Analysis of the Spatial Extent, Intensity, and Duration of Rainfall Events in the Netherlands

A Statistical Approach Using Radar Data (1998-2023)

A.M. Primavera

Delft University of Technology

Analysis of the Spatial Extent, Intensity, and Duration of Rainfall Events in the Netherlands

A Statistical Approach Using Radar Data (1998-2023)

by

A.M. Primavera

to obtain the degree of Master of Science in Applied Mathematics
at the Delft University of Technology,
to be defended publicly on Thursday April 17 at 10:00 AM.

Student number:	4942035		
Project duration:	September, 2024 – April, 2025		
Thesis committee:	Dr. F. Mies,	TU Delft,	daily supervisor
	Dr. C. Kraaikamp,	TU Delft,	responsible supervisor
	Ir. D. Lugt,	HKV,	external supervisor
	Dr. R.P. Nicolai,	HKV,	external supervisor
	Prof. Dr. Ir. G. Jongbloed,	TU Delft	

Cover: Visualisation of large-scale rainfall events over the Netherlands on July 31, 2023, at 05:25, based on radar data.

Preface

This thesis represents the final step in completing my Master's in Applied Mathematics at Delft University of Technology and marks the conclusion of my journey as a student. Conducted in collaboration with HKV, a Dutch consultancy specialising in water and flood risk management, this research focuses on analysing the spatial extent, intensity, and duration of rainfall events in the Netherlands.

While a basic understanding of mathematics is assumed, no prior knowledge of meteorology or climatology is required to read this document. I hope that this thesis provides useful insights and contributes to ongoing discussions in this field.

I am very grateful to my supervisors for their support and guidance throughout this process. At EEMCS, Department of Statistics at TU Delft, Fabian Mies provided thoughtful feedback and helped steer my research in the right direction. At HKV, Dorien Lugt and Robin Nicolai shared their knowledge, expertise, and experience, offering insights that improved my work. I would also like to thank everyone at HKV for making my time at the company both enjoyable and rewarding. Their welcoming atmosphere and collaborative spirit made a positive impact on my experience.

Additionally, I sincerely appreciate the discussions I had with Aart Overeem, Hylke de Vries, and Geert Lenderink from KNMI. Their insights from a meteorological perspective and their assistance in constructing the rainfall event dataset were extremely helpful to this research.

Finally, I would like to thank Cornelis Kraaikamp and Geurt Jongbloed for being part of my thesis committee and for taking the time to review my research.

A.M. Primavera
Delft, April 2025

Summary

This thesis investigates rainfall events in the Netherlands over a 26-year period (1998–2023) using radar data to understand how rainfall event dynamics have evolved. Extreme precipitation is a main contributor to floods, posing risks to human life, infrastructure, and ecosystems. This study takes a life-cycle-based approach, distinct from point-based rainfall measurements, utilising radar data to track rainfall events in both time and space. The goal of using this approach is to gain a better understanding of the spatial extent of rainfall events.

The research primarily focuses on the relationships between the spatial extent, duration, and intensity of rainfall events and explores how these relationships as well as the individual variables may have changed over time. The study utilises KNMI's NL21 radar-derived precipitation data, which consists of 5-minute precipitation accumulations with a spatial resolution of 2.4 km. Using the Celltrack event-tracking algorithm developed by KNMI, rainfall events are identified and the corresponding characteristics describing their spatial extent (i.e. area of the event), duration and intensity are calculated. This results in a dataset of approximately 112,000 events.

Statistical trend detection techniques, specifically the Mann-Kendall test and the Theil-Sen slope estimator, are employed to identify significant trends in rainfall characteristics over the study period. Additionally, the relationships between the rainfall variables are explored using correlation analysis and a binned analysis, providing insights into the dependencies between area, duration, intensity, and dew point temperature.

Over the study period, there has been a significant increase in the spatial extent (as can be seen from Figure 1) and duration of rainfall events, a trend that persists across most seasons.

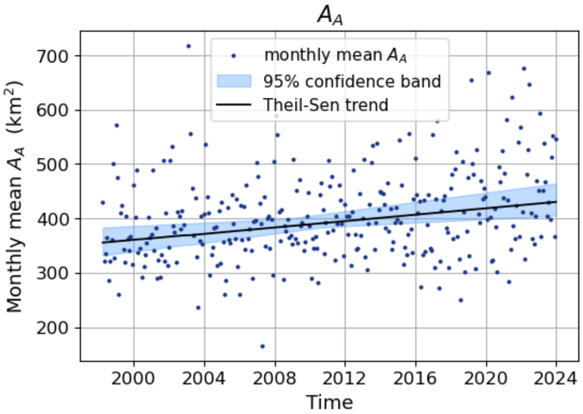

Figure 1: Monthly mean values of the average area (A_A) of rainfall events over the study period. The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.

The changes in rainfall intensities are less consistent. Shorter-duration events (less than 105 minutes) exhibit decreased intensity, which contradicts previous research suggesting that extreme rainfall events are becoming more intense. One possible explanation for this discrepancy is that the thresholds used in the event-tracking methodology may filter out convective events, potentially highlighting a distinction between average and extreme rainfall events. While overall rainfall intensity may have decreased on average, the most extreme events could have become more intense, as suggested by other studies on extreme precipitation trends. Indeed, when we study smaller-scale shorter-duration higher-intensity events separately, an increasing trend in intensity is observed, reinforcing findings from existing literature. Additionally, the analysis of extreme rainfall behaviour reveals no significant changes in the

extreme value index for any of the variables, suggesting that the shapes of the distributions of extreme rainfall events have remained stable over the study period.

Regarding the relationships between rainfall variables, the study finds a strong positive correlation between rainfall duration and area, while intensity is moderately correlated with both duration and dew point temperature, as illustrated in the correlation matrix in Figure 2. The latter supports the theoretical understanding of moisture dynamics and aligns with the Clausius-Clapeyron equation, which explains how rising dew point temperatures enhance precipitation intensity. A statistical test confirms that these correlations are not constant over time. Specifically, towards the end of the study period, longer events more frequently cover larger areas, and higher-intensity events are increasingly characterised by longer durations and larger spatial extent.

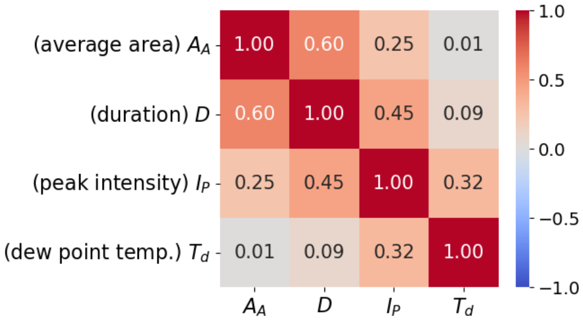

Figure 2: Heatmap showing the correlations between key rainfall event variables. The plot highlights a strong positive correlation between rainfall duration and area, as well as moderate correlations between intensity and both duration and dew point temperature.

In conclusion, this study provides a comprehensive analysis of rainfall event characteristics and trends for rainfall events in the Netherlands over the past 26 years. While notable trends in spatial extent and duration have been identified, intensity changes exhibit more complex patterns. This work lays a solid foundation for future research, which could focus on refining intensity metrics, analysing convective and stratiform events separately, applying multivariate methods to better capture rainfall interactions, extending the analysis to other regions, and refining event-tracking methodologies. These approaches would further enhance our understanding of potential flood risks.

Contents

1

Introduction

Understanding extreme precipitation is important for assessing flood risk and mitigating its effects. This introduction first outlines the motivation for studying rainfall events using an event tracking approach. It then presents the research objectives and questions before discussing methods for identifying and tracking rainfall events from radar data. Finally, the structure of the remainder of this thesis is presented.

1.1. Motivation

Extreme precipitation is one of the primary drivers of floods, posing significant threats to human life, infrastructure, agriculture and ecosystems. When intense rainfall cannot be absorbed by the land or transported by rivers, the excess water can quickly lead to flash floods, river floods, and urban flooding. Such events often cause severe socio-economic damage, including infrastructure damage, disruption of public services and health impacts [1, 15]. In recent years, climate change has increased the frequency and intensity of extreme rainfall events, making floods a growing concern for many regions worldwide [17, 18].

In the Netherlands, the threat of flooding is particularly pronounced due to its low-lying geography and dense network of rivers and canals [43]. As a country where a significant portion of land lies below sea level, the Netherlands has a long history of flood management and water control. However, the increasing frequency and intensity of extreme weather events challenge the capacity of traditional flood defences to mitigate these risks. As such, there is a growing need to better understand rainfall dynamics, trends, and their implications for flood management in the Netherlands.

Effective flood risk management requires not only analysing the occurrence of extreme rainfall but also understanding the relationships between its defining characteristics – namely, spatial extent (or area), duration, and intensity [39]. These attributes often interact in ways that amplify flood risks, yet they are rarely examined together. This study addresses this gap by statistically analysing rainfall events in the Netherlands from 1998 to 2023 using radar data. By examining the relationships between rainfall characteristics, this study provides insights into how precipitation patterns might be evolving.

Traditional approaches to analysing rainfall trends and extremes have relied on rain gauge measurements and statistical models that focus on location-based statistics, total precipitation, or long-term averages (e.g. [31, 8, 46]). While these methods provide valuable insights, they often lack the spatial and temporal resolution necessary to capture the complexity of individual rainfall events. Advances in radar and satellite technology have introduced new possibilities, enabling life-cycle-based analyses that track rainfall events in both time and space. Such analyses can provide data on rainfall areas, movement, and intensity at a level of detail previously unattainable [37, 38, 2]. This shift from point-based to life-cycle-based approaches represents a novel perspective in rainfall analysis, offering the potential for a deeper understanding of how rainfall event dynamics contribute to extreme precipitation and flooding.

The Netherlands, characterised by its temperate maritime climate, experiences frequent rainfall distributed relatively evenly throughout the year, with some seasonal variations [7]. However, studies suggest that a warming climate is likely to increase the occurrence of heavy rainfall and more extreme precipitation events [40, 14, 19]. Although prior research on Dutch precipitation trends has primarily focused on individual aspects of rainfall, such as total precipitation or peak intensity (e.g., [11, 29]), these studies often do not explore the interactions between rainfall characteristics. Notable studies, such as [28], have examined relationships between the area and intensity of rainfall, showing that higher rainfall intensities are often linked to larger areas. Building on such insights, this study seeks to investigate not only these relationships but also their evolution over time, focusing on area, duration and intensity.

To quantify and analyse extremes in rainfall event characteristics, this study employs Extreme Value Theory (EVT), a statistical framework designed to model rare and extreme phenomena [10]. EVT has been widely used in meteorology to analyse extreme rainfall events, offering insights into their probability and intensity. By applying EVT to radar-derived precipitation data, this study investigates the extremes of the different rainfall characteristics.

1.2. Research Objectives

The primary aim of this study is to determine whether a change can be detected in the relationships among the spatial extent, duration, and intensity of rainfall events in the Netherlands over the period 1998–2023. To fully understand these relationships, it is also essential to examine trends in each variable individually, as changes in one may influence the others. Identifying these trends provides a more comprehensive picture of how rainfall event dynamics have evolved over time, particularly in the context of a changing climate. To achieve this, the study addresses the following research questions:

1. How can (extreme) rainfall events be mathematically described?
2. What trends in the spatial extent, duration and intensity of rainfall events can be identified from historical radar data?
3. Do the trends in spatial extent, duration, and intensity persist when considering only rainfall events occurring at high dew point temperatures (i.e., $> 15^\circ C$)?
   Since higher dew point temperatures are associated with increased moisture availability, examining trends under these conditions helps assess whether changes in rainfall patterns are linked to atmospheric moisture content.
4. How can the relationships between spatial extent, duration and intensity of rainfall events be described?
5. Have the pairwise relationships between spatial extent, duration, and intensity of rainfall events changed from 1998 to 2023?

1.3. Identifying Rainfall Events from Radar Data

To derive information about the characteristics of rainfall events – such as spatial extent, duration and intensity – it is desirable to automatically extract these events from radar or satellite observations. There exist several approaches to achieve this. One common method is a threshold-based approach, as demonstrated by [5]. In this approach, a rainfall event is defined as the fraction of the total radar area where the rain rate exceeds a specified threshold. Using this method, characteristics like area, start time, and end time of the rainfall event can be identified. However, a key limitation of this approach is that at most one event can be defined per time step.

Another approach involves the use of object tracking algorithms, which generally consists of two main steps. First, 'rainfall objects' are identified at each individual time step. In the second step, these identified rainfall objects are tracked and linked across consecutive time steps to form one 'rainfall event'. In this research, a rainfall event is interpreted as a precipitating area governed by its own distinct physical processes. If multiple, disjoint precipitating areas exist – each with its own unique processes – then it is important to detect multiple events accordingly. For this reason, an object tracking-based approach is preferred over the previously described threshold-based approach.

Many object tracking methods identify rainfall objects as continuous areas of pixels where the rain rate exceeds a certain threshold. Ebert and McBride introduced the concept of a contiguous rain area

(CRA), defining it as a region of observed and/or forecast rainfall enclosed within a specified isohyet [13]. Building on this, Davis et al. developed an object tracking method, known as MODE, to identify CRAs from satellite observations [12]. Subsequent methods extended this approach by incorporating the temporal dimension, linking spatial objects across consecutive time steps to track rainfall events [42, 34, 9].

In the second step of object tracking, where objects are linked across different time steps, two common approaches are typically used. The threshold-distance approach links objects in consecutive images if the distance between their centres is below a specified threshold [12, 21, 33]. Another widely used method is the overlap approach, which links objects from consecutive images if they share overlapping pixels [34, 22, 16].

Linking objects across time steps can present challenges when tracking fast-moving convective events, particularly when the temporal resolution is low, as objects from consecutive time steps may be too far apart [35, 22]. One way to address this is to advect objects at time step t before comparing them with those at time step $t + 1$. The advection can be determined using climatological data, velocity from the previous time step, or by iteratively calculating the velocity field between time steps using the already linked objects from t and $t + 1$ [34].

In this study, we employ an event tracking approach, called 'Celltrack', developed by Lochbihler et al. [28], which is based on the iterative rain cell tracking (IRT) method introduced by Moseley et al. [34]. This approach links objects across time steps using an overlap-based method. The Celltrack method features an iterative process that advects objects by one time step using a velocity field derived from previously detected overlaps, checks for new overlaps, updates the velocity field, and repeats the advection. This iterative cycle allows for the detection of fast-moving events. The Celltrack algorithm will be described in detail in Section 2.2.

1.4. Thesis Structure

The remainder of this thesis is structured as follows. Section 2 describes the radar data and explains the event tracking approach used in this study. Additionally, this section describes the event data, including key characteristic variables of the events. Section 3 continues by presenting the analytical methods used to examine the event data. It covers trend detection techniques, the application of EVT, methods for exploring the relationships among area, duration, and intensity, and a method for detecting changes in pairwise correlations between these event variables. Section 4 focuses on the results derived from the methods outlined in Section 3. Section 5 concludes the report with a comprehensive discussion of the results. Additionally, future research directions are suggested to build upon the insights gained from this study.

2

Radar Data and Event Tracking

2.1. Radar Data

In this study, rainfall events will be extracted from radar images, utilizing the NL21 KNMI radar rainfall dataset [25] containing 5-minute precipitation accumulations across the Netherlands. The radar pixels have a 2.4-km spatial resolution. In total, the dataset spans 26 years, from 1998 to 2023. To improve accuracy, KNMI corrected the radar data using measurements from both their automatic and manual rain gauge networks [6]. However, since the distance between rain gauges increases outside the Netherlands, the correction is effective only for radar pixels over or near the Netherlands' land surface. Beyond this range, particularly beyond 30 km from the land boundaries, rainfall intensities are significantly underestimated. Therefore, for this study, the data are masked outside the area where the rain gauge correction is effective. The complete preprocessing procedure for the radar data is described in Appendix B. The area for which precipitation events will be studied covers a total of 60,400 km^2 and is shown in Figure 2.1 .

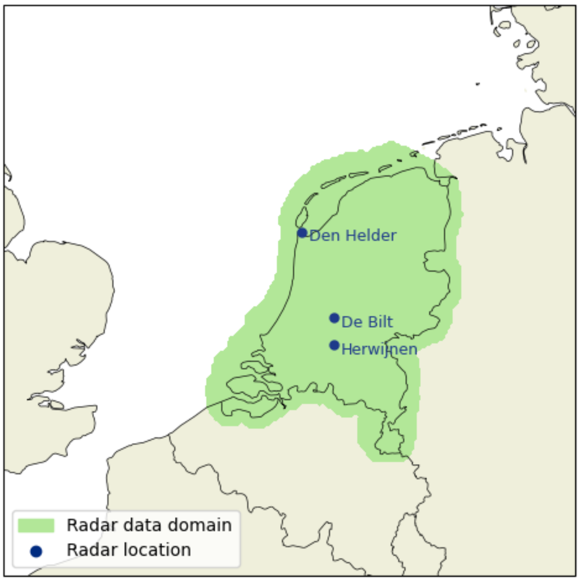

Figure 2.1: The spatial extent of the radar data used in this study, including the locations of the radars contributing to the NL21 radar-derived precipitation dataset [25] from KNMI.

Precipitation depths are derived from reflectivity data coming from KNMI radars, for which the locations are marked on the map in Figure 2.1. From 1998 until January 2001, only the De Bilt radar was used. From January 2001 until January 2017, both the De Bilt radar and the Den Helder radar were used to derive precipitation data. At the end of January 2017, the De Bilt radar was replaced by the Herwijnen radar. The radars measure reflectivity Z (dBZ), which is converted to rainfall intensity R (mm) using the relation described in Equation 2.1.

$$Z = 200 \cdot R^{1.6} \tag{2.1}$$

The radar data has an average yearly availability of 91.7%, representing the percentage of 5-minute radar images available. Figure 2.2 illustrates the availability for the 26 years, showing overall strong coverage with slightly lower availability in the dataset's earlier years.

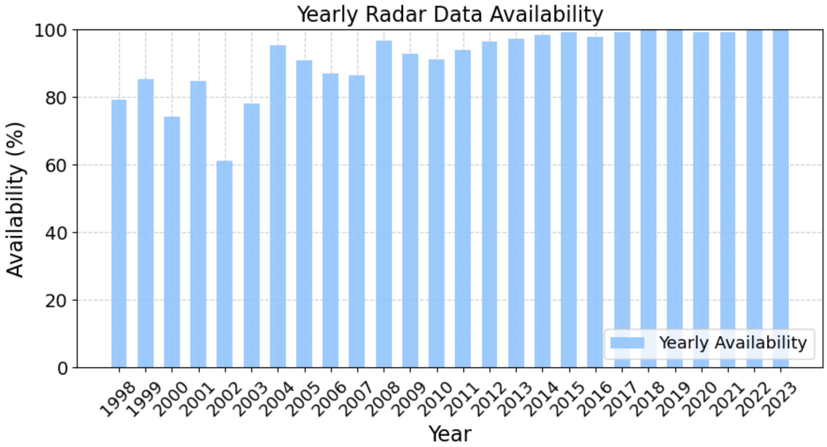

Figure 2.2: Annual availability of radar data, expressed as the percentage of available 5-minute radar images per year.

It is important to note that, while the radar data used in this study are validated and of consistently high quality for a radar-based dataset, they do not provide a perfect, direct representation of ground-level precipitation. Radar measurements can be influenced by various factors, including hardware calibration errors, fluctuations in received power, and changes in surrounding infrastructure, such as the construction of new buildings near the radar over time. In fact, such infrastructure changes contributed to the replacement of the De Bilt radar with the Herwijnen radar in 2017. These factors can introduce some variability in radar-derived precipitation estimates over time, which should be considered, especially when interpreting trends.

2.2. Event Tracking Algorithm

The algorithm 'Celltrack' as implemented by Lochbihler et al. [28] will be used to track rainfall events from the radar data as described in Section 2.1. Celltrack can be described as follows. For each time step, rain cells are detected. A rain cell is defined as a group of at least four directly neighbouring grid points exceeding an intensity threshold of 0.1 mm/5min. This threshold was selected based on precipitation statistics from STOWA [4] and to ensure that a representative range of 'regular' rainfall events is retained for the analysis.

Diagonally neighbouring grid points are not considered directly neighbouring. After the rain cell is detected, several characteristics, e.g. area, maximum intensity and the intensity weighted centre of mass, are determined by the algorithm. The algorithmic procedure, along with the accompanying decision-making process, is detailed in Appendix C.

The obtained rain cell images from consecutive time steps are checked for partially overlapping cells. If two rain cells from consecutive time steps overlap, they are considered to belong to the same 'track'. This can be used to determine the velocity and movement direction of rain cells. The obtained velocities of cells that neither split nor merge are averaged over a grid of 160 by 160 km, resulting in a large scale velocity field. Using an iterative procedure, rain cells are advected using the previously obtained velocity field, a check for new overlaps is performed, and the velocity field is then updated. This step is repeated until the procedure converges, resulting in a final set of tracks.

Defining the start and end of a track is crucial for understanding its dynamics. By definition, a track cannot contain more than one rain cell at any given time step. A track may begin in one of three ways: by splitting from a rain cell in the previous time step, by merging with other rain cells from the previous time step, or by emerging independently with no prior overlap. Similarly, a track may end by splitting into multiple rain cells in the subsequent time step, merging with other rain cells, or dissolving entirely, with no overlap in the next time step. Conceptually, tracks can be viewed as path-like structures, where nodes represent rain cells at each time step, and arcs represent temporal connections – always spanning a single time step – between these cells.

As tracks grow longer, spanning more time steps, their termination through splitting or merging becomes more likely. To account for this, Lochbihler et al. [28] introduced the concept of a 'meta-track', which combines interconnected tracks linked through splitting or merging. Meta-tracks can also be represented by graphs, where arcs correspond to individual tracks, and nodes (with a degree of at least three) signify connections between these tracks. Figure 2.3 illustrates such a graph for an example a meta-track.

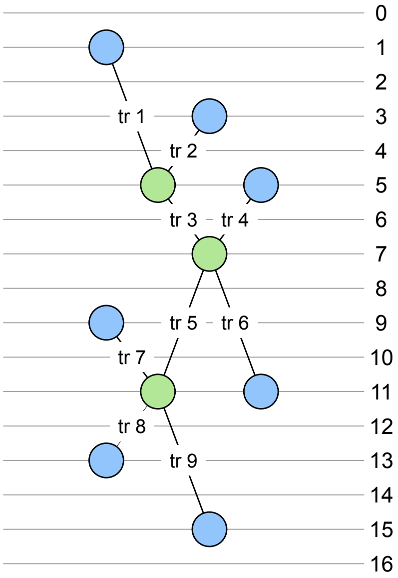

Figure 2.3: Graph construction of an example meta-track. Grey horizontal lines represent time steps. Green nodes indicate splitting or merging points (nodes with a degree of at least three). Track numbers are labeled on the arcs. Blue nodes mark the start or end of a track.

In this research, meta-tracks serve as the primary representation of rainfall events. Additionally, we include 'clean tracks' in the rainfall event dataset. These are tracks that exhibit no splitting or merging throughout their lifetime and are not part of any meta-track. Nevertheless, they are considered independent rainfall events.

Given the low thresholds set for tracking – 4 grid cells and 0.1 mm/5min – further filtering is necessary to refine the dataset. First, events lasting less than 20 minutes are excluded. Second, only events with an average area exceeding 15 grid cells (equivalent to 86.4 km^2) are included in the analysis. After applying these filters, various characteristics of the meta-tracks, including their area, duration and intensity, are evaluated. The final rainfall event dataset and its properties are detailed in Section 2.3.

2.3. Event Data

The dataset used for this analysis consists of rainfall events tracked using the Celltrack algorithm, as described in Section 2.2. These events were identified across 26 years (1998–2023) of radar-derived precipitation data (Section 2.1). After applying the filtering criteria, the dataset includes approximately 112,000 rainfall events, each characterised by a set of variables. The number of events detected per year varies, as shown in Figure 2.4, which displays the distribution of events over the study period. Notably, the percentage of events intersecting the boundaries of the radar domain at some point during their lifetime equals 2.47%. This is expected to have minimal impact on the results of the subsequent analysis.

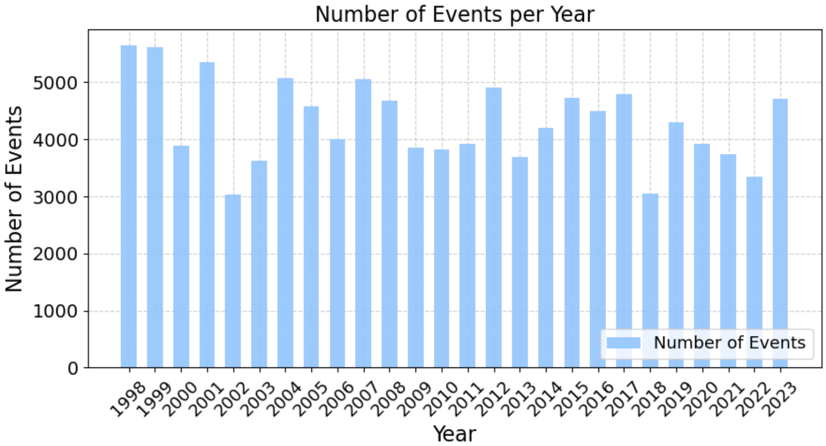

Figure 2.4: Number of detected rainfall events per year, identified using the Celltrack algorithm.

Comparing Figure 2.4 to Figure 2.2, we observe that despite lower radar data availability in 1998 and 1999 (~80%), a high number of events were detected. This may be due to radar data limitations (Section 2.1) or simply coincidence. In 2002, for instance, the low availability of data (61%) is reflected in a reduced number of detected events.

Each event is described by several variables capturing its spatial extent (area), duration and intensity. Table 2.1 lists these variables, with intensity-related variables highlighted in blue, duration-related variables in red and area-related variables in green for clarity. To ensure precision in their definitions, the following mathematical notation is introduced to model rainfall events.

- Let E denote the set of rainfall events, where $E = \{e_1, e_2, e_3, \dots\}$ and each $e \in E$ represents a distinct rainfall event.
- $T \subset \mathbb{N}$ denotes the set of time steps, with $T(e) \subseteq T$ representing the time steps during which event e occurs.
- Each event e consists of a set of rain cells $C(e) = \{c_1, c_2, c_3, \dots\}$, and $C_t(e) \subseteq C(e)$ denotes the rain cells present at time t.
- Each rain cell $c \in C_t(e)$ corresponds to a set of radar grid cells $R_t(c)$, where each radar grid cell $r \in R_t(c)$ has a precipitation intensity, or value, $p(r)$, measured in mm per 5 minutes.

Using this notation, we define the following variables for each event e.

- **Average area**:

$$A_{\mathrm{A}}(e) := \frac{1}{|T(e)|} \sum_{t \in T(e)} \left(\sum_{c \in C_t(e)} |R_t(c)| \right) \tag{2.2}$$

  This quantifies the average spatial coverage of rainfall across the event's duration.

- **Peak area**:

$$A_{\mathrm{P}}(e) := \max_{t \in T(e)} \left(\sum_{c \in C_t(e)} |R_t(c)| \right) \tag{2.3}$$

  This measures the largest spatial coverage (number of radar grid cells) observed during the event.

- **Duration**:

$$D(e) := 5 \cdot |T(e)| \tag{2.4}$$

  This represents the duration of the event in minutes (each time step being 5 minutes).

- **Average intensity**:

$$I_{\mathrm{A}}(e) := \frac{1}{|T(e)|} \sum_{t \in T(e)} \left(\frac{1}{|C_t(e)|} \sum_{c \in C_t(e)} \left(\frac{1}{|R_t(c)|} \sum_{r \in R_t(c)} p(r) \right) \right) \tag{2.5}$$

  This captures the average precipitation intensity, first averaging over grid cells in each rain cell, then over rain cells per time step, and finally over all time steps.

- **Peak intensity**:

$$I_{\mathrm{P}}(e) := \max_{t \in T(e), c \in C_t(e), r \in R_t(c)} p(r) \tag{2.6}$$

  This represents the maximum precipitation intensity observed across all grid cells and time steps of the event.

- **Peak average intensity**:

$$I_{\mathrm{PA}}(e) := \max_{t \in T(e)} \left(\frac{1}{|C_t(e)|} \sum_{c \in C_t(e)} \left(\frac{1}{|R_t(c)|} \sum_{r \in R_t(c)} p(r) \right) \right) \tag{2.7}$$

  This represents the maximum average precipitation intensity across all time steps of the event.

The above metrics are computed during a post-processing procedure following the execution of the Celltrack algorithm. This process involves iterating through each event's time steps and calculating the relevant properties for rain cells. Event-level metrics, as described in Table 2.1, such as maxima or averages, are then derived by aggregating values over all time steps.

Additionally, daily dew point temperatures are matched to the events, so that later on, a selection of events during high dew points can be made. The dew point is an indicator of the atmospheric moisture content, which is closely related to rainfall intensity. According to the Clausius-Clapeyron (CC) relation, the atmosphere's moisture holding capacity increases with temperature [45]. Since higher moisture levels are typically associated with more intense rainfall, this relationship makes the dew point an important factor to consider in the analysis.

The process for matching the dew point temperature to the events is as follows. Daily quality controlled climate data from KNMI [24] is used to select the ground temperature (T_g ($^\circ$C)) and the relative humidity (RH (%)) at the day of $t(I_{\mathrm{PA}})$ of the event. To determine then the dew point temperature (T_d ($^\circ$C)), we use the approximation as introduced by Mark G. Lawrence [26]:

$$T_d \approx T_g - \left(\frac{100 - RH}{5} \right) \tag{2.8}$$

For this study, events with $T_d \geq$15°C will be considered high dew point events.

With the event-level variables computed, the dataset provides a solid basis for analysing the characteristics and patterns of rainfall events. An examination of the variable distributions reveals that they are highly skewed, as can be seen from Figure A.1. A majority of events exhibit lower intensity, shorter duration, and smaller areas – an expected outcome for rainfall events. To address this skewness and gain a clearer understanding of the distributions, log-transformations of the variables are applied. Figure 2.5 presents histograms with 30 bins, showcasing the log-transformed distributions of several key variables from the event data.

Across the entire dataset, a small number of events exhibit extremely high intensity values that appear unrealistic. Despite the application of rain gauge corrections to the radar data, these outliers persist. However, when visualising these specific events, they resemble typical rainfall patterns (i.e. moving clusters of radar cells rather than isolated, sudden peaks). We suspect these anomalies are caused by hail, which is known to produce exceptionally high reflectivity values that may not be fully corrected by rain gauge adjustments. Since these cases are extremely rare and unlikely to significantly impact the overall results, they have been retained in the dataset.

Table 2.1: Names, descriptions and units of the variables in the rainfall event dataset. The table lists the key variables used to characterise rainfall events, describing their intensity, duration and area.

Variable	Description	Unit
e	Indicator variable of the event	-
$I_{\mathrm{P}}(e)$	Peak intensity	mm/5min
$t(I_{\mathrm{P}}(e))$	Time that $I_{\mathrm{P}}(e)$ occurs	-
$I_{\mathrm{A}}(e)$	Average intensity	mm/5min
$I_{\mathrm{PA}}(e)$	Peak average intensity	mm/5min
$t(I_{\mathrm{PA}}(e))$	Time that $I_{\mathrm{PA}}(e)$ occurs	-
$t_s(e)$	Start time	-
$D(e)$	Duration	min
$A_{\mathrm{P}}(e)$	Peak area	km^2
$A_{\mathrm{A}}(e)$	Average area	km^2
$A_{I_{\mathrm{P}}}(e)$	Area of the event at $t(I_{\mathrm{P}}(e))$	km^2
$A_{I_{\mathrm{PA}}}(e)$	Area of the event at $t(I_{\mathrm{PA}}(e))$	km^2
$T_d(e)$	Dew point temperature at $t(I_{\mathrm{PA}}(e))$	°C

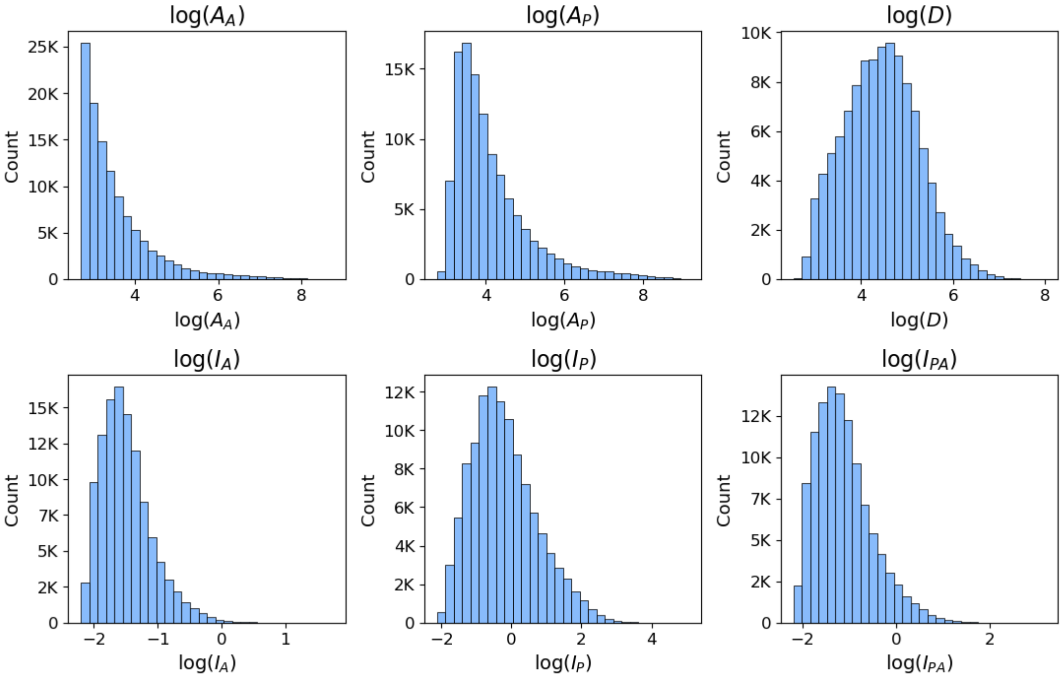

Figure 2.5: Histograms of the log-transformed event variables. Each histogram is plotted using 30 bins.

3

Statistical Methods

3.1. Trends in Individual Variables

To analyse trends in individual rainfall event variables, we begin by calculating their monthly mean values. This step helps mitigate the effects of short-term variability and missing data, while ensuring that observations are sufficiently spaced in time. This spacing is important because the Mann-Kendall test, which we use for trend detection, assumes that observations are independent and not serially correlated.

Once the monthly means are computed, we apply the Mann-Kendall test, a widely used nonparametric method for detecting trends in time series [30, 23]. The test operates under the assumption that, in the absence of a trend, the observations over time $(X_1, X_2, \ldots, X_n)$ are independent and identically distributed (i.i.d.). The test evaluates the trend by comparing all possible pairs of data points. For each pair (X_i, X_j), with $i < j$, the difference between them is calculated. The statistic S is then defined as:

$$S = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} \operatorname{sign}(X_j - X_i), \qquad \text{where } \operatorname{sign}(X_j - X_i) = \begin{cases} 1, & \text{if } X_j > X_i \\ 0, & \text{if } X_j = X_i \\ -1, & \text{if } X_j < X_i \end{cases} \tag{3.1}$$

A positive value of S indicates that later observations tend to be larger than earlier ones, while a negative value suggests that later observations tend to be smaller than earlier ones. $S = 0$ indicates no monotonic trend. The variance of S is given by:

$$\operatorname{Var}(S) = \frac{1}{18}\Big(n(n-1)(2n+5) - \sum_i r_i(r_i - 1)(2r_i + 5)\Big), \tag{3.2}$$

where i varies over the set of tied ranks and r_i is the frequency that the rank i appears.

Next, the Mann-Kendall test statistic (or Z-score) is computed as:

$$Z = \begin{cases} \frac{S-1}{\sqrt{\operatorname{Var}(S)}}, & \text{if } S > 0 \\ 0, & \text{if } S = 0 \\ \frac{S+1}{\sqrt{\operatorname{Var}(S)}}, & \text{if } S < 0 \end{cases} \tag{3.3}$$

A positive value of Z indicates an increasing trend, while a negative value suggests a decreasing trend. The hypothesis can be formulated as:

- H_0: There is no monotonic trend in the data.
- H_1: A monotonic trend exists in the data (either increasing or decreasing).

The null hypothesis H_0 is rejected if $|Z| > Z_{1-\alpha/2}$, where $Z_{1-\alpha/2}$ is the critical value from the standard normal distribution. Here we use the normality of the Z statistic, following from the central limit theorem. In this study, the significance level is $\alpha = 0.05$, meaning H_0 is rejected if $|Z| > 1.96$.

If the Mann-Kendall test indicates a monotonic trend in the data, we estimate the slope using the Theil-Sen estimator, a robust method proposed by [44] and later extended by [41]. This estimator is particularly useful in the context of precipitation data, which often exhibit high variability and extreme values, as it is less sensitive to outliers and non-normally distributed data compared to ordinary least squares (OLS) regression. In this approach, a linear model as in Equation (3.4) is fitted to the data.

$$X_t = Qt + B + \epsilon_t, \tag{3.4}$$

where Q represents the slope, B the intercept, and ϵ_t is the error term at time t. The slope Q is estimated by taking the median of all pairwise slopes between data points. The slope between any pair (X_i, X_j), with $i < j$, is defined as:

$$Q_{ij} = \frac{X_j - X_i}{j - i}, \text{ for } 1 \leq i < j \leq n. \tag{3.5}$$

Let N represent the number of pairs of data. These N values of Q_{ij} are ranked from smallest to largest and the Theil-Sen slope estimator $\hat{Q}_s$ is defined as the median of these slopes:

$$\hat{Q}_s = \begin{cases} Q_{[(N+1)/2]}, & \text{if } N \text{ is odd,} \\ \frac{Q_{[N/2]} + Q_{[(N+2)/2]}}{2}, & \text{if } N \text{ is even.} \end{cases} \tag{3.6}$$

To assess whether this slope is statistically different from zero, we obtain a confidence interval for $\hat{Q}_s$, which can be determined as follows.

$$C_\alpha = Z_{1-\alpha/2}\sqrt{\operatorname{Var}(S)}, \tag{3.7}$$

where $\operatorname{Var}(S)$ is the variance of S (from Equation 3.2) and $Z_{1-\alpha/2}$ is the $(1 - \alpha/2)$-quantile from the standard normal distribution. We then determine $M_1 = \frac{N - C_\alpha}{2}$ and $M_2 = \frac{N + C_\alpha}{2}$. The lower and upper bounds of the confidence interval, Q_L and Q_U, are the M_1-th and $(M_2 + 1)$-th largest values of the N ordered slopes, respectively.

To compute the intercept B, we estimate it by taking the median of the residuals:

$$\hat{B}(\hat{Q}_s) = \text{median}\left(X_i - \hat{Q}_s t_i\right), \tag{3.8}$$

where t_i is the time point corresponding to X_i.

To describe the uncertainty around the Theil-Sen estimated trend line, we construct confidence bands by combining the confidence intervals for both the slope Q and the intercept B. To ensure that the overall confidence band remains at the desired level $1 - \alpha$, we introduce $\beta = \alpha/2$.

For both the lower bound Q_L and the upper bound Q_U of the slope estimate, the confidence bounds for the intercept are computed using bootstrapping. This allows us to estimate the confidence around the median (see Equation 3.8), especially in cases with limited data. The bootstrapping procedure involves generating 10^4 resampled datasets from the observed data by sampling with replacement. For each resample, the median intercept $\hat{B}(Q)$ is computed. The $1 - \beta$ confidence interval for the median is then determined from the lower and upper percentiles of the bootstrapped samples.

The confidence interval for the slope is also determined at a confidence level β. Consequently, for each of Q_L and Q_U, we compute the corresponding confidence intervals for the intercept:

$$[B_L(Q_L), B_U(Q_L)] \text{ and } [B_L(Q_U), B_U(Q_U)]. \tag{3.9}$$

This leads to the following proposition for the construction of the joint confidence region for the slope and intercept:

Proposition 1 *Let Q_L and Q_U denote the lower and upper bounds of the confidence interval for the slope estimate, and let $B_L(Q)$ and $B_U(Q)$ denote the lower and upper bounds for the intercept estimate given a fixed slope Q. Then, the confidence region for the slope and intercept is given by:*

$$C = \{(Q, B) | Q \in [Q_L, Q_U], B \in [B_L(Q_U), B_U(Q_L)]\}. \tag{3.10}$$

Furthermore, the confidence bands for the trend line at time t are given by:

$$\sup_{(Q,B)\in C} \{B + Qt\}, \quad \inf_{(Q,B)\in C} \{B + Qt\}, \tag{3.11}$$

Using the fact that we construct a $1 - \beta$ confidence interval for both the slope and the intercept, we obtain the following probability bounds:

$$P\left(B \in [B_L(Q), B_U(Q)]\right) \geq 1 - \beta. \tag{3.12}$$

Additionally, the confidence region for both the slope and the intercept satisfies:

$$P\left(Q \in [Q_L, Q_U], B \in [B_L(Q_U), B_U(Q_L)]\right) \geq 1 - \alpha \tag{3.13}$$

The confidence bands derived from this proposition thus provide the $1 - \alpha$ confidence intervals around the estimated trend line for each time point t.

The results from the Mann-Kendall test and Theil-Sen estimator can be found in Section 4.1.

3.2. Trend Detection in the Extreme Value Index

Extreme value analysis is a framework for making statistical inferences about the tail of a distribution, which describes the behaviour of rare events, such as extreme rainfall. To model the frequency and magnitude of such extreme events, the data can be described using an extreme value distribution. Specifically, extreme observations exceeding a high threshold approximately follow a generalised Pareto distribution (GPD) [3].

The GPD is characterised by a shape parameter, commonly referred to as the extreme value index (EVI), which indicates the heaviness of the tail of the distribution. While classical extreme value analysis assumes that observations are independent and identically distributed (iid), in this study, we relax this assumption. The EVI is assumed to vary continuously over time, and the goal is to detect potential trends in the extremal behaviour of the size, duration, and intensity of events.

To investigate trends in the EVI for each variable, we apply the method introduced by De Haan and Zhou [20]. This approach involves estimating the EVI locally and then a global estimator for the trend is introduced to test whether the EVI remains constant over time. A description of the method follows.

We consider a sequence of observed values X_i, for $i = 1, \dots, n$ and a set of distribution functions $F_s(x)$ for $s \in [0, 1]$, such that $X_i \sim F_{\frac{i}{n}}(x)$ for $i = 1, \dots, n$. Here we assume that $F_s \in \mathcal{D}_{\gamma(s)}$, where $\mathcal{D}_{\gamma(s)}$ denotes the max-domain of attraction with EVI $\gamma(s)$. In the context of this study, and as suggested by [20], we assume that γ is positive and continuous on $[0, 1]$.

The first step of the method involves estimating $\gamma(s)$ locally for different values of $s \in [0, 1]$. To segment the data, we consider a bandwidth $h := h(n)$, such that as $n \to \infty$, $h \to 0$ and $nh \to \infty$. Only the observations X_i in the h-neighbourhood of s will be taken into account for estimating $\gamma(s)$. Specifically, we consider all X_i where $i \in I_n(s) = \left\{i : |\frac{i}{n} - s| \le h\right\}$. This results in $[2nh]$ observations for $s \in [h, 1 - h]$.

Next, $\gamma(s)$ will be estimated using the local Hill estimator. Let $k := k(n)$ such that as $n \to \infty$, $k \to \infty$ and $k/n \to 0$. Out of the $[2nh]$ local observations in the h-neighborhood of s, the top $[2kh]$ order statistics are used to estimate $\gamma(s)$ using the local Hill estimator, which is defined as follows. First, ranking the $[2nh]$ observations X_i, $i \in I_n(s)$, as

$$X^{(s)}_{1,[2nh]} \le \dots \le X^{(s)}_{[2nh],[2nh]}, \tag{3.14}$$

the local Hill estimator is defined as:

$$\hat{\gamma}_H(s) := \frac{1}{[2kh]} \sum_{i \in I_n(s)} \left(\log X_i - \log X_{[2nh]-[2kh],[2nh]}\right)^+ \tag{3.15}$$

To test whether $\gamma(s)$ remains constant across all observations, we estimate the function

$$\Gamma(s) = \int_0^s \gamma(u) du \tag{3.16}$$

and test whether it remains constant. Γ is estimated by aggregating the local estimators of $\gamma(s)$ to a ”global estimator” as follows. First, a discretised version of $\hat{\gamma}_H(s)$ is considered:

$$\hat{\gamma}_H\left(\left(2[\frac{s}{2h}] + 1\right) h\right). \tag{3.17}$$

The estimator of $\Gamma(s)$ is then defined as in Equation (3.18).

$$\hat{\Gamma}_H(s) = \int_0^s \hat{\gamma}_H\left(\left(2[\frac{u}{2h}] + 1\right) h\right) du \tag{3.18}$$

Since $\hat{\gamma}_H(s)$ is only defined for $s \leq 1-h$, its range is extended for $u > 1-h$ as follows:

$$\hat{\gamma}_H\left(\left(2[\frac{1}{2h}]+1\right)h\right) := \hat{\gamma}_H\left(\left(2[\frac{1}{2h}]-1\right)h\right). \tag{3.19}$$

To test whether the EVI remains constant over time, we use the null hypothesis $H_0 : \gamma(s) = \gamma$, for all $s \in [0,1]$, for some γ. The corresponding test statistic is defined as:

$$\tilde{T} := \sup_{s\in[0,1]} \left| \frac{\hat{\Gamma}_H(s)}{\hat{\Gamma}_H(1)} - s \right|, \tag{3.20}$$

where $\hat{\Gamma}_H(1)$ is the estimator of the constant EVI γ. Under the null hypothesis, as $n \to \infty$, we have:

$$\sqrt{k}\tilde{T} \xrightarrow{d} \sup_{s\in[0,1]} |B(s)| =: K, \tag{3.21}$$

where $B(s)$ is a standard Brownian bridge on $[0,1]$. The distribution of K is given by:

$$P(K \leq x) = 1 - 2\sum_{k=1}^{\infty}(-1)^{k-1}e^{-2k^2x^2}. \tag{3.22}$$

The null hypothesis is rejected if $p \leq 0.05$, where $p = P(K \geq \tilde{T}) = 1 - P(K \leq \tilde{T})$, which can be computed as:

$$p = 2\sum_{k=1}^{\infty}(-1)^{k-1}e^{-2k^2\tilde{T}^2} \tag{3.23}$$

If the null hypothesis is rejected for a sequence of values of k, we conclude that the EVI does not remain at a constant level over time. If the null hypothesis is not rejected, the estimator of the constant EVI $\hat{\Gamma}_H(1)$ will be determined. The results from this testing procedure applied to the area, duration, and intensity variables are presented in Section 4.2.

3.3. Relationships Between Area, Duration, and Intensity of Events

To analyse the relationships between the area, duration, and intensity of rainfall events, we first conduct a general exploration of the data. This is followed by a binned analysis to examine conditional distributions and trends. Additionally, possible changes in the pairwise correlation functions are investigated using a change point test.

3.3.1. Initial Data Exploration

The initial exploration of the relationships between variables is carried out using pairwise Pearson correlation coefficients and pair plots. Given the highly skewed distributions of the variables, rank transformations are applied to compute Spearman's rank correlations and generate corresponding pair plots of the ranked data. These steps provide a preliminary understanding of the pairwise dependencies among intensity, duration, and area.

3.3.2. Binned Analysis of Intensity Given Duration and Area

To investigate the conditional distribution of intensity given duration and spatial extent, we employ a binning approach. The key objective of this analysis is to assess how intensity varies with duration and area across different quantiles of the data. For this part of the analysis, we take the average area (A_A) as the variable describing the event area.

- **Binning process**: Duration and area are divided into bins using quantile binning, an approach that ensures each bin contains an approximately equal number of observations for each individual variable. In this study, both duration and area are divided into 5 quantiles, resulting in 25 bins (5 bins for duration and 5 for area). However, because we are binning two variables together, there can be variability in the number of events per bin. Quantile binning helps to ensure that each bin represents a decent range of values, but the exact number of events in each bin can still vary.
- **Trend analysis**: For each bin, we evaluate trends in intensity using the Mann-Kendall test. Additionally, to gain more knowledge about the behaviour of extreme events, we assess trends in the extreme value index (EVI) within each bin, using the method outlined in Section 3.2.

3.3.3. Changes in Correlation

To detect changes in the correlation between rainfall event variables, we apply the method described in [32]. Specifically, we analyse changes in the correlation between two variables, $X_{t,n}$ and $Z_{t,n}$, where $t = 1, \ldots, n$, and n is the number of observations. The analysis follows the framework introduced in [32], adapted to our specific case.

We define the vector:

$$Y_{t,n} = \left(X_{t,n}, Z_{t,n}, X_{t,n}^2, Z_{t,n}^2, X_{t,n} Z_{t,n}\right) \tag{3.24}$$

We are interested in changes in the correlation function between $X_{t,n}$ and $Z_{t,n}$, which can be expressed as:

$$\rho_{t,n} := \mathrm{Cor}(X_{t,n}, Z_{t,n}) = f(\mathbb{E}Y_{t,n}) = f(\mu_{t,n}), \tag{3.25}$$

where $\mu_{t,n}$ is the expectation of $Y_{t,n}$ at time t. The function f is defined as:

$$f(x) = \frac{x_5 - x_1 x_2}{\sqrt{(x_3 - x_1^2)(x_4 - x_2^2)}}, \tag{3.26}$$

where x_1, x_2, x_3, x_4, x_5 are the elements of $Y_{t,n}$. This function determines the correlation based on the expectations of the variables in $Y_{t,n}$.

To estimate $\mu_{t,n}$, we use a rolling window average:

$$\hat{\mu}_{t,n} = \frac{1}{k} \sum_{i=t-k+1}^{t} Y_{i,n}, \tag{3.27}$$

where k is the window bandwidth.

Following [32], we define the integrated correlation function:

$$F_n(u) = \int_0^u f(\mu_{v,n})dv \tag{3.28}$$

which captures changes in the correlation while accounting for variations in mean and variance. $F_n(u)$ is then estimated by $M_n(u)$:

$$M_n(u) = \frac{1}{n}\sum_{t=1}^{\lfloor nu \rfloor} \left[\hat{\rho}_{t-1,n} + D\hat{\rho}_{t-1,n}(Y_{t,n} - \hat{\mu}_{t-1,n})\right], \quad u \in [0,1], \tag{3.29}$$

where $\hat{\rho}_{t-1,n} = f(\hat{\mu}_{t-1,n})$, $D\hat{\rho}_{t-1,n}$ is the derivative of $\hat{\rho}_{t-1,n}$, and $Y_{t,n} - \hat{\mu}_{t-1,n}$ represents residuals (i.e., deviations from the smoothed estimate). Note that, in this study, we assume independent observations, which is not covered in [32]. However, a formal proof that the method remains valid in this setting is beyond the scope of this work. Under the assumption of independent observations, we set the lag L and initial offset τ from the framework of [32] to $L = 1$ and $\tau = 0$.

To formally assess whether the correlation remains constant over time, we test the hypothesis:

$$H_0 : \rho_{u,n} = \rho_{0,n} \text{ for all } u \in [0,1] \quad \leftrightarrow \quad H_1 : \rho_{u,n} \neq \rho_{0,n} \text{ for some } u \in [0,1]. \tag{3.30}$$

To evaluate this hypothesis, we compute the CUSUM statistic:

$$T_n^* = \sup_{u\in[0,1]} |T_n(u)|, \quad \text{where } T_n(u) = M_n(u) - uM_n(1) \tag{3.31}$$

The critical value for this test is derived through bootstrapping to obtain the limit distribution of T_n^*. For this, we first resample $M_n(u)$ as:

$$\hat{M}_n(u) = \frac{1}{\sqrt{n}}\sum_{t=1}^{\lfloor nu \rfloor} N_t \left[D\hat{\rho}_{t-1,n}(Y_{t,n} - \hat{\mu}_{t-1,n})\right], \tag{3.32}$$

where $N_t \sim \mathcal{N}(0,1)$ are i.i.d. standard normal variables. From Theorem 3.4 of [32], it follows that the conditional distribution of $\hat{M}_n$ given the data converges in probability to $M(u)$, the limit process of $M_n(u)$.

Next, we determine:

$$\hat{T}_n^* = \sup_{u\in[0,1]} |\hat{M}_n(u) - u\hat{M}_n(1)|. \tag{3.33}$$

From Theorem 3.4, we know that $\hat{T}_n^*$ given the data converges in distribution to the limit distribution T^*. We denote by t_α the $1-\alpha$ quantile of T^*, and by $t_{\alpha,n}$ the $1-\alpha$ quantile of $\hat{T}_n^*$ conditioned on the data. The quantile $t_{\alpha,n}$ can be approximated by sampling from the conditional distribution of $\hat{T}_n^*$, as described in Proposition 2 (which corresponds to Proposition 4.1 from [32]), and we obtain the critical value $t_{\alpha,n}$ for our CUSUM test:

Proposition 2 *Let the conditions of Theorem 3.4 hold, and denote by $t_{\alpha,n}$ the $1-\alpha$ quantile of the conditional distribution $T_n^* \mid X_n$. If the null hypothesis holds, and if Var$(M(u_0)) > 0$ for some $u_0 \in (0,1)$, then*

$$\lim_{n\to\infty} P(T_n^* > t_{\alpha,n}) = \alpha, \quad n \to \infty.$$

Hence, rejecting H_0 if the test statistic T_n^* exceeds the critical value $t_{\alpha,n}$ leads to a test with nominal size $\alpha \in (0,1)$ asymptotically.

In our analysis, we use 10^4 bootstrap samples to obtain the critical values.

4

Results

4.1. Trends in Individual Variables

This section presents the results of trend detection for each individual variable (area, duration, and intensity) using the Mann-Kendall test and the Theil-Sen estimator for the slope, as described in Section 3.1. The results are organised by variable type, and time series of the monthly means for each variable are considered. These time series are visualised in Figure 4.1.

For all variables, we observe high variability in the data, which is partly attributed to seasonality. For this reason, the variables will also be studied separately for each season.

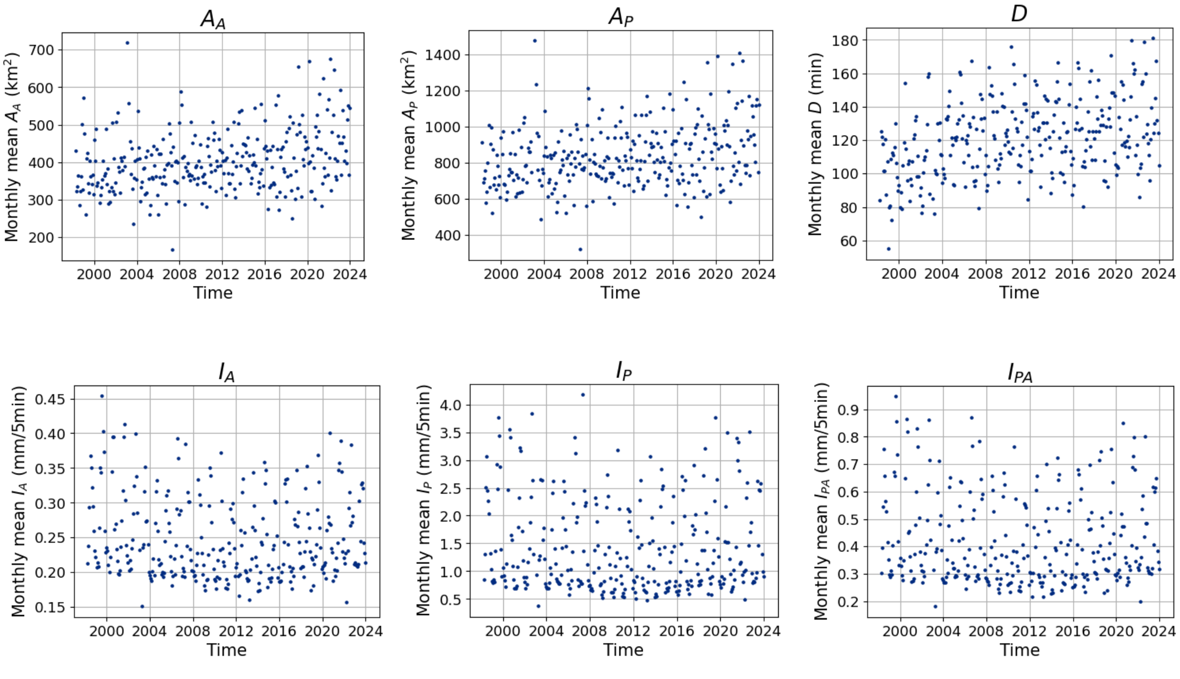

Figure 4.1: Time series of monthly means of the event variables.

4.1.1. Area

First, we analyse the variable A_{A} (average area) for trends. The time series of the monthly means of A_{A} is displayed in Figure 4.2. Since a similar approach is applied to the other variables, we describe the trend detection methods in more detail here for A_{A}.

Application of the Mann-Kendall test yields a Z-score of 4.90, leading to rejection of the null hypothesis (no monotonic trend) at the 0.05 significance level. This result indicates an increasing trend in the monthly means of A_{A}.

Next, the slope of the trend is estimated using the Theil-Sen estimator, yielding a slope of 0.266. The 95% confidence interval for the slope is [0.161, 0.369], indicating a statistically significant trend (since both bounds are positive). These results are summarised in Table 4.1, and the time series, including the trend line, is presented in Figure 4.2.

Table 4.1: Results of the Mann-Kendall test and Theil-Sen's slope estimator for the monthly means of A_{A} (average area). The table shows the test statistic $\boldsymbol{Z}$, the estimated slope $\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$, and the lower ($\boldsymbol{Q}_{\boldsymbol{L}}$) and upper ($\boldsymbol{Q}_{\boldsymbol{U}}$) 95% confidence bounds for the slope estimate.

$\boldsymbol{Z}$	$\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$	$\boldsymbol{Q}_{\boldsymbol{L}}$	$\boldsymbol{Q}_{\boldsymbol{U}}$
4.90	0.266	0.161	0.369

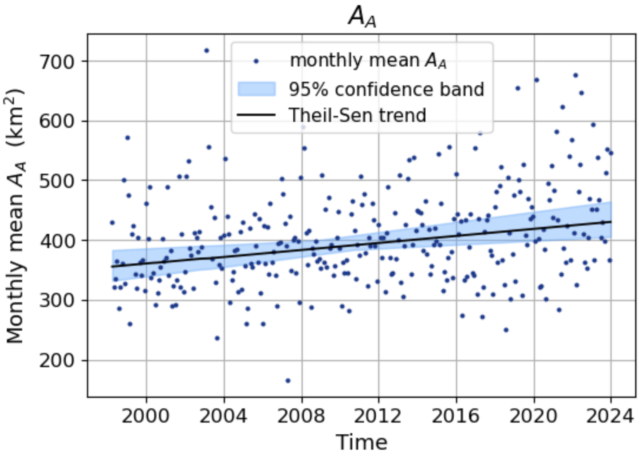

Figure 4.2: Monthly mean values of the average area (A_A) of rainfall events over the study period. The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.

To examine trends for specific seasons, four new time series are obtained by calculating seasonal means instead of monthly averages. For example, to compute the summer means, the mean of A_{A} is calculated for each summer, resulting in a time series with one data point per year. In addition, trends in the A_{A} of high-intensity events – defined as those in the top 10% of I_{PA} (peak average intensity) – were tested to understand how the area of more intense events evolves. The results of the trend detection methods for each data subset are presented in Table 4.2. Data subsets for which a significant trend is detected are marked with an asterisk (*) in the table. The time series for the seasonal subsets, including trends, can be found in Appendix A.2 (Figure A.2).

We note significantly increasing trends in the monthly means, as well as in the seasonal subsets, except for winter events. For high-intensity events, only the trend in summer means is significantly increasing.

Results for the other area variable, A_{P} (peak area), can be found in Table 4.3. Data subsets with significant trends are marked with an asterisk (*) in the table. The monthly means of A_{P}, including the trend, are shown in Figure 4.3. For all subsets, the time series with trend can be found in Appendix A.2 (Figure A.3). Significantly increasing trends are found in the monthly means and in all seasons separately. For high-intensity events, only the summer events exhibit a significantly increasing trend in A_{P}.

Table 4.2: Results of the Mann-Kendall test and Theil-Sen's slope estimator for different seasons and subsets of A_{A} (average area). The table includes the test statistic ($\boldsymbol{Z}$), the Theil-Sen slope estimator ($\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$), and the 95% confidence bounds for the slope estimate ($\boldsymbol{Q_L}$, $\boldsymbol{Q_U}$). * : Significant trend at the 0.05 significance level.

Data subset	$\boldsymbol{Z}$	$\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$	$\boldsymbol{Q_L}$	$\boldsymbol{Q_U}$
Monthly means	4.90*	0.245	0.149	0.340
Summer means	3.88*	4.019	2.281	5.836
Autumn means	2.42*	2.882	0.732	5.010
Winter means	1.59	2.993	−1.028	6.530
Spring means	2.12*	2.770	0.226	5.155
Monthly means (high-intensity events)	1.72	0.539	−0.081	1.157
Summer means (high-intensity events)	2.87*	8.589	2.827	16.42
Autumn means (high-intensity events)	0.04	0.211	−13.03	14.56
Winter means (high-intensity events)	0.35	9.306	−47.30	58.99
Spring means (high-intensity events)	1.59	13.58	−1.452	30.17

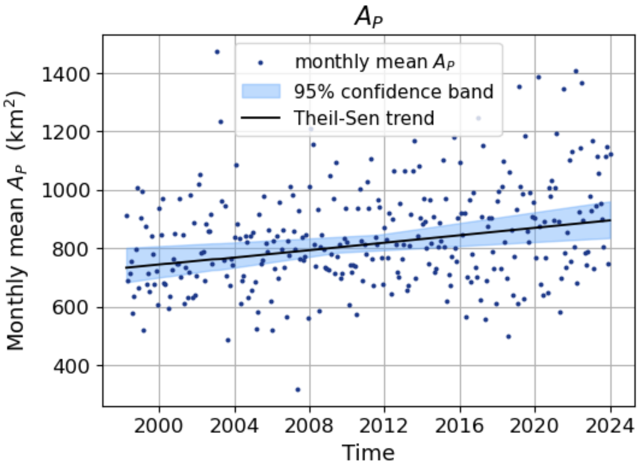

Figure 4.3: Monthly mean values of A_{P} (peak area) of rainfall events over the study period. The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.

Table 4.3: Results of the Mann-Kendall test and Theil-Sen's slope estimator for different seasons and subsets of A_{P} (peak area). The table includes the test statistic ($\boldsymbol{Z}$), the Theil-Sen slope estimator ($\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$), and the 95% confidence bounds for the slope estimate ($\boldsymbol{Q_L}$, $\boldsymbol{Q_U}$). * : Significant trend at the 0.05 significance level.

Data subset	$\boldsymbol{Z}$	$\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$	$\boldsymbol{Q_L}$	$\boldsymbol{Q_U}$
monthly means	4.83*	0.531	0.312	0.739
summer means	3.48*	8.193	4.117	12.65
autumn means	2.60*	7.913	1.892	11.45
winter means	2.03*	6.887	0.862	12.95
spring means	2.20*	5.784	0.453	10.63
monthly means (high-intensity events)	1.67	1.265	−0.216	2.647
summer means (high-intensity events)	3.09*	18.17	8.383	40.04
autumn means (high-intensity events)	0.26	5.039	−20.97	41.45
winter means (high-intensity events)	0.48	35.08	−91.27	147.5
spring means (high-intensity events)	1.59	28.30	−4.133	70.10

4.1.2. Duration

Results from the Mann-Kendall test and Theil-Sen estimator for the slope, regarding the D (duration) variable, are shown in Table 4.4. The monthly means of D, including the trend, are given in Figure 4.4. Data subsets for which a significant trend is detected are marked with an asterisk (*) in the table. For these subsets, the time series, including the detected trends, are given in Appendix A.2 (Figure A.4). We note a significantly increasing trend in the monthly means, as well as for all separate seasons. For the high-intensity events, we note significantly increasing trends for the monthly means and for the summer and autumn. High-intensity events in winter and spring seasons do not exhibit significant trends according to the Mann-Kendall test.

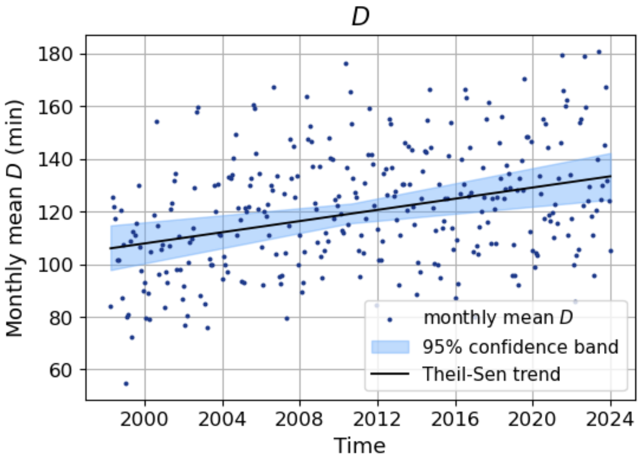

Figure 4.4: Monthly mean values of D (duration) of rainfall events over the study period. The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.

Table 4.4: Results of the Mann-Kendall test and Theil-Sen's slope estimator for different seasons and subsets of D (duration). The table includes the test statistic ($\boldsymbol{Z}$), the Theil-Sen slope estimator ($\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$), and the 95% confidence bounds for the slope estimate ($\boldsymbol{Q_L}$, $\boldsymbol{Q_U}$). * : Significant trend at the 0.05 significance level.

Data subset	$\boldsymbol{Z}$	$\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$	$\boldsymbol{Q_L}$	$\boldsymbol{Q_U}$
monthly means	5.95^{*}	0.090	0.061	0.118
summer means	2.38^{*}	0.916	0.243	1.632
autumn means	3.53^{*}	1.095	0.653	1.517
winter means	3.00^{*}	0.879	0.363	1.492
spring means	2.82^{*}	1.190	0.442	1.880
monthly means (high-intensity events)	2.87^{*}	0.128	0.041	0.208
summer means (high-intensity events)	3.09^{*}	1.497	0.693	2.539
autumn means (high-intensity events)	2.82^{*}	2.299	0.850	3.694
winter means (high-intensity events)	0.71	1.042	-2.863	4.802
spring means (high-intensity events)	1.41	1.818	-1.057	4.451

4.1.3. Intensity

The results of the Mann-Kendall test and Theil-Sen estimator for the intensity-related variables – I_{A}, I_{P}, and I_{PA} – are presented in Tables 4.5, 4.6 and 4.7, respectively. The monthly means of the variables, including the trend, are given in Figure 4.5, where the confidence region is plotted in grey if there is no significant trend according to the Mann-Kendall test.

From the Z-scores in Table 4.6, it is evident that none of the subsets of I_{P} exhibit significant trends. However, for the I_{A} and I_{PA} variables, significant decreasing trends are observed for means of events occurring in the summer season. For high-intensity events, we notice no significant trends for I_{P} and I_{A} for any season. For I_{PA}, there is a significant decreasing trend in the summer season, for high-intensity events. The corresponding time series, along with the trend lines, are visualised in Appendix A.2 in Figures A.5, A.6 and A.7 for the variables I_{A}, I_{P} and I_{PA}, respectively.

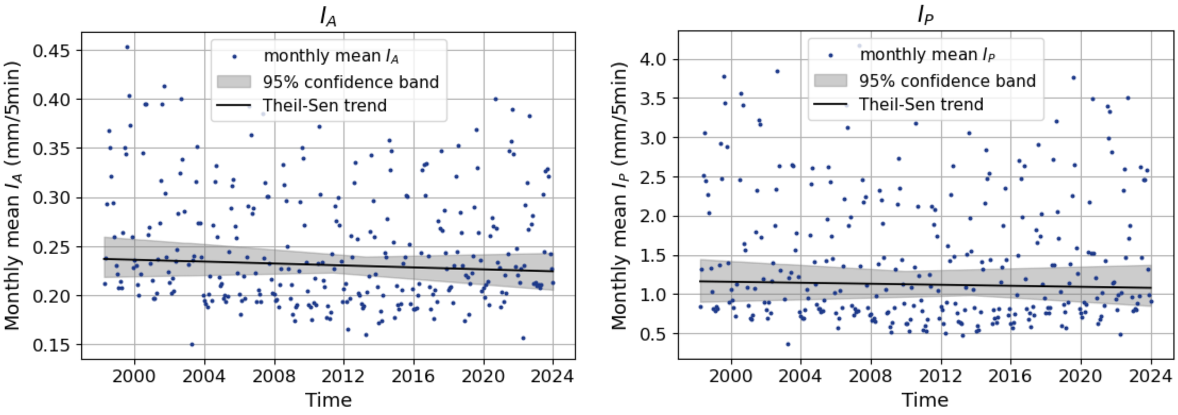

(a) I_{A} (average intensity)

(b) I_{P} (peak intensity)

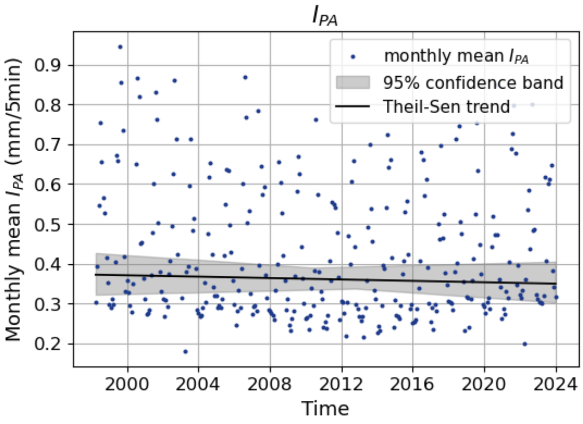

(c) I_{PA} (peak average intensity)

Figure 4.5: Monthly mean values of the intensity variables over the study period. The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.

Table 4.5: Results of the Mann-Kendall test and Theil-Sen's slope estimator for different seasons and subsets of I_{A} (average intensity). The table includes the test statistic ($\boldsymbol{Z}$), the Theil-Sen slope estimator ($\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$), and the 95% confidence bounds for the slope estimate ($\boldsymbol{Q_L}$, $\boldsymbol{Q_U}$). * : Significant trend at the 0.05 significance level.

Data subset	$\boldsymbol{Z}$	$\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$	$\boldsymbol{Q_L}$	$\boldsymbol{Q_U}$
monthly means	-1.33	$-4.11 \cdot 10^{-5}$	$-1.05 \cdot 10^{-4}$	$2.06 \cdot 10^{-5}$
summer means	-2.07^{*}	$-1.53 \cdot 10^{-3}$	$-3.18 \cdot 10^{-3}$	$-7.81 \cdot 10^{-5}$
autumn means	-1.15	$-5.07 \cdot 10^{-4}$	$-1.77 \cdot 10^{-3}$	$6.84 \cdot 10^{-4}$
winter means	-0.97	$-3.19 \cdot 10^{-4}$	$-1.27 \cdot 10^{-3}$	$5.39 \cdot 10^{-4}$
spring means	0.31	$1.35 \cdot 10^{-4}$	$-9.07 \cdot 10^{-4}$	$8.96 \cdot 10^{-4}$
monthly means (high-intensity events)	0.10	$5.42 \cdot 10^{-6}$	$-1.08 \cdot 10^{-4}$	$1.20 \cdot 10^{-4}$
summer means (high-intensity events)	-1.63	$-1.78 \cdot 10^{-3}$	$-3.18 \cdot 10^{-3}$	$7.27 \cdot 10^{-4}$
autumn means (high-intensity events)	-0.22	$-3.04 \cdot 10^{-4}$	$-2.28 \cdot 10^{-3}$	$2.08 \cdot 10^{-3}$
winter means (high-intensity events)	1.19	$2.57 \cdot 10^{-3}$	$-1.29 \cdot 10^{-3}$	$5.89 \cdot 10^{-3}$
spring means (high-intensity events)	0.44	$5.76 \cdot 10^{-4}$	$-2.17 \cdot 10^{-3}$	$2.59 \cdot 10^{-3}$

Table 4.6: Results of the Mann-Kendall test and Theil-Sen's slope estimator for different seasons and subsets of I_{P} (peak intensity). The table includes the test statistic ($\boldsymbol{Z}$), the Theil-Sen slope estimator ($\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$), and the 95% confidence bounds for the slope estimate ($\boldsymbol{Q_L}$, $\boldsymbol{Q_U}$). * : Significant trend at the 0.05 significance level.

Data subset	$\boldsymbol{Z}$	$\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$	$\boldsymbol{Q_L}$	$\boldsymbol{Q_U}$
monthly means	-0.84	$2.73 \cdot 10^{-4}$	$-9.99 \cdot 10^{-4}$	$3.82 \cdot 10^{-4}$
summer means	-1.76	$-1.66 \cdot 10^{-2}$	$-3.65 \cdot 10^{-2}$	$4.05 \cdot 10^{-2}$
autumn means	-0.75	$4.88 \cdot 10^{-3}$	$-1.71 \cdot 10^{-2}$	$9.52 \cdot 10^{-3}$
winter means	-0.53	$-2.00 \cdot 10^{-3}$	$-7.91 \cdot 10^{-3}$	$4.06 \cdot 10^{-3}$
spring means	0.53	$2.00 \cdot 10^{-3}$	$-8.73 \cdot 10^{-3}$	$1.29 \cdot 10^{-3}$
monthly means (high-intensity events)	-1.20	$-1.22 \cdot 10^{-3}$	$-3.32 \cdot 10^{-3}$	$7.69 \cdot 10^{-4}$
summer means (high-intensity events)	-0.93	$-1.45 \cdot 10^{-2}$	$-5.23 \cdot 10^{-2}$	$1.83 \cdot 10^{-2}$
autumn means (high-intensity events)	-0.40	$-6.25 \cdot 10^{-3}$	$-4.55 \cdot 10^{-2}$	$2.55 \cdot 10^{-2}$
winter means (high-intensity events)	0.22	$5.05 \cdot 10^{-3}$	$-3.85 \cdot 10^{-2}$	$4.83 \cdot 10^{-2}$
spring means (high-intensity events)	-0.35	$-7.88 \cdot 10^{-3}$	$-5.23 \cdot 10^{-2}$	$4.42 \cdot 10^{-2}$

Table 4.7: Results of the Mann-Kendall test and Theil-Sen's slope estimator for different seasons and subsets of I_{PA} (peak average intensity). The table includes the test statistic ($\boldsymbol{Z}$), the Theil-Sen slope estimator ($\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$), and the 95% confidence bounds for the slope estimate ($\boldsymbol{Q_L}$, $\boldsymbol{Q_U}$). * : Significant trend at the 0.05 significance level.

Data subset	$\boldsymbol{Z}$	$\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$	$\boldsymbol{Q_L}$	$\boldsymbol{Q_U}$
monthly means	0.97	$-7.36 \cdot 10^{-5}$	$-2.33 \cdot 10^{-4}$	$6.89 \cdot 10^{-5}$
summer means	-2.47^{*}	$-6.02 \cdot 10^{-3}$	$-8.34 \cdot 10^{-3}$	$-1.02 \cdot 10^{-3}$
autumn means	-1.01	$-1.22 \cdot 10^{-3}$	$-3.62 \cdot 10^{-3}$	$1.59 \cdot 10^{-3}$
winter means	-0.53	$-3.28 \cdot 10^{-4}$	$-1.62 \cdot 10^{-3}$	$9.55 \cdot 10^{-4}$
spring means	0.22	$1.83 \cdot 10^{-4}$	$-1.48 \cdot 10^{-3}$	$2.36 \cdot 10^{-3}$
monthly means (high-intensity events)	-1.27	$-1.78 \cdot 10^{-4}$	$-4.41 \cdot 10^{-4}$	$9.54 \cdot 10^{-5}$
summer means (high-intensity events)	-2.60^{*}	$-8.23 \cdot 10^{-3}$	$-1.24 \cdot 10^{-2}$	$-1.74 \cdot 10^{-3}$
autumn means (high-intensity events)	-0.53	$-1.83 \cdot 10^{-3}$	$-8.24 \cdot 10^{-3}$	$4.05 \cdot 10^{-3}$
winter means (high-intensity events)	1.50	$3.62 \cdot 10^{-3}$	$-8.73 \cdot 10^{-4}$	$8.13 \cdot 10^{-3}$
spring means (high-intensity events)	-0.13	$-2.53 \cdot 10^{-4}$	$-6.04 \cdot 10^{-3}$	$5.11 \cdot 10^{-3}$

Upon further analysis, we consider whether the initial filtering criteria, particularly the area constraint (events with an average area larger than 15 grid cells), might have been too restrictive. This constraint may have filtered out some convective events, which are known for their high intensity but smaller spatial extent and shorter duration. To address this, the event data is filtered again, focusing specifically on smaller-scale shorter-duration higher-intensity (SSHI) events characterised specifically by durations of less than 300 minutes (but longer than 10 minutes to minimise noise), areas smaller than 30 grid cells (172.8 km^2), and an average intensity greater than 0.4 mm/5min. This results in a new dataset consisting of 18,600 SSHI events. Due to the limited dataset length, trends were analysed using only monthly averages rather than seasonal subsets.

The results of this analysis, including the trends in the intensity variables for these events, are presented in Table 4.8. Notably, significant increasing trends in monthly averages of I_{A} and I_{PA} were observed for these events. The corresponding time series and trend lines for the intensity variables with significant trend (i.e. I_{A} and I_{PA}) are visualised in Figure 4.6. The time series and trends for the other event variables for these SSHI events are given in Appendix A.2.4 in Figure A.8. We furthermore note that, in contrast to what we previously saw, there are no significant trends in the area for this subset of events. There is, however, a significantly increasing trend in D.

Table 4.8: Results of the Mann-Kendall test and Theil-Sen slope estimator for rainfall event variables, for smaller-scale shorter-duration higher-intensity (SSHI) events (i.e. durations shorter than 300 minutes, areas smaller than 30 grid cells, and average intensity greater than 0.4 mm/5min). The table includes the test statistic ($\boldsymbol{Z}$), the Theil-Sen slope estimator ($\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$), and the 95% confidence bounds for the slope estimate ($\boldsymbol{Q_L}$, $\boldsymbol{Q_U}$). * : Significant trend at the 0.05 significance level.

Variable	$\boldsymbol{Z}$	$\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$	$\boldsymbol{Q_L}$	$\boldsymbol{Q_U}$
I_{A} (average intensity)	2.01^*	$9.38 \cdot 10^{-5}$	$2.42 \cdot 10^{-6}$	$1.91 \cdot 10^{-4}$
I_{P} (peak intensity)	0.87	$4.74 \cdot 10^{-4}$	$-5.76 \cdot 10^{-4}$	$1.60 \cdot 10^{-3}$
I_{PA} (peak average intensity)	2.51^*	$3.37 \cdot 10^{-4}$	$7.75 \cdot 10^{-5}$	$6.06 \cdot 10^{-4}$
A_{A} (average area)	-0.33	$-2.13 \cdot 10^{-3}$	$-1.57 \cdot 10^{-2}$	$1.17 \cdot 10^{-2}$
A_{P} (peak area)	-0.41	$-5.33 \cdot 10^{-3}$	$-2.99 \cdot 10^{-2}$	$2.05 \cdot 10^{-2}$
D (duration)	4.15^*	$4.01 \cdot 10^{-2}$	$2.14 \cdot 10^{-2}$	$5.78 \cdot 10^{-2}$

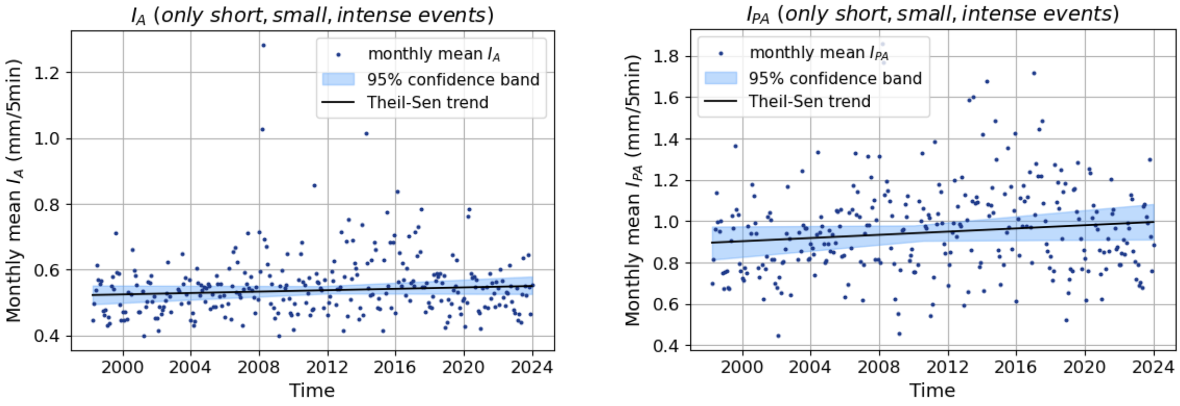

(a) **(b)**

Figure 4.6: Monthly mean values of (a) I_{A} and (b) I_{PA} for smaller-scale shorter-duration higher-intensity (SSHI) events (durations shorter than 300 minutes, areas smaller than 30 grid cells and average intensity greater than 0.4 mm/5min). The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.

4.1.4. Trends for Events During High Dew Points

Dew point temperature is an indicator of atmospheric moisture, with higher dew points signifying higher water vapor availability. Rainfall events occurring under high dew point conditions are more likely to produce heavy precipitation due to the increased potential for condensation. In the context of climate change, rising temperatures increase the atmosphere's capacity to hold moisture, leading to more frequent occurrences of high dew point conditions. For this study, we take as a subset the events with $T_d \geq 15\,°\mathrm{C}$. This leads to a total of 5941 events across all 26 years. For trend analysis within this data selection, only monthly means are considered.

For A_A and A_P, the results of the Mann-Kendall test are given in Table 4.9. We note that there are significantly increasing trends in both variables. The time series for both variables, including their trend estimated with the Theil-Sen estimator are visualised in Figure 4.7.

Table 4.9: Results of the Mann-Kendall test and Theil-Sen's slope estimator for area variables when $T_d \geq 15\,°\mathrm{C}$. The table includes the test statistic ($\boldsymbol{Z}$), the Theil-Sen slope estimator ($\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$), and the 95% confidence bounds for the slope estimate ($\boldsymbol{Q_L}$, $\boldsymbol{Q_U}$). * : Significant trend at the 0.05 significance level.

Variable	$\boldsymbol{Z}$	$\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$	$\boldsymbol{Q_L}$	$\boldsymbol{Q_U}$
A_A (average area)	2.23^*	1.149	0.132	2.218
A_P (peak area)	2.08^*	2.473	0.210	4.970

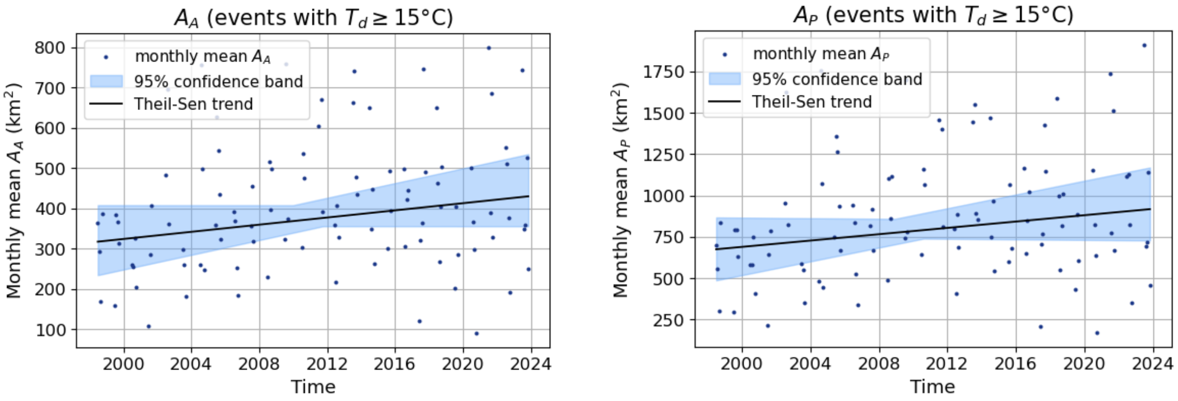

(a) monthly means of A_A (average area) with $T_d \geq 15\,°\mathrm{C}$ **(b)** monthly means of A_P (peak area) with $T_d \geq 15\,°\mathrm{C}$

Figure 4.7: Monthly mean values of area variables when $T_d \geq 15\,°\mathrm{C}$. The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.

The results of the Mann-Kendall test for the duration are given in Table 4.10 and the corresponding time series is visualised in Figure 4.8. We find a significantly increasing trend in the monthly means of the duration.

Table 4.10: Results of the Mann-Kendall test and Theil-Sen's slope estimator for D (duration) when $T_d \geq 15\,°\mathrm{C}$. The table includes the test statistic ($\boldsymbol{Z}$), the Theil-Sen slope estimator ($\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$), and the 95% confidence bounds for the slope estimate ($\boldsymbol{Q_L}$, $\boldsymbol{Q_U}$). * : Significant trend at the 0.05 significance level.

Variable	$\boldsymbol{Z}$	$\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$	$\boldsymbol{Q_L}$	$\boldsymbol{Q_U}$
D (duration)	3.92^*	0.476	0.255	0.697

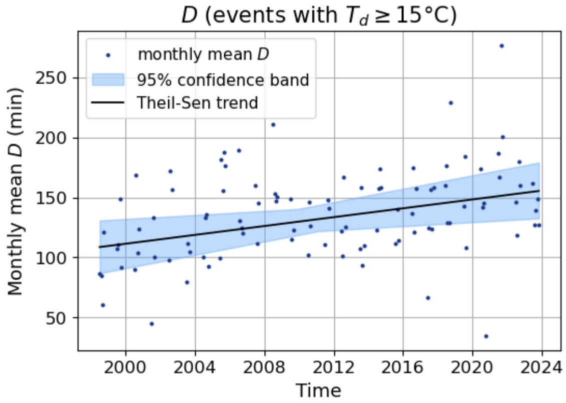

Figure 4.8: Monthly mean values of D (duration) when $T_d \geq 15$ °C. The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.

For the intensity variables, the results are summarised in Table 4.11 and the time series are visualised in Figure 4.9. There are no significant trends in any of the variables, according to the Mann-Kendall test.

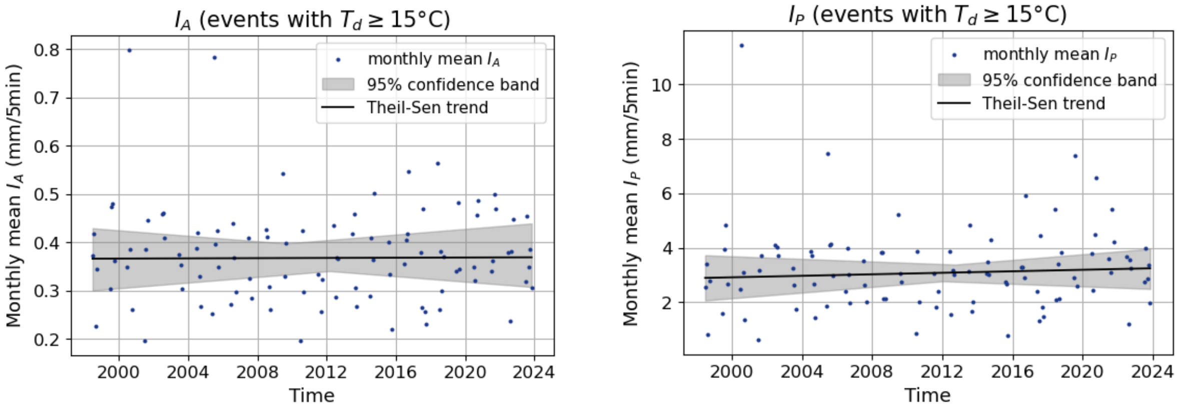

(a) monthly means of I_{A} (average intensity) with $T_d \geq 15$ °C

(b) monthly means of I_{P} (peak intensity) with $T_d \geq 15$ °C

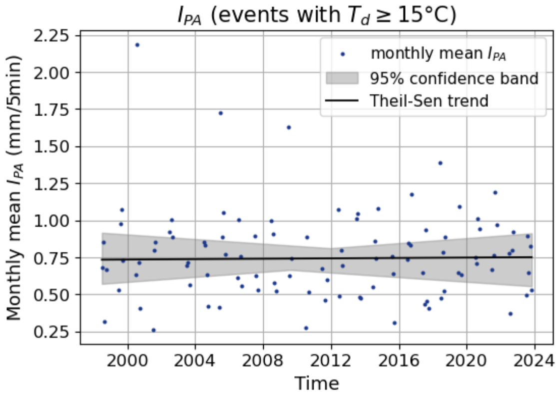

(c) monthly means of I_{PA} (peak average intensity) with $T_d \geq 15$ °C

Figure 4.9: Monthly mean values of intensity variables when $T_d \geq 15$ °C. The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.

Table 4.11: Results of the Mann-Kendall test and Theil-Sen's slope estimator for intensity variables when $T_d \geq 15$ °C. The table includes the test statistic ($\boldsymbol{Z}$), the Theil-Sen slope estimator ($\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$), and the 95% confidence bounds for the slope estimate ($\boldsymbol{Q_L}$, $\boldsymbol{Q_U}$). * : Significant trend at the 0.05 significance level.

Variable	$\boldsymbol{Z}$	$\hat{\boldsymbol{Q}}_{\boldsymbol{s}}$	$\boldsymbol{Q_L}$	$\boldsymbol{Q_U}$
I_{A} (average intensity)	0.07	$2.84 \cdot 10^{-5}$	$-6.23 \cdot 10^{-4}$	$6.83 \cdot 10^{-4}$
I_{P} (peak intensity)	0.76	$3.58 \cdot 10^{-3}$	$-5.04 \cdot 10^{-3}$	$1.23 \cdot 10^{-2}$
I_{PA} (peak average intensity)	0.18	$1.17 \cdot 10^{-4}$	$-1.72 \cdot 10^{-3}$	$1.91 \cdot 10^{-3}$

4.2. Trend Detection in the Extreme Value Index

This section presents the results of the method for detecting trends in the extreme value index (EVI), as described in Section 3.2. The method is applied to each rainfall event variable to determine whether to reject the null hypothesis H_0 : the EVI remains constant over time. For this, p-values are plotted against various values of k. If H_0 cannot be rejected for a range of k values, the constant EVI is estimated using $\hat{\Gamma}_H(1)$.

4.2.1. Area

Figures 4.10a and 4.10b display the p-values for varying k for the variables A_A and A_P, respectively. Based on these plots, H_0 cannot be rejected under the 0.05 significance level for either variable.

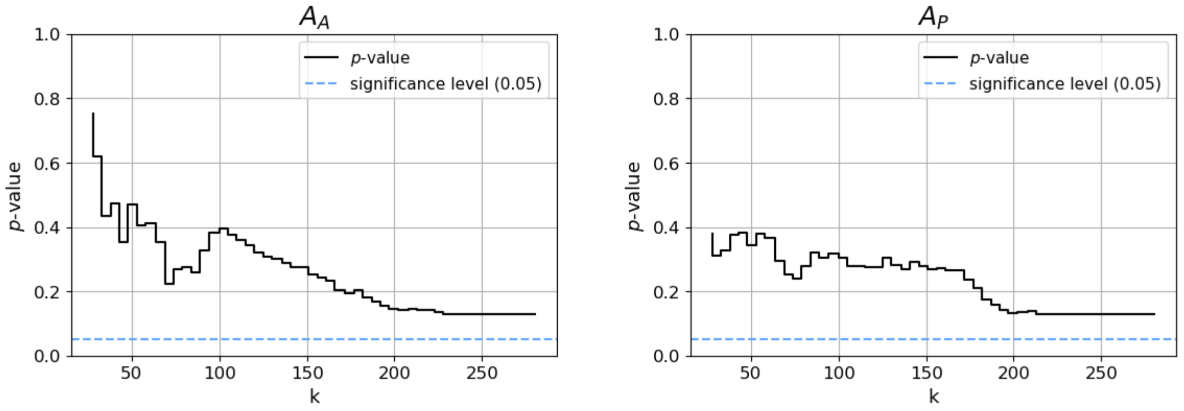

(a) p-values for A_A (average area). **(b)** p-values for A_P (peak area).

Figure 4.10: p-values for area-related event variables across different values of k.

For $k = 100$, the estimate of the constant EVI for A_A is equal to 0.592. For A_P, the estimate of the constant EVI is 0.594.

4.2.2. Duration

In Figure 4.11 the p-values from applying the test to D (duration) are shown for various values of k. As with the other variables, H_0 cannot be rejected at the 0.05 significance level. For $k = 100$, the constant EVI is estimated at 0.314.

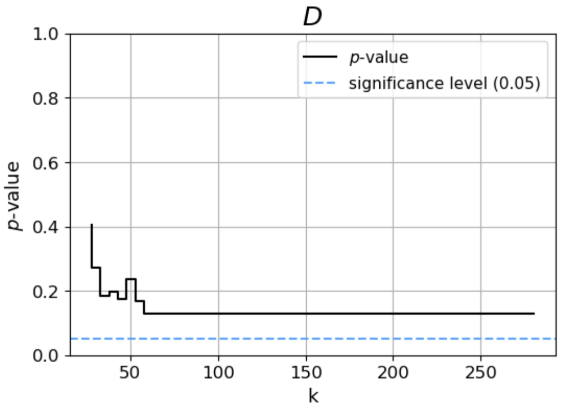

Figure 4.11: p-values for D (duration) across different values of k.

4.2.3. Intensity

The p-values for intensity-related variables (I_{A}, I_{P}, and I_{PA}) are presented in Figures 4.12a, 4.12b and 4.12c, respectively. Across all three variables, H_0 is not rejected at the 0.05 significance level.

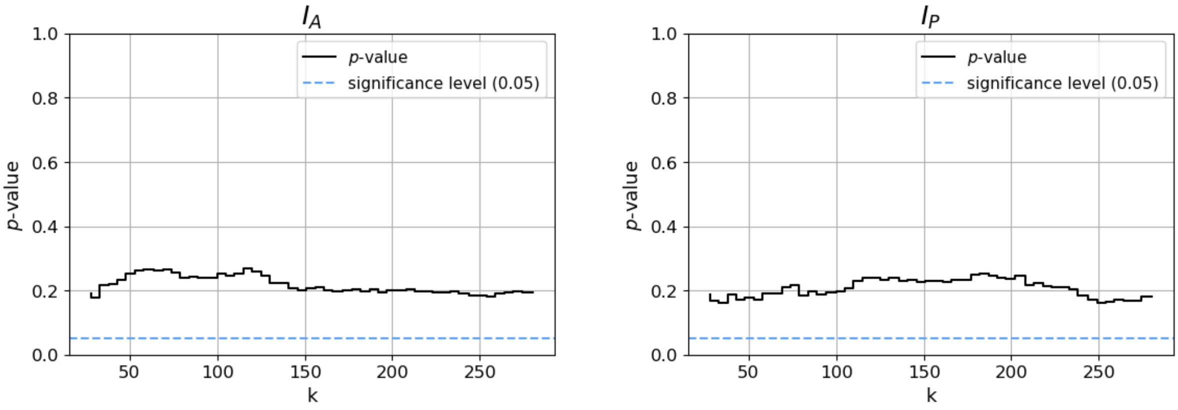

(a) p-values for I_{A} (average intensity).

(b) p-values for I_{P} (peak intensity).

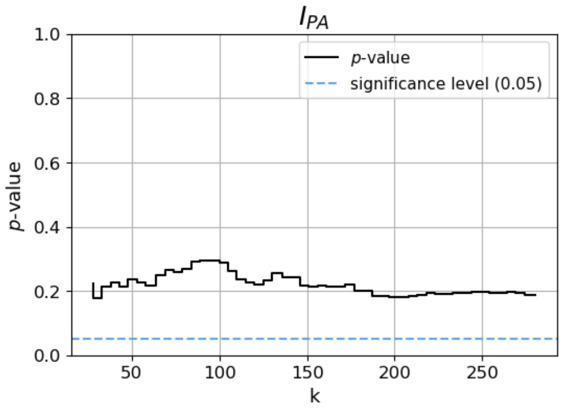

(c) p-values for I_{PA} (peak average intensity).

Figure 4.12: p-values for intensity-related event variables across different values of k.

The constant EVI for I_{A} , for $k = 100$, is estimated at 0.224. For I_{P}, the estimate of the constant EVI equals 0.327 and for I_{PA} this is 0.289.

4.3. Relationships Between Area, Duration, and Intensity of Events

4.3.1. Correlations and Pair Plots

To explore the relationships between rainfall event variables, pairwise Pearson correlation coefficients are visualised in Figure 4.13. The analysis reveals strong correlations among variables within the same category, namely between intensity-related variables (e.g., I_{A} and I_{PA}) and area-related variables (A_{A} and A_{P}). Additionally, notable correlations are observed between event duration (D) and the area variables (A_{A} and A_{P}), as well as between D and the intensity variables. The dewpoint temperature (T_d) also shows moderate correlations with the intensity variables. Outside of these relationships, most other correlations are weak but positive, and no negative correlations are found in the dataset.

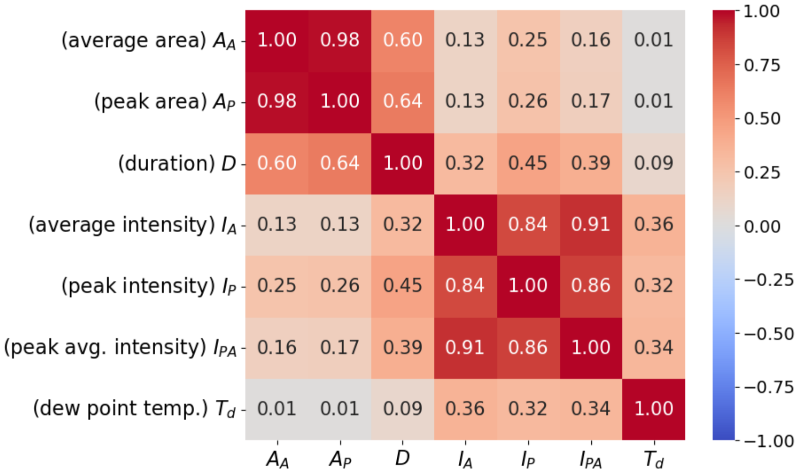

Figure 4.13: Pearson correlation matrix showing the relationships between rainfall event variables, including area, duration, intensity and dewpoint temperature.

Figure 4.14 provides pair plots of the data, where the lower diagonal displays scatterplots of variable pairs, and the upper diagonal presents the corresponding Pearson correlation coefficients. Cumulative distributions for each variable are plotted along the diagonal. These pair plots further illustrate the strong intra-category relationships and highlight the significant correlations between D and A_{A}, as well as between D and I_{PA}. Additionally, the moderate correlations between T_d and intensity variables are visually evident.

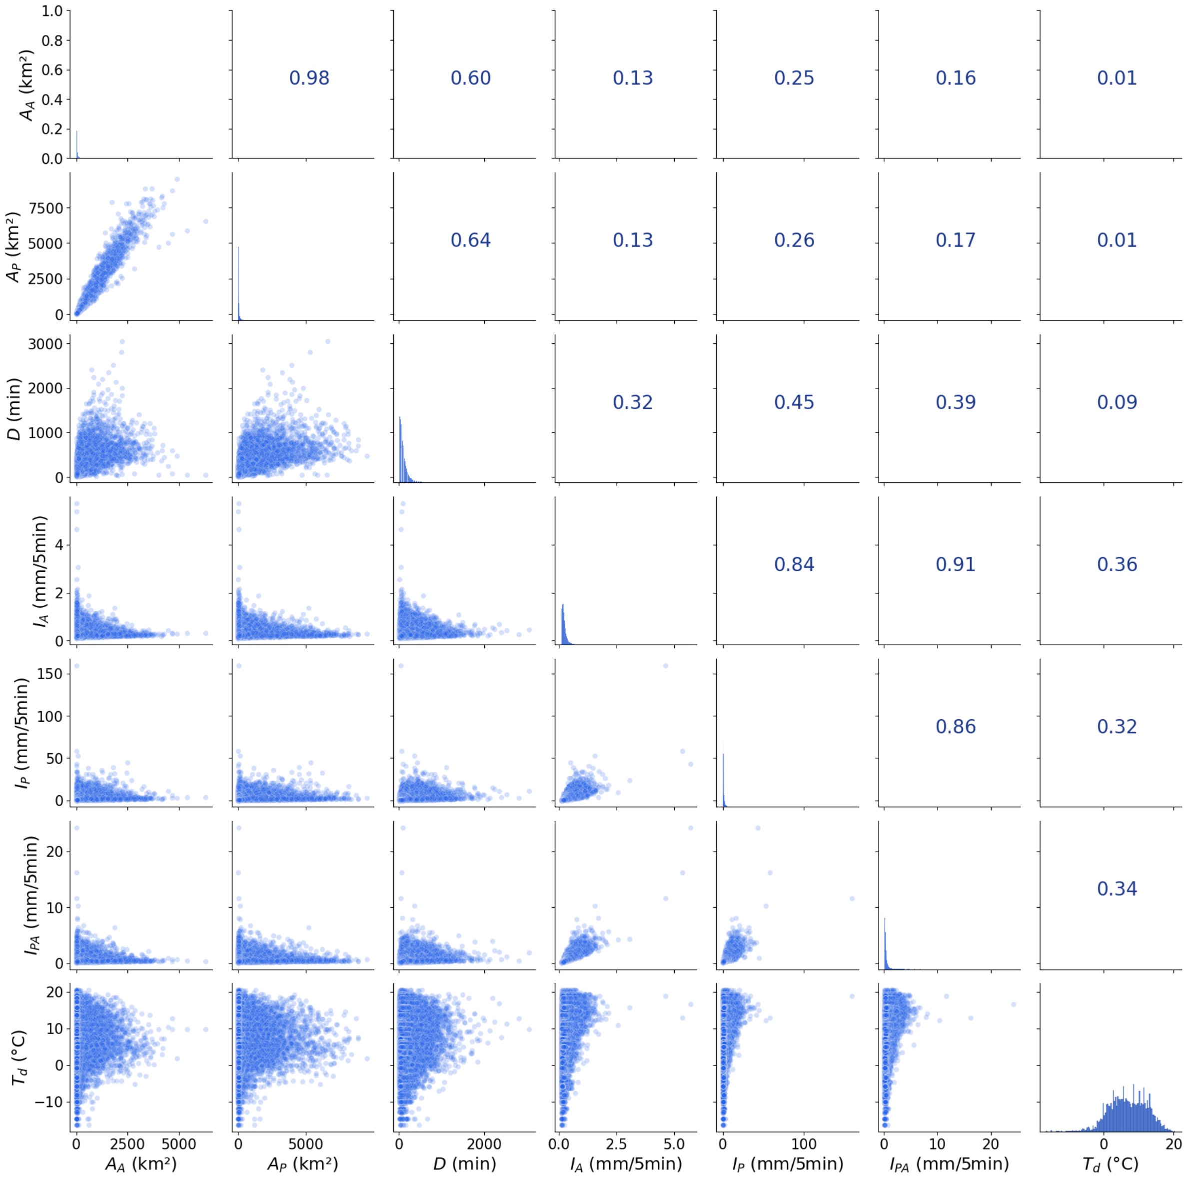

Figure 4.14: Pair plots and correlation coefficients of the rainfall event variables.

To better understand dependencies between variables while minimizing the influence of outliers, we rank the data and focus on four key variables: duration (D), peak average intensity (I_{PA}), average area (A_{A}) and the dewpoint temperature (T_d). Ranking the data emphasises relative relationships rather than absolute magnitudes, improving robustness against anomalies. Figure 4.15 presents these ranked pair plots alongside Spearman's rank correlation coefficients, providing a clearer view of the interactions between the selected variables. This analysis reveals that the relationship between T_d and I_{PA} is particularly pronounced at the higher ends of both variables, which indicates that high dew point temperatures are associated with high peak average intensities.

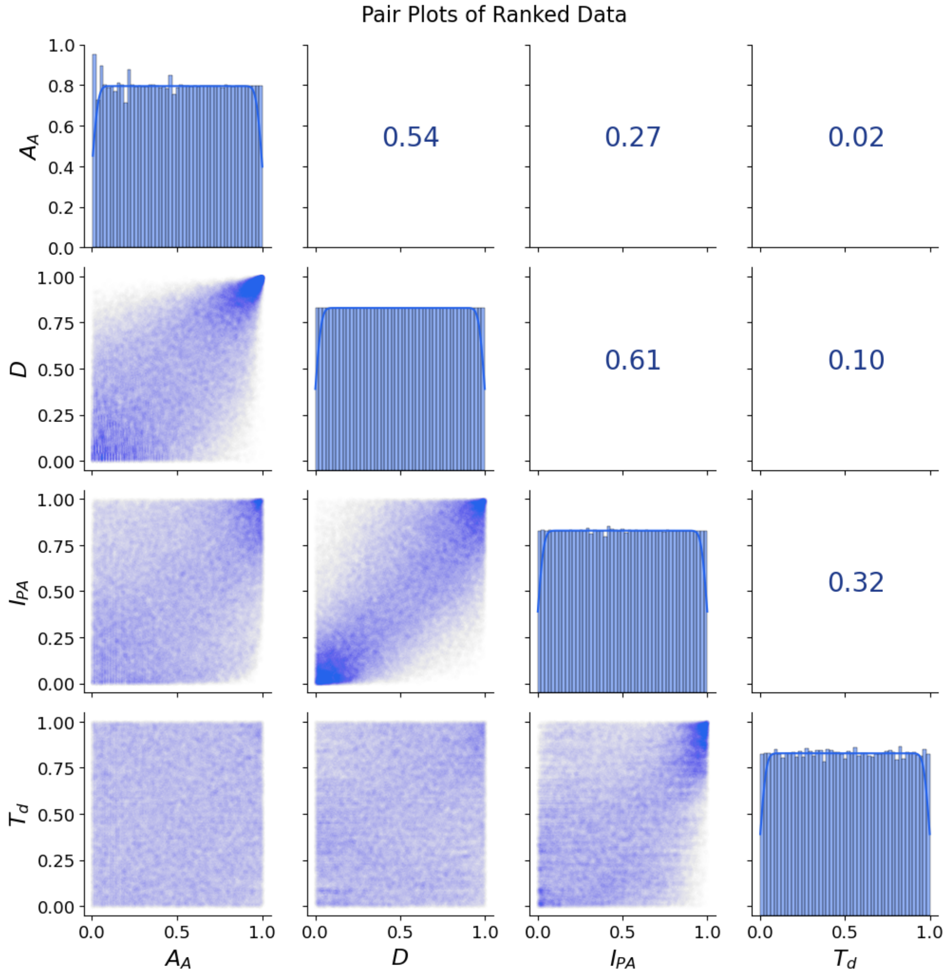

Figure 4.15: Pair plots and Spearman's correlation coefficients for ranked rainfall event variables.

4.3.2. Binned Analysis of Intensity Given Duration and Area

To analyse trends in intensity variables across different event durations and spatial extents, the data is first binned using quantile binning. The bin edges and the number of events in each bin are displayed in Figure 4.16. The variability in event counts across the bins is evident, with certain bins containing more events than others.

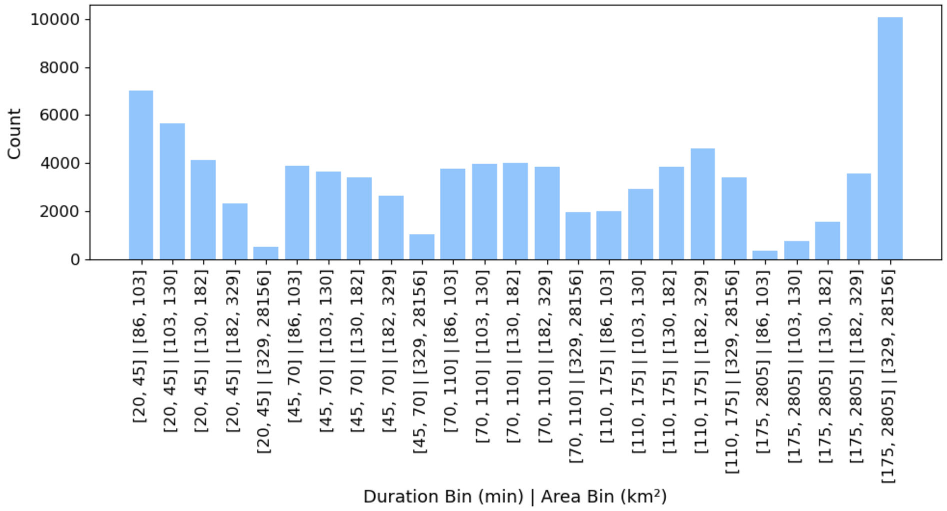

Figure 4.16: Bin counts of the duration-area bins. The edges of each bin are mentioned on the horizontal axis, where the first interval corresponds to the duration (min) and the second interval denotes the area (km^2).

Trends within each bin are then evaluated. Table A.1 in Appendix A.3 summarises the results of the Mann-Kendall test for the three intensity variables across these bins.

The analysis reveals that for shorter-duration bins (i.e., durations up to 105 minutes), the majority of cases exhibit statistically significant decreasing trends in intensity variables. In contrast, for the longest two duration bins (i.e., durations greater than 105 minutes), few cases show statistically significant trends according to the Mann-Kendall test. Notably, no statistically significant increasing trends are observed in any of the bins.

To further examine the behaviour of intensity variables within each bin, the method by De Haan and Zhou (discussed in Section 3.2) was applied. This method assesses the behaviour of the EVI of the extreme value distribution for the intensity variables. Across all bins and intensity variables, the null hypothesis (H_0) could not be rejected at the 0.05 significance level. This indicates that no significant changes were detected in the EVI of intensity variables when analysed within the bins.

4.3.3. Changes in Correlation

This section presents the results of the change point test procedure described in Section 3.3.3. We first analyse the correlation between the duration (X_t) and the average intensity (Z_t) in detail before summarizing the results for other variable pairs. The corresponding figures for these additional results can be found in Section A.4 of the Appendix.

For this analysis, a bandwidth of $k = \lfloor n^{5/8} \rfloor = 1431$ is used, which is in line with the framework of [32]. Additionally, the procedure was done using bandwidths of $k = \lfloor n^{4/7} \rfloor = 768$ and $k = \lfloor n^{2/3} \rfloor = 2324$, which did not significantly change the results. The results of the procedure for these values of k can be found in Appendix A.4.2.

Figure 4.17 displays the local estimate of the correlation function $\hat{\rho}_{t,n}$ over time. We observe noticeable variation in the correlation, suggesting potential structural changes.

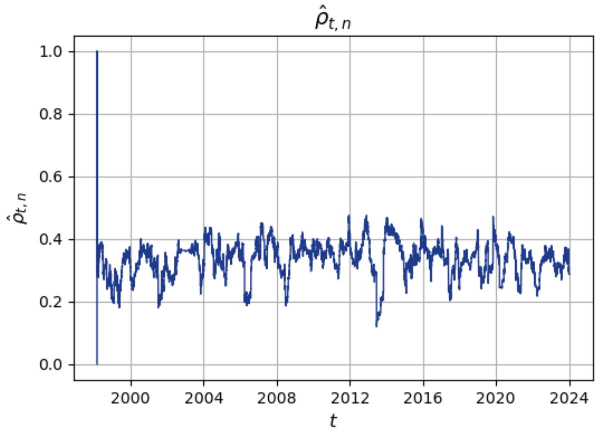

Figure 4.17: Local estimate of the correlation function $\hat{\rho}_{t,n}$ between event duration (D) and average intensity (I_A) over time.

Next, we estimate the integrated correlation function $F_n(u)$ using $M_n(u)$, with the results shown in Figure 4.18 (left panel). The general increase in values indicates that, on average, the correlation is positive. The corresponding CUSUM process $T_n(u)$ is visualised in Figure 4.18 (right panel). The lower values of $T_n(u)$ for lower values of u indicate that in approximately the first half of the study period the correlation was lower than towards the end of the study period.

Critical values for the test statistic are obtained via the bootstrapping procedure as outlined in Section 3.3.3, using a total of 10^4 bootstrap samples. The critical thresholds are determined as $t_{0.05} = \pm 2.34$ and $t_{0.10} = \pm 2.15$. Since the CUSUM process crosses the critical boundary at the 5% significance level, we reject the null hypothesis of constant correlation between D and I_A. This suggests a significant shift in the correlation structure over time.

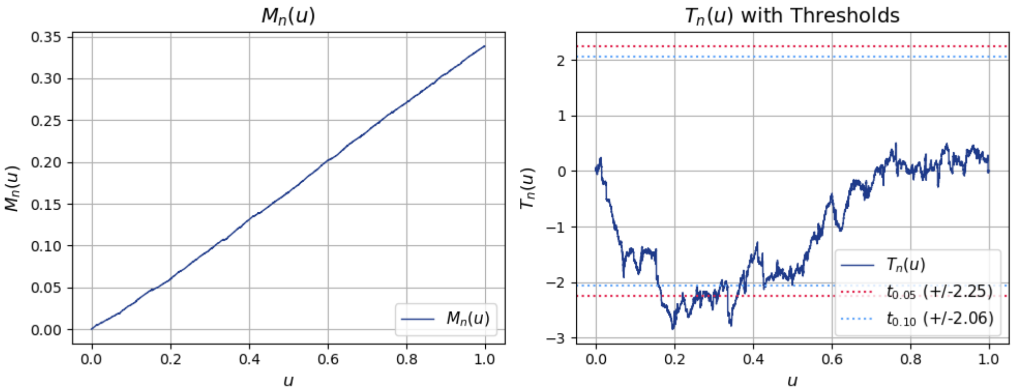

Figure 4.18: The estimator $M_n(u)$ of the integrated correlation between event duration (D) and average intensity (I_A) (left), and the corresponding CUSUM process $T_n(u)$ (right), with critical thresholds at the 5% and 10% significance levels.

The density function of the bootstrapped $\hat{T}_n^*$, along with the critical value $t_{0.05}$, is presented in Figure 4.19.

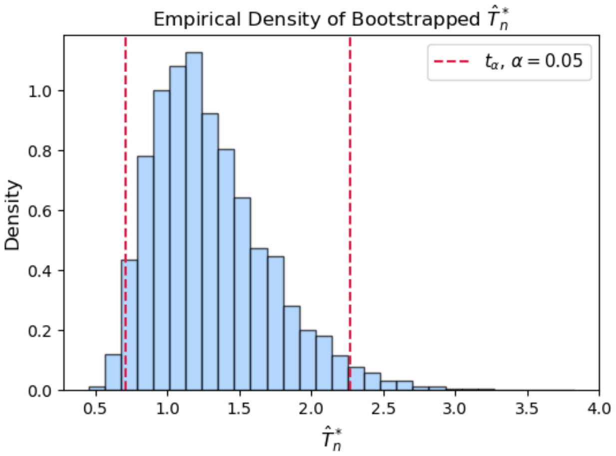

Figure 4.19: Density of the bootstrapped $\hat{T}_n^*$, with critical values at the 5% significance level.

We apply the same procedure to all pairs of rainfall event variables. The resulting p-values are summarised in Table 4.12. Significant p-values at the 10% level are marked with * and significance at the 5% level is marked with **.

Table 4.12: p-values from the change point test procedure ($k = \lfloor n^{5/8} \rfloor = 1431$) for pairs of rainfall event variables. *: Significant at 10% level. **: Significant at 5% level.

$k = \lfloor n^{5/8} \rfloor = 1431$		D	A_A	A_P	I_A	I_P	I_PA
(duration)	D		<0.001**	0.023**	0.002**	0.73	0.80
(avg. area)	A_A				0.026*	0.19	0.18
(peak area)	A_P				0.039*	0.34	0.40
(avg. intensity)	I_A						
(peak intensity)	I_P						
(peak avg. intensity)	I_PA						

5

Conclusion and Discussion

This study utilises over two decades (1998–2023) of radar data of the Netherlands to analyse rainfall events as spatiotemporal objects, a novel approach to studying rainfall characteristics. Using the Celltrack algorithm, we characterise rainfall events in terms of attributes such as area, duration, and intensity. This life-cycle-based approach offers a different perspective than point-based studies by treating rainfall events as dynamic objects. This process addresses the first research question by demonstrating how rainfall events can be mathematically described.

5.1. Key Findings

Trends in Rainfall Event Characteristics

To address the second research question on what trends in the event variables can be identified, trends in the spatial extent, duration, and intensity of rainfall events were evaluated using the Mann-Kendall test and Sen's slope estimator. The results reveal clear trends:

- **Area**: Significant increases in the spatial extent of rainfall events are observed in the monthly means of both average and maximum areas, suggesting that over time, rainfall events have been covering larger regions and/or events with larger regions are occurring more frequently. This trend holds across most seasons, indicating a broader scale increase of the spatial extent of events. For high-intensity events (top 10% of peak average intensity), area increases are significant only during summer, indicating an increase in areas of intense summer events.

- **Duration**: Event durations exhibit an upward trend in monthly averages as well as across all individual seasons. This reflects a shift toward more sustained rainfall events. Among the high-intensity events, significant increases in duration are observed in summer and autumn, indicating more persistent heavy rainfall in those seasons.

- **Intensity**: Intensity trends present a more nuanced picture. Shorter-duration events (up to 105 minutes) show significant decreases in intensity, while longer-duration events exhibit few significant trends. However, when focusing on smaller-scale shorter-duration higher-intensity (SSHI) events (lasting less than 300 minutes, covering fewer than 30 grid cells, and with an average intensity above 0.4 mm/5min), both average and peak average intensities display increasing trends. This suggests that high-intensity events have become more extreme over the study period.

  In contrast, rainfall events occurring under high dew point conditions ($\geq 15^\circ$C) do not show significant intensity trends. However, this does not necessarily imply that dew point temperature plays no role in intensity changes. Since our correlation analysis suggests that higher dew points are linked to more intense rainfall, it is possible that increasing dew point temperatures have contributed to intensity changes. However, when taking monthly averages, this relationship may be smoothed out, making a trend difficult to detect.

Extremal Behaviour

The application of the De Haan and Zhou method revealed no significant changes in the extreme value index (EVI) for any variable. This suggests that the shape parameter of the extreme value distribution has remained stable over the study period, although further investigation into the location and scale parameters could provide additional insights into changes in extreme rainfall characteristics.

Relationships Between Rainfall Event Variables

Correlations between event variables (area, duration and intensity) provide insights into their pairwise relationships, addressing the third research question on how the relationships between the variables can be described.

- **Duration and Area**: Strong correlations (0.60 – 0.64) between duration and area suggest that longer events typically cover larger spatial extents.
- **Duration and Intensity**: Moderate correlations (0.32 – 0.45) between duration and intensity suggest that while longer events may exhibit more intense rainfall, the relationship is not as straightforward.
- **Intensity and Dew Point Temperature**: Moderate correlations (0.32 – 0.46) between intensity and dew point temperature indicate a connection between atmospheric moisture content and rainfall intensities.
- **Intensity and Area**: Weak correlations (0.13 - 0.26) between intensity and area suggest that event size does not strongly influence intensity when considering both convective and stratiform events together. Distinguishing between these event types might provide clearer patterns, as convective events tend to be short and intense, while stratiform events are longer and less intense.

To address the final research question on whether pairwise relationships between event variables have changed over time, a statistical test for changes in the correlation function is applied. The results provide significant evidence for changes in correlation between the duration and the area of rainfall events, where the correlation towards the end of the period is larger than in the beginning of the period. Similarly, correlations between average intensity and both duration and area have also increased, indicating that events with high average intensity are more frequently associated with longer and larger events. In contrast, peak intensity and peak average intensity showed no significant changes in correlation with other variables.

5.2. Discussion

While the results of this study provide informative answers to the research questions, there are some aspects that can be improved upon and would result in more in-depth answers to these questions.

Implications of a Life-cycle-Based Approach

The life-cycle-based methodology of this study provides a fundamentally different view compared to point-based studies. By capturing rainfall as spatio-temporal phenomena, it enables the identification of patterns and trends that might otherwise be overlooked. However, this approach also presents challenges. For instance, definitions of metrics such as duration and intensity differ between these two approaches, which can complicate direct comparisons. While point-based studies often focus on intensity as a local measure, the life-cycle-based approach considers intensities across the spatial extent and life-cycle of events. While this approach offers a detailed perspective, it also introduces interpretative challenges that require careful contextualization.

Improving Intensity Metrics

The intensity metrics used in this study, while informative, could benefit from refinements to better capture rainfall dynamics. For example, the average intensity smooths values across entire events, potentially under-representing localised high-intensity periods. Similarly, peak intensity metrics (peak intensity, peak average intensity) reflect only single time steps, neglecting sustained high-intensity episodes. Alternative metrics, such as peak average intensity within a specified period (longer than 5 minutes) or a restricted area would potentially provide a more meaningful representation. Another option would be to incorporate the life-cycle-summed precipitation. These refinements may involve greater computational complexity but would enhance the understanding of rainfall dynamics.

Distinguishing Convective and Stratiform Events

The combination of convective and stratiform events in the analysis may mask distinct trends due to their differing characteristics. Convective events, typically associated with intense, short-duration rainfall, and stratiform events, characterised by prolonged and less intense precipitation, may exhibit different trends and relationships. Our analysis of intensity trends suggests that some convective events may have been unintentionally filtered out, as increasing trends emerged when focusing on smaller-scale shorter-duration higher-intensity (SSHI) events separately.

To better distinguish between these event types, a more inclusive dataset – using less restrictive filtering conditions to ensure convective events are retained – should first be created. Only then would separating convective and stratiform events, for example through clustering or classification techniques, provide clearer physical insights into their respective trends.

5.3. Recommendations for Future Research

Building on the findings of this study, I propose several directions for future research:

- **Refinement of Metrics**: Incorporate metrics such as total precipitation sum or modified intensity metrics to better capture rainfall characteristics.
- **Extreme Event Analysis**: Investigate trends in the location and scale parameters of extreme value distributions to better understand changes in extremes.
- **Multivariate Analysis**: Employ techniques such as copulas, multivariate extreme value theory, or machine learning to analyse the joint behaviour of and dependencies among rainfall variables.
- **Convective vs. Stratiform Events**: First, develop a more inclusive dataset with less restrictive filtering to ensure smaller-scale shorter-duration higher-intensity (SSHI) events are retained. Then, try separating convective and stratiform events to uncover distinct trends and improve the clarity of relationships among event characteristics.
- **Geographical Expansion**: Extending the study to other regions or climates could help identify broader patterns and improve understanding of rainfall dynamics under varying conditions.

References

[1] M. Allaire. “Socio-economic impacts of flooding: A review of the empirical literature”. In: *Water Security* 3 (2018), pp. 18–26. DOI: `10.1016/j.wasec.2018.09.002`.

[2] N. Ansh Srivastava and G. Mascaro. “Improving the utility of weather radar for the spatial frequency analysis of extreme precipitation”. In: *Journal of Hydrology* 624 (2023), p. 129902. DOI: `10.1016/j.jhydrol.2023.129902`.

[3] A. A. Balkema and L. de Haan. “Residual Life Time at Great Age”. In: *The Annals of Probability* 2.5 (Oct. 1974). Publisher: Institute of Mathematical Statistics, pp. 792–804. DOI: `10.1214/aop/1176996548`.

[4] Jules Beersma et al. *Neerslagstatistiek en-reeksen voor het waterbeheer 2019*. Vol. 19. STOWA, 2019.

[5] A.J. Bowden, C. Jakob, and J. Soderholm. “Identification of Rainfall Events and Heavy Rainfall Events From Radar Measurements in Southeastern Australia”. In: *Journal of Geophysical Research: Atmospheres* 129.2 (2024). DOI: `10.1029/2023JD039253`.

[6] Theo Brandsma. *Comparison of automatic and manual precipitation networks in the Netherlands*. KNMI, 2014.

[7] T.A. Buishand et al. “Homogeneity of precipitation series in the Netherlands and their trends in the past century”. In: *International Journal of Climatology* 33.4 (2013), pp. 815–833. DOI: `10.1002/joc.3471`.

[8] P. Chu et al. “Extreme Rainfall Events in the Hawaiian Islands”. In: (2009). DOI: `10.1175/2008JAMC1829.1`.

[9] A.J. Clark et al. “Application of Object-Based Time-Domain Diagnostics for Tracking Precipitation Systems in Convection-Allowing Models”. In: (2014). DOI: `10.1175/WAF-D-13-00098.1`.

[10] S. Coles. *An Introduction to Statistical Modeling of Extreme Values*. Springer Series in Statistics. London: Springer, 2001. DOI: `10.1007/978-1-4471-3675-0`. URL: `http://link.springer.com/10.1007/978-1-4471-3675-0` (visited on 11/05/2024).

[11] E.E. Daniels et al. “Spatial precipitation patterns and trends in The Netherlands during 1951–2009”. In: *International Journal of Climatology* 34.6 (2014), pp. 1773–1784. DOI: `10.1002/joc.3800`.

[12] C. Davis, B. Brown, and R. Bullock. “Object-Based Verification of Precipitation Forecasts. Part I: Methodology and Application to Mesoscale Rain Areas”. In: *Monthly Weather Review* (2006). DOI: `10.1175/MWR3145.1`.

[13] E.E. Ebert and J.L. McBride. “Verification of precipitation in weather systems: determination of systematic errors”. In: *Journal of Hydrology* 239.1-4 (2000), pp. 179–202. DOI: `10.1016/S0022-1694(00)00343-7`.

[14] J.M. Eden et al. “Extreme precipitation in the Netherlands: An event attribution case study”. In: *Weather and Climate Extremes* 21 (2018), pp. 90–101. DOI: `10.1016/j.wace.2018.07.003`.

[15] S. Fadhel, M.A. Rico-Ramirez, and D. Han. “Sensitivity of peak flow to the change of rainfall temporal pattern due to warmer climate”. In: *Journal of Hydrology* 560 (2018), pp. 546–559. DOI: `10.1016/j.jhydrol.2018.03.041`.

[16] Z. Feng et al. “A Global High-Resolution Mesoscale Convective System Database Using Satellite-Derived Cloud Tops, Surface Precipitation, and Tracking”. In: *Journal of Geophysical Research: Atmospheres* 126.8 (2021). DOI: `10.1029/2020JD034202`.

[17] E.M. Fischer and R. Knutti. “Observed heavy precipitation increase confirms theory and early models”. In: *Nature Climate Change* 6.11 (2016), pp. 986–991. DOI: `10.1038/nclimate3110`.

[18] H.J. Fowler et al. “Anthropogenic intensification of short-duration rainfall extremes”. In: *Nature Reviews Earth & Environment* 2.2 (2021), pp. 107–122. DOI: `10.1038/s43017-020-00128-6`.

[19] F. Giorgi, F. Raffaele, and E. Coppola. “The response of precipitation characteristics to global warming from climate projections”. In: *Earth System Dynamics* 10.1 (2019), pp. 73–89. DOI: `10.5194/esd-10-73-2019`.

[20] Laurens de Haan and Chen Zhou. “Trends in Extreme Value Indices”. In: *Journal of the American Statistical Association* 116.535 (2021), pp. 1265–1279. ISSN: 0162-1459. DOI: `10.1080/01621459.2019.1705307`.

[21] L. Hayden, C. Liu, and N. Liu. “Properties of Mesoscale Convective Systems Throughout Their Lifetimes Using IMERG, GPM, WWLLN, and a Simplified Tracking Algorithm”. In: *Journal of Geophysical Research: Atmospheres* 126.20 (2021). DOI: `10.1029/2021JD035264`.

[22] X. Huang et al. “A long-term tropical mesoscale convective systems dataset based on a novel objective automatic tracking algorithm”. In: *Climate Dynamics* 51.7 (2018), pp. 3145–3159. DOI: `10.1007/s00382-018-4071-0`.

[23] M.G. Kendall. *Rank correlation methods*. Rank correlation methods. Oxford, England: Griffin, 1948.

[24] *KNMI Data Platform: Meteo data - daily quality controlled climate data KNMI, the Netherlands*. `https://dataplatform.knmi.nl/dataset/etmaalgegevensknmistations-1`.

[25] *KNMI Data Platform: Precipitation - 5 minute precipitation accumulations from climatological gauge-adjusted radar dataset for The Netherlands (2.4 km) in KNMI HDF5 format*. `https://dataplatform.knmi.nl/dataset/rad-nl21-rac-mfbs-5min-2-0`.

[26] Mark G. Lawrence. “The Relationship between Relative Humidity and the Dewpoint Temperature in Moist Air: A Simple Conversion and Applications”. In: (2005). DOI: `10.1175/BAMS-86-2-225`.

[27] K. Lochbihler. *Celltrack - A 2D cell tracking algorithm*. Version 0.8. 2020.

[28] K. Lochbihler, G. Lenderink, and A.P. Siebesma. “The spatial extent of rainfall events and its relation to precipitation scaling”. In: *Geophysical Research Letters* 44.16 (2017), pp. 8629–8636. DOI: `10.1002/2017GL074857`.

[29] T. Lukić et al. “Rainfall erosivity and extreme precipitation in the Netherlands”. In: *Időjárás* 122.4 (2018), pp. 409–432. DOI: `10.28974/idojaras.2018.4.4`.

[30] Henry B. Mann. “Nonparametric Tests Against Trend”. In: *Econometrica* 13.3 (1945). Publisher: [Wiley, Econometric Society], pp. 245–259. DOI: `10.2307/1907187`.

[31] S.J. Mason et al. “Changes in Extreme Rainfall Events in South Africa”. In: *Climatic Change* 41.2 (1999), pp. 249–257. DOI: `10.1023/A:1005450924499`.

[32] Fabian Mies. “Functional estimation and change detection for nonstationary time series”. In: *Journal of the American Statistical Association* 118.542 (2023), pp. 1011–1022.

[33] A. del Moral, T. Rigo, and M.C. Llasat. “A radar-based centroid tracking algorithm for severe weather surveillance: identifying split/merge processes in convective systems”. In: *Atmospheric Research* 213 (2018), pp. 110–120. DOI: `10.1016/j.atmosres.2018.05.030`.

[34] C. Moseley, P. Berg, and J.O. Haerter. “Probing the precipitation life cycle by iterative rain cell tracking”. In: *Journal of Geophysical Research: Atmospheres* 118.24 (2013), pp. 13, 361–13, 370. DOI: `10.1002/2013JD020868`.

[35] C. Muñoz, L. Wang, and P. Willems. “Enhanced object-based tracking algorithm for convective rain storms and cells”. In: *Atmospheric Research* 201 (2018), pp. 144–158. DOI: `10.1016/j.atmosres.2017.10.027`.

[36] *NASA GISS: Panoply 5 netCDF, HDF and GRIB Data Viewer*. National Aeronautics and Space Administration. URL: `https://www.giss.nasa.gov/tools/panoply/`.

[37] A. Overeem, T.A. Buishand, and I. Holleman. “Extreme rainfall analysis and estimation of depth-duration-frequency curves using weather radar”. In: *Water Resources Research* 45.10 (2009). DOI: `10.1029/2009WR007869`.

[38] L. Panziera et al. “A radar-based regional extreme rainfall analysis to derive the thresholds for a novel automatic alert system in Switzerland”. In: *Hydrology and Earth System Sciences* 20.6 (2016), pp. 2317–2332. DOI: `10.5194/hess-20-2317-2016`.

[39] T. Rosin, F. Marra, and E. Morin. “Exploring patterns in precipitation intensity–duration–area–frequency relationships using weather radar data”. In: *Hydrology and Earth System Sciences* 28.15 (2024), pp. 3549–3566. DOI: `10.5194/hess-28-3549-2024`.

[40] T. Semmler and D. Jacob. “Modeling extreme precipitation events—a climate change simulation for Europe”. In: *Global and Planetary Change*. Extreme climatic events 44.1 (2004), pp. 119–127. DOI: `10.1016/j.gloplacha.2004.06.008`.

[41] Pranab Kumar Sen. “Estimates of the regression coefficient based on Kendall’s tau”. In: *Journal of the American statistical association* 63.324 (1968), pp. 1379–1389.

[42] G. Skok et al. “Object-Based Analysis of Satellite-Derived Precipitation Systems over the Low and Midlatitude Pacific Ocean”. In: *Monthly Weather Review* 137 (2009), pp. 3196–3218. DOI: `10.1175/2009MWR2900.1`.

[43] H.T.C. van Stokkom, A.J.M. Smits, and R.S.E.W. Leuven. “Flood Defense in The Netherlands”. In: *Water International* 30.1 (2005). Publisher: Routledge, pp. 76–87. DOI: `10.1080/02508060508691839`.

[44] Henri Theil. “A rank-invariant method of linear and polynomial regression analysis”. In: *Indagationes mathematicae* 12.85 (1950), p. 173.

[45] Kevin E Trenberth et al. “The changing character of precipitation”. In: *Bulletin of the American Meteorological Society* 84.9 (2003), pp. 1205–1218.

[46] Y. Zheng et al. “Spatial characteristics of extreme rainfall over China with hourly through 24-hour accumulation periods based on national-level hourly rain gauge data”. In: *Advances in Atmospheric Sciences* 33.11 (2016), pp. 1218–1232. DOI: `10.1007/s00376-016-6128-5`.

A

Figure and Table Supplement

In this appendix, we provide additional figures and tables that supplement the main text. These items offer further details on the results presented in the core of the report. Each figure and table is described below with a brief explanation of its relevance to the overall findings of the study.

A.1. Event Data

Figure A.1 presents histograms of the key event variables, illustrating the distribution of values within the dataset. Each histogram is constructed with 50 bins. As shown, all distributions are highly skewed, with many smaller areas, shorter durations, and lower intensities in the dataset. This pattern is typical for rainfall event data, where extreme values (e.g., very large areas, long durations, or intense events) are much less frequent. Because of this skewness, we also investigated the log-transformed versions of these variables in the main text. However, for completeness, we include the histograms of the untransformed variables here.

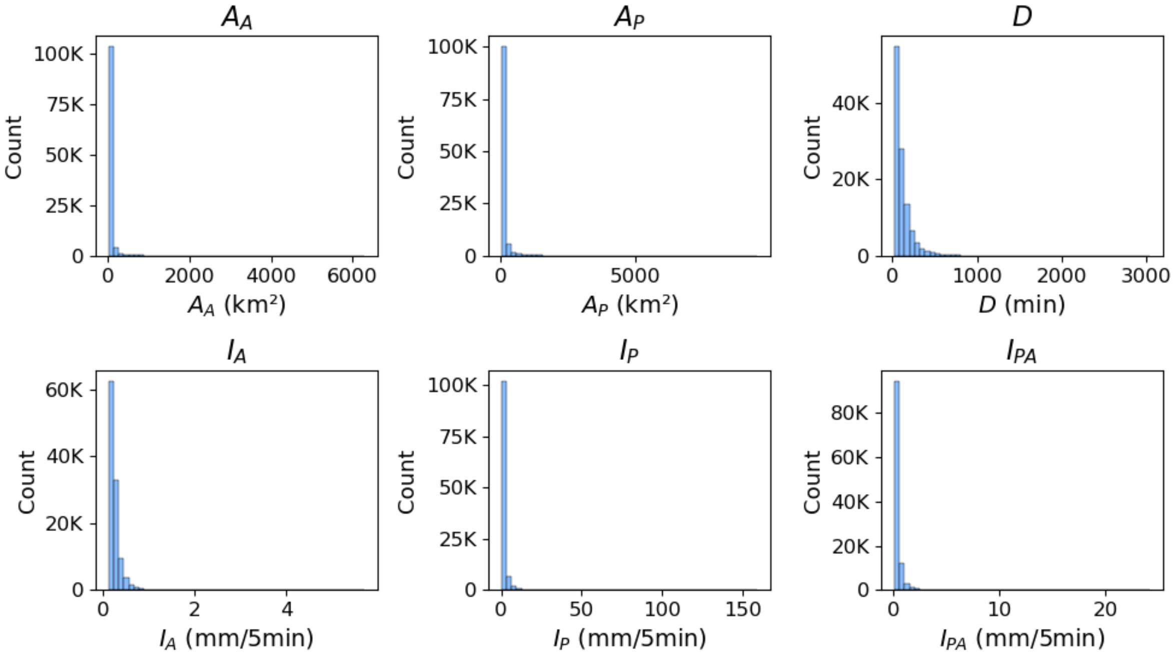

Figure A.1: Histograms showing the distributions of event variables. Each histogram is plotted using 50 bins.

A.2. Trends in Event Variables

In this section, we provide additional figures that detail the results of the Mann-Kendall test and the Theil-Sen slope estimator applied to the event variables. The results are given for different subsets of the data, including seasonal categories and subsets focusing on high-intensity events. While the main text highlights the most important trends, this appendix includes all figures, providing a comprehensive overview of trends across all subsets of the data, including those that did not show significant results.

A.2.1. Area

This section presents time series analyses of the average area ($\boldsymbol{A}_{\mathbf{A}}$) and peak area ($\boldsymbol{A}_{\mathbf{P}}$) of rainfall events, which were assessed across different seasons as well as for high-intensity (top 10% I_{PA}) events. Each figure shows the results of the Theil-Sen slope estimator with the 95% confidence interval around the trend.

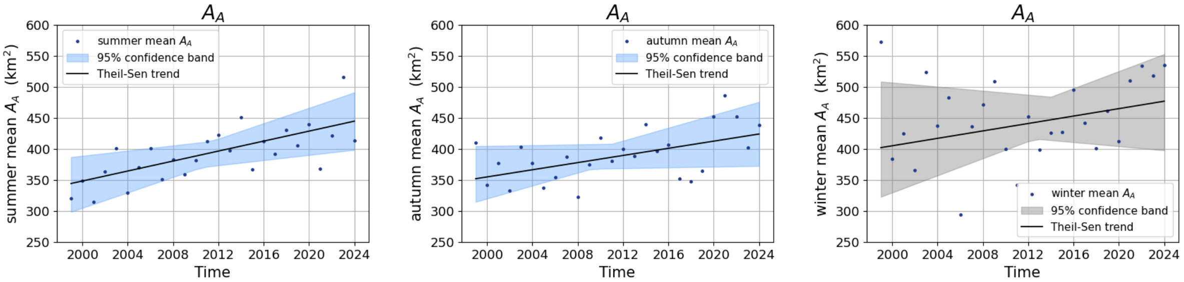

(a) Summer means*

(b) Autumn means*

(c) Winter means

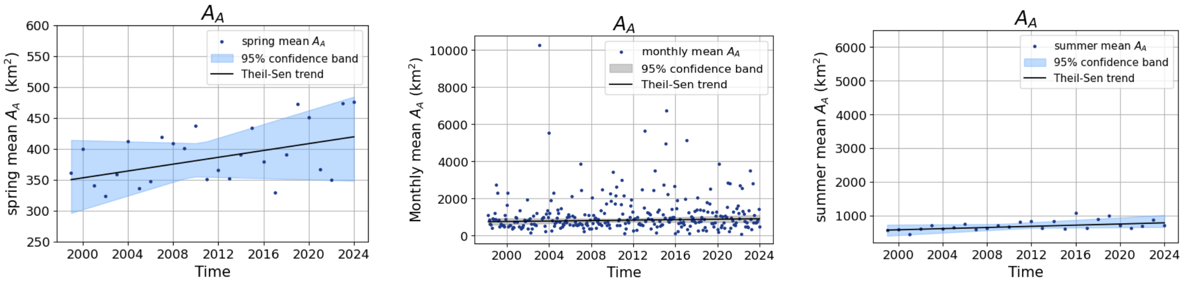

(d) Spring means*

(e) Monthly means for high-intensity (top 10% I_{PA}) events

(f) Summer means for high-intensity (top 10% I_{PA}) events*

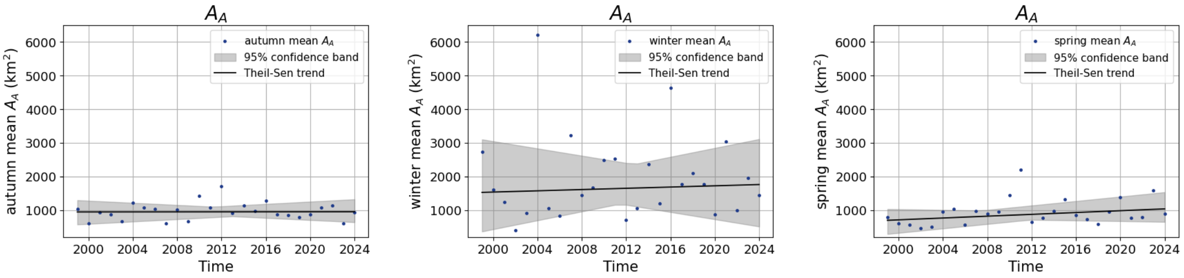

(g) Autumn means for high-intensity (top 10% I_{PA}) events

(h) Winter means for high-intensity (top 10 % I_{PA}) events

(i) Spring means for high-intensity (top 10 % I_{PA}) events

Figure A.2: Time series of (seasonal) subsets of $\boldsymbol{A}_{\mathbf{A}}$ (average area) of rainfall events over the study period. The solid line represents the trend estimated using the Theil-Sen estimated method, and the shaded area indicates the 95% confidence interval around the trend.

* : Significant trend at the 0.05 significance level. Confidence bands for variables where a significant trend is detected are shown in blue (otherwise grey).

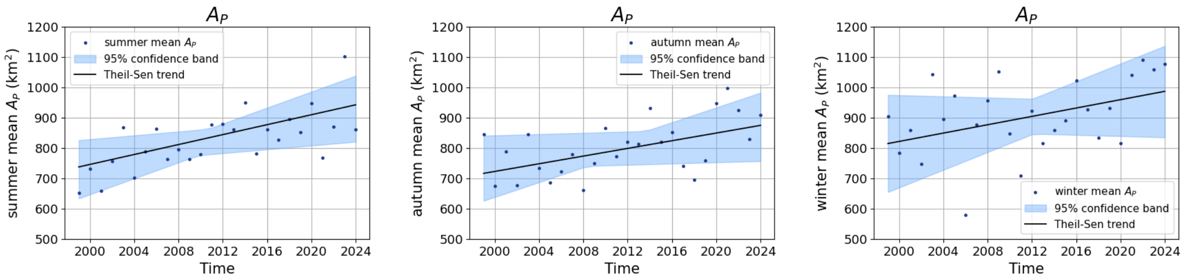

(a) Summer means*

(b) Autumn means*

(c) Winter means*

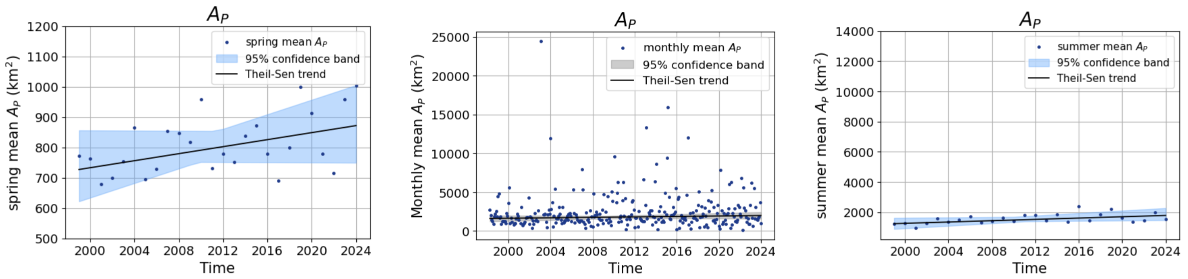

(d) Spring means*

(e) Monthly means for high-intensity (top 10% I_{PA}) events

(f) Summer means for high-intensity (top 10% I_{PA}) events*

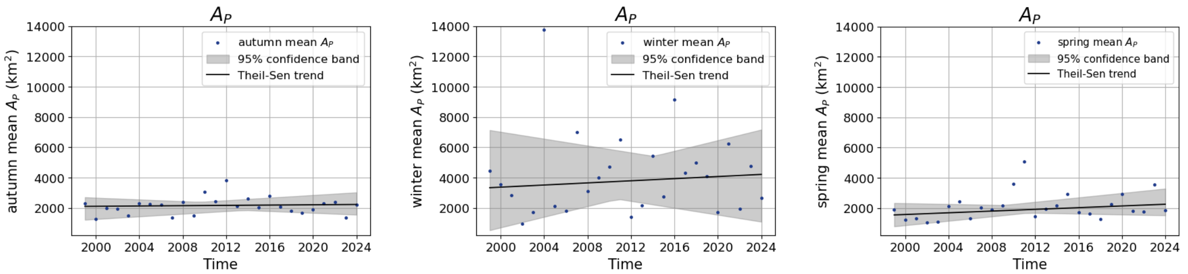

(g) Autumn means for high-intensity (top 10% I_{PA}) events

(h) Winter means for high-intensity (top 10 % I_{PA}) events

(i) Spring means for high-intensity (top 10 % I_{PA}) events

Figure A.3: Time series of (seasonal) subsets of A_{P} (peak area) of rainfall events over the study period. The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.
* : Significant trend at the 0.05 significance level. Confidence bands for variables where a significant trend is detected are shown in blue (otherwise grey).

A.2.2. Duration

In this section, we present the time series analyses of the duration (D) of rainfall events, which were assessed across different seasons as well as for high-intensity (top 10% I_{PA}) events. Each figure shows the trend estimated using the Theil-Sen method, along with the 95% confidence interval around the trend.

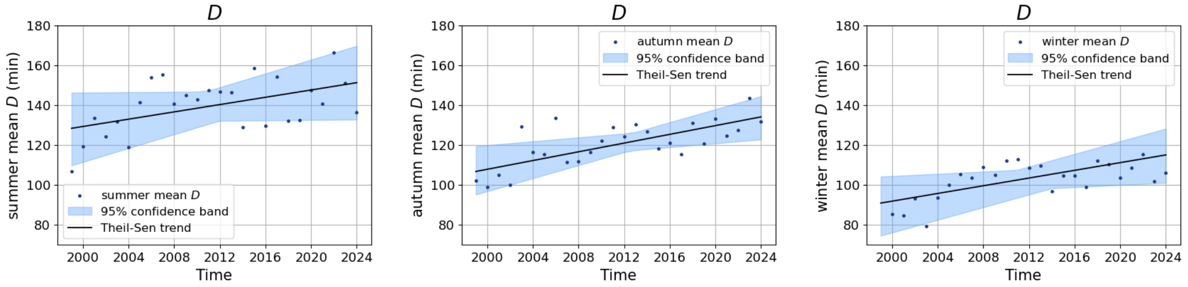

(a) Summer means*

(b) Autumn means*

(c) Winter means*

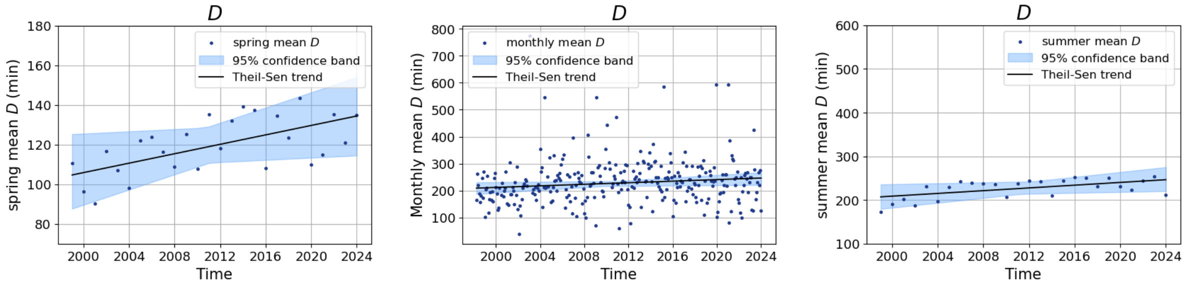

(d) Spring means*

(e) Monthly means for high-intensity (top 10% I_{PA}) events*

(f) Summer means for high-intensity (top 10% I_{PA}) events*

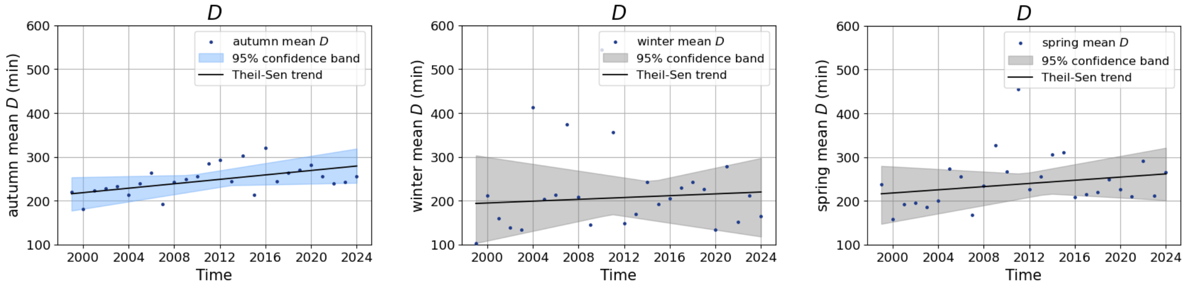

(g) Autumn means for high-intensity (top 10% I_{PA}) events*

(h) Winter means for high-intensity (top 10 % I_{PA}) events

(i) Spring means for high-intensity (top 10 % I_{PA}) events

Figure A.4: Time series of (seasonal) subsets of D (duration) of rainfall events over the study period. The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend. * : Significant trend at the 0.05 significance level. Confidence bands for variables where a significant trend is detected are shown in blue (otherwise grey).

A.2.3. Intensity

In this section, we present the time series analyses of the intensity variables of rainfall events, assessed across different seasons as well as for high-intensity (top 10% I_{PA}) events. Each figure shows the trend estimated using the Theil-Sen method, along with the 95% confidence interval around the trend.

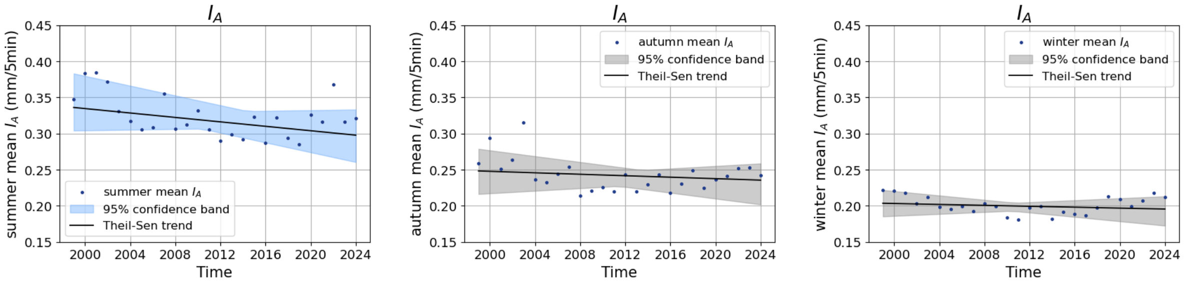

(a) Summer means* **(b)** Autumn means **(c)** Winter means

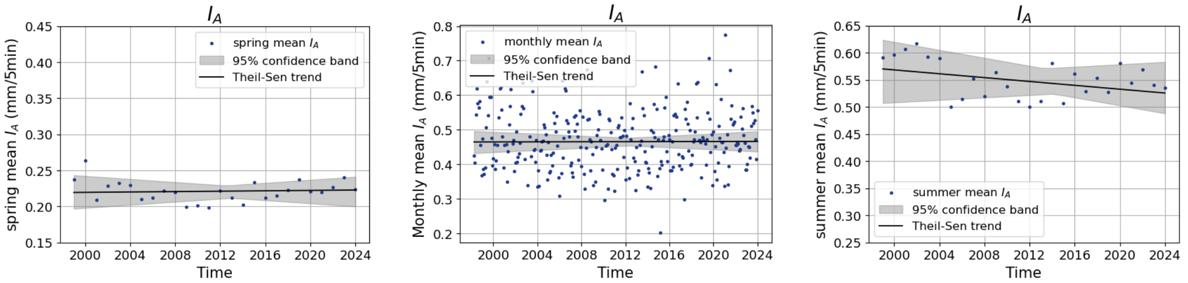

(d) Spring means **(e)** Monthly means for high-intensity (top 10% I_{PA}) events **(f)** Summer means for high-intensity (top 10% I_{PA}) events

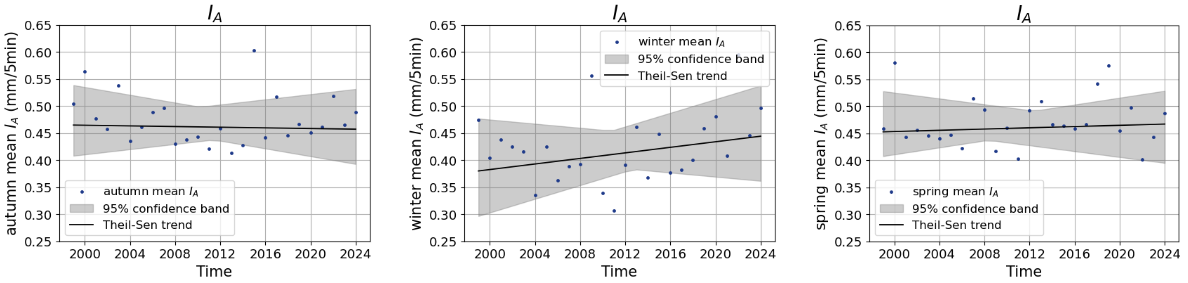

(g) Autumn means for high-intensity (top 10% I_{PA}) events **(h)** Winter means for high-intensity (top 10% I_{PA}) events **(i)** Spring means for high-intensity (top 10% I_{PA}) events

Figure A.5: Time series of (seasonal) subsets of $\boldsymbol{I_{\mathrm{A}}}$ (average intensity) of rainfall events over the study period. The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.
* : Significant trend at the 0.05 significance level. Confidence bands for variables where a significant trend is detected are shown in blue (otherwise grey).

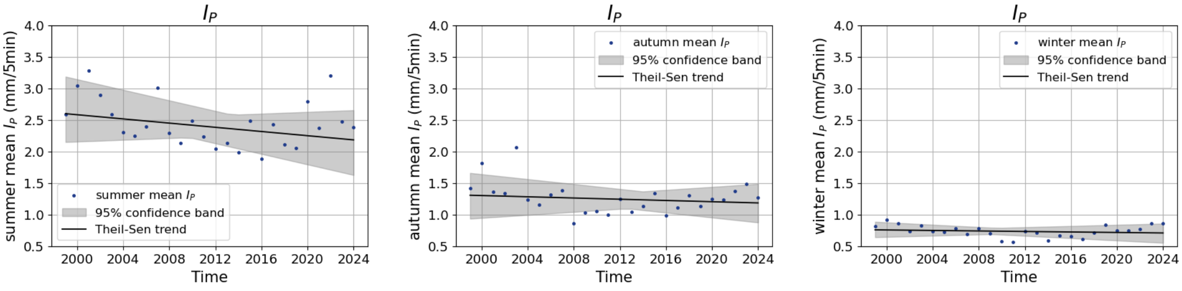

(a) Summer means

(b) Autumn means

(c) Winter means

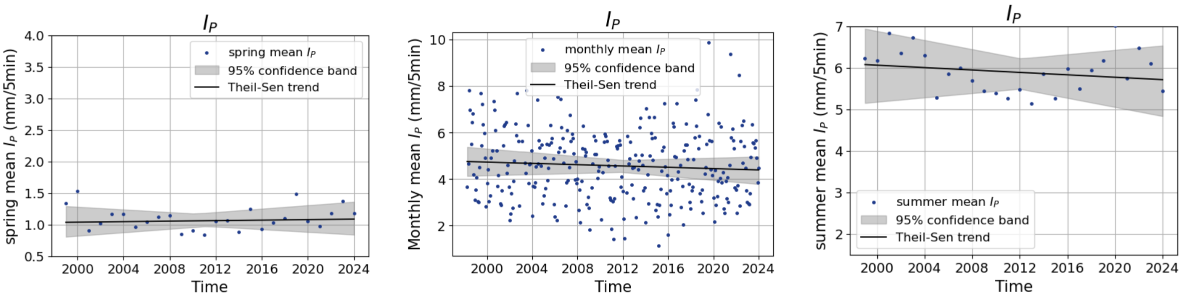

(d) Spring means

(e) Monthly means for high-intensity (top 10% I_{PA}) events

(f) Summer means for high-intensity (top 10% I_{PA}) events

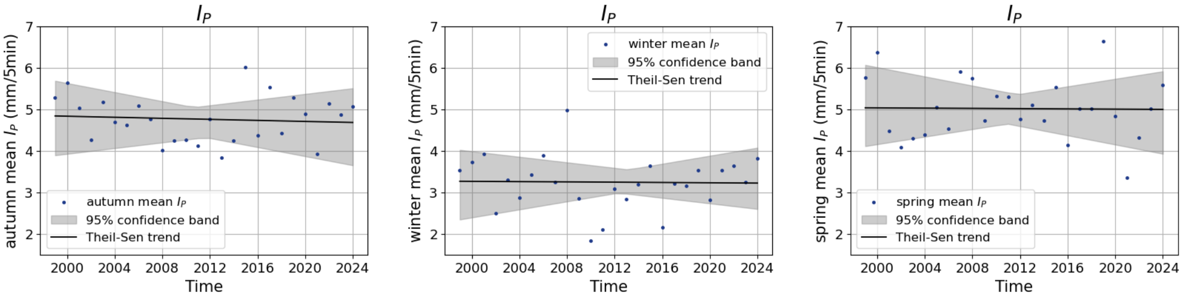

(g) Autumn means for high-intensity (top 10% I_{PA}) events

(h) Winter means for high-intensity (top 10% I_{PA}) events

(i) Spring means for high-intensity (top 10% I_{PA}) events

Figure A.6: Time series of (seasonal) subsets of $\boldsymbol{I_{\mathrm{P}}}$ (peak intensity) of rainfall events over the study period. The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.
* : Significant trend at the 0.05 significance level. Confidence bands for variables where a significant trend is detected are shown in blue (otherwise grey).

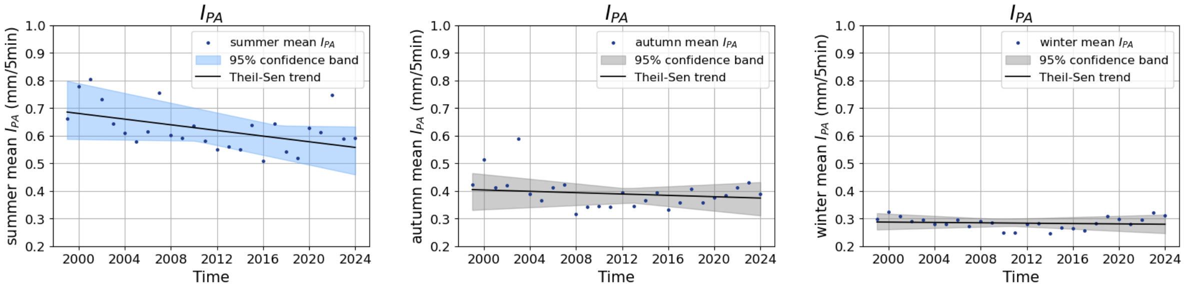

(a) Summer means*

(b) Autumn means

(c) Winter means

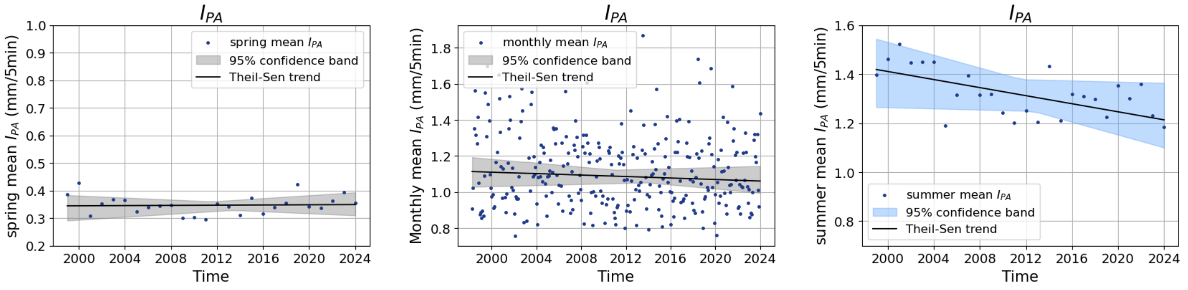

(d) Spring means

(e) Monthly means for high-intensity (top 10% I_{PA}) events

(f) Summer means for high-intensity (top 10% I_{PA}) events*

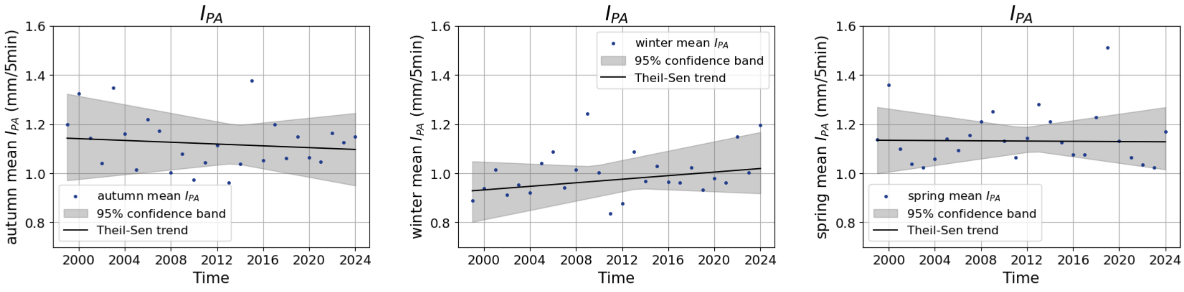

(g) Autumn means for high-intensity (top 10% I_{PA}) events

(h) Winter means for high-intensity (top 10% I_{PA}) events

(i) Spring means for high-intensity (top 10% I_{PA}) events

Figure A.7: Time series of (seasonal) subsets of $\boldsymbol{I_{\mathrm{PA}}}$ (peak average intensity) of rainfall events over the study period. The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.
* : Significant trend at the 0.05 significance level. Confidence bands for variables where a significant trend is detected are shown in blue (otherwise grey).

A.2.4. Trends for Small, Short, Intense Events

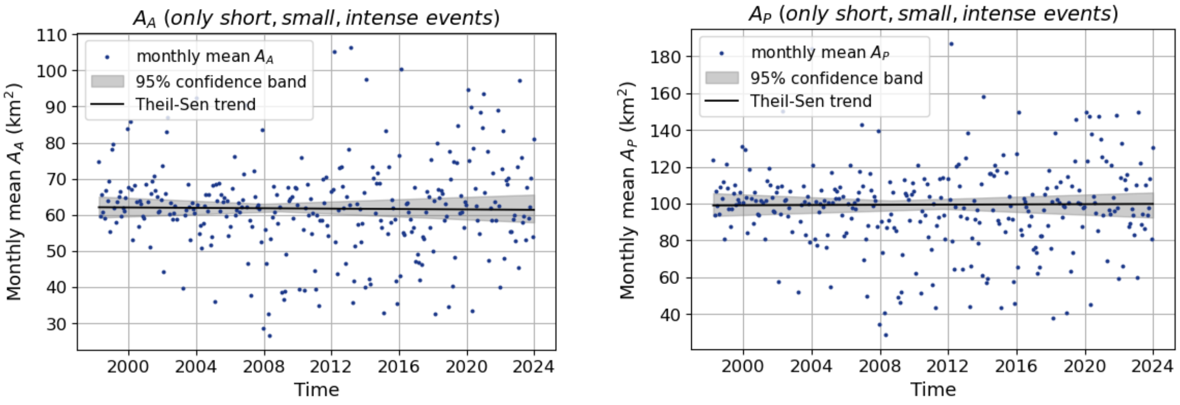

(a) monthly means of A_A (average area) for short, intense events.

(b) monthly means of A_P (peak area) for short, intense events.

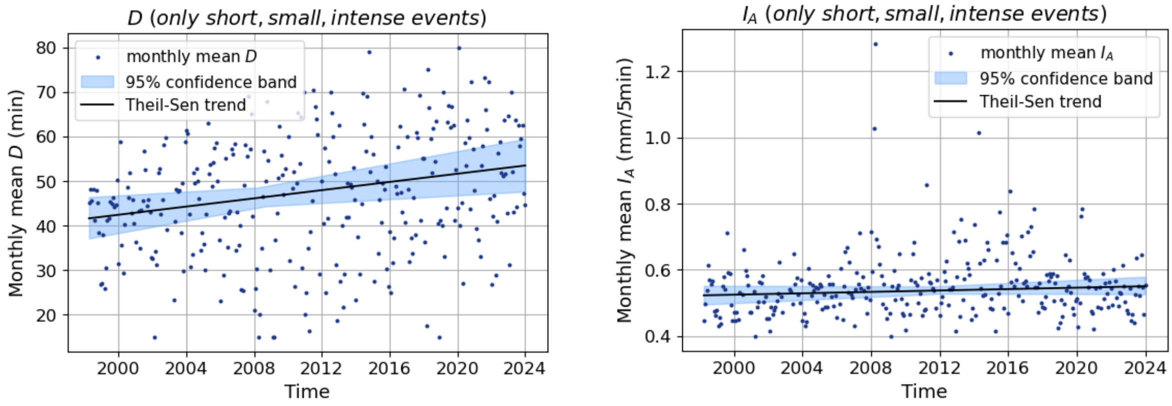

(c) monthly means of D (duration) for short, intense events.

(d) Monthly means of I_A (average intensity) for short, intense events.

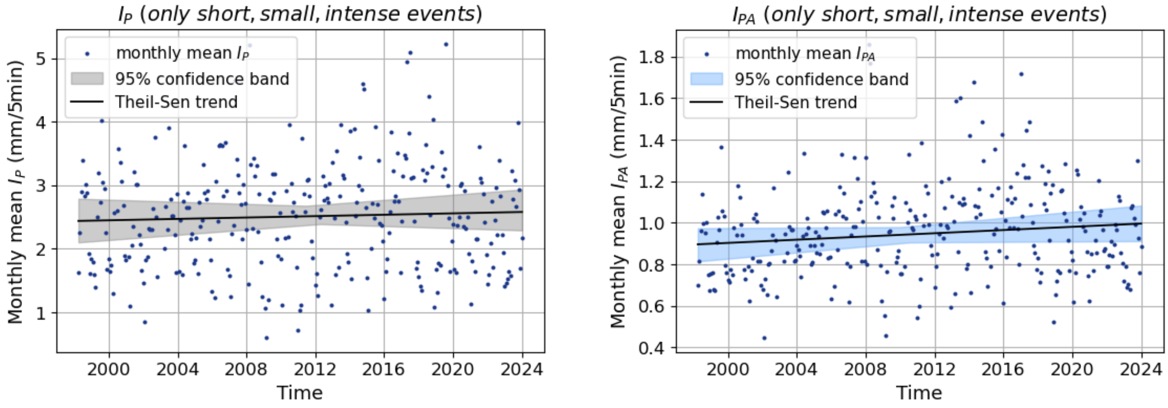

(e) Monthly means of I_P (peak intensity) for short, intense events.

(f) Monthly means of I_PA (peak average intensity) for short, intense events.

Figure A.8: Time series of monthly means of event variables for smaller-scale shorter-duration higher-intensity (SSHI) events (i.e. durations shorter than 300 minutes, areas smaller than 30 grid cells and average intensity greater than 0.4 mm/5min). The solid line represents the trend estimated using the Theil-Sen method, and the shaded area indicates the 95% confidence interval around the trend.
* : Significant trend at the 0.05 significance level. Confidence bands for variables where a significant trend is detected are shown in blue (otherwise grey).

A.3. Results Binned Analysis of Intensity Given Duration and Area

This section covers the results of a binned analysis of intensity, based on duration and average area. The duration and area were categorised into 25 bins using quantile binning, and trends in the three intensity variables – average intensity, peak intensity, and peak average intensity – were evaluated for each bin. Table A.1 displays the results of this analysis, including whether the trend is significantly increasing or decreasing (or not significant), the associated p-value, and the estimated Theil-Sen slope for each bin.

Table A.1: Results of the Mann-Kendall test for all three intensity variables, categorised by duration-area bin.

	D bin	**A bin**	**Var**	**Trend**	**p-value**	$\hat{Q}_s$
0	[20, 40]	[94, 111]	peak intensity	decreasing	<0.001	-0.0005
1	[20, 40]	[94, 111]	avg. intensity	decreasing	<0.001	-0.0001
2	[20, 40]	[94, 111]	peak avg. intensity	decreasing	<0.001	-0.0002
3	[20, 40]	[111, 141]	peak intensity	decreasing	<0.001	-0.0005
4	[20, 40]	[111, 141]	avg. intensity	decreasing	<0.001	-0.0001
5	[20, 40]	[111, 141]	peak avg. intensity	decreasing	<0.001	-0.0001
6	[20, 40]	[141, 195]	peak intensity	decreasing	<0.001	-0.0005
7	[20, 40]	[141, 195]	avg. intensity	decreasing	<0.001	-0.0001
8	[20, 40]	[141, 195]	peak avg. intensity	decreasing	<0.001	-0.0001
9	[20, 40]	[195, 351]	peak intensity	decreasing	<0.001	-0.0006
10	[20, 40]	[195, 351]	avg. intensity	decreasing	<0.001	-0.0001
11	[20, 40]	[195, 351]	peak avg. intensity	decreasing	<0.001	-0.0002
12	[20, 40]	[351, 39299]	peak intensity	decreasing	<0.001	-0.0009
13	[20, 40]	[351, 39299]	avg. intensity	decreasing	<0.001	-0.0002
14	[20, 40]	[351, 39299]	peak avg. intensity	decreasing	0.0011	-0.0002
15	[40, 65]	[94, 111]	peak intensity	decreasing	0.0045	-0.0005
16	[40, 65]	[94, 111]	avg. intensity	decreasing	0.0017	-0.0001
17	[40, 65]	[94, 111]	peak avg. intensity	decreasing	0.0031	-0.0002
18	[40, 65]	[111, 141]	peak intensity	decreasing	<0.001	-0.0006
19	[40, 65]	[111, 141]	avg. intensity	decreasing	0.0019	-0.0001
20	[40, 65]	[111, 141]	peak avg. intensity	decreasing	<0.001	-0.0002
21	[40, 65]	[141, 195]	peak intensity	decreasing	<0.001	-0.0008
22	[40, 65]	[141, 195]	avg. intensity	decreasing	<0.001	-0.0001
23	[40, 65]	[141, 195]	peak avg. intensity	decreasing	<0.001	-0.0002
24	[40, 65]	[195, 351]	peak intensity	decreasing	<0.001	-0.0009
25	[40, 65]	[195, 351]	avg. intensity	decreasing	<0.001	-0.0001
26	[40, 65]	[195, 351]	peak avg. intensity	decreasing	<0.001	-0.0002
27	[40, 65]	[351, 39299]	peak intensity	decreasing	<0.001	-0.0008
28	[40, 65]	[351, 39299]	avg. intensity	decreasing	<0.001	-0.0001
29	[40, 65]	[351, 39299]	peak avg. intensity	decreasing	0.0030	-0.0002
30	[65, 105]	[94, 111]	peak intensity	decreasing	0.019	-0.0007
31	[65, 105]	[94, 111]	avg. intensity	decreasing	0.039	-0.0001
32	[65, 105]	[94, 111]	peak avg. intensity	decreasing	0.019	-0.0002
33	[65, 105]	[111, 141]	peak intensity	decreasing	0.018	-0.0006
34	[65, 105]	[111, 141]	avg. intensity	decreasing	0.025	-0.0001
35	[65, 105]	[111, 141]	peak avg. intensity	decreasing	0.0094	-0.0002
36	[65, 105]	[141, 195]	peak intensity	decreasing	0.011	-0.0008
37	[65, 105]	[141, 195]	avg. intensity	decreasing	0.018	-0.0001
38	[65, 105]	[141, 195]	peak avg. intensity	decreasing	0.0074	-0.0002
39	[65, 105]	[195, 351]	peak intensity	decreasing	0.018	-0.0007
40	[65, 105]	[195, 351]	avg. intensity	no trend	0.056	-0.0001
41	[65, 105]	[195, 351]	peak avg. intensity	decreasing	0.021	-0.0002
42	[65, 105]	[351, 39299]	peak intensity	decreasing	<0.001	-0.0011
43	[65, 105]	[351, 39299]	avg. intensity	decreasing	0.0079	-0.0001

Table A.1 continued from previous page

44	[65, 105]	[351, 39299]	peak avg. intensity	decreasing	0.0081	-0.0002
45	[105, 165]	[94, 111]	peak intensity	no trend	0.069	-0.0007
46	[105, 165]	[94, 111]	avg. intensity	no trend	0.23	-0.0000
47	[105, 165]	[94, 111]	peak avg. intensity	no trend	0.13	-0.0002
48	[105, 165]	[111, 141]	peak intensity	no trend	0.073	-0.0007
49	[105, 165]	[111, 141]	avg. intensity	no trend	0.15	-0.0001
50	[105, 165]	[111, 141]	peak avg. intensity	decreasing	0.043	-0.0002
51	[105, 165]	[141, 195]	peak intensity	no trend	0.24	-0.0005
52	[105, 165]	[141, 195]	avg. intensity	no trend	0.20	-0.0001
53	[105, 165]	[141, 195]	peak avg. intensity	no trend	0.20	-0.0001
54	[105, 165]	[195, 351]	peak intensity	decreasing	0.030	-0.0010
55	[105, 165]	[195, 351]	avg. intensity	no trend	0.15	-0.0001
56	[105, 165]	[195, 351]	peak avg. intensity	no trend	0.064	-0.0002
57	[105, 165]	[351, 39299]	peak intensity	decreasing	<0.001	-0.0017
58	[105, 165]	[351, 39299]	avg. intensity	decreasing	0.0012	-0.0001
59	[105, 165]	[351, 39299]	peak avg. intensity	decreasing	<0.001	-0.0003
60	[165, 3040]	[94, 111]	peak intensity	no trend	0.76	-0.0003
61	[165, 3040]	[94, 111]	avg. intensity	no trend	0.81	-0.0000
62	[165, 3040]	[94, 111]	peak avg. intensity	no trend	0.71	-0.0001
63	[165, 3040]	[111, 141]	peak intensity	no trend	0.31	-0.0006
64	[165, 3040]	[111, 141]	avg. intensity	no trend	0.87	-0.0000
65	[165, 3040]	[111, 141]	peak avg. intensity	no trend	0.98	0.0000
66	[165, 3040]	[141, 195]	peak intensity	no trend	0.71	-0.0002
67	[165, 3040]	[141, 195]	avg. intensity	no trend	0.64	-0.0000
68	[165, 3040]	[141, 195]	peak avg. intensity	no trend	0.75	-0.0000
69	[165, 3040]	[195, 351]	peak intensity	no trend	0.16	-0.0009
70	[165, 3040]	[195, 351]	avg. intensity	no trend	0.12	-0.0001
71	[165, 3040]	[195, 351]	peak avg. intensity	decreasing	0.043	-0.0003
72	[165, 3040]	[351, 39299]	peak intensity	no trend	0.10	-0.0014
73	[165, 3040]	[351, 39299]	avg. intensity	no trend	0.16	-0.0001
74	[165, 3040]	[351, 39299]	peak avg. intensity	no trend	0.066	-0.0003

A.4. Results of the Change Point Test for Correlation

In this section, we present the results of a change point test for correlation, which was applied to examine changes in the correlation functions between pairs of event variables. The test helps to identify potential shifts in the relationships between the variables over the study period.

A.4.1. Results for Analysis with Bandwidth k=1431

Here we present the results of the change point test applied with a bandwidth $k = \lfloor n^{5/8} \rfloor = 1431$. For each pair of event variables, plots of the estimator $M_n(u)$ of the integrated correlation function alongside the corresponding CUSUM process $M_n(u)$ are shown.

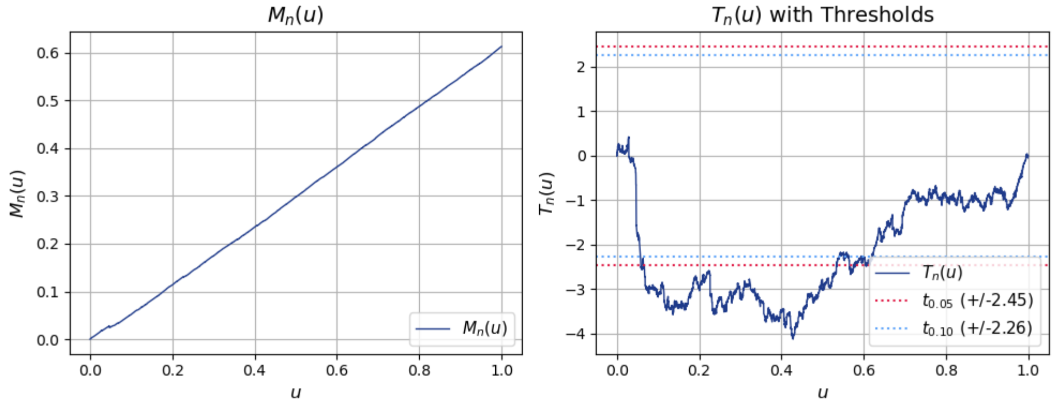

Figure A.9: The estimator $M_n(u)$ of the integrated correlation (left) between $\boldsymbol{D}$ and $\boldsymbol{A}_{\mathbf{A}}$ and the corresponding CUSUM process $T_n(u)$ (right) with critical thresholds at the 5% and 10% level.

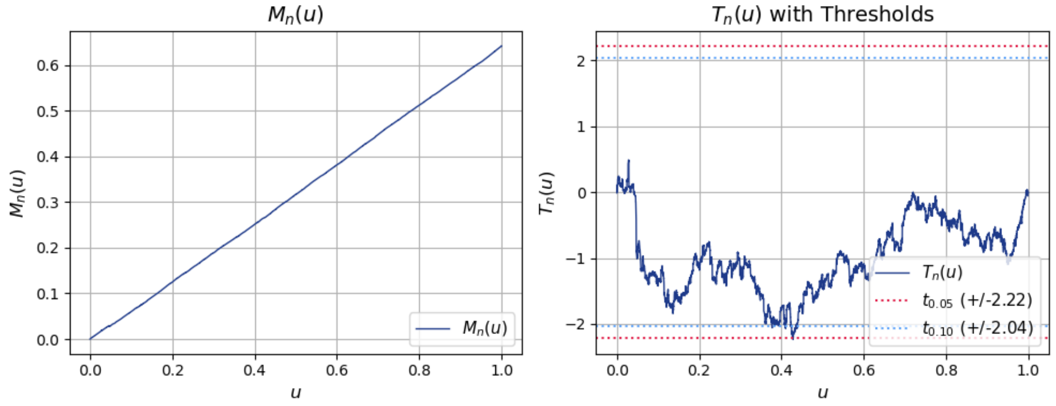

Figure A.10: The estimator $M_n(u)$ of the integrated correlation (left) between $\boldsymbol{D}$ and $\boldsymbol{A}_{\mathbf{P}}$ and the corresponding CUSUM process $T_n(u)$ (right) with critical thresholds at the 5% and 10% level.

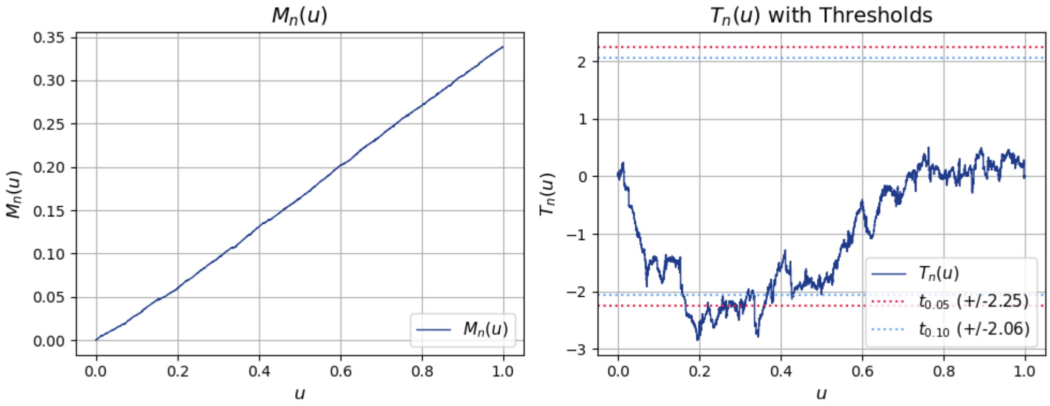

Figure A.11: The estimator $M_n(u)$ of the integrated correlation (left) between $\boldsymbol{D}$ and $\boldsymbol{I}_{\mathbf{A}}$ and the corresponding CUSUM process $T_n(u)$ (right) with critical thresholds at the 5% and 10% level.

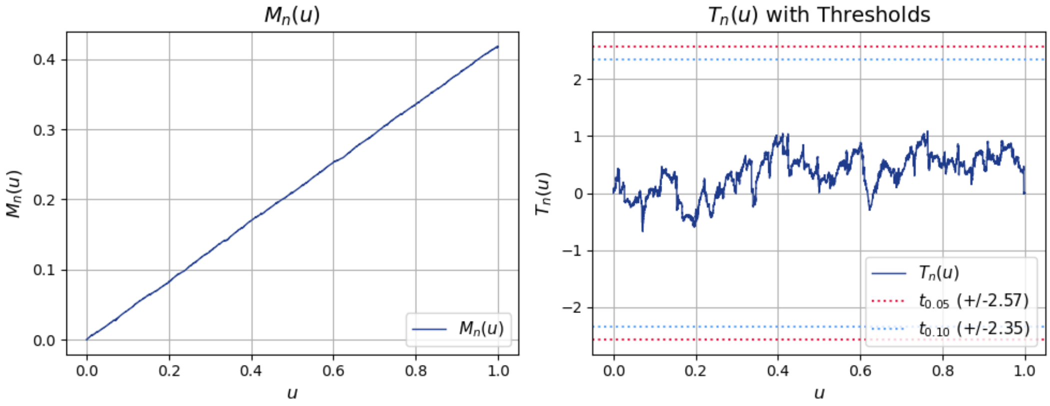

Figure A.12: The estimator $M_n(u)$ of the integrated correlation (left) between $\boldsymbol{D}$ and $\boldsymbol{I}_{\mathbf{PA}}$ and the corresponding CUSUM process $T_n(u)$ (right) with critical thresholds at the 5% and 10% level.

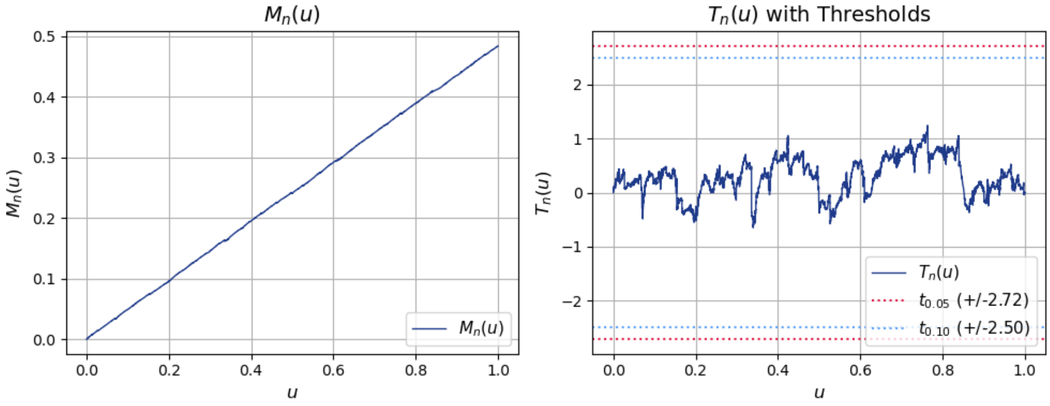

Figure A.13: The estimator $M_n(u)$ of the integrated correlation (left) between $\boldsymbol{D}$ and $\boldsymbol{I}_{\mathbf{P}}$ and the corresponding CUSUM process $T_n(u)$ (right) with critical thresholds at the 5% and 10% level.

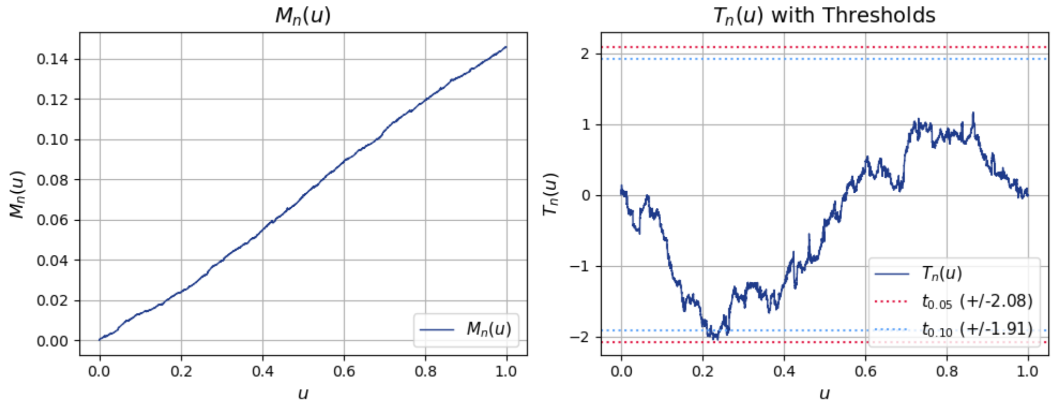

Figure A.14: The estimator $M_n(u)$ of the integrated correlation (left) between $\boldsymbol{A}_{\mathbf{A}}$ and $\boldsymbol{I}_{\mathbf{A}}$ and the corresponding CUSUM process $T_n(u)$ (right) with critical thresholds at the 5% and 10% level.

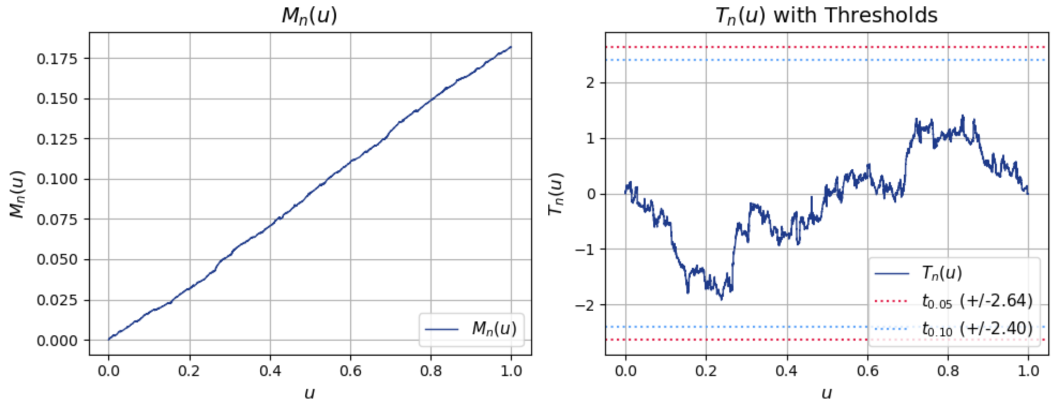

Figure A.15: The estimator $M_n(u)$ of the integrated correlation (left) between $\boldsymbol{A}_{\mathbf{A}}$ and $\boldsymbol{I}_{\mathbf{PA}}$ and the corresponding CUSUM process $T_n(u)$ (right) with critical thresholds at the 5% and 10% level.

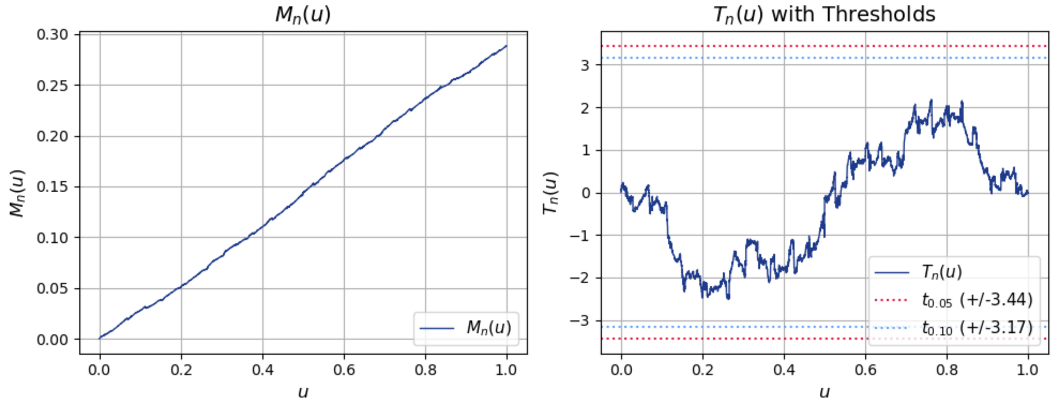

Figure A.16: The estimator $M_n(u)$ of the integrated correlation (left) between $\boldsymbol{A}_{\mathbf{A}}$ and $\boldsymbol{I}_{\mathbf{P}}$ and the corresponding CUSUM process $T_n(u)$ (right) with critical thresholds at the 5% and 10% level.

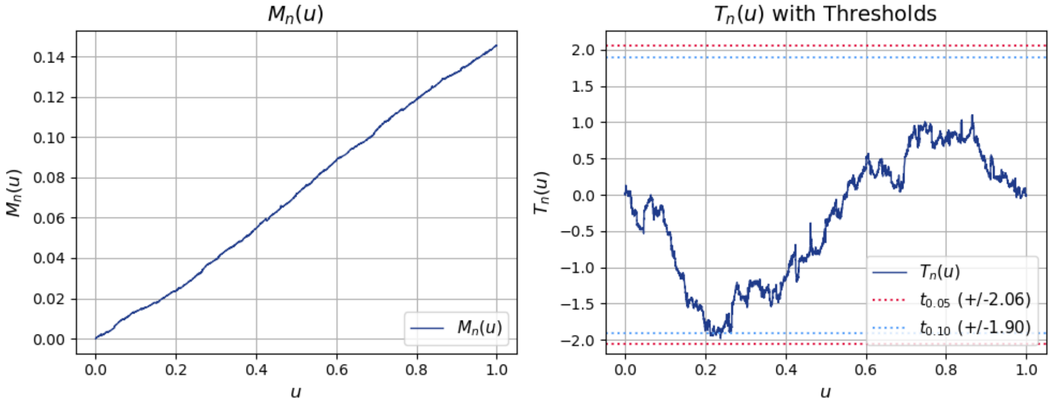

Figure A.17: The estimator $M_n(u)$ of the integrated correlation (left) between $\boldsymbol{A}_{\mathbf{P}}$ and $\boldsymbol{I}_{\mathbf{A}}$ and the corresponding CUSUM process $T_n(u)$ (right) with critical thresholds at the 5% and 10% level.

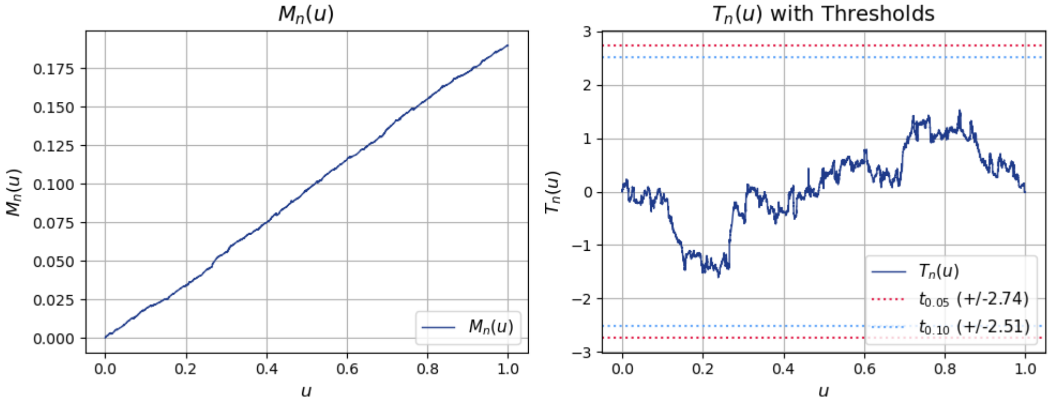

Figure A.18: The estimator $M_n(u)$ of the integrated correlation (left) between $\boldsymbol{A}_{\mathbf{P}}$ and $\boldsymbol{I}_{\mathbf{PA}}$ and the corresponding CUSUM process $T_n(u)$ (right) with critical thresholds at the 5% and 10% level.

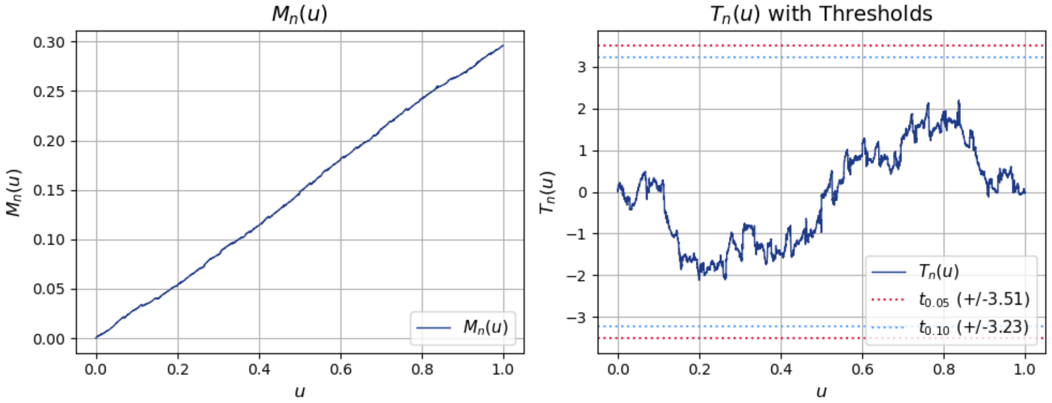

Figure A.19: The estimator $M_n(u)$ of the integrated correlation (left) between $\boldsymbol{A}_{\mathbf{P}}$ and $\boldsymbol{I}_{\mathbf{P}}$ and the corresponding CUSUM process $T_n(u)$ (right) with critical thresholds at the 5% and 10% level.

A.4.2. Results for Other Bandwidths k

In this section, the change point analysis is extended to two additional bandwidths k. Tables A.2 and A.3 show the results of the change point test applied to these different bandwidths, allowing for a comparison of how varying the bandwidth affects the outcome of the test. The results for these alternative bandwidths show that the identified changes in correlations do not vary significantly across different bandwidths.

Table A.2: p-values from the change point test procedure ($k = \lfloor n^{4/7} \rfloor = 768$) for pairs of rainfall event variables. *: Significant at 10% level. **: Significant at 5% level.

$k = \lfloor n^{4/7} \rfloor = 768$		D	A_A	A_P	I_A	I_P	I_PA
(duration)	D		<0.001**	0.135	<0.001**	0.56	0.17
(avg. area)	A_A				0.004**	0.014**	0.14
(peak area)	A_P				0.015**	0.054	0.40
(avg. intensity)	I_A						
(peak intensity)	I_P						
(peak avg. intensity)	I_PA						

Table A.3: p-values from the change point test procedure ($k = \lfloor n^{2/3} \rfloor = 2324$) for pairs of rainfall event variables. *: Significant at 10% level. **: Significant at 5% level.

$k = \lfloor n^{2/3} \rfloor = 2324$		D	A_A	A_P	I_A	I_P	I_PA
(duration)	D		<0.001**	0.038*	0.014**	0.26	0.48
(avg. area)	A_A				0.036*	0.23	0.23
(peak area)	A_P				0.054	0.41	0.47
(avg. intensity)	I_A						
(peak intensity)	I_P						
(peak avg. intensity)	I_PA						

B

Preprocessing of the Radar Data

The uncropped radar dataset was received from KNMI in the native HDF5 format. Even though uncropped, all data outside a radius of 200 kilometers from De Bilt is masked. Since the Celltrack software deals only with NetCDF4 or GRIB format, some preprocessing is necessary before Celltrack can be used. Moreover, the original data has no temporal dimension, but this is required for merging the files into seasonal files, which is desired for running Celltrack. Additionally, the radar data is corrected using data coming from rain gauges. Since these gauges are located on the land surface of the Netherlands only, the correction will not work for locations far away (approximately 30 kilometers) from the land borders. Therefore, a mask needs to be applied to the data. Hence, the steps of preprocessing the radar files are as follows:

1. convert all 5-minute HDF5 files to NetCDF
2. add temporal dimension to every file
3. apply an extended landmask to the data
4. combine the masked NetCDF files into single files per season (i.e. 3 months)

Note that when we speak of a season, we refer to meteorological seasons. For example, one season is Summer and it includes the complete months June, July and August. In Figure B.1 an example of one time step of the preprocessed data is shown.

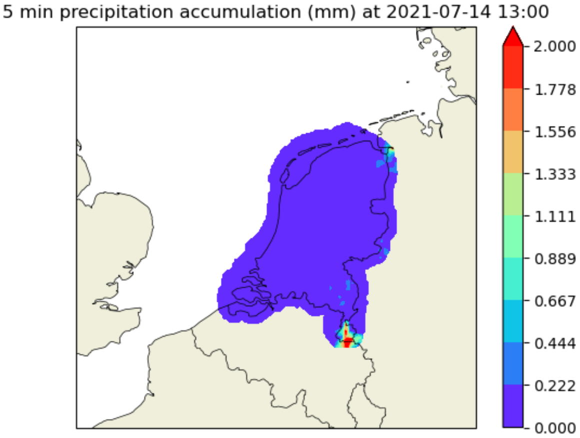

Figure B.1: Example of one time step of a preprocessed seasonal radar file. This figure shows precipitation depths in mm between 12:55 and 13:00 on July 14th, which was during the extreme precipitation events in Limburg during July 2021.

The runtime of the preprocessing of the radar data was in total around 3 to 3.5 days, running on a regular computer. Making quick visualizations of the data can be done in the NetCDF viewer software Panoply [36].

C

Running Celltrack

There are many possible choices for the settings of Celltrack, which are described in [27]. Hence, some choices need to be made, mainly regarding the area and intensity thresholds. For this study, the choice was made to choose quite low thresholds and filter out unnecessary or unwanted tracks at a later stage. With this strategy, we avoid eliminating tracks that might turn out to be important later on. The thresholds with which Celltrack was run are 0.1 mm/5min in terms of minimum intensity and 4 grid cells (equaling 25 km^2) in terms of minimum area. The runtime of Celltrack for all seasons during the 26 years is approximately 3 days.

The complete list of output files produced by Celltrack is described in [27]. The files that will be the main files for analysis are tracks_all_summary.txt and the tracks_all_stats.txt. The idea of the post-processing is to filter the tracks in tracks_all_summary.txt and append information, computed from tracks_all_stats.txt, to this file. Filtering of the tracks is done based on the following criteria, where different choices for m, n, and i can be considered.

1. Tracks must have duration longer than m minutes.
2. Tracks must have an average cell area larger than n grid cells.
3. Tracks must have an average intensity larger than i mm/5min.

Information appended to the summary file includes the following variables.

1. maximum average intensity. This is the maximum of rain cell average intensities over all rain cells belonging to the track.
2. average area. This is the average of rain cell areas over all rain cells in the track.
3. area at peak value. Consider the very maximum of grid cell values of the whole track. The area at peak value is the area at the time this maximum occurs.
4. x coordinate of the center of mass at peak time. This is the x coordinate of the center of mass of the cell where the peak intensity grid cell occurs.
5. y coordinate of the center of mass at peak time.

The runtime of filtering for all of the data can be up to 12 hours, depending on the choices made for filtering and additional variables that need to be determined.