weather radar estimates of precipitation intensity as additional variable. Therefore, the step of estimating the regression model is omitted.

2.4.4 Cross-validation for performance assessment

Cross-validation is a common method for the derivation of interpolation performance measures. In cross-validation, the data set is split into training and validation data and the estimation method is trained on the training data and tested on the validation data. In interpolation, measurements are removed from the data and estimated using the remaining data by interpolation. This process is repeated for different training and validation sets. Then, the differences between estimated and measured values are aggregated to measures of prediction performance of the interpolation method. A special case of cross-validation is *leave-one-out* cross-validation, where only a single measurement is removed at a time and estimated using the remaining measurements (or subsets of it) and the process is repeated for every measurement in the data set. With time series data, e.g. for rain gauge time series, assessing the prediction performance is often necessary for locations where there is no past data as well. Then, a common method is to remove not only the single current measurement but the whole rain gauge time series at that location, which is the approach followed in this work.

2.5 Optical flow

2.5.1 Optical flow intensity conservation

Optical flow methods, such as Lucas et al. (1981) or Horn & Schunck (1981), are usually employed for estimating pixel displacement (motion) in between two images. The assumption underlying

most optical flow approaches is that the intensity (pixel/field values) remains constant in between the sampling periods and a change in values for a particular location solely comes from field motion. Formally, this means that there exists a vector in the space-time cube $h = [\Delta x, \Delta y, \Delta t]^T$ such that Equation (2.37) holds.

$$Z(u) = Z(u + h) \quad (2.37)$$

Optical flow methods can be classified into two categories: those that use partial derivatives in space and time for estimation (termed *gradient-based methods*) and those that use matching of image parts, e.g. pixels within a window, which are commonly known as *block matching methods* (Black, 1992). Figure 2.8 displays an example of the basic optical flow intensity conservation assumption of Equation 2.37.




Figure 2.8: Optical flow example. The intensity structure of Image 1 translates to the intensity structure of image 2 along a particular spatio temporal vector, here: 1 pixel in X ($\Delta x = 1$), 1 pixel in Y ($\Delta y = 1$) per 1 time unit ($\Delta t = 1$), i.e.: $h = [1, 1, 1]^T$

The intensity structure of image 1 translates to image 2 along a constant motion vector. For example, the field values in the areas enclosed by the blue rectangles is equal in both images. Block matching methods compute correlations for different spatial lags and blocks of image parts, e.g. for the two blocks highlighted in Figure 2.8. However, in this work, optical flow methods are used on irregular data provided by in-situ sensors (Section 6). The application of block matching methodologies for such irregular data would require interpolation of the data to regular grids, which would introduce additional assumptions about the field structure. Instead, a gradient-based approach is chosen, which will be introduced in the next section, and which is considered more suitable for irregular data.

2.5.2 Gradient-based optical flow

Gradient-based methods are based on the assumption that a first-order Taylor series expansion of the values is adequate, as displayed in Equation 2.38.

$$Z(u + h) \cong Z(u) + Z'_X(u)\Delta x + Z'_Y(u)\Delta y + Z'_T(u)\Delta t \quad (2.38)$$

where Z'_X , Z'_Y and Z'_T are the partial derivatives in the space-time cube. Equation (2.38) is called the linearity assumption of optical flow, as higher order terms are ignored. Combining (2.37) and (2.38) and dividing by Δt then results in the *gradient constraint* (GC) Equation (2.39).

$$Z'_X(u)\Delta x/\Delta t + Z'_Y(u)\Delta y/\Delta t + Z'_T(u) \cong 0 \quad (2.39)$$

Instantiating Equation (2.39) requires estimates of the partial derivatives. Usually, they are estimated with numerical differentiation using neighboring pixels in space and time. Estimating optical flow $v_X = \Delta x/\Delta t$ and $v_Y = \Delta y/\Delta t$ then requires at least two gradient constraints to be integrated. The linearity assumption of Equation (2.38) implies that the length of the displacement h has an influence on motion estimation accuracy: the shorter the displacement, the more valid the expansion of Equation (2.38) and, thus, the more accurate are the motion estimates. There is also an interplay between the degree of linearity of the field and the motion estimation accuracy: for a linear field, Equation (2.38) holds exactly (\cong can be replaced by $=$).

The different gradient-based optical flow methods then differ in how motion is estimated from a set of gradient constraints, which are linear equations with two unknowns. For example, Lucas et al. (1981) integrate constraints over a pixel neighborhood using least squares adjustment. In Section 6.4 the methodology for estimation used in this work is described, which reformulates the estimation problem as a regression problem.

2.5.3 Probabilistic optical flow

In order to derive probability distributions for the motion vector components v_X and v_Y , Simoncelli et al. (1991) introduced a probabilistic model for the gradient constraint (see also his PhD thesis Simoncelli (1993)). For doing this, he assumed that all variables involved are random and that their estimates are related to the true values via some Gaussian noise terms, as displayed in Equation 2.40.

$$\begin{aligned} \hat{z}'_T(u) &= Z'_T(u) + n_{z'_T}(u) \\ \hat{z}'_X(u) &= Z'_X(u) + n_{z'_X}(u) \\ \hat{z}'_Y(u) &= Z'_Y(u) + n_{z'_Y}(u) \end{aligned} \quad (2.40)$$

where $n_{z'_X}$, $n_{z'_Y}$ and $n_{z'_T}$ are the Gaussian noise variables. Then, using these relationships, the GC of Equation 2.39 can be reformulated and a conditional probability distribution can be calculated for the motion components. This probabilistic model is used in this work for calculating error measures from gradient constraint provided by in-situ geosensors (details on the approach are provided Section 6).

3 Related Work

In this chapter, the relevant related work is introduced for the main topics covered in this thesis, which are the *quantitative precipitation estimation* from weather radar, rain gauges and new types of measurement devices (3.1) and the *decentralized estimation* of environmental phenomena with geosensor networks (3.2).

3.1 Quantitative precipitation estimation from rain gauges, weather radar and other data sources

The sensing and communication technologies for in-situ rain gauges developed over the years and the scan rates of weather radar steadily increased. For example, recently in 2012, a new strategy has been introduced by the DWD switching from 15 min to 5 min volume scans (Seltmann et al., 2013). While radar provides a high spatial coverage of rainfall, the indirect estimation of rainfall from radar reflectivity often results in large errors concerning spatial or temporal bias or rainfall rate (Wilson & Brandes, 1979; Berndt et al., 2014; Legates, 2000), an issue that has been described in detail in Section 2.1.3. Further, despite the network extension of automatically recording rain gauges in recent years¹, it is still not dense enough for some use cases in hydrology (Haberlandt & Sester, 2010), e.g. for assessing areal rainfall within a river catchment area with a low number or even no rain gauges (Rabiei et al., 2016). Further, a dense network of gauges is still needed for the adjustment of radar (Haberlandt & Sester, 2010). Thus, the derivation of accurate quantitative precipitation estimates at high spatial and temporal resolutions from weather radar, rain gauges and other data sources is an active research field.

3.1.1 Precipitation estimation with weather radar

The estimation of quantitative precipitation rates with weather radar begins with the determination of the relationship between radar reflectivity and precipitation intensity. As described in Section 2.1.3, the Z-R relationship depends on drop-size distribution, is influenced by hail and snow and varies with spatial configuration: far away from the radar station location, the reflectivity might be attenuated and the Z-R relationship might be different from the one close by. Different works address these problems. For example, Legates (2000) adds a parameter to the standard Z-R relationship that represents the distance between the estimation point and the radar station. The relationship is calibrated using tuples of radar, rain gauge values and rain gauge locations. The goal is to account for the increasing overshoot and attenuation of the radar beam with increasing distance but, as mentioned, the new relationship also accounts for other effects induced by spatial

¹This is only valid for Germany (Figure 1.1) Globally, the density of rain gauges is even declining as pointed out by Overeem et al. (2013).

distance to the radar station. In order to reduce the under- or overestimation depending on the event and intensity, different Z-R relationships are applied in a number of works depending on the current measured reflectivity values (Weigl, 2015; Bartels et al., 2004). More recent works study the adaptive, sometimes real-time calibration of the Z-R relationship. For example, the study provided by Alfieri et al. (2010) considers a time window surrounding a particular (hourly) time step and calibrates the Z-R relationship using tuples of hourly aggregations of rain gauge and radar values of pixels in a $1.5 \times 1.5 \text{ km}$ window surrounding the rain gauge, collected within the time window. Both, a real-time setting is tested where the time window only includes past data for a current time step, as well a post-processing setting, termed *continuous-time* by the authors, where future data for a particular time step is assumed to be available. Time-windows tested range from 1h to 24h. The mentioned approach is similar to the one presented in this thesis in that a temporal window is considered for calibration (Section 4). However, the temporal resolution is hours and for 1-min resolution, the spatio-temporal bias between radar pixels and ground precipitation measured by rain gauges gains importance, which is, due to the hourly resolution, not accounted for in Alfieri et al. (2010). Another recent work is provided by Libertino et al. (2015) who adopt the real-time version of the time-window approach of Alfieri et al. (2010) and add a spatial calibration domain by considering radar-rain gauge tuples of the nearest neighbors (rain gauges) only, for a specific estimation location. The consideration of a spatial calibration neighborhood is similar to our approach but, again, no spatio-temporal bias is considered in Libertino et al. (2015).

As described previously in Section 2.1.3, radar estimates might not necessarily correspond to ground precipitation, especially in the presence of strong advection which might result in a spatio-temporal bias in between radar pixels and gauge measurements. This bias is more significant, the smaller the aggregation period and might disappear for very long periods such as days. Most of the work on quantitative precipitation estimation using radar and rain gauges start with temporal aggregations of hours and therefore either neglect spatio-temporal biases, as e.g. the works of Legates (2000), Alfieri et al. (2010) or Libertino et al. (2015). Or, they only consider a spatial bias. For example, the hourly radar-rain gauge merging for the RADOLAN product of the DWD (Winterrath et al., 2012) relies on the identification of the radar pixel in a 3×3 neighborhood around a rain gauge location whose hourly sum is closest to the hourly sum of precipitation values provided by that particular rain gauge. The selected radar pixel is then used to calculate a difference or ratio with the rain gauge value. The differences or ratios in values at all gauge locations are then interpolated and used to adjust the radar grid. While this might be a valid approach at hourly resolution, it can not be expected to work well for higher temporal resolutions such as 1-min, where the corresponding radar pixel can be expected to be separated from the spatio-temporal location of the gauge measurement not only in space but also in time due to the effects described in Section 2.1.3.

3.1.2 Precipitation estimation by interpolation of rain gauges measurements

Most of the approaches that employ interpolation methods for precipitation estimation consider spatial snapshot interpolation only. This is adequate, as long as (a) the sampling density both in space and time is high enough and (b) the data is aggregated sufficiently for the gauge density at hand. For example, the current rain gauge density might be sufficient for the interpolation

of hourly rainfall with kriging using an intrinsic stationary variogram (Definition 6). However, especially in convective rainfall events, precipitation clouds are small scale and precipitation rate changes quickly in space and time. At higher temporal resolutions, this leads to strong deviations from the classical assumptions underlying geostatistical methods, such as the mentioned intrinsic hypothesis. For example, while the assumption of constant mean might be valid for hourly or daily rainfall at a particular spatial scale and terrain, e.g. in a city with rather flat terrain, it is clearly violated for 1-min resolution, where areas of high rainfall exist together with areas of no rain. *Isotropy*, i.e. directional independence, is another property that is often assumed but does not always hold at high temporal resolutions. There, narrow rainfall bands might occur that show spatial correlation over long distances in a particular direction only. Finally, the distribution of rainfall strongly deviates from the Gaussian distribution. In fact, it is *bimodal* and *skewed*: there may be a large number of zero rainfall measurements, a few high values and a peak in between. This makes inferences problematic, for example those based on the kriging variance, as described by Heuvelink & Pebesma (2002). This is why classical deterministic or geostatistical interpolation methods cannot be expected to perform well at 1-min temporal resolution.

Further, the spatial snapshot approaches to rain gauge interpolation disregard time, leading to the loss of any temporal dependence structure of the field. Taking temporal persistence into account when interpolating is the more relevant, the larger the gaps between the samples in space and time are and the smaller the required temporal resolution. There are geostatistical interpolation approaches that integrate the temporal dimension into the interpolation. However, they often consider space-time interactions as being symmetric, leading to the fact that e.g. the two point samples '30 km east, 10 minutes earlier' and '30 km east, 10 minutes later' are treated, e.g. weighted, equally (see Definition 2.18). The studies of Gneiting (2002) or the book of Cressie & Wikle (2011) provide in-depth discussions of this topic. Symmetric space-time interactions can be considered adequate when there is no field advection present, e.g. when there is no wind, or no knowledge about asymmetry. Also, it may be adequate when aggregates over longer time periods such as hourly, daily, monthly or even yearly time periods are considered. However, for precipitation with a high temporal resolution, it is observable that there is an interplay between rain and atmospheric motion induced by wind. There are geostatistical approaches that take motion behavior into account by using non-symmetric spatio-temporal dependence functions. However, then, besides the problem of estimating the non-symmetric dependence structures, a particular difficulty is to ensure that the resulting spatio-temporal dependence functions (i.e. covariance functions or variograms) are valid (Cressie & Wikle, 2011). A whole line of geostatistical research studies is concerned with the development of valid spatio-temporal covariance functions. Examples are those provided by Cressie & Huang (1999), Cesare et al. (2001) or Gneiting (2002). Cressie & Wikle (2011) and Gneiting et al. (2006) also describe the 'frozen field' covariance of Equation 2.21.

There exists a significant body of work investigating the Taylor hypothesis for precipitation fields without using the spatial and temporal covariances for estimation. One of the early studies provided by Zawadzki (Zawadzki, 1973) concludes that "the Taylor hypothesis holds for time periods shorter than 40 min." The more recent study of Li et al. (2009) employed statistical hypothesis testing for investigating the validity of the hypothesis for precipitation fields estimated from radar snapshots with a spatial resolution of 4km at a temporal sampling rate of 15min . They conclude

that the Taylor hypothesis does not hold for their spatial and temporal scales but might hold for shorter scales. However, to the best of the author's knowledge, the 'frozen field' model of Equation 2.21 derived from the Taylor hypothesis has not been applied to precipitation estimation before in a similar way as it has been done in this work.

A geostatistical approach that has certain similarities to the work presented here is provided by Amani & Lebel (1997). The authors explicitly consider non-symmetric spatio-temporal geostatistical interpolation for precipitation fields. The introduced approach, termed Lagrangian kriging, uses the starting times of the gauge time series, defined as the time where the gauge measurements start deviating from zero, to estimate the starting time at the ungauged location where rainfall is to be estimated. Then, based on the differences in starting times between estimation and gauge locations, the time series of gauge measurements at each gauge are shifted in time and a spatial interpolation is performed on the shifted data. The method is applied to Sahelian rainfields at 5-min to 10-min resolution and is shown to outperform ordinary kriging. The authors mention that the approach assumes "that the observation window is of relatively small size as compared with the size of the MCSs, so that mature MCSs are homogeneously hitting the network." (MSC: Mesoscale convective system). This means, the approach might not work well for other types of precipitation events, such as those composed of a set of rain clouds with significantly smaller extent than the observation window. Further, the approach disregards weather radar as an additional important source of motion and also precipitation information.

3.1.3 Geostatistical merging of radar and rain gauge data

Merging methods derive quantitative precipitation estimates from some combination of radar- and rain gauge measurements. Geostatistical radar-rain gauge merging methods rely on radar data as an additional variable in the rain gauge interpolation process. Common geostatistical merging methods are external drift kriging (EDK) and regression kriging (RK), which have been successfully applied for precipitation data, e.g. in the studies of Haberlandt (2007), Verworn & Haberlandt (2011) or Berndt et al. (2014). However, most of the merging approaches provide coarse temporal resolutions such as hours or days and have not been applied to measurements at 1-min temporal resolutions so far.

3.1.4 Motion-based methods used in nowcasting

In nowcasting, motion vectors are estimated from radar images and used for extrapolating pixel values to produce forecasts in time. e.g. for short-time periods up to 6 hours ahead. Such methods usually rely on motion fields estimated from the current and recent past data by computer vision techniques such as optical flow (OF), and on the subsequent extrapolation in time using the motion information, see e.g. (Bowler et al., 2004). Thus, the focus of these methods is not the accuracy of point estimates of precipitation intensity, but the prediction of the spatial structure of the precipitation field to the future. For regular grids, the advection or motion estimation algorithms usually result in dense fields of motion vectors and the nowcasting then extrapolates pixel values along their respective motion vectors (Pfaff, 2013; Berenguer et al., 2005). Often, forward and backward schemes are distinguished (Berenguer et al. (2005); see also resampling schemes in Image

processing, e.g. described by McGlone (2013)). Forward schemes extrapolate pixel values along the motion vector to the spatio-temporal location of the predicted grid. This results in some pixels being populated more than once due to convergence of motion and some pixels being not populated at all, since they are not pointed to by any motion vector (Pfaff, 2013). In contrast, backward schemes start with the predicted grid and identify the corresponding pixel value in the original image. This is accomplished e.g. by inverting the latest motion vectors which are available for each pixel and 'pulling' values to the predicted grid (see the studies of Germann & Zawadzki (2002) and Pfaff (2013) for a more detailed description of the two methods). Backward schemes therefore lead to the fact that every pixel in the forecast grid is assigned a value, while forward schemes may produce gaps and thus require spatial interpolation in the forecast grid (Pfaff, 2013). Sometimes, these radar and advection based methods are not dealing with forecasting but with upsampling of the radar images in time. For example, Nielsen et al. (2014) use motion estimates from optical flow to generate precipitation grids in between measured radar grids, i.e. to upsample the radar grids in time. This is also the approach followed here (see Section 4.1.2) for radar upsampling.

However, the methods for spatio-temporal, motion-based interpolation of rain gauge data used in this work differ from nowcasting methods or the approach of Nielsen et al. (2014) in that there is the additional step of estimating dense grids from the non-gridded, potentially sparse and irregular rain gauge data. For this reason backward schemes used in nowcasting are not applicable since for most of the estimation locations, there are no samples available to "pull" data from. Therefore, forward schemes are used by relocating samples along a constant motion vector and then performing a pure spatial interpolation on the resulting temporal snapshot to close the gaps. This relocation of samples is not restricted to past samples: 'future' samples are relocated as well, backwards in time. Gupta & Waymire (1987) termed this approach the 'frozen field' model as a certain spatial structure is assumed, e.g. a spatial covariance, which is extended to space-time by relocating samples according to constant motion.

3.1.5 New data sources for precipitation estimation

The lack of weather radar and dense rain gauge networks in many countries and the problems with radar rainfall estimates has lead researchers to investigate new sensing technologies that complement radar and rain gauges or even replace them when they are not available. A promising approach concerns rainfall estimation by using microwave links of cellular phone networks, an idea first proposed in Messer et al. (2006). The observation that the microwave signal is attenuated by precipitation lead to a number of studies investigating the accuracy that is achievable with such a method. Messer et al. (2006) compare the new method with rain gauge and radar data and report promising results: the rainfall estimated using the communication network corresponds better with the rain gauge measurements, and hence with precipitation on the ground, than weather radar. Starting with the initial study of Messer et al. (2006), several works investigated the idea on a larger scale with promising results as well, for example for whole countries such as for the Netherlands in the study provided in Overeem et al. (2013).

Other recent works investigate acoustic rainfall measurements (Gaucherel & Grimaldi, 2015). Hut et al. (2014) even propose to attach these acoustic rain gauges to umbrellas to enable the provision

of crowd-sourced rainfall data. Another operational approach for collecting and providing crowd-sourced precipitation data is the website Weather Underground.² Users can register their personal weather station, e.g. located in their backyard, and deliver the data to Weather Underground. However, most of the data on the website concerns temperature, as this is a variable that is more easily measured. If precipitation data is provided, the sensor type and set-up such as the height of the sensor above ground, the closeness to wall, is largely unknown as well as the accuracy of the provided spatio-temporal location information. Accuracy is an important topic in crowd-sourced data in general (Heipke, 2010). Nonetheless, due to its potentially big amount of possibly inaccurate precipitation measurements, crowd-sourced precipitation data has the potential to provide better spatial and temporal resolutions than current methods or at least to be a valuable complement, in the future.

Other works approach the problem of timely provision of precipitation information from a technical point of view and deal with communication infrastructures and user interfaces for data access. For example, Gaitan et al. (2014) describe an infrastructure and user interface for the integrated provision of precipitation information from a number of different sources. The Hydrology Working Group of the Open Geospatial Consortium³ deals with standards, best practices and service interfaces for exchanging hydrological information through the World Wide Web.

All of the methods described above complement each other and the established methods and have the potential to lead to highly accurate rainfall estimates at high spatial and temporal resolutions in the future. Rainfall estimation with cars, an idea which has first been introduced in Haberlandt & Sester (2010) and investigated in this work, is considered part of these new developments in sensing technology for rainfall estimation. It is a crowd-based approach with the common advantages and disadvantages described previously, such as low accuracy of individual measurements but the potential of large amounts of data.

3.2 Decentralized estimation with geosensor networks

The development of wireless sensor network technologies (2.2) has lead to the new research field of decentralized spatial computing (Duckham, 2012) that studies algorithms for the decentralized processing of spatial data with geosensor networks. Those algorithms are operated by the geosensor network nodes that are located in space and that process information that is spatially located as well, e.g. information on moving objects or properties of spatial or spatio-temporal fields. In the following sections, those works from the literature are described that are most closely related to the content of this thesis.

3.2.1 Estimation of spatio-temporal field properties with GSN

There exists a significant amount of work on the estimation of the properties of spatio-temporal fields with GSN. Problems include the estimation of region boundaries, e.g. boundaries of rain

²<https://www.wunderground.com/>

³<http://www.opengeospatial.org/projects/groups/hydrologydwg>

clouds, or the estimation of peaks and pits within a field, e.g. the maximum and minimum precipitation rates within a precipitation field. For example, Duckham (2012) presents different algorithms for estimating whether a node lies inside, outside or on the boundary of a region. In addition, he provides an algorithm for determining the intersection of two spatial regions in a decentralized way. (Sester, 2009) extends the boundary estimation by boundary approximation with mobile nodes. Her approach is based on self-organizing maps and, when applied to a certain region, results in a spatial configuration of the network where certain node locations approximate the boundary of the region. Other works deal with the identification of critical points of a field, such as the peaks and pits (Jeong et al., 2014), or even spatial interpolation of point measurements in the network (Umer et al., 2010). The book of Duckham (2012) provides a thorough overview of this topic, as well as a description of the advantages of decentralized computation in the network, which also apply to the work presented in this paper. However, the problem of estimating motion properties of a spatio-temporal field with a GSN in a decentralized way has not been tackled so far.

3.2.2 Object-tracking with GSN

Another research line related to our work is the tracking (following) of advected spatio-temporal features by mobile nodes (Brink & Pebesma, 2014; Das et al., 2012). These works concern the recognition of object motion and the adaptation of the node motion to the object motion. While in these works, mobile nodes are assumed that can either move by themselves or move with the field such as buoys moved by ocean currents, the work presented in this thesis assumes a network of stationary nodes or nodes (cars) whose movement is predetermined by the road network. Then, the goal is to estimate the motion from the time series of sensor measurements collected by the nodes. There exists another significant body of work on object tracking with GSNs, i.e., generating information on the trajectory of a mobile object without necessarily following it, such as Tsai et al. (2007). However, the problem of estimating field- (not object-) motion with a GSN has not been tackled so far.

4 Methodology for precipitation intensity estimation at 1-min resolution from radar and rain gauges

This section deals with the investigation of research hypothesis 1 ("The integration of weather radar and rain gauges is beneficial for accurate 1-min point estimates of precipitation at locations with no rain gauge compared to estimates provided by radar or rain gauges only") and research hypothesis 2 ("The integration of motion information improves the quality of point estimates of precipitation at 1-min resolution compared to solely spatial or spatio-temporal interpolation methods that do not include motion information"). For doing this, weather radar data at a sampling rate of 5-min together with 1-min rain gauge data is available. In this section, the principles behind the investigation are described, Chapter 7 provides the results. Figure 4.1 displays the process that is carried out for each 1-min time step.




Figure 4.1: Flow chart for the estimation of precipitation at 1-min resolution.

Check for significant precipitation: First, if the previous radar image at a full 5-min time step shows less rainfall than 1 mm/h on average (over all pixels in the study area window), the time step is disregarded and the process repeats with the next time step.

Full 5-min timestep: Otherwise, if the time step is a full 5-min time step, motion is estimated from the radar image at that time step and the radar image at the previous 5-min time step.

1-min timestep in between two 5-min timesteps: If the current time step is not a full 5-min time step, a radar image is generated for that particular 1-min time step using the algorithm described in Listing 1 in Section 4.1.2.

Variogram estimation: The process proceeds with the estimation of the spatial and temporal variograms as described in Section 4.1.3.

Precipitation estimation: Finally, the 1-min precipitation intensity at locations without a rain gauge are estimated using different methods (Section 4.2

After finishing the estimation of precipitation, the process repeats with the next 1-min time step.

4.1 Time-window approach for estimation

Following Alfieri et al. (2010), the methodology for 1-min precipitation intensity estimation is a time-window approach that employs temporal windows over which data is integrated (Figure 4.2).




Figure 4.2: Integration of rain gauge and radar samples in time. Two parameters, T_{past} and T_{future} determine the temporal window over which data is integrated for a particular time step $t_{current}$. Large bars indicate 5-min time steps where both, weather radar and rain gauge data is available. Small bars indicate 1-min time steps where only rain gauge data is available.

The temporal integration window for a time step $t_{current}$ is determined by a temporal look-back parameter $T_{past} \geq 0$ ($T_{past} \in \mathbb{N}^0$) and a temporal look-forward parameter $T_{future} \geq 0$ ($T_{future} \in \mathbb{N}^0$), where the positive integers represent the number of minutes too look forward or backward in time. When $T_{past} > 0$ and $T_{future} > 0$, the proposed methodology can only be applied in post-processing, when future data for a certain time step is available. This is the approach followed in this work as the main goal is the derivation valid 'ground truth' rainfall for historic (i.e. past) car data to derive the W-R relationship (Chapter 5). When $T_{future} = 0$, the spatio-temporal estimation methods can theoretically be applied for the online (or real-time) estimation. Then, of course, other issues appear that are not subject of this work, such as those related to the technical infrastructure for data access.

In this work, the data integrated within the temporal window determined by both parameters is mainly used for the 'frozen field' models. Thus, the decision on the size of the temporal window determined by the parameters T_{past} and T_{future} is solely based on considerations related to time periods for which the 'frozen field' model (and thus, the Taylor hypothesis) is considered valid for precipitation fields (see also Section 3.1). Here, the observation of Zawadzki (1973), that "the Taylor hypothesis holds for time periods shorter than 40 min." is used as the basis for determining the size of the temporal window, which is set to 40 min, so $T_{past} = T_{future} = 20$.

4.1.1 Estimation of field motion

The DWD precipitation product DX (Section 2.1.3) provides the basis from which field motion information is derived. The product comprises a radar scan every 5 minutes with a horizontal resolution of $1km \times 1km$ and a range of $128km$. For this study, the raw reflectivities are transformed to rainfall intensities by applying the standard DWD parameters $a = 256$ and $b = 1.42$ (see Riedl (1986) or Seltmann (1997)). Further, regular 230×230 pixel images are generated. This pre-processing is done by the Institute of Hydrology of the Leibniz University of Hanover and is not part of this work. For motion estimation, the OpenCV implementation of the optical flow algorithm

provided by Farnebäck (2003) is used. The algorithm accepts two images as parameters, here two 5-minute radar images, and returns a motion field in pixel coordinates. An example of the derived motion vectors is displayed in Figure 1.2 in Section 1.

4.1.2 Weather radar upsampling

The motion field provides the basis (Figure 4.3) for the generation of 1-min images from two adjacent measured 5-min radar grids. For a particular pixel (red), the field motion at that pixel is used to identify the pixel locations in the previous and next 5-min radar snapshots (grey).




Figure 4.3: Creation of 1-min radar snapshots out of a pair of adjacent 5-min snapshots.

The pixel value is computed as the inverse temporal weighted average of both pixels. Listing 1 displays the pseudo-code for upsampling.

Pseudo-code: generation of 1-min- from the surrounding 5-min radar snapshots

PARAMETERS:

previous: the measured radar image of the previous 5-min time step

next: the measured radar image of the next 5-min time step

result: the radar image to be generated

flow: the flow image computed from the images **next** and **previous**.

$\text{flow}(x_i, y_i).x$ denotes the 5-min pixel flow at pixel location (x_i, y_i) along the X-axis.

$\text{flow}(x_i, y_i).y$ denotes the 5-min flow at pixel location (x_i, y_i) along the Y-axis.

time: parameter time $\in [1,4] \subseteq \mathbb{N}$ denoting the time step within the 5-min interval.

FUNCTION upsample(previous, next, result, flow, time)

```

(1)  FOR EACH radar pixel  $p_i$  with pixel position  $x_i, y_i$ 
(2)     $p_{previous} := (p_i.x - \text{flow}(x_i, y_i).x * (\text{time}/5), p_i.y - \text{flow}(x_i, y_i).y * (\text{time}/5))$ 
(3)     $p_{next} := (p_i.x + \text{flow}(x_i, y_i).x * (1 - \text{time}/5), p_i.y + \text{flow}(x_i, y_i).y * (1 - \text{time}/5))$ 
(4)    IF previous( $p_{previous}$ ).valid AND next( $p_{next}$ ).valid
(5)       $\text{result}(x_i, y_i) := (5 - \text{time}/5) * \text{previous}(p_{previous}) + (\text{time}/5) * \text{next}(p_{next})$ 
(6)    ELSE IF (NOT previous( $p_{previous}$ ).valid) AND next( $p_{next}$ ).valid
(7)       $\text{result}(x_i, y_i) := \text{next}(p_{next})$ 
(8)    ELSE IF previous( $p_{previous}$ ).valid AND (NOT next( $p_{next}$ ).valid)
(9)       $\text{result}(x_i, y_i) := \text{previous}(p_{previous})$ 
(10)   ELSE
(11)      $\text{result}(x_i, y_i).valid := \text{false};$ 

```

Code Listing 1: Pseudo-code for radar upsampling.

In short, for each pixel in the image to be generated, the current motion at that particular pixel is derived from the OpenCV flow field. Using this motion, the pixels in the next and previous images are identified (lines 2 and 3). If both pixels are valid such that they lie within the radar image boundaries, the new pixel value is calculated as an inverse temporal difference weighted average (lines 4 and 5). If only one of both pixels is valid, that one is used as the estimate (lines 6-9). If none of both pixels is valid, the pixel to be estimated is marked as invalid (lines 10 and 11).

4.1.3 Variogram estimation

The geostatistical methods based on the ordinary kriging (OK) interpolation method (Section 2.4.2) require the estimation of empirical spatial and temporal variograms. In this work, the *empirical spatial variograms* of Equation 2.24 required for OK are estimated from the weather radar images due to the sparse spatial distribution of the rain gauges (only 10 rain gauges). This is a common approach and is done in a similar way in related work such as Haberlandt (2007). Sample locations are generated on a regular grid only within the study area and not within the whole radar window, as the spatial correlation within the window is the one of interest. The samples are generated with a grid spacing of $2km$. Further, only those tuples of radar pixel values are used for variogram estimation that are spatially closer than a maximum spacing of $30km$. Although the radar data has a spatial resolution of $1km^2$ and an extent larger than $30km$, these two parameters have been chosen in order to reduce processing time when computing the empirical variograms. They are considered a good tradeoff between a small grid spacing, e.g. $1km$, and a large maximum tuple distance. Distances between tuples are rounded to bins of discrete (integer) km values. For a stable estimation, the radar samples collected over the whole temporal integration window are used for estimation, but the tuples of radar samples $z_R(u_i)$ and $z_R(u_j)$ from which the empirical variogram is estimated are always collected at the same time step, i.e. $t_i = t_j$.¹ The *temporal empirical variogram* for the spatio-temporal method ST-OK is estimated from the rain gauge measurements using Equation 2.25. The tuples of rain gauge samples $z_S(u_i)$ and $z_S(u_j)$ from which the variogram is estimated are always collected by the same rain gauge, i.e. $s_i = s_j$, at different time steps.

It has turned out that exponential models of the form of Equation 2.26 mostly fit perfectly to the empirical spatial and temporal variograms. Thus, only exponential models are used, which are fitted to the empirical variograms automatically using the Levenberg-Marquardt algorithm for non-linear least squares problems (Section 2.3.2).² In Figure 4.4, examples of empirical and theoretical exponential variograms are displayed for two different precipitation events.

¹In total, the study area of $70km \times 70km$ results in 1156 samples per radar image. Pairing all samples and selecting those that are spatially closer than $30km$ results in ≈ 250000 tuples per radar image. Multiplication with the number of time steps in the window then gives the final number of tuples for empirical variogram computation. On average, the resulting number of tuples per $1km$ distance bin is > 300000 , which is far more than sufficient for reliable variogram estimation.

²For model fitting, the Apache Commons library (<http://commons.apache.org/>) is used.

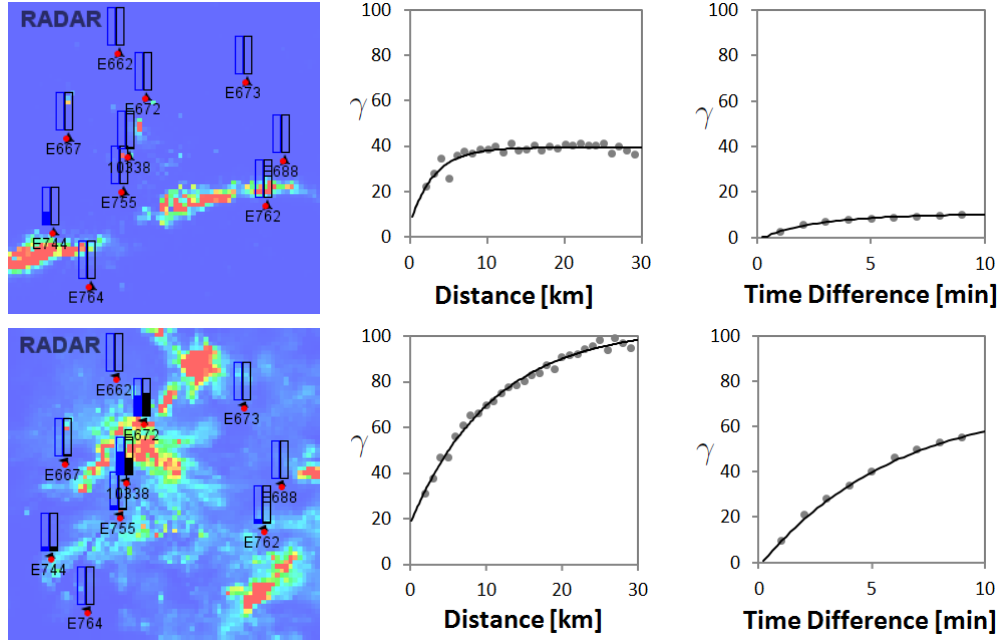


Figure 4.4: Examples of radar images and spatial and temporal variograms for two events. Radar images of the two events (left), showing current rain gauge measurements (blue bars, left) and current radar estimate (black bars, right). Corresponding spatial and temporal empirical and theoretical variograms (right).

Often, the estimation of variograms is done for whole events and then applied to every time step. As the spatial and temporal dependencies that underlie the variograms might change during an event, new variograms are estimated on a regular basis in this work. In order to reduce processing time for the processing of each 1-min time step, new spatial empirical variograms are estimated and new theoretical variograms are fitted every 10 minutes only, which is considered sufficient.

4.2 Estimation methods

The estimation methods investigated are those described in Section 2.4: inverse-distance weighted (IDW), ordinary kriging (OK), space-time symmetric ordinary kriging (ST-OK), 'frozen field' IDW (FF-IDW), 'frozen field' ordinary kriging (FF-OK) and regression kriging (RK). All of these estimation methods (except FF-IDW) have been described in the literature but have not been applied in a similar way to precipitation estimation at 1-min temporal resolution so far, as described in Section 3.1. The three new methods that are introduced are termed "radar dynamic-bias correction" (RADAR-DBC), 'frozen field' IDW (FF-IDW) and 'frozen field' regression kriging (FF-RK). In the following sections, the new methods are described in detail and, where necessary, more details are provided for the methods already described in Section 2.4, such as the method for variogram estimation. The methods investigated can be classified into spatial interpolation methods (IDW, OK), spatio-temporal interpolation methods (ST-OK, FF-IDW, FF-OK) and radar-rain gauge merging methods (RK, FF-RK, RADAR-DBC). While the spatial and spatio-temporal interpolation methods estimate precipitation from the rain gauge measurements only, the merging methods compute an estimate by some combination of radar and rain gauge values. Finally, estimation methods

solely based on the radar values are tested and described in the following sections. Chapter 7 then provides the results of the experiments.

In the following, if not indicated otherwise, the spatio-temporal estimation methods integrate all data within the temporal window specified by the parameters T_{past} and T_{future} . In the following sections, rain gauge measurements at a particular spatio-temporal location $u_i = (x_i, y_i, t_i)$ are denoted with $z_G(u_i)$, radar samples with $z_R(u_i)$. The pure spatial, snapshot-based interpolation methods compute estimates from rain gauge measurements of the current time step $t_{current}$ only. In the study presented in Section 7, data of only 10 rain gauges is available. Thus, only 10 samples of the 10 gauges are used for spatial interpolation at each 1-min time step. In cross-validation, this number reduces to 9 ("leave-one-out"). The spatio-temporal methods compute estimates from all rain gauge measurements within the temporal integration window determined by T_{past} and T_{future} . For a time window of ± 20 (and thus, 41 minutes) and a set of 10 rain gauge, the number of samples for spatio-temporal interpolation is $n = 410$. In leave-one-out cross-validation, this number reduces to $n = 9 \times 41 = 369$.

4.2.1 Spatial rain gauge interpolation methods

In the experiments of Section 7, the spatial interpolation methods **IDW** and **OK** described previously in Section 2.4 are compared. For **IDW** interpolation, the power parameter of $pp = 2$ is used, which is a common choice. For **OK**, the exponential spatial variogram model is computed as described in the previous section.

4.2.2 Space-time symmetric rain gauge interpolation method

Space-time symmetric interpolation assumes that the spatio-temporal behavior of the phenomenon is space-time symmetric (Equation 2.18). This assumption is clearly violated for precipitation intensities at 1-min resolution, where wind plays a dominant role and leads to clear asymmetries in the spatio-temporal behavior of the intensity field. Nonetheless, a space-time symmetric method is applied as a kind of benchmark for comparison with the other methods. For the space-time symmetric spatio-temporal interpolation (**ST-OK**), the ordinary kriging equations are used with the symmetric spatio-temporal variogram model of Equation 2.27. The theoretical variograms involved in Equation 2.27 are computed as described previously.

4.2.3 Space-time asymmetric rain gauge interpolation methods

Space-time asymmetric methods no longer assume symmetric space-time behavior of the precipitation field (Section 2.3.5). In this work, two space-time asymmetric methods are compared. Both methods rely on the 'frozen field' distance function d_{ff} of Equation 2.23. d_{ff} is either used directly for weight computation in an **IDW** setting and replaces the Euclidean distance function in Equation 2.29. In this case, the method is termed **FF-IDW**. Or, it replaces the Euclidean distance as a parameter in the spatial theoretical variogram estimated previously for **OK**. Then, the method is termed **FF-OK**. The motion vector V_f required for d_{ff} is derived from the optical

flow field by averaging the vectors over all sensor locations (here: rain gauge locations) within the study area and by transforming pixel motion to motion in the underlying spatial and temporal coordinate system (metric Gauss-Krüger system for space and minutes for time).

4.2.4 Radar-rain gauge merging methods

While the previously described methods of IDW, OK, ST-OK, FF-IDW and FF-OK provided precipitation estimates based on rain gauge data only and used radar data only for the motion information, the radar-rain gauge merging methods have in common that an estimate is computed by some combination of radar and rain gauge values.

Regression kriging (RK)

The method regression kriging (**RK**) introduced previously in Section 2.4.3 is used as one of the radar-rain gauge merging method. As predictor variable, the average of the current radar values within a 3×3 pixel window surrounding the rain gauge location are used. As both variables, radar and rain gauge estimates, already represent precipitation intensity in the unit mm/h , the estimation of the regression coefficients can be omitted. The estimation then consists of the following remaining steps (Haberlandt, 2007):

- 1.) The residuals at a particular gauge location s_i are computed as $R(s_i) = z_G(s_i) - z_{R,3 \times 3}(s_i)$.
- 2.) The residuals $R(s_i)$ of all rain gauges and rain gauge locations s_i are interpolated with OK using the spatial variogram, which has been estimated as described previously, providing an estimate of the residual at estimation location s_0 , $\hat{R}(s_0)$.
- 3.) The residuals are added to the trend, providing the final RK estimate of precipitation intensity at the location of estimation s_0 : $\hat{z}_{RK}(s_0) = z_{R,3 \times 3}(s_0) + \hat{R}(s_0)$.

For a statistically sound methodology, the variogram has to be calculated from the regression residuals and not from the original values. However, here it is assumed that the residuals show similar spatial correlation as the original value. Since the kriging variance, which depends on the variogram values, is not used in this work, this is considered a valid approach.

Dynamic radar-based bias-correction method (RADAR-DBC)

In this section, a new radar-rain gauge merging method is introduced, termed RADAR-DBC, that consists of four steps (Figure 4.5).




Figure 4.5: Process scheme for the RADAR-DBC method.

First, at each time step and for each rain gauge, a spatio-temporal radar-rain gauge bias (lag) is calculated using cross-correlations between 1-min radar and the rain gauge time series within the time-window of 40 minutes. Second, a radar-rain gauge adjustment coefficient is computed as a simple regression function without intercept using tuples of values of the highest correlated radar time series and the gauge time series within the time-window. Third, for each estimation location, e.g. a location without a rain gauge, the spatio-temporal bias and adjustment coefficient is spatially interpolated over the whole study area and used for estimation.

RADAR-DBC: Computation of space-time lag

The 1-min radar time series at each pixel location enable the calculation of cross-correlations between radar time series R and rain gauge time series for any gauge G_i at gauge location $s_i = (x_i, y_i)$ for different spatial and temporal lags using Equation 2.9. For the method RADAR-DBC, this is done at each 1-min time step for all 121 radar time series in a spatial 11×11 -pixel window surrounding each rain gauge pixel location, including radar and rain gauge data within the time window of $\pm 20min$. This results in correlation values for different spatial lags. In addition, cross-correlations are computed for each spatial lag and different temporal lags of $\pm 5min$. Thus, in total, cross-correlations $corr_{GR}(\Delta x, \Delta y, \Delta t)$ are computed for all $\Delta x \in [-5km, 5km]$, $\Delta y \in [-5km, 5km]$ and $\Delta t \in [-5min, 5min]$, resulting in $11 \times 11 \times 11$ correlation values at each time step and gauge location. Finally, the space-time lag is selected that provides maximum cross-correlation:

$$\begin{bmatrix} \Delta x_{max} \\ \Delta y_{max} \\ \Delta t_{max} \end{bmatrix} = \arg \max_{\Delta x, \Delta y, \Delta t} corr_{GR}(\Delta x, \Delta y, \Delta t) \quad (4.1)$$

RADAR-DBC: Computation of adjustment coefficient

After the identification of the best space-time lag $(\Delta x_{max}, \Delta y_{max}, \Delta t_{max})$, the remaining tuples of rain gauge, $z_G(x_i, y_i, t_i)$, and radar values, $Z_R(x_i + \Delta x_{max}, y_i + \Delta y_{max}, t_i + \Delta t_{max})$ within the current temporal observation window determined by T_{past} and T_{future} can be used to calibrate a current radar-rain gauge adjustment coefficient at the spatial gauge location $s_i = (x_i, y_i)$. Here, the model of a simple linear regression without intercept is used:

$$z_G = a \times z_R \quad (4.2)$$

For calibration, tuples are only used if both values of a tuple, radar and rain gauge, are larger zero. Further, the coefficient a is restricted to the interval $[0.5, 2]$, preventing a too-strong deviation of estimates. The results in Section 7 show that this interval is sufficient for correcting the underestimation of radar. However, a coefficient of 0.5 does not seem to be sufficient for correcting the radar overestimation, especially in the winter months. More sophisticated adjustment methods for dynamic, possibly non-linear adjustment computation are left for future work.

RADAR-DBC: Interpolation of spatio-temporal bias and adjustment coefficient

The method for radar-rain gauge bias identification and adjustment coefficient computation results in a spatio-temporal bias $\Delta x, \Delta y, \Delta t$ and adjustment coefficient a at every rain gauge location at every 1-min time step. In order to use the information at locations with no rain gauge, spatial interpolation is required. As described previously, the spatio-temporal bias between radar pixels and ground locations can have several reasons, ranging from the geometric configuration such as the distance to the radar to influences of the current weather situation, such as advection effects.

The assumption underlying the proposed interpolation scheme is that the bias and the Z-R relationship is (a) correlated spatially such that the closer two spatial locations are, the more similar is the spatio-temporal bias and coefficient a . And (b), the information is correlated with the distance to the radar station such that the more similar this distance is, the more similar is the spatio-temporal bias and adjustment coefficient. Figure 4.6 displays a motivating example for a relationship between radar estimates and rain gauges that depends on the spatial locations.




Figure 4.6: Motivating example for RADAR DBC Estimation method. The example shows a time series of radar images from left to right separated by 10 minutes collected in January 2012. Blue color indicates zero, green color low precipitation intensity ($\approx 12.5\text{mm/h}$). Blue bars show the current rain gauge measurement with a maximum of 25mm/h , black bars the current radar estimate at the gauge location with a maximum of 25mm/h . The black triangle denotes the radar station location. A clear pattern in the radar estimates can be recognized that depends on the spatial configuration: the larger the distance to the radar station, the larger the overestimation of precipitation intensity by radar, possibly due to ice in the atmosphere and the increased beam elevation.

The example shows that far away from the radar station, the radar is heavily overestimating precipitation intensity by a factor of 2 up to ≈ 10 at some locations (e.g., gauge E688 at the last time step). Thus, the radar-rain gauge adjustment coefficient at these locations must be different from the coefficient close to the radar station. In addition, due to the reasons described in Sections 2.1.3 and 3.1.1, it is likely that the spatio-temporal bias between radar pixels and ground precipitation is related to the spatial configuration in a similar way, i.e. spatially closer locations show a similar bias as well as locations with a similar distance to the radar station.

Inspired by these considerations, the spatio-temporal biases computed at a gauge location at a particular 1-min time step are spatially interpolated in a similar way as the radar-rain gauge adjustment coefficient. Thus, in the interpolation, the weight of the information generated at gauge location s_i for a spatio-temporal estimation location s_0 is computed from a combination of Euclidean distance between s_i and s_0 and difference in Euclidean distances to the radar station location s_R . Here, the simple distance function of Equation 4.3 is used in an IDW interpolation.

$$w_i(s_0) = (|d_{eucl}(s_R, s_0) - d_{eucl}(s_R, s_i)| + d_{eucl}(s_0, s_i))^{(-0.5)} \quad (4.3)$$

where s_R is the spatial location of the weather radar and $|\cdot|$ denotes absolute value. The power parameter of 0.5 is used to generate smooth estimates. Figure 4.7 displays the involved Euclidean distances visually for an estimation location s_0 , a rain gauge location s_i and the weather radar location s_R .




Figure 4.7: Distances involved in IDW weight computation for the method RADAR DBC.

RADAR-DBC: Precipitation intensity estimation

The interpolation of each of the three biases and the adjustment coefficient results in four spatial fields at every 1-min time step: $\Delta x(s)$, $\Delta y(s)$, $\Delta t(s)$ and $a(s)$. The **RADAR-DBC** estimate at a specific spatio-temporal location $u_0 = (s_0, t_0)$ is then computed as:

$$\hat{z}_{DBC}(u_0) = a(s_0) \times z_{R,3 \times 3}(x_0 + \Delta x(s_0), y_0 + \Delta y(s_0), t_0 + \Delta t(s_0)) \quad (4.4)$$

where $z_{R,3 \times 3}$ is the average of radar pixels in a 3×3 window surrounding the given radar pixel location.

'Frozen field' regression kriging (FF-RK)

Instead of computing the bias adjustment coefficient a , another method termed 'frozen field' regression kriging (**FF-RK**) is tested. Here, as additional variable at a spatio-temporal estimation location s_0 , the radar pixel average $z_{R,3 \times 3}(x_0 + \Delta x(s_0), y_0 + \Delta y(s_0), t_0 + \Delta t(s_0))$ identified previously for the RADAR-DBC method is used and the residuals are interpolated using the FF-OK distance function (OK with the 'frozen field' distance function). This way, the assumed superior performance of FF-OK over OK is used together with the identified 'best fitting' radar pixels as trend.

4.2.5 Estimation methods solely based on radar

Radar aggregates in different spatial and temporal windows around the estimation location are calculated. The spatial window ranges from the single radar pixel at the estimation location to a $11 \times 11 = 121$ pixel window (Figure 4.8, left). The size of the temporal 1-min window is $\pm 10min$ and estimates for a particular time step are averaged over a spatial pixel neighborhood of $3 \times 3 = 9$ pixels (Figure 4.8, right).




Figure 4.8: Radar windows for the radar-based estimation methods. Spatial windows (left) and temporal windows (right) surrounding a particular estimation location (black dot).

4.3 Summary

This section has introduced the different methods for the estimation of precipitation intensity at 1-min temporal resolution at arbitrary locations in the study area. The different methods can be classified into spatial rain gauge interpolation methods, spatio-temporal rain gauge interpolation methods, radar-rain gauge merging methods and estimation methods solely based on radar. The comparison of the performance of the different methods using data from the Hanover weather radar and rain gauges is provided in Chapter 7.

5 Methodology for precipitation intensity estimation with car sensors

In order to estimate precipitation intensity with cars, e.g. using the number of windscreen wipes per minute as a predictor variable, a functional relationship between the sensor readings and rainfall needs to be established, termed Wiper-Rainfall (W-R) relationship in the following. The nature of such a relationship ranges from a simple linear regression with a single predictor variable, e.g. the wiper frequency per minute, to more complex non-linear models with multiple parameters. Calibration can either be performed in an off-line fashion in a controlled environment such as a laboratory, as done in Rabiei et al. (2013), or also 'in the field'.

Calibration of W-R relationships using data collected 'in the field' is the goal of this work. Test cars have been equipped with sensors for recording wiper activity, position and time and have been compared with some '*ground truth*' derived from rain gauges and weather radar. In this way, it is theoretically possible to investigate influencing factors that are not easily testable in the laboratory, such as car speed, wind speed or local environment. However, it has turned out that the main problem in deriving accurate W-R relationships is the availability of some suitable and correct '*ground truth*', i.e. reference rainfall at the car positions.

This work investigates three different variables as the main predictors in a W-R relationship. These are: the wiper frequency per minute of cars with **manually-operated windscreen wipers**, the wiper frequency per minute of cars with **automatically-operated windscreen wipers** and the readings of **optical rain sensors** detecting water on the windscreen. In the following, the car sensor prototypes, the experimental set-up and the pre-processing of the collected data is described that is performed before the data is integrated with the reference provided by weather radar and rain gauges (Section 5.1). Then, a more theoretical investigation of W-R relationships, their influences and calibration is given (Section 5.2), followed by a discussion on the influence of the car speed and windscreen angle on the sensor readings (Section 5.3). The results of the experiments are provided in Section 7.2 of Chapter 7.

5.1 Car sensors

5.1.1 Wiper Frequency Sensor

The wiper frequency sensor is based on an ATMEGA324P microcontroller connected to a Navilock GPS-receiver and a magnetic sensor. For power supply, the microcontroller is connected to the car power. The magnetic sensor is placed behind the windscreen and triggers each time a magnet attached to the windscreen wipers (small disk 8x5mm of strength N45) passes, which happens two times each single wipe. Each time the sensor triggers, a NMEA String (i.e. position, time and additional information) is recorded in an ASCII text file on an SD card with the current date as filename. The SD card content is manually copied to a PC for further processing on a regular

basis. The raw data is then parsed into the spatial data base PostgreSQL with enabled PostGIS extension and a coordinate transformation to the Gauss-Krüger projection is executed using the built-in functions of PostGIS. In Figure 5.1, the sensor and an extract of the raw ASCII files from the data logger is displayed.




Figure 5.1: Image of sensors and data logger (left, power supply is not displayed) and extract of raw ASCII data. Time in UTC (first column from left), WGS84 latitude (second column) and WGS84 longitude (third column).

As mentioned before, the data from the wiper frequency sensor is recorded every time, the wipers pass the magnet, which can also happen more than once within a second, as displayed.

5.1.2 Xanonex optical sensor

The Xanonex sensor is developed by the company rain-electronics GmbH (www.rain-electronics.de) and has especially been designed for equipping cars with manually-operated wipers with an automatic wiper option. The sensor can be attached to the windscreen from the inside of a car. Eight LEDs, arranged in a circle around the sensor in the center send near infrared light at a wavelength of 740 nm through the windscreen glass, which is reflected at the outside surface of the windscreen. In case of water on the windscreen surface, the angle of refraction of the infrared ray changes and thus also the amount of light reflecting back to the sensor. The sensor gives a signal whose length is an indication of the amount of rain on the windscreen surface. The duration of the signal depends on the drop size and is in the range of 0.01 s to 2.2 s . The sensor specification lists a possible drop size range of $[100\mu\text{m}, \infty\mu\text{m}]$. For analyzing the signal, the Xanonex has been connected to an ATMEGA324P and the data is logged every 10 seconds in the unit *signal length/10 seconds*. The logging frequency of 10 seconds is a trade-off between the two contradicting requirements that (a) the data should not be logged too often, as writing on the SD card requires a certain amount of time during which new incoming sensor data can not be recorded and (b) the possibility of reliable estimation of the average car speed, which would be more difficult at 1 minute logging frequency. Again, the raw data is further processed using a PostgreSQL spatially-enabled database.¹

The Xanonex rain sensor is still not completely independent of the windscreen wipers, as the sensing area is located within the area cleaned by the windscreen wipers. However, the direct

¹The hardware set-up for both sensor types is kindly provided by Bastian Heinrich from the Institute of Hydrology, who also implemented the data logging for the frequency sensor. The data logging for the Xanonex sensor was done by the author.



Figure 5.2: Left: The Xanonex optical rain sensor, which is placed behind the windscreen (image courtesy of B.Heinrich, Institute of Hydrology, Leibniz University Hanover). Right: Extract of raw ASCII data with time in UTC (first column from left), WGS84 latitude (second column), WGS84 longitude (third column) and the sensor reading in the unit $[s/10s]$ (fourth column).

influence of the wipers on the Xanonex sensor readings has been tested in a dry setting and has turned out to be non-existent.

5.1.3 Other sensors investigated

Another optical sensor provided by Volkswagen was tested. For doing this, the LIN-Bus data logger BabyLin of the company Lipowsky (<http://www.lipowsky.de>) was connected to the sensor. However, it turned out that the the raw sensor data from the optical sensor is not available via the LIN Bus interface of the car but only the wiper frequency and driver settings such as the sensitivity. While this is still valuable information, the approach has not been followed further as the operation of the data logger required more hardware within the cars, such as a laptop, and was restricted to VW test cars with that particular optical sensor type.

5.1.4 Experimental setup and preprocessing

Different cars in the Hanover area have been equipped with either with a wiper frequency sensor or the optical rain sensor. In total, 9 cars have been equipped with wiper frequency sensors. Five of the equipped cars belonged to the Hanover Taxi company "Hallo Taxi!". Those cars collected by far most of the wiper frequency data. However, the optical sensors could not be installed within taxi cars as the prototype required a set-up that was incompatible with the requirements of a taxi car (cables on the dashboard, operation by driver necessary, different drivers of the taxi, etc.). Instead, four private cars have been equipped with an optical rain sensor. Section 7.2 provides the results of the experiments and a more detailed description of the underlying data.

After data collection and integration into a PostGIS database, the wiper frequency per minute is derived from the wiper counts collected by the cars by counting wipe events within each full minute. The Xanonex signal length per 10 seconds that is logged every 10 seconds is accumulated to 1-min values by summing up all data within a full minute. In order to derive a spatial location for the aggregated data, the car positions within the particular minute are averaged by computing the arithmetic mean of the X and Y metric Gauss-Krüger coordinates. Further, for Xanonex

data, the car speed is calculated as the arithmetic mean of all car speeds calculated from adjacent 10-second samples within that minute. For wiper frequency data, the car speed is calculated in a similar way. In Figure 5.3, the aggregation methodology is displayed visually for the Xanonex readings.




Figure 5.3: Derivation of 1-minute from 10-second samples - Schematic Example (T: Timestamp, R: Sensor Reading, S: Car Speed).

This preprocessing then results in tuples of (car, position, time, speed, sensor reading)-information within the PostGIS database, for example, the tuple ("Skoda Fabia", 3549253.269, 5805271.104, 03.07.2013 14:23, 40 km/h, 1.45). The derived 1-minute point samples of wiper frequency or Xanonex readings then provide the data basis for the further analyses provided in Section 7.2 to derive the W-R relationship.

5.2 Theoretical considerations for the calibration of the W-R relationship in the field

W-R relationships can be expected to be influenced by numerous factors such as car types, drivers, the road surface or current wind conditions, etc. In general, predictors that are measured can function as parameters of the W-R relationship and hence their influence, i.e. model coefficients, can be calibrated. In general, highly accurate W-R functions would require the integration of all influencing factors, which then have to be either:

1. dynamic and measured (i.e. a model parameter such as the car speed), or
2. static (constant) for a particular car, e. g. car type or windscreen angle, and not necessarily measured

However, not all influencing factors that are dynamic can be measured and it is most likely not feasible to perform the calibration of the W-R relationship for each car, car type or driver individually. Thus, the idea that has been explored theoretically was the automatic, decentralized calibration of each car individually by assuming wireless communication with stationary rain gauges transmitting precipitation intensity measurements to the cars. This way, over time, each car would generate a W-R relationship that works best for the car specific constant factors such as the windscreen

angle, the particular wiper frequencies or the driver (Schulze et al. (2010), Fitzner et al. (2012) or Fitzner et al. (2013) provide deeper discussions on this topic). However, due to the limiting factor of accurate 'ground truth' precipitation intensities, which has been described previously, this idea has not been further explored with real data and thus remained largely theoretical.

As mentioned, the dynamic factors influencing the W-R relationship can not always be measured. For example, in an area where the car is covered by trees or bridges that prevent the rain from reaching the car sensor, any W-R function disregarding this information will underestimate rainfall. Or, it is likely that there will be differences in the sensor readings, e.g. the wiper frequency, when the car is driving with or opposite the wind direction, even if the precipitation intensity is equal. These dynamic and sometimes not easily measurable factors might include the current car environment such as **road surface** that influences road spray, **tunnels**, the **wind speed and direction** or the erratic wiper operation by the **driver**. As these unmodelled factors are often spatially (car environment), temporally (current driver), or spatio-temporally correlated (wind speed and direction), they result in autocorrelated errors. For example, if a car is traveling through a dense forest and hence underestimates the rainfall, it is quite likely that at a subsequent measurement, the car is still traveling under similar conditions and still underestimates. Or, if a particular driver A prefers faster wiper frequencies than a driver B who provided the data for calibration, the W-R function of some car will always overestimate the rainfall, when driver A is driving. Thus, in the field experiments where the car data is used for W-R function calibration, methods accounting for these autocorrelations in the errors would theoretically be required, such as Generalized Least Squares (GLS), in order to receive reliable error statistics. The calibration of a W-R relationship using GLS has been explored in a master's thesis (Bittner, 2012).

However, in the scope of work, it has turned out that the available 'ground truth' rainfall at the car positions is the most limiting factor in the analyses. The limited accuracy of the reference prevents any fine-grained analysis of the influence of single factors. Investigating e.g. the influence of road spray, local car environment, driver or the mentioned autocorrelation in the errors would require much more accurate reference data, e.g. derived from a denser network of rain gauges. In addition, much larger sample sizes than those that could be collected in the scope of this work would be required. Therefore, this work focuses on the analysis of the influence of the two important variables **car speed** and **windscreen angle** on the sensor readings of the cars.

5.3 Dependency between car speed, windscreen angle and sensor readings

Among physicists, it is agreed upon that the speed of a moving object influences the amount of rain per time unit hitting the object. For example, Bocci (2012) investigates the ratio between the amount of water hitting an object under stationary and moving conditions and under the assumption of vertical rain. He assumed different shapes of the moving objects, including flat surfaces with certain angles, and shows that the relationship between speed and rain amount is linear but certainly depends on the angle of the surface in motion direction. Rabiei et al. (2013) employed these findings of Bocci (2012) to calculate the ratio for different assumed rain drop falling speeds and car windscreen angles and compared them to those obtained from experiments with an optical rain sensor moving through rain simulated in a laboratory. The empirical results confirm

the theoretical ratios. Thus, the theoretical ratios are also employed in this work and compared with the data collected by optical rain sensors attached to the windscreen of cars driving in the field. However, these relationships can not directly be transferred to the wiper frequencies of cars, as these depend on the wiper system (manual or automatic) and the car type. The car speed-wiper-frequency relationship for the two types of wiper systems, manual and automatic, are discussed in the next sections.

5.3.1 Manually-operated windscreen wipers

The speed-angle-rain amount relationship described in the previous section is certainly also known by car manufacturers who couple the windscreen wiping behaviors of the cars to the car speed such that the wiping activities intensify with increasing car speed automatically, which means that the pause intervals in between single wipes get shorter. Thus, there is a **causal relationship** between car speed and wiper frequency, which is implemented by the car manufacturers. Therefore, certain wiper frequencies occur preferably (or probably even only) at certain car speeds. This is a behavior that can be observed in most cars and that has been confirmed in a personal communication with an expert on windscreen wipers of Volkswagen. Thus, at least for those cars with manually adjusted windscreen wipers, the wiper frequency is not the variable of interest but the **wiper frequency class** that is provided by the human operator by adjusting the lever controlling the wipers. Then, it becomes clear that an analysis of (wiper frequency, rainfall)-tuples, e.g. a simple regression, is not suitable. Simply adding car speed as an additional predictor variable in the relationship is also not beneficial due to the causal relationship and the resulting assumed strong co-linearity between car speed and wiper frequency, which would render the results less reliable (see Section 2.3.2).

Thus, in this work, in order to eliminate the car speed influence on the wiper frequency of cars without a rain sensor, the current wiper class / wiper setting is identified and the rainfall distribution is calculated for each class. If the wiper class is known, e.g. via a connection to the CAN or LIN buses, the car speed and the wiper frequency of such cars can be neglected. The underlying assumption is that the current rainfall uniquely determines the wiper setting and that the car speed-wiper frequency relationship implemented by the car manufacturers is valid and works well for any driver. However, in this work, a connection to the car electronics is not available but only wiper frequency sensors. Thus, the wiper class has to be derived from information on the frequency and current car speed. However, since most cars still allow the driver to influence the wiper frequency of a particular wiper class by an additional lever², this is a difficult task and another potential source of inaccuracy in the W-R relationships derived from such cars.

5.3.2 Automatically-operated windscreen wipers

The analysis of wiper frequencies that are automatically adjusted by optical rain sensors in the windscreen are even more complicated. For such systems, there is an interplay between the optical rain sensor detecting rain on the windscreen, the car speed and the human operator. The

²For example, the author's Skoda Fabia allows the setting of four different frequencies for each of the base classes set by the wiper lever.

algorithms implemented in the car for deriving a wiper frequency from optical sensor readings and car speed are unknown and likely to differ between car manufacturers and car types. Further, at least for those cars tested, there is still the possibility for the human operator to adjust the sensitivity of the system, which has a direct impact on the systems wiper frequency. Finally, the wiper frequencies that are automatically derived can still be completely overruled by the human operator by switching the wiper automatic off. Thus, a high number of factors influence the wiper frequency with rainfall being only one of them. Therefore, for the W-R calibration of cars with automatically adjusted wiper frequencies, a connection to the car bus and the accessibility of all of this information is considered of vital importance, if the raw data from the optical sensor is not available.

For this work, only the wiper frequencies are available as the main predictor variable. Following the laboratory experiments of Rabiei et al. (2013) and in order to eliminate the car speed influence from the wiper frequency, a model is calibrated that describes the relationship between car speed and frequency, disregarding rainfall. As advocated by the theoretical ratios derived from Bocci (2012), the relationship seems to be roughly linear. Non-linearities in the data, e.g. for low frequencies, are discussed in Section 7.2. Therefore, the model used in this work is a simple linear regression model. After calibration, the model is used to eliminate the car speed influence from the frequency. The resulting variable is then related to the reference rainfall rainfall.

5.3.3 Xanonex optical sensor

Although, as described previously, the Xanonex sensor is placed behind the windscreen in an area where the wipers clean the windscreen, the wiper frequency does not have an influence on the Xanonex readings. Thus, the increased magnitude of sensor readings with increasing car speed observed in the experiments (Section 7.2) can be expected to result from the increased amount of water on the windscreen. The elimination of this influence from the sensor readings is performed in a similar way as for the wiper frequency data of the cars with automatically-operated windscreen wipers: first, a linear regression model is calibrated between car speed and sensor readings, disregarding precipitation intensity. Afterward, this model is used to eliminate the car speed influence from the sensor readings. The resulting values are then related to the reference.

5.4 Summary

This section has addressed the precipitation intensity estimation with car sensors. Three different types of car sensors were introduced and the preprocessing of the data before the integration with reference rainfall was described. Section 7.2 provides the results of the experiments relating the car measurements to reference precipitation intensities derived with the methods of Section 4 to derive the W-R relationship.

6 Methodology for motion field estimation with a geosensor network

This part of work investigates the motion of spatio-temporal dynamic fields by in situ sensors that are irregularly distributed in space. The sensors are assumed to be attached to computing and communication facilities and are therefore considered as the communicating nodes of a geosensor network. One of the well-known optical flow algorithms is used as a basis and adjusted to the specifics of GSNs and spatio-temporal fields, such as the irregularity of samples, the strong constraints on communication and computation and the assumed motion constancy over sampling periods. This way, the pixels in images translate to field samples collected by a node, and the pixel neighborhood translates to the node neighborhood in the GSN.

In common applications of optical flow on arbitrary images such as images of road intersections, moving objects can occur that change their direction, stop or accelerate rather quickly, e.g., within a couple of frames. In this case, temporal continuity usually only holds for very short time periods (Black, 1992). In addition, motion in arbitrary images often exhibits sharp motion discontinuities, e.g., at boundaries of the moving objects such as cars with opposing motion passing by each other. Computing the motion of atmospheric or oceanographic fields from images or irregular samples is different. First, a certain degree of a priori knowledge of the field motion properties can be assumed, such as knowledge on wind speed or on the advection of rainclouds. This domain knowledge can be used to specify the required parameters of the optical flow algorithms. Further, for atmospheric or oceanographic fields, it is likely that the motion is rather consistent over long sampling periods, and therefore, integration of motion information over time gains importance. In addition, sharp motion discontinuities can be considered unlikely: they might only arise at boundaries of the field, e.g., at the boundaries of rain clouds. However, then, they do not exist in reality since the atmosphere surrounding the rain cloud also moves, but affect the motion estimation algorithm that is based on field measurements. For example, while the atmosphere still moves when there is no rain, the motion can only be estimated when there is some rain that is not uniform. This is also known as the "blank wall" problem in optical flow literature.

All of these considerations provide the basis for the design of the decentralized motion estimation algorithm presented in this section. Further, some of the main differences of the proposed GSN-based motion estimation to image-based optical flow are summarized as follows:

Controlled node deployment, sampling rate and irregular data: In image-based optical flow, the pixels determine the sampling locations, and the motion speed and direction are influenced by numerous factors, such as the spatial resolution and orientation of the image, sampling rate, the spatial distance of the camera to the moving object, as well as the speed and direction of the moving object. When motion is to be estimated with a GSN deployed within and sensing the field, the node deployment and sampling rate can be controlled to a certain extent, for example, the node spacing relative to the assumed field motion. However, while image-based optical flow

approaches usually rely on regular grids (images), a grid-like deployment of the nodes might not be possible, resulting in irregular sensor neighborhoods.

Decentralized estimation: Image-based OF usually relies on a single computer that estimates the motion field from the images, as e.g. done with weather radar. In a similar way, the motion estimation using sensor data collected by a GSN could be performed by a central node, e.g. a sink node, where all data is routed to and integrated. In this work, a decentralized algorithm is proposed where no central node for data collection and processing exists but every node holds a current motion estimate and communicates only locally with neighbors. Although a centralized solution integrating all data always outperforms decentralized solutions in terms of accuracy, the latter has certain benefits. Duckham (2012) lists the following ones: *information overload*, *scalability*, *communication latency* and *information and location privacy*. With thousands or even millions of nodes, a decentralized approach helps managing the huge amounts of data generated by the nodes and makes the whole process independent of the number of nodes: it allows adding and removing nodes on-the-fly, since the whole system behavior is solely specified by node interactions and thus, micro-controller program code, which can remain unchanged (Duckham, 2012). Further, the *communication latency* can be expected to be decreased by decentralization, especially if the information is used in a *sensor/actuator setting* close the spatial location of where it is generated. An example is the application of the proposed algorithm for precipitation field motion estimation by Vehicular Ad-hoc NETWORKS (VANETs), where the information on the motion of strong rainfall is important for the cars in the area, e.g. to initiate warning messages on approaching strong rainfall. Especially in such vehicular networks, *information privacy* and *location privacy* is an important topic. Communication within the car’s neighborhood only can help protecting both (Duckham, 2012).

The proposed algorithm accounts for all these differences to standard optical flow algorithms. An overview of the algorithm is provided in Section 6.1, the algorithm itself is then described in the following subsections.

6.1 Algorithm overview

Figure 6.1 displays the process carried out by each node, assuming that the bidirectional node communication in between neighboring nodes is already established.

As stationary nodes are assumed, a particular node broadcasts its spatial location to the neighbors and collects all spatial locations of the neighbors at the initialization phase of the algorithm. Then, at each time step, there are two rounds of communication. In the first round, each node broadcasts its current measurement of the field and collects all measurements of neighboring nodes. When finished, a node estimates a gradient constraint of the form of Equation 2.39. In a second round of communication, a node broadcasts the gradient constraint and collects all gradient constraints of neighbors. Finally, the node estimates the motion. Thus, besides communication, the proposed decentralized algorithm for motion estimation consists of two distinct computations:




Figure 6.1: Flow chart for the motion estimation carried out by each (stationary) node. The flow chart is a visual illustration of protocol 2.

- 1.) **Compute gradient constraint:** Estimates of a gradient constraint at a node are computed from sensor measurements of neighboring nodes. The error of a GC is derived from the spatial configuration of the node neighborhood. Details are given in Section 6.3.
- 2.) **Estimate motion:** A set of estimates of gradient constraints is integrated by each node over its direct 1-hop neighborhood to solve for the motion components. Further details are provided in Section 6.4

For centralized motion estimation, either all measurements are directly routed to a central node integrating all data. Or, the gradient constraints generated after the first round of communication could be send to the central node. In this case, the node that integrates all data could use standard algorithms for computing motion from gradient constraints, e.g. those provided by Horn & Schunck (1981) or Lucas et al. (1981).

The second step of the algorithm carried out after the second round of communication is concerned with the decentralized motion estimation by each node. The algorithm, and especially the second step of GC integration builds on the image-based OF method presented in Lucas et al. (1981) in that information is integrated over the neighborhood of a node (resp. *pixel* in Lucas et al. (1981)). However, instead of integrating the gradient constraints centrally using a single computer as done in Lucas et al. (1981), in this work, every node integrates the gradient constraints and estimates the motion. Nonetheless, the assumption of Lucas et al. (1981) that the motion is translational within the node (pixel) neighborhood still applies. Possible extensions to this assumption are known and e.g. described in Black (1992). The spatial irregularity of data is accounted for by a probabilistic error model. The motion estimation from gradient constraints is then formalized as a regression model. In order to account for motion coherence over sampling periods and to allow efficient processing, a recursive formalization of the regression equations is used in the form of a Kalman filter (Kalman, 1960), which has been introduced previously in Section 2.3.4. The proposed algorithm requires the specification of only three parameters with unambiguous interpretations: the maximum communication distance in between the nodes, which is, in a real deployment of a GSN,

determined by the physical hardware. Further parameters are the Kalman filter prediction and measurement noise covariances.

6.2 Network and field model

A GSN is modeled as a graph $G = (V, E)$ where V is the set of nodes distributed in the plane and E is the set of communication links between nodes. The allowed bidirectional communication links are solely determined by a maximum Euclidean communication distance r (unit distance) in the plane, and hence, G is a unit disk graph (UDG, Duckham (2012)). A node $n_i \in V$ knows its position $s_i = (x_i, y_i)$ on the plane, e.g., by using the Global Positioning System (GPS). Further, the nodes are able to sense a real-valued scalar spatio-temporal field $Z(u) : \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}$ where $u = [x, y, t]^T$ is a location in the space-time cube (and $[]^T$ indicates the matrix transposition). A particular sensor measurement of node n_i at time step t is denoted with $z(u_{i,t})$ where $u_{i,t} = [x_i, y_i, t]^T$. Partial derivatives of the field along spatial axes X and Y and temporal axis T at a particular spatio-temporal location u are written as $Z'_X(u)$, $Z'_Y(u)$ and $Z'_T(u)$. Their estimates provided by a node n_i at time t are denoted with $\hat{z}'_X(u_{i,t})$, $\hat{z}'_Y(u_{i,t})$ and $\hat{z}'_T(u_{i,t})$. The column vector of estimated partial derivatives is then written as $\hat{z}'(u_{i,t}) = [\hat{z}'_X(u_{i,t}), \hat{z}'_Y(u_{i,t}), \hat{z}'_T(u_{i,t})]^T$.

6.3 Gradient constraint estimation in the network

In this section, the methodology for the estimation of optical flow gradient constraints by each node from neighboring node measurements is described. First, the equations for the estimation of partial from directional derivatives are provided (6.3.1). Then, in Section 6.3.2, the requirements of node stationarity and sampling synchronicity are discussed. Subsection 6.3.3 provides the methodology for the derivation of error measures for the estimated partial derivatives and Section 6.3.4 describes, how gradient constraints for motion estimation are selected and how a single scalar gradient constraint error measure is derived to be used later in the motion estimation method provided in Section 6.4.

6.3.1 Gradient constraint estimation from irregular data

When the data is irregularly sampled in space and time, the estimation of the partial derivatives directly is impossible, as there might be no samples aligned along the coordinate axes (Figure 6.2).

In addition, the sampling in a GSN might not be synchronized such that the neighbors sample the field at a different rate, e.g. some every five minutes, some others every minute. Then, the samples differ not only in space but also in time. Thus, this completely irregular sampling of the data in space and time requires the indirect estimation of $Z'_X(u)$, $Z'_Y(u)$ and $Z'_T(u)$ by numerical estimates of *directional derivatives* in the space-time cube. A directional derivative along a particular direction d in the space-time cube is defined as:

$$Z'_d(u) = \lim_{d \rightarrow 0} \frac{z(u + d) - z(u)}{|d|} \quad (6.1)$$




Figure 6.2: In images, the partial spatial derivatives at a pixel location (red) can be estimated from neighboring pixel values, e.g. via central differencing (left). In GSN, the data might be irregularly distributed such that the direct estimation at a node location (red) is impossible (right).

An estimate of $Z'_d(u)$ can then be calculated from sensor samples of the field (Equation (6.2)).

$$\hat{z}'_{d_{ij,tr}}(u_{i,t}) = \frac{z(u_{j,r}) - z(u_{i,t})}{|d_{ij,tr}|} \quad (6.2)$$

where $d_{ij,tr}$ is the spatio-temporal distance vector between the spatio-temporal node locations $u_{i,t}$ and $u_{j,r}$ (the locations of node i at time t and node j at time r), and $z(u_{i,t})$ and $z(u_{j,r})$ are field samples at these locations, collected by the nodes. The functional relationship between the estimate of such a directional derivative and the partial derivatives along the space-time cube coordinate axes is displayed in Equation (6.3) (for easing the readability, the spatio-temporal location index $u_{i,t}$ is omitted in the following equation).

$$\hat{z}'_d = \hat{z}'_X \hat{u}_{X,d} + \hat{z}'_Y \hat{u}_{Y,d} + \hat{z}'_T \hat{u}_{T,d} \quad (6.3)$$

where $\hat{u}_{X,d}$, $\hat{u}_{Y,d}$ and $\hat{u}_{T,d}$ are the unit vector components in the particular spatio-temporal direction d , \hat{z}'_d is a particular estimate of a directional derivative calculated with Equation (6.2) and \hat{z}'_X , \hat{z}'_Y and \hat{z}'_T are estimates of the partial derivatives required for optical flow. A set of such linear equations available at a specific spatio-temporal node position $u_{i,t}$ then forms a linear system of the form of Equation 6.4.

$$A(u_{i,t}) \hat{z}'(u_{i,t}) = \hat{b}(u_{i,t}) \quad (6.4)$$

where $A(u_{i,t})$ is the matrix containing the unit vectors as rows, $\hat{z}'(u_{i,t}) = [\hat{z}'_X(u_{i,t}), \hat{z}'_Y(u_{i,t}), \hat{z}'_T(u_{i,t})]^T$ is the column vector of partial derivatives to be estimated, i.e., $\hat{z}'(u_{i,t})$ is the unknown estimate of the gradient constraint of Equation 2.39. $\hat{b}(u_{i,t})$ is the known column vector of estimates of directional derivatives. The system then allows solving for $\hat{z}'(u_{i,t})$. An approximate solution to the equation is given by the least-squares estimator displayed in Equation 6.5.

$$\arg \min_{\hat{z}'(u_{i,t})} \|A(u_{i,t}) \hat{z}'(u_{i,t}) - \hat{b}(u_{i,t})\|^2 \quad (6.5)$$

where $\|\cdot\|$ is the Euclidean norm. The minimization could then be performed using different methods. Here, it is performed by solving the normal equations for the vector of partial derivatives

$\hat{z}'(u_{i,t})$, including a weight matrix W , as displayed in Equation 6.6 (again, spatio-temporal locations are omitted for easing readability).

$$\hat{z}' = (A^T W A)^{-1} A^T W \hat{b} \quad (6.6)$$

The weight matrix W contains a weight for each estimate of a directional derivative. Certainly, the weight should be a function of field properties, e.g. the degree of field linearity at the site of estimation. Further, it should be influenced by the distance $|d|$ between the pair of sensor samples $z(u_{j,r})$ and $z(u_{i,t})$ used for estimation, as the accuracy of numerical estimates of derivatives decreases with increasing distance.

6.3.2 Requirements on node stationarity and sampling synchronicity

As the distance vector $d_{ij,tr}$ between any two spatio-temporal sample locations is a vector in space and time, the calculation of the length $|d_{ij,tr}|$ that is needed in Equation 6.2 requires the a priori specification of a spatio-temporal anisotropy factor, such as a decision on the unit of measures. For example, it requires a decision on whether the temporal unit of *minutes* is suitable, when the spatial unit is *kilometres*, or if *seconds* or *hours* should be used. This factor is certainly driven by domain knowledge on the behavior of the phenomenon in space and time. If this is available a priori, e.g. by statistical investigations, the nodes might be allowed to move and sample the field asynchronously, i.e., at different time steps. In this work, the underlying design decision is that the nodes are stationary and sample the field at equal time steps. Then, the anisotropy factor is not required anymore. In Figure 6.3, a visualization of this assumption is displayed.




Figure 6.3: Assumption of node stationarity and sampling synchronicity. Red dots represent field samples provided by the sensors (measurements), which are taken at fixed spatial sensor positions at synchronized time steps / periods.

In this case, space can be treated separated from time and there are always other sensor samples of the same time step available, although they might not always be accessible to a node, e.g. when a node does not have any neighbors in the communication graph. The linear system then only includes estimates of spatial directional derivatives and the distance $d_{ij,tr}$ becomes spatial distance d_{ij} between nodes n_i and n_j only. Due to stationarity of nodes, temporal change can be estimated by each node individually from the difference of the current and previous field sample

measured by that node: $\hat{z}'_T(u_{i,t}) = z(u_{i,t}) - z(u_{i,t-1})$. With this fixed spatial configuration, the matrix $A(u_{i,t})$ in the linear system of Equation 6.4 is time-independent and constant for each node n_i and is therefore denoted with A_i in the following. When the weight matrix W is constant as well, the normal equations for solving the system (Equation 6.6) can be pre-computed at the initialization phase of the algorithm (Section 6.5). It is clear that this assumption does not hold for non-stationary, moving nodes, e.g., cars. Nonetheless, in Section 7.3, the algorithm is applied in a non-stationary setting with simulated cars, with only slight modifications detailed in Section 6.5.

6.3.3 Estimation of partial derivative error

Due to node stationarity, the change along the time axis can be directly estimated by each node by the so-called backward differencing: $\hat{z}'_T(u_{i,t}) = z(u_{i,t}) - z(u_{i,t-1})$. Therefore, no error for the estimated temporal derivative $\hat{z}'_T(u_{i,t})$ is assumed. For estimating the error and therefore, the weights of matrix W in the WLS adjustment, associated with the estimate of a particular directional derivative in space, the heuristic displayed in Equation (6.7) is used:

$$\epsilon_{ij} = \frac{d_{ij}}{r} \quad (6.7)$$

where r is the maximum possible communication range in between the nodes and d_{ij} is spatial distance between the two nodes n_i and n_j . From the definition of directional derivatives provided in Equation 6.1, it becomes clear that the closer two measurements are, the better does the numerical differentiation of Equation (6.2) approximate the true directional derivative. This particular weight of Equation 6.7 accounts for this by down-weighting measurements with increasing spatial distance from the estimating node. Dividing distance by communication range ensures that the down-weighting is significant and independent of the unit of measure of spatial distance. The usefulness of this weighting compared to a uniform weighting is shown in the evaluation provided in Section 7.3.4.

The weight matrix for a particular node n_i , W_i , in the WLS adjustment then contains the inverses of ϵ_{ij} for each neighboring node n_j as the weight on the diagonal and zeros elsewhere, as no dependency between the observations is assumed. Due to stationarity of nodes, W_i is constant for each node and can also be precomputed.

Derivation of partial derivative error via variance propagation

The error associated with the derived partial derivatives can then be estimated from individual directional derivative errors using the law of propagation of error (see, e.g., Langley (1999); Niemeier (2002) or any geodetic textbook):

$$C_{GC_i} = (A_i^T W_i A_i)^{-1} \quad (6.8)$$

where C_{GC_i} is the 2×2 so-called cofactor matrix of the 2×1 vector of estimated spatial partial derivatives $\hat{z}'(u_{i,t}) = [\hat{z}'_X(u_{i,t}), \hat{z}'_Y(u_{i,t})]^T$. Since the nodes are assumed to be stationary, matrix A_i and weight matrix W_i are constant and hence C_{GC_i} is constant for a particular node n_i and can be pre-computed as well (Section 6.5). Since field properties have been ignored in Equation (6.8),

C_{GC_i} is not a proper covariance matrix, and the diagonal entries of C_{GC_i} are not absolute, but only relative representations of the errors of the estimated partial derivatives $\hat{z}'_X(u_{i,t})$ and $\hat{z}'_Y(u_{i,t})$ estimated by that node (Niemeier, 2002).

Therefore, in order to derive a valid covariance matrix for each $\hat{z}'(u_{i,t})$ provided by node n_i , the entries of C_{GC_i} have to be transformed to the proper level of error for the specific field under consideration and for the specific time step t . A common approach in least squares adjustment is to derive a so-called variance-factor from the least squares adjustment and to multiply it with C_{GC_i} to derive the *empirical covariance matrix* $\Sigma_{\hat{z}'}(u_{i,t})$ for the estimated partial derivatives $\hat{z}'(u_{i,t})$ at node n_i at time t :

$$\Sigma_{\hat{z}'}(u_{i,t}) = \hat{\sigma}^2(u_{i,t})C_{GC_i} \quad (6.9)$$

The required variance factor $\hat{\sigma}^2(u_{i,t})$ can be calculated with:

$$\hat{\sigma}^2(u_{i,t}) = \frac{v(u_{i,t})^T W_i v(u_{i,t})}{f} \quad (6.10)$$

where $v(u_{i,t})$ is the column vector of deviations of the least squares adjustment for node n_i at time t . It contains the result of $\hat{z}'_X \hat{u}_{X,d} + \hat{z}'_Y \hat{u}_{Y,d} - \hat{z}'_d$ as row entries for each directional derivative/neighbor node at time t . $f = m - p$ is the degrees-of-freedom with m being the number of neighbors of a node and p the number of parameters ($p = 2$ for the stationary case).

Calculating the variance factor is a common approach in least squares adjustment, e.g., for GPS positioning (Langley, 1999), in order to quantify the error of the derived parameters/coefficients; in this case, the error of the least squares result $\hat{z}'(u_{i,t})$. However, calculating a stable $\hat{\sigma}^2(u_{i,t})$ at each time step t usually requires sufficient degrees-of-freedom in the order of p to $2p$ (Niemeier, 2002). Thus, for a sufficiently large f , a node is required to have 4 to 6 (resp. 6 to 9, when temporal derivatives are included into the system) neighbors. Therefore, in this work, an a priori variance factor is used in the following, which is equal to all nodes and denoted with $\hat{\sigma}^2$, which is, however, derived from observed values for $\hat{\sigma}^2(u_{i,t})$. A particular advantage of using a pre-calculated variance factor is that in a stationary setting, the variance factor $\hat{\sigma}^2(u_{i,t})$ for a node n_i is constant and therefore, the covariance matrix of the partial derivative calculated by the node $\Sigma_{\hat{z}'}(u_{i,t})$ is constant as well (written as $\Sigma_{\hat{z}',i}$ in the following) and can also be pre-computed (see Appendix 6.5).

6.3.4 Gradient constraint selection and derivation of gradient constraint error

The methodology described in the previous sections delivers gradient constraint estimates $\hat{z}'_X(u_{i,t})$, $\hat{z}'_Y(u_{i,t})$, $\hat{z}'_T(u_{i,t})$ of the true gradient constraints $Z'_X(u_{i,t})$, $Z'_Y(u_{i,t})$, $Z'_T(u_{i,t})$ for each node n_i and time step t together with error estimates in the form of the covariance matrix $\Sigma_{\hat{z}'}(u_{i,t})$ containing the variances of the estimates at the diagonal.

However, the motion estimation methodology described in Section 6.4 requires an estimate of the error associated with a particular estimate of a GC, $\hat{z}'(u_{i,t})$, in the form of a **scalar** error variance. In order to derive such a measure for a GC, the probabilistic approach of Simoncelli (1993) is used, who derives a probability distribution for the motion components from the gradient constraint.

Simoncelli (1993) argues that the gradient constraint of Equation (2.39) is an idealization for several reasons: the true spatial and temporal derivatives are not available, but their estimates are. Further, the gradient constraint is a constraint on the (optical) flow and not on the true motion of the field. For example, when using gradient constraints for estimating the motion of a precipitation field, the motion can be determined only at locations where there is some precipitation. In those regions where there is no rainfall, the motion estimate will be zero, as the spatial and temporal derivatives are zero, although there might still be motion in the atmosphere. Since the optical flow gradient constraints are usually integrated in space, a spatial neighborhood that includes rainfall and no rainfall will exhibit different 'optical motions' and, therefore, provides an averaged motion estimate (see also Section 1.2 of Simoncelli (1993), where this is called the 'blank-wall' problem). Therefore, Simoncelli et al. (1991) introduced a probabilistic model for the gradient constraint (further details can be found in the doctoral thesis Simoncelli (1993)) including three possible sources of error, each characterized by Gaussian random noise variables. Here, the model of Simoncelli (1993) is used, but it is assumed that the estimation of the partial derivatives in space is the dominant source of error and that there is no error in the estimated temporal derivative.

Gradient constraint selection

Before the approach of Simoncelli (1993) is used to compute the error of an estimated gradient constraint, certain configurations of measurements prevent a node from estimating a gradient constraint at all:

- 1.) **Zero-field values:** When the field values used for derivative calculation are zero and, hence, all derivatives are zero, the GC is not estimated at all. Often, measurements of spatio-temporal fields are zero-inflated, meaning that the majority of samples are zero. In such a case, no motion can be estimated, and node energy can be saved.
- 2.) **Node neighborhood extending field boundary:** When only one of the sensor samples for estimating the GC is zero, the GC is not estimated at all, as the node neighborhood extends over the field boundaries. While the GC could still be estimated using the remaining sensor samples, this is not done, as the least squares matrices for estimating the partial derivatives are pre-computed (Section 6.5).
- 3.) **Non-zero, but equal field values:** When the field is completely flat, corresponding to the 'blank wall' problem described previously, the GC is not estimated at all. This case can be recognized, when all of the neighboring sensor samples of a time step t are larger than zero, but equal. Then, the derivatives are zero, and the GC does not contribute to the motion estimation.

Gradient constraint error derivation

For deriving a gradient constraint error measure for the remaining gradient constraints, the approach of Simoncelli (1993) is used: it is assumed that the true but unknown spatial derivatives are

related to the estimated spatial derivatives via some random noise variables n_{z_X} and n_{z_Y} (again, the spatio-temporal location parameter $u_{i,t}$ is skipped for easing readability).

$$\begin{aligned}\hat{z}'_X &= Z'_X + n_{z'_X} \\ \hat{z}'_Y &= Z'_Y + n_{z'_Y}\end{aligned}$$

where \hat{z}'_X and \hat{z}'_Y are the estimated partial derivatives in X and Y direction, Z'_X and Z'_Y are the true derivatives and $n_{z'_X}$ and $n_{z'_Y}$ are the noise variables. Since no error for the estimate of the temporal derivative is assumed, a GC (Equation (2.39)) is then reformulated as follows, following Simoncelli (1993):

$$\begin{aligned}0 &= Z'_X v_X + Z'_Y v_Y + Z'_T \\ &= (\hat{z}'_X - n_{z'_X}) v_X + (\hat{z}'_Y - n_{z'_Y}) v_Y + \hat{z}'_T \\ &= \hat{z}'_X v_X - n_{z'_X} v_X + \hat{z}'_Y v_Y - n_{z'_Y} v_Y + \hat{z}'_T\end{aligned}$$

From this, the probabilistic relationship between the estimates of the partial derivatives and the true motion that is to be estimated can be derived (Simoncelli, 1993):

$$\Rightarrow \hat{z}'_X v_X + \hat{z}'_Y v_Y + \hat{z}'_T = n_{z'_X} v_X + n_{z'_Y} v_Y \quad (6.11)$$

Under the assumption of zero-mean independent Gaussian noises, the right-hand term of Equation (6.11) is a zero-mean random variable with variance:

$$\sigma_{GC}^2 = \sigma_{z'_X}^2 v_X^2 + \sigma_{z'_Y}^2 v_Y^2 \quad (6.12)$$

where $\sigma_{z'_X}^2$ and $\sigma_{z'_Y}^2$ are the variances of $n_{z'_X}$ and $n_{z'_Y}$, resp., and v_X and v_Y are the field motion variables in both spatial directions. The statistical independence of $n_{z'_X}$ and $n_{z'_Y}$ cannot be expected to hold in reality, but as no information on the dependence (i.e., covariance) can be made, this is a necessary simplification.

The relationship of Equation (6.12) makes intuitive sense, as a gradient constraint can be understood as a regression equation without intercept, the coefficients v_X and v_Y , the independent variables $-Z'_X$ and $-Z'_Y$ and the dependent variable Z'_T . This is also the Kalman filter formalization described in the next section. The tuples $(\hat{z}'_X, \hat{z}'_Y, \hat{z}'_T)$ are used for calibrating the coefficients. Then, it is clear that the smaller the regression coefficients v_X and v_Y , the smaller the influence of $\sigma_{z'_X}^2$ and $\sigma_{z'_Y}^2$, disappearing completely for a flat regression plane, i.e., $v_X = v_Y = 0$. The terms $\sigma_{z'_X}^2$ and $\sigma_{z'_Y}^2$ in Equation (6.12) can be derived from the covariance matrix of the vector of partial derivatives (the diagonal entries of $\Sigma_{\hat{z}'}$ of Equation (6.9)). The terms v_X and v_Y are unknown and have to be predefined and are therefore a parameter of the algorithm. In this case, the assumed error of a gradient constraint provided by a node n_i is constant and solely determined by spatial configuration. Therefore, it can be precomputed at the initialization phase of the algorithm (Section 6.5) and is denoted with $\sigma_{GC_i}^2$ in the following.

In case of negative $n_{z'_X}$ and $n_{z'_Y}$, the estimated spatial derivatives are underestimates of the true spatial derivatives. Then, the motion estimates will overestimate the true motion speed. This can be explained by means of a simple example: with a fixed temporal absolute derivative and low absolute spatial derivatives, the field must move a larger distance in space in order to render the GC equation valid, i.e., the absolute values of v_X and v_Y must be larger. Conversely, larger absolute spatial derivatives result in smaller speed.

6.4 Temporal coherence: Kalman filter for recursive motion estimation

The gradient constraints can be recast into the linear regression form. Then, the problem of motion estimation can be considered as the calibration of a multivariate, intercept-free linear regression model with the motion components in both directions as the coefficients, the spatial partial derivatives as the regressors (independent variables) and the partial derivative in time as the random response variable (Equation 6.13):

$$E(Z'_T | Z'_X = z'_X, Z'_Y = z'_Y) = -z'_X v_X - z'_Y v_Y \quad (6.13)$$

Analogous to Equation 2.5, for particular tuples of realizations of the partial derivatives, here, the estimates \hat{z}'_X , \hat{z}'_Y and \hat{z}'_T computed previously, the regression equation has to be reformulated as follows.

$$\hat{z}'_T = -\hat{z}'_X v_X - \hat{z}'_Y v_Y + \epsilon \quad (6.14)$$

With this formulation, the linear regression assumptions apply, namely that $\epsilon \sim NID(0, \sigma^2)$ and that the partial derivatives z'_X and z'_Y are *deterministic* instead of random variables.

However, here, the response variable is assumed to be error-free, while the regressors are subject to error (see Section 6.3.4). Therefore, for a statistically sound methodology, estimation methods accounting for these errors in the data would be required, for example errors-in-variables regression models, such as total least squares. However, we leave the investigation of such estimation methods to future work and use standard ordinary least squares (OLS) for solving the regression model. For calibration with GCs collected over space and time, a decision is required on which GCs are integrated to estimate motion at a particular node n_i and time step t . Here, it is assumed that within the one-hop node neighborhood, the motion is uniform. Therefore, a node integrates the GCs provided by its direct neighbors. For integrating GCs in time, a recursive least-squares estimation method is used in the form of a Kalman filter (Kalman (1960), see also Särkkä (2013) or Section 2.3.4 for an application of Kalman filtering to regression problems).

The Kalman filter has several advantages over least-squares methods for time series data, as it does not require the storage of large amounts of past data, reduces the computational costs when updating the regression coefficients and comes in a predict-and-update formalization that fits well to the problem. The Kalman filter requires an initial decision on the state variables. The most simple form for the problem at hand is a Kalman state with only the two motion components. In this case, motion is considered constant, and motion change is solely modeled by setting the prediction error variance Q larger than zero. Including motion change variables in the form of first derivatives of motion with respect to time into the Kalman state is slightly more realistic

by assuming motion change constancy. In this case, the Kalman state $v_i(t)$ including motion derivatives at a node i and time step t is:

$$v_i(t) = \begin{bmatrix} v_X \\ v_Y \\ \dot{v}_X \\ \dot{v}_Y \end{bmatrix} \quad (6.15)$$

where v_X and v_Y are the motion vector components in the directions X and Y , resp. \dot{v}_X and \dot{v}_Y are the motion derivatives, i.e., acceleration. The prediction matrix F is:

$$F = \begin{bmatrix} 1 & 0 & \Delta t & 0 \\ 0 & 1 & 0 & \Delta t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (6.16)$$

where $\Delta t = t - (t - 1)$, i.e., the difference in between filter steps. In the regression formalization, the temporal derivative $Z'_T(t)$ is considered a linear function of the motion components and the partial derivatives in space. The Kalman filter measurement matrix then contains the current partial derivatives in space (Equation 6.17).

$$H_i(t) = \begin{bmatrix} -Z'_X(t) & -Z'_Y(t) & 0 & 0 \end{bmatrix} \quad (6.17)$$

This way, the motion state of a single node n_i is updated with new GCs, which can be derived from the node n_i itself or neighboring nodes. Since the measurement $Z'_T(t)$ is scalar, the measurement error variance R is scalar, and therefore, solving the Kalman update equations does not require matrix inversion (Särkkä, 2013), which is advantageous for the amount of processing required. The specification of an a priori state is required, which can only be set to the zero vector. The a priori uncertainty, the initial P , should contain large values on the diagonal, indicating the low confidence in the initial state.

6.4.1 Estimation of process noise Q

The uncertainty associated with prediction Q depends on the temporal sampling rate, i.e. the time difference between Kalman filter steps, as well as the assumed motion and motion change constancy of the spatio-temporal field. Therefore, domain knowledge could be employed for setting the required variances. For example, when sampling an atmospheric field such as precipitation at a sampling rate of 1-minute, the difference in motion in between adjacent 1-minute time steps can be expected to be in a very low km/h range.

6.4.2 Estimation of measurement noise R

The scalar measurement noise variance R is a direct function of GC accuracy. GC accuracy in turn depends on the accuracy of the estimated derivatives and the spatial distance between the node

estimating the motion and the node estimating the GC. The assumption is: the closer the two nodes are, the more similar is the motion. In addition, as described in work on image-based optical flow, such as Lucas et al. (1981), the GC accuracy depends on field properties at the site where it is constructed. However, in this work and due to the irregularity of data, the main source of error is considered to be the partial derivative error, and therefore, the measurement noise variance R is solely calculated from the derived error measures. Thus, it is calculated using Equation (6.12) with a pre-computed variance factor.

6.4.3 Difference to common Kalman filtering problems

In most Kalman filter applications such as the position estimation of moving vehicles, there are more or less regular measurement updates, e.g. by GPS. Here, there are no updates when there are no field values larger zero, which, depending on the field, could happen rather often. Thus, there is the risk of filter deviation, when no measurement updates occur, especially since the filter state encompasses motion change variables, which are applied to the motion variables at every time step via the dynamic model. Deviation could especially happen when the field is composed of a set of distinct regions with zero field values in between, for example precipitation clouds in rain events. Then, filter updates are only computed when the node and its neighborhood are under such a cloud. If an area of no values, such as in between clouds, appears over the node at a stage where the state variables show large errors, these errors result in potential strong deviations before the next update occurs. Thus, when evaluating the algorithm in Section 7.3, the initial filter error for the acceleration variables is set to zero as well as the prediction error for the acceleration variables. Thus, the filter holds the motion constant, when no updates occur.

6.5 Algorithm protocol

This section introduces the algorithm protocol using the structure and formalism of Duckham (2012). The protocol provides the pseudo-code for the flow chart of Figure 6.1. As the base case, stationary nodes are assumed that sample the field at synchronized time periods. Further, the nodes are assumed to communicate in a synchronous way (Schulz, 2007), such that after the initialization phase (state INIT), there are rounds of communication where each node sends a message first, followed by a processing step and another message exchange (Protocol 2, also Figure 6.1).

Protocol: Field Motion Estimation - Stationary Nodes

Restrictions: Graph $G = (V, E)$. Nodes $n_i \in V$ with constant locations and comm.links E .
Function $nbr : V \rightarrow V$ returning set of neighbors of a node
Function $nbr_{>1} : V \rightarrow V$ returning set of neighbors with > 1 neighbors themselves

Init.: All nodes in state INIT

Data: Each node implements the Kalman filter
Each node stores the normal equations M_i
Each node stores the errors associated with the gradient constraints $\sigma_{GC_j}^2$ of all neighboring nodes $n_j \in nbr_{>1}(n_i)$ and itself $\sigma_{GC_i}^2$ in a dictionary D_σ mapping node IDs to GC error variances

Parameters: Communication range r , Kalman prediction noise parameter Q , variance factor $\hat{\sigma}^2$, field motion parameters v_X and v_Y for Equation 6.12.

INIT

```

broadcast  $(x_i, y_i)$ 
if  $|nbr(n_i)| < 2$ 
  become PROCESSING
Receiving neighboring node position  $(x_j, y_j)$ 
if  $|nbr(n_i)| > 1$ 
  compute  $d_{ij}$  and unit vector  $\hat{u}_{ij}$ 
  add  $\hat{u}_{ij}$  as new row to  $A_i$ 
  add  $r/|d_{ij}|$  as new diagonal entry to  $W_i$ 
  if  $|rows(A_i)| = |nbr(n_i)|$ 
    compute normal equations  $M_i = (A_i^T W_i A_i)^{-1} A_i^T W_i$ 
    compute  $\sigma_{GC_i}^2$  (Equation (6.12))
    add  $(i, \sigma_{GC_i}^2)$  to  $D_\sigma$ 
    broadcast  $\sigma_{GC_i}^2$  to  $nbr_{>1}(n_i)$ 
Receiving  $\sigma_{GC_j}^2$  of neighboring node  $n_j$ 
  add  $(j, \sigma_{GC_j}^2)$  to  $D_\sigma$ 
  if  $|D_\sigma| = |nbr_{>1}(n_i)| + 1$ 
    become PROCESSING

```

PROCESSING at time step t

```

(1) Whenever new sample  $z(u_{i,t})$  is available
  if  $z(u_{i,t}) > 0$ 
    broadcast  $(t, z(u_{i,t}))$ 
    if  $|nbr(n_i)| > 1$ 
      compute Kalman prediction step
      compute current temporal derivative  $\hat{z}'_T(u_{i,t}) = z(u_{i,t}) - z(u_{i,t-1})$ 
(2) Receiving sample  $(t, z(u_{j,t}))$  (with  $z(u_{j,t}) > 0$ ) from neighbor  $n_j$ 
  if  $|nbr(n_i)| > 1$  and own measurement  $z(u_{i,t}) > 0$ 
    compute estimate of directional derivative  $\hat{z}'_{d_{ij}}(u_{i,t}) = (z(u_{j,t}) - z(u_{i,t})) / |d_{ij}|$ 
    add  $\hat{z}'_{d_{ij}}(u_{i,t})$  as new entry to  $\hat{b}(u_{i,t})$ 
    if  $|rows(\hat{b}(u_{i,t}))| = |nbr(n_i)|$ 
      compute  $\hat{z}'(u_{i,t}) = M_i \times \hat{b}(u_{i,t})$ 
      compute Kalman update with  $\hat{z}'(u_{i,t})$  and stored  $\sigma_{GC_i}^2$ 
      broadcast  $(t, \hat{z}'(u_{i,t}))$  to  $nbr_{>1}(n_i)$ 
(3) Receiving gradient constraint  $(t, \hat{z}'(u_{j,t}))$  from neighbor  $n_j$ 
  if  $|nbr(n_i)| > 1$ 
    compute Kalman update with  $\hat{z}'(u_{j,t})$  and stored  $\sigma_{GC_j}^2$ 

```

In the INIT step, a node n_i distributes its position to all neighbors. If it has zero or only a single neighbor, it is not able to estimate partial derivatives and, hence, proceeds to state **PROCESSING**. Otherwise, it waits for receiving the positions of all neighbors. If it has done so, the matrix M_i can be computed and the error associated with the node can be stored and broadcast. When the node receives the GC error $\sigma_{GC_j}^2$ of neighboring node n_j , it is stored. When all neighbors provided their GC error measures, the node is ready for motion estimation, i.e., proceeds to state **PROCESSING**. There, at each new sampling step (1) of a node, a Kalman prediction on the motion is performed and the temporal derivative is calculated. Further, the sensor measurement is transmitted to the neighbors. When a new sample from a neighboring node arrives (2), the current directional derivative is estimated and added to the derivative vector \hat{b} . If \hat{b} is filled with all data from participating neighbors, the gradient constraint is computed, which is then used for the Kalman update and transmitted to all neighbors participating in motion estimation. When a gradient constraint from a neighboring node arrives (3), the Kalman update equations are executed, as well. The complexity of the proposed algorithm along with the complexity measures of communication complexity/node count, load balance (Duckham, 2012), as well as computational complexity in terms of floating point operations (Golub & Loan, 1996) are provided in Section 6.6.

For non-stationary nodes, such as cars, an initial broadcast of the node position is not sufficient, and hence, the pre-computation of the matrices for derivative calculation at the INIT state is not possible. Instead, the protocol is adjusted so that the position broadcast and all computations are performed at each time step.

6.6 Algorithm complexity

This section contains a description of the complexity of the proposed decentralized algorithm of Section 6.5. The analysis of communicational complexity and load balance follows those of Duckham (2012). The analysis of computational complexity uses the notion of *flops* (floating point operations) as e.g. described in Golub & Loan (1996).

6.6.1 Communication complexity

In terms of communication complexity, the algorithm requires an initial broadcast of the sensor positions and the error associated with the gradient constraints provided by a sensor. Therefore, $|V| + |V'|$ messages are sent at the INIT state, where V is the whole set of sensor nodes and $V' \subseteq V$ is the subset of sensor nodes with at least two neighboring nodes. Then, each **PROCESSING** step t , each sensor sends the sensed value if it is larger zero, resulting in $|V_{t,>0}|$ messages, where $V_{t,>0}$ is the set of sensor node with a measurement at time step t that is larger zero. Further, each node with at least two neighboring nodes and with a measurement larger zero sends the gradient constraint. Therefore, additional $|V_{t,>0} \cap V'|$ messages are sent, resulting in a total of $|V_{t,>0}| + |V_{t,>0} \cap V'|$ messages per **PROCESSING** step t .

6.6.2 Load balance

The set of nodes with at least two neighbors V' bears the greater communication load while members of the set of all other nodes $V \setminus V'$ only send their position once at the INIT step and then only provide larger zero measurements to their neighbors at each PROCESSING step t . Therefore, the load is unevenly balanced between the set of nodes V' and $V \setminus V'$. Therefore, in a real deployment, the nodes V' could be equipped with additional resources as mentioned by Duckham (2012). Or, the network could be deployed in a way such that all nodes have at least two neighbors ($V' = V$) and thus, comparable communication load.

6.6.3 Computational complexity of partial derivative estimation

Estimating the partial derivatives in space from neighboring sensor measurements requires the computation of $M = (A^T W A)^{-1} A^T W$ once at the INIT state of the algorithm, plus multiplication with the vector of the estimates of directional spatial derivatives at each time step. Estimating the partial derivative along the temporal axis requires a single subtraction of values and no multiplications / divisions at each time step. The following sections provide the number of flops for these computations. This number mainly depends on the number of neighbors m of a particular node and is therefore given as a function of m .

Processing costs at INIT state:

Building the matrices: The estimation of a directional derivative requires computing the Euclidean distance to each neighboring node. This requires 2 subtractions, 2 multiplications and a square root. Building the unit vectors then requires another 2 divisions. In addition, building the weight matrix W requires m divisions, where m is the number of neighbors. Hence, in total $m \times 7$ flops plus a square root are required. For example, a sensor with 5 neighbors requires 35 flops and a square root computation for building the required matrices and vectors.

Computing $M = (A^T W A)^{-1} A^T W$: A is a $m \times 2$ matrix, where m is the number of neighboring sensors. W is a $m \times m$ diagonal matrix of weights. Therefore, due to W being diagonal, matrix multiplication $A^T W$ requires $2 \times m$ multiplications and no additions and results in a $2 \times m$ matrix. Subsequent multiplication with A requires additional $2 \times 2 \times m$ multiplications and $2 \times 2 \times (m - 1)$ additions. Inverting the resulting 2×2 matrix $(A^T W A)$ requires 2 multiplications and one subtraction for the determinant plus 4 divisions. Multiplication with the already computed $A^T W$ then requires $2 \times 2 \times m$ multiplications and $2 \times m$ additions and results in the $2 \times m$ matrix M to be available for partial derivative estimation at each time step.

Computing the gradient constraint error σ_{GC}^2 from the already computed $(A^T W A)^{-1}$ requires two multiplications with the previously computed variance factor (Equation 6.9), two multiplications and a single addition of Equation 6.12.

In total, the computation of M and σ_{GC}^2 requires $12m + 4(m - 1) + 12$ flops. For example, a sensor with 5 neighbors requires 88 flops for computing M and σ_{GC}^2 at the INIT state.

Processing costs at PROCESSING state:

At each time step, estimating the current temporal derivative requires a single subtraction. Estimating spatial derivatives requires multiplication of the $2 \times m$ matrix M with the directional derivative $m \times 1$ vector \hat{b} , which needs $2 \times m$ multiplications and $2 \times (m - 1)$ additions.

Thus, in total, $2m + 2(m - 1) + 1$ flops are required for derivative estimation. For example, a sensor with 5 neighbors requires 19 flops for estimating the required partial derivatives at the PROCESSING state.

6.6.4 Computational complexity of motion estimation

The computations within in the Kalman prediction and update steps are independent from the number of neighbors m : the prediction step is executed once per PROCESSING step, the update step is executed m times per time step. In the following, the computational complexity in terms of flops is discussed for the Kalman filter implementation described in Appendix 9.2.

Prediction step:

The Kalman prediction step involves the application of the Kalman prediction matrix of Equation 6.16 to the Kalman state. This requires 6 flops (2 additions and 4 multiplications). In addition, the state covariance is predicted according to Equation 9.4. This requires 32 flops (16 additions, 16 multiplications) for the term $FP_i(k-1)^+F^T$ 4 additions for adding diagonal matrix Q . In total, the Kalman prediction step requires 42 flops.

Update step:

Computing the Kalman gain (Equation 9.7): Computing $P_i(k)^-H_i(k)^T$ requires 8 multiplications and 4 additions. Subsequent multiplication with $H_i(k)$ requires 4 multiplications and three additions and results in a scalar value to be added to measurement error variance R . The resulting value is inverted, requiring a single division. Multiplying with the already computed $P_i(k)^-H_i(k)^T$ requires 4 multiplications and results in the 4×1 Kalman gain matrix K . In total, **25** flops are required for Kalman gain computation.

Measurement Update (Equation 9.5): Updating with a new measurement requires 2 multiplications and one addition for $H_i(k)\hat{v}_i(k)^-$, a single subtraction and the subsequent scalar multiplication with the Kalman gain vector, requiring 4 multiplications. The resulting vector is added to the state, requiring four additions. In total, **12** flops are required for updating the state with a gradient constraint.

Error Covariance Update (Equation 9.8): Updating the error covariance matrix is the most costly part in the computations. It requires 8 multiplications, 4 subtractions for $I - K_i(k)H_i(k)$ plus another 32 multiplications and 24 additions for the multiplication with the a-priori error covariance $P_i(k)$. In total, 68 flops are required for updating $P_i(k)^-$ to $P_i(k)^+$.

In total, the motion estimation update step requires 105 flops per neighbor. The processing costs can be significantly reduced when the Kalman state is only a 2-state vector and is not augmented with motion derivatives. Then, all involved matrices reduce from 4×4 to 2×2 matrices. However, then, possible motion changes are solely modeled by setting the prediction error to larger zero.

6.7 Summary

This section has introduced an algorithm for the decentralized estimation of the motion of spatio-temporal moving fields with a geosensor network. A well-known optical flow algorithm was adapted to the specifics of geosensor networks, namely the lack of a centralized control unit and the strong constraints on communication and computation. A specification of the algorithm and an analysis of its complexity was provided. Section 7.3 contains the results of the evaluation of the algorithm performance on different types of simulated geosensor networks and moving fields.

7 Results

This chapter provides the results for the precipitation estimation methods (Section 7.1), the car data analyses (Section 7.2) and the experiments of GSN-based motion estimation (Section 7.3). Every section begins with a description of the study area and data basis before the results of the analyses are provided.

7.1 Precipitation intensity estimation at 1-min resolution from radar and rain gauges

7.1.1 Study area and data basis

For an evaluation of the different rainfall estimation methods presented in Chapter 4, 1-min rain gauge data from the rain gauge network of the German Weather Service (Deutscher Wetterdienst, DWD) is available, consisting of 10 stations in the greater Hanover area (Figure 7.1).




Figure 7.1: Distribution of rain gauges in Hanover area (left) and Hanover radar coverage (right))

The rain gauges have an average pairwise distance of around 32 km, delivering discrete (integer) rainfall measurements in the range $[0, 99]$ and unit of measure 0.01 mm/min . The values are converted to mm/h^1 , resulting in discrete mm/h values that are multiples of 0.6. The radar data from a $70\text{km} \times 70\text{km}$ (4900 pixels) window covering the Hanover region is used for evaluation. Two month of data provide the basis for the analyzes: a typical winter month (January 2012) and a typical summer month (June 2012).

¹This is done by multiplication with $60/100 = 0.6$.

7.1.2 Performance assessment via cross-validation

For assessing the performance of each estimation method, the 1-min rain gauge measurements are used as 'ground truth', which is a common approach and e.g. done in a similar way in related works such as the one provided by Berndt et al. (2014) or the DWD RADOLAN product (Bartels et al., 2004). The estimation of performance measures is done by the method of *leave-one-out cross validation* that is carried out for each rain gauge at each time step (Figure 7.2).

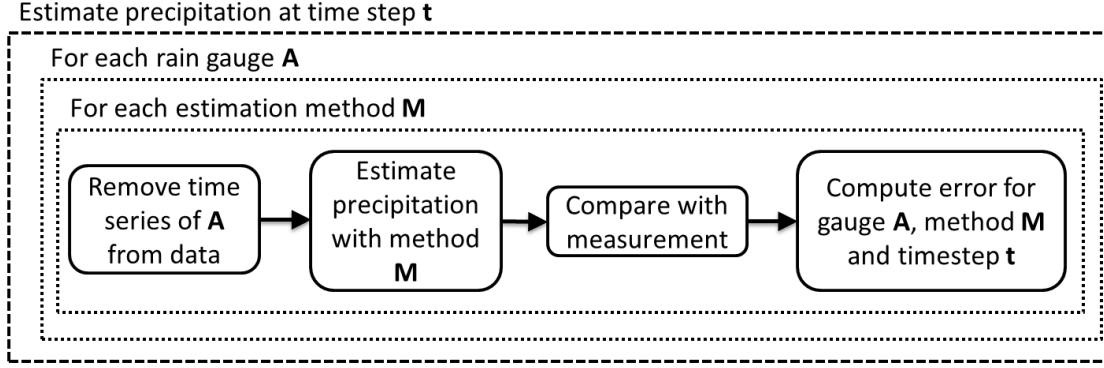


Figure 7.2: The process of cross-validation

First, the data of the whole rain gauge is removed from the data set. Then, the 1-min precipitation intensity is estimated using one of the estimation methods described in Section 4. After comparison with the original measurement, performance measures are computed for that particular gauge, time step and method.

Error measures

Following the suggestions of Haberlandt (2011), the following performance measures are applied. The coefficient of correlation is used as an important measure of statistical dependency between estimates and measurements, calculated with Equation 2.3 from pairs of measurements $z_S(u_{i,t})$ and estimates $\hat{z}(u_{i,t})$. In addition, the *bias* or *mean error* (ME) criterion is used:

$$bias \equiv \frac{1}{T \times n} \sum_{t=1}^T \sum_{i=1}^{10} [\hat{z}(u_{i,t}) - z_S(u_{i,t})] \quad (7.1)$$

where $z_S(u_{i,t})$ is the measurement of rain gauge i at time step t and $\hat{z}(u_{i,t})$ is the estimate at that location and time. Then, the *relative standard error* (RSE) (or normalized root-mean-square error, given as a percentage of the observed mean value) is used:

$$RSE \equiv \frac{100}{\bar{z}_S} \sqrt{\frac{1}{T \times 10} \sum_{t=1}^T \sum_{i=1}^{10} [\hat{z}(u_{i,t}) - z_S(u_{i,t})]^2} \quad (7.2)$$

where \bar{z}_S is the mean of rain gauge measurements for all 10 rain gauges and all time steps T of the evaluated events. As noted by Haberlandt (2011), these error measures give complementary information: the correlation coefficient delivers information on the linear dependency of the estimates and rain gauge measurements, while the *bias* is an indicator on whether a particular method tends

to under- or overestimate rainfall. In contrast, the RSA is a measure that indicates the achievable accuracy of the provided 1-min point estimates of precipitation.

7.1.3 Exploratory and visual data analysis

For both periods, only those time steps are considered where the mean rainfall estimated by radar exceeds 1 mm/h (i.e., the arithmetic mean of all 4900 pixels, including zero pixels). For characterizing rainfall within both periods based on radar data, the characterizing properties introduced by Ehret (2003) are used. These are the *Wetted Area Ratio* (WAR), defined as the percentage of radar pixels showing rainfall above 1 mm/h , the *Image Mean Flux* (IMF), defined as the mean rainfall in mm/h over all pixels and the percentage of WAR values showing rainfall above 10 mm/h (*10AR*). In addition, event characteristics are calculated from the rain gauge measurements. Table 7.1 displays the results.

Table 7.1: Characterization of the evaluation periods.

	Winter Event 1/2012	Summer Event 6/2012
1-min time steps	2745	915
Radar characteristics		
WAR (%)	31.44	16.53
IMF (mm/h)	2.22	1.95
10AR (%)	19.5	24.50
Rain gauge characteristics		
Mean intensity incl. zero measurements (mm/h)	1.3	2.0
Mean intensity non-zero measurements (mm/h)	3.19	7.98
Percentage of non-zero measurements (%)	41.05	25.2
Motion characteristics		
Mean speed ($\text{m}/5\text{min}$)	627	682
West to east mean speed ($\text{m}/5\text{min}$)	445.02	602.43
South to north mean speed ($\text{m}/5\text{min}$)	-329.42	69.06

The winter month shows a larger number of 1-min time steps with significant rainfall. The rain events in winter show lower rain intensities but are more widespread compared to the summer events, which can be recognized from the WAR values. The rain gauge analyzes confirm this: in winter, the mean rainfall is lower and the percentage of non-zero measurements is higher. The motion information is derived from the optical flow algorithm described in Section 4.² Concerning motion, the mean speed is comparable but the direction differs: in summer, the prevailing direction is west-to-east with only slight south to north motion. In winter, the precipitation fields seem to

²The OpenCV algorithm requires two images, here: two adjacent 5-min radar images, and additional parameters. They have been specified based on the visual comparison of the resulting motion fields with animated radar time series. The parameter values are (using the OpenCV parameter names): `pyr_scale=0.1`, `levels=100`, `winsize=50`, `iterations=10`, `poly_n=7`, `poly_sigma=1.5`, `flags=org.opencv.video.Video.OPTFLOW_FARNEBACK_GAUSSIAN`.

increasingly move towards east-south-east, indicated by the negative south-to-north mean speed and the lower west-to-east mean speed. The difference in intensities is confirmed in the histograms for both months displayed in Figure 7.3.

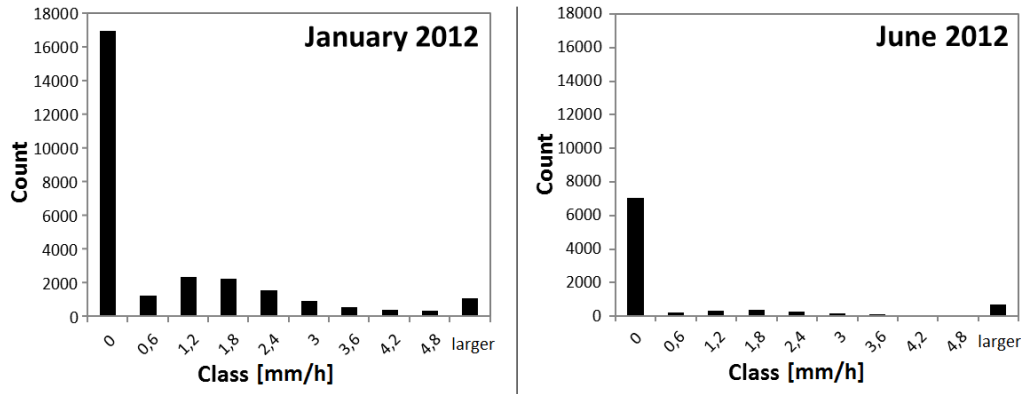


Figure 7.3: Histograms of rain gauge measurements for January 2012 (left) and June 2012 (right)

The histograms show the expected large number of zero measurements, a few high values and a peak in between. Larger sample sizes for January can be recognized.

Examples of a radar measurement and the fields estimated with the different methods are displayed in the following figures. Blue color indicates zero precipitation intensity, red color indicates high intensity (values larger 25 mm/h). Blue bars show the current 1-min rain gauge measurement, black bars the estimate of the particular method, when the rain gauge measurement or, in case of spatio-temporal interpolation, the whole time series of that rain gauge is removed from the data set used for estimation.

Figure 7.4 displays the radar measurement of a 1-min time step and the motion vectors estimated from optical flow. Due to the chosen parameters, the estimated motion vectors in the example are rather similar for all gauge locations, which corresponds well with the visual impression from animated radar time series. The figure also shows the fields interpolated with the different spatial interpolation methods Inverse-Distance Weighted (IDW) and ordinary kriging (OK) and the spatio-temporal interpolation method spatio-temporal ordinary kriging (ST-OK). All of the methods displayed estimate precipitation intensity solely from the rain gauge data.

All three estimation methods show the typical smoothing behavior of interpolation methods: a single current measurement at gauge E755 is 'distributed' in space. Figure 7.5 displays the results for the same 1-min time step and the two spatio-temporal methods 'frozen field' IDW (FF-IDW) and 'frozen field' ordinary kriging (FF-OK) that use the motion information derived from optical flow for interpolation. In addition, the results for the radar-rain gauge merging methods regression kriging (RK) and 'frozen field' regression kriging (FF-RK) are shown.

FF-IDW and FF-OK show similar fields resulting from the displacement of past and previous rain gauge samples in the direction of motion. The difference, e.g. the structures in motion direction are clearly visible in the FF-IDW surface, result from the two different dependency functions (Euclidean distance vs. variogram) of both methods. The methods RK and FF-RK show a clear dependency on the radar image.




Figure 7.4: Example of a snapshot of current radar image (left) and surfaces interpolated with the different methods. Blue corresponds to no precipitation, red to values equal to or larger 25 mm/h. Current rain gauge measurements displayed with **blue bars**, estimates with **black bars** with maximum corresponding to 25 mm/h. Black bar in the left most image shows the current radar measurement, black bars in the other images show cross-validation results for the different methods IDW, OK and ST-OK. Motion vectors (left image) estimated using optical flow show the field motion direction and speed in 5 min.




Figure 7.5: Example of a snapshot of the non-symmetric spatio-temporal interpolation and radar-rain gauge merging methods.

Finally, Figure 7.6 displays the estimate provided by the method RADAR-DBC introduced in Section 4.2.4. The method computes a space-time lag between radar and rain gauges for each rain gauge first from a temporal window of $\pm 20 \text{ min}$ and potential temporal lags of $\pm 5 \text{ min}$ and spatial lags of $\pm 5 \text{ km}$, resulting in the best fitting radar time series for a particular rain gauge and time. Then, a dynamic radar-rain gauge adjustment coefficient is computed for each rain gauge from both time series (Equation 4.2). Finally, the spatio-temporal lag and adjustment coefficient to be used for intensity estimation are interpolated. Figure 7.6 displays examples of the result of interpolation of the spatial and temporal lags.

The lag surfaces indicate a slightly positive lag in both spatial directions on average and a stronger negative lag in time. Thus, for most of the spatial locations, the method RADAR-DBC selects a radar image of a previous time step for estimation. The impact can be seen when comparing the estimated surface in Figure 7.6, left, with the current radar image displayed in Figure 7.4, left: The surface estimated by RADAR-DBC is slightly displaced towards the negative direction of field motion, due to the selection of previous radar measurements.




Figure 7.6: Example of the surface estimated by the method RADAR-DBC (left) and surfaces displaying the current estimated spatial lag in X (SPACE-LAG-X-GRID), Y (SPACE-LAG-Y-GRID) and temporal lag (TIME-LAG-GRID). Red corresponds to negative lag, blue to positive. Color intensity indicates lag magnitude in the interval $[-5\text{km}, 5\text{km}]$ for space and $[-5\text{min}, 5\text{min}]$ for time.




Figure 7.7: Example of the regression surface estimated by the method RADAR-DBC. Blue indicates an adjustment coefficient $a > 1$ and thus, a recognized underestimation of precipitation by radar. Red color indicates an adjustment coefficient $a < 1$ and thus, a radar-based overestimation of precipitation.

Figure 7.7 displays an example of the result of adjustment coefficient interpolation. The method seems to learn that there is a current underestimation of precipitation intensity by radar for locations close to rain gauge 10338. At all other locations (except E688), the radar seems to slightly overestimate precipitation intensity.

7.1.4 Radar estimation and rain gauge cross-validation results

This section provides the cross-validation results for all estimation methods, computed as described in Section 7.1.2. In Figure 7.8, the correlation coefficient is displayed for the spatial interpolation methods (black bars), the spatio-temporal interpolation methods (grey bars), the radar-rain gauge merging methods (blue bars), the radar measurements of the same time step aggregated over the given number of pixels (darker red bars) and the radar measurements in a temporal window surrounding the gauge measurement (lighter red bars).

The pure spatial interpolation methods IDW and OK, the symmetric spatio-temporal method ST-OK and the single radar pixel (RADAR 1) at the gauge location yield lowest values. They are outperformed by the non-symmetric spatio-temporal methods FF-IDW and FF-OK and the radar rain gauge merging methods. In winter, the radar-based methods seem to perform worse, which might be due to the ice in the atmosphere resulting in erroneous radar reflectivities. Contrary, FF-IDW and FF-OK perform better in winter, possibly due to the more homogeneous and large

scale nature of the rain events. In general, FF-IDW performs better than FF-OK, despite the dynamic variogram estimation of FF-OK, which is a surprising result. A reason could be that FF-OK is more vulnerable to slow precipitation field speeds, a topic that will be discussed in more detail later.




Figure 7.8: Correlation of estimates with rain gauges measurements for January 2012 and June 2012. Black bars: pure spatial interpolation methods. Grey bars: spatio-temporal interpolation methods. Blue bars: radar-rain gauge merging methods. Red bars: radar measurements integrated over the given pixel neighborhood. Light red: radar measurements at the spatial gauge location separated by the given number of minutes (negative values indicate previous radar values).

In summer, the radar-based methods clearly outperform the rain gauge interpolation methods. Concerning the merging methods: the RADAR-DBC method always outperforms the other two methods RK and FF-RK. Smoothing radar is also beneficial: the very large radar window of 121 pixels provides highest correlations. An interesting observation can be made on the radar time series data: the highest correlation occurs for the radar image 6 minutes previously (RADAR T-6), an effect that seems to be more pronounced in summer. The reason for the time lag might be the difference in measurement procedures: radar estimates in the atmosphere, rain gauge on the ground, and precipitation requires time to reach the ground. In Section 2.1.3, it was shown that at a distance of 100 km from the radar station, structures of small rain drops with a low falling speed of 2 m/s are detected by radar 8 minutes earlier. Since all of the rain gauges are spatially closer to the radar station, a time lag smaller than 8 minutes seems realistic. In summer, the correlation peak at RADAR T-6 is higher, probably due to the influence of 'bright band' or hail and snow effects being lower in summer.

Figure 7.9 displays another error metric, the mean bias. Those methods that estimate precipitation from rain gauges only show biases close to zero. All other methods seem to overestimate rainfall on average. In summer, the radar methods show better performance in terms of lower positive biases, i.e. a lower overestimation of precipitation in general. Again, the reason might be ice in the atmosphere in winter.




Figure 7.9: Bias of estimates of different methods for January 2012 and June 2012.

The relative standard error, given as a percentage of the observed mean value, is displayed in Figure 7.10. All methods show high errors above 200 % of the mean value, indicating a rather low predictive performance of all methods. As displayed in Table 7.1, the mean precipitation intensities of the rain gauge measurements are 1.3 mm/h in January 2012 and 2.0 mm/h in June 2012, resp. Thus, a RSE of 200 % corresponds to a root-mean-square error (RMSE) of $1.3 \text{ mm/h} \times 200/100 = 2.6 \text{ mm/h}$ in January and $2.0 \text{ mm/h} \times 200/100 = 4.0 \text{ mm/h}$ in June.




Figure 7.10: RSE of estimates for January 2012 and June 2012.

In general, the rain gauge interpolation methods yield lowest RSE values compared to the merging or radar-only methods.

Results for rain gauge measurements larger zero

The different methods for 1-min precipitation estimation are applied to the collected car data in Section 7.2. Then, when the car sensors report some sensor reading larger zero, e.g. a certain wiper frequency per minute, it is very likely that there is rainfall at the car position. The effect of the windscreen cleaning using the wipers and wiper washer spray when there is no rain, is neglected. It is assumed that this does not occur very often and thus, does not influence the statistics. Therefore, another comparative analysis is done for rain gauge measurements larger zero. Here, the cross-validation statistics are only computed for those rain gauge measurements that are known to be larger zero, motivated by the use case of attaching 'ground truth' rainfall to the car measurements, as described previously. Figure 7.11 displays the resulting correlation coefficient.




Figure 7.11: Correlation of estimates with rain gauges measurements for January 2012 and June 2012 for rain gauge samples larger zero.

Although the correlation is slightly lower for all methods compared to the results that included zero measurements displayed in Figure 7.8, the results are very similar. However, the bias measure shows a completely different picture (Figure 7.12).

Now, all rain gauge interpolation methods underestimate the precipitation intensity, with a slight performance advantage for the methods FF-IDW and FF-OK. The underestimation is more severe in summer due to the larger rain intensities. The pure radar-based methods still overestimate in winter but tend to underestimate in summer, which is an interesting result that supports the known property of radar to underestimate strong, but overestimate light rainfall. It is also likely that ice or melting ice in the atmosphere is still contributing to this overestimation in winter ('bright band' effect).

For samples larger zero and the bias measure, the merging methods mostly outperform the rain gauge interpolation and radar-only methods. The dynamic bias-correction method RADAR-DBC does not seem to prevent the radar overestimation in winter, while in summer, it delivers the lowest bias of all methods. The reason is not completely clear but could lie in the requirement of the




Figure 7.12: Bias of estimates for larger-zero rain gauge samples for January 2012 and June 2012.

method that the adjustment coefficient a must be in the interval $[0.5, 2]$ in order to prevent it from deviation. In summer, it might be possible to correct the underestimation of radar by a coefficient $1 \leq a \leq 2$. In winter, the overestimation might be more severe to be corrected by a coefficient $1 \geq a \geq 0.5$. Thus, future work requires more investigation on the computation of the dynamic adjustment coefficient. A good trade-off between over- and underestimation seems to be provided by the method FF-RK.

The RSE displayed in Figure 7.13 shows significantly lower values compared to the previous results, due to the larger mean values. Here, as displayed in Table 7.1, the mean intensities excluding zero measurements were 3.19 mm/h in January 2012 and 7.98 mm/h in June. Thus, this time, a RSE of 200 % corresponds to a RMSE of $3.19 \text{ mm/h} \times 200/100 = 6.38 \text{ mm/h}$ in January, and $7.98 \text{ mm/h} \times 200/100 = 15.96 \text{ mm/h}$ in June. Still, these values indicate a low predictive performance of all methods and no method provides significantly lower RSE values compared the others.

The results presented in this section serve as a basis for the selection of the reference methods for the car data analyses presented in Section 7.2.




Figure 7.13: RSE of estimates for January 2012 and June 2012 for rain gauge samples larger zero.

Dependency on motion speed

Another analysis is provided for different classes of field speed for the winter month (Figure 7.14). The results show that the performance of the interpolation methods IDW and OK (and also ST-OK) strongly depends on the field speed: for fast fields, their performance is heavily decreased. An explanation could be that slow field speeds are associated with a rather homogeneous rainfall over larger distances, which is clearly beneficial for standard interpolation methods. As expected, the motion-based non-symmetric spatio-temporal methods perform well under fast speed. Surprisingly, the performance of FF-OK is decreased for small field speeds and is no longer better than IDW, OK or ST-OK. This could be explained by considering the extreme case of zero speed. Then, all samples of a particular rain gauge that are used for FF-IDW or FF-OK (41 samples) have roughly the same location due to the 'frozen field' distance function. At zero field speed, the samples are not moved at all but are still used for interpolation. Thus, all samples result in exactly equal variogram values when paired with the estimation location or with other measurements. Thus, the resulting kriging matrix of Equation 2.33 might become instable, resulting in less accurate estimates. The method FF-IDW seems to be more stable and independent of field motion.

Another interesting observation can be made related to the radar data: the slower the speed, the less the correlation with radar. This is a surprising result and the reasons are not obvious. It could be that the event selection method (radar values larger 1 mm/h on average) selects a number of rain events with large radar errors based on the spatial configuration and possibly due to ice in the atmosphere, similar to those displayed in Figure 4.6. Further, as the motion is calculated from the radar images, such events also result in low field speeds and thus, the combination of low accuracy of radar with low field speeds. However, this is a hypothesis that has to be investigated further in the future.




Figure 7.14: Correlation of estimates with rain gauges measurements for January 2012 for two classes of field speed.

Individual rain gauge analyses

The analysis of the correlation results for each rain gauge individually and for the winter month is displayed in Figure 7.15.

As expected, the spatial methods IDW and OK perform better in areas with larger rain gauge density, such as at rain gauges E667, 10338 or E755. In comparison, the 'frozen field' methods show better performance everywhere but especially at locations where there is a rain gauge location not too far away in a direction roughly corresponding to the prevailing motion direction of the rain field (445 m/5min from west to east and 330 m/5min from north to south), as displayed in Table 7.1. It is likely that the good performance of the 'frozen field' interpolation methods at the rain gauges E667 and 10338 is due to their spatial separation being close to this direction of mean motion. The radar data shows a clear performance decrease at the rain gauge 10338 that is close to the weather radar station at Hanover airport, possibly due to increased errors in radar reflectivity due to ground clutter. Further, the temporal bias of 6 minutes of the radar estimates seems to be more pronounced in areas far away from the radar station. Most likely, this is due to the increased radar beam elevation at these locations. However, this requires further investigations, especially larger study areas and more rain gauges.




Figure 7.15: Map of correlations of estimates with rain gauges measurements for January 2012 for all rain gauges individually.

7.1.5 Summary

The pure spatial-interpolation methods IDW and OK and the space-time symmetric spatio-temporal method ST-OK are outperformed by the non-symmetric spatio-temporal methods FF-IDW and FF-OK for almost any analysis and error measure. Thus, integration of motion information is beneficial. Surprisingly, the dynamic variogram estimation for OK, ST-OK and FF-OK did not provide any significant improvement compared to the deterministic IDW and FF-IDW methods relying on Euclidean distance. The motion-based methods seem to perform nearly as well as radar-based methods concerning correlations. Further, in contrast to weather radar estimates, they have the property that they still provide unbiased estimates, i.e. the average bias is zero when all measurements are considered, including those showing zero intensity. The field speed analysis shows that especially for slowly moving precipitation fields, the performance of FF-OK decreases, which might be due to the instability of the kriging system when a high number of data points with almost equal locations is involved. Thus, the simple FF-IDW model seems to be preferable over

FF-OK, as it does not require variogram estimation and model fitting, is therefore computationally efficient and easy to implement and seems to be more robust.

Merging rain gauges and radar showed promising results. There, for all evaluation criteria, the methods FF-RK and RADAR-DBC mostly outperformed RK. Especially for samples larger zero, the merging methods are capable of providing a smaller bias. Nonetheless, no merging method was capable of removing the radar-induced overestimation of precipitation in winter. Mostly, the RADAR-DBC method provided highest correlations among the merging methods tested but showed almost no superior performance concerning the other error measures. Thus, more work is required on this method.

The radar-only methods showed interesting results as well: as expected and described in the literature, the single radar pixel at the spatio-temporal estimation location does not provide a good estimate. Smoothing radar, i.e. aggregating over pixel neighborhoods, is always beneficial. However, the results show that there is a strong temporal lag between weather radar snapshots and rain gauge measurements: the radar snapshot 6 minutes before the rain gauge 1-min accumulation provided highest correlations. This temporal lag can have several reasons, ranging from the different measurement procedures to the difference between ground and atmospheric measurements, or a combination of these. In the case that the reason for the lag is the elevated radar beam, it is likely that the lag intensifies with increasing radial distance from the radar station.

The cross-validation statistics for larger-zero measurements have shown that all rain gauge interpolation methods underestimate the precipitation intensity with FF-IDW showing best performance. In general, the merging methods yield the lowest absolute bias. The results are important for the selection of the reference method for the car data analyses provided in the next section.

The correlation measure for different field speed categories shows that the accuracy of standard spatial and space-time symmetric spatio-temporal rain gauge interpolation methods decreases with increasing field speed. The motion based methods still perform well for fast speed values. The correlation between radar and rain gauges seems to be higher for faster moving field. This could be an artifact of the processing methodology: since motion is estimated from weather radar, large errors in weather radar, e.g. errors due to the 'bright band' effect, result in slower field speed values and in turn, a lower correlation between rain gauges and the erroneous radar estimates.

The analysis of the errors of individual rain gauges showed poor radar performance in areas close to the radar station, especially when there is no radar smoothing involved. Further, the motion based interpolation methods outperform the other interpolation methods at all locations but especially at those where there is another rain gauge close by in the direction of the prevailing field motion.

The results show the general behavior of the different methods and can be used for the selection of a particular estimation method for particular use cases, for example, for the derivation of the reference intensities for the car data. However, since the methods have been evaluated with gauge measurements that are subject to errors as well (Section 2.1.2), the results do not represent absolute accuracy of estimates of ground precipitation.

7.2 Precipitation intensity estimation with cars

7.2.1 Study area and data basis

For the investigation of the estimation of 1-min precipitation intensity with cars, a set of test cars has been equipped with wiper frequency sensors or the Xanonex optical rain sensor. In Figure 7.16, the distribution of samples in Hanover city is displayed. Most of the wiper frequency data (left) is collected by cars from a Hanover Taxi company, resulting in most of the data being located in the very center of the city. The Xanonex data is mainly collected in close proximity to the rain gauge E755 in Hanover Herrenhausen by the author. The existence of a road with a maximum speed of 100km/h in that area was beneficial for the car speed analyses.




Figure 7.16: Heat maps and car measurements for the wiper frequency (left) and Xanonex sensors (right).

For the analyses presented in this section, a subset of the collected data is used, which consists of 2437 minutes of wiper frequency data of two VW Touran with manually controlled wipers. Further, 1028 minutes of data collected by a car with automatically controlled wipers (Ford S-Max) and 657 minutes of data collected by a Skoda Fabia equipped with the Xanonex sensor is analyzed.

Figure 7.17 displays two examples of Xanonex 1-min samples collected at different days in June 2014. The example shows some of the difficulties that occur when attaching 'ground truth' precipitation intensity data to the car samples: in Figure 7.17, left, a pure spatial interpolation of the current rain gauge samples would result in zero precipitation intensity at the car position, although there is rainfall at the car position. Attaching the current radar pixel, which seems to be zero as well, is also not beneficial as there seems to be a spatio-temporal bias in between ground precipitation and radar. Thus, more sophisticated methods are required. Figure 7.17, right, shows an example where the precipitation is more widespread and spatial interpolation of rain gauge samples might be valid.

7.2.2 Selection of the reference method

The analysis of the errors of individual rain gauges provided in the previous section showed poor radar performance in areas close to the radar station, especially when there is no radar smoothing




Figure 7.17: Extract of the visualization of the car data for two Skoda Fabia Xanone samples collected in June 2014. The red bar indicates the current Xanone sensor reading. For display purposes, the maximum here corresponds to 15 [s/min]. Current radar and rain gauge estimates of precipitation and field motion per 5 min.

involved. Since, as displayed, most of the car data is collected in the Hanover city center close to the radar station, only smoothed radar data is used as the reference. Further, from each of the three categories rain gauge interpolation-, radar-rain gauge merging- and radar-only methods, one method is selected and used as a reference for the car data. The results of the cross-validation analyses for samples larger zero presented in Section 7.1.4 provide the basis for the selection. For the rain gauge interpolation methods, the method FF-IDW showed best performance for all measures. Further, the method RADAR-DBC showed highest correlations among the radar rain gauge merging methods. Finally, the method RADAR 121 showed best performance of all snapshot-based radar-only methods. Since it is likely that the time lag between rain gauges and weather radar is due to the increased beam elevation and since most of the car data is collected close to the weather radar station, the radar-only time-lag methods are disregarded in the car analyses. Thus, for the car data analyses, the methods FF-IDW, RADAR 121 and RADAR-DBC are selected as the reference. RADAR-DBC is expected to yield least smoothed estimates among the three methods, as radar values are averaged in a pixel window of 9 pixels only (Equation 4.4). Further, as described previously, it can be expected to provide the lowest bias among all methods tested (as displayed in the previous section in Figure 7.12). The rain gauge interpolation FF-IDW can be expected to yield the W-R relationships with the lowest slope, since it tends to underestimate precipitation intensity, as described in the previous section.

7.2.3 Manually-operated windscreen wipers

The analysis of wiper frequencies for cars with manually controlled wipers provided in this section is based on data of two Volkswagen Touran of a Taxi company. As described earlier in Section 5.3, for a particular wiper setting using the lever, e.g. "periodic low wiping", the base frequency at zero car speed (the number of windscreen wipes per minute) increases with increasing car speed until it reaches an upper bound for that particular setting of the lever. Thus, in a scatter plot of car speed vs. wiper frequency, a pattern of increasing wiper frequency with increasing car speed should be recognizable. Figure 7.18 displays the 1-min wiper frequency samples vs. the car's average speed within that particular minute. There seem to be three different wiper classes starting with a certain wiper frequency at car speed zero. With increasing car speed, the frequency increases and reaches an upper bound for each class, which is exactly what was expected after the theoretical

considerations of Section 5. The mean car speed for a particular wiper frequency (colored dots) give a more clear picture of the pattern than the individual 1-min data points.




Figure 7.18: Car speed vs. wiper frequency for two VW Touran (grey dots represent 1-min samples) and mean car speed for a particular wiper frequency (colored dots). Colors indicate different (assumed) wiper classes.

The deviations of the individual 1-min data points from the mean speed value for a particular frequency can be explained by the data aggregation, i.e. errors in the car speed calculation, the existence of different car speeds within the 1-minute aggregations as well as the manual switching of the wiping class within a minute by the driver. Above a frequency of 20, a similar pattern can be observed, but the data is too sparse for the calculation of the mean speed for each frequency. Thus, every frequency above 20 is considered wiper class 4 in the following. Under the assumption that this separation of wiper classes is valid, the distribution of precipitation intensity is calculated for each using different reference intensities (Figure 7.19).




Figure 7.19: Wiper classes vs. rainfall for two VW Touran and the different references FF-IDW (left), RADAR-DBC (center) and RADAR 121 (right). Bars represent standard deviations, dots arithmetic mean values and filled triangles the medians.

As expected, there is an increase in mean and median rainfall intensities with increasing wiper class. The medians always lie below the mean values, indicating a distribution of values for each class that is skewed towards zero values. The reason for the large standard deviations is a mixture of inaccurate reference intensity, the low number of discrete wiper classes in such cars and the

possibly erratic operation of the wipers by the drivers. As expected, the method FF-IDW yields lowest precipitation intensities and thus, the lowest increase in intensity with increasing wiper class. It is likely that the underestimation of intensity by the rain gauge interpolation methods described previously is the reason.

7.2.4 Automatically-operated windscreen wipers

Different cars with wipers automatically operated by an optical sensor in the windscreen have been equipped with frequency sensors. Here, an analysis is provided for one of them, a Ford S-MAX, which has also been tested in the laboratory (Rabiei et al., 2013). In Figure 7.20, left, the relationship between the sensor readings and the car speed is displayed. In order to reduce the precipitation intensity influence, only a subset of 1-min samples is selected where the attached RADAR 121 rainfall is in the interval $[0.5 \text{ mm/h}, 1.5 \text{ mm/h}]$.




Figure 7.20: Car speed vs. 1-min wiper frequency (grey dots) for a single Ford S-Max with automatically controlled wipers (left) and derived ratio (right). Red dots represent the mean wiper frequency for car speed classes of 0-10, 10-20, etc.

Most likely, the car speed influence on the wiper frequency is a combination of the increased water on the windscreen surface detected by the sensor controlling the wipers and the dependency between car speed and wiper frequency implemented in the car, which has been described previously. This might be the reason for the nonlinearity of the relationship at car speeds below 10 km/h and also the jump in conditional mean values at a car speed of 50. At a low car speed, the wipers might be set to a low frequency automatically, independent of rainfall. The derived ratio displayed in Figure 7.20, right, is used to eliminate the car speed influence from the sensor readings.

The correlation between the car-speed corrected wiper frequencies and the reference rainfall intensity of the different methods is displayed in Figure 7.21.




Figure 7.21: Correlation of speed-corrected wiper frequencies of a single Ford S-Max with automatically controlled wipers with rainfall estimated with the different methods.

Interestingly, the correlation coefficients are mostly higher compared to the cross-validation statistics computed from the rain gauges and presented in the previous section. On average, the car data is spatially much closer to the next rain gauge than any of the rain gauges to its nearest neighboring rain gauge. Thus, the rain gauge interpolation methods can be expected to perform better for the car data. Further, one reason for the temporal bias in between radar and rain gauge data identified in the previous section (the fact, that RADAR T-6 provided highest correlations in most cases) could be the distance of the rain gauges to the radar station described previously: with increasing distance from the station, the radar beam is higher elevated, resulting in a longer time, the precipitation detected by radar requires to reach the ground. It is likely that, since the car data is also closer to the radar station, this is the reason for the disappearance of the correlation peak at RADAR T-6 (which moved towards RADAR T0) for the car data and the higher correlations for radar data of the same 1-min time step (RADAR 1, RADAR 9, etc.).

The displayed results can give an indication on the degree of linear relationship between both variables and thus, can help in selecting the variable for which to calculate prediction functions. It is clear that nonlinear dependencies can not be detected this way and that the existence of strong non-linear dependencies might still result in low correlation values, which seems to be the reason for the low correlation values for the RADAR-DBC method. Thus, in the following, as selection of reference methods, the methods FF-IDW, RADAR DBC and RADAR 121 are used. Figure 7.22 displays the number of 1-min samples for classes of speed corrected sensor readings available.




Figure 7.22: Histogram of speed-corrected readings (wiper frequencies) of the Ford S-Max.

The data does not contain sufficient samples for readings larger 10. Thus, only readings up to 10 are included in the following analyses. For the same reason, the coefficients of determination for the linear functions provided in the scatter plots do not equal the squared correlation coefficients as displayed in Figure 7.21. The W-R relationships for the speed-corrected readings (Figure 7.23) show a clear dependency of conditional mean precipitation intensities on the speed-corrected readings.




Figure 7.23: Speed-corrected 1-min sensor readings vs. two different references (grey dots) for a single Ford S-Max with automatically controlled wipers. Red dots represent the mean precipitation intensity for reading classes of 0-2, 2-4, etc. Black filled triangles represent median precipitation intensities for these classes. Red and blue lines (left figure) represent the 95% confidence intervals for the conditional mean (red) and prediction (dotted blue). Grey bars represent standard deviations for the reading classes 0-2, 2-4, etc.

While the reference FF-IDW shows a clear linear dependency, RADAR 121 shows a more exponential increase. The reason is currently unclear. Nonetheless, the intensities are comparable. The scatter plots show a clear indication for heteroscedasticity (Section 2.3.2). However, due to the large deviations, heteroscedasticity is considered a problem that is negligible.




Figure 7.24: Speed-corrected 1-min wiper frequencies vs. RADAR DBC reference (grey dots) for a single Ford S-Max with automatically controlled wipers. Red dots represent the mean precipitation intensity for reading classes of 0-2, 2-4, etc. Black filled triangles represent median precipitation intensities for these classes.

Again, similar to the results for the two VW Touran, the method RADAR-DBC shows the strongest increase in precipitation intensity with increasing sensor reading (Figure 7.24) compared to the other two methods. Again, the reason might be the non-smoothing behavior compared to RADAR 121 and the underestimation of intensities by FF-IDW.

7.2.5 Xanonex optical rain sensor

This section provides the results for the Xanonex sensor installed in a Skoda Fabia. As described previously, although the Xanonex sensor has been placed behind the windscreen in an area where the wipers clean the windscreen, the wiper influence on the Xanonex readings is considered low. Therefore, the data displayed in Figure 7.25, left, is considered to be the result of an increasing amount of water on the windscreen surface with increasing car speed.




Figure 7.25: Car speed vs. 1-min sensor readings (left) for a single Skoda Fabia with the Xanonex sensor (grey dots) and derived ratio (right). Red dots (left figure) represent mean values for the speed classes 0-10, 10-20, etc. The theoretical ratios for rain drop falling speeds of 5m/s and 2m/s for the windscreen angle of 32.5° of the Skoda Fabia are calculated using Equation (6) of Rabiei et al. (2013).

In order to reduce the precipitation intensity influence, only a subset of 1-min samples is selected where the attached RADAR 121 rainfall is in the interval $[0.5 \text{ mm/h}, 1.5 \text{ mm/h}]$. Thus, the displayed car-speed-sensor reading relationship is considered valid for low rain intensities. In fact, the derived ratio displayed in Figure 7.25, right, shows a strong similarity to the theoretical ratio for low rain intensities (falling speed of 2m/s) after Bocci (2012). The speed-corrected Xanonex readings show high correlation with the references (Figure 7.26).

Again, the correlations are much higher compared to the correlations determined by rain gauge cross-validation described in the previous section. Most of the Xanonex data has been collected by the author in close proximity to gauge E755 (Figure 7.16), which is beneficial for the rain gauge interpolation methods. Again, an interesting observation can be made on past radar data: the time lag in between radar and rain gauge measurements (the radar pixel 6 min earlier RADAR T-6 provided highest correlations in cross-validation) disappears for the car data. Again, the reason could be that on average, the car data is spatially much closer to the radar station than most of the rain gauges and thus, the problem of increased radar beam elevation at the car position is reduced. However, the analyses later show that the measurements of the nearest rain gauge lag behind the




Figure 7.26: Correlation of speed-corrected Xanonex readings installed in Skoda Fabia with rainfall estimated with the different methods.

car measurement a number of minutes, which suggests a general time lag in gauge measurements. Figure 7.27 displays the number of 1-min samples for classes of speed corrected Xanonex readings.




Figure 7.27: Histogram of speed-corrected Xanonex readings of the Skoda Fabia.

Here, the number of readings above 5 is considered too low for reliable analyses. Thus, Figure 7.28 displays the W-R relationships for the Xanonex for speed-corrected readings smaller 5.

Again, the mean values show a linear dependency and both references show comparable values. Since, as described in Section 5.3.3, the influence of the wipers on the Xanonex readings is assumed to be neglectable, the **uncorrected** Xanonex readings are assumed to be a direct measure of water on the windscreen, which increases with increasing car speed. Therefore, the **speed-corrected** Xanonex readings are considered to be an indicator for precipitation intensities. The rather large residuals in intensity derive from a combination of factors that are not modeled but influence the Xanonex readings (such as road spray, etc.) and errors in the references. Thus, the rather large 95% confidence bounds for the predictions are assumed to be upper bounds for prediction accuracy (conservative estimates of prediction accuracy) and should be taken with care. In Figure 7.29, the relationship is displayed for the RADAR DBC method.




Figure 7.28: Speed-corrected 1-min sensor readings (smaller 5) vs. references FF-IDW (left) and RADAR 121 (right) for a single Skoda Fabia (grey dots). Red dots represent conditional mean intensity for the reading classes 0-1, 1-2, etc. Red and blue lines represent the 95% confidence intervals for the conditional mean (red) and prediction (dotted blue).




Figure 7.29: Speed-corrected 1-min sensor readings vs. RADAR DBC reference for a single Skoda Fabia and sensor readings smaller 5 (grey dots). Red dots represent conditional mean intensity for the reading classes 0-1, 1-2, etc. Red and blue lines represent the 95% confidence intervals for the conditional mean (red) and prediction (dotted blue).

Again, a similar relationship can be observed which shows a slightly larger slope, potentially due to the non-smoothing behavior of the method. In order to reduce the smoothing effect of the reference method in general, another analysis is performed for samples closer than $2km$ to the next rain gauge, which, for the Xanonex data, is solely rain gauge E755 (Hannover-Herrenhausen). An interesting observation can be made when investigating the correlation of the speed-corrected Xanonex readings with precipitation intensity of E755 in a time window of $\pm 10min$.

Highest correlations can be observed for rain gauge measurements a couple of minutes later with a peak for a time lag of $+7min$. This result supports the previous findings that the rain gauge measurements seem to lag behind radar and cars. The dependency for the FF-IDW method and the nearest neighbor (NN) rain gauge measurement 7 min later (Figure 7.31) shows a strong non-linear increase, which, however, has to be taken with care due to the low number of samples).




Figure 7.30: Correlation of sensor readings with rainfall estimated from rain gauges for sensor samples with a spatial distance closer 2km to next rain gauge.




Figure 7.31: Speed-corrected 1-min Xanonex readings vs. FF-IDW for samples closer 2 km to nearest rain gauge (left) and vs. the measurement of nearest neighbor 7 minutes later (right).

The comparison with the non-interpolated measurement of the nearest rain gauge shows higher values than before. Since here, the reference estimates are the original non-interpolated values, there is no smoothing involved and the magnitude of values might be considered more accurate than for any other reference before. However, still, the number of data points is considered too low to make reliable inferences.

7.2.6 Results of experiments on the VW rain track

During a two-day experiment on the Volkswagen rain track in Wolfsburg, the Skoda Fabia test car with the Xanonex installation was tested under laboratory conditions. The test track allows the simulation of homogeneous rainfall throughout the whole track. Another Xanonex sensor was installed on the car roof, with the possibility of adjusting the angle in driving direction and hence the simulation of different car windshield angles. Inside the car, a wiper frequency sensor was installed measuring the frequency of the manually operated wipers. The rainfall on the track was

thoroughly analysed using the tipping bucket and disdrometer sensors. However, the produced rainfall turned out to be much stronger than natural rain. Therefore, all car sensors, including the wiper frequency sensor, delivered maximum readings for all rain intensities above the lowest possible intensity. For this reason, the analysis of the data concentrated on the influence of the car speed and the windscreen angle on the sensor readings. Four speed classes were investigated: 0, 20, 40, 60 (the maximum possible speed on the track was around 60 km/h). During an experiment for a particular car speed, the speed was driven constantly and the acceleration and deceleration phase were deleted from the data either directly or in post-processing. Four classes of windscreen angles were tested: 0° (horizontal), 30° from horizontal, 60° from horizontal and 90° (vertical). Figure 7.32 displays the results.




Figure 7.32: Xanonex sensor readings vs. windscreen angle (conditional mean values, left) for zero car speed and static-dynamic ratio for different windscreen angles (right).

The results show an almost perfect linear relationship of the mean values of the sensor readings conditioned on the windscreen angle for a car speed of zero (Figure 7.32, left). The results for different windscreen angles at different car speeds show that the water amount on the windscreen is independent of car speed for a horizontal windscreen. As expected, the car speed influence increases with increasing angle towards vertical. The ratios are lower compared to those derived in the field experiments, most likely as the precipitation intensity and thus, the drop size was significantly higher on the test track.

7.2.7 Summary

For cars with manually operated windscreen wipers, the results show a clear dependency between car speed and wiper frequencies implemented in the cars. After the elimination of the car speed influence, a positive dependence to the reference precipitation intensity could be observed. However, the prediction accuracy for the wiper frequencies was rather low, possibly due to the large number of influencing factors, the erratic operation by the driver and the inaccurate reference intensities. The results for cars with automatically operated windscreen wipers showed higher accuracies after the elimination of the car speed influence. As expected, the optical rain sensor Xanonex showed best performance among the sensors tested and surprisingly high correlations with the reference intensities.

For all sensor types, the RADAR-DBC method showed the strongest increase in precipitation intensities with increasing sensor readings. It is likely that the non-smoothing behavior of the method and the low absolute bias compared e.g. to FF-IDW is the reason for this. Thus, the W-R relationships provided by this method might be considered as the most reliable ones among the reference methods tested.

However, most of the data was collected at rather low rain intensities and the amount of data at high intensities was too sparse for reliable analyses. Thus, more data collected at high rain intensities together with better reference intensities will be required in order to make reliable inferences concerning the applicability of cars for precipitation intensity estimation in the envisioned hydrological use cases.

7.3 Field motion estimation with a geosensor network

In this section, the performance of the decentralized algorithm for field motion estimation described in Chapter 6 is evaluated. The goal of the evaluation is to illustrate the general ability of the algorithm to estimate the motion of spatio-temporal moving fields with a GSN and to show the algorithm behavior under different conditions such as fields, motion speed and network configurations.

7.3.1 Study Area, sensor network and deployment strategies

A 10 km \times 10 km area in the center of Hanover city is chosen as the study area for evaluating the proposed motion estimation algorithm (Figure 7.33).




Figure 7.33: The 10 km \times 10 km study area for the motion estimation algorithms. OSM, Open Street Map.

This particular study area has been chosen due to the availability of weather radar data. As described in the previous section, the weather radar covers a much larger area, but current GSN technology usually restricts communication distances to ranges of tens to hundreds of meters, in order to reduce energy demand. Therefore, in order to account for the low communication ranges, a small urban sub-area of the radar-covered area is chosen. Different distributions of sensors, fields, motion speed and parameter set-ups are evaluated.

A GSN with a UDG 1-hop communication model is simulated, and different deployment strategies are tested (Figure 7.34). First, nodes are distributed in the study area in a completely random way (Figure 7.34a), potentially resulting in a disconnected graph where only a subset of nodes is able to estimate motion. Second, a regular grid-based deployment is tested (Figure 7.34b). Then, as two more realistic cases, a diffusion-based deployment is simulated based on a source node in the center of the study area, for example a base station or node collecting the results (termed *sink* (Duckham, 2012)), where the density is controlled by two parameters, the maximum and minimum distance to neighboring nodes (Figure 7.34c). The diffusion-based strategy is more realistic in the sense that the resulting graph is a connected graph where every node has at least one neighboring node to communicate with. Such a setting is considered more likely to be deployed in reality than the completely random or grid-based deployment. As last strategy, cars are generated randomly

on a road network derived from Open Street Map (OSM). A snapshot of the generated car nodes is displayed in Figure 7.34d.




Figure 7.34: Deployment strategies. (a) Completely random deployment; (b) grid deployment; (c) diffusion deployment around a central node; (d) vehicular ad hoc network (VANET).

Throughout the experiments, two different node densities and communication ranges are simulated: (1) a low-density, large-distance communication deployment with $n = 25$ nodes and a communication range of $r = 2000$ m; and (2) a high-density, short-distance communication deployment with $n = 400$ and $r = 500$ m.

Car motion model

The cars move according to a simple random walk model on the road network with a certain speed, which varies throughout the experiments. The speed is specified as the Euclidean distance per time step that a car can move. In the following, the car motion model is described that is executed for each car in between two steps of the simulation, in order to generate a car position for the upcoming simulation step.

In between two time steps, the cars move from node to node in the road network and constantly decrement the remaining Euclidean distance that the car still has to move by the Euclidean distance between the nodes. In order to prevent unrealistic jumping of the car position in between two nodes, a car cannot directly revisit a node it has visited previously. If the distance to the next node in the road network is larger than the remaining distance to move, the car position for the upcoming time step of the simulation is generated by linear interpolation in the direction of the next node. In addition, the communication graph (UDG) is updated after all car positions have been computed. Due to the update of the communication graph, the INIT step of the algorithm (Section 6.5) is now executed every time step, which results in a heavy increase of the costs for communication and computation for moving nodes.

7.3.2 Error measures

Following Barron et al. (1994), the error of the motion estimates is quantified using the angular difference between true motion V and estimated motion \hat{V} , called **angular error**:

$$\arccos \left(\frac{\hat{V}(u_{i,t})}{\|\hat{V}(u_{i,t})\|} \cdot \frac{V(u_{i,t})}{\|V(u_{i,t})\|} \right) \in [0, 180] \quad (7.3)$$

where $\|\cdot\|$ is the vector norm, \cdot is the dot-product, $V(u_{i,t}) = (v_X, v_Y)^T$ is the true motion at the spatio-temporal node location $u_{i,t}$ and $\hat{V}(u_{i,t}) = (\hat{v}_X, \hat{v}_Y)^T$ is the estimated motion, i.e., the Kalman filter state of node i at time t (excluding the acceleration variables). In addition, the difference in motion speed is calculated, called **speed bias**:

$$\left\| \hat{V}(u_{i,t}) \right\| - \|V(u_{i,t})\| \quad (7.4)$$

In the case of a simulated field, the true field motion is known, as it is a parameter of the simulation. For weather radar fields, the motion estimated by the algorithm of Farnebäck (2003) is used as true motion, which is estimated from the weather radar images.

7.3.3 Setting the filter parameters

The initial state of the filter at time step zero is set to the zero vector: $v_i(0) = 0_{4,1}$. The initial state covariance P is set to large values on the diagonal, indicating the uncertainty in the initial filter knowledge on the involved variables.

Kalman filter prediction and measurement noise

The Kalman filter requires two parameters: the prediction noise and the measurement noise. As in any Kalman filtering problem, if the measurement noise is large compared to the prediction noise, the filter will only slowly follow the measurements and ‘trust’ the predictions more. Here, the individual measurements (gradient constraints) are considered to be rather inaccurate, since they are estimated from the irregular data. Moreover, as described in the introductory parts of Chapter 6, the motion of the spatio-temporal fields is likely to be rather constant over sampling periods (i.e., the Kalman filter steps). Therefore, the Kalman measurement noise has to be larger than the prediction noise, especially as also multiple filter updates per time step occur. Therefore, and since only constant motion is simulated in the following evaluations, the prediction noise Q is set to the zero matrix in the following evaluations, i.e. $Q = 0_{4,4}$.

If not indicated otherwise, the scalar measurement noise of the Kalman filter is calculated with Equation (6.12). In order to receive conservative estimates of the gradient constraint errors, the motion parameters could be set to the maximum possible motion, a parameter that could be derived from domain knowledge, e.g., knowledge on the maximum possible speed of the field. In this work, the parameters are set to the current, simulated motion. Thus, for example, for a field motion of $v_X = -1000m/min$ and $v_Y = 1000m/min$ per time step, a gradient constraint error is computed as $\sigma_{GC}^2 = \sigma_{z'_X}^2 (-1000)^2 + \sigma_{z'_Y}^2 (1000)^2$, where $\sigma_{z'_X}^2$ and $\sigma_{z'_Y}^2$ are derived via variance propagation as described in Section 6.3.3. The variance factor of Equation 6.10 that is required for computing $\sigma_{z'_X}^2$ and $\sigma_{z'_Y}^2$ is estimated by pre-running the filter and observing the value calculated by Equation (6.10) for each node. Then, an average variance factor is used in the simulations. The main factor influencing the variance factor is the field. For the fields simulated in in this work, the pre-running of the filter has delivered a low variance factor of 0.0001, which is used in all simulation presented in subsequent sections.

7.3.4 Results - simulated field

Simulating realistic fields is a complex task. Therefore, the approach taken here is to first simulate a rather simplistic spatio-temporal moving field in order to illustrate the basic properties of the proposed algorithm, such as the behavior under different motion speeds. The simulated field is a Gaussian mixture model (Equation (7.5)).

$$Z(x,y) = \sum_{i=1}^{n_G} e^{-\frac{(x-\mu_{i,x})^2+(y-\mu_{i,y})^2}{2\sigma_i^2}} \quad (7.5)$$

where n_G is the number of Gaussians, $\mu_{i,x}$ is the x-coordinate, $\mu_{i,y}$ is the y-coordinate of the center of Gaussian i , which are drawn randomly from the coordinate space of the study area. σ_i^2 is the variance, which is drawn randomly from the interval $[0, \sigma_{max}^2]$ for each Gaussian. In reality, the fields can be expected to exhibit large parts of zero field values, e.g., periods of no rainfall in between rain clouds. In order to reflect this in the simulations, a minimum possible field value is assumed, and the field is set to zero, if the simulated values fall below this threshold. Due to the linearity assumption of optical flow of Equation (2.38) and the methodology for derivative estimation, the motion estimation works well when the field is approximately linear at the sites where the gradient constraints are constructed, and the reachable accuracy depends largely on the degree of field linearity. Therefore, two different fields are simulated that differ in their degree of linearity: the more linear one is simulated using a low number of Gaussians with a large maximum variance σ_{max}^2 . A more diverse, non-homogenous field is simulated using a high number of Gaussians with a small maximum variance σ_{max}^2 . Examples of the resulting fields are displayed in Figure 7.35a,b.

Different motion behaviors of the field can be simulated by moving the center coordinates of the Gaussians ($\mu_{i,x}$ and $\mu_{i,y}$) through the study area. In order to be able to simulate an unlimited number of time steps, a Gaussian that leaves the area to the west re-appears to the right, etc. In general, the motion behavior can be distinguished along the properties of motion behavior in time (coherence) and motion behavior in space (uniformity). Although the algorithm allows for changing motion in space and time, in this work it is assumed that the motion is both constant in time and uniform space.

The visualization displayed in Figure 7.35 shows a single snapshot of the simulations. Two different fields are displayed, a more homogenous field (a) and a more heterogeneous field (b). Figure 7.35(c) displays the sensor nodes (white color) and the motion vectors (black arrows) per time step, estimated with the proposed algorithm. Figure 7.35(d) displays the true constant motion vector at the node position underlying the simulated displacement of the field.




Figure 7.35: Homogenous field (a); and more heterogeneous, non-linear field (b). Randomly-distributed set of nodes with simulated spatio-temporal field and estimated (c) and true (d) field motion per time step. Red indicates high, blue zero field values.

The results provided in the subsequent sections show time series of the Kalman filter motion state variables. The state variables are not shown directly but the error measures introduced previously in Section 7.3.2. For investigating the performance of the algorithm, 100 time steps of a constantly- and uniformly moving field are simulated, and the estimation errors are aggregated over all nodes for each time step. The goal of the simulations is to illustrate the general behavior of the Kalman filter over time, which is a common strategy in the analysis of Kalman filter applications. However, it is clear that under different settings (fields, node densities, motion characteristics, filter parameters), the filter might behave in a different way and converge faster or slower to the true constant motion, or does not converge at all, when the motion is not constant in time.

Influence of field linearity, field speed and node density

In the first experiment, the influence of the field, the motion speed and the node density is investigated. In Figure 7.36, the motion estimation error is displayed as a function of time for two different spatio-temporal fields and the two different motion speeds: fast motion ($V_x = -1000$ m/min and $V_y = 1000$ m/min) and slow motion ($V_x = -100$ m/min and $V_y = 100$ m/min), as well as the two node network densities in a random node deployment and different UDG communication ranges r (the deployment strategy displayed in Figure 7.34a).

Over time, the motion estimation error for angle and speed decreases for both deployments, dense and sparse (upper and lower row). For slowly moving fields (black and red solid lines) and a dense deployment (upper row), the angular and speed errors decrease fast to low values. For a sparse deployment (lower row), the speed error decreases fast while the angular error is still rather large after the 100 time steps. Thus, a denser deployment with shorter communication ranges is beneficial for motion estimation (upper row in Figure 7.36). When the field is moving fast (dotted lines), the angular error and speed bias increases and the deployment strategy does not have a large influence anymore.

Further, the more homogeneous and linear the field (black lines), the lower the errors (the black lines always show lower errors than the red lines). Both influences, field linearity and motion speed, are strongly related: the more homogeneous and linear the field, the less influential is the motion speed on the estimation accuracy. Only when the field speed is large and the field is heterogeneous

and strongly non-linear (dotted red lines), the estimation accuracy decreases and even a dense node deployment can not prevent large errors.




Figure 7.36: Comparison of mean angular error and mean speed bias for different motion speeds, two field types and two different deployment densities and communication ranges. (a) Mean angular error per time step for $n = 400$ and $r = 500$ m; (b) mean speed bias per time step for $n = 400$ and $r = 500$ m; (c) mean angular error per time step for $n = 25$ and $r = 2000$ m; (d) mean speed bias per time step for $n = 25$ and $r = 2000$ m.

It is likely that the underestimation of speed for large motions is due to the increased invalidity of the Taylor expansion of Equation (2.38) for large motion vectors. Then, the estimated spatial derivatives seem to overestimate the change in field values for the length of the displacement (see also the last paragraph of Section 6.3.4 for a more detailed elaboration on this topic).

Influence of deployment strategy

For evaluating the influence of node deployment on motion estimation results, a homogeneous field is simulated moving slowly over the study area and a dense node deployment, i.e., $n = 400$, with a communication range of $r = 500$ m. The performance of the different deployment strategies shown previously in Figure 7.37 is compared.

The results show that the algorithm performance does not depend on the deployment of stationary nodes. Although, as expected, a grid-based deployment is slightly advantageous, the difference is not considered significant, at least for the rather high number of nodes that has been tested. With a lower number of nodes and the increasing risk of a disconnected graph when the nodes are randomly deployed, a grid-based deployment might become advantageous. Motion estimation




Figure 7.37: Comparison of mean angular error and mean speed bias for different geosensor network (GSN) deployments. (a) Mean angular error per time step; (b) mean speed bias per time step.

with a VANET results in significantly larger errors, potentially due to the increased invalidity of the estimate of the temporal derivative: since the cars move, the simple estimate of a temporal derivative as the difference between current and previous measurement of a single node (car) gets increasingly invalid with increasing car speed. Figure 7.37 displays this effect: the slower the cars (the lighter the blue curves), the more accurate the motion estimates. Surprisingly, stationary cars with zero speed (light blue curve) also provide large errors, although they are quite similar to a random deployment of stationary nodes. This can be explained by the spatial configuration and the sampling on the road network: due to the assumed communication range of $r = 500$ m, cars usually communicate only when driving on the same road. Therefore, the samples for derivative estimation are often aligned linearly along the road, which is disadvantageous for motion estimation. This is a problem that is very likely to occur in reality, as well.

Influence of Kalman measurement noise

For investigating the Kalman measurement noise parameter, a GSN of $n = 25$ nodes with a UDG communication distance of $r = 2000$ m is simulated. The nodes are distributed according to a diffusion-based deployment with a minimum allowed node distance of 1000 m and a maximum node distance equal to the UDG communication range $r = 2000$ m. The field motion behavior is randomly chosen from the small interval $V_X \in [-500m/t, 500m/t]$ and $V_Y \in [-500m/t, 500m/t]$. Different fixed Kalman measurement noises are tested, with an equal weighting of the directional derivatives (i.e., $\epsilon_{ij} = 1$ in Equation (6.7) for all pairs of nodes). Further, the Kalman measurement noise estimated with the methodology described in Section 6.3.4 is tested. In order to quantify the terms of Equation (6.12), maximum possible field motions $V_X = 1000$ and $V_Y = 1000$ are assumed. As described previously, the required variance factor is estimated by pre-running the simulations and calculating an average variance factor using Equation (6.10). The results presented in Figure 7.38 are averages of ten runs of simulations with a different spatio-temporal field, different motion and different network each time.




Figure 7.38: Comparison of the angular error and speed bias for different Kalman filter measurement noises. (a) Mean angular error; (b) mean speed bias; (c) standard deviation angular error; (d) standard deviation speed bias.

The derivation of gradient constraint errors from spatial configuration seems to be beneficial for motion estimation: the red and black curves mostly lie below the others for both error measures and both, mean values and standard deviations. Further, the provided formula for deriving the gradient constraint error (Equation (6.12)) gives a straight-forward way of calculating the required Kalman measurement noise parameter. Without it, the specification of the measurement noise has to rely on rather arbitrary considerations. It can be seen from the comparison of the gray and dotted curves that the estimation accuracy does not change much with different fixed Kalman measurement noises, except for very large noise values. The weighting of the directional derivatives seems to provide a slight advantage over a uniform weighting: the red curve is mostly below the black curve. The theoretical considerations of Section 6.3.4 provided a discussion on the influence of the derivatives on the motion estimates and concluded that, when the estimated spatial derivatives are underestimates of the true spatial derivatives, the motion speed will be overestimated. Since the speed bias is positive (Figure 7.38b), it is likely that this underestimation of the true derivatives occurs here. However, whether the method under- or overestimates the derivatives on average

depends on numerous factors such as the nature of the field, the node distance, etc. Therefore, it is not further discussed here.

7.3.5 Results - radar field

In a second experiment, the algorithm is evaluated on a real precipitation field derived from the Hanover weather radar for the same study area as before (Figure 7.33). Six hours of a period of rather strong rainfall (19 July 2012, 4 a.m. to 10 a.m.) are chosen. The raw radar data (i.e., reflectivities) were preprocessed by using the methodology described previously and in Berndt et al. (2014), resulting in precipitation values in the unit mm/h on a regular grid with a resolution of $1 km^2$ (10×10 pixels in the study area). For this experiment, some more preprocessing is done on the grid. First, 1-min grids are calculated from the 5-min snapshots using the methodology described in Section 4.1.2, because the motion in between two 5-min snapshots would be too large for the sensor spacings in the study area. In reality, this would correspond to an adjustment of the sampling rate of the sensor nodes, e.g. from 5-min sampling to 1-min sampling. Further, the 1-min grids are smoothed by a 3×3 arithmetic mean filter, because the previous experiments of Section 7.1 have shown that the radar smoothing seems to provide more realistic precipitation fields that are higher correlated with ground measurements. The sparse deployment introduced previously is used with $n = 25$ nodes and a communication distance of $r = 2000m$. A characterization of the event and the derived motion information is displayed in Figure 7.39. The plots display mean values of rainfall (a), motion direction (b) and motion speed (c) in the study area.




Figure 7.39: Characterization of the radar event. (a) Time series of radar estimated rainfall averaged over the study area; (b) time series of mean motion direction calculated by image-based optical flow; (c) time series of the mean motion speed calculated by optical flow. Time steps of no data values result from the inability of the optical flow algorithm to detect field motion when there is no rainfall.

Figure 7.39a shows that a number of rain clouds passes the study area within the six hours, with two very strong rainfall periods. The field moves with a rather constant motion from west to (slightly north) east (a sudden change of direction in the middle of the period occurs, which is potentially due to radar clutter). The average motion speed is at around 1000 m per minute, but decreases for periods of low rainfall, which is an artifact of the optical flow algorithm that, by computing motion from a pair of adjacent 5-min radar snapshots, does not consider temporal coherence of motion. In Figure 7.40, time series of 1-min radar snapshots and motion vectors estimated with the proposed algorithm are displayed.




Figure 7.40: Time series of radar images separated by 1 min from the upper left to the lower right for the first rain cloud in the event (1-min time steps 14 to 23). Simulated sensor network of $n = 25$ nodes, a communication range of $r = 2000$ m and estimated motion vectors.

Before the first cloud reaches the study area, no motion estimation is possible. Over time, the motion vectors of the nodes approach the cloud motion direction of east-north-east (Figure 7.39b) and some of the nodes also the cloud speed of 1000 m/min (Figure 7.39c), which is equal to the width or height of a pixel and thus, easily recognizable in Figure 7.40. This example shows, that the algorithm is theoretically capable of estimating the motion of a real field. However, only a real deployment of a GSN measuring a moving field can prove the applicability of the algorithm in practice.

7.3.6 Summary

The results show that the decentralized field motion estimation by the irregularly-distributed nodes of a GSN is indeed possible.

The simulations have shown that a dense node deployment is beneficial for motion estimation. Further, the combination of field linearity and motion speed is an important factor and the reachable accuracy decreases with increasing non-linearity and motion speed (Figure 7.36). This is a fact that is known from work on image-based optical flow, but has direct implications in a GSN setting where the deployment and sampling rate of the nodes can be controlled to a certain extent.

When the number of nodes and communication distance are held fixed, the deployment strategy of stationary nodes does not have a large influence on the motion estimation results (Figure 7.37). However, with decreasing node density, the performance of a random deployment will certainly decrease, since there will be disconnected nodes. Further, the algorithm has been developed for stationary nodes. Nonetheless, it can also be applied in a non-stationary setting, e.g., for cars, with slight modifications. However, then the motion estimation accuracy decreases (Figure 7.37), most likely due to the increased error in the temporal derivative estimate and the linear alignment of the cars along roads. Thus, for cars or non-stationary nodes in general, the algorithm requires adjustments, e.g. an integration of the estimation of the temporal derivative into the least squares process for derivative estimation.

In addition, it could be shown that the computation of error measures from spatial node configuration improves the motion estimation accuracy (Figure 7.38).

Certainly, the algorithm performance is limited when: (a) the field does not exhibit sufficient intensity structure or intensity changes (the ‘blank wall’ problem described previously); (b) exhibits too much intensity change or too fast motion to be approximated linearly over the length of the displacement by Equation (2.38); (c) changes its structure rapidly in between consecutive sampling periods / time steps, i.e., deviates too much from the basic intensity conservation constraint of OF of Equation (2.37); or (d) the sensor measurements are corrupted by too much noise.

The results of the experiments provided in this section are not considered to be absolute measures of algorithm accuracy due to the numerous influencing factors like field structure, motion behavior and node density and sampling rate. However, they could provide the basis for further investigations, especially when the algorithm is to be applied in a real deployment of a GSN for proving its applicability in practice.

8 Summary and discussion of the research hypotheses

The ongoing advances in hardware, software, communication and positioning technologies described in the introductory parts of this thesis currently lead to increasing amounts of data and ultimately to the possibility of "*knowing where everything is, at all times*" (Goodchild, 2010). While, when phrasing this vision, Goodchild (2010) had single entities or *objects* in mind, such as mobile phones or vehicles, or even animals, the sentence can also be applied to *fields*: with the increasing density of sensors that measure environmental fields at high temporal resolutions and accuracies and the technologies mentioned previously, the chance for knowing the field at all times increases, e.g. knowing precipitation intensities at all times, locations and aggregation periods or data resolutions. However, as long as use cases exist whose requirements on the spatial and temporal coverage, the resolutions and the accuracies can not be met by the underlying sensors and data, sophisticated methods for the estimation of the field at locations without a sensor are necessary.

The first two parts of this thesis dealt with improving the estimation of precipitation intensity at locations without a rain gauge, either by methods of integrating existing data from weather radar and rain gauges or by augmenting them with cars as mobile rain gauges. The third part dealt with estimating the motion behavior of a dynamic field from the sensor data. The following sections provide a concluding discussion on the research hypotheses provided in the introductory parts in Section 1 as well as ideas for future developments. Section 8.4 provides a more generic outlook on future developments and further applications of the proposed methods to other use cases.

8.1 Discussion of research hypotheses 1 and 2: 1-min precipitation intensity estimation

For investigating research hypothesis 1 that dealt with the integration of weather radar and rain gauges at 1-min resolution, different methods for merging radar and rain gauges were compared with methods that estimated precipitation intensity from rain gauges or weather radar only. The merging methods yield higher correlations compared to the rain gauge interpolation methods. The two newly introduced merging methods RADAR-DBC and FF-RK outperform the standard merging method RK. Surprisingly, two radar-only methods provided high correlations in cross-validation: the rather strong smoothing of the radar image over a large pixel window and a time-lagged radar estimate.

However, concerning the other two error measures, the merging methods did not provide a significant improvement compared to the rain gauge interpolation or radar-only methods. Nonetheless, for samples larger-zero, the cross-validation results show that the merging methods are able to reduce the strong underestimation of intensities by the rain gauge interpolation methods and pro-

vided the lowest bias of all three categories. Again, the two new methods RADAR-DBC and FF-RK mostly outperformed the standard merging method RK.

Thus, although promising, the merging methods and especially the newly introduced methods RADAR-DBC and FF-RK require deeper investigations. Both methods have been a first step towards the dynamic, time-dependent computation of space-time lags and radar-rain gauge relationships to be applied in an operational (automatic) setting. However, both relied on the generation of 1-min radar snapshots out of 5-min snapshots and on the computation of cross-correlations between the rain gauges and the resulting 1-min radar time series. As the 5-min to 1-min up-sampling process might introduce errors to the 1-min snapshots, depending on, for example, the performance and accuracy of the used optical flow algorithm, another promising approach could be to accumulate the 1-min rain gauge values to 5-min values over different accumulation intervals (same as radar, 1-min after radar, 2-min after radar, etc) and then compute correlations for the 5-min accumulations.¹ In addition, existing radar pre-processing algorithms could be applied beforehand, for example to reduce the effects of orographic shading. Finally, the methods based on statistical models and correlations presented in this work could be augmented with physical models, forming what is called a scientific-statistical model by Cressie & Wikle (2011).

For the investigation of research hypothesis 2 that was concerned with the integration of motion information into the estimation process, three different motion-based estimation methods were tested. All of them relied on motion information estimated by optical flow from the weather radar images. Two of them, FF-IDW and FF-OK were based on rain gauge measurements only and applied a pure spatial interpolation to rain gauge samples relocated in space according to the motion. The method FF-RK combined this relocation with a radar-rain gauge merging method. The two rain gauge interpolation methods FF-IDW and FF-OK show that the integration of motion information into the process is beneficial: the two methods mostly outperformed standard spatial or space-time symmetric spatio-temporal rain gauge interpolation methods. Surprisingly, the deterministic FF-IDW method showed better performance than its geostatistical counterpart FF-OK. The motion-based merging method FF-RK showed promising results and outperformed standard regression kriging for most of the error measures.

The analysis of the radar-only methods has shown that smoothing the radar images is always beneficial. Further, a strong temporal lag between radar and rain gauges could be identified: the highest correlations between radar and rain gauges could be observed for rain gauge measurements 6 minutes after the radar snapshot. Most likely, the lag is due to the differences in measurement procedures and seems to intensify with increasing distance from the radar due to beam elevation. Therefore, the identified time lag of 6 minutes is only valid for the radar-rain gauge configuration that has been investigated in this work and might be different for other study areas and spatial radar-gauge configurations. These results show that a simple snapshot-based merging of weather radar and rain gauges is not sufficient at 1-min temporal resolution. This finding is supported by the rather poor performance of the snapshot-based radar-rain gauge merging method RK.

¹ An idea which has been proposed to the author by Dr. Stefan Krämer from the Institute for Technical and Scientific Hydrology, Hannover, (itwh GmbH) in a personal communication and which will be investigated further in future projects which will be carried out together.

Since the methods have been evaluated with gauge measurements that are subject to errors as well, the results do not represent absolute accuracy of estimates of ground precipitation and have to be taken with care. Further, the rather low prediction accuracy of all methods show that the 1-min point estimation estimation of precipitation intensities is a difficult task. A promising approach is to apply further preprocessing to the weather radar images based on known physical properties of the measurement procedure. Nonetheless, the best performing prediction methods have been selected and used as 'ground truth' in the field experiments with cars, which are summarized in the next section.

8.2 Discussion of research hypothesis 3: precipitation estimation with cars

In this part of work, the estimation of precipitation intensity by cars was investigated in field experiments. Cars were equipped with sensors and the collected data was compared with reference precipitation intensities estimated from rain gauges, weather radar or both. Establishing such W-R relationships in the field turned out to be difficult due to the high space-time variability of rainfall, especially strong rainfall, and the errors associated with the reference.

The influence of car speed on the sensor readings could be determined with a high degree of certainty. It turned out that the wiper frequency of cars with discrete wiper classes (operated manually) is strongly coupled with the car speed and seems to be only marginally influenced by the rain amount, resulting in very large standard deviations of rainfall for the different wiper classes. Therefore, such cars are currently not considered suitable for those use cases that require quantitative estimates of precipitation and not just Boolean information on whether it rains or not. The optical sensor readings and the wiper frequency of wipers controlled by an optical sensor showed a clear linear relationship to the car speed. The optical sensor readings, corrected by the car speed influence, show high correlations to reference rainfall estimated from rain gauges or weather radar. For selected references, a W-R relationship and confidence bounds for precipitation intensity predictions could be established. Due to the inaccuracy of the references, the prediction bounds are considered lower bounds for prediction accuracy. Most of the car measurements were collected at rather low rain intensities. Therefore, a reliable analysis could only be made for these low intensities and investigations on the ability of cars for the estimation of high rain intensities remains subject to future investigations.

Interestingly, the temporal lag between cars and radar was significantly different from the one between rain gauges and radar. The reason could be that on average, the locations of the car measurements were much closer to the radar station than the locations of the rain gauges. However, the correlation in between the car measurements and the nearest rain gauge closer than 2 km showed that the rain gauge measurement temporally lagged behind the car measurements roughly the same amount of time (6 minutes) that they lagged behind the radar estimates. Thus, it could also be that the time lags lie in the nature of the measurement process of the rain gauges, a topic that requires further investigation.

Several ideas for future work emerge from the research on cars for precipitation intensity estimation. First, the investigation of the other factors influencing car measurements of precipitation (besides car speed and windscreen angle), described in Section 5, is required. For doing this, a better

'ground truth' is necessary, e.g. in the form of a denser network of rain gauges in conjunction with significantly more data of a large-scale experiment, especially at high rain intensities relevant for the imagined hydrological use cases. For investigating the relationship between car and radar data and potential time lags between both, sufficient car data at different radial distances from the radar station would be beneficial.

Different methods for investigating the performance of cars compared to a conventional rain gauge network for areal rainfall estimation and hydrological use cases can be imagined. One possibility is already provided by Rabiei et al. (2016), who investigated this question in computer simulations and compared the performance of a conventional rain gauge network with the performance of cars simulated on a road network, where both, rain gauge and car measurements were sampled from weather radar estimates as 'ground truth' precipitation intensities. Another approach would be to record data from a sufficient number of cars driving *simultaneously* in a certain study area, e.g., a catchment area, and to repeat the studies of Rabiei et al. (2016) using real data of cars, rain gauges and weather radar.

8.3 Discussion of research hypothesis 4: decentralized motion estimation

This part of work was concerned with the development of an algorithm for the estimation of the motion of spatio-temporal moving fields by the nodes of a GSN. A well-known optical flow algorithm was used as the basis and has been adjusted to the specifics of GSN, e.g., the irregular distribution of nodes and the strong resource constraints. The proposed algorithm has been formalized in a decentralized fashion, and node pseudo-code has been provided in the form of a decentralized algorithm protocol (Section 6.5). Due to its two-part nature, the algorithm can also be adjusted such that all information is routed to a central node responsible for the motion estimation. The performance of the algorithm has been evaluated using simulated fields, as well as precipitation fields derived from weather radar. The simulations, especially on the real precipitation field measured by weather radar, have shown that the decentralized field motion estimation by the irregularly-distributed nodes of a GSN is indeed possible.

Future work on the decentralized motion estimation includes the analysis of the motion state error of the Kalman filter, which is an important measure in practice. In addition, error-in-variables models could be investigated for replacing the Kalman filter for motion estimation, e.g., a total least squares method accounting for errors in the estimates of spatial derivatives directly. Further, simulations with weather radar or other measured fields for longer time periods, such as days or weeks, are necessary in order to evaluate the behavior of the algorithm under changing motion conditions. The simulations of VANETs have shown that in order to achieve satisfying and reliable results with moving sensors, more work is required on the algorithm. Finally, the ultimate goal is to replace the simulated network with a real deployment of a GSN and real measurements of a moving atmospheric or oceanographic field.

8.4 Outlook

The main goal of the methods for the 1-min point estimates of precipitation intensity was to derive valid reference intensities for the collected car data. However, the methods for integrating data from the three different types of precipitation sensors (rain gauges, weather radar and cars) might provide valuable information for other use cases that require point estimates of precipitation intensity at high resolutions. In addition, not only the methodologies but also the results provided in Chapter 7.1 could be valuable input to further investigations in other hydrological application areas that require precipitation intensity estimations at high spatial and temporal resolutions, such as urban drainage management or flood forecasting. The results of the car data analyses might be valuable for future efforts that investigate the idea of using cars as mobile rain gauges on a larger scale.

Further, the observed increased accuracy of motion-based rain gauge interpolation methods is promising also for use cases and areas where no weather radar but only rain gauges exist, and where the motion information could still be derived by other means than optical flow, e.g. by wind measurements. It is also promising for other fields such as those where remote sensing of the field is not possible, for example, fields of air pollutants. Furthermore, the thesis has delivered a possibility of estimating the motion behavior of the field from the in-situ sensor data directly. When this is done, the motion information could be estimated first and then be used to improve estimation accuracy afterwards.

Finally, the ultimate goal for future applications of the methods can be summarized as follows. First, the number of possibly cheap and inaccurate sensors measuring a field in-situ could be increased, which then provide the data for the estimation of properties of the field, such as information on the motion behavior. Finally, this information could be used to reconstruct the field with a higher accuracy at locations without a sensor.

9 Appendix

This section contains a discussion on the properties of 'frozen field' distance function that is applied at different places throughout this work (Section 9.1). Then, the executable Kalman filter equations for the motion estimation algorithm are provided (Section 9.2) and the properties of controllability and observability of the filter are discussed (Section 9.3).

9.1 Discussion on the 'frozen field' distance function

In this section some properties of the 'frozen field' distance function are discussed. The 'frozen field' distance function as used in this work is defined as:

$$d_{ff}(u_i, u_j) \equiv \left\| \begin{bmatrix} \Delta x_{ij} \\ \Delta y_{ij} \end{bmatrix} - V_f \Delta t_{ij} \right\| \quad (9.1)$$

where $\| \cdot \|$ is the Euclidean metric on \mathbb{R}^2 and V_f is a known motion vector of the field to be interpolated. The requirement of motion constancy when $d_{ff}(u_i, u_j)$ is used as parameter in spatial covariance functions is known and widely used in geostatistical literature, e.g. in the book of Cressie & Wike (2011) or by Gneiting et al. (2006). However, to the best of the author's knowledge, there has been no deeper or formal discussion on the model concerning non-constant motion vectors. In the following, a formal (mathematical) discussion on the properties of d_{ff} is provided with two main outcomes:

- d_{ff} violates the requirement for metrics that $d_{ff}(u_i, u_j) = 0 \Leftrightarrow u_i = u_j$, even when the motion is constant.
- d_{ff} violates different requirements for a metric if used with non-constant motion. Which requirements are violated depends on the motion model, i.e. how motion vectors are assigned to spatio-temporal locations.

These results are important when the d_{ff} is used for prediction with IDW or OK. Concerning OK, the results on the properties of d_{ff} concerning metrics gain importance in conjunction with the findings of Curriero (2006), who shows that common geostatistical covariance functions might become invalid when a distance function is used as parameter that does not satisfy the conditions of a metric.

First, the definition of a metric repeated: d_{ff} is called a metric iff (if and only if):

- a) $(d_{ff}(u_i, u_j) \geq 0) \wedge (d_{ff}(u_i, u_j) = 0 \Leftrightarrow u_i = u_j)$
- b) $d_{ff}(u_i, u_j) = d_{ff}(u_j, u_i)$ (Symmetry)

$$c) d_{ff}(u_i, u_j) \leq d_{ff}(u_i, u_k) + d_{ff}(u_k, u_j) \text{ (Triangle Inequality)}$$

Using the reformulation of Equation 2.23 that $d_{ff}(u_i, u_j) = d_{eucl}(s_i, s_j + V_f \Delta t_{ij})$, the conditions can be reformulated:

$$a) d_{eucl}(s_i, s_j + V_f \Delta t_{ij}) \geq 0) \wedge (d_{eucl}(s_i, s_j + V_f \Delta t_{ij}) = 0 \Leftrightarrow s_i = s_j + V_f \Delta t_{ij})$$

$$b) d_{eucl}(s_i, s_j + V_f \Delta t_{ij}) = d_{eucl}(s_j + V_f \Delta t_{ij}, s_i) \text{ (Symmetry)}$$

$$c) d_{eucl}(s_i, s_j + V_f \Delta t_{ij}) \leq d_{eucl}(s_i, s_k + V_f \Delta t_{ik}) + d_{eucl}(s_k, s_j + V_f \Delta t_{kj}) \text{ (Triangle Inequality)}$$

The validity of each of these conditions is discussed in the following.

Condition a)

With the reformulation, it is obvious that $d_{ff}(u_i, u_j) \geq 0$ holds. The condition $(d_{ff}(u_i, u_j) = 0 \Leftrightarrow u_i = u_j)$ or reformulated $(d_{eucl}(s_i, s_j + V_f \Delta t_{ij}) = 0 \Leftrightarrow s_i = s_j + V_f \Delta t_{ij})$ holds for two cases:

1. $u_i = u_j$, then $\Delta t_{ij} = 0$ and $d_{eucl}(s_i, s_j + V_f \Delta t_{ij}) = d_{eucl}(s_i, s_j)$ and since $s_i = s_j$ (follows from $u_i = u_j$), $d_{eucl}(s_i, s_j) = d_{eucl}(s_i, s_i) = 0$.
2. $u_i \neq u_j \wedge s_i = s_j + V_f \Delta t_{ij}$, then $d_{eucl}(s_i, s_j + V_f \Delta t_{ij}) = d_{eucl}(s_i, s_i) = 0$

The fact that these two cases exist makes $d_{ff}(u_i, u_j) = 0 \Leftrightarrow u_i = u_j$ valid but renders $d_{ff}(u_i, u_j) = 0 \Rightarrow u_i = u_j$ invalid. Thus, it could happen that for a particular estimation location, both cases hold. this is exactly the case when estimation is to be performed at a measurement location where case 2 holds as well, i.e. there exists another measurement that is relocated to that exact location by d_{ff} . Then, the formulation of the IDW interpolation of Equation 2.28 is no longer valid as the second case in the IDW case distinction (i.e., that there is some location for which the distance is zero) holds for both. The method must then decide, which value to use for estimation. In the case that V_f is not constant, case 2 could also hold for more than one measurement location u_j . This could also happen when the motion vector V_f is constant but the measurement locations move. Therefore, theoretically, the IDW method must be reformulated and be more explicit on how to cope with multiple measurements for which the distance becomes zero. For OK, this is more involved, as the distance function is used for populating the ordinary kriging matrices. Thus, when ordinary kriging is applied with d_{ff} as distance function, the interplay between spatial configuration of the sampling locations and the motion vector requires some attention.

Condition b), Symmetry

The symmetry of the 'frozen field' distance function $d_{ff}(u_i, u_j)$ is shown in the following:

$$d_{st,ff}(u_i, u_j) = d_{eucl}(s_i, s_j + V_f \Delta t_{ij})$$

$$= d_{eucl}(s_i - V_f \Delta t_{ij}, s_j) // \text{equality straightforward to show geometrically or algebraically}$$

$$= d_{eucl}(s_j, s_i - V_f \Delta t_{ij}) // \text{Symmetry of Euclidean metric}$$

$$= d_{eucl}(s_j, s_i + V_f (-\Delta t_{ij}))$$

$$\begin{aligned}
&= d_{eucl}(s_j, s_i + V_f(\Delta t_{ji})) // \text{using the equality } -\Delta t_{ij} = \Delta t_{ji} \\
&= d_{st,ff}(u_j, u_i)
\end{aligned}$$

For non-constant vectors, this symmetry not necessarily holds. For example, when the motion is associated with one of the two spatio-temporal locations, the motion vector used for $d_{ff}(u_i, u_j)$ might differ from the motion vector used for $d_{ff}(u_j, u_i)$, implying that $d_{ff}(u_i, u_j) \neq d_{ff}(u_j, u_i)$.

Condition c), Triangle Inequality

In order to prove $d_{eucl}(s_i, s_j + V_f \Delta t_{ij}) \leq d_{eucl}(s_i, s_k + V_f \Delta t_{ik}) + d_{eucl}(s_k, s_j + V_f \Delta t_{kj})$, it is sufficient to show that: $d_{eucl}(s_k, s_j + V_f \Delta t_{kj}) = d_{eucl}(s_k + V_f \Delta t_{ik}, s_j + V_f \Delta t_{ij})$. Then, the relevant points are equal in \mathbb{R}^2 and the triangle inequality concerning spatio-temporal locations is reduced to the triangle inequality of the spatial Euclidean metric, which obviously holds for constant motion V_f : $d_{eucl}(s_i, s_j + V_f \Delta t_{ij}) \leq d_{eucl}(s_i, s_k + V_f \Delta t_{ik}) + d_{eucl}(s_k + V_f \Delta t_{ik}, s_j + V_f \Delta t_{ij})$.

From the definition of temporal differences $\Delta t_{kj} = t_k - t_j$, $\Delta t_{ik} = t_i - t_k \Leftrightarrow t_k = t_i - \Delta t_{ik}$ and $\Delta t_{ij} = t_i - t_j \Leftrightarrow t_j = t_i - \Delta t_{ij}$, it follows that $\Delta t_{kj} = (t_i - \Delta t_{ik}) - (t_i - \Delta t_{ij}) = t_i - \Delta t_{ik} - t_i + \Delta t_{ij} = -\Delta t_{ik} + \Delta t_{ij}$. With this equality, the following transformation of $d_{eucl}(s_k, s_j + V_f \Delta t_{kj})$ can be performed:

$$\begin{aligned}
d_{eucl}(s_k, s_j + V_f \Delta t_{kj}) &= \sqrt{(x_k - (x_j + v_x \Delta t_{kj}))^2 + (y_k - (y_j + v_y \Delta t_{kj}))^2} \\
&= \sqrt{(x_k - x_j - v_x \Delta t_{kj})^2 + (y_k - y_j - v_y \Delta t_{kj})^2} // \text{expanding} \\
&= \sqrt{(x_k - x_j - v_x (-\Delta t_{ik} + \Delta t_{ij}))^2 + (y_k - y_j - v_y (-\Delta t_{ik} + \Delta t_{ij}))^2} // \text{using the equality } \Delta t_{kj} = -\Delta t_{ik} + \Delta t_{ij} \\
&= \sqrt{(x_k - x_j + v_x \Delta t_{ik} - v_x \Delta t_{ij})^2 + (y_k - y_j + v_y \Delta t_{ik} - v_y \Delta t_{ij})^2} // \text{expanding} \\
&= \sqrt{(x_k + v_x \Delta t_{ik} - x_j - v_x \Delta t_{ij})^2 + (y_k + v_y \Delta t_{ik} - y_j - v_y \Delta t_{ij})^2} // \text{rearranging} \\
\|(s_k + V_f \Delta t_{ik}) - (s_j + V_f \Delta t_{ij})\| &= d_{eucl}(s_k + V_f \Delta t_{ik}, s_j + V_f \Delta t_{ij})
\end{aligned}$$

Therefore, the triangle inequality holds for constant motion vectors. However, already the definition of the triangle inequality shows that when arbitrary motion vectors are allowed, the inequality not necessarily holds anymore.

Summary

This section has shown that under assumption of a constant motion vector V_f , the 'frozen field' distance function violates only the first requirement on metrics and satisfies the others. Thus, when the unlikely case of violation does not occur (i.e. the existence of a measurement that is relocated to another measurement location), d_{ff} satisfies the conditions of a metric. Further, the

metric conditions are violated, when non-constant motion vectors are used. Therefore, together with the results of Curriero (2006), the covariance functions might become invalid.

9.2 Executable Kalman filter equations for the motion estimation algorithm

This section introduces the executable Kalman filter equations for the motion estimation. It follows standard implementations as e.g. described in Welch & Bishop (2006). The Kalman state of a node n_i at time t contains the estimates of motion and acceleration variables:

$$\hat{v}_i(t) = \begin{bmatrix} \hat{v}_X \\ \hat{v}_Y \\ \hat{\dot{v}}_X \\ \hat{\dot{v}}_Y \end{bmatrix} \quad (9.2)$$

Prediction

Predicting the motion for a time step t from a previous time step $t - 1$ involves applying the prediction matrix F of Equation 6.16 to the a-posteriori state (i.e., the state after measurement update) of time step $t - 1$.

$$\hat{v}_i(t)^- = F\hat{v}_i(t-1)^+ \quad (9.3)$$

Further, the state error covariance is predicted according to:

$$P_i(t)^- = FP_i(t-1)^+F^T + Q \quad (9.4)$$

Update

Updating the Kalman state involves the calculation of a weighted average of prediction and measurement:

$$\hat{v}_i(t)^+ = \hat{v}_i(t)^- + K_i(t)(\hat{z}'_{i,T}(t) - H_i(t)\hat{v}_i(t)^-) \quad (9.5)$$

where $\hat{v}_i(t)^-$ is the predicted state at time t before update, $\hat{z}'_{i,T}(t)$ is the estimated temporal derivative at time step t and $H_i(t)$ is the matrix of estimated partial derivatives:

$$H_i(t) = \begin{bmatrix} \hat{z}'_{i,X} & \hat{z}'_{i,Y} & 0 & 0 \end{bmatrix} \quad (9.6)$$

The Kalman gain matrix K is computed as:

$$K_i(t) = P_i(t)^- H_i(t)^T (H_i(t) P_i(t)^- H_i(t)^T + R)^{-1} \quad (9.7)$$

Finally, the state covariance is updated according to:

$$P_i(t)^+ = (I - K_i(t)H_i(t))P_i(t)^- \quad (9.8)$$

9.3 Controllability and Observability of the Kalman filter for motion estimation

For linear dynamic systems such as the Kalman filter, the properties of controllability and observability first introduced by Kalman (1960) and extensively discussed e.g. in Gopal (1993) are essential. While a system is controllable if the control inputs can be used to transform the system to a desired state (Southall et al., 1998) via the dynamic model, observability concerns the actual possibility to estimate all state variables from the measurements. Since the Kalman filter used for motion estimation differs from standard Kalman filter applications, it is not a-priori clear if the proposed filter is controllable and observable. Thus, the following sections contain an informal discussion on controllability and observability of the proposed Kalman filter for field motion estimation.

Controllability

Here the informal characterization of controllability of Gopal (1993) or Southall et al. (1998) is used:

"A system is said to be controllable if every state vector can be transformed to a desired state by the application of unconstrained control inputs." (Southall et al., 1998).

As no control input for the filter is yet considered, the system is not controllable. However, control inputs could be added to the filter, e.g. in the form of wind speed measurements. As mentioned by Southall et al. (1998), controllability also refers to the noise:

For the current filter at hand, the update of state covariance matrix P using $FP_i(t-1)^+F^T$ of Equation 9.4 does not alter the variances of the acceleration variables (i.e., the elements on the diagonal of P containing the variances of the acceleration variables). Then, if the matrix Q does not contain noise variances for the acceleration variables, those variances will not increase (Southall et al., 1998) and the filter will eventually "drive these elements to zero, after which point new observations will not alter the estimate of any uncontrollable states." (Southall et al., 1998). This effect is used in the application of the filter by setting the variances of the acceleration variables to zero.

Observability

From the set-up of the Kalman filter, it is clear that the motion components can only indirectly be "measured" by the gradient constraints, and the change in motion (the acceleration variables) can not be measured at all. Therefore, the question arises whether the proposed 4-element Kalman state vector consisting of the motion components and the derivatives is actually observable, i.e. whether all of the state variables can actually be estimated from the available measurements. While observability for time-invariant systems can be defined independent of concrete time steps,

observability for time-varying systems is defined for certain time periods, as it depends on the time-dependent prediction and update matrices. The common definitions of observability concern systems where either both matrices F and H are time-invariant or time-varying and only single updates per time step occur. In this work, the prediction matrix F is time-invariant, the update matrix H is time-dependent and multiple updates per time step occur. Therefore, here, the common definition of observability for time-independent systems, e.g. the one provided in Southall et al. (1998), is used and adjusted to this specific case. The definition provided by Southall et al. (1998) reads:

"A system is said to be observable at a time step t_0 if for a state $x(t_0)$ at that time, there is a finite $t_1 > t_0$ such that knowledge of the outputs z from t_0 to t_1 are sufficient to determine state t_0 ." (Southall et al., 1998)

Usually, proving observability then uses the so-called observability matrix and defines:

"A system with state vector x of dimension n is observable if the observability matrix O has row rank n (i.e. n linearly independent rows)" (Southall et al., 1998)

The observability matrix O is built from the measurement and prediction matrices. For example, for time invariant systems (those where measurement matrix H and prediction matrix F are constant in time), the observability matrix is (Southall et al., 1998):

$$O = \begin{bmatrix} H \\ HF \\ \dots \\ HF^i \\ \dots \\ HF^{n-1} \end{bmatrix} \quad (9.9)$$

where n is the number of state variables. The observability matrix O relates current measurements to the current state via the measurement matrix H and relates future measurements to that state via the prediction matrix. It is then proven, e.g. in Gopal (1993), that if O has n linearly independent rows, the system is observable, i.e. all state variables can be reconstructed from the measurements.

Observability of the Kalman filter for motion estimation

The form of the Kalman filter proposed in this work differs from the time-invariant case in that the measurement matrix is time-dependent since it contains the currently measured partial derivatives. Further, multiple measurement updates per time step occur, one for each neighboring node. Thus, observability can not be proven for the whole system but only for particular time steps and nodes and clearly depends on the actual partial derivatives to be integrated. Therefore, the goal is to show that it is theoretically possible to reconstruct the state at a particular time step t_0 and node n_i from the measurements collected within a particular time interval defined by time steps t_1 and t_0 with $t_1 > t_0$.

For showing this, the observability matrix O_i of a particular sensor node n_i is reformulated: it is now a function of time and includes the partial derivatives measured by the neighboring nodes n_j to n_m as well. For example, for a Kalman state with only the two motion components:

$$\hat{v}_i(t) = \begin{bmatrix} \hat{v}_X \\ \hat{v}_Y \end{bmatrix} \quad (9.10)$$

and no motion change variables (motion derivatives), the number of state variables is $n = 2$ and the prediction matrix F is the identity matrix ($F = I$). Then, the observability matrix O_i of a node n_i can be constructed in a similar way as e.g. done in Southall et al. (1998) and described above. However, here, it additionally contains the measurement matrices of all neighboring nodes n_j to n_m . For a particular time step t_0 and another time step $t_1 \geq t_0$, the observability matrix becomes.

$$O_i(t_0) = \begin{bmatrix} H_i(t_0) \\ H_j(t_0) \\ \dots \\ H_m(t_0) \\ \dots \\ H_i(t_1) \\ H_j(t_1) \\ \dots \\ H_m(t_1) \end{bmatrix} = \begin{bmatrix} -\hat{z}'_{i,X}(t_0) & -\hat{z}'_{i,Y}(t_0) \\ -\hat{z}'_{j,X}(t_0) & -\hat{z}'_{j,Y}(t_0) \\ \dots & \dots \\ -\hat{z}'_{m,X}(t_0) & -\hat{z}'_{m,Y}(t_0) \\ \dots & \dots \\ -\hat{z}'_{i,X}(t_1) & -\hat{z}'_{i,Y}(t_1) \\ -\hat{z}'_{j,X}(t_1) & -\hat{z}'_{j,Y}(t_1) \\ \dots & \dots \\ -\hat{z}'_{m,X}(t_1) & -\hat{z}'_{m,Y}(t_1) \end{bmatrix} \quad (9.11)$$

This forms a linear system of the form $O_i(t_0) \times \hat{v}_i(t_0) = \hat{z}'_T$, where \hat{z}'_T is the vector of partial derivatives in time. Now it becomes obvious that even when only data from a single time step is available, i.e. $t_1 = t_0$, the system might be observable at t_0 , i.e. the two-element state vector $\hat{v}_i(t_0)$ can be estimated, as long as there are at least two linearly independent rows available.

For the 4-element Kalman state including motion acceleration variables, this is different. Then, the state is:

$$\hat{v}_i(t) = \begin{bmatrix} \hat{v}_X \\ \hat{v}_Y \\ \hat{\dot{v}}_X \\ \hat{\dot{v}}_Y \end{bmatrix} \quad (9.12)$$

And the prediction matrix F predicting the state of time step t from time step $t - 1$ becomes:

$$F(t, t-1) = \begin{bmatrix} 1 & 0 & \Delta t & 0 \\ 0 & 1 & 0 & \Delta t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (9.13)$$

From this, the prediction matrix $F(t_1, t_0)$ predicting the state at a time step t_1 from an arbitrary previous time step t_0 (not necessarily the direct predecessor) can be derived:

$$\Rightarrow F(t_1, t_0) = F(t, t-1)^{t_1-t_0} = \begin{bmatrix} 1 & 0 & (t_1 - t_0)\Delta t & 0 \\ 0 & 1 & 0 & (t_1 - t_0)\Delta t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (9.14)$$

The measurement matrix shows that not all of the state variables have a connection to the measurement:

$$H_i(t) = \begin{bmatrix} \hat{z}'_{i,X} & \hat{z}'_{i,Y} & 0 & 0 \end{bmatrix} \quad (9.15)$$

Using the measurement and prediction matrices, the observability matrix for a particular node n_i with neighbors n_j to n_m at a particular time step t_0 is then again reformulated as:

$$O(t_1, t_0) = \begin{bmatrix} \hat{z}'_{i,X}(t_0) & \hat{z}'_{i,Y}(t_0) & 0 & 0 \\ \hat{z}'_{j,X}(t_0) & \hat{z}'_{j,Y}(t_0) & 0 & 0 \\ \dots & \dots & \dots & \dots \\ \hat{z}'_{m,X}(t_0) & \hat{z}'_{m,Y}(t_0) & 0 & 0 \\ \dots & \dots & \dots & \dots \\ \hat{z}'_{i,X}(t_1) & \hat{z}'_{i,Y}(t_1) & (t_1 - t_0)\Delta t \hat{z}'_{i,X}(t_1) & (t_1 - t_0)\Delta t \hat{z}'_{i,Y}(t_1) \\ \hat{z}'_{j,X}(t_1) & \hat{z}'_{j,Y}(t_1) & (t_1 - t_0)\Delta t \hat{z}'_{j,X}(t_1) & (t_1 - t_0)\Delta t \hat{z}'_{j,Y}(t_1) \\ \dots & \dots & \dots & \dots \\ \hat{z}'_{m,X}(t_1) & \hat{z}'_{m,Y}(t_1) & (t_1 - t_0)\Delta t \hat{z}'_{m,X}(t_1) & (t_1 - t_0)\Delta t \hat{z}'_{m,Y}(t_1) \end{bmatrix} \quad (9.16)$$

It can be recognized that the state $v_0(t_0)$ can no longer be reconstructed from data from a single time step (when $t_1 = t_0$) since in that case, there is no connection to the acceleration variables. This is what has been expected intuitively. When including the next time step, i.e. $t_1 = t_0 + 1$, the system could become observable, since then, all rows might become linearly independent (of course, depending on the entries, i.e. the linear independence of the partial derivatives).

List of Figures

1.1	DWD rain gauges starting operation since 1990.	12
1.2	Example of radar image and motion vector time series.	13
2.1	DWD standard Z-R relationship	21
2.2	Weather radar principle.	21
2.3	Visual examples of wireless sensor networks and communication graph.	24
2.4	Possible violations of linear regression assumptions.	29
2.5	One-dimensional 'frozen field' example.	36
2.6	Empirical and theoretical variogram example	39
2.7	FF-IDW Example.	41
2.8	Optical flow example	44
4.1	Flow chart for the estimation of precipitation at 1-min resolution	55
4.2	Time-window approach for the estimation of precipitation intensity at 1-min resolution.	56
4.3	Creation of 1-min radar snapshots out of a pair of adjacent 5-min snapshots	57
4.4	Examples of radar image and spatial and temporal theoretical variograms for two events.	59
4.5	Process scheme for the RADAR-DBC method.	62
4.6	Motivating example for RADAR-DBC method.	63
4.7	Distances involved in IDW weight computation for the method RADAR DBC. . .	64
4.8	Radar windows for the radar-based estimation methods.	65
5.1	Image of sensor and data logger and extract of raw car data provided by the sensors.	68
5.2	Image of Xanonex sensor and data logger and extract of raw ASCII data provided by the data logger.	69
5.3	Aggregation of Xanonex samples.	70
6.1	Flow chart for the motion estimation algorithm	77
6.2	Estimation of partial derivatives from irregular data.	79
6.3	Assumption of node stationarity and sampling synchronicity.	80
7.1	Study area for the evaluation of the 1-min precipitation estimation methods. . . .	93
7.2	Process of cross-validation	94
7.3	Histogram of rain gauge measurements for January 2012 and June 2012.	96
7.4	Example of a snapshot of current radar image and surfaces interpolated by IDW, OK and ST-OK.	97
7.5	Example of a snapshot of the non-symmetric spatio-temporal interpolation and radar-rain gauge merging methods.	97
7.6	RADAR-DBC example and visualization of computed spatio-temporal lags. . . .	98
7.7	RADAR-DBC regression surface.	98
7.8	Correlation of estimates with rain gauges measurements for January 2012 and June 2012.	99
7.9	Bias of estimates of different methods for January 2012 and June 2012.	100
7.10	RSE of estimates for January 2012 and June 2012.	100

7.11	Correlation of estimates with rain gauges measurements for January 2012 and June 2012 for rain gauge samples larger zero.	101
7.12	Bias of estimates for larger-zero rain gauge samples for January 2012 and June 2012.	102
7.13	RSE of estimates for January 2012 and June 2012 for rain gauge samples larger zero.	103
7.14	Correlation of estimates with rain gauges measurements for January 2012 for two classes of field speed.	104
7.15	Map of correlations of estimates with rain gauges measurements for January 2012 for all rain gauges individually.	105
7.16	Heat maps of collected car data.	107
7.17	Example from the integrated visualization of car data, weather radar, rain gauges and motion vectors.	108
7.18	Car speed vs. 1-min wiper frequency of two VW Touran.	109
7.19	Wiper classes of two VW Touran vs. the references FF-IDW, RADAR-DBC and RADAR 121.	109
7.20	Car speed vs. 1-min wiper frequency of a Ford S-Max and dynamic-static ratio. . .	110
7.21	Correlation of speed-corrected wiper frequencies of a single Ford S-Max with automatically controlled wipers with rainfall estimated with the different methods. . .	111
7.22	Histogram of speed-corrected readings (wiper frequencies) of the Ford S-Max. . .	111
7.23	Speed-corrected 1-min wiper frequencies vs. two different references FF-IDW and RADAR 121.	112
7.24	Speed-corrected 1-min wiper frequencies of a Ford S-Max with automatically controlled wipers vs. RADAR DBC.	112
7.25	Car speed vs. 1-min sensor readings of Xanonex sensor installed in Skoda Fabia and dynamic-static ratio.	113
7.26	Correlation of speed-corrected Xanonex readings installed in Skoda Fabia with rainfall estimated with the different methods.	114
7.27	Histogram of speed-corrected Xanonex readings of the Skoda Fabia.	114
7.28	Speed-corrected 1-min Xanonex readings vs. references FF-IDW and RADAR 121.	115
7.29	Speed-corrected 1-min Xanonex readings vs. RADAR DBC.	115
7.30	Correlation of sensor readings with rainfall estimated from rain gauges for sensor samples with a spatial distance closer 2km to next rain gauge.	116
7.31	Speed-corrected 1-min Xanonex readings vs. FF-IDW and NN for samples closer 2km to nearest neighbor.	116
7.32	Analysis of the Xanonex data collected at VW rain track.	117
7.33	Study area for the decentralized motion estimation algorithm.	119
7.34	Deployment strategies for the analysis of the motion estimation algorithm.	120
7.35	The two field types underlying the simulation and a visualization of a GSN and the assumed true and estimated motion vectors.	123
7.36	Comparison of mean angular error and mean speed bias for different motion speeds, deployment densities and two field types.	124
7.37	Comparison of mean angular error and mean speed bias for different geosensor network (GSN) deployments.	125
7.38	Comparison of mean angular error and mean speed bias for different Kalman filter measurement noises.	126
7.39	Characterization of the radar event for evaluation of the motion estimation algorithm.	127
7.40	Time series of radar images and motion vectors estimated by the proposed algorithm.	128

List of Tables

2.1	Classification of liquid precipitation provided by the DWD.	18
7.1	Characterization of the evaluation periods.	95

Bibliography

- L. Alfieri, et al. (2010). ‘Time-dependent ZR relationships for estimating rainfall fields from radar measurements’. *Natural Hazards and Earth System Sciences* **10**(1):149–158.
- A. Amani & T. Lebel (1997). ‘Lagrangian kriging for the estimation of Sahelian rainfall at small time steps’. *Journal of Hydrology* **192**(1):125–157.
- J. L. Barron, et al. (1994). ‘Performance of optical flow techniques’. *International journal of computer vision* **12**(1):43–77.
- H. Bartels, et al. (2004). ‘Projekt RADOLAN - Routineverfahren zur Online-Aneichung der Radarniederschlagsdaten mit Hilfe von automatischen Bodenniederschlagsstationen (Ombrometer) - Zusammenfassender Abschlussbericht für die Projektlaufzeit von 1997 bis 2004’. Final Report, Deutscher Wetterdienst (DWD).
- M. Berenguer, et al. (2005). ‘Hydrological Validation of a Radar-Based Nowcasting Technique’. *Journal of Hydrometeorology* **6**(4):532–549.
- C. Berndt, et al. (2014). ‘Geostatistical merging of rain gauge and radar data for high temporal resolutions and various station density scenarios’. *Journal of Hydrology* **508**:88–101.
- F. Bittner (2012). ‘Untersuchung externer Faktoren bei der Regenmessung mittels Scheibenwischern’. Master thesis, Leibniz Universität Hannover, Hannover.
- M. J. Black (1992). *Robust incremental optical flow*. Ph.D. thesis, PhD thesis, Yale university, New Haven, CT, USA.
- F. Bocci (2012). ‘Whether or not to run in the rain’. *European Journal of Physics* **33**(5):1321–1332.
- A. Boukerche (2008). *Algorithms and protocols for wireless sensor networks*, vol. 62. John Wiley & Sons.
- N. E. Bowler, et al. (2004). ‘Development of a precipitation nowcasting algorithm based upon optical flow techniques’. *Journal of Hydrology* **288**(1–2):74–91.
- J. Brink & E. Pebesma (2014). ‘Plume Tracking with a Mobile Sensor Based on Incomplete and Imprecise Information’. *Transactions in GIS* **18**(5):740–766.
- L. D. Cesare, et al. (2001). ‘Estimating and modeling space-time correlation structures’. *Statistics & Probability Letters* **51**(1):9–14.
- N. Cressie (1993). *Statistics for spatial data*. Wiley.
- N. Cressie & H.-C. Huang (1999). ‘Classes of Nonseparable, Spatio-Temporal Stationary Covariance Functions’. *Journal of the American Statistical Association* **94**(448):1330–1339.
- N. Cressie & C. K. Wikle (2011). *Statistics for Spatio-Temporal Data*. John Wiley & Sons.
- F. C. Curriero (2006). ‘On the Use of Non-Euclidean Distance Measures in Geostatistics’. *Mathematical Geology* **38**(8):907–926.
- J. Das, et al. (2012). ‘Coordinated sampling of dynamic oceanographic features with underwater vehicles and drifters’. *The International Journal of Robotics Research* **31**(5):626–646.
- R. J. Doviak & D. S. Zrnic (1993). *Doppler Radar and Weather Observations*. Dover Publications, Mineola, N.Y.

- M. Duckham (2012). *Decentralized Spatial Computing: Foundations of Geosensor Networks*. Springer, Berlin Heidelberg, Germany, 2013 edn.
- M. Duckham, et al. (2015). ‘Challenges to using decentralized spatial algorithms in the field: The RISERnet geosensor network case study’. *SIGSPATIAL Special* **7**(2):14–21.
- U. Ehret (2003). *Rainfall and flood nowcasting in small catchments using weather radar*. Ph.D. thesis, Inst. für Wasserbau an der Univ. Stuttgart, Stuttgart.
- G. Farnebäck (2003). ‘Two-Frame Motion Estimation Based on Polynomial Expansion’. In J. Bigun & T. Gustavsson (eds.), *Image Analysis*, no. 2749 in Lecture Notes in Computer Science, pp. 363–370. Springer Berlin Heidelberg.
- D. Fitzner, et al. (2012). ‘Online Calibration of Measurement Devices in Geosensor Networks - An Example Application’. In *Extended Abstracts - Proceedings of the 7th International Conference on Geographic Information Science*, Columbus, Ohio.
- D. Fitzner, et al. (2013). ‘Rainfall Estimation with a Geosensor Network of Cars Theoretical Considerations and First Results’. *Photogrammetrie - Fernerkundung - Geoinformation* **2013**(2):93–103.
- S. Gaitan, et al. (2014). ‘From sensing to action: Quick and reliable access to information in cities vulnerable to heavy rain’. *IEEE Sensors Journal* **14**(12):4175–4184.
- C. Gaucherel & V. Grimaldi (2015). ‘The Pluviophone: Measuring Rainfall by Its Sound’. *Journal of Vibration and Acoustics* **137**(3):034504.
- U. Germann & I. Zawadzki (2002). ‘Scale-Dependence of the Predictability of Precipitation from Continental Radar Images. Part I: Description of the Methodology’. *Monthly Weather Review* **130**(12):2859–2873.
- T. Gneiting (2002). ‘Nonseparable, Stationary Covariance Functions for Space–Time Data’. *Journal of the American Statistical Association* **97**(458):590–600.
- T. Gneiting, et al. (2006). ‘Geostatistical space-time models, stationarity, separability, and full symmetry’. *Monographs On Statistics and Applied Probability* **107**:151.
- T. Gneiting, et al. (2001). ‘Analogies and correspondences between variograms and covariance functions’. *Advances in Applied Probability* **33**(3):617–630.
- G. H. Golub & C. F. V. Loan (1996). *Matrix Computations*. JHU Press, Baltimore London, USA/UK.
- M. F. Goodchild (2010). ‘Twenty years of progress: GIScience in 2010’. *Journal of spatial information science* **2010**(1):3–20.
- M. Gopal (1993). *Modern control system theory*. New Age International.
- B. Gräler, et al. (2012). ‘Spatio-temporal analysis and interpolation of PM10 measurements in Europe for 2009’. *ETC/ACM Technical Paper* **8**:1–29.
- V. K. Gupta & E. Waymire (1987). ‘On Taylor’s hypothesis and dissipation in rainfall’. *Journal of Geophysical Research: Atmospheres* **92**(D8):9657–9660.
- C. Guy (2006). ‘Wireless sensor networks’. In *Sixth International Symposium on Instrumentation and Control Technology: Signal Analysis, Measurement Theory, Photo-Electronic technology, and Artificial Intelligence*, pp. 63571I–63571I. International Society for Optics and Photonics.
- U. Haberlandt (2007). ‘Geostatistical interpolation of hourly precipitation from rain gauges and radar for a large-scale extreme rainfall event’. *Journal of Hydrology* **332**(1–2):144–157.
- U. Haberlandt (2011). ‘Interpolation of precipitation for flood modelling’. In *Flood Risk Assessment and Management*, pp. 35–52. Springer.

- U. Haberlandt & M. Sester (2010). ‘Areal rainfall estimation using moving cars as rain gauges – a modelling study’. *Hydrol. Earth Syst. Sci.* **14**(7):1139–1151.
- C. Heipke (2010). ‘Crowdsourcing geospatial data’. *ISPRS Journal of Photogrammetry and Remote Sensing* **65**(6):550–557.
- G. Heuvelink & E. Pebesma (2002). ‘Is the ordinary kriging variance a proper measure of interpolation error?’. In *Proceedings of the fifth international symposium on spatial accuracy assessment in natural resources and environmental sciences*, pp. 179–186.
- B. K. Horn & B. G. Schunck (1981). ‘Determining optical flow’. *Artificial Intelligence* **17**(1–3):185–203.
- R. A. Houze (1981). ‘Structures of atmospheric precipitation systems: A global survey’. *Radio Science* **16**(5):671–689.
- R. Hut, et al. (2014). ‘Using umbrellas as mobile rain gauges: prototype demonstration’. In *EGU General Assembly Conference Abstracts*, vol. 16, p. 16418.
- M.-H. Jeong, et al. (2014). ‘Decentralized and coordinate-free computation of critical points and surface networks in a discretized scalar field’. *International Journal of Geographical Information Science* **28**(1):1–21.
- R. E. Kalman (1960). ‘A New Approach to Linear Filtering and Prediction Problems’. *Transactions of the ASME – Journal of Basic Engineering* **82**(1):35–45.
- R. Kuner (2015). ‘DWD Vorschriften und Betriebsunterlagen Nr. 3 (VuB 3) - Beobachterhandbuch’. Tech. rep., Deutscher Wetterdienst.
- R. B. Langley (1999). ‘Dilution of precision’. *GPS world* **10**(5):52–59.
- D. R. Legates (2000). ‘Real-time calibration of radar precipitation estimates’. *The Professional Geographer* **52**(2):235–246.
- F. L. Lewis et al. (2004). ‘Wireless sensor networks’. *Smart environments: technologies, protocols, and applications* pp. 11–46.
- B. Li, et al. (2009). ‘Statistical tests of Taylor’s hypothesis: an application to precipitation fields’. *Journal of Hydrometeorology* **10**(1):254–265.
- A. Libertino, et al. (2015). ‘Radar estimation of intense rainfall rates through adaptive calibration of the ZR relation’. *Atmosphere* **6**(10):1559–1577.
- B. D. Lucas, et al. (1981). ‘An iterative image registration technique with an application to stereo vision.’. In *7th International Joint Conference on Artificial Intelligence (IJCAI)*, vol. 81, pp. 674–679, Vancouver, Canada.
- N. A. Lynch (1996). *Distributed algorithms*. Morgan Kaufmann.
- J. S. Marshall & W. M. K. Palmer (1948). ‘The distribution of raindrops with size’. *Journal of Meteorology* **5**(4):165–166.
- G. Matheron (1971). *The theory of regionalized variables and its applications*, vol. 5. École nationale supérieure des mines.
- P. S. Maybeck (1979). *Stochastic models, estimation, and control - Volume 1*. Mathematics in Science and Engineering. Academic Press, New York.
- C. McGlone (2013). *Manual of Photogrammetry, Sixth Edition*. American Society for Photogrammetry and Remote Sensing (ASPRS).
- H. Messer, et al. (2006). ‘Environmental Monitoring by Wireless Communication Networks’. *Science* **312**(5774):713–713.

- J. J. Moré (1978). 'The Levenberg-Marquardt algorithm: implementation and theory'. In *Numerical analysis*, pp. 105–116. Springer.
- J. E. Nielsen, et al. (2014). 'A numerical method to generate high temporal resolution precipitation time series by combining weather radar measurements with a nowcast model'. *Atmospheric Research* **138**:1–12.
- W. Niemeier (2002). *Ausgleichungsrechnung: eine Einführung für Studierende und Praktiker des Vermessungs- und Geoinformationswesens*. De-Gruyter-Lehrbuch. de Gruyter, Berlin New York.
- A. Overeem, et al. (2013). 'Country-wide rainfall maps from cellular communication networks'. *Proceedings of the National Academy of Sciences* **110**(8):2741–2745.
- T. Pfaff (2013). *Processing and analysis of weather radar data for use in hydrology*. Ph.D. thesis, Universität Stuttgart.
- E. Rabiei, et al. (2013). 'Rainfall estimation using moving cars as rain gauges – laboratory experiments'. *Hydrol. Earth Syst. Sci.* **17**(11):4701–4712.
- E. Rabiei, et al. (2016). 'Areal rainfall estimation using moving cars – computer experiments including hydrological modeling'. *Hydrology and Earth System Sciences* **20**(9):3907–3922.
- J. Riedl (1986). 'Radar-Flächenniederschlagsmessung'. *Promet* **2**(3):20–23.
- S. Särkkä (2013). *Bayesian Filtering and Smoothing*. Cambridge University Press, Cambridge, UK.
- S. Schmid & R. Wattenhofer (2008). 'Modeling sensor networks'. *Algorithms and Protocols for Wireless Sensor Networks* **62**:77.
- F. Schulz (2007). 'Modeling sensor and ad hoc networks'. In *Algorithms for Sensor and Ad Hoc Networks*, pp. 21–36. Springer.
- M.-J. Schulze, et al. (2010). 'Cooperative information augmentation in a geosensor network'. In W. Shi, M. Goodchild, B. Lees, & Y. Leung (eds.), *Advances in Geo-Spatial Information Science*, pp. 444–449. CRC Press.
- J. Seltmann (1997). 'Radarforschung im DWD: Vom Scan zum Produkt'. *Promet* **26**:32–52.
- J. E. Seltmann, et al. (2013). 'DWD's new operational scan strategy'. In *36rd AMS Conf. on Radar Meteorology, Breckenridge, CO*.
- M. Sester (2009). 'Cooperative Boundary Detection in a Geosensor Network usng a SOM'. In *Proceedings of the International Csrtographic Conference, Santiago, Chile*.
- E. P. Simoncelli (1993). *Distributed representation and analysis of visual motion*. Ph.D. thesis, Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science.
- E. P. Simoncelli, et al. (1991). 'Probability distributions of optical flow'. In *Computer Vision and Pattern Recognition, 1991. Proceedings CVPR'91., IEEE Computer Society Conference on*, pp. 310–315, Loas Alamos, CA, USA. IEEE.
- K. Sohrawy, et al. (2007). *Wireless sensor networks: technology, protocols, and applications*. John Wiley & Sons.
- B. Southall, et al. (1998). 'Controllability and observability: Tools for Kalman filter design'. In *BMVC 98: Proceedings of the Ninth British Machine Vision Conference*, pp. 164–173. British Machine Vision Association.
- A. O. Sykes (1993). 'An Introduction to Regression Analysis'. Coase-Sandor Institute for Law & Economics Working Paper 20, Coase-Sandor Institute for Law & Economics.
- G. I. Taylor (1938). 'The Spectrum of Turbulence'. *Proceedings of the Royal Society of London. Series A - Mathematical and Physical Sciences* **164**(919):476–490.

- H.-W. Tsai, et al. (2007). ‘Mobile object tracking in wireless sensor networks’. *Computer Communications* **30**(8):1811–1825.
- M. Umer, et al. (2010). ‘Spatial interpolation in wireless sensor networks: localized algorithms for variogram modeling and Kriging’. *GeoInformatica* **14**(1):101–134.
- A. Verworn & U. Haberlandt (2011). ‘Spatial interpolation of hourly rainfall – effect of additional information, variogram inference and storm properties’. *Hydrol. Earth Syst. Sci.* **15**(2):569–584.
- P. K. Watson (1983). ‘Kalman filtering as an alternative to Ordinary Least Squares — Some theoretical considerations and empirical results’. *Empirical Economics* **8**(2):71–85.
- E. Weigl (2015). ‘Radarniederschlag - Prinzip der Niederschlagsbestimmung mit Radar inkl. Umrechnung der Radarreflektivitäten in Momentanwerte des Niederschlags - Version 1.1’. Tech. rep., Deutscher Wetterdienst.
- G. Welch & G. Bishop (2006). *An introduction to the kalman filter.*, vol. Department of Computer Science, University of North Carolina. Chapel Hill, NC, unpublished manuscript.
- S. L. Willis & C. J. Kikkert (2005). ‘Radio propagation model for long-range AD HOC wireless sensor network’. In *Wireless Networks, Communications and Mobile Computing, 2005 International Conference on*, vol. 1, pp. 826–832. IEEE.
- J. W. Wilson & E. A. Brandes (1979). ‘Radar Measurement of Rainfall—A Summary’. *Bulletin of the American Meteorological Society* **60**(9):1048–1058.
- T. Winterrath, et al. (2012). ‘On the DWD quantitative precipitation analysis and nowcasting system for real-time application in German flood risk management’. *IAHS-AISH publication* pp. 323–329.
- World Meteorological Organization (2008). ‘Guide to Hydrological Practices - Volume I - Hydrology - From Measurement to Hydrological Information’. Tech. Rep. WMO-No. 168, World Meteorological Organization.
- W. Yen, et al. (2005). ‘DWD Radar Products for Model Evaluation in SPP 1167’. In *Poster at the Kolloquium SPP 1167: Quantitative Precipitation Forecast*, Bad Honnef.
- M. Younis & K. Akkaya (2008). ‘Node Positioning for Increased Dependability of Wireless Sensor’. *Algorithms and Protocols for Wireless Sensor Networks* **62**:225.
- S. Yousefi, et al. (2006). ‘Vehicular ad hoc networks (VANETs): challenges and perspectives’. In *ITS Telecommunications Proceedings, 2006 6th International Conference on*, pp. 761–766. IEEE.
- I. I. Zawadzki (1973). ‘Statistical Properties of Precipitation Patterns’. *Journal of Applied Meteorology* **12**(3):459–472.

Danksagung

Die vorliegende Dissertation ist während meiner Zeit am Institut für Kartographie und Geoinformatik der Leibniz Universität Hannover entstanden. Viele Personen haben zu dieser Arbeit in vielfältiger Form beigetragen.

An erster Stelle möchte ich mich herzlich bei Frau Prof. Dr.-Ing. habil. Monika Sester bedanken. Sie hatte, zusammen mit Herrn Prof. Dr. Uwe Haberlandt, die grundlegende Idee, Autos als mobile Regenmessstationen zu verwenden. Beide haben dann die Möglichkeit geschaffen, die Idee im DFG Forschungsprojekt RainCars näher zu untersuchen und somit auch die Grundlage gelegt für das Promotionsprojekt, dessen Ergebnisse in diesem Dokument vorliegen. Während des Arbeitsprozesses hat Frau Sester kontinuierlich wertvolle Unterstützung geleistet, für die ich ihr außerordentlich dankbar bin. Zudem haben ihre Hinweise und Korrekturvorschläge die vorliegende Arbeit äußerst positiv beeinflusst.

Herrn Prof. Dr.-Ing. Uwe Haberlandt danke ich ebenfalls für die Möglichkeit zur Mitarbeit im RainCars Projekt und somit die Möglichkeit der Promotion. Zudem möchte ich ihm für die vielen Anregungen danken sowie für die Qualitätskontrolle der im Zuge der Projektarbeit entstandenen Ergebnisse, von denen sich nun auch Teile in dieser Arbeit wiederfinden. Außerdem danke ich ihm herzlich für die Unterstützung bei den Feldexperimenten.

Mein Dank gilt Herrn Prof. Dr. Edzer Pebesma und Herrn Prof. Dr.-Ing. Steffen Schön für die Übernahme des Korreferats dieser Arbeit sowie Herrn Prof. Dr.-Ing Ingo Neumann für die Übernahme des Vorsitzes der Prüfungskommission.

Besonders danken möchte ich auch Bastian Heinrich für seine Unterstützung, insbesondere für die Entwicklung der Autosensoren. Des Weiteren möchte ich Herrn Dr. Ehsan Rabiei und Herrn Dr. Christian Berndt, Frau Bora Shehu, Herrn Prof. Dr. Claus Brenner, Herrn Dr. Stefan Krämer sowie Herrn Prof. Dr. Edzer Pebesma für Ideen und Diskussionen danken, die an verschiedenen Stellen in diese Arbeit eingeflossen sind und sie positiv beeinflusst haben.

Mein Dank gilt auch allen Kolleginnen und Kollegen am Institut für Kartographie und Geoinformatik für die fruchtbare Zusammenarbeit der letzten Jahre und das angenehme Arbeitsumfeld. Viele Kollegen haben ihr Auto als Testfahrzeug zur Verfügung gestellt, ihnen gilt mein Dank für die Unterstützung bei den Feldexperimenten.

Meiner Familie danke ich für die positive Unterstützung und Begleitung während der gesamten Promotionsphase.

Wissenschaftliche Arbeiten der Fachrichtung Geodäsie und Geoinformatik der Leibniz Universität Hannover

(Eine vollständige Liste der Wiss. Arb. ist beim Geodätischen Institut, Nienburger Str. 1, 30167 Hannover erhältlich.)

Nr. 311	SCHUNERT, Alexander:	Assignment of Persistent Scatterers to Buildings (Diss. 2014)
Nr. 312	GUERCKE, Richard:	Optimization Aspects in the Generalization of 3D Building Models (Diss. 2014)
Nr. 313	ZIEMS, Marcel:	Automatic verification of road databases using multiple road models (Diss. 2014)
Nr. 314	DINI, Gholam Reza:	Toward an automatic solution for updating building databases using high resolution space-borne stereo imaging (Diss. 2014)
Nr. 315	KERSTEN, Tobias:	Bestimmung von Codephasen-Variationen bei GNSS-Empfangsantennen und deren Einfluss auf Positionierung, Navigation und Zeitübertragung (Diss. 2014)
Nr. 316	BISKUPEK, Liliane:	Bestimmung der Erdorientierung mit Lunar Laser Ranging (Diss. 2015)
Nr. 317	STEINER, Christina:	Highspeed Stereo-Endoskopie für eng begrenzte Messvolumina (Diss. 2015)
Nr. 318	BANDIKOVA, Tamara:	The role of attitude determination for inter-satellite ranging (Diss. 2015)
Nr. 319	LIN, Miao:	Regional gravity field recovery using the point mass method (Diss. 2015)
Nr. 320	ZHANG, Lijuan:	Mining GPS-Trajectory Data for Map Refinement and Behavior Detection (Diss. 2015)
Nr. 321	ZADDACH, Sebastian:	Zum Beitrag Bayesscher Schätzverfahren in der Vergleichswertermittlung (Diss. 2016)
Nr. 322	SMYRNAIOS, Marios:	Carrier-phase Multipath in Satellite-based Positioning (Diss. 2016)
Nr. 323	MENZE, Moritz:	Object Scene Flow (Diss. 2016)
Nr. 324	WU, Hu:	Gravity field recovery from GOCE observations (Diss. 2016)
Nr. 325	XU, Xiangyang:	Terrestrial Laser Scanning for the Generation and Calibration of Finite Element Models (Diss. 2016)
Nr. 326	SAYYAD, Muhammad Naeem Shahzad:	Joint use and mutual control of terrestrial laser scans and digital images for accurate 3D measurements (Diss. 2016)
Nr. 327	SCHACK, Lukas:	Object-based matching of Persistent Scatterers to Optical Oblique Images (Diss. 2016)
Nr. 328	REICH, Martin:	Global Image Orientation from Pairwise Relative Orientations (Diss. 2016)
Nr. 329	KLINGER, Tobias:	Probabilistic multi-person localisation and tracking (Diss. 2016)
Nr. 330	SCHMIDT, Alena:	Markierte Punktprozesse für die automatische Extraktion von Liniennetzen in Rasterdaten (Diss. 2016)
Nr. 331	HOFMANN, Franz:	Lunar Laser Ranging - verbesserte Modellierung der Monddynamik und Schätzung relativistischer Parameter (Diss. 2017)
Nr. 332	BRIEDEN, Phillip:	Validierung von GOCE-Gravitationsgradienten in Kreuzungspunkten und Zukunftsperspektiven der Satellitengradiometrie (Diss. 2017)
Nr. 333	VON GÖSSELN, Ilka:	Simulationsbasierte Effizienzoptimierung von Messprozessen am Beispiel der tachymetrischen Netzmessung (Diss. 2017)
Nr. 334	HOFMANN, Sabine:	Potential von LiDAR Mobile Mapping für hochgenaue Karten (Diss. 2017)
Nr. 335	ALBERT, Lena:	Simultane Klassifikation der Bodenbedeckung und Landnutzung unter Verwendung von Conditional Random Fields (Diss. 2017)
Nr. 336	NIEMEYER, Joachim:	Verwendung von Kontext zur Klassifikation luftgestützter Laserdaten urbaner Gebiete (Diss. 2017)
Nr. 337	EGGERT, Daniel:	Effiziente Verarbeitung und Visualisierung von Mobile Mapping Daten (Diss. 2017)
Nr. 338	FITZNER, Daniel:	Estimation of Spatio-Temporal Moving Fields at High Resolution (Diss. 2017)

Die Arbeiten werden im Rahmen des wissenschaftlichen Schriftenaustausches verteilt und sind nicht im Buchhandel erhältlich. Der Erwerb ist zu einem Stückpreis von € 25,00 bei den herausgebenden Instituten möglich.