

Multivariate Methods for Coastal and Offshore Risks

Jaeger, Wiebke

DOI

[10.4233/uuid:4be7ac1d-8232-457c-8ad6-73ec23baf9ce](https://doi.org/10.4233/uuid:4be7ac1d-8232-457c-8ad6-73ec23baf9ce)

Publication date

2018

Document Version

Final published version

Citation (APA)

Jaeger, W. (2018). *Multivariate Methods for Coastal and Offshore Risks*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:4be7ac1d-8232-457c-8ad6-73ec23baf9ce>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

MULTIVARIATE METHODS FOR COASTAL AND OFFSHORE RISKS

MULTIVARIATE METHODS FOR COASTAL AND OFFSHORE RISKS

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology,
by the authority of the Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen,
chair of the Board for Doctorates,
to be defended publicly on Monday 11 June 2018 at 10:00 o'clock

by

Wiebke Solvig JÄGER

Master of Science in Engineering and Policy Analysis and in Applied Mathematics,
Delft University of Technology, the Netherlands,
born in Hamburg, Germany.

This dissertation has been approved by the promotor.

Composition of the doctoral committee:

Rector Magnificus,	chairman
Prof. dr. ir. M.J.F. Stive,	Delft University of Technology, promotor
Dr. ir. A.M. Hanea,	The University of Melbourne, promotor
Dr. ir. O. Morales Nápoles,	Delft University of Technology, copromotor

Independent members:

Prof. dr. J.A. Jiménez Quintana	BarcelonaTech, Spain
Prof. dr. ir. P.H.A.J.M. van Gelder	Delft University of Technology
Prof. dr. ir. S.N. Jonkman,	Delft University of Technology
Dr. ir. A.R. van Dongeren,	Deltares
Prof. dr. A. Metrikine	Delft University of Technology, reserve member



Keywords: Probabilistic simulation, hazard analysis, source-pathway-receptor concept, Bayesian network, vine-copula, time series modeling

Printed by: Ridderprint BV

Front & Back: Phaedra Oikonomopoulou

Copyright © 2018 by W.S. Jäger

ISBN 978-94-6375-009-7

An electronic version of this dissertation is available at
<http://repository.tudelft.nl/>.

CONTENTS

Summary	vii
Samenvatting	ix
I Introduction	1
1 Research Context	3
1.1 Examples of Risks in Coastal and Offshore Environments	3
1.2 Decision Making on Risk Reduction Efforts in Coastal Environments	4
1.3 Statistical Simulation Methods for Time Series of Metocean Variables	7
1.4 Research Objectives and Scope	10
1.5 Thesis Structure.	12
2 Probabilistic Preliminaries	15
2.1 Discrete Bayesian Networks (BNs)	15
2.2 Copulas and Vines	16
2.3 Autoregressive Moving-Average (ARMA) Processes	18
II BN Framework for Decision Support on Risk Reduction Efforts	21
3 Framework Design	23
3.1 Methodological Background	23
3.2 Design of the Decision Support System	26
3.3 Example of Application	29
3.4 Integration into Operational Forecasting Systems.	39
3.5 Key Points.	40
4 Recommended Extensions	41
4.1 Use of the BN as an Ensemble Mean Predictor	42
4.2 A Metric for BN Validation	43
4.3 Incorporation of Training Data Uncertainties	46
4.4 Choice of Training Data and Model Structure	47
4.5 Optimization of Spatial Zones.	48
4.6 Optimization of Boundary Node Discretization	50
4.7 Key Points.	52
III Statistical Simulation Methods for Time Series of Wave Conditions	53
5 Temporal Vine-Copula Method	55
5.1 Data.	56
5.2 Representation of Time Series as Vines	58

5.3	Results	61
5.4	Key Points.	70
6	ARMA Method	71
6.1	Data and Regime Definition.	72
6.2	Model Development	74
6.3	Simulation results.	88
6.4	Key Points.	92
7	Comparison and Limitations	95
7.1	Model Comparison	95
7.2	Model Limitations	102
7.3	Key Points.	104
IV	Conclusion	105
8	Main Contributions	107
9	Recommendations	111
	Appendices	113
A	Learning Algorithm for the Conditional Probability Tables of Hazard Nodes	115
B	BN for Wells-next-the-Sea in GeNIe	117
C	Sampling Algorithms for Vine-Copulas	121
D	Background on Selected Copula Families	123
D.1	Tawn	123
D.2	Gamma 1-Factor Model.	123
D.3	Skew-t	123
	References	125
	Acknowledgements	141
	Curriculum Vitæ	143
	List of Publications	145
	Conferences and Lectures	147

SUMMARY

This thesis investigates how selected multivariate probabilistic methods can be adapted for risk analysis and decision making in coastal and offshore environments. In particular, the thesis makes a contribution to decision support tools for risk reduction efforts in coastal environments and to statistical simulation methods for wave conditions.

Generally, very few observations on negative impacts in coastal or offshore environments are available for risk analysis or decision making due to the rare nature of extreme events. However, synthetic impact data can be generated by propagating relevant hydro-meteorological conditions to the environment of interest through a chain of multiple models. Especially in coastal environments, this chain often includes computationally intensive models.

Two of the current challenges are to instantly predict storm impacts in coastal areas and to make the propagation of risks transparent. Both aspects are relevant for emergency management as well as long-term planning. Inspired by numerous applications of Bayesian networks (BNs) as tools for decision-making in a risk context, a framework for decision support on risk reduction efforts in coastal areas has been developed in this thesis.

A main contribution is a learning algorithm for estimating the parameters of the BN. The algorithm can integrate simulations of synthetic or historical hydro- and morphodynamic storm simulations with land-use data, vulnerability relationships (e.g., depth damage curves) and risk reduction measures. The core of the algorithm builds on SMILE, which is a reasoning engine for graphical models.¹ The algorithm is programmed in C++ and openly available at: <https://github.com/openearth/coastal-dss>.

The BN framework is applicable to any coastal site and models constructed accordingly have the potential to:

- predict impacts for any relevant storm events instantly,
- reflect the diversity of impacts (e.g., socio-economic and environmental aspects),
- reflect the spatial variability of impacts,
- and
- evaluate the performance of risk reduction measures under different events.

This thesis illustrated the application to the town Wells-next-the-Sea in North Norfolk, UK, which is vulnerable to storm surges.² BN output showed, for example, that extending the length of the existing flood wall can be more effective for impact reduction across the variety of storm events that were considered than increasing its height.

In addition to the learning algorithm, detailed recommendations are given for mathematical validation and optimization methods. These include metrics for 'prediction

¹SMILE is available free of charge for academic research and teaching use from BayesFusion, LCC, <http://bayesfusion.com/>.

²Within the RISC-KIT project, funded under the Seventh Framework Programme of the European Commission, applications have been made to nine other sites at all different European Seas (see www.riskkit.eu).

precision' and 'spatial informativeness' that are specific to the developed framework rather than BNs in general. Nonetheless, it is advised to involve intended end users in further model development to ensure that their needs are met, rather than solely relying on mathematical validation and optimization methods.

A third challenge is to determine which hydro-meteorological conditions are relevant for impact assessment and should be propagated through the model train. If relevant events are omitted in the assessment, risks will be underestimated. A complicating issue is that storm impacts can also depend on previous hydro-meteorological conditions, as they may have weakened the environment. Examples are progressive beach and dune erosion or structural fatigue.

For this reason many applications call for the generation of time series of hydro-meteorological conditions, often at temporal resolution of approximately one hour. In the literature, so-called vine-copulas and autoregressive moving-average (ARMA) processes have been applied to multivariate time series, also of hydro-meteorological variables. Nonetheless, high temporal resolutions and the accurate description of dependency between two variables which are elementary for characterizing wave conditions, significant wave height and mean zero-crossing period, are still under-explored.

In response, this thesis developed two statistical simulation methods based on a data set of wave conditions in the North Sea and compared their results. One method was based on vine-copulas and the other on ARMA processes. In the first method, the skew-t copula, which is a flexible 4-parameter family, was key component for capturing the dependence between significant wave height and mean zero-crossing period. This copula can also be applied to construct a bivariate distribution of significant wave height and mean zero-crossing period outside the time series context. In the second method, a data-driven equation for a maximum wave-steepness condition has been derived for the location of the data. The dependence is mainly captured by this equation and a copula on the ARMA residuals.

The scope of the two methods is different. The method based on vine-copulas focuses on bivariate time series during oceanographic winter periods assuming stationarity. The method based on ARMA processes additionally considers non-stationaries on annual and inter-annual scale as well as the influence of the wave direction on the joint behavior of significant wave height and mean zero-crossing period.

Both methods are suitable for generating synthetic time series of significant wave heights and mean zero-crossing periods. As they emulate storms as well as calm periods well, which is shown by analyzing their persistency above and below selected thresholds, they can be flexibly applied in different contexts. However, an uncertainty and sensitivity analysis would be required to validate models built according to either of the methods for any given application. One application is coastal risk analysis, for example via the BN framework developed in this thesis, for which synthetic storm events can be generated. A second application is the estimation of project duration for coastal or offshore operations that are sensitive to particular wave height and period limits.

SAMENVATTING

Dit proefschrift onderzoekt op welke wijze bepaalde multivariate probabilistische methoden geschikt kunnen worden gemaakt voor risicoanalyse en besluitvorming in kustgebieden of een offshore omgeving. In het bijzonder draagt dit proefschrift bij aan het construeren van besluitondersteunende gereedschappen voor het reduceren van risico's in een kustomgeving en van statistische simulatiemethoden voor golfcondities.

In het algemeen zijn er slechts weinig observaties beschikbaar van negatieve impact op kust of offshore omgevingen voor risicoanalyse of besluitvorming vanwege het zeldzame voorkomen van extreme gebeurtenissen. Echter, synthetische impactdata kunnen worden gegenereerd door relevante hydro-meteorologische condities te propageren naar het aandachtsgebied door middel van modellen. Specifiek in kustgebieden moet hierbij gebruik gemaakt worden van een keten van verschillende -soms rekenintensieve- modellen om de hydro-meteorologische condities te koppelen aan impacts.

Twee van de huidige uitdagingen betreffen de instantane voorspelling van de impact van stormen in kustgebieden en het inzichtelijk maken van het doorwerken van risico's. Beide aspecten zijn relevant voor zowel rampenbeheersing als voor langetermijnplanning. Geïnspireerd door vele toepassingen van Bayesiaanse netwerken (BNs) als gereedschap voor besluitvorming in een risicocontext, is een raamwerk ontwikkeld voor besluitvormende ondersteuning voor risico verminderende maatregelen in kustgebieden.

Een belangrijke bijdrage van dit proefschrift betreft een lerend BN algoritme dat synthetische of historische hydro- en morfodynamische storm simulaties kan integreren met landgebruik data, kwetsbaarheidsrelaties (bijvoorbeeld waterdiepte-schade diagrammen) en risico reducerende maatregelen. De kern van het algoritme is gebaseerd op SMILE, een reasoning engine voor grafische modellen.³ Het algoritme is geprogrammeerd in C++ en vrij verkrijgbaar op: <https://github.com/openearth/coastal-dss>.

Het BN-raamwerk is toepasbaar op een willekeurig kustgebied en hierop gebaseerde modellen hebben de mogelijkheid om:

- instantane impact te voorspellen van een willekeurige storm
- de variabiliteit aan impact weer te geven (bijv. socio-economische en milieu aspecten)
- de ruimtelijke variabiliteit van impacts te laten zien, en
- het gedrag van risico reducerende maatregelen te evalueren voor verschillende stormen.

Dit proefschrift illustreert de toepassing op de stad Wells-next-the-Sea in North Norfolk, UK, die kwetsbaar is voor stormvoeden.⁴ BN output laat bijvoorbeeld zien dat het

³SMILE is vrij beschikbaar voor academisch onderzoek en onderwijs via BayesFusion, LCC, <http://bayesfusion.com/>.

⁴Toepassingen voor negen andere gebieden aan verschillende Europese zeeën zijn onderdeel van het RISC-KIT project, welke gefinancierd is binnen het Seventh Framework Programme van de Europese Commissie (www.risckit.eu).

verlengen van de bestaande hoogwaterkering meer effectief kan zijn voor impact reductie onder een variëteit van beschouwde storm scenario's dan het vergroten van de hoogte.

In aanvulling op het lerend algoritme zijn gedetailleerde aanbevelingen geformuleerd voor de wiskundige validatie en voor optimalisatie methoden. Deze aanbevelingen bevatten metriecken voor 'voorspellingsprecisie' en 'ruimtelijk informatie gehalte'. Desalniettemin wordt het aanbevolen om mogelijke eindgebruikers te betrekken bij de verdere modelontwikkeling om er zeker van te zijn dat hun behoeften worden gedekt, eerder dan alleen te vertrouwen op wiskundige validatie en optimalisatie methoden.

Een derde uitdaging is vast te stellen welke hydro-meteorologische condities relevant zijn voor impactduiding en zouden moeten worden doorgeleid via de modelketen. Indien relevante stormen worden genegeerd in de evaluatie, zullen risico's worden ondergewaardeerd. Een complicerend aspect is dat storm impacts ook kunnen afhangen van voorgaande hydro-meteorologische condities, omdat zij de omgeving hebben kunnen aantasten. Voorbeelden zijn voortschrijdende strand- en duinerosie of vermoeiingsschade aan constructies.

Om deze reden vragen vele toepassingen om het genereren van tijdreeksen van hydro-meteorologische condities, veelal met een temporele resolutie van een uur. In de literatuur zijn zogenaamde vine-copulas en autoregressieve moving-average (ARMA) processen toegepast op multi-variate tijdreeksen, ook van hydro-meteorologische variabelen. Niettegenstaande, een hoge temporele resolutie en een nauwkeurige afhankelijkheidsbeschrijving tussen twee variabelen die elementair zijn om golfcondities te beschrijven, de significante golfhoogte en de gemiddelde zero-crossing periode, zijn onderbelicht.

In antwoord hierop ontwikkelt dit proefschrift twee statistische simulatiemethoden gebaseerd op een dataset van golfcondities in de Noordzee en heeft de uitkomst van deze methoden vergeleken. Eén methode was gebaseerd op vine-copulas en één op ARMA-processen. In de eerste methode is het sleutelcomponent voor het vangen van de afhankelijkheid tussen de significante golfhoogte en gemiddelde zero-crossing periode een skew-t copula, welke een flexibele 4-parameterfamilie is. Deze copula's kunnen ook worden toegepast op de constructie van een bivariate verdeling van significante golfhoogtes en gemiddelde zero-crossing periodes, los van de tijdreekscontext. De tweede methode maakt gebruik van een datagedreven vergelijking voor de maximale golfstijlheidconditie, die speciaal ontwikkeld is voor de locatie van de data. De afhankelijkheid is vooral beschreven door deze vergelijking en een copula voor de ARMA-residuen.

De twee methodes verschillen in scope. De vine-copula methode focust op bivariate tijdreeksen gedurende oceanografische winterperiodes en veronderstelt een stationair proces. De ARMA-methode houdt rekening met niet-stationariteiten op jaarlijkse en interjaarlijkse tijdschalen evenals de invloed van de golfrichting op het gemeenschappelijke gedrag van de significante golfhoogte en de gemiddelde zero-crossing periode.

Beide methodes zijn geschikt voor het genereren van synthetische tijdreeksen van significante golfhoogtes en gemiddelde zero-crossing periodes. De methodes kunnen flexibel worden toegepast in verschillende contexten, omdat de tijdreeksen zowel stormen als rustige condities weergeven. Dit blijkt uit de analyse van de persistentie boven hoge en onder lage drempelwaarden. Onzekerheids- en sensitiviteitsanalyses zijn echter noodzakelijk om de modellen voor een specifieke toepassing te valideren. Een toepas-

sing is de risicoanalyse in kustgebieden, bijvoorbeeld via het BN-raamwerk ontwikkeld in dit proefschrift, waarvoor synthetische stormen worden gegenereerd als input. Een tweede toepassing is het schatten van de projectduur voor operaties in kust- of offshore-gebieden die vatbaar zijn voor bepaalde golfcondities.

I

INTRODUCTION

RESEARCH CONTEXT

1.1. EXAMPLES OF RISKS IN COASTAL AND OFFSHORE ENVIRONMENTS

Coastal and offshore environments are exposed to a variety of hydro-meteorological forces that originate, for instance, from winds, waves, surges and precipitation. Coastal environments include beaches and dunes, engineered coastal defense structures (e.g., dikes, flood walls or revetments), and ecosystems (e.g., salt marshes and mangroves). Offshore environments include marine structures (e.g., breakwater, floating or bottom-founded platforms and quays/jetties/docks), vessels and ecosystems (e.g., coral reefs). Each of these coastal and offshore environments responds in its own way to hydro-meteorological forces, which can pose severe threats when reaching extraordinary magnitudes during storms.

In the past, the biggest threat from hydro-meteorological forces has been coastal flooding, causing communities to suffer from material damages, economic, political and social disruption, health issues, or damaged ecosystems [91, and references therein]. Extreme events have caused societal disasters and had ripple effects on regional or national scale. For instance, approximately 2000 people died in 1953 in The Netherlands, Belgium and the UK [88] during the North Sea storm surge. Less than two decades ago, in 2005, Hurricane Katrina caused more than 1800 fatalities the United States [96]. In countries with lower protection standards, death tolls have been even higher. A recent example is Typhoon Haiyan killing more then 6300 people in the Philippines in 2013 [128].

Flooding usually arises from high still water levels in combination with high wave heights and long wave periods when they cause:

- overtopping or overflow due to high absolute run-up levels [76],
- (catastrophic) failure of engineered flood defenses due to extreme hydrodynamic loads [153, 180],
- or
- breaches in dunes due to intensified erosion rates [171, 174, 176].

On one hand erosion increases the risk of flooding, because beaches and dunes often act as primary flood defenses to coastal developments and this role can be compromised if the dune and beach width are decreased [e.g., 163, and references therein]. As a consequence dunes could breach during a storm which would not have caused them to breach otherwise. On the other hand, erosion is also problematic in itself. It can destroy ecological habitats [55] as well as recreational areas [133], and threaten structures that are built close to the shoreline [108].

Offshore, the combined action of waves, wind and currents can damage or destroy marine structures or vessels. Issues include:

- structural failure under extreme loads from combined contributions of waves, wind and currents [125],
- structural failure due to fatigue from oscillatory stresses [188],
- scour development creating foundation instability or cable exposure [187], and
- extreme motion responses due to wave periods close to a resonant heave period [54].

For example, during Hurricane Katrina, 44 oil and gas platforms in the Gulf of Mexico were completely destroyed and 21 others severely damaged [40].

Furthermore, weather-related down-time poses a significant financial risk to offshore operations, which usually consists of multiple sequential activities with different durations and executed by different vessels [e.g., 102, for an example of the cable installation for offshore wind farms]. In general, such activities are weather-sensitive and can only be executed as long as hydro-meteorological variables stay within certain limits. Acero et al. [3] provide a recent overview of such limits developed for various operations. If the limits are exceeded, operations cannot start or may need to be interrupted leading to unexpected down-time causing project delays and related costs.

The above examples of risks in coastal and offshore environments show that the field of risk analysis and risk management in coastal and offshore environments is broad. This doctoral thesis concentrates on the multivariate nature of the risks. More specifically, the focus will be on two particular topics within the field: decision making on risk reduction efforts in coastal areas and statistical simulation techniques for time series of hydro-meteorological variables. The topics will be elaborated upon in the next two sections.

1.2. DECISION MAKING ON RISK REDUCTION EFFORTS IN COASTAL ENVIRONMENTS

Past coastal floods, for example the ones mentioned in the previous section, have emphasized a continuing need for effective coastal risk management. This is becoming all the more important as risks are projected to increase globally, due to growing populations and assets, accelerated sea level rise and potential increases in storminess (both tropical and extra-tropical) [70, 74, 177].

Coastal risk management essentially includes two types of activities: (1) taking prompt actions in the face of an impending storm and (2) long-term planning. Accordingly, we distinguish between a *hot phase* and a *cold phase*. In the hot phase, emergency managers depend on real-time and reliable predictions of the expected conditions in the coastal zone, as they attempt to select mitigation measures and allocate limited resources in order to keep negative impacts as small as possible. In the cold phase, multiple actors, including politicians, local stakeholders and scientists, cooperate to determine sensible strategies for reducing risks in an uncertain future, while considering their constraints [52]. To evaluate these strategies against historical and conceivable future storms, they turn to impact assessments.

Cutting across numerous disciplines, including oceanography, coastal science and engineering, statistics, economics, and social and political science, coastal risk assessment is highly challenging. Each field has complex models which target individual elements of the risk process. For example, multivariate statistical models estimate the return periods of extreme hydro-meteorological events [45, 184], while numerical models, based on, for instance, hydro- and morphodynamic processes, determine the respective natural responses of the coast and extent of flooding [13, 73, 142, 185]. Finally, behavioral or statistical models estimate the diverse and complex consequences onshore [4, 69, 93, 116]. However, risk management requires a framework that integrates the individual elements of the risk process [28]. Thus, various model outputs need to be integrated in order to understand what drives losses and how they could be reduced.

A number of issues arise when attempting to link offshore sea conditions to their expected onshore hazards and impacts into a single model for operational use. Numerical models, being computationally expensive, often have a long run time, while instant assessments are needed for any conceivable hazard event during both the hot and cold phases. On the other hand, the spatial and temporal scales of numerical and impact models differ from one another and need to be integrated. Numerical models have grids whose sizes depend on the physical properties of the area under consideration, whereas impact models usually operate on the level of individual receptors. Furthermore, there are various types of impact, such as property damages, lost lives or disruption of infrastructures, each of which has a unique spatial distribution under each storm event. To evade the high-dimensionality, impact data can be aggregated into a single total estimate for each storm event, for instance in a monetary unit, but this can be disadvantageous for designing risk reduction measures.

To address some of these challenges, Jonkman et al. [91] developed an integrated model to assess the flood hazard and corresponding damages to the built environment, loss of life, as well as indirect economic impacts (e.g., the interruption of production flows) for one hypothetical extreme event. The fundamental element of this approach is a spatial database through which the authors connect output and input of the individual models according to a common spatial attribute. In this way the Jonkman et al.'s model addressed the challenges of different scales. While the study was limited to a single storm, the databased could be extended to multiple storm events. In Portugal, Poelhekke et al. [136] integrated a wide range of simulated storm events in a discrete Bayesian network (BN), which is a graphical model that describes system relations in probabilistic terms and can give instantaneous predictions. Using this model the authors assessed related flooding and erosion hazards in Praia de Faro.

Nevertheless, Poelhekke et al.'s approach did not estimate impacts and neither model provides insight into the outcomes of risk reduction measures. However, for decision makers it is fundamental to understand the impacts under various events and how they would respond to human interventions in order to design effective risk reduction strategies and to avoid unforeseen consequences [e.g., 58]. Thus, further integration of various models is needed to support decision making on risk reduction efforts in coastal areas.

While exploratory analyses of relevant storm events and effects of risk reduction measures are fundamental, more information may be required regarding a particular measure's performance before it is finally selected and implemented in practice. For ex-

ample, an important criterion in the evaluation of risk reduction strategies is the achievable *safety level*. This term can be used synonymous to risk level, but is often more specifically defined in terms of *individual risk* or *societal risk* [e.g., 5]. In general, *risk* can be defined as the expected value of loss¹ for a given time horizon [15]. In this definition loss could refer to a specific type of damage that is relevant for the application at hand, or refer to the sum of all adverse consequences aggregated into a single value, for example, in a monetary unit. On the other hand, individual risk refers to the annual probability of death of an average, unprotected person and societal risk represents the annual probability of multi-fatality events. For example, in the Netherlands, individual and societal risk criteria legally apply to installations of hazardous industries, such as chemical plants or LPG-fueling stations [6], and similar safety levels are targeted by the Dutch government for people living behind primary flood defenses [47].

Probabilistic risk analysis (PRA) is the norm for assessing safety levels and for determining corresponding requirements for structural designs. The method extends exploratory analyses, in that it not only deals with the identification of events and their consequences, but also with the quantification of their probabilities [15]. PRA is also the foundation of an economic optimization, which aims to be an objective approach to selecting the most beneficial risk reduction strategy from a set of candidate strategies [90]. In this case, the most beneficial strategy is the one that minimizes risk, when the initial investment and maintenance costs are included in the calculation as losses that occur with certainty.

A challenge lies in the quantification of probability density functions for losses. As mentioned before, empirical loss data are hardly available, but synthetic data could be obtained, for example, by performing a Monte Carlo simulation of hydro-meteorological forcing conditions and propagating them through a train of models, which target the individual elements of the risk process (Figure 1.1). The probability density function can then be estimated from the synthetic loss data. In this case, a prerequisite for high-quality loss data is a high-quality Monte Carlo sample of hydro-meteorological forcing conditions. Also the quality of the individual models in the train is fundamental. However, understanding and evaluating them requires discipline-specific knowledge of each domain and is beyond the scope of this research effort.




Figure 1.1: Schematic illustration of a model train for coastal risk analysis.

It is important that the Monte Carlo simulation of forcing events is ‘complete’, because there is no one-to-one mapping between the probabilities of forcing events and the probabilities of loss events. Instead, very similar loss events, could arise from different hydro-meteorological forcing events. To illustrate this, we suppose that the sample space of forcing events (i.e., the set of all possible forcing events) can be partitioned into $\{s_1, \dots, s_n\}$ and that l is a fixed loss value. Then, the exceedance probability for this loss

¹This definition is often less formally stated as $\text{risk} = \text{probability} \times \text{loss}$.

can be calculated with the law of total probability:

$$P(L > l) = \sum_{i=1}^n P(L > l \mid S = s_i)P(S = s_i). \quad (1.1)$$

The equation shows that omitting any summand i leads to an underestimation of $P(L > l)$, which ultimately leads to an underestimation of risk. For practical applications this means that the sample space of forcing events needs to be well approximated in order to well estimate $P(L > l)$, at least for those events in which $P(L > l \mid S = s_i)$ is non-zero. This issue has also been pointed out by Jonkman et al. [92]. To limit a possible underestimation of risk, they chose a number of "representative" events for their flood risk assessment based on the highest probabilities of occurrence.

While illustrative, equation (1.1) is an over-simplification of the problem. A first issue is that it is difficult to define S . In fact, each s_i can be thought of as a realization of a multivariate, spatio-temporal stochastic process of water levels, wave heights, wave periods, incident wave angles, wind speeds, wind directions, precipitation, and other variables of interest. A second issue is that the extent of damages may depend on many past forcing events, as they could have weakened the environment. For example, frequent storms can have a great impact on the beach and dune profile [8, 95]. Similarly, structures and vessels can fatigue, if repeatedly subjected to high hydrodynamic loads [188]. Therefore, it is important to account for storm sequencing in risk analyses.

1.3. STATISTICAL SIMULATION METHODS FOR TIME SERIES OF METOCEAN VARIABLES

Sea storms are segments of multivariate temporal processes of hydro-meteorological variables that pose a hazard to the coastal and offshore environment or operations. Typically, these processes are described by hourly statistics, for example, the significant wave height, which is computed from a spectrum of individual waves during a one hour record of individual waves. The processes exhibit strong state-to-state autocorrelation on short time scales, seasonal cycles on annual and multi-annual time scales, inter-series dependences and, potentially, long-term trends [e.g., 114, 115, 158]. These statistical features make it challenging to model time series of hydro-meteorological variables, including sea storms.

Many simulation methods are based on binary renewal processes to model alternating sequences of storm and calm durations [31, 38, 43, 46, 104, 184]. A renewal process is an idealized description of events that occur randomly in time. In the studies, storms are defined in terms of the duration for which one or more variables exceed predefined thresholds, together with the values of all variables of interest at the peak of the storm. Furthermore, these are modeled as interdependent, for example with copulas in above studies.

Copulas isolate the marginal properties from the dependence structure of random variables. Genest and Favre [61], for example, provide a short introduction to copula modeling. Comprehensive texts include Nelsen [129] and Joe [85]. A combination of any copula with any marginal distribution leads to a valid specification of a joint distribution, enabling representations of a wide range of complex multivariate behaviors. In the




Figure 1.2: Illustration of a "triangular equivalent storm" in terms of significant wave height (H_{m0}) and storm duration (D).

bivariate case, many different copula families have proven to be useful (e.g., [111, 145], and references therein). For more than two oceanographic variables, nested (also called hierarchical) Archimedean copulas [38, 106, 183] and elliptical copulas, such as Gaussian or t , [103, 143, 184] have been implemented and found valuable. Dependence trees [136] and vines [45, 122], which are a generalization thereof, are alternative proposals. Callaghan et al. [31] and Serafin and Ruggiero [150] adopt another approach to model dependencies. They use a bivariate logistics model [165] and they specify parameters for a conditional distribution of one variable based on the value of the conditioning variable.

Not only the dependencies between variables are important for impact assessment, but also their temporal evolution; impacts amass during long-lasting or recurring extreme conditions [e.g., 8, 95, and references therein for impacts on a sandy beach]. For the storm periods, high temporal resolution time series of the relevant hydrological and meteorological variables are then derived from an idealized 'storm shape'. For most applications hourly values are needed. A typical assumption is that each univariate time series segment corresponds to two sides of a symmetric triangle whose height determines the peak value and whose base is defined by the storm duration. Figure 1.2 provides an illustration. For instance, Corbella and Stretch [37] and Poelhekke et al. [136] have used these triangles and the dependencies between variables at the peak to derive idealized storm time series with high resolution (~ 1 h) to force numerical, physics-based models that compute resulting erosion and flooding. Differently, Wahl et al. [182] applied a linear regression to parameterize and simulate the temporal evolution of total water levels during storm surges. Similar schematizations with alternative geometrical shapes have also been suggested [110, 159].

Simulating storm sequences as renewal processes by adopting a geometrical storm shape based on pre-defined thresholds is a powerful approach to assess risks and, if seasonal climatic and long-term non-stationarities are represented, to identify future trends of risk. The main advantage is a minimized modeling effort, because features of the hydro-meteorological time series that are less relevant for the application are not resolved. An example are serial dependence or dependencies between the variables during calm periods.

Nonetheless, the simplification to a standard storm shape implies a loss of accuracy and it should be investigated how that affects risk estimates. Suppose s_i^* is the equivalent 'standard shape storm' of s_i . Is $P(L > l \mid S = s_i^*) \approx P(L > l \mid S = s_i)$? The answer will depend on the joint distribution of L and S and on "how close" s^* and s are. Furthermore, the choice of shape and threshold ought to depend on the application, because different coastal and offshore environments are vulnerable to other storm characteristics. For instance, an operating vessel can be sensitive to hydro-meteorological conditions that a sandy beach is not. Consequently, such storm models are application-specific and may not easily be transferable to a different environment. An alternative to the renewal process based approach is the modeling of complete time series. While increasing the modeling burden, it allows for more flexibility in terms of potential applications.

Currently, three lines of research concentrate on simulating multivariate time series of hydro-meteorological variables with high temporal resolution. Guanche et al. [63, and references therein] developed a simulation method based on statistical downscaling of synoptic scale weather patterns. The authors statistically simulate time series of larger-scale sea level pressure fields with autoregressive moving average (ARMA) models from which they then derive local sea state time series.

Furthermore, ARMA models have been used to directly represent time series of hydro-meteorological variables at a location, most of them at three-hourly scales. Multiple studies exist on univariate time series of significant wave heights [9, 65, 66, 149, 162]. Extension to bivariate processes have been made by including the mean wave periods [64] and by including surges [30]. In addition to significant wave height and peak period, Solari and van Gelder [158] incorporated parameters related to wind speed, wind direction and wave direction, thus simulating five interrelated processes. The bivariate and multivariate approaches used so-called vector ARMA models, which are able to capture linear interdependencies between multiple time series. However, Solari and van Gelder reported that dependencies could not always be adequately represented.

Finally, copulas and vine-copulas have been adopted to model both serial dependence as well as inter-series dependencies of hydro-meteorological processes. For instance, Leontaris et al. [102] simulated wind speeds and significant wave heights. Solari and Losada [157] compared a copula-based serial dependence model to an ARMA model for significant wave height time series. They found that storm frequency and persistence of storms were better represented by their copula-based model, whereas longterm autocorrelation was better represented by their ARMA model.

Different techniques have been used to account for non-stationarities. The simplest approach has been to focus on the most important season or to piecewise model seasons or months [102, 103, 184]. Other studies have used linear or cyclic functions of time [9, 31, 37, 66, 158, 162] and climate indices as co-variates [43, 114, 115, 150] to represent trends or seasonal cycles on annual to decadal time scales. Climate indices under consideration were the North Atlantic Oscillation (NAO), the Southern Oscillation Index (SOI), the Pacific-North America (PNA) and the El Niño-Southern Oscillation (ENSO) index. Another difference between the techniques lies in the treatment of non-stationarities. Some studies assumed that the hydro-meteorological processes can be decomposed into seasonal mean and standard deviation processes as well as a stationary process [9, 66, 162, 184], while others applied non-stationary probability distribu-

tions (i.e., with time-varying parameters) [31, 37, 43, 102, 103, 114, 115, 150, 158].

Statistical simulation models for time series of hydro-meteorological variables are now being developed for several decades and have become increasingly sophisticated. Nonetheless, accurate characterization of dependencies between hydro-meteorological variables still appears to be a challenge, especially, when not limited to joint extremal behavior. In particular, the relationship between significant wave height and peak or mean zero-crossing wave period is complex. Besides originating from a common meteorological system, there is a deterministic limit to maximum wave heights at given wave periods [for background see e.g., 77]. This so-called maximum steepness limit postulates that too steep waves break and reduce in height.

1.4. RESEARCH OBJECTIVES AND SCOPE

The overarching aim of this research is to advance risk analysis and risk management in coastal and offshore environments through problem-specific adaptation of multivariate probabilistic methods. More specifically, this thesis makes a contribution to decision support tools for risk reduction efforts in coastal environments and to statistical simulation methods for time series of hydro-meteorological variables. The research is guided by four key objectives. The first two are addressed in Part II and the last two in Part III, as described below. Key objectives will be treated in separate chapters. An exception is key objective 3, which is split into two sub-objectives treated in two chapters.

PART II: BN FRAMEWORK FOR DECISION SUPPORT ON RISK REDUCTION EFFORTS

In the past BNs have been applied numerous times as tools for decision-making under uncertainty in other fields. For instance, Henriksen et al. [72] concluded that they are very valuable for negotiations and discussions between managers, experts, stakeholders and representatives of the general public, among others, because they are transparent and flexible models. They can handle various sources and types of data enabling us to combine information on the topography and assets of the potentially affected area with simulation data of storm events and damage estimations from single discipline models. This motivates the first key objective of this thesis:

1. Design a BN framework for decision support on risk reduction efforts in coastal environments

The output will be a framework and a BN learning algorithm that is applicable to any coastal site. BNs constructed according to this framework will be able to

- predict impacts for any relevant storm event instantly,
- reflect the diversity of impacts (e.g., socio-economic and environmental aspects),
- reflect the spatial variability of impacts,
- and
- evaluate the performance of risk reduction measures under different storm events.

Once the framework is developed, BN prototypes will be built for 10 case study sites within the RISC-KIT project as part of the European Commission's 7th framework programme. The case study sites are located at different European regional seas and have

diverse geomorphic and ecological characteristics, land uses, storm climates and hazard types, as well as socio-economic and cultural aspects. This thesis will describe the application to a site in North Norfolk, UK, as an example.

Within the project, models will be built and validated through the expert knowledge of project partners. Mathematical methods for model validation and optimization will not yet be developed, but are the topic of the second key objective.

2. Indicate limitations of the BN framework and methods for improvement

The output will be a list of main limitations of the framework together with an account of how they could be remedied.

PART III: STATISTICAL SIMULATION METHODS FOR TIME SERIES OF WAVE CONDITIONS

Many engineering applications call for the generation of synthetic time series of wave conditions at high temporal resolution, including the BN framework of Part II. Approaches based on temporal vine-copulas as well as ARMA processes have shown potential for the simulation of time series of hydro-meteorological variables in several studies, as described in the previous section. Nonetheless, high temporal resolutions and the accurate description of the dependency between significant wave height and mean zero-crossing period are still under-explored in the scientific literature. This motivates the next key objective, which is split into two sub-objectives.

3. Develop a statistical simulation model for joint time series of wave conditions at hourly resolution based on (a) temporal vine-copulas and (b) ARMA processes

The output will be similar for both sub-objectives (a) and (b). The statistical simulation models will be developed for a data set in the Southern North Sea, which contains time series data of significant wave height, mean zero-crossing period and mean wave direction for 24 years with high coverage. The results will show how modeling choices have to be guided by insights from exploratory data analysis. Besides the models for serial dependence being different for sub-objective (a) and (b), the scope is unequal. The simulation model based on temporal vine-copulas focuses on bivariate time series of significant wave heights and mean zero-crossing periods during oceanographic winter periods assuming stationarity. The simulation model based on ARMA processes additionally accounts for non-stationary behavior on annual and inter-annual scale as well as a proxy for the wave direction.

4. Compare performance of the simulation models

The output will be an indication of the methods' potential for applications based on a comparison of simulation results from the two methods. Because different applications have a need for accurate reflection of different statistical properties in the time series, a comparison will be made for three of them. They are (1) the bivariate distribution of significant wave height and mean-zero crossing periods, (2) persistence of storm and calm periods, and (3) the reproduction of annual maxima. Furthermore, current limitations will be pointed out.

1.5. THESIS STRUCTURE

This doctoral thesis is divided into four parts and the structure is illustrated in Figure 1.3. Parts I and IV constitute the introduction and the conclusion. Parts II and III represent the body matter of the thesis and are built around two published articles [78, 80] and a manuscript submitted for publication [81]. Nonetheless, the thesis also contains content that has previously not been published.

Part I is the introduction, including the current Chapter 1 on the research context as well as Chapter 2, which provides the theoretical background to the probabilistic models that will be applied later in this thesis.

Part II, which consists of Chapters 3 and 4, addresses key objectives 1 and 2. This part focuses on the development of a decision support system for coastal risk management. Chapter 3 is based on [78] and develops the framework of the decision support tool using a Bayesian network. The chapter describes the design of the decision support tool and illustrates an application to a case study site in North Norfolk, UK. Chapter 4 outlines opportunities for further optimizing the performance and usefulness of the tool. To this end, a measure for model validity and an improved learning algorithm are suggested.

Part III, which consists of Chapters 5-7, addresses key objectives 3 and 4. This part explores two different methods to simulate high-frequency time series of wave conditions. Chapter 5 is based on [80] and develops a temporal vine-copula method to simulate stationary time series of hourly significant wave heights and mean zero-crossing periods during oceanographic winters. Chapter 6 is based on [81] and establishes an alternative approach based on autoregressive moving-average models. This approach is non-stationary in that it accounts for seasonalities and inter-year differences. Furthermore, the influence of the incident wave angle on significant wave heights and mean zero-crossing periods is modeled. Chapter 7 compares the two methods and discusses their potential for different applications.

Part IV contains Chapter 8 and 9 and concludes the thesis. Chapter 8 summarizes the main findings and reflects on them, while Chapter 9 suggests future research directions and provides recommendations.




Figure 1.3: Visual Outline of Thesis.

PROBABILISTIC PRELIMINARIES

The core of this research, as introduced in the previous chapter, is the problem specific adaption of selected multivariate probabilistic methods to create a decision support tool for risk reduction efforts and simulation models for time series of wave parameters as input for risk analyses. The selected methods are Bayesian networks (BNs), copulas, vines and autoregressive moving-average (ARMA) models. The purpose of this chapter is to provide the theoretical background of these methods for their adaptation, application and evaluation in Part II and III of the thesis. Section 2.1 explains the basic theory of discrete Bayesian networks, which will be used to develop a decision support tool in Part II. Section 2.2 reviews copulas and vines and Section 2.3 introduces autoregressive moving-average (ARMA) processes. The methods described in the last two sections are applied to simulating bivariate time series of significant wave height and mean zero-crossing periods in Part III.

2.1. DISCRETE BAYESIAN NETWORKS (BNs)

In this section we explain the basic theory of BNs. They represent a joint probability distribution over a set of random variables. If one or more variables are observed, a BN can evaluate the influence of this new evidence on the distributions of all other variables. If the model is “small enough”, it can predict changes in distributions instantly and can be interpreted intuitively. For this reason, BNs have been used as early warning systems for natural hazards [18, 60, 136] and as input for negotiations and discussions between experts, managers, stakeholders and citizens [72, 113, 192]. In coastal settings, such models have been shown to successfully predict erosion and shoreline retreat [48, 67, 71]. BNs have also proved to be valuable for estimating damages to residential buildings after hurricanes [172] and to evaluate the risk to nuclear facilities from coastal hazards [167].

A discrete¹ BN represents the joint probability mass function of a set of random variables $\mathbf{X} = \{X_1, \dots, X_n\}$ as a directed acyclic graph [83, 131]. Each variable constitutes a *node* in the graph. The nodes are connected by *arcs* which indicate potential dependence between variables. The direction of an arc, from so-called *parent* to *child*, signifies the direction of influence. The arcs must not form a cycle; no path $X_i \rightarrow \dots \rightarrow X_i$ may exist for any $i = 1, \dots, n$. Figure 2.1 illustrates such a graph structure.

The semantics of the graph stipulate that each X_i is conditionally independent of all predecessors given its parents. Therefore, a joint probability distribution $P(X_1, \dots, X_n)$

¹We do not consider continuous BNs in this thesis.




Figure 2.1: Example of a directed acyclic graph on six variables

can be economically factorized through the chain rule:

$$P(X_1, \dots, X_n) = \prod_{i=1}^n P(X_i \mid pa(X_i)), \quad (2.1)$$

where $pa(X_i)$ denotes the set of parent nodes of X_i . The factors $P(X_i \mid pa(X_i))$ on the right hand side of the equation are stored as *conditional probability tables* (CPTs), or in case of no parents as probability tables (PTs), and associated with each node X_i . Together, the graph semantics and all CPTs uniquely specify the joint probability mass function of $\mathbf{X}$.

A BN's computing algorithm uses Bayes' theorem. Lauritzen and Spiegelhalter [101] developed exact algorithms for high dimensions, which are implemented in most BN software. In two dimensions the theorem is given by

$$P(X_1 \mid X_2) = \frac{P(X_2 \mid X_1)P(X_1)}{P(X_2)}. \quad (2.2)$$

$P(X_2 \mid X_1)$ is the CPT of node X_2 . $P(X_1 \mid X_2)$, which is computed, is called the posterior distribution, and can be interpreted as the updated distribution of X_1 taking into account new evidence on X_2 . The CPTs and PTs can be learned from data, specified based on experts' estimates or derived from equations.

2.2. COPULAS AND VINES

A copula is a specific type of joint distribution function that fully characterizes the joint dependence between random variables, separately from their respective marginal behaviors.

Definition 1 For a n -variate distribution function F with univariate margins $F_1, \dots, F_n$, the **copula** associated with F is a distribution function $C : [0, 1]^n \rightarrow [0, 1]$ with uniform margins on $[0, 1]$ that satisfies

$$F(\mathbf{x}) = C(F_1(x_1), \dots, F_n(x_n)), \quad \mathbf{x} \in \mathbb{R}^n. \quad (2.3)$$

Theorem 1 [154]

If F is a continuous n -variate distribution function with univariate margins $F_1, \dots, F_n$ and

quantile functions $F_1^{-1}, \dots, F_n^{-1}$, then the copula

$$C(\mathbf{u}) = F(F_1^{-1}(u_1), \dots, F_n^{-1}(u_n)), \quad \mathbf{u} \in [0, 1]^n, \quad (2.4)$$

is unique.

A valid parametric model for F arises when $F_1, F_2, \dots, F_n$ and C are chosen from appropriate parametric families of distributions. In the bivariate case, $n = 2$, many parametric families have been proposed, covering a wide range of dependence structures. Joe [85] and Nelsen [129] provide comprehensive overviews.

Constructing higher dimensional families of copulas has proven to be difficult and existing models, for example, multivariate elliptical or Archimedean copulas, can be too restrictive for many applications. Montes-Iturrizaga and Heredia-Zavoni [122] discuss this with respect to environmental variables. A more flexible approach to modeling multivariate dependencies is offered by R-vines [1, 16, 36, 84]. They are graphical models consisting of a nested set of trees whose nodes in the first tree represent the random variables $X_1, \dots, X_n$. A formal definition is given below.

Definition 2 *Vine, R-vine* [100]

$\mathcal{V}$ is a **vine** on n elements if

1. $\mathcal{V} = (T_1, \dots, T_{n-1})$
2. T_1 is a connected tree with nodes $N_1 = \{1, \dots, n-1\}$, and edges E_1 ; for $i = 2, \dots, n-2$, T_i is a connected tree with nodes $N_i = E_{i-1}$.

and $\mathcal{V}$ is a **regular vine**, or **R-vine**, on n elements if additionally

- 3 (**proximity**) For $i = 2, \dots, n-2$, if $\{a, b\} \in E_i$, then $\#a\Delta b = 2$, where Δ denotes the symmetric difference. In other words if a and b are nodes of T_i connected by an edge in T_i , where $a = \{a_1, a_2\}$, $b = \{b_1, b_2\}$, then exactly one of the a_i equals one of the b_i .

A special class of R-vines that are considered in this research are drawable vines, or D-vines, for which the maximal number of edges attached to any node in the first tree is 2. The nodes reachable from a given edge in a R-vine are called the *constraint set* of that edge. When two edges are joined by an edge in tree T_i , the intersection of the respective constraint sets form the *conditioning set*. The symmetric difference of the constraint sets form the *conditioned set*. R-vines can be used to specify a joint density through a decomposition into univariate densities and (conditional) bivariate copulas:

Theorem 2 *R-Vine density* [100]

Let $\mathcal{V} = (T_1, \dots, T_{n-1})$ be an R-vine on n elements. For each edge $e(j, k) \in E_i$, $i = 1, \dots, n-1$ with conditioned set $\{j, k\}$ and conditioning set D_e , let the conditional copula and copula density be $C_{jk|D_e}$ and $c_{jk|D_e}$. Let the marginal distributions F_i , $i = 1, \dots, n$ with densities f_i , $i = 1, \dots, n$ be given. Then the vine-dependent distribution is uniquely determined and has a density given by

$$f_{1\dots n} = f_1 \cdots f_n \prod_{i=1}^n \prod_{e(j,k) \in E_i} c_{jk|D_e}(F_{j|D_e}, F_{k|D_e}). \quad (2.5)$$

For a given R-vine this density is unique. The product on the right hand side contains $n(n-1)/2$ copulas and conditional copulas, which is the exact number of ways in which n elements can be coupled. This property is one reason why a vine-copula is more flexible than a fully nested Archimedean copula. Only $n-1$ bivariate margins can be modeled distinctively with the latter, while all others are recurrent (e.g., [151], for an exemplification). Furthermore, the families of the bivariate copulas in an R-vine are not restricted to the Archimedean class.

A practical difficulty arises from the many possible different R-vine structures when attempting to estimate a suitable vine-copula, especially if the dimension is high. On n variables there are in total

$$\binom{n}{2} \times (n-2)! \times 2^{\binom{n-2}{2}} \quad (2.6)$$

labeled R-vines [123, 124]. We address the issue of selecting a suitable structure for time series in the next section.

In this thesis we make the simplifying assumption that copulas of conditional distributions do not directly depend on the conditioning variable in order to keep inference and model selection fast and robust. While [68] showed that a simplified pair copula decomposition can be a good approximation even when the assumption is far from being fulfilled, [2] illustrated that it can also be misleading. To simulate time series, we sample recursively according to well known algorithms using the inverse conditional copulas corresponding to the R-vine density decomposition in (2.5) [1, 100].

2.3. AUTOREGRESSIVE MOVING-AVERAGE (ARMA) PROCESSES

An important concept in time series modeling is weak stationarity. A stochastic process $\{X_t : t = 1, 2, 3, \dots\}$ is considered to be weakly stationary if all its moments up to the order of two do not vary in time. Thus, the mean and the variance of random variable X_t is equal to a constant and the covariance between any pair $\{X_t, X_{t+k}\}$, $\forall k \in \mathbb{N}$, only depends on k but not on t .

ARMA models can describe weakly stationary processes. For a comprehensive introduction to the topic see, for example, [25], [27] or [152]. A process X_t is called ARMA, if it can be expressed as the following function of past observations, $X_{t-1}, \dots, X_{t-p}$, and past residuals, $\epsilon_{t-1}, \dots, \epsilon_{t-q}$:

$$X_t = c + \sum_{j=1}^p \phi_j X_{t-j} + \epsilon_t + \sum_{j=1}^q \theta_j \epsilon_{t-j}, \quad (2.7)$$

where c is a constant intercept term, ϕ_j and θ_j are non-zero constants, and the residuals ϵ_t are independent and identically distributed (i.i.d.) with zero mean. If every ϕ_j is zero, the process is said to be a moving average process of order q , $MA(q)$, and if every θ_j is zero, then it is called an autoregressive process of order p , $AR(p)$.

For given orders p and q , the model parameters, ϕ_j and θ_j , can be estimated by maximum likelihood or by minimizing the conditional sum of squares of the fitted residuals. An indication for suitable orders can usually be found by inspecting the autocorrelation function (ACF) and the partial ACF (PACF). The ACF at lag k is defined as

$$\rho(k) = \text{corr}(X_{t+k}, X_t), \quad (2.8)$$

where corr denotes the product moment correlation. In contrast, the PACF measures the correlation between X_{t+k} and X_t , for $k \geq 2$, with the linear effects of $X_{t+1}, \dots, X_{t+k-1}$ removed. In order to define the PACF, let $\hat{X}_{t+k}$ denote the estimated mean from a regression of X_{t+k} on $\{X_{t+k-1}, \dots, X_{t+1}\}$ and $\hat{X}_t$ denote the estimated mean from a regression of X_t on $\{X_{t+1}, \dots, X_{t+k-1}\}$. The PACF for lag k can then be defined as:

$$\phi_{kk} = \begin{cases} \text{corr}(X_1, X_0), & k = 1 \\ \text{corr}(X_{t+k} - \hat{X}_{t+k}, X_t - \hat{X}_t), & k \geq 2 \end{cases} \quad (2.9)$$

ARMA models with different orders have distinctive ACF and PACF behaviors. The ACF of an $AR(p)$ process decays slowly, while its PACF has a cut off at lag p . Conversely, the ACF of an $MA(q)$ process has a cut off at lag q , but its PACF decays more slowly. Finally, both ACF and PACF tail off in $ARMA(p, q)$ processes and are dominated by mixtures of exponentials and damped sine waves after the first $q - p$ lags and $p - q$ lags respectively.

II

BN FRAMEWORK FOR DECISION SUPPORT ON RISK REDUCTION EFFORTS

3

FRAMEWORK DESIGN

The introduction highlighted that emergency management and long-term planning in coastal areas depend on detailed assessments (meter scale) of flood and erosion risks. For these, multiple models from different disciplines need to be linked in a so-called model train to account for various relevant physical processes and damage mechanisms. A recognized challenge is the integration of the high-dimensional model output in a clear manner, so that the information is easily accessible to and usable by decision makers.

In this chapter, a Bayesian network (BN) approach is developed to integrate the separate model results. Background on BNs was given in Section 2.1. We build on the widely recognized *source-pathway-receptor* (SPR) concept and attempt to extend and generalize the work of [136]. The BN framework is part of a suite of tools, developed in the RISC-KIT project, whose purpose is to help effective disaster risk reduction (DRR) management at coasts [170]. For different extreme events, the BN predicts percentages of affected receptors in terms of the hazards experienced and their impacts in real-time. Moreover, the BN can evaluate the effects of potential DRR measures. Although our focus is on marine storms, which are the primary threat to coastline stability, the approach is broader. It is also possible to include, or even solely concentrate on, other types of natural disasters, such as extreme river discharges or exceptional rainfall events in this model.

The chapter is organized as follows. In Section 3.1, we introduce the methodological background. We explain the SPR-concept and review how risk reduction measures have been incorporated into models. In Section 3.2, we describe the design of the BN framework, followed by examples from the case study site of Wells-next-the-Sea, Norfolk, UK, in Section 3.3. Section 3.4 describes the integration with operational forecasting systems. Finally, Section 3.5, summarizes the key points of the chapter.

3.1. METHODOLOGICAL BACKGROUND

In this section we provide an overview of models for the different elements in the risk chain, following the logic of the SPR concept, as well as an approach to quantitatively assess the effect of DRR measures. After that we describe the method we use to integrate the various models and DRR measures: BNs.

Except for section 3.4, the content of this chapters have been published within [78]. Section 3.3 is largely a contribution from co-author Elisabeth Christie and published here with her permission.

THE SOURCE-PATHWAY-RECEPTOR CONCEPT

The SPR concept is a high-level framework to evaluate risks. It was first used to describe the possible movements of a pollutant from its source to a receptor [75] and is now well established in coastal risk management [29, 53, 127, 147].

In its basic form, the framework characterizes a causal chain of processes and events in terms of sources, pathways and receptors (Figure 3.1). When considering coastal storms, the chain reaches from offshore to onshore. The *source* is the offshore marine environment. Typical source variables, or *boundary conditions*, are peak water level, maximum wave height and peak period, and storm duration. The storm threat can affect onshore areas through *pathways*. They are the interaction of water levels and waves with coastal landforms and ecosystems, coastal infrastructure and low-lying coastal hinterlands. Finally, *receptors* are the entities at risk, such as people, built environments or ecosystems.




Figure 3.1: Illustration of the source-pathway-receptor (SPR) concept for coastal storms

Sometimes, the framework explicitly includes *consequences* (C) as a fourth term (SPRC). Any receptor can experience them, if affected by a *hazard*. Gouldby and Samuels [62] have defined a hazard as the triple: source, pathway and receptor. However, we consider a hazard to be a local condition directly affecting the receptors. Examples are flood depth, flow velocity and erosion, which can, for instance, cause structural damage or injuries.

Coastal risk assessments often follow this concept. The general idea is to generate a set of representative extreme events, model the pathways, and estimate the resultant impact [e.g., 130]. More specifically, detailed and specific models are applied to various individual processes in the SPRC chain and then linked together. However, to the best of our knowledge, a single model that captures the entire chain does not exist yet.

Source Models A set of events that are representative for the storm climate at a given site can be derived from a statistical analysis. Often, storms are characterized by the values of hydraulic variables in deep water at the peak of the storm along with its duration. In the past decade, copula-based models have become increasingly popular to estimate dependencies between (some) such variables [e.g., 37, 45, 145, 184]. As discussed in Section 2.2, copula models are a specific type of probability distribution that characterizes the dependence structure between random variables irrespective of their marginal behavior. The temporal evolution of the variables, which is typically required as input for pathway models, is often idealized as a so-called *equivalent triangle* [19]. Nonetheless, a couple of studies model times series explicitly [80, 183].

Pathway Models The response of coasts to storms and the extent of flooding can be assessed with computational models, which numerically solve the physical equations that govern the motion of water and sediment in the nearshore and the hinterland.

A number of different models [e.g., 13, 73, 142, 185], varying in their numerical solutions, spatial dimensions and the range of physical processes included, have been successfully applied to the problem of hazard modeling in different coastal settings [e.g., 12, 44, 112]. The models are driven by time series of meteorological or hydraulic variables at the offshore boundary of their domain. Their output contains time series of hydraulic or morphological variables, which are potential hazards at the shore and in the hinterland, on a structured or unstructured numerical grid.

Probabilistic pathway models also exist, though they are not considered in this thesis, as they have not been used for any of the RISC-KIT case study sites. An example are levee failure models, which play an important role for flood risk assessment in the Netherlands [e.g., 87, 92, 180].

Consequence Models for Receptors Diverse and complex consequences can arise from flooding or erosion. Separate approaches can estimate economic, political, social, cultural, environmental or health-related consequences. In general, these approaches operate on the receptor level. Most commonly, they are functions, often referred to as *vulnerability relationships*, which map one or multiple hazards to consequences for a specific type of receptor. Literature reviews on vulnerability relationships exist, for instance, for economic damage [116], health impacts [4, 69], and the loss of life [94] due to flooding.

DISASTER RISK REDUCTION MEASURES IN MODELS

According to the terminology of the United Nations Office for Disaster Risk Reduction, DRR measures reduce the exposure to hazards or lessen the vulnerability of receptors [168]. For modeling, a similar categorization into three types is useful, as we explain below. *Exposure-reducing* measures move receptors out of high risk areas, for instance, by temporarily evacuating people or permanently relocating residential areas. *Pathway-obstructing* measures change the bathymetry and hence its interactions with waves and water levels. Examples are beach or dune nourishment, revetments and floodwalls. The third type are *vulnerability reducing* measures, which include for instance flood protection for individual receptors. Also raising the awareness of potential flooding and flood impacts amongst inhabitants belongs to this category.

Modeling the effect of DRR measures belonging to one of the first two types is straightforward. For exposure-reducing measures, receptors are excluded from the model or different types of land-use can be assumed for high risk areas. For pathway-obstructing measures, the pathway models can be modified and the effect simulated. Modeling vulnerability-reducing measures is more intricate. In principle, their effect is assessed by modifying consequence models [e.g., 166]. Some such measures such as early warning or awareness raising, depend on effective uptake or operation by people. To accommodate this, [41] developed a methodology to quantify and aggregate factors that influence uptake and operation.

3.2. DESIGN OF THE DECISION SUPPORT SYSTEM

OVERVIEW

The BN that we designed as a decision support system (DSS) for coastal risk management has five categories of variables: *Boundary condition* (BC), *receptor type* (R), *hazard* (H), *consequence* (C), and *DRR measure* (M). Figure 3.2 shows a high-level framework of the DSS and illustrates the influences from variables of one category onto another. If in an application each category had only one variable, then this representation would correspond to the BN graph. There are no arcs between variables of the same category. As a consequence, all boundary conditions, receptor types and DRR measures are mutually independent, while hazards and consequences are conditionally independent of each other given their parents. Admittedly, this assumption may appear unnatural for boundary conditions and DRR measures. We reflect on this issue when describing those categories later in this Section.




Figure 3.2: High-level framework of the DSS for coastal risk management illustrating the influences between boundary conditions (BC), receptors (R), hazards (H), consequences (C), and risk reduction measures (M). The solid arc indicates that all variables of the parent category influence all variables of the child category. If this is not the case, the arc is dashed.

All boundary conditions influence all hazards, which is indicated by the solid arc in Figure 3.2. In contrast, each type of receptor (e.g., people, buildings, infrastructure, and ecosystems) has a sub-module in the BN. It consists of an R node (representing the locations of receptors on the site) as well as H nodes (representing the hazards given the receptors' locations) and C nodes (representing the consequences given (some of) the receptors' hazards). The dashed arcs in Figure 3.2 represent the fact that the sub-modules are not directly inter-connected. Nevertheless, dependencies arise from the common parents, which are boundary conditions and, possibly, DRR measures.

While the high-level representation is generic, the BN is tailored to case study sites through the choice of variables and the supplied training data, based on the characteristics of the site under investigation. Owing to the generic structure, the process of constructing the BN can be automated and, if desired, integrated with open shell systems that manage forecasting processes (Section 3.4). We developed specific file formats in which variable definitions and training data need to be provided, as well as a C++ program to read them and create a BN. The source code is open and available at <https://github.com/openearth/coastal-dss>. An executable for Windows and documen-

tation is also provided. The core of our program builds on SMILE, which is a reasoning engine for graphical models. It is available free of charge for academic research and teaching use from BayesFusion, LLC, <http://www.bayesfusion.com/>.

TRAINING DATA

In essence, the training data comprises a set of storm simulations, cadastral information on the case study site, and vulnerability relationships (cf. section 3.1). Additional data may be required to include risk reduction measures, as described in Section 3.2.

The set of simulations should reflect the storm climate of the site. Each storm is defined by an offshore time series of waves and water levels, which is typically assumed to be uniform along the offshore boundary, and the simulation shows its propagation into the hinterland. For simplicity, we use statistics of the offshore time series, such as maxima and averages, to characterize the storm. In our decision support framework, these statistics are the variables of the *BC* category. Ideally, the storm events are derived from a multivariate statistical analysis of measured wave and water level time series [e.g., 136, 183] and can be a combination of historical and synthetic events. However, if adequate data or analysis tools are not available, the events can also be derived from expert opinion (e.g., with the *classical model* for structured expert judgment and extensions thereof [35, 186]).

As the *BC* category is related to input of the simulations, the *H* category is related to output. Statistics of the gridded time series of hydraulic or morphological variables in the hinterland constitute the *H* variables. Finally, the data for the *R* variables stem from the cadastral information and the data for the *C* variables are estimated from the *H* variables using vulnerability relationships. In each category, the CPTs or PTs of the variables are learned differently, as will be explained in the following section.

INCLUDING DRR MEASURES

DRR measures do not have PTs. They are implemented as decision nodes¹ and not as random nodes, because their states represent the actions of decision makers. For the same reason, measures are not interconnected with arcs. To avoid situations in which incompatible measures are simultaneously selected, they can be treated as different states of the same variable. When there are no conflicts, the variables typically have two states ("in place" and "not in place").

The three types of measures described in Section 3.1 can be incorporated into the BN. To demonstrate the outcome from exposure-reducing measures, we conceptually move receptors to a *safe zone*. To show the effect of the pathway obstructing measures, we simulate different storm events with the original and the adapted bathymetry. Finally, we modify vulnerability relationships and, if applicable, we also include a *measure effectiveness* factor, which can account for the fact that parts of the population may not receive a warning, or be able or willing to take appropriate actions (e.g., to evacuate or flood-proof their home).

¹While we include decision nodes, our model is not an influence diagram, because it does not solve for an optimal decision in terms of a maximum expected utility.

QUANTIFICATION OF (CONDITIONAL) PROBABILITY TABLES

Boundary Conditions The BN learns the PTs of BC nodes from storm simulations. The PT entries are set to the relative frequencies of observed values in all simulations. Although there may be notable dependencies between hydraulic boundary conditions (cf. Section 3.1), these are not modeled with the current version of the program; boundary conditions do not have parents. We reflect on the implications of this assumption in the next Chapter.

We use a simple *counting-learning* algorithm (e.g., [97]). Consider a data set with simulations $n = 1, \dots, N$ and a node, X_{bc} , $bc = 1, \dots, N_{bc}$, where N_{bc} is the number of boundary conditions in the BN. X_{bc} has states $s_{bc} = 1, \dots, S_{bc}$. Before learning, the PTs are uniform, $Pr^{(0)}(X_{bc} = s_{bc}) = \frac{1}{S_{bc}}$ for all s_{bc} , and the so-called *experience* is zero, $e^{(0)} = 0$. Then, for each simulation $n = 1, \dots, N$, we compute

$$e^{(n)} = e^{(n-1)} + 1, \quad (3.1)$$

$$Pr^{(n)}(X_{bc} = s_{bc}^{(n)}) = \frac{Pr^{(n-1)}(X_{bc} = s_{bc}^{(n)}) \cdot e^{(n-1)} + 1}{e^{(n)}}, \quad (3.2)$$

and

$$Pr^{(n)}(X_{bc} = \bar{s}_{bc}^{(n)}) = \frac{Pr^{(n-1)}(X_{bc} = \bar{s}_{bc}^{(n)}) \cdot e^{(n-1)}}{e^{(n)}}, \quad (3.3)$$

where $s_{bc}^{(n)}$ is the state that corresponds to simulation n and $\bar{s}_{bc}^{(n)} = 1, \dots, S_{bc}$, $\bar{s}_{bc}^{(n)} \neq s_{bc}^{(n)}$.

Receptors Each receptor type, such as building, people or infrastructure, is associated with a node that characterizes the location of a randomly selected object or individual. To this end, the case study site is divided into zones (including the external safe area, if applicable). A division is subjective and depends on features of the site; natural topological or political boundaries can be used. The zones are identical for all receptor types.

Figure 3.3 shows an illustration site with 26 houses across two zones: the beachfront and inland. We denote the node for the receptor type house as X_{house} . To understand why we consider the location of an individual house as uncertain, imagine selecting one of them randomly, just like drawing a ball from an urn. The probability for a house to be within a given zone is proportional to the total number of houses in that zone. Because 10 houses are located at the beachfront and 16 inland at the illustration site, $P(X_{house} = \text{beachfront}) = \frac{5}{13}$ and $P(X_{house} = \text{inland}) = \frac{8}{13}$. The same entries of the PT can be thought of as a spatial distribution of the houses in terms of the zones.

Note that the number of receptors in a zone could change for different combinations of exposure influencing DRR measures. In that case X_{house} would have a CPT and not a PT, because DRR measure variables would be parent nodes.

Hazards The BN learns the CPTs of hazards from storm simulations as well, however in a different manner than the PTs of boundary conditions. Now, the CPT entries represent an estimate for the proportion of affected objects or individuals, for different degrees of hazard severity. (They can also be interpreted as the probability of degree of hazard




Figure 3.3: Illustration site with 26 houses distributed across 2 zones. The color of a house indicates its degree of flooding: none is white, medium is white-blue, and high is blue.

3

severity for a randomly selected receptor.) Of course the proportions vary depending on the zone, the boundary conditions and the DRR measures in place.

We illustrate the learning with an example and refer the reader to A for the precise algorithm. Consider again the illustration in Figure 3.3. We call the combination of boundary conditions and DRR measures an *scenario*. The hazard variable, flood depth, has three states: none, medium and high. The flood depth at each house is indicated by its color: white, white-blue or blue. In the beachfront zone, $\frac{3}{10}$ houses experience a high flood depth, $\frac{6}{10}$ a medium depth, and $\frac{1}{10}$ none. These three numbers are the CPT entries given the scenario and for the zone *beachfront*. Similarly, the entries are $\frac{3}{16}$, $\frac{6}{16}$ and $\frac{7}{16}$, for the zone *inland* under the same scenario.

Because the local features of a site influence the flood flow, the degree to which a receptor is affected depends on its exact location. Without resolving to the spatial scale of individual receptors, this is reflected by the above fractions. As other types of receptors have a different spatial configuration (and total number), their fractions are different. Therefore, we model hazards separately for each receptor type.

Finally, due to the discretization, it is not possible to distinguish storms whose boundary conditions fall into the same pre-defined states. In such cases, we average the observed fractions.

Consequences Consequences have truth tables, which are a special case of CPT. In a truth table each combination of parent states corresponds to a single child state with probability 1. Hence, the relationship of consequences to their parents is deterministic. In this BN, we use vulnerability relationships, which are commonly used in the field (cf. section 3.1), to compute consequences as a function of hazards and DRR measures.

3.3. EXAMPLE OF APPLICATION

CASE STUDY SITE

The North Norfolk coast is a north-facing coastline, characterized by both gravel and sand barriers, with an extensive (>2000ha) saltmarsh area behind barrier islands, spits and areas of low gradient sand flats on open coasts. In this area the natural environment is a major source of revenue for the local economy via its contribution to nature-based tourism and recreational uses. Analysis of coastal hazards along this stretch of coast has

highlighted Wells-next-the-Sea as a risk hot spot [33].

Wells is a small coastal town (population 2165, 2011 Census). It is the largest urban center on the North Norfolk coast. The main industry of Wells is tourism, but there is also a small fishing and a wind farm servicing industry. The town has a long history of flooding due to storm surges. As such, the coastal defenses have been improved over time. The embankment forming the western side of the Wells Harbour Channel breached in the catastrophic storm of 1953 [14] and again in the storm surge of 1978 [161], leading to considerable areas being flooded on the western side of the town. Following the storm surge of 1978, this embankment was re-built to a much higher specification and, as a consequence, withstood the December 2013 surge with minimal damage [160]. In addition, the raising of the flood wall to the south of this embankment and the construction of a movable barrier between this defense and the building line at the back of the Wells Quay has reduced the risk of flooding to the low-lying western part of the town. Some individual residential and commercial properties have also implemented property level protection measures.

BN SPECIFICATION

Following the general structure introduced in Figure 3.2, the BN for Wells has boundary conditions, DRR measures, receptors, hazards and consequences. Figure 3.4 shows the graph of the BN. The boundary conditions are maximum water level and maximum wave height. There are four receptor types, each with its own hazards and consequences. For residential and commercial properties we include flood depth and absolute monetary damage. For people we assess the maximum depth-velocity product experienced and the risk to life. For the saltmarsh, we include flood depths and wave height, and damage in terms of an ecosystem vulnerability indicator (cf. Table 3.2). Finally, there are two pathway obstructing measures (extending the sea wall and increasing its height) and a vulnerability reducing measure (raising awareness through display boards).

The BN is trained with 85 storm events representing the range of potential extreme event conditions, which are generated from coastal scale modeling. They include historical storms (8 events), climate change (18 events) and synthetic events (59 events).

Boundary Conditions We defined the extreme events at Wells in terms of the peak water level and maximum significant wave height. We calculated the boundary conditions from a model train (Figure 3.5), which transforms coarse external model conditions to detailed model results at the BN boundary condition location (Wells harbour channel entrance, 52.993N 0.853E) and generates the input conditions for the model train to calculate the subsequent flooding.

The model train uses predicted hindcast data at 12 locations from the CS3 (before 2007) and CS3X (post 2007) tidal surge model run by the National Oceanography Centre, UK (NOC). The model predicts hourly tide and surge residual data with a resolution of approximately 12km which are used to drive a 2D TELEMAC model [73]. The TELEMAC model extends 50km offshore from the case study site, over an area of 4750km^2 . The unstructured grid consists of 10,704 elements with 5759 nodes. The largest elements have a resolution of 12km but the grid resolution becomes finer towards the coast with




Figure 3.4: Graph of the Bayesian Network for Wells-next-the-Sea




Figure 3.5: Framework of the Model train to assess the flood hazard in Wells-next-the-Sea, North Norfolk, UK.

the smallest resolutions in the Wells Harbour channel of approximately 15m.

The wave conditions are generated at the entrance to the harbour channel by 2 nested SWAN models [23]. The nested SWAN models are driven by 2D spectra from the UK Met

Office Wavewatch III (WWIII) North Atlantic European model with a resolution of approximately 12km. The Met Office model has a spectral resolution of 25 frequency bins and 24 directional bins. It is forced by a 10m wind field from the Met Office Numerical Weather Prediction models and the Global Wave Model. The wind characteristics are obtained from the 10m wind field from the Met Office Numerical Weather Prediction model. The water level and flow velocities are obtained from the TELEMAC model and interpolated onto the the nested SWAN model grids. The largest of those grids covers an area of 90 x 60km with a resolution of 5km, the smaller grid covers an area of 32 x 13km with a resolution of 500m. Wave growth from wind, whitecapping, quadruplets, breaking waves, diffraction, triad interactions and setup are included in the models using the present default settings (SWAN version 41.01 [22]).

The historic storm events were generated by external model hindcast storms selected where boundary conditions were available from external models. In total eight historic storm surge events were selected: 20-21 February 1993, 10 January 1995, 19-20 January 1996, 14 December 2003, 31 October-3 November 2006, 17-20 March 2007, 7-10 November 2007, and 4-7 December 2013. Synthetic storm events were generated to produce a fuller range of potential storms based on historic surge and wave conditions. The synthetic storms were generated using a typical spring tide, with length of 55 hours, with the addition of a storm surge residual from one of eight historic storm surges. The peak storm surge residual coincided with the peak tidal water level. To create a greater range of synthetic storms the spring tide or storm surge residuals were multiplied by a factor ranging from 0.7 to 1.3. Analysis of the historic surge events found no clear dependencies between the surge and wave characteristics. At this location it is clear that it is the high water levels, usually caused by the coincidence of a surge event with high water on a high spring tide, which generates the flooding. In order to create representative wave conditions which are likely to occur with a surge, the synthetic storm wave and wind conditions were taken from the historic storm events. The synthetic storms cover a range of conditions up to a return period of 400 years, based on peak water level. Only those synthetic storms which had an impact are included in the BN.

Climate change events were generated by modifying the boundary conditions from the historical storm event hindcast model to include a sea level rise prediction. The climate change event sea level and surge data were based on the IPCC (2013) RCP8.5 (Representative Concentration Pathway) projections for 2060. The future extreme storm surge levels (SSL) along the European coasts have been predicted by [179], using a DELFT3D-Flow model forced by an 8 member climate model ensemble. The authors predict the RCP8.5 2060 relative sea level rise (RSLR) at the case study site (52.98N 1.228E) to be +39cm and the predicted change in surge height in 2060 for a 20 year return period is -4cm, giving an overall climate effect of +35cm.

Receptors We divided the case study site into 6 zones, based on topographic features and key current flood prevention measures, such as the flood wall and movable flood barrier. The receptor types are residential property, commercial property, people and saltmarsh. Their locations are shown in Figure 3.6. We did not include caravans as receptors themselves, but we did include the people living inside the caravans in the analysis. We considered people to be inside their homes/caravans at the time of the flooding

impact. The model assumed a ratio of 3:1 for the number of people in houses to caravans.




Figure 3.6: Defined zones in Wells-next-the-Sea and receptor locations

Hazards We obtained detailed hazard values (inundation depth, flow velocity, wave height) by modeling the flood inundation in the study area and by mapping the receptor locations at their nearest grid point.

A small SWAN model (4 x 4km with a resolution of 15m) calculates the wave height within the harbor channel. It includes the wave energy dissipation due to the saltmarsh vegetation adjacent to the Wells Harbor (plant height=0.11m, field data from Stiffkey, North Norfolk [121]; plant diameter=0.00125m [120], plant density=1061 plants per m^2 , field data at Tillingham, Essex [119]). From the SWAN model we obtained water level and wave conditions at 64 locations along the Wells Harbor Channel. We calculated overtopping rates at transects at these 64 locations using EurOtop [138] and used the resulting overtopping discharge to drive a flood inundation model for the town. The inundation model covers an area of $7.5km^2$, the whole of Wells-next-the-Sea and some of the surrounding countryside. The landward margin of the flood grid is defined by the 10m Ordinance Datum Newlyn (ODN) contour. The majority of the grid has an approximately 12.5m resolution. A finer resolution is needed to represent the current flood wall and DRR measures in the mesh; therefore a resolution of 0.5m is used at the wall.

Consequences We determined the impact of the flooding with vulnerability relationships. For residential properties, we estimated the potential absolute damage from the inundation depth. To do that, we applied UK depth-damage curves for semi-detached houses [132], the most common property type in Wells. The residential properties are divided into those with flood protection and those without protection. From a site and Google Street View survey only 17 of 960 houses below 10m ODN were found to have any form of property level protection measures in place. We assumed that if flood resistance measures are present the house had up to date protection to the current standards. UK industry standards assume that resistance measures protect up to 0.6m above the threshold of the house [132]. The depth-damage curves were therefore modified for properties with flood protection, so that no damage occurred below 0.6m.

Commercial properties are also divided into those that have flood resistance measures and those that are unprotected (surveys found 10 properties with flood protection out of a total of 140 surveyed). Commercial properties within the case study site are typically small shops, restaurants, pubs, cafés and small warehouses. We assumed that the depth-damage curve data for retail adequately represents the commercial properties in the case study site. We obtained the depth-damage curve data from [132] for retail properties with no cellar, a short duration flood and a mean area of $140.45m^2$. We modified the curve following the same principle as for the residential properties with flood resistance measures.

As is common in the field, we calculated risk to life through the matrix developed by [137], which is based on the depth velocity product experienced and the vulnerability of the area (Table 3.1). Vulnerability of the area is based on the type of buildings and the construction methods at three levels: low vulnerability is applied to masonry, concrete and brick buildings; medium vulnerability is applied to mixed building types and high vulnerability is applied to caravans, campsites, bungalows and poorly constructed buildings. We assumed that the area is largely medium vulnerability, as it is a typical residential area with mixed type of properties. An exception is zone 1 to the west of the earthen embankment at the edge of the Wells Channel, which we defined as a high vulnerability area due to the large number of caravans (568 caravans in zone 1).

We estimated the damage to the saltmarsh by an ecosystem vulnerability indicator (Table 3.2), which relates damage to inundation depth and maximum wave height. We developed the saltmarsh vulnerability indicator from the work of [189] and [120]. We assumed the saltmarsh at the case study site is an open coast marsh in a mesotidal area. For inclusion in the BN, the saltmarsh was divided into 372 individual units, based on the grid node locations of the TELEMAC model.

DRR Measures We tested two structural DRR measures: i) an extended flood wall and ii) increasing the height of the flood wall in combination with a movable barrier. Figure 3.7 shows the location of the current structural and natural coastal protection in Wells-next-the-Sea as well as the possible flood wall extension.

We ran modified versions of the TELEMAC flood model to assess the change in the hazard experienced by the receptors. The extended flood wall lengthens the existing wall along the front of the harbor quay, protecting the properties in Area 3. The higher flood wall was chosen as water levels for the recent 2013 storm surge event almost reached the

Depth Velocity product $m^2 s^{-1}$	Nature of the Area		
	Low Vulnerability	Medium Vulnerability	High vulnerability
<0.25	Low risk	Low risk	Low risk
0.25-0.50	Low risk	Medium risk	Medium risk
0.50-1.10	Medium risk	Medium risk	High risk
1.10-7	Medium risk	High risk	Extreme risk
>7	Extreme risk	Extreme risk	Extreme risk

Table 3.1: Risk to Life matrix [137], relating the depth-velocity product experienced during a flood event and the area vulnerability to the risk to life.

Depth (m)	Wave Height (m)				
	<0.3	0.3-0.6	0.6-1	1-2	>2
0 to 1	0	2	2	3	3
1 to 2	0	1	2	2	3
2 to 3	0	1	1	2	2
3 to 4	0	0	1	1	2

0: No effect

1: Changes within normal seasonal variation

2: Changes beyond normal seasonal variation,
but partial/total recovery

3: Irreversible change

Table 3.2: Saltmarsh Ecosystem Vulnerability Indicator [175]

top of the existing defenses.

Additionally, we included iii) a vulnerability influencing DRR measure in the BN. This takes the form of a series of display boards containing cartographic information on former shoreline positions to demonstrate coastal dynamics, and images of flood markers along the coast to indicate elevations reached by historic storm surges. The aim of the display boards is to increase awareness of flood risk, with a hoped-for increase in property level protection.

To determine the effectiveness of the display boards we employed the method of [41], which assumes that the effectiveness of DRR measure depends on three factors: uptake, operator and performance. The values for each of these factors are displayed in Table 3.3.

- Uptake is the percentage of the population who will adopt the measure. We calculated a value of 9% based on 50% of the population will see the display board, 30% of the population will want to take action as the threat and coping appraisal level are high and 60% of the population can afford the measures due to an above average income level ($9\% = 50\% \cdot 30\% \cdot 60\%$).
- The operator factor is the percentage of the population that will operate the measure before a flood, assumed as 77.3%. The Environment Agency [50] gives an operator value of 86.1%, which was adjusted down as 24% of the population within the North Norfolk Area of Natural Beauty (AONB) are second home owners. We assumed that a third of the second homes would have someone available to operate the flood protection (due to the probability of an event occurring at a weekend or holiday when second home owners are more likely to be in residence). This leads an operator factor of 72.3% ($= 86.1\% - 86.1\% \cdot 24\% \cdot (100\% - 33.3\%)$). However, with the use of the display boards it is expected that more of the population will be reminded to check they know how to operate the measures they have. Therefore, post DRR operation measure of 77.3% is assumed.




Figure 3.7: Current structural and natural coastal protection in Wells-next-the-Sea and possible flood wall extension.

- The performance factor is the percentage of the population who will operate the measure effectively. We calculated a value of 73% from the assumptions that 95% of equipment are not lost or misplaced, 95% of the product are in good working order, 90% of the population receive a flood warning, and that there is a 90% chance that flood heights do not exceed flood protection level [82]. ($73\% = 95\% \cdot 95\% \cdot 90\% \cdot 90\%$)

By multiplying uptake, operator and performance factors, we estimate an increase of property level protection with the display board DRR of 5.08%. While it is possible to include the increase in protection as a measure effectiveness node, in this instance, we have directly modified the vulnerability relationships accordingly.

Influencing Factors	Post DRR measure
Uptake: % of the population who will adopt the measure	9%
Operator: % of the population that will operate the measure before a flood	77.3%
Performance: % of the population who will operate the measure effectively	73%

Table 3.3: Influencing factors for the effectiveness of the display board DRR measure

Results We could predict hazards and impacts in Wells with the BN in real-time without the need for further detailed flood modeling, as the BN is trained with a range of potential storms. The BN also allowed the impact and cost-effectiveness of the DRR

measures to be studied. The quantified model with its prior probabilities can be seen in Figure B.1 in Appendix B. In this section, we describe the BN results for a storm of 4.41 m water level and 2.17 m significant wave height. This storm was calculated as the 1 in 100 year return period storm based on the maximum water level only. For reference, the BN constraint on these bins is given in Figure B.2 in Appendix B.

The majority, 93.55%, of residential properties had no flood inundation and therefore sustained no damages (Figure 3.8a). Maximum inundation depths with an associated absolute damage of £30,000-55,000 reached 0.49% of the properties. The higher sea wall had little effect on the absolute damage distribution (Figure 3.8c). The extended sea wall shifted the distribution of damage to lower values (Figure 3.8b). Display boards can be applied together with any of the structural measure options. Whilst the display boards had little influence on the largest damage bins, due to the property level protection only protecting up to 0.6m, the distribution of damage shifted to lower levels below 0.6m (Figure 3.8d, 3.8e, 3.8f).

Commercial properties were at a greater risk in this event with 36.6% of commercial properties experiencing some flooding with no DRR measures in place (Figure 3.9a). 5.91% of commercial properties were inundated by 1-3m resulting in absolute damage of £120,000 to 200,000. Again, the higher sea wall, showed little damage reduction, suggesting that this is not a suitable measure for this storm event (Figure 3.9c). The extended sea wall reduced the number of commercial properties experiencing the highest levels of absolute damage (Figure 3.9b). The Display Board showed a shift to lower levels of damage for all the events tested. However, they had no influence on the higher damage levels due to the flood protection height limit (Figure 3.9d, 3.9e, 3.9f).




Figure 3.8: Absolute Residential Damage values for the 1 in 100 year return period storm with a) No DRR measures, b) Extended sea wall c) Higher sea wall d) Display boards e) Display boards and extended sea wall and f) display boards and higher sea wall




Figure 3.9: Absolute Commercial Damage values for the 1 in 100 year return period storm with a) No DRR measures, b) Extended sea wall c) Higher sea wall d) Display boards e) Display boards and extended sea wall and f) Display boards and higher sea wall

The expected value of any of the shown conditional damage distributions represents the total absolute damage corresponding to the selected 1 in 100 year return period storm and under the considered DRR measure. Table 3.4 shows these values for absolute residential and commercial damage for different DRR measures. For this storm event, a combination of display boards and an extended sea wall would reduce the damage most.

Table 3.4: Total absolute damages in 1000 GBP for the 1 in 100 year storm with a) No DRR measures, b) Extended sea wall c) Higher sea wall d) Display boards e) Display boards and extended sea wall and f) Display boards and higher sea wall

Type of damage	DRR measure					
	a)	b)	c)	d)	e)	f)
Residential	126.25	97.15	126.63	113.35	88.65	113.2
Commercial	2948.55	2567.45	2948.7	2688.15	2213.15	2677.9

97.6% of the people were not at risk to life, 2.3% at a low risk and 0.1% at a medium risk. The risk to life changed through the hazard influencing DRR measures, as the measures modified the maximum depth-velocity product in Wells. With an extended sea wall, risk to life was reduced compared to the no DRR case. 98.3% of the population was at no risk and 1.6% at a low risk. The higher sea wall had negligible influence on the risk to life. 97.5% of the population was at no risk, 2.3% at low and 0.1% at medium risk.

The saltmarsh was not influenced by the DRR measures due to its location seaward. For this 1 in 100 year storm, 98.4% of saltmarsh was undamaged, while 1.54% of the saltmarsh experienced changes within the normal seasonal variation of saltmarsh con-

dition.

We observed similar effects of the DRR measures across all the storm events of the BN. The higher sea wall showed a very similar distribution of damage to the no DRR measure option for all the storm events, suggesting that the higher sea wall is not a suitable DRR measure at this location. The extended sea wall led to a slight reduction in damage to residential properties, a slight reduction in risk to life and a larger reduction in damage to commercial property. The display board DRR measure showed a general reduction in commercial and property damage for those properties with low inundation depth, however, at large flood depths this DRR measure had no effect.

3.4. INTEGRATION INTO OPERATIONAL FORECASTING SYSTEMS

Delft-FEWS is a generic tool to set-up forecasting and early warning systems, originally developed for inland flooding, which has been adapted to accommodate applications in coastal areas during the RISC-KIT project [20]. Essentially, Delft-FEWS is an open data handling platform that can connect different models by automatizing the data transfer between them. It defines and executes a sequence of models through which a data set passes.

Figure 3.10 shows a typical configuration for a coastal application with the final model being the Bayesian network tool. The sequence starts with numerical predictions of global weather and waves. Then, it nests different models for water levels, currents and waves to downscale the predictions from a large offshore domain with a coarse resolution to a small nearshore domain with high resolution. Next, a hydrodynamic, and possibly morphodynamic, model computes the hazards at the coast and in the hinterland. Finally, the data is transferred to the Bayesian network tool to estimate the consequences. Figure 3.5 is a specific example for such a configuration for the case study site North Norfolk.




Figure 3.10: Typical model sequence in a coastal FEWS (adapted from: [20])

Data is not directly transferred from one model to another, but imported and exported by Delft-FEWS with so-called model adapters. This is necessary to accommodate the specific data formats required by each of the models. The model adapter for the BN computes all conditional probability tables from Delft-FEWS-exported data and updates them as a new complete storm simulations become available. However, conditioning of the BN always takes place in the software GeNIe. For instance, when nearshore boundary conditions for an impending storm have been forecast, but the nearshore model is still computing, the BN can provide useful estimation of hazards and impacts.

3.5. KEY POINTS

We developed a BN approach to support decision making in coastal risk management. An important contribution is the CPT learning algorithm for the BN, which integrates output from storm simulations with land use data, vulnerability relationships and DRR measures. The algorithm is programmed in C++ and openly available at

<https://github.com/openearth/coastal-dss> .

3

We described the application to a small town in North Norfolk, UK, which is a risk hot spot during coastal storms. The case study demonstrates how information flows through the BN and how it can predict onshore hazards and impacts, when provided with evidence of the offshore boundary conditions of a storm event. Because detailed data of severe storms including their hazards and impacts are almost never available, the BN learns from a database of hindcast and synthetic events. Each event in the database is simulated with a 2D physics-based numerical model that covers the hot spot area. This way, the training data set captures the dependence between boundary conditions and various hazards, such as erosion and flood velocity, and reflects the complex influence of the local bathymetry. At this stage, we take the precise locations of receptors into account. We include the dependence between hazards and impacts via vulnerability relationships, such as depth-damage-curves. Finally, we can also incorporate DRR measures. If they are structural, they are added through additional simulations with altered bathymetry. Otherwise, we assign them modified spatial distributions of receptors or vulnerability relationships.

The resulting BN forms a comprehensive and concise representation of risk propagation in a complex system of physical, ecological and human components. It can also be integrated in operational forecasting systems. From a practical point of view, this integrative character, together with the capability to predict in real-time, makes the BN a helpful tool for decision makers. From a scientific point of view, the model development approach emphasizes how results from multiple disciplines must be connected in order to understand risks and can provide an objective basis for choosing DRR measures.

4

RECOMMENDED EXTENSIONS

The previous chapter developed a BN framework for decision support on risk reduction efforts in coastal areas and presented an application to a site in North Norfolk, UK. The BN enables a comprehensive analysis of storm impacts on coastal communities by integrating hydrodynamic, morphological, socio-economic and environmental data. The BN can also inform on the estimated effects of DRR measures, accounting for human behavior before and during extreme events.

Within the RISC-KIT project, additional applications have been made to nine sites at the different European regional seas, which have diverse geomorphic and ecological characteristics, land uses, storm climates and hazard types, as well as socio-economic and cultural aspects [56]. Detailed descriptions of the application have been published for coastal sites in Portugal, Spain, Italy, Bulgaria and Belgium [21, 135, 146, 169].

During the project, the BN models were built and validated through the expert knowledge of project partners from each case study site and mathematical methods for model optimization have not yet been developed. This chapter discusses how the BN framework could be further advanced. In particular, we reflect on current limitations and elaborate on possible model optimization and validation.

Section 4.1 explains why the updated distributions of hazards and consequences conditional on a set of boundary conditions can be interpreted as ensemble mean predictions of affected receptor proportions. Section 4.2 proposes a metric for model validation. The metric assesses prediction performance for forward reasoning, which is the primary mode of use of the BNs. For the time being, backward reasoning (i.e., conditioning on hazards or impact and inferring storm characteristics) is not considered. Section 4.3 briefly outlines how this metric could be used to represent uncertainties in training data.

Then the focus will be on model optimization. In general, the performance of a discrete BN depends on three main factors: (1) quality of training data, (2) model structure, and (3) node discretization. Section 4.4 discusses the first and the second factor. Section 4.5 proposes a discretization strategy for receptor nodes and Section 4.6 proposes a discretization strategy for boundary condition nodes. Instead of also developing discretization strategies for hazard and consequence nodes, we recommend to elicit them from the model users. Therefore, there is no section on discretization for these node categories. Also note that DRR measure nodes do not need to be discretized, since they constitute actions of decision makers rather than continuous random variables. Finally, Section 4.7 summarizes the key points of this chapter.

4.1. USE OF THE BN AS AN ENSEMBLE MEAN PREDICTOR¹

A main feature of the BNs is the immediate prediction of receptor proportions that experience hazards and consequences (to different degrees of severity) in distinct spatial zones of a site when imposing evidence for the offshore boundary conditions of a storm event. This property is crucial for operational forecasting and early warning systems (cf. Section 3.4). Many of them rely on ensemble forecasting to account for (some of the) uncertainties in initial conditions and model formulations of numerical weather predictions as well as regional hydrodynamics and wave models (e.g., see references in [7]). Hence, the boundary conditions that the BN should evaluate in the hot phase could be an ensemble instead of a point forecast. The BN could deal with such forecasts.

We consider the simple case of a BN with one boundary condition variable influencing one hazard variable. Given a regular point estimate for the boundary condition value, the BN can be updated to present the associated conditional hazard distribution. We will refer to it as the ‘predicted hazard distribution’ and interpret it as an estimate of the proportions of receptors that are affected to different degrees.

Since boundary condition nodes are discretized representations of continuous variables, the predicted hazard distribution will not correspond directly to the one of the point estimate of the boundary condition. Rather will it be the mean of hazard distributions of all training storms whose boundary condition value falls into the same bin as the point estimate. Similarly, an ensemble of boundary conditions would link to an ensemble of hazard distributions. In the BN, we can set the distribution of the boundary condition node so that it represents the ensemble members. The predicted hazard distribution would then be an estimate of the ensemble mean hazard distribution.

In principle, the ensemble mean can be estimated for multiple boundary conditions as well, but we might introduce error, because they are modeled as independent in the BN. This model assumption is generally unproblematic for predictions based on fixed values of all boundary conditions. However, neglecting those dependencies can limit the BN’s applicability for other purposes. In particular, this simplification can introduce error to the estimates of unconditional hazard and consequence probabilities, as the BN evaluates the law of total probability.

Accurate estimates of unconditional hazard and consequence probabilities would be necessary for risk assessments and to find the most economical suite of DRR measures [e.g., 89, 190]. The situation is similar for hazard and consequence probabilities that are computed for updated distributions of boundary conditions, for instance to represent an ensemble forecast. Finally, the assumption could also have effects on the diagnostic reasoning (i.e., in opposite arc direction) of the BN, because then the unconditional hazard and impact probabilities appear in the denominator of Bayes’ theorem.

To expand the BN’s potential for wider applications, the graph structure dictated by the framework should be adapted so that it can represent the multivariate distribution of boundary conditions at case study sites more realistically. We also recommend, to investigate the influence of different discretizations on the dependence structure and, if necessary, to optimize it per site. Of course, a prerequisite for these extensions would

¹The content of this section has been previously published as part of [78]. Minor modifications have been made the text consistent with other sections of this chapter.

be a multivariate statistical analysis of available boundary condition measurements or hindcast for the location of interest, for instance with copula-based approaches [45, 184].

4.2. A METRIC FOR BN VALIDATION

Two important questions have not yet been answered with formal methods. The first question is ‘How precisely do the BNs built (according to the prescribed framework from Chapter 3) describe the training data?’ and the second one is ‘How precisely would the models predict new data?’. In other words, ‘How good is the fit?’ and ‘Did we over-fit?’. To be able to answer these question, a metric that measures prediction performance needs to be defined, which is the aim of this section.

Primarily, we designed the BN to predict hazards and impacts for given storm events. Therefore, we focus, for the time being, on a validation measure in terms of *BN precision* in predictions, though the BN can also be used to infer likely boundary conditions for given hazards and impacts. The objective is to assess how precisely the model predicts in practice, when it is given an independent case that has not been used for training. Especially, in the operational early warning mode when decision makers may need to solely depend on the estimates of the BN, it is important to understand the associated uncertainties.

Earlier studies on BN applications in coastal science either computed log-likelihood ratios [134, 136, 172] or a so-called skill metric [67, 134]. Neither method is suitable for BNs developed according to the proposed framework, because they evaluate the mean or mode of an updated BN distribution, obtained by conditioning on other variables, against a single observation corresponding to the conditioning values. In our case, we need to compare the updated BN distribution against a spatial distribution of observations. Because this is different to any other BN application we are aware of, we need to define a new validation metric.

Before we define the metric, we describe the particular structure of the training data and how the hazard CPTs are learned once more. To this end, we consider a simple BN with one boundary condition node and one hazard node, as displayed in Figure 4.1. We assume that there is only one type of receptor and do not consider any division into zones. The boundary condition node is *storm magnitude* with the three states *low*, *medium* and *high*. The hazard node is *flood depth* and has states *none*, *medium* and *high*.




Figure 4.1: A simple BN with one boundary node and one hazard node.

In terms of training data, each storm simulation contains one boundary condition value and as many hazard values as there are individual receptors. In fact, hazard values are obtained at multiple locations in the hinterland, while the boundary conditions are taken from a single location offshore. Instead of all individual hazard values, the pro-

portions of receptors that are affected to predefined degrees of severity are stored in the BN (cf. section 3.2). For simplicity, we assume that the BN has been “well-trained”, that is with a sufficient variety of storm simulations to capture the variability in storm climate at the study site. The BN, as conditioned in the Figure, predicts that, on average, across all high magnitude storms, 23% of the receptors would be flooded to a high degree, 46% to a medium degree, and 31% not at all. In other words, the BN functions as an ensemble mean predictor, although making simplifying assumption on the probability distribution of storm magnitude. However, the BN does not reflect how much the three fractions vary between different high-magnitude storms.

The variation constitutes the uncertainty associated with its ensemble mean prediction. How much the fractions vary depends on how different the effects are of storms that fall into the same boundary condition bin in terms of hazard. Figures 4.2a and 4.2b illustrate ensemble mean predictions that are associated with different degrees of uncertainty. This functioning as ensemble mean predictor of a spatial distribution, makes this BN different to conventional applications, where the nodes’ distributions are interpreted as an uncertain estimate for a fixed, but unknown value. This is another reason why we need to develop a new validation metric, instead of using already existing approaches from the literature.

For decision making, it is desirable to optimize the BN in a such way that the prediction uncertainties are minimized. In this section, we propose a metric to quantify the uncertainty, and, hence, help to validate the model.

Validating the BN against field data is hardly possible. Indeed, boundary conditions of past storm events have often been recorded for years, but detailed observations of hazards or consequences, are rare. Moreover, two identical storm events now and in the past could cause very different hazards and impacts, because characteristics of the affected area may have changed. For example, in 2013, cyclone Xaver did not lead to flooding in the Netherlands, although its surge levels were comparable to the disastrous 1953 event, which caused extensive flooding and resulted in 1836 casualties [109]. The advanced flood defense system, built in response to the disaster, accounts for the difference. On the other hand, land use evolves continuously. A similar flood event as in 1953 is likely to lead to far more casualties, since the population density has increased by approximately 70% since then [32]. To take such changing characteristics into account, we have to validate the BN against modeled data.

Because generating training data with computer simulations is an expensive process, a k -fold cross-validation seems useful. Nonetheless, an out-of-sample validation is preferable, if sufficient data can be made available. In a k -fold cross-validation, we partition the data into k equal-sized sets of subsamples and perform k rounds of cross validation. In each round, $k - 1$ sets of subsamples are used for training and 1 set of subsamples is used for testing the prediction accuracy. The process is repeated until every sub-sample has been used for testing exactly once. Common choices are between 2 and 10, depending on the sample size [57].

To define the metric, we denote the predicted distribution of hazard H_R , tested on storm bc , as $\hat{p}_{H_R|BC=bc}$. Again, the predicted distribution of hazard H_R refers to an ensemble mean prediction of fractions of receptors experiencing different degrees of hazard severity, as defined by the discretization of H_R . Similarly, $p_{H_R|BC=bc}$ is the simu-




Figure 4.2: Illustration of two identical ensemble mean predictions for the node *flood depth* for multiple storm simulations classified as high-magnitude. In (a) the predicted proportions of receptors that experience a given depth of flooding vary more across the simulations than in (b).

lated (true) distribution, adhering to the same discretization as the predicted one. In terms of validation, we would like to assess the distance between the predicted and simulated hazard distributions. Therefore, we can compute the maximum over- and under-prediction for each discrete state s of H_R , as

$$A_{H_R}^+(s) = \max_{1 \leq bc \leq n_{BC}} (\hat{p}_{H_R|BC=bc}(s) - p_{H_R|BC=bc}(s), 0) \quad (4.1)$$

and

$$A_{H_R}^-(s) = \max_{1 \leq bc \leq n_{BC}} (p_{H_R|BC=bc}(s) - \hat{p}_{H_R|BC=bc}(s), 0), \quad (4.2)$$

respectively. Here, n_{BC} is the number of storms in the complete data set. The maximum over- and under-prediction can be used to estimate error intervals for the prediction of any new storm $bc + i$, $i \geq 1$, which has not been part of the validation:

$$\mathcal{E}_{H_R, BC=bc+i}(s) = \left[\hat{p}_{H_R|BC=bc+i}(s) - A_{H_R}^-(s), \hat{p}_{H_R|BC=bc+i}(s) + A_{H_R}^+(s) \right]. \quad (4.3)$$

Figure 4.3a) illustrates the error band for a hypothetical prediction example with $H_R = \text{flood level}$ and $s \in [\text{none}, \text{medium}, \text{high}]$.

Computing such error bands is currently not implemented in the c++ program, but technically straightforward. However, it would be difficult to communicate all these intervals to the (intended) end users of the BN without developing a new graphical interface that displays them. A simpler alternative could be to separately maximize over- and under-prediction over all states and hazard variables so that we arrive at two single estimates. Error bands can then be calculated by replacing $A_{H_R}^+(s)$ and $A_{H_R}^-(s)$ in equation (4.3) by taking their respective maxima. As illustrated in Figures 4.3b) and c), this generally results in more conservative error bands. Note, that the error bands in (b) and (c) would be the same, if *flood depth* had a higher over- and under-prediction than any of the other hazards. The advantage is that the maximum over- and under-prediction encountered during the BN validation process could be reported without an upgraded graphical user interface.

4




Figure 4.3: Illustration of error bands for a *flood depth* prediction of a hypothetical BN: (a) individual error bands for each state of *flood depth*, (b) identical error bands for each state of *flood depth* and (c) identical error bands for each state of any hazard variable in the BN.

4.3. INCORPORATION OF TRAINING DATA UNCERTAINTIES

A general limitation of the proposed BN approach is its reliance on synthetic data. While we need modeled data to gain insight into hazards and impacts of storms when field observations are lacking, the validity of the BN depends on the validity of the underlying data-generating models. At present, we do not account for imperfections in those models, but future research should explore the sensitivity of BN predictions to errors or uncertainties in the data generation process.

If quantifiable, errors and uncertainties of the underlying models could be incorporated in the BN, for instance, by using Monte Carlo sampling, and be represented by showing error bands, as defined in equation (4.3), for both hazard and consequence distributions. To outline the approach, we use the simple BN in Figure 4.4.




Figure 4.4: A simple BN

First we consider the nearshore model and suppose that it is forced by a boundary condition value bc_1 . The model computes erroneous hazard values $\tilde{h}^i$ for each location i of interest. If the ‘true’ values h^i are given by some functional relationship $f(\tilde{h}^i, \epsilon)$, where ϵ is an error term with known probability distribution function, we can generate multiple samples $h_1^i, \dots, h_l^i$ through Monte Carlo simulation. Next, we use an equivalent approach

to account for uncertainty in vulnerability relationships. However, now we create a set of samples $c_1^i, \dots, c_m^i$ for each $h_1^i, \dots, h_l^i$. Finally, we repeat this process for all remaining boundary conditions $bc_2, \dots, bc_k$ in the storm data set. This way, we obtain $k \cdot l \cdot m$ cases to train the BN with.

The inherent problem in this approach is that f and ϵ are generally unknown for any model. A practical answer to this problem could be to apply a number of alternative models to compute the hazards and consequences of interest and to weigh their results equally. Such an approach has been used by [181] to analyze uncertainty in flood damage estimates and its potential effect on investment decisions.

If we would condition the BN on a single boundary condition value, for instance bc_1 , it would display the mean hazard distribution of the corresponding l training samples and the mean consequence distribution of the corresponding $l \cdot m$ training samples. While the BN would not reflect how much the distributions could differ from one another according to the uncertainties in the underlying models, this could be represented by showing error bands for both hazard and consequence distributions.

4.4. CHOICE OF TRAINING DATA AND MODEL STRUCTURE

This section deals with the choice of training data. As mentioned in previous Sections, training and cross-validation is largely based on synthetic storm event data, since observations are scarce. A danger is that future real storms will not be similar enough to any of the synthetic training storms for the BN to make useful predictions.

To minimize this risk, we advise to use statistical models to create synthetic storm events based on past observations of the relevant hydro-meteorological variables. In a subsequent step, these events should be validated by local experts as likely representations of possible storms at the site. In the previous chapter we called such statistical models ‘source models’ as they describe the source component of the SPRC concept (Section 2.1.1). Additionally, it is recommended to explore, if there are other events which would not be foreseeable based on past data, but could be physically possible. Lin and Emanuel [105] introduced such an approach for the modeling of extreme tropical cyclones.

An approach of statistic simulation has not been taken the BN of the case study site in North Norfolk, which was presented in the previous Chapter, but for the BN of the case study site in Ria Formosa, Portugal [135, 136]. Nonetheless, storms were not selected by simply simulating the desired number of events. This would result in proportionally more mild than severe storms, though the latter are of particular interest for decision making. Instead, a large set of 10^6 storm events has been simulated from a joint probability distribution constructed via a dependence tree with bivariate copulas. Each sample consisted of a values on significant wave height, peak period, surge level and storm duration. To train the BN with roughly as many severe as mild storms, 100 events have been selected from all simulations so that significant wave heights were distributed uniformly over their range of possible values [136]. Hence, the PTs only provide an overview of which storm events have been used for training and cannot be linked to actual occurrence probabilities of the storms.

The statistical model has only been used to create storms that represent a variety of

possible storms with realistic dependencies between the individual boundary condition (BC) variables. Nonetheless, the PTs could be modified after training to represent actual occurrence probabilities. If that were done, the BN would represent the annual expected spatial distributions of consequences, when boundary nodes are unconditioned. In a subsequent step, these could be converted into annual expected total damages, which is a common metric for overall risk, as illustrated in Section 3.3 for a 1 in 100 year storm. Hence, the joint probability of boundary condition variables must be estimated and incorporated into the BN in order to perform a formal risk assessment, for example, similar to the ones by Jonkman et al. [92] or Jongejan and Maaskant [87] for case studies in the Netherlands.

To achieve a greater accuracy in representation of actual occurrence probabilities, it is also recommended to adapt the model structure. In our case, the BN structure adheres to the decision support framework illustrated in Figure 3.2. Therefore the high level structure is fixed, while the exact implementation is case study site specific. A current limitation of this high level structure is that boundary conditions are modeled as independent of one another and hazards are represented as mutually conditionally independent given their boundary conditions.

Thus, one way to improve the framework would be to introduce arcs among boundary conditions and hazards. However, a prerequisite is a statistical analysis of the joint behavior of the boundary condition variables. Furthermore, it is difficult to provide general guidelines without such an analysis, because dependencies are likely to vary across study sites.

4.5. OPTIMIZATION OF SPATIAL ZONES

A design objective for the BN was to aggregate high-dimensional data in such a way that relevant information is easily accessible to decision makers. We recall that receptor nodes characterize the location of individual receptors by zones that have been established in advance. Discretization of receptor nodes is thus equivalent to definition of zones. A BN whose receptor nodes contain as many states as there are receptors, hence reflects each individual receptor location, is confusing and overloads model users with unnecessarily detailed information they cannot easily process. For this reason, a zoning is required. Aggregating receptor locations into zones, even though the individual locations are known, compacts the data and inevitably involves a loss of information. Hence, there is a trade-off between how clearly presented and how detailed the information is.

When conditioned on individual zones, or combinations thereof, the hazard distributions depend on the discretization of receptor nodes. However, when receptor nodes are left unconditioned, their discretization does not influence hazard predictions: Across the entire site the same proportions of receptors are affected in the same way, no matter how they have been assigned to zones. In this sense, the discretization of receptor nodes does not affect model validity in terms of BN precision, but the *spatial informativeness*.

A simple strategy to define a discretization is to ask model users or local experts to decide on the zoning. For Wells-next-the-Sea, the zones were derived based on topographic features and key current flood prevention measures, such as the flood wall and movable flood barrier. Similar strategies have been employed for other case studies

within the RISC-KIT project [11, 21, 135, 169].

Often there are natural zones, for instance due to administrative, geographical or built boundaries, that decision makers in the area are already working with. In this case, imposing a new zoning could result in a model that is not intuitive to the user. An alternative strategy could be to seek a discretization that maximizes spatial informativeness while keeping the number of zones ‘manageable’ for model users.

The purpose of the receptor nodes is to provide information on the spatial distribution of hazards and consequences. Hence, the zoning should help identify vulnerable and safe areas. To illustrate the issue, we consider the most simple case of a division into two zones, a single storm event, a single type of receptor and the two hazard categories ‘affected’ and ‘not affected’. The least informative division would be the one according to which the proportions of affected and unaffected in each zone are identical. The most informative division would be the one according to which the first zone contains all affected receptors, whereas the second zone contains all the not affected receptors.

Now we consider any given number of potential zones, n_R . Since there is also a known and fixed number of receptors and it is sufficient to define the zones in terms of which receptors they contain (instead of exact geographical coordinates), the number of ways in which zones can be defined is limited as well.² We denote the predicted distribution of the hazard H_R conditioned on zone r_i under discretization λ as $\hat{p}_{H_R|R=r_i}^{(\lambda)}$. An optimal discretization could then be found by maximizing as follows:

$$\lambda_{optimal} = \underset{\lambda}{\operatorname{argmax}} \mathcal{J}(\lambda), \quad (4.4)$$

with $\mathcal{J}$ being the informativeness criterion defined as

$$\mathcal{J}(\lambda) = \sum_{r_i=1}^{n_R} \sum_{r_j=1}^{n_R} \|\hat{p}_{H_R|R=r_j}^{(\lambda)} - \hat{p}_{H_R|R=r_i}^{(\lambda)}\|_2. \quad (4.5)$$

Essentially, the resulting discretizations group together receptors which tend to experience similar hazard intensities.

However, in practice finding an optimal discretization is more complicated than outlined above. First, different storm events likely result in different spatial distributions of hazard intensities. As a consequence the most informative division for one storm, may not be the most informative for another. A simple way to address this would be to use the ensemble mean prediction for the hazard in equation (4.5) considering all storm events. However, the affected areas might vary strongly, for example, for storm events from different directions. By doing so, the resulting discretization may not be very useful when turning to predictions of individual events. Another way to account for multiple storm events could be to condition on and sum over the storm events in equation (4.5). Both ways could be further refined, by introducing a weighing function to define the relative importance of having spatial information on more or less severely affected receptors.

Secondly, the BN might include different types of hazards, which affect different areas. For instance, flow velocities might be critical in sloped areas, whereas they are less

²Note that strictly speaking the division in zones is not a discretization problem, but a clustering problem with the constraint that only spatially adjacent receptors can be assigned to the same group.

likely to be relevant in low lying plains. The opposite holds for flood depth. Again, there are different approaches to address the issue. On one hand, one could elicit from the model user for which hazard spatial informativeness is most important and optimize the zoning according to that. On the other hand, one could add additional double sum terms for each hazard in the BN to equation (4.5). Further, weighing functions could be added not only for different hazard intensities, but also for different types of hazards.

Finally, there are likely to be several types of receptors in the BN. While one could find an optimal zoning for each one of them, we assume that working with a single definition of zones is significantly easier for model users. Similar to above, one could ask the model user for a most relevant receptor type or incorporate additional double sum terms in equation (4.5). Weighting functions could also be added.

In the above discussion, we assumed that the number of zones is fixed, for instance, according to the preference of a model user. Consequently, equation (4.5) finds a discretization that maximizes informativeness relative to the number of zones. Adding zones will likely increase informativeness. However, we would expect the effect to diminish once a certain number of zones is exceeded, because historic events have shown that there are spatial dependencies that cause groups of receptors to experience very similar hazards.

Based on this assumption, another promising strategy would be to iteratively add zones until the increase in informativeness is 'small enough', where a suitable threshold needs to be defined in advance. Developing this even further one could attempt to find an optimal division for a single receptor and hazard, and then iteratively include other receptors and hazards, while further dividing the study site.

This section outlined potential strategies to select a discretization for receptor nodes and highlighted some of their advantages and disadvantages. Nonetheless, further research is needed to refine and evaluate the strategies. This research should start by addressing remaining uncertainties about the model users' needs.

4.6. OPTIMIZATION OF BOUNDARY NODE DISCRETIZATION

In the above section, we described the discretization of receptor nodes, which affects spatial informativeness rather than BN precision. The present framework has nodes of four more categories. Of these, DRR measures nodes do not need to be discretized, as they represent distinct actions by decision makers, but boundary conditions, hazards and consequences do. We can simplify the task, by eliciting the discretization of hazard and consequence nodes from model users. They are the ones taking decisions based on model predictions and could indicate which threshold levels are relevant to that end. Hence, the task at hand is to devise a strategy for discretizing boundary condition nodes.

Competent handling of continuous variables is a key challenge in discrete BNs. The discretization of a continuous variable is the division of its range in mutually exclusive and collectively exhaustive intervals and their mapping to discrete values. Since variables in a BN interact, the optimal discretization for one node is likely to depend on the ones of other nodes. In the literature, multiple formal approaches have been suggested for learning the discretization of a BN from data (e.g., [59, 98, 99, 141]). They have in common that they search for a discretization, which maximizes a predefined performance

metric, while trying to avoid over-fitting. However, they are based on different theoretical principles and differ on the search strategy, the choice of performance metric and the approach to avoid over-fitting. As far as we can tell there are no clear recommendations on what strategy is most appropriate under given circumstances.

A more practical approach than the formal ones cited above, has been put forth by Fienen and Plant [57]. They propose to evaluate descriptive and predictive model performance through cross-validation and to choose a discretization that balances them using a visual diagnostic method. The underlying idea is that a BN should be able to accurately describe the data set it has been trained with as well as accurately predict new data. For instance, a BN whose nodes have a bin for each data point, would precisely describe the training data set, but lack predictive power. By reducing the number of bins, predictive performance can be increased, but at a loss of descriptive performance. We can easily adapt Fienen and Plant's method to our type of BN application by using the BN precision metric, defined in Section 4.2, to evaluate model performance.

Another question is how to deal with the interactions of the discretizations of various BC nodes, meaning that the discretization of one influences the optimal discretization of another (see e.g., [59]). Hence, we should search the space of possible discretizations of all the boundary condition variables at once. The number of possible discretizations is finite in our case, since values of variables are provided with a finite resolution (e.g., significant wave height data are usually available with a resolution of 1 cm). A wider resolution could also be chosen, for instance, based on the argument that differences below a certain threshold are not of practical relevance. A disadvantage is that this strategy can be computationally intensive depending on the number of BC nodes. On the other hand, one could iteratively discretize of one BC variable at a time, while keeping the discretizations of all others, until the discretization of any variable cannot be improved anymore. Whereas this approach is less computationally intensive, it does not guarantee to find a global optimum discretization.

Less computationally intensive strategies to search for discretizations of all interacting variables simultaneously have been proposed. These strategies reduce computational time by only considering a subset of the space of possible discretizations. Hence, their number decreases. For example, Kozlov and Koller [98] suggested an algorithm for non-uniform discretization based on so-called binary split trees. The concept is to iteratively split bins of any variable in half. For our application, the discretization would be executed as follows. At the beginning all BC nodes would have a single bin. Then, one of the bins is split in half. Here, the split, of all possible ones, which increases predictive BN precision most is chosen. Next, a second bin is split in half according to the same rule. Note that this could result in two nodes having two equally large bins or in one node having a bin that spans half of its value range and two bins that each span a quarter of its value range. The splitting procedure is continued until predictive BN precision decreases.

If computational intensity is not a concern, the first search strategy is the best option. In the other case, we propose to further investigate the potential of the approach of iterative discretization of single variables at the time or the one based on binary split trees for the present type of BN application. Further investigations should be carried out for a case study, as real data is needed to refine and test the approach.

4.7. KEY POINTS

This chapter discussed how the BN framework for decision making on risk reduction efforts in coastal areas could be further advanced, thus indicated a number of directions for future research. However, a prerequisite for implementing and testing the recommendations is a case study with data of sufficient quality and quantity. Another crucial step is to involve the intended model users in the development process. Many components of the BN can be optimized via mathematical methods. However, these require a predefined optimization criterion and it is unclear how to choose it. This issue can only be resolved by communicating with end users and understanding their needs. Additionally, end users may not want certain aspects of the BN to be mathematically optimized. An example could be the division into zones. End users might prefer to adhere to known administrative zones, even though the BN could make better predictions for a different division of the area.

The chapter focused on the optimization of the BN via mathematical methods when used as an ensemble mean predictor for affected receptor proportions under uncertain storm boundary conditions. The fundamental recommendation is the development of a metric for the prediction performance of the BNs, which computes error bands for the predicted proportions of affected receptors through cross-validation. The metric can be used to (1) validate the BNs for application in emergency management or long-term planning (if a maximum error has been specified by the model user), (2) to represent uncertainties in the underlying training data (if these have been made available to the BN in the training phase), and (3) to optimize the discretization of boundary condition nodes. Differently, we suggested to discretize receptor nodes (i.e., the division of the area into zones) by compromising between clarity of and level of detail of spatial information.

Finally, we recommended to adapt the graph structure of the framework for each application individually so that it can represent the multivariate distribution of boundary conditions at case study sites more realistically. This will expand the BN's potential to a wider range of applications, for instance, improve ensemble forecasts and diagnostic reasoning. A prerequisite for finding a better graph structure, would be analysis of the dependencies between boundary condition variables for the location of interest (for instance with copula-based approaches [45, 184]). A multivariate probabilistic model should also be used to generate a representative set of synthetic storm events, as to be able to estimate overall risk.

III

STATISTICAL SIMULATION METHODS FOR TIME SERIES OF WAVE CONDITIONS

TEMPORAL VINE-COPULA METHOD

The previous chapter discussed the importance of selecting storm events that are representative for the case study site when building a decision support system based on the proposed Bayesian network framework. As explained in the introduction, one approach to create synthetic events is to simulated complete time series with a probabilistic model and to extract storms from the simulated series. An open problem is the accurate representation of the dependency between significant wave height and mean zero-crossing wave period. Their relationship is particularly complex, because of the limiting wave steepness condition, which causes waves that are too high for their period to break and decrease in height.

This chapter explores a temporal vine-copula approach to simulate time series of significant wave heights and corresponding mean zero-crossing periods with an hourly resolution. Background on the mathematical methods is provided in Section 2.2. Vine-copulas have recently been suggested for time series modeling in the financial field, for energy research and in the social sciences [26, 126, 155, 156]. Dependence trees, the simplest form of a vine-copula, have been applied for environmental time series before: in a preliminary study for this chapter [79] and to compute workability windows for the installation of offshore wind farms based on wind speeds and significant wave heights [102]. The dependencies between wave height and period statistics have also received attention before. Repko et al. [140] compared a physics-based model to different statistical models for analyses of extremes and Vanem [173] proposed the construction of skewed copulas by extra-parameterization.

With the temporal vine-copula approach we model only one time step explicitly, but simulate time series of arbitrary lengths by assuming stationarity and a first order Markov property. In this chapter the time series modeling is restricted to oceanographic winter periods in order to bypass seasonal effects. The approach is tested for a case study data set in the Southern North Sea. Simulated time series can, for instance, serve as input for coastal hazard models or offshore operations.

The aim of this chapter is three-fold: (1) to present an overview of vine-copula models for time series, (2) to point out useful copula families for the instantaneous relationship of significant wave height and mean zero-crossing period, which safeguard the physical maximum wave steepness, and (3) to set forth a methodology for parsimoniously representing sea waves as a stochastic process.

This chapters is based on [78]

Section 5.1 describes the data set and the preprocessing that has been carried out prior to the analysis. Section 5.2 discusses the representation of time series by vine-copulas and Section 5.3 presents the results. Section 5.4 concludes the chapter by summarizing the key points.

5.1. DATA

Hourly values of the significant wave height, H_{m0} , and the mean zero-crossing wave period, T_{m02} , are available for a period of 27 years, starting in April 1989 and ending in February 2016. The data has been collected at the Europlatform, a Dutch offshore station located 58km off the coast of Rotterdam ($52^{\circ}00'N$, $03^{\circ}17'E$) at a water depth of approximately 26.5m, which is operated by Rijkswaterstaat. The location is shown on the map in Figure 5.1.

To bypass seasonal effects, the study is limited to the wave climate during the oceanographic winter period (1 November - 15 April) in which storm events are more frequent and heavier than during the summer period, see Figure 5.2. An analysis for the summer season may be conducted similar to the one presented in this chapter.

5




Figure 5.1: Location of the offshore measuring station “Europlatform” indicated as gray circle.

Both time series exhibit strong temporal serial correlation. Figure 5.3 shows the serial correlation, in terms of Spearman’s rank correlation coefficient, and indicates that it only drops below 0.1 for a time lag of at least 200 hours for significant wave heights. A similar relationship arises for mean zero-crossing periods (Figure 5.3). Here a periodic variation is notable in the order of 12 hours. In the following analysis this variation is neglected to keep the modeling parsimonious.

While we try to model the serial correlation structure, ignoring it in the inference process can cause unwanted bias. In particular, standard statistical inference procedures require independent and identically distributed (i.i.d.) observations. To obtain (approximately) i.i.d. observations, we have sub-sampled the data. More precisely, we extracted a comparatively small number of time series fragments ($N = 506$, $T = 2$). These are sufficient for the inference of the vine-copula models under consideration, because we assume stationarity and a first order Markov property, as explained in the previous section.




Figure 5.2: Hourly (a) significant wave heights and (b) mean zero-crossing periods at the Europlatform (April 1989 - February 2016). Winters are depicted in black, summers in gray.

In the time series fragments, we denote the respective j -th sample of the t -th time point of significant wave height and mean zero-crossing period as $h_{m0j,t}$ and $t_{m02j,t}$, where $j = 1, \dots, N$ and $t = 1, T$. The N joint samples are (approximately) *i.i.d* and can be used for inference of the marginal distributions and bivariate copulas.

Because formal hypothesis testing for stationarity yielded contradictory results, we followed a diagnostic approach. QQ-plots attest that the univariate distributions do not depend on t (not shown) and the observed differences in rank correlation coefficients for the same variable pairs at different points in time are very small. Table 5.1 reports all coefficients for a time series fragment of length $T = 4$ to indicate the order of magnitude in the differences. Thus, we concluded that it is reasonable to model the processes, in a first instance, as stationary. Nevertheless, more research in this direction is desirable.

Another aspect that has to be dealt with before the data analysis, is the limited instrumental resolution, which causes the available measurements to be discrete, while the physical variables are continuous. Mainly in the case of the mean zero-crossing period this phenomenon causes a high number of ties, imposing additional challenges for the statistical analysis. To overcome these, the data has been randomized following the extensive guidelines in [145]: Suppose T_{m02j}^* is a discrete measurement of T_{m02j} . It is transformed to

$$\tilde{T}_{m02j} = T_{m02j}^* + \Delta \cdot W_j, \quad (5.1)$$

where $j = 1, \dots, N$, Δ is the approximate resolution of the measurement instrument and




Figure 5.3: Serial correlation in terms of Spearman's rank correlation coefficient for measured significant wave heights and mean zero-crossing periods

5

W_j has a uniform distribution on $[0, 1]$. Thus, T_{m02j}^* is re-sampled uniformly over the resolution interval and the resulting $\tilde{T}_{m02j}$ is assumed to be statistically equivalent to the real T_{m02j} . We did not find our results to be sensitive to the choice of Δ .

5.2. REPRESENTATION OF TIME SERIES AS VINES

In this section we consider n univariate time series jointly observed at time points $t = 1, \dots, T$, which we denote by $\{X_t^{(1)}\}, \dots, \{X_t^{(n)}\}$. In our application, $n = 2$, namely significant wave heights and mean zero-crossing periods. A time series with $T = 30$ and $n = 2$ could be represented as a vine on 60 nodes, which represent $X_1^{(1)}, \dots, X_{30}^{(1)}$ and $X_1^{(2)}, \dots, X_{30}^{(2)}$. According to equation (2.6) the number of possible R-vines on 60 nodes is incredibly high, each requiring the specification of 60 marginal distributions and 1770 pair copulas.

A significant simplification of this problem can be achieved by making stationarity and Markov assumptions. For stationary data, the marginal distributions for each variable do not change with time. In the same way, the bivariate copulas (including the conditional ones) joining any two variables are preserved in time. A Markov property of

Table 5.1: Spearman's rank correlation coefficients for time series fragments with $T = 4$ ($N = 506$)

	H_{m0_t}	T_{m02_t}	$H_{m0_{t+1}}$	$T_{m02_{t+1}}$	$H_{m0_{t+2}}$	$T_{m02_{t+2}}$	$H_{m0_{t+3}}$	$T_{m02_{t+3}}$
H_{m0_t}	1	0.80	0.99	0.81	0.98	0.80	0.95	0.79
T_{m02_t}	0.80	1	0.78	0.97	0.76	0.92	0.73	0.85
$H_{m0_{t+1}}$	0.99	0.78	1	0.80	0.99	0.80	0.98	0.79
$T_{m02_{t+1}}$	0.81	0.97	0.80	1	0.78	0.97	0.76	0.90
$H_{m0_{t+2}}$	0.97	0.76	0.99	0.78	1	0.79	0.99	0.79
$T_{m02_{t+2}}$	0.80	0.92	0.80	0.97	0.79	1	0.78	0.96
$H_{m0_{t+3}}$	0.95	0.73	0.98	0.76	0.99	0.78	1	0.79
$T_{m02_{t+3}}$	0.79	0.85	0.79	0.90	0.79	0.96	0.79	1

order k implies that all copulas with a lag length greater than k in the conditioned nodes can be set to the independence copula. As a consequence, the number of bivariate copulas that need to be estimated no longer depends on T , but on k and the vine structure. The simplest case is, of course, $k = 1$. When developing the vine-copula model for significant wave heights and mean zero-crossing periods, we limit ourselves to this case.


To achieve a desirable parsimonious model for the time series data, the vine structure should be selected so that (1) the number of unique copulas to be used is minimized, (2) the number of independence copulas in the model is maximized, and (3) the model is easily extendable to arbitrarily many time steps. The last point automatically leads to a structure that represents the flow of time in an intuitive manner. Structures that fulfill these requirements are usually not obtained when using the popular structure selection algorithm by [49], which allocates the strongest pairwise dependencies to the first tree.

Advantageous structures for stationary and Markovian multivariate time series have been proposed by Brechmann and Czado [26], Smith [156] and Nai Ruscone and Osmetti [126]. Brechmann and Czado [26] call their approach COPAR, which stands for COPula AutoRegressive model, and investigate inflation effects on industrial production, stock returns and interest rates. In COPAR serial dependence of $X_t^{(1)}$ is modeled unconditionally of $X_t^{(2)} \dots X_t^{(d)}$ and the serial dependence of $X_t^{(i)}$, $i \geq 2$, is modeled conditionally on all $X_t^{(j)}$, $j < i$. In two dimensions, this approach is most appropriate when modeling the dependence of one time series onto another. The first three trees of a modified version of the COPAR model for $n = 2$ and $T = 3$ are depicted in Figure 5.4a. In this version, serial dependence of $\{X_t^{(1)}\}$ is also modeled conditionally, in this case on $\{X_t^{(2)}\}$. We made the modification for two reasons. On one hand, there is no obvious cause and effect relationship between significant wave heights and mean zero-crossing periods. On the other hand, all non-independence copulas are now included in the first three trees yielding a clearer graphical representation. Just as for the COPAR, 5 unique copulas are needed to quantify this R-vine. These are listed in the third column of Table 5.4. It is evident that we can sample $X_{t+s}^{(1)}$ and $X_{t+s}^{(2)}$ iteratively for all $s > 1$ in the same way we sampled $X_{t+1}^{(1)}$ and $X_{t+1}^{(2)}$.


Smith [156] propose a D-vine structure of a univariate series $\mathbf{Y} = (Y_1, \dots, Y_N)$, $N = Tn$, into which the elements of a multivariate time series have been re-ordered. His article examines five dimensional time series from the Australian electricity market. Figure 5.4b shows the first three trees of this model for $n = 2$ and $T = 4$. From now on, we call this model *alternating* D-vine, because $X^{(1)}$ and $X^{(2)}$ alternate in the first tree. Similarly to the modified COPAR example, 5 unique bivariate copulas are needed to quantify it (column 4 in Table 5.2), although they differ except for $C_{X_t^{(1)}, X_t^{(2)}}$. Again, it is sufficient to create an alternating D-vine up to $t = 2$. After that, we can iteratively sample $X_{t+s}^{(1)}$ and $X_{t+s}^{(2)}$ for $s > 1$ (Appendix C).

A third alternative is given in Figure 5.4c with unique copulas reported in column 5 of Table 5.2. In the 2D case it coincides with a proposal by [126]. This structure is referred to as *branching* D-vine, as $X^{(1)}$ and $X^{(2)}$ branch out. In higher dimensions, the first tree is star-like. Note that 6 instead of 5 unique copulas are needed, with one of them being in the highest order tree. The main difference with the previous approaches is that serial correlations in each univariate series are modeled unconditionally on the other series.


While cross-series dependence is modeled conditionally on previous time points, except of course for $t = 1$. Because of the high number of independence copulas, we can again sample $X_{t+s}^{(1)}$ and $X_{t+s}^{(2)}$ iteratively, this time for $s > 2$ (Appendix C).



(a) First three trees of the modified COPAR model



(b) First three trees of the alternating D-Vine Structure



(c) Branching D-vine Structure

Figure 5.4: Examples of (truncated) vines for a bivariate stochastic process $\{X_t^{(1)}, X_t^{(2)}\}_{t=1,2,3}$. Edges of the same tree have the same line type and edges signifying serial dependence with Markov property of order $k > 1$ are depicted in gray.

In theory, the three vine structures define different decompositions for the same joint density. In other words, they are equivalent representations of a bivariate stationary stochastic process with first order Markov property. However, in practice, it may be easier to find suitable bivariate copula models for one structure than for another. This will affect the overall performance of the vine-copula, which is discussed in section 5.3.

Although Smith [156] and Brechmann and Czado [26] propose integral vine-copula estimation methods, the sea waves application relies on a sequential estimation approach, because fast algorithms are not available for some of the copula models we consider. Estimation and simulation efforts are comparable for the three vine-copulas.

Table 5.2: Overview of bivariate copulas other than the independence copula that are needed for each of the three vines to specify a stationary bivariate stochastic process $\{X_t^{(1)}, X_t^{(2)}\}_{t=1,2,\dots}$ with 1st order Markov property.

Tree	Copula	modified COPAR	Alternating D-Vine	Branching D-Vine
1	$C_{X_t^{(1)}, X_{t+1}^{(1)}}$	x		x
	$C_{X_t^{(1)}, X_t^{(2)}}$	x	x	x
	$C_{X_t^{(2)}, X_{t+1}^{(1)}}$		x	
	$C_{X_t^{(2)}, X_{t+1}^{(2)}}$			x
2	$C_{X_t^{(2)}, X_{t+1}^{(1)} X_t^{(1)}}$	x		x
	$C_{X_t^{(1)}, X_{t+1}^{(2)} X_t^{(1)}}$	x		
	$C_{X_t^{(1)}, X_{t+1}^{(2)} X_t^{(2)}}$			x
	$C_{X_t^{(2)}, X_{t+1}^{(2)} X_t^{(1)}}$		x	
	$C_{X_t^{(1)}, X_{t+1}^{(1)} X_t^{(2)}}$		x	
3 or higher	$C_{X_t^{(2)}, X_{t+1}^{(2)} X_t^{(1)}, X_{t+1}^{(1)}}$	x		
	$C_{X_t^{(1)}, X_{t+1}^{(2)} X_t^{(2)}, X_{t+1}^{(1)}}$		x	
	$C_{X_{t+1}^{(1)}, X_{t+1}^{(2)} X_t^{(1)}, X_t^{(2)}}$			x

5.3. RESULTS

UNIVARIATE MARGINS

The marginal distribution of the significant wave height is assumed to follow a Weibull distribution with shape parameter 1.78 and scale parameter 167.21. The mean zero-crossing period is approximated by a gamma distribution with shape parameter 30.08 and rate parameter 6.51. Both distributions are estimated from sub-sampled data, as described in section 5.1. Histograms and the maximum likelihood fits are displayed in Figure 5.5. The adequate fit of these distributions is attested by the QQ-plots in Figure 5.6. Moreover, we sampled 506 random numbers from the specified distributions and performed two-sample Kolmogorov-Smirnov tests. These did not reject the null hypotheses that the observed and sampled data stem from the same distributions. We performed such tests 100 times and obtained p-values ranging from 0.07 to 0.90 for significant wave heights and from 0.08 to 0.97 for mean zero-crossing periods.

DEPENDENCE STRUCTURE

To be consistent with the mathematical notation of section 5.2, we denote significant wave heights and mean zero-crossing periods now also as $X^{(1)}$ and $X^{(2)}$ (cf. Table 5.3).

We estimated the vines sequentially, that is, we estimated the bivariate copulas in lower order trees before the ones in higher order trees. To quantify the three candidate




Figure 5.5: Histograms and maximum likelihood fit of (a) significant wave heights and (b) mean zero-crossing periods




Figure 5.6: QQ-plots showing the fit of the univariate models for (a) significant wave heights and (b) mean zero-crossing periods

Table 5.3: Nomenclature for variables

Variable	Standard name	Name in this chapter
Significant wave height [cm]	H_{m0}	$X^{(1)}$
Mean period [s]	T_{m02}	$X^{(2)}$

models, 12 unique bivariate copulas need to be estimated. They are listed in the first column of Table 5.4. Inference on the bivariate copula models is based on ranks and not by using the parametric univariate distributions. They give a faithful representation of dependence, irrespective of the marginal behavior of the variables. The rank of $x_{j,t}^{(i)}$

among $x_{1,t}^{(i)}, \dots, x_{N,t}^{(i)}$ divided by $N + 1$ will be denoted as $u_{j,t}^{(i)}$. Before copula selection, we performed an independence test based on Kendall's τ . This test is implemented in the VineCopula package [148] and described in Genest and Favre [61]. If the null hypothesis of independence could not be rejected, we chose the independence copula. In general, we estimated model parameters by pseudo maximum likelihood and selected families according to the AIC criterion with the VineCopula package. As potential candidate copula models, we considered the following commonly known families, which are included in the package: Gaussian, t, Frank, Gumbel, Clayton, Joe, BB1, BB6, BB7, BB8 families and, if applicable, rotated versions thereof. All selected models as well as their parameters are listed in Table 5.4. The independence copula is represented by Π .

Table 5.4: Copula families and parameters for vines

Copula	Family	Parameters
$C_{X_t^{(1)}, X_{t+1}^{(1)}}$	t	$\rho = 0.99, \nu = 11.18$
$C_{X_t^{(1)}, X_t^{(2)}}$	skew- t	$(\rho, \delta_1, \delta_2, \nu) = (0.73, -0.85, -0.27, 30)$
$C_{X_t^{(2)}, X_{t+1}^{(1)}}$	skew- t	$(\rho, \delta_1, \delta_2, \nu) = (0.7, -0.85, -0.26, 30)$
$C_{X_t^{(2)}, X_{t+1}^{(2)}}$	t	$\rho = 0.96, \nu = 12.06$
$C_{X_t^{(2)}, X_{t+1}^{(1)} X_t^{(1)}}$	Π	—
$C_{X_t^{(1)}, X_{t+1}^{(2)} X_{t+1}^{(1)}}$	Π	—
$C_{X_t^{(1)}, X_{t+1}^{(2)} X_t^{(2)}}$	Frank	$\theta = 1.35$
$C_{X_t^{(2)}, X_{t+1}^{(2)} X_{t+1}^{(1)}}$	Gumbel	$\theta = 3.01$
$C_{X_t^{(1)}, X_{t+1}^{(1)} X_t^{(2)}}$	Survival Gumbel	$\theta = 5.21$
$C_{X_t^{(2)}, X_{t+1}^{(2)} X_t^{(1)}, X_{t+1}^{(1)}}$	Gumbel	$\theta = 3.17$
$C_{X_t^{(1)}, X_{t+1}^{(2)} X_t^{(2)}, X_{t+1}^{(1)}}$	90° rotated BB8	$(\theta_1, \theta_2) = (-1.47, -0.96)$
$C_{X_{t+1}^{(1)}, X_{t+1}^{(2)} X_t^{(1)}, X_t^{(2)}}$	t	$(\rho, \nu) = (0.36, 17.05)$

We made an exception to this procedure for $C_{X_t^{(1)}, X_t^{(2)}}$ and $C_{X_t^{(2)}, X_{t+1}^{(1)}}$, because the corresponding data exhibits a skewness feature that none of the above models can represent. The observations of $(U_t^{(1)}, U_t^{(2)})$ are shown in Figure 5.7a). The variables are positively associated. Very notable is an asymmetry with respect to the main diagonal, or skewness, and a sharp line bounding the maximum significant wave height for given mean zero-crossing period values. Considering the water depth at the measuring location, this feature of the data can be attributed to steepness-induced wave breaking ("white-capping"). The observations of $(U_t^{(2)}, U_{t+1}^{(1)})$ are visually not distinguishable from the ones of $(U_t^{(2)}, U_t^{(1)})$. For $C_{X_t^{(1)}, X_t^{(2)}}$ and $C_{X_t^{(2)}, X_{t+1}^{(1)}}$, we considered the Tawn, the skew-t and the gamma 1-factor model (GFM) families. Again, we estimated parameters by pseudo maximum likelihood. Appendix D provides more details on these families and shows samples of the respective maximum likelihood fits for $(U_t^{(1)}, U_t^{(2)})$.

The Tawn copula does not capture the sharp line which limits the wave height with respect to the period, and indicates the threshold for wave breaking (Figure 5.8a). This is an important physical process and, therefore, we did not select this family. Both, the




Figure 5.7: Empirical normalized ranks of (a) significant wave heights at t and mean zero-crossing periods at t , (b) significant wave heights at t and at $t + 1$, (c) mean zero-crossing periods at t and at $t + 1$.




Figure 5.8: Comparison of observations $(U_t^{(1)}, U_t^{(2)})$ with simulated points from (a) Tawn type 2 copula, (b) gamma-1-factor model and (c) skew-t copula. Parameters are estimated via maximum likelihood.

GFM and the skew-t family are able to capture the wave breaking (Figure 5.8b and Figure 5.8c). Besides visually comparing the simulated points from the three copulas to the data, we computed semi-correlations, in terms of Pearson's product moment correlation, for their normal scores ($N_{sim} = 10^5$). These are listed in Table 5.5. Semi-correlations are correlations of the lower and upper quadrants, ρ_N^+ and ρ_N^- , respectively. The empirical ρ_N^+ and ρ_N^- are notably different, indicating stronger upper than lower tail-dependence in the data. Because, the GFM has more symmetric semi-correlations than the skew-t copula, the latter has been selected for further analysis. A disadvantage of the skew-t family is that closed form conditional and inverse conditional copula are not available and need to be approximated numerically.

The three vines are different approaches to specifying the dependence structure between significant wave heights and mean zero-crossing periods. We investigated their relative performance on three bivariate margins, namely on $(U_t^{(1)}, U_t^{(2)})$, $(U_t^{(1)}, U_{t+1}^{(1)})$ and $(U_t^{(2)}, U_{t+1}^{(2)})$. To this end, we simulated 506 bivariate time series with $T = 100$. The scatter plots for the empirical pairs of ranks are depicted in Figure 5.7. Figure 5.7b exposes the dependence between significant wave heights at subsequent hours, while Figure 5.7c shows the dependence between mean zero-crossing periods at subsequent hours. In

Table 5.5: Lower and upper semi-correlations, ρ_N^+ and ρ_N^- , for the pair $(X_t^{(1)}, X_t^{(2)})$

Copula	ρ_N^+	ρ_N^-
Empirical	0.80	0.45
Tawn type 2	0.79	0.54
GFM	0.66	0.59
Skew-t	0.72	0.56

both cases the points are distributed closely around the main diagonal indicating a very strong positive association, which is in agreement with Figure 5.3.

Figure 5.9 shows the simulated ranks for $t = 1$ as well as $t = 99$. Results of the modified COPAR model are promising (Figure 5.9 (a) - (c)). Pairwise scatter plots are very similar to the ones of the data (Figure 5.7 (a) - (c)) and the dependence structure does not appear to change significantly over time. Moreover, considering a small time series fragment of $T = 4$, the simulated and observed rank correlation coefficients are comparable (cf. Table 5.6 and Table 5.1). The maximum difference we observe is 0.047.

Table 5.6: Spearman's rank correlation coefficients for time series fragments simulated with the modified COPAR ($T = 4$, $N_{sim} = 506$)

	$\hat{H}_{m0_t}$	$\hat{T}_{m02_t}$	$\hat{H}_{m0_{t+1}}$	$\hat{T}_{m02_{t+1}}$	$\hat{H}_{m0_{t+2}}$	$\hat{T}_{m02_{t+2}}$	$\hat{H}_{m0_{t+3}}$	$\hat{T}_{m02_{t+3}}$
$\hat{H}_{m0_t}$	1	0.79	0.99	0.76	0.98	0.76	0.97	0.77
$\hat{T}_{m02_t}$	0.79	1	0.79	0.93	0.78	0.88	0.78	0.83
$\hat{H}_{m0_{t+1}}$	0.99	0.79	1	0.78	0.99	0.78	0.98	0.78
$\hat{T}_{m02_{t+1}}$	0.76	0.93	0.78	1	0.77	0.94	0.76	0.88
$\hat{H}_{m0_{t+2}}$	0.98	0.78	0.99	0.77	1	0.78	0.99	0.78
$\hat{T}_{m02_{t+2}}$	0.76	0.88	0.78	0.94	0.78	1	0.78	0.93
$\hat{H}_{m0_{t+3}}$	0.97	0.78	0.98	0.76	0.99	0.78	1	0.79
$\hat{T}_{m02_{t+3}}$	0.77	0.83	0.78	0.88	0.78	0.93	0.79	1

The simulated ranks from the alternating D-vine are plotted in Figure 5.9 (d) - (f). This model does not capture the strong serial correlation between mean zero-crossing periods; $\rho_{U_t^{(2)}, U_{t+1}^{(2)}}$ of the simulated data is around 0.61, while we estimated a value of 0.97 for the data (Table 5.1). In the branching D-vine the skewness of the $(U_t^{(1)}, U_t^{(2)})$ pair deteriorates in time (actually already within a few steps), which is not surprising given the skew-t copula is only used to model the very first point in time (Figure 5.9 (d) - (f)).

Hence, we chose the modified COPAR structure for time series simulation of significant wave heights and mean zero-crossing periods. Copula samples can be transformed to the original (H_{m0}, T_{m02}) -space by applying the respective inverse marginal distribution functions.




Figure 5.9: Simulated points for significant wave heights at t and mean zero-crossing periods at t , significant wave heights at t and at $t+1$, and mean zero-crossing periods at t and at $t+1$: (a) - (c) modified COPAR, (d) - (f) alternating D-vine, and (g) - (i) branching D-vine.

BIVARIATE DISTRIBUTION OF SIGNIFICANT WAVE HEIGHTS AND MEAN ZERO-CROSSING PERIODS

The specification of the modified COPAR includes a specification of the bivariate distribution of significant wave height and corresponding mean zero-crossing period. Samples of this distribution are shown in Figure 5.10. A visual comparison of observed and simulated points indicates that the model is valuable. In particular, it captures the physical limitations on wave steepness due to wave breaking. Wave steepness can be computed as

$$S_{m02} = \frac{2\pi}{g} \frac{H_{m0}}{T_{m02}^2}. \quad (5.2)$$

$S_{p_{max}} = 0.07$ is thought to be an upper limit [77], and is represented as a line in Figure 5.10. A histogram of wave steepness for observations and simulated points confirms that the distribution of S_p is well described by the model (Figure 5.11). The *skew-t* family is a promising, although computationally more intensive, alternative to the skewed families constructed by extra-parameterization in [173].




Figure 5.10: Simulated and observed significant wave heights and mean zero-crossing periods




Figure 5.11: Histogram of wave steepness for simulated and observed significant wave heights and mean zero-crossing periods

SIMULATED TIME SERIES

We simulated 100 time series of 3984 hours, which is equivalent to an oceanographic winter, with the modified COPAR. However, these synthetic series contain notably fewer storm events than observed. An exploratory analysis of the root cause of this phenomenon shows that the long term serial correlation is very sensitive to the choice of the ρ parameter in the $C_{X_t^{(1)}, X_{t+1}^{(1)}}$ copula, now denoted by $\rho_{t,t+1}^{(1)}$.

We conducted a sensitivity analysis, assuming values 0.98 and 0.95 besides the estimated 0.99 for $\rho_{t,t+1}^{(1)}$. Figure 5.12 shows examples of simulated time series as well as the observations from the winter 2013/2014. The three simulated time series are based on the same random seed. They exhibit the same pattern, however, the spikes in significant wave height and mean zero-crossing period are higher for lower values of $\rho_{t,t+1}^{(1)}$. The boxplot in Figure 5.13a provides more insight into the total hours with significant wave heights higher than 400 cm, $N_{H_{m0} \geq 400}$. Recall that $N = 27$ time series have been observed and $N = 100$ have been simulated. For smaller $\rho_{t,t+1}^{(1)}$, the median increases, while the variance decreases. The reason for this model behavior is likely to lie in the faster decay of the serial correlation for smaller $\rho_{t,t+1}^{(1)}$ (Figure 5.14), increasing the likelihood of sudden, extreme changes in H_{m0} and T_{m02} . The serial correlation function of both variables is most closely approximated when using $\rho_{t,t+1}^{(1)} = 0.98$. Results for the number of hours with significant wave heights smaller than 50 cm, $N_{H_{m0} \leq 50}$, are different (Figure 5.13b). In this case the median is higher in the simulations than in observations.

To investigate whether the observed $N_{H_{m0} \geq 400}$ and $N_{H_{m0} \leq 50}$ are statistically different to the ones simulated, we conducted an analysis of variance (ANOVA) and performed a two-sample Kolmogorov-Smirnov (KS) test. At the 5%-level the Null-hypothesis of the ANOVA, the mean of observed and simulated $N_{H_{m0} \geq 400}$ or $N_{H_{m0} \leq 50}$ are equal, is not rejected for any of the values of $\rho_{t,t+1}^{(1)}$. For the KS-test with the Null-hypothesis that observed and simulated $N_{H_{m0} \geq 400}$ are equal in distribution, $\rho_{t,t+1}^{(1)} = 0.99$ is rejected with a p-value close to 0, but $\rho_{t,t+1}^{(1)} = 0.98$ and $\rho_{t,t+1}^{(1)} = 0.95$ cannot be rejected at the 5%-level. The results are different when testing the Null-hypothesis that observed and simulated $N_{H_{m0} \leq 50}$ are equal in distribution. $\rho_{t,t+1}^{(1)} = 0.99$ cannot be rejected at the 5%-level, $\rho_{t,t+1}^{(1)} = 0.98$ is rejected with a p-value of 0.04 and $\rho_{t,t+1}^{(1)} = 0.95$ is rejected with a p-value close to 0.

Considering these findings, the modified COPAR is useful for applications that require extreme wave time series, such as coastal risk analyses, as well as for applications that need estimations of quiet seas, such as scheduling of offshore operations. It is not surprising that very small differences in $\rho_{t,t+1}^{(1)}$, which could even be sample fluctuations, gain importance when simulating longer time frames, such as $T = 3984$. It is therefore essential to validate the model in view of the application and time frame of interest.




Figure 5.12: Hourly measurements of (a) significant wave height and (b) mean zero-crossing period during winter 2013/2014 and examples of simulated time series using the same random seed, but a different parameter in $C_{X_t^{(1)}, X_{t+1}^{(1)}}$ denoted as $\rho_{t,t+1}^{(1)}$.




Figure 5.13: Boxplots for (a) $N_{H_{m0} \geq 400}$ and (b) $N_{H_{m0} \leq 50}$ in the measured and simulated data. $\rho_{t,t+1}^{(1)}$ denotes one of the two parameters of $C_{X_t^{(1)}, X_{t+1}^{(1)}}$. The whiskers represent the 5th percentile and the 95th percentile of the distribution of interest.




Figure 5.14: Serial correlation of measured and simulated (a) significant wave heights and (b) mean zero-crossing periods. $\rho_{t,t+1}^{(1)}$ denotes one of the two parameters of $C_{X_t^{(1)}, X_{t+1}^{(1)}}$.

5.4. KEY POINTS

In this chapter, we presented a vine-copula approach to simulate joint time series of significant wave heights and mean zero-crossing periods. First, we reviewed existing vine-copula models for time series. While many different vine structures are possible, a few have recently been proposed that make efficient use of stationarity and Markov assumptions and yield parsimonious representations of such stochastic processes. In two dimensions, depending on the structure, no more than 5 or 6 unique copulas have to be specified to characterize arbitrarily long, stationary time series with first order Markov property.

Furthermore, as a building block for the vine-copula, we investigated the dependence structure of significant wave heights and mean zero-crossing periods at the same time. We approximated the dependence with the flexible 4-parameter skew-t copula family. It captures asymmetric tail dependence as well as skewness patterns of the data, and, as a result, preserves the limiting wave steepness condition in simulated samples of the two variables.

Finally, we found an R-vine structure similar to the COPAR introduced by Brechmann and Czado [26] to be most valuable for this application. Based on observations of two successive pairs of significant wave height and mean zero-crossing period, we constructed a vine-copula on four variables. These observations are a small i.i.d. subset of measurements (506 out of 53492), collected during the oceanographic winter period at a single location in the North Sea. We simulated time series with the duration of an oceanographic winter and compared them against all measured winters. While results are sensitive to small changes in the copula parameters, the presented model has potential to provide valuable input for various applications, such as coastal risk analysis, design of marine structures and offshore operations.

6

ARMA METHOD

The previous chapter explored a temporal vine-copula approach for time series of significant wave heights and corresponding mean zero-crossing periods with an hourly resolution. While the method could emulate storms as well as daily conditions during oceanographic winter periods, the validity of the model is highly sensitive to the selected vine-structure, the chosen copula families and the parameter estimates. Moreover, one copula needs to be from a high-parameter family, making the model complex and requiring long simulation times. In this chapter, we explore a different, less complex time series approach, based on ARMA models. To account for complex dependencies, such as the one caused by the limiting wave steepness condition, we apply additional transformations. In addition, the model includes non-stationary behavior and also demonstrates an influence of the wave direction on the time series behavior. While we reapply many techniques that have been suggested in the literature cited in Section 1.3, the method presented here distinguishes itself from others on five principle points:

1. We develop a data-driven equation for the limiting wave steepness condition at the study location and use it for an initial variable transformation. In this way, we separate the deterministic part of the relationship between significant wave height and mean zero-crossing periods, which is steepness-induced wave breaking, from the stochastic part, which is due to common meteorological and geographical factors, in our modeling.
2. Modeling of the mean wave direction is simplified by assuming a categorical variable with two possible values, north and southwest, which we refer to as wave direction regime. The assumption is reasonable given the geographical context of the measurement station. The main advantage is that we circumvent challenges related to modeling a circular variable and avoid inaccuracies that could arise from ignoring the circular aspect. Time series of the wave direction regime are modeled as a seasonal alternating renewal process, inspired by [144].
3. Instead of applying a VARMA model with joint-normally distributed residuals to the bivariate time series of significant wave heights and mean zero-crossing periods, we estimate two univariate ARMA models with a non-normal joint residual distribution constructed via a copula. Recent examples of such an approach for other types of environmental time series can be found in [51] and [17].

The content of this chapter is based on the manuscript [81], which has been submitted for publication.

4. The wave direction regime is used to trigger regime switches in the joint residual distribution to account for possible differences in the statistical characteristics of northern and south-western waves. Differences are expected, because south-western waves are mostly wind-sea due to a limited fetch length, while northern waves can be a mixture of swells and wind-seas.
5. Similar to existing studies, we use Fourier series to characterize a seasonal mean process and a standard deviation process. However, we assume that the Fourier coefficients are random variables, potentially dependent, instead of constants in order to represent inter-year differences and dependencies between the processes on yearly time scales.

The outline of the chapter is as follows. Section 6.1 introduces a data set from the measuring station Europlatform in the Dutch Southern North Sea. This data set will be used to develop and illustrate the simulation method. The section also shows how the wave direction affects the bivariate distribution of significant wave heights and mean zero-crossing periods and motivates why wave directions are clustered into two regimes. Section 6.2 develops the methodology for jointly modeling time series of the wave parameters. Section 6.3 shows simulation results. Finally, Section 6.4 contains the discussion and conclusions.

6

6.1. DATA AND REGIME DEFINITION

As in chapter 5, the data stems from the Europlatform (Figure 5.1). The data set consists of hourly measurements of three wave statistics, spectral significant wave height, mean zero-crossing period and mean wave direction, for a period of 24 years (1 Jan 1991 - 31 Dec 2015).

The existence of leap years complicates time series modeling on an hourly scale, because the number of hours per year varies. Annual seasonal processes would thus have a period that differs by 24 hours if the year is a leap year as compared to when it is not. To avoid this, we introduced a new calendar for modeling purposes, in which all years are equally long. This calendar assumes that all Februaries consist of 28 days and 6 additional hours. Thus, each year has 8766 hours, instead of non-leap years having 8760 hours and leap years having 8784 hours. Of course, days in different years start at different times of the day, but this is not relevant for our application and has no effect on the results. However, the reader should be aware that dates mentioned and displayed do not exactly correspond to actual dates and times.

The data coverage is higher than 94%; most missing values arise before 2003. The period 1 Jan 2003 - 31 Dec 2014, according to the new calendar, has only five instances in which values of the three wave statistics are jointly missing. We have filled these by linear interpolation. If a component of the simulation model required a complete time series record without missing values for parameter estimation, this shorter-length record was used and this is mentioned in the corresponding section. Otherwise, models were estimated based on the full-length record.

The observed time series for significant wave height and mean zero-crossing period are shown in Figures 6.1a and 6.1b. Most waves either originate from distinctly northern

or south-western directions (Figure 6.2). Because of the geographical characteristics of the location (cf. Figure 5.1), waves from northern directions can be swells, wind seas or a mixed sea state, while waves from south-western directions are mainly wind seas. For this reason, we expect differences in the statistical properties of northern and south-western waves.




Figure 6.1: Time series of (a) H_{m0} and (b) T_{m02} from January 2003 to December 2014. Observations are 1 hour apart.




Figure 6.2: Circular histogram of direction wave directions. Green and yellow indicate the two wave direction regimes.

Issues can arise when ignoring the circular nature of a variable. For a variable in polar coordinates 0° and 360° are identical directions. Furthermore, 358° cannot be called larger than 2° and neither is 180° their meaningful average. To analyze and model such a variable, many standard statistical tools or measures for ratio variables, such as significant wave height and mean zero-crossing period, might not be suitable. For example, it would be difficult to interpret a rank correlation of the mean wave direction with another variable. If we ignore the circular nature of the mean wave direction and compute its rank correlation with significant wave height, we obtain a value of 0.11. The calculation assumes that 0° is the smallest and 360° the largest possible direction. Now suppose that, we would have defined mean wave direction as the direction into which the waves

are traveling, contrary to the norm, which is the direction of origin. All our measurements would be shifted by 180° , but the dependence between mean wave direction and significant wave height would remain the same. However, we compute a rank correlation value of -0.21 for this case. The example shows the danger of obtaining misleading results, when neglecting the circular nature of mean wave direction.

In the present geographical context it seems natural to represent the wave direction as a categorical variable with two states in order to circumvent issues related to circularity. Hence, the data have been partitioned into two clusters representing the two main directional sectors and a new variable, the wave direction regime, is defined as follows:

$$\Theta_t = \begin{cases} 0, & \text{mean wave direction at time } t \in (304^\circ, 48^\circ), \\ 1, & \text{mean wave direction at time } t \in [48^\circ, 304^\circ]. \end{cases} \quad (6.1)$$

In the remainder of this chapter we use the standard notation for significant wave heights, H_{m0} , and mean zero-crossing periods T_{m02} . An overview is given in Table 6.1.

Table 6.1: Overview of Variables

Variable	Unit	Name	Sample space
Significant wave height	cm	H_{m0}	$\{\mathbb{R}^+\}$
Mean zero-crossing period	s	T_{m02}	$\{\mathbb{R}^+\}$
Wave direction regime	–	Θ	$\{0, 1\}$

6.2. MODEL DEVELOPMENT

Two types of probabilistic models are at the core of this simulation method: ARMA models and copulas. A short introduction to them can be found in Sections 2.2 and 2.3. The statistical simulation method developed here consists of several steps shown in the flowchart in Figure 6.3. First, directional regime time series are derived according to equation (6.1) from the mean wave direction series. Then, the durations for which the directions remain in each regime before switching to the other are represented by a seasonal renewal process. Independently thereof, a limiting wave steepness condition is estimated for the collected data and used to remove the effects of steepness-induced wave breaking. Next, the data of significant wave height and mean zero-crossing period are normalized and decomposed into stationary and non-stationary processes. The non-stationary processes are modeled using Fourier series with random coefficients, while the stationary processes are modeled as ARMA using a regime-switching joint residual distribution. The regime switches are triggered by the directional regimes.

SEASONAL MODEL FOR REGIME SWITCHES

The process $\{\Theta_t\}_t$ is modeled as an alternating binary renewal process, following [144], who described wet and dry periods of precipitation in this way. For this application, the durations for which waves are coming from one of the two directions are random variables, N and SW . For example, the initial wave direction regime is 0 and remains




Figure 6.3: Overview of the steps of the statistical simulation method for wave parameter time series

6

that for a time SW_1 . Then it switches to 1 and remains that for a time N_1 . It is then 0 for time SW_2 , and so on. Figure 6.4 illustrates the process.




Figure 6.4: Illustration of renewal process

An alternating renewal process supposes that both sequences $\{N_n\}$ and $\{SW_n\}$, $n \geq 1$, are independent and identically distributed. However, N_n and SW_n may be dependent. We made two modifications to this set up. On one hand, we allow SW_n to depend on N_{n-1} , $n \geq 2$. On the other hand, we assume that both N_n and SW_n , $n \geq 1$, depend on the time of the year as well. Thus, we estimated a bivariate distribution for (N_{n-1}, SW_n) and for (SW_n, N_n) for each season, using a decomposition into univariate distributions and a copula.

We could not find adequate parametric univariate distributions among well-known families and suspect that this is caused by the many one-valued observations (i.e., many

durations are 1 hour). However, we did not investigate if and how so-called zero-inflated probability distributions could be adapted to this problem [e.g., 193]. For simplicity, and since we do not need to extrapolate beyond the range of the observations, we used the empirical distribution functions. Figure 6.5 shows box plots of the observations for N and SW for the four seasons. Seasonal differences are more pronounced for SW .




Figure 6.5: Boxplots of the observed durations (a) N and (b) SW for the four seasons.

Bivariate copulas were selected according to the AIC criterion using the VineCopula package [148], which compares twelve different families and, if applicable, their rotated versions. The corresponding parameters were estimated by maximum likelihoods. The selected families, parameter estimates and Kendall's τ for (N_{n-1}, SW_n) , $n \geq 2$, and for (SW_n, N_n) , $n \geq 1$, are given in Table 6.2. Comparing values of Kendall's τ , both variable pairs are positively associated. Thus, a wave direction regime tends to persist longer when the duration in the preceding regime was long than when it was short. This tendency is stronger and seasonal differences more distinct for (N_{n-1}, SW_n) than for (SW_n, N_n) .

Table 6.2: Selected copula families, parameter estimates (denoted par1 and par2) and Kendall's τ for the pairs (N_{n-1}, SW_n) and (SW_n, N_n) .

Variable pair	Season	copula family	par1	par2	τ
(N_{n-1}, SW_n)	Spring	BB8	1.86	0.68	0.13
	Summer	BB8	1.52	0.81	0.11
	Autumn	Frank	1.77	–	0.19
	Winter	Survival BB8	2.74	0.59	0.19
(SW_n, N_n)	Spring	Frank	1.0.7	–	0.08
	Summer	BB8	1.41	0.85	0.1
	Autumn	Frank	0.99	–	0.11
	Winter	Frank	1.02	–	0.11

SIGNIFICANT WAVE HEIGHTS AND MEAN ZERO-CROSSING PERIODS

Model for the limiting wave steepness condition While the relationship between significant wave heights and mean zero-crossing periods is in large part stochastic, there is a physical limit on the maximum steepness that individual waves can attain. As soon as waves approach this limit, they break¹. Wave steepness is defined as wave height divided by wave length, but can be formulated as a function of wave height and wave period. In terms of significant wave height and mean zero-crossing period, it is

$$s_{m02} = \frac{2\pi}{g} \frac{H_{m0}}{T_{m02}^2}. \quad (6.2)$$

The limiting steepness condition is clearly visible in the scatter plot of H_{m0} and T_{m02} (Figure 6.6a): For a given T_{m02} the corresponding H_{m0} cannot exceed a certain upper limit, or equivalently, for a given H_{m0} the corresponding T_{m02} is bounded from below. Nonetheless, we observed a few data points that appear to be unusually distant from the others (gray crosses in Figure 6.6a). We suspect that these are anomalies in the measurements and substituted them by missing values before proceeding with the data analysis.

Recent studies showed that bivariate distributions constructed with 3- or 4-parameter copula families can be suitable to reflect the limiting steepness condition and to represent the joint distribution of H_{m0} and T_{m02} [80, 173]. However, we cannot directly apply such approaches to the ARMA modeling in this study. Instead, we separate the deterministic part and the stochastic part of the relationship between H_{m0} and T_{m02} to model them individually. The idea is to remove the deterministic part by subtracting the lower bound from T_{m02} :

$$\tilde{T}_{m02} = T_{m02} - T_{m02_{min}}, \quad (6.3)$$

where the lower bound can be determined from H_{m0} and $s_{m02_{max}}$ as

$$T_{m02_{min}} = \sqrt{\frac{2\pi}{g} \frac{H_{m0}}{s_{m02_{max}}}}. \quad (6.4)$$

¹ This is different in shallow water, where depth-induced breaking occurs before steepness-induced breaking.




Figure 6.6: Scatter plots of (a) significant wave height and mean zero-crossing period and of (b) significant wave height and steepness. Different proposed wave steepness limits are indicated in both plots.

After this transformation, we proceed with time series modeling of the pair $(H_{m0}, \tilde{T}_{m02})$.

A difficulty that arises is the choice of $s_{m02_{max}}$. Different values have been proposed in the literature. For the North Sea, Holthuijsen [77] states $s_{m02_{max}} = 0.067$, while Michel [117] derives $s_{m02_{max}} = 0.098$. However, by visual comparison neither of the two estimates seems to fit our data set. Figure 6.6b shows a bivariate scatter plot of (H_{m0}, s_{m02}) and the proposed constant values for $s_{m02_{max}}$ as purple and green lines. The corresponding lines for $T_{m02_{min}}$ are plotted in Figure 6.6a, also in purple and green. Michel's proposed value is too far away from the data, while Holthuijsen's value would cut off a notable part of the data.

To achieve a better fit, we set out to estimate a maximum steepness condition from the data. From the scatter plot in Figure 6.6b it appears that observations of $s_{m02_{max}}$ are not constant, but depend on H_{m0} . More precisely, there appears to be a horizontal asymptote roughly below $s_{m02_{max}} = 0.08$, while the observed $s_{m02_{max}}$ is rapidly decreasing for small H_{m0} . To account for this behavior we fit the curve

$$s_{m02_{max}}(H_{m0}) = a \left(\frac{H_{m0}}{b} \right)^{\left(\frac{c}{H_{m0}} \right)}, \quad a, b, c > 0 \quad (6.5)$$

to the data. The reasoning behind selecting this functional form is the following. Suppose b is a large value that cannot be attained by measurements of H_{m0} at the Europlatform. Then, a can be interpreted as the value defining the horizontal asymptote, since $s_{m02_{max}} \rightarrow a$, as $H_{m0} \rightarrow b$. Finally, c affects the slope of $s_{m02_{max}}$ for smaller values of H_{m0} .

The procedure to fit the $s_{m02_{max}}$ -curve was the following. First, we discretized the H_{m0} data into 108 discrete bins. These were not equally spaced, but contained an equal number of data points (1870). Most bins cover a range in height of 1cm or 2cm. An exception is the widest bin, which spans from 274cm to 656cm. Next, we computed the maximum value of $s_{m02_{max}}$ in each bin and associated it with the value of H_{m0} at the

bin center. These data points are shown as orange circles in Figure 6.6b. Finally, we estimated the coefficients a , b and c using nonlinear least-squares.

The resulting estimates are $a = 0.0782$, $b = 999.4\text{cm}$ (the upper bound was fixed at 1000) and $c = 7.674\text{cm}$. The coefficient of determination is $R^2 = 0.858$ and the root mean square error is $RMSE = 0.0051$. The fitted $s_{m02_{max}}$ -curve is also shown in Figure 6.6b and the corresponding $T_{m02_{min}}$ -curve is depicted in Figure 6.6a (orange lines). According to this limiting steepness condition, 160 measurements are classified as anomalies, because they are too steep. These are 67 more than initially identified by visual inspection and amount to less than 0.08% of the data. It should be noted that while this relation fits the observational data well in statistical sense, we do not propose to use this formulation to better describe the physics of wave steepness-related breaking.

After substituting the 160 measurements classified with missing values, T_{m02} is transformed to $\tilde{T}_{m02}$ (equation 6.3). Figure 6.7 shows a scatter plot of the pair $(H_{m0}, \tilde{T}_{m02})$. In the remainder, we will denote $X^{(H_{m0})} = H_{m0}$, $X^{(\tilde{T}_{m02})} = \tilde{T}_{m02}$ to simplify notation.




Figure 6.7: Scatter plot of significant wave height and mean zero-crossing period subtracted by its steepness induced lower bound.

Normalization and decomposition into non-stationary and stationary components

Neither of the processes $\{X_t^{(H_{m0})}\}_t$ and $\{X_t^{(\tilde{T}_{m02})}\}_t$ is stationary. Inspection of the data suggests strong seasonal behavior (cf. Figures 6.1a and 6.1b). Trends do not seem apparent. Moreover, the data are notably right skewed and strictly positive. Different methods can be used to analyze and model such time series data [e.g., 25]. The prevailing approach is to deseasonalize the data and to model the seasonal components and the stationary component separately.

We identified two main procedures for deseasonalization in the literature on wave parameter modeling and simulation that appeared successful. Suppose x_t , $t = 1, \dots, T$, is a time series for an arbitrary wave parameter. The first procedure involves two steps [162]. One step is to transform the data to reduce skewness:

$$y_t = f(x_t), \quad t = 1, \dots, T, \quad (6.6)$$

with f being a suitable monotone transformation function. This facilitates finding a fitting distribution for the ARMA residuals of the stationary component, but is also important for simulations, as will be explained in the next paragraph. A second step is to represent the transformed time series data as a realization of the following process:

$$y_t = \mu_t + \sigma_t z_t, \quad t = 1, \dots, T, \quad (6.7)$$

where μ_t and σ_t are slowly changing non-stationary components exhibiting seasonal features and z_t is a high-frequency, stationary component. Each of the components is then modeled and y_t is obtained by combining them according to equation (6.7). Guedes Soares et al. [66] followed a similar procedure, however they first decomposed the data into seasonal and stationary components and then applied a skewness reducing transformation to the stationary data series. The second approach relies on using a non-stationary distribution function to transform the data to standard normal [158]. With this approach no additional transformation is necessary.

We used a method in line with the first approach to develop the simulation model. As potential transformations we considered the Box-Cox family [24]

$$f(x, \lambda) = \begin{cases} \frac{x^\lambda - 1}{\lambda}, & \lambda \neq 0, \\ \log(x), & \lambda = 0 \end{cases}, \quad (6.8)$$

shifted logarithms

$$f(x, c) = \log(x + c), \quad c \geq 0, \quad (6.9)$$

and a transformation to standard normal using the empirical distribution function in the probability integral transform

$$f(x) = \Phi^{-1}(\hat{F}_n(x)), \quad (6.10)$$

where Φ^{-1} is the inverse of the standard normal cumulative distribution function and $\hat{F}_n$ is the empirical distribution function estimated for the data $x_1, \dots, x_T$. All three transformations, each denoted as f for simplicity, have proven valuable in similar applications [42, 66, 162].

The choice of transformation strongly influenced the simulation results that would be obtained at later stages. Box-Cox transformations whose λ was estimated by maximum likelihood or a (not shifted) logarithm resulted in unrealistically high simulated values. For instance, simulated significant wave heights were in the order of 30m, while the highest observed is below 7m. In contrast, choosing higher values of λ or shifting the logarithm by a positive constant c would result in reasonable maximum heights, but also in negative ones. These issues arose both when transforming the data before the decomposition as well as when transforming the stationary component.

The transformation to standard normal via the probability integral transform resulted in simulation values that were representative of real values, at least when applied before the decomposition. When transforming after the decomposition, the simulated values of μ_t , σ_t and Z_t would sometimes combine to a negative value. This is not surprising, since we modeled μ_t and σ_t independent of Z_t , as will be explained in the next sections. Hence, the transformation to standard normal via the probability integral transform before a decomposition was chosen. Figures 6.8a and 6.8b show the relationship

between transformed and original variables. Time series of the transformed variables are shown in Figure 6.9.




Figure 6.8: Normalizing transformation for (a) significant wave heights and (b) mean zero-crossing periods




Figure 6.9: Time series of normalized (a) significant wave heights and (b) mean zero-crossing periods from January 1989 to December 2014.

After the time series x_t , $t = 1, \dots, T$, has been transformed to the series y_t , the seasonal components μ_t and σ_t of equation 6.7 are extracted using a smoothing technique. μ_t can be interpreted as a local mean and σ_t as a local standard deviation. Both were computed with sliding windows and an Epanechnikov kernel as weighting function for smoothing, as follows:

$$\mu_t = \frac{1}{2t'} \sum_{k=t-t'}^{t+t'} K_{2t'}(x_t - x_k) \quad (6.11)$$

and

$$\sigma_t = \sqrt{\frac{1}{2t'} \sum_{k=t-t'}^{t+t'} K_{2t'}((x_t - \mu_t)^2 - (x_k - \mu_k)^2)}, \quad (6.12)$$

where K_{2k} is the Epanechnikov (parabolic) kernel. The bandwidth, t' , was set to 720 hours, which amounts to 30 days and is in line with the common practice to deseasonalize oceanographic variables via monthly statistics [recent examples are 43, 184].

Model for the non-stationary components The non-stationary processes exhibit seasonal patterns, but also inter-year variability (Figures 6.10a-6.10d). Such variabilities are inherent in the climatic system [e.g., 107, 118], and thus reflected in wave parameter time series. Méndez et al. [114, 115] as well as Serafin and Ruggiero [150] described seasonality of wave or water level parameters as a superposition of an annual cycle and a semiannual cycle using a Fourier series representation. Further, they approximated inter-year variability by regressing on time series of regional climate indices related to El Niño - Southern oscillation (ENSO), such as the POI, PNA or MEI. Solari and van Gelder [158] took another approach to modeling inter-year variability. Instead of studying the physical basis and using climatic covariates, they identified a 5-year and 26-year cycle in their observed data and included these in a Fourier series representation together with seasonal cycles.




Figure 6.10: Seasonal mean and standard deviation processes: (a) $\mu_t^{(Hm0)}$, (b) $\mu_t^{(Tm02)}$, (c) $\sigma_t^{(Hm0)}$, and (d) $\sigma_t^{(Tm02)}$.

For this chapter, we neither studied climatological covariates nor cycles longer than 1 year. We developed a somewhat different approach assuming that it is more important to represent the range of inter-year differences than their temporal sequencing. While we use a 5-factor Fourier series in order to represent both annual and semiannual cycles following earlier studies [114, 115, 150], we let the Fourier coefficients vary each year to produce inter-year variations.

For simplicity, we let the coefficients vary randomly according to a multivariate Gaussian distribution that accounts for dependencies between the four local mean and standard deviation processes. In the future, it might be worthwhile to further investigate, if the coefficients could be predicted by climatic covariates.

The first step of this approach involves dissecting the 24-year data series $\mu_t^{(Hm0)}$, $\mu_t^{(Tm02)}$, $\sigma_t^{(Hm0)}$ and $\sigma_t^{(Tm02)}$ into 1-year segments. To easily identify them, we re-index

the series using a double index, $k = 1, \dots, 24$ for the year and $\tau = 1, \dots, 8766$ for the hour within the year². Then, we estimate 5-factor Fourier series for each segment:

$$f^{(\mu_k^{(i)})}(\tau) = a_0^{(\mu_k^{(i)})} + a_1^{(\mu_k^{(i)})} \cos\left(\frac{2\pi\tau}{T}\right) + a_2^{(\mu_k^{(i)})} \sin\left(\frac{2\pi\tau}{T}\right) + a_3^{(\mu_k^{(i)})} \cos\left(\frac{4\pi\tau}{T}\right) + a_4^{(\mu_k^{(i)})} \sin\left(\frac{4\pi\tau}{T}\right) \quad (6.13)$$

and

$$f^{(\sigma_k^{(i)})}(\tau) = a_0^{(\sigma_k^{(i)})} + a_1^{(\sigma_k^{(i)})} \cos\left(\frac{2\pi\tau}{T}\right) + a_2^{(\sigma_k^{(i)})} \sin\left(\frac{2\pi\tau}{T}\right) + a_3^{(\sigma_k^{(i)})} \cos\left(\frac{4\pi\tau}{T}\right) + a_4^{(\sigma_k^{(i)})} \sin\left(\frac{4\pi\tau}{T}\right) \quad (6.14)$$

where $T = 8766$ and $i = \{H_{m0}, \tilde{T}_{m02}\}$. To obtain the 24-year series, the fitted 1-year segments are concatenated, for example $f^{(\mu^{(i)})} = [f^{(\mu_1^{(i)})}, \dots, f^{(\mu_{24}^{(i)})}]$, and discontinuities at the transitions from one year to another are smoothed out using cubic spline interpolation. The continuous fitted series obtained in this way explain more than 90% of the variance in the corresponding data series (Table 6.3).

Table 6.3: Coefficient of determination R^2 for the fitted seasonal mean and standard deviation series.

$\mu_t^{(H_{m0})}$	$\mu_t^{(\tilde{T}_{m02})}$	$\sigma_t^{(H_{m0})}$	$\sigma_t^{(\tilde{T}_{m02})}$
0.96	0.91	0.98	0.94

In the next step, we assume that the estimated coefficients $a_m^{(i,k)}$ are i.i.d. observations of random variables $A_m^{(i)}$, $m = 0, \dots, 4$, and estimate a joint distribution for them. Since the sample size is small ($N = 24$) compared to the dimension of the random vector ($d = 20$), an extensive analysis of its distribution seems infeasible and for simplicity we assumed it to be multivariate Gaussian.

We modeled the distribution in terms of univariate Gaussian marginals and a multivariate Gaussian copula. Despite the relatively high dimensionality, the correlation matrix is sparse. On one hand, the basis functions of a Fourier series are mutually orthogonal, hence their coefficients uncorrelated leading to many zero-valued entries. On the other hand, we performed the bivariate asymptotic independence test based on Kendall's τ for the remaining pairs of coefficients, which is implemented in the VineCopula Package [148]. According to the test, most pairs are independent. Only three correlations were found to be significant: $\rho(A_0^{(\sigma^{(H_{m0})})}, A_0^{(\sigma^{(\tilde{T}_{m02})})}) = -0.7$, $\rho(A_1^{(\mu^{(\tilde{T}_{m02})})}, A_1^{(\sigma^{(\tilde{T}_{m02})})}) = 0.6$, and $\rho(A_2^{(\mu^{(H_{m0})})}, A_2^{(\mu^{(\tilde{T}_{m02})})}) = -0.5$, where ρ denotes the product moment correlation. Thus, these correlation values correspond to the only three non-zero off-diagonal elements of the correlation matrix parameterizing the Gaussian copula. The parameters of the univariate distributions were estimated by maximum likelihood and can be found in Table 6.4.

²8766 hours correspond to one year when "correcting" for leap years.

Table 6.4: Parameters (mean, standard deviation) of the univariate normal distributions of the Fourier coefficients

	A_0	A_1	A_2	A_3	A_4
$\mu^{(H_{m0})}$	(0, 0.09)	(0.39, 0.09)	(-0.10, 0.12)	(0.03, 0.08)	(-0.03, 0.10)
$\mu^{(\tilde{T}_{m02})}$	(0, 0.10)	(-0.21, 0.06)	(0.14, 0.1)	(-0.03, 0.04)	(0, 0.07)
$\sigma^{(H_{m0})}$	(0.91, 0.05)	(0.01, 0.05)	(0, 0.06)	(-0.01, 0.03)	(0.01, 0.05)
$\sigma^{(\tilde{T}_{m02})}$	(0.96, 0.04)	(0, 0.04)	(0.03, 0.03)	(-0.01, 0.02)	(0, 0.03)

Model for the stationary components The deseasonalised processes, obtained through

$$z_t^{(i)} = \frac{y_t^{(i)} - \mu_t^{(i)}}{\sigma_t^{(i)}}, \quad (6.15)$$

are shown in Figure 6.11. These will be represented as ARMA process. To this end, we first inspect the autocorrelation functions (ACFs) and partial autocorrelation functions (PACFs) to make a first guess on the form of the ARMA model (see Section 2.3). We iteratively search for suitable orders p and q , starting from our first guess. We choose the $ARMA(p, q)$ model with the least parameters $p + q$ that can produce simulated series with ACFs and PACFs that closely resemble the ones of the process $z_t^{(i)}$ and for which the residuals are white noise processes. The latter is checked by computing ACFs and PACFs of the residuals and squared residuals. In each iteration parameters are estimated based on maximum likelihood, as implemented in the *arima()* function of R's stats package [39]. The parameter estimates finally obtained are listed in Table 6.5.




Figure 6.11: Stationary components of normalized (a) significant wave height and (b) mean zero-crossing period time series as obtained after the decomposition.

Once adequate orders p and q have been found, the residuals, $\epsilon_t^{(i)}$, of the fitted $ARMA(p, q)$ model for $z_t^{(i)}$ should be (almost) i.i.d. and following an arbitrary distribution with zero mean and constant variance. Because we expect differences across the wave direction regimes, we estimate conditional distributions for the residuals given the regime, $\mathcal{E}^{(i)} | R^\Theta = k$, for each $k \in \{0, 1\}$. We cannot expect these distributions to be Gaussian and we need to find a suitable parametric family of distributions. We considered the following families in this study: normal, skew-normal, t and skew-t. We select the family and attest its goodness of fit with visual diagnostic tools. We examine qq-plots and check if applying a probability integral transform (PIT) to the residuals via the selected

family results in uniformly distributed values. We refer to these as PIT residuals and visually assess their uniformity with histograms. All four univariate distributions are best approximated by a skew-t family. The estimated parameters can be found in Table 6.6. Diagnostic plots to verify the adequacy of the selected orders p and q and the marginal distribution are shown in Figures 6.12 and 6.13.

Table 6.5: Coefficients of ARMA models for $Z_t^{(H_{m0})}$ and $Z_t^{(T_{m02})}$. Standard errors are given in parenthesis.

Process	ar1	ar2	ar3	ma1	ma2	intercept
$Z_t^{(H_{m0})}$	1.07 (0.00)	0.10 (0.01)	-0.18 (0.00)	-	-	0.00 (0.03)
$Z_t^{(T_{m02})}$	2.63 (0.00)	-2.54 (0.00)	0.89 (0.00)	-1.62 (0.00)	0.83 (0.00)	0.00 (0.01)

Table 6.6: Parameters (mu, sigma, skew, shape) of regime-dependent skew-t distributions for residuals $\mathcal{E}^{(H_{m0})}$ and $\mathcal{E}^{(T_{m02})}$.

k	$\hat{F}_{\mathcal{E}^{(H_{m0})} R^\Theta=k}$	$\hat{F}_{\mathcal{E}^{(T_{m02})} R^\Theta=k}$
0	(-0.01, 0.16, 1.07, 4.80)	(0.05, 0.35, 0.87, 5.58)
1	(0.01, 0.17, 1.13, 5.36)	(-0.04, 0.43, 0.90, 5.52)

Finally, the residual processes, $\epsilon_t^{(1)}$ and $\epsilon_t^{(2)}$ could depend on each other, because the original time series are interrelated. Therefore, we construct regime-dependent bivariate residual distributions via copulas. We fit them on the empirical ranks of the conditional residuals normalized to (0, 1). As before, we use the AIC criteria for model selection and estimate the parameters by maximum likelihood. Table 6.7 contains the selected bivariate copula families, estimated parameters and Kendall's τ . Figure 6.14 shows the bivariate density contours for observed and simulated residuals.

Table 6.7: Bi-variate copula families, their parameters and Kendall's τ for regime-dependent copulas for residuals $\mathcal{E}^{(H_{m0})}$ and $\mathcal{E}^{(T_{m02})}$.

k	$\mathcal{E}^{(H_{m0})}, \mathcal{E}^{(T_{m02})} R = k$			
	family	par	par2	tau
0	t	-0.09	5.53	-0.06
1	t	-0.23	6.36	-0.14




Figure 6.12: Diagnostic plots for $\epsilon_t^{(Hm0)}$: (a) ACF of residuals, (b) PACF of residuals, (c) ACF of squared residuals, (d) PACF of squared residuals, (e) histogram of PIT residuals for northern directions, (f) histogram of PIT residuals for southern directions.




Figure 6.13: Diagnostic plots for $\epsilon_t^{(Tm02)}$: (a) ACF of residuals, (b) PACF of residuals, (c) ACF of squared residuals, (d) PACF of squared residuals, (e) histogram of fitted quantiles for northern directions, (f) histogram of fitted quantiles for southern directions.




Figure 6.14: Bivariate density contour plots residuals $\mathcal{E}^{(Hm0)}$ and $\mathcal{E}^{(Tm02)}$ (a) corresponding to northern wave directions and (b) corresponding to south-western wave directions. Densities are estimated from 24 observed years (black) and 24 simulated years (gray).

6.3. SIMULATION RESULTS

WAVE ANGLE REGIMES

The developed renewal process model for wave angle regimes captures persistency and seasonal differences well. Figure 6.15 shows an example of a observed and a simulated wave direction regime time series. Figure 6.16 compares the percentage of time per year in which waves come from the Southwest for observed data (12 years) and simulated data (10^5 years). The model correctly reflects that the highest percentage of south-western waves occurs in spring, while the lowest percentage of south-western waves occurs in autumn. All except one measured percentage fall within 1.5 times the interquartile range from the lower and higher quartile, respectively. Also a two-sided, sample-based Kolmogorov-Smirnov (KS) test does not reject the null hypothesis that observed and simulated percentages are equal in distribution for any of the seasons. P-values range from 0.52 for winter to 0.96 for autumn.




Figure 6.15: Wave direction regime time series (a) observed in 2003 and (b) for a simulated year. Light gray represents northern wave directions and dark gray represents south-western wave directions.




Figure 6.16: Percentages of time per season with south-westerly wave directions during each of the four seasons. Boxplots summarize data of 10000 simulated years. The whiskers represent the lowest and highest simulated data point still within 1.5 times the interquartile range from the lower and higher quartile, respectively. The open circles represent all remaining data points. Colored points represent observed percentages during 12 years.

SIGNIFICANT WAVE HEIGHTS AND MEAN ZERO-CROSSING PERIODS

To assess how well the simulation method replicates the characteristics of significant wave height and mean zero-crossing periods time series, we rely on a visual diagnostics.

We do not apply statistical goodness of fit tests, because any model would be rejected for a sample size this large. There are 201960 observations and 210384 simulated values. (The discrepancy is due to the fact that the observed time series contain a number of missing values.)

Figures 6.17a and 6.17b show simulated time series for H_{m0} and T_{m02} . These time series results from a combination of all model components described previously. The simulated time series of H_{m0} are visually not distinguishable from the equally long observed time series in Figure 6.1a. In particular, the model reflects annual seasonality and inter-year differences well. Nonetheless, potential multi-annual cycles have not been accounted for. The simulated time series of T_{m02} also reflect the main characteristics of the corresponding observed time series in Figure 6.1b.




Figure 6.17: Simulated time series of (a) significant wave height and (b) mean zero-crossing period for a duration of 24 years.

Figure 6.18a shows boxplots of observed and simulated significant wave heights for different wave angle directions. The model approximates well the distribution of observed data (based on visual comparison of first quartile, median and third quartile). The model also reflects well that northern waves tend to be lower than southwestern waves, and that their distribution is more narrow. Analogously, Figure 6.18b shows boxplots of observed and simulated mean zero-crossing periods. Again, the model well represents the overall distribution of the observed data. Differences between the regimes, which are less pronounced than in the case of significant wave heights are also adequately captured. Mean zero crossing-periods tend to be higher for northern waves than for southwestern waves.

Figure 6.19 shows probability density contours for simulated and observed pairs of H_{m0} and T_{m02} . In both cases the density contours have been estimated from time series data of a length of 24 years. A visual comparison of contours shows that the model well approximates the bivariate distribution of H_{m0} and T_{m02} . Especially the maximum steepness condition is well-represented.

Figure 6.20 compares individual contour lines for the bivariate distribution conditional on northern and southwestern directions. The contour levels are the same as in Figure 6.19. The unconditional contours are included for reference. From top to bottom, the $25 \cdot 10^{-4}$, $5 \cdot 10^{-4}$ and 10^{-4} density contours are shown in separate panels. The left three panels show bivariate density contours of observed data, while the three right panels show bivariate density contours of simulated data. For the highest density level (top pair of panels), the model well approximates the bivariate density for southwest-




Figure 6.18: Boxplots of observed and simulated data for (a) significant wave height and (b) mean zero-crossing period.

ern directions, but underestimates the degree of positive dependence between the two variables for northern directions. For the intermediate density level (middle pair of panels), the opposite is true. The model adequately approximates the bivariate density for northern directions, but underestimates the degree of positive dependence between the two variables for southwestern directions. For the lowest density level (bottom pair of panels), the model still underestimates the degree of positive dependence between the two variables for southwestern directions. Moreover, for high mean zero-crossing periods, significant wave heights for northern directions are underestimated, while they are overestimated for southwestern directions. An overall tendency to underestimate the degree of positive dependence for southwestern directions is also reflected by values of Spearman's ρ (Table 6.8).

Table 6.8: Values of Spearman's ρ for observed and simulated significant wave height and mean zero-crossing periods for different wave angle directions.

Wave direction	Observed data	Simulated data
All	0.74	0.73
Northern	0.70	0.70
Southwestern	0.88	0.80

Next, we investigate the capability of the model to simulate sequences of storms. As mentioned in the introduction, a potential advantage of simulating full time series is that




Figure 6.19: Bivariate density contours for significant wave heights and mean zero-crossing periods. Densities are estimated from 24 observed years (black) and 24 simulated years (gray).

no prior assumptions have to be made on critical threshold values or storm shape. This is useful, because different applications might require different assumptions. Here, we explore the model's capability to reproduce storm durations and storm interarrival time for two different storm definitions. Storm I events are defined as periods in which the significant wave height exceeds 300cm and the mean zero-crossing period exceeds 6s.

This storm event definition is inspired by Li et al. [103] who simulated sea storms for a location in the Dutch North sea with comparable geographical properties. They use a significant wave height threshold of 300cm in combination with a surge threshold of 50cm. These thresholds are chosen, because more severe conditions are likely to cause morphological change [139]. Since 300cm significant wave height corresponds to the 96.5% quantile in our data set, we combined it with the 96.5% mean zero-crossing period quantile, which is 6s. Storm II events are defined as more severe. We defined a period as a storm II events when both variables exceed their respective 99.5% quantiles, which are 407cm and 6.7s. The threshold values for both event definitions are summarized in Table 6.9.

Table 6.9: Critical threshold values for Storm I and Storm II events.

Storm	H_{m0}	T_{m02}
I	300cm	6s
II	407cm	6.7s

Figure 6.21 compares storm durations and interarrival times for observed and simulated time series. The observed time series is the segment from 2003 to 2014, because this period does not contain missing values. The simulated time series is also 12 years long. The simulation model produces realistic storm durations and storm interarrival times for both event definitions, storm I and storm II. However, the model appears to slightly overestimate storm durations and produces three storm II interarrival times that

are longer than a year.

Furthermore, the model simulates a realistic number of storms, when comparing the observed 12 year period to 100 simulated 12 year periods. Of course, this is expected given the results on storm duration and interarrival time. From 2003 to 2014, 402 storm I events and 69 storm II events have been observed. The respective 5th percentile and 95th percentile of the corresponding simulated storm events are 291 and 434 as well as 47 and 104.

6.4. KEY POINTS

In this chapter, we presented a simulation method for joint time series of significant wave height, mean zero-crossing period and a directional regime variable. The latter distinguishes between northern and southwestern waves. Time series can be simulated at a high resolution of 1 hour, which is useful for risk analyses in various coastal and offshore applications. The method has been applied to a data set in the Dutch North sea.

The method contains several modeling steps and relies on renewal processes, Fourier series with random coefficients, ARMA processes, copulas, and regime-switching. A particular feature is a data-driven estimate for a wave height-dependent limiting wave-steepness condition, which we use to describe part of the dependence between significant wave heights and mean zero-crossing periods and which facilitates the copula-based dependence modeling later on. Similar to many other studies, annual seasonality is represented by Fourier series. The coefficients are modeled as inter-dependent random variables to account for inter-year differences. At this point we did not consider climatic covariates and recommend to examine if they have predictive skill for inter-year differences.

The stationary components of the two processes are represented as ARMA with regime-switching joint residual distribution, constructed using copulas. The regime-switches are triggered by changes in wave direction from North to Southwest. While these regime-switches result in differences between bivariate distributions of simulated significant wave height and mean zero-crossing period when conditioned on the northern and southwestern regime, they are not as pronounced as in observed data. We recommend that future research is directed at improving and extending the simulation method as to better capture these differences. Nonetheless, the unconditioned bivariate distribution of significant wave height and mean zero-crossing periods appears to be represented adequately.

Moreover, storm durations and storm interarrival times are well captured for two different storm definitions, which rest on different critical threshold values for the two significant wave height and mean-zero crossing periods. As storm sequences are adequately represented, the model has potential value for applications in coastal and offshore engineering, such as the prediction of long-term morphological changes, or the planning and budgeting of offshore operations.






Figure 6.20: Individual contour lines for the bivariate distribution of significant wave heights and mean zero-crossing periods conditional on northern and southwestern wave directions. From top to bottom the two plots in each panel show the 0.0025, 0.0005 and the 0.0001 contour.



(a)



(b)

Figure 6.21: Boxplots of storm durations and storm interarrival times for storm I and storm II events. The underlying observed and simulated time series from which the durations and interarrival times were derived were both 12 years long.

COMPARISON AND LIMITATIONS

The two preceding chapters developed methods for the statistical simulation of time series of wave parameters at an hourly scale. The main difference between the two methods is how serial dependence is modeled. The first method (Chapter 5) applied vine-copulas, while the second method (Chapter 6) relied on ARMA processes. Furthermore, the data pre-processing was different. The purpose of this chapter is to evaluate the practical applicability of the methods in different contexts. To this end, the chapter compares the methods in terms of underlying assumptions, model performance, extendability to higher dimensions and transferability to other locations. For easy reference, the simulation models proposed in Chapter 5 will be referred to as ‘model 1a-c’ and the one in Chapter 6 as ‘model 2’ (cf. Table 7.1).

Table 7.1: Overview on Models.

Name	Model
Model 1a	Vine-Copula based with $\rho_{t,t+1} = 0.99$
Model 1b	Vine-Copula based with $\rho_{t,t+1} = 0.98$
Model 1c	Vine-Copula based with $\rho_{t,t+1} = 0.95$
Model 2	ARMA based

Section 7.1 compares simulation results obtained from the two methods. In particular, the section compares (1) the bivariate density for simulated time series data of significant wave height and mean zero-crossing periods during the oceanographic winter period, (2) their persistency characteristics, and (3) return levels of significant wave heights. Section 7.2 points toward model limitations. Topics are (1) model uncertainty and sensitivity, (2) extendability of either method to higher dimensions, that is to include additional variables, and (3) transferability to other locations. Finally, Section 7.3 summarizes the findings of the chapter.

7.1. MODEL COMPARISON

The main differences between the methods in terms of underlying assumptions and data preprocessing are listed in Table 7.2. The comparison of observed and simulated data from the models developed in the previous two chapters is not exhaustive; the focus is on three selected topics: (1) bivariate density contours of the time series data, (2)

persistence characteristics, and (3) return level calculations for significant wave heights. Because models 1a-c are restricted to oceanographic winter periods (1 November - 15 April) and only includes significant wave height and mean zero-crossing periods, the comparison is only performed on according time series segments for model 2.

Table 7.2: Comparison of selected properties of the methods as applied in this thesis.

Aspect	Vine-Copula method	ARMA method
Markov order	1	3 for the H_{m0} ; infinite for T_{m02}
Treatment of seasonality	Restriction to oceanographic winter season	Decomposition and separate modeling of seasonal and stationary processes
Inter-year differences	Not explicitly accounted for	Explicitly accounted for through randomly varying Fourier series coefficients
Modeling of serial correlation	Sub-sampling & vine-copula	ARMA process
Modeling of cross-series dependence	Vine-copula	Wave steepness-based transformation & copulas
Gaps in time series record	Allowed	Not allowed

BIVARIATE DENSITY CONTOURS OF TIME SERIES DATA

Observed data are available for 23 winter seasons (1991/1992 through 2013/2014). Figure 7.1 compares bivariate density contours of the time series data observed during these periods with time series data simulated for 23 winter seasons. The contours are obtained via bivariate kernel density estimation on a square grid with 50 grid points in each direction. For all four models, bivariate densities of simulated and observed data are comparable. In particular, a limiting steepness condition is well represented in all cases. Discrepancy mainly exist in the outermost contour when both variables are small and when they are large for model 1a-c.

For a more objective comparison, we computed the sum of squared distances between the distribution of observed and simulated points, which is the Cramer-van-Mises test statistic:

$$\mathcal{C.M}_n = \int_A \sqrt{n} (F_n(x, y) - F_{\theta_n}(x, y))^2 dx dy, \quad (7.1)$$

where F_n denotes the empirical bivariate cumulative distribution function obtained from observed data and F_{θ_n} the one obtained from simulated data and A is the domain of X, Y .¹ We computed the CvM statistic numerically for five simulations of 23 winter seasons with each of the models. The number five is a heuristic choice to explore how large differences between simulation runs are. Just as before, we use a square grid with 50 points in each direction.

¹Here, θ_n refers to the parameter vector of the distribution estimated from n data points. Note that we previously use θ to denote copula parameters in Chapter 5 and Θ for the wave direction regime in Chapter 6.




Figure 7.1: Bivariate contours for observed and simulated time series data of significant wave height and mean zero-crossing periods

The results are listed in Table 7.3. The statistics for model 1a-c are of similar order of magnitudes, with the smallest for different simulations either corresponding to model 1a or 1c. The statistics for model 2 are notably smaller than for models 1a-c. Thus, model 2 better approximates the bivariate distribution according to the CvM criterion. However, this increased accuracy may not be of practical relevance.

PERSISTENCY

This section investigates the models' capability to represent sequences of (1) storms that are expected to have morphological impact on the coast and (2) calm weather windows during which specific offshore operations can be executed.

Storm events are defined as periods in which the significant wave height and mean

Table 7.3: Sum of squared distances for distributions of simulated and observed time series data.

Model	Simulation 1	Simulation 2	Simulation 3	Simulation 4	Simulation 5
1a	154.6	403.8	577.9	270.4	510.4
1b	184.7	416.5	527.6	304.3	476.8
1c	246.5	418.2	477.0	347.2	440.5
2	97.7	40.0	61.0	19.0	31.4

zero-crossing period exceed 300cm and 6s, respectively. The motivation behind this storm definition is given in the previous Chapter in Section 6.3. Weather windows are defined as periods in which the significant wave heights remains below 150cm and the mean zero-crossing periods remains below 4.7s. The selected significant wave height limit has been reported in several studies on operability of marine operations [3, 102] and the mean zero-crossing period limit has been obtained in a similar fashion as for the storm. Note, that in practice, the limits are usually not fixed and independent, but one is dependent on the other [3].


Figures 7.2a and 7.2b show storm duration and interarrival time for observed and simulated time series for a period of 11 years. The period for which data are compared here is shorter than in the previous section, because persistency cannot be accurately computed when missing values are present. All four models capture that storm durations are approximately an order of magnitude shorter than storm interarrival times. Simulated storm durations from models 1a-c are more narrowly distributed and tend to underestimate the ones observed. The contrary is true for durations from model 2. In terms of storm interarrival time, simulated values from model 1a and b have a lower median and third quartile and the ones from model 1c and 2 have a higher median and 3rd quartile than observed values. When comparing models 1a, b and c with each other, skewness of the distribution of storm interarrival time seems to increase with decreasing values for $\rho_{t,t+1}$ and at the same time the tail shortens.

Similarly, Figures 7.3a and 7.3b show weather window duration and interarrival time. In this case the duration is of a similar order of magnitude as the interarrival time, which is reflected by all models. Nonetheless, they tend to underestimate both. The distribution of duration and interarrival time obtained via model 2 is closest to the empirical one. Furthermore, a similar trend for skewness and tail behaviors for models 1a-c can be observed.


Overall, model 2 seems to be more versatile than models 1a-c, because its performance in terms of persistence is more consistent across different thresholds. Nonetheless, risks could be underestimated when applying the model, since storm interarrival times tend to be longer than observed, while weather window interarrival times tend to be shorter.

ANNUAL MAXIMA AND RETURN LEVELS

Many engineering applications require multivariate return levels or simulations of annual maxima of hydro-meteorological variables, for instance design and reliability analysis of structures. A natural question is therefore, if the time series models could be valu-




(a)




(b)

Figure 7.2: Boxplots for storm duration and interarrival time of simulated and observed time series.



(a)



(b)

Figure 7.3: Boxplots for weather window duration and interarrival time for simulated and observed time series.

able for the estimation of return levels. For a general introduction to statistical modeling of extreme values the reader is referred to [34].

Note that in most years (19 out of 23) the observed annual maximum occurs in winter. The instances, for which this is not the case are listed in table 7.4. All them occur during the second half of October, where the wave climate is similar to the one in winter, indicating that using winter time series to estimate annual maxima is justified. Nonetheless, we would recommend further research on the definition of the winter period and its effects on simulated annual maxima.

Table 7.4: Annual maxima that did not occur during the oceanographic winter period

Date	Value
25/10/1998	567cm
27/10/2002	636cm
28/10/2013	533cm
21/10/2014	538cm

As an example, Figure 7.4 compares return levels estimated from observed and simulated annual winter maxima of significant wave heights up to a return period of 100 years. The observed data consists of 23 and the simulated data of 115 winter segments. The latter correspond to the $5 \cdot 23$ simulated winter seasons of section 7.1. The Figure shows return levels estimated empirically (dots) and by fitting the generalized extreme value (GEV) distribution (lines).




Figure 7.4: Return levels (RLs) estimated empirically and via a GEV distribution for an observed winter series ($N = 23$) and four simulated winter series generated by the different models ($N = 100$).

Return levels estimated from time series simulated by models 1a-b are lower than observed. This could indicate that the models tend to underestimate annual maxima. Levels from model 1a are the lowest, followed by levels from model 1b, then 1c. Differences between the models are larger for lower return periods than for higher ones. For 100 years, the return levels seem to approach the same value. Differently, return levels estimated from time series simulated by model 2 are slightly lower than observed below return periods of 12 years and slightly higher above.

While return level estimates from models 1c and 2 seem realistic, the application of the time series models for extreme value problems is not straightforward and requires further research. One issue related to model 2 is the use of the empirical cumulative distribution function for an initial transformation of significant wave heights (Section 6.2). This limits simulated annual maxima to the highest observed in simulated time series and may artificially bound return level estimates for high return periods (e.g., 1000 years) from above. Finding a different suitable variable transformation is therefore fundamental before estimating return values for practical applications.

It is important to realized that the simulated time series are five times longer than the observed series. Due to this difference, uncertainty on the parameters of the GEV distribution can be expected to be lower, leading to narrower uncertainty bounds for the return levels. As an example, Figure 7.5 shows the 90% uncertainty bands for return levels of series that have been observed and simulated by model 1c. As the bands do not account for uncertainty related to model 1c, they are likely overconfident estimates of the uncertainty bands for return levels. Hence, further research is required on time series model uncertainties before applying them to practical extreme value problems.




Figure 7.5: 90% uncertainty interval on GEV estimates for return levels based on the observed winter series ($N = 23$) and the simulated winter series from model 1c ($N = 115$).

7.2. MODEL LIMITATIONS

MODEL UNCERTAINTY AND SENSITIVITY

Temporal Vine-copula Method Sections 5.3 and 7.1 have shown that the simulation results are highly sensitive to the choice of the parameter $\rho_{t,t+1}$. Other parameters, also of other vine configurations, have not been investigated. Neither did we explore sampling fluctuations in the estimates for $\rho_{t,t+1}$ or any of the other parameters. These also depend on the number of samples from which they are estimated. Two factors influence the number of samples: (1) length of the available data record and (2) the selected sub-sampling interval (cf. Section 5.1). Both could be varied to investigate the influence of sample size. Furthermore, the model structure, including the assumption of 1st order Markov property, is likely to an influence. For instance, Solari and Losada [157] found that a temporal vine-copula model with 2nd order Markov property outperformed an ARMA(23,23) model in terms of storm persistency when applied to a time series data set of three-hourly significant wave heights in Spain. However, the results from Solari and Losada's study are not directly translatable to the present research and further research in this direction is recommended.

ARMA Method While developing the method, we noticed that results are very sensitive to the choice of initial variable transformation. For example, a logarithmic variable transformation would result in simulated wave heights in the order of magnitude of 30m. This was mentioned in Section 6.2, although results were not presented at length. Sensitivity to other modeling choices or to parameter uncertainty has not yet been explored and remains a topic for further research.

Parameter uncertainty itself has not been investigated in depth either, but preliminary results are available for the autoregressive parameters of the AR(3) model for significant wave heights. In Figure 7.6 parameter estimates obtained individually for each year in the period 2003 - 2014 are indicated as circles in different colors. (Table 6.5 showed the parameters as estimated over the entire period.) Additionally, 90% uncertainty intervals for the parameter estimates of the year 2003 are shown indicating that differences between years are larger than sampling fluctuations. The uncertainty interval is obtained via 100 bootstrap samples of the same length as the observed series. Main open questions are (1) how sensitive simulation results would be to variations in the parameters in the order of magnitude displayed in Figure 7.6 and (2) what the underlying drivers are.

INCORPORATION OF ADDITIONAL VARIABLES

While the two methods have been specifically developed to simulate time series of significant wave height and mean zero-crossing period, they could also be adapted for other or extended to include additional variables. In fact, the present research was inspired by previous applications of similar methods to multiple financial variables [e.g., 26, 164]. Hence, examples of similar methods to other sets of variables already exist.

For offshore environments, wind speed would be a valuable addition. For coastal environments, surges as well as wave direction with a resolution of nautical degrees would be of interest [see 178, and references therein]. However, adaptation of one of the meth-




Figure 7.6: Parameter estimates of $AR(3)$ models for significant wave heights for individual years in the period 2003 - 2014 (colored circles) and 90% bootstrap uncertainty band for the estimates of the year 2003 (boxplots with whiskers indicating the 90% uncertainty band).

ods to other or additional variables may be a research task of similar scope as presented for that method in this thesis, at least for hydro-meteorological variables. The reason is that each new type of process will come with its own physical properties and relations to other variables, posing unique challenges for the adaptation of a statistical model.

TRANSFERABILITY TO OTHER LOCATIONS

The methods could be transferred to other locations, but, again, this is not necessarily straight forward. Waves are generated by physical processes on different spatial scales. Local and regional meteorological variables and bathymetry play a role in the wave generation and propagation. For locations with similar properties, in terms of climate and bathymetry, the adaptation of the methods can be expected to be easier than for locations with very different properties. For instance, we have used the wave direction as

a predictor for swell-dominant or wind-sea dominant waves, which was unique for the location in the Southern North Sea (Figure 5.1).

7.3. KEY POINTS

This chapter compared the two simulation methods developed in the previous two chapters and pointed towards some of their limitations. Models based on either method have the potential to generate realistic time series of significant wave height and mean zero-crossing period. More precise results of the comparison are the following:

- The representation of the bivariate distribution of time series data is adequate in both cases.
- In terms of persistency, model 2 seems to be more versatile than models 1a-c, because its performance is more consistent for different thresholds.
- Return level estimates of model 1c and 2 for annual winter maxima of significant wave heights are realistic as well. Although model 2 should not be applied for extreme value purposes before substituting the empirical distribution function used in the simulation methods by a suitable parametric alternative. When simulating longer synthetic time series than observed, the uncertainty bound for return level estimate may decrease. However, these bounds are likely to be over-confident, if model uncertainty of the simulation method is not taken into account. In this context an uncertainty and sensitivity analysis is recommended.

While both methods are generally suitable to generate time series that exhibit statistical features relevant for coastal and offshore risk analysis, additional variables would need to be included for many applications. For offshore operation, wind speeds are often central, while coastal risk analysis or the prediction of longterm morphological changes require surges and precise wave angles. Furthermore, the methods are location-specific, although they potentially can be adapted to other locations.

IV

CONCLUSION

MAIN CONTRIBUTIONS

The present study investigated the adaptation of selected multivariate probabilistic methods for risk analysis and decision making in coastal and offshore environments. In particular, this thesis developed a BN framework for decision support on risk reduction efforts in coastal areas and statistical simulation methods for wave conditions using temporal vine-copulas or ARMA processes. Detailed conclusions to the five research objectives, which were introduced in at the beginning of this thesis in Section 1.4 and numbered 1, 2, 3*a*, 3*b* and 4, are given below.

1. Design a BN framework for decision support on risk reduction efforts in coastal environments

Emergency management and long-term planning in coastal areas depend on detailed assessments (meter scale) of flood and erosion risks. Typically, models of the risk chain are fragmented into smaller parts, because the physical processes involved are very complex and consequences can be diverse. For example, multivariate probability models can be used to estimate the return periods of extreme storms, numerical models, based on hydro- and morphodynamic processes, to assess the respective natural responses of the coast and extent of flooding, and behavior or statistical models to estimate the diverse and complex consequences onshore. Chapter 3 developed an approach to integrate results from the separate models using a BN.

An important contribution is the learning algorithm for the BN, which integrates output from synthetic and historical hydro-and morphodynamic storm simulations (2D) with land use data, vulnerability relationships (e.g., depth-damage curves) and risk reduction measures. An executable for Windows systems and documentation is also provided. The core of the algorithm builds on SMILE, which is a reasoning engine for graphical models. It is available free of charge for academic research and teaching use from BayesFusion, LLC, <http://www.bayesfusion.com/>. The algorithm is programmed in C++ and openly available at: <https://github.com/openearth/coastal-dss>.

As part of the RISC-KIT (Resilience-Increasing Strategies for Coasts toolKIT) project, we successfully tested the approach and algorithm in a range of morphological settings. In this thesis, we describe the application to the town of Wells-next-the-Sea, Norfolk, UK, which is vulnerable to storm surges. For any storm input event, the BN estimated the percentage of affected receptors in different zones of the site by predicting their hazards and damages. Receptor types under consideration were people, residential and commercial properties, and a salt marsh ecosystem. Additionally, the BN displayed the out-

come of two structural risk reduction measures, the raising and extending of a flood wall, and a soft measure, the setting up of educational display boards. BN output showed, for example, that extending the length of the existing flood wall can be more effective for impact reduction across the variety of storm events that were considered than increasing its height.

The application illustrates that the developed BN framework fulfills its design objectives and makes it possible to

- apply it to any coastal site,
- predict impacts for any relevant storm event instantly,
- reflect the diversity of impacts (e.g., socio-economic and environmental aspects),
- reflect the spatial variability of impacts,
- and
- evaluate the performance of risk reduction measures under different storm events.

With these capabilities, BNs can form a comprehensive and concise representation of risk propagation in a complex system of physical, ecological and human components. BNs can also be integrated in operational forecasting systems. From a practical point of view, their integrative character, together with the capability to predict in real-time, makes the BN a helpful tool for decision makers. From a scientific point of view, the model development approach emphasizes how results from multiple disciplines must be connected in order to understand risks and can provide an objective basis for choosing risk reduction measures.

2. Indicate limitations of the BN framework and methods for improvement

Currently, the BN framework has a number of limitations, which could be remedied. Important cases were identified in Chapter 4 and are enumerated below:

1. The basic BN graph structure is generic for all applications and makes simplifying assumptions on independence and conditional independence between variables. While they have no effect on hazard and impact predictions for a given storm event, they may limit the BNs applicability to ensemble forecasting or diagnostic reasoning.
2. In each application, the discretization of boundary condition (BC) and receptor (R) nodes has been chosen ad hoc by local experts. The discretization's effect on predictions for hazard (H) and consequence (C) nodes has not been assessed systematically.
3. Model validation has been performed through expert knowledge. A formal method has not yet been implemented.
4. The BNs are built from synthetic data, since field observations are generally lacking due to rare occurrences of relevant events. Therefore, the validity of BN predictions depends on the validity of the underlying data-generating models. At present, imperfections of those models have not been considered.

In Chapter 4 we reasoned that limitations 2-4 could mostly be addressed by implementing a metric for *BN precision*, assuming the primary use of the BN would be ensemble mean prediction of affected receptor proportions under uncertain storm conditions. The metric would represent error bands for predicted proportions of affected receptors through cross-validation. The metric could

- assess a BN's validity for application in emergency management or long-term planning (if a maximum error has been specified by the model user),
- represent uncertainties in the underlying training data (if these have been made available to the BN in the training phase),
- and
- constitute an optimality criterion for the discretization of BC nodes.

A second metric of spatial informativeness has been proposed to serve as an optimality criterion for the discretization of R nodes. According to this criterion, receptors are grouped into the same spatial zone, when they tend to experience similar hazard intensities. This maximizes spatial informativeness, even though the BN's resolution is limited to a few zones rather than the individual locations of hundreds or thousands of receptors.

3.a) Develop a statistical simulation model for joint time series of wave parameters at hourly resolution based on temporal vine-copulas

Many engineering applications call for the generation of synthetic time series of wave conditions at high temporal resolution (in the order of one hour). Chapter 5 showed for data set in the Southern North Sea that temporal vine-copulas can be a suitable method to simulate time series of significant wave height and mean zero-crossing period. Nonetheless, selecting a vine structure, finding appropriate copula families and estimating parameters is not straightforward. The validity of the model, as well as the conclusions that can be drawn from it, are sensitive to these choices.

The simulation model was developed for oceanographic winter periods to bypass considerations on seasonal effects. The assumption on stationarity, together with the first order Markov property, enabled a parsimonious representation of the temporal bivariate process as a vine-copula based on five bivariate copulas and two marginal distributions. In the model the serial dependence between subsequent significant wave heights as well as the cross-serial dependence between significant wave heights and mean zero-crossing periods at the same time were modeled unconditionally. Differently, the serial dependence between subsequent mean zero-crossing periods was modeled conditionally on the significant wave heights at the corresponding times.

The family of skew-t copulas has been found to adequately capture cross-serial dependence, in particular tail dependence and asymmetry resulting from the limiting wave steepness condition. While this bivariate copula is a building block for the temporal vine-copula, it could also be applied outside a time series context to construct a bivariate distribution of significant wave height and mean zero-crossing period. This is a valuable by-product of the approach.

3.b) Develop a statistical simulation model for joint time series of wave parameters at hourly resolution based on ARMA models

For the same data set in the Southern North Sea, Chapter 6 showed that ARMA models are also useful for simulating bivariate time series of significant wave heights and mean zero-crossing periods. Nevertheless, the modeling scope was different to the vine-copula approach on two aspects: (1) non-stationary behavior on annual and inter-annual scale was addressed and (2) the influence of the wave direction on the marginal as well

as joint behavior of significant wave height and mean zero-crossing periods has been included. However, the modeling of the wave direction was simplified by assuming a categorical variable with two possible values, north and southwest. Hence, the entire angular domain was represented by two pre-defined directional sectors.

The data analysis confirmed that Fourier series are suitable to represent seasonal cycles in mean and standard deviation on annual and semi-annual scales, as previously reported by other studies. Additionally, variations in the cycles across years as well as a relation between the seasonal cycles of significant wave height and mean zero-crossing could be detected, when estimating Fourier series for individual years, despite the sample size being small ($N=24$, which is the length of the data record in years). These variations and dependencies were reproduced in the simulated time series by using a multivariate Gaussian distribution on the Fourier coefficients.

Univariate time series were modeled as ARMA processes and cross-serial dependencies were represented with additional methods. For the present data set, cross-serial dependence could be adequately captured by modeling parts of it in three different steps:

- A data-driven equation for the limiting wave steepness condition,
- a multivariate Gaussian distribution of Fourier coefficients,
- and
- a bivariate distribution, constructed via a copula, of ARMA residuals related to significant wave height and mean zero-crossing period.

Time series of the wave direction regime could be represented by a binary seasonal renewal process. The wave direction regimes were used to trigger regime switches in the joint distribution of ARMA residuals. While this resulted in differences between the bivariate distributions of simulated significant wave height and mean zero-crossing period when conditioned on the northern and southwestern regime, they were not as pronounced as in observed data. Nonetheless, the unconditioned bivariate distribution of significant wave height and mean zero-crossing periods was represented adequately.

8

4. Compare performance of the simulation models Both methods are suitable for generating synthetic time series of significant wave height and mean zero-crossing period. The methods can be flexibly applied in different contexts, as simulated time series emulate storms as well as calm periods. This is shown by analysing their persistency proper above high and below thresholds. For example, such methods could be applied to generate synthetic storm events as input for BN framework developed in this thesis or other approaches to coastal risk analysis. The methods could also be applied to estimate the duration of coastal or offshore operations that are sensitive to particular wave height and period limits. Nonetheless, an uncertainty and sensitivity analysis is required to validate models built according to either of the methods for a given application.

9

RECOMMENDATIONS

Recommendations are given separately for the BN framework for decision support on risk reduction efforts in coastal areas and the simulation methods for time series of wave conditions.

BN FRAMEWORK FOR DECISION SUPPORT ON RISK REDUCTION EFFORT

A number of possible future research directions were given as main contribution of Chapter 4, whose objective was to indicated limitations of the BN framework and methods for improvement, and summarized in Chapter 8. We have the following main recommendations for implementing and verifying those methods effectively:

1. Intended model end users should be involved in all stages of the model development process to ensure that their needs are met. A number of aspects can be optimized via mathematical models. These are the choice of model structure, the generation of a representative set of synthetic storm events for BN training, and the discretization of boundary condition or receptor nodes. However, especially the latter two require the definition of an optimization criterion and it can be debatable how to choose it. This issue can be resolved by eliciting such criteria from end users. This may also clarify instances in which mathematical optimization of some sense is not desired due to external factors. For example, end users might prefer to adhere to known administrative zones, even though the BN could make "better" predictions when a different geographic division is used.
2. Especially for verifying the proposed methods for improvement, it is essential to use a case study with data of high quality and quantity.
3. As the validity of the BN (and of any similar tool) depends on the validity of underlying data-generating models, further research should be directed at improving them and assessing their uncertainties. Models of interest include numerical hydro- and morphodynamic models as well as the vulnerability relationships.

STATISTICAL SIMULATION METHODS FOR TIME SERIES OF WAVE CONDITIONS

The main recommendations pertaining to Part III of this thesis are the following:

1. The model development process has shown that simulation results can be sensitive to existing structural uncertainty and parameter uncertainty. A sensitivity

and uncertainty analysis is therefore recommended. For the temporal vine-copula method, a particular question of interest is, if a model with 2nd order Markov property would be less sensitive to parameter uncertainty and have improved persistency characteristics. For the ARMA method, the choice of initial variable transformation should be further investigated, including a probability integral transform (PIT) with a non-stationary distribution function. A second order Markov property as well as a PIT with a non-stationary distribution function have been reported useful by Solari and Losada [157] for a similar univariate application.

2. Time series of significant wave height and mean zero-crossing periods are often necessary, but not sufficient, as input for risk analysis in coastal and offshore environments. Further research should be directed at the the inclusion of additional hydro-meteorological variables. For example, wind speeds are of interest for offshore operations, while surges and precise wave directions central to coastal risk analysis and predictions of long-term morphological change.
3. Drivers for inter-year variations in wave climate have not been investigated. We recommend to examine, if climatic covariates have predictive skill for inter-year differences, as this was the case in similar studies [43, 114, 115, 150].
4. For the data set in the Southern North Sea, the joint behavior of significant wave height and mean zero-crossing period depended on the directional regime (North or Southwest). By introducing regime-switches in the joint residual distribution only minor differences in joint behavior of simulated data could be produced. Further research should be directed at extending the simulation method as to better capture these differences.
5. It is recommended to explicitly model wind sea significant wave height and mean zero-crossing period as well as swell significant wave height and mean zero-crossing period, if the corresponding data can be obtained. We would expect this to be more accurate than using a wave direction regime as an indicator for wind sea and swell conditions and corresponding differences in the joint distribution. Furthermore, the approach could be more easily transferable to locations where the wave direction has little predictive skill on wind sea and swell conditions.

APPENDICES

A

LEARNING ALGORITHM FOR THE CONDITIONAL PROBABILITY TABLES OF HAZARD NODES

Table A.1: Nomenclature for CPT learning

<i>Super- or Subscripts</i>	
n	simulation index
bc	boundary condition index
m	DRR measure index
r	receptor type index
h	hazard index
pa	parent set
<i>Variables</i>	
X_{bc}	boundary condition bc
X_m	DRR measure m
X_r	receptor type r
X_{h_r}	hazard h for receptor type r
E	experience matrix
<i>Constants</i>	
S_{bc}	number of states of X_{bc}
S_m	number of states of X_m
S_r	number of states of X_r
S_{h_r}	number of states of X_{h_r}

Let X_{h_r} denote a hazard node for receptor type r with states $s_{h_r} = 1, \dots, S_{h_r}$. We denote DRR measures, boundary conditions and receptors analogously using the subscripts m , bc and r (cf. Table A.1).

Before learning the CPT, we introduce an experience matrix, E , of size $\mathbf{S}_{bc} \times \mathbf{S}_{\mathbf{m} \in pa(h_r)}$. $\mathbf{S}_{bc}$ is the vector of number of states of all boundary conditions and $\mathbf{S}_{\mathbf{m} \in pa(h_r)}$ is the vector of the number of states of DRR measures in the parent set of X_{h_r} . Before learning, all

A

entries are 0. For each simulation n we set

$$E^{(n)}[\mathbf{s}_{\mathbf{bc}}^{(n)}, \mathbf{s}_{\mathbf{m} \in pa(h_r)}^{(n)}] = E^{(n-1)}[\mathbf{s}_{\mathbf{bc}}^{(n)}, \mathbf{s}_{\mathbf{m} \in pa(h_r)}^{(n)}] + 1. \quad (\text{A.1})$$

Here, $\mathbf{s}_{\mathbf{bc}}^{(n)}$ and $\mathbf{s}_{\mathbf{m} \in pa(h_r)}^{(n)}$ denote the states of $\mathbf{X}_{\mathbf{bc}}^{(n)}$ and $\mathbf{X}_{\mathbf{m} \in pa(h_r)}^{(n)}$ in simulation n . Then, we compute, for each s_r and s_{h_r} ,

$$\begin{aligned} P r^{(n)}(X_{h_r} = s_{h_r} \mid \mathbf{X}_{\mathbf{bc}} = \mathbf{s}_{\mathbf{bc}}^{(n)}, \mathbf{X}_{\mathbf{m} \in pa(h_r)} = \mathbf{s}_{\mathbf{m} \in pa(h_r)}^{(n)}, X_r = s_r) \\ = \frac{\frac{m_{s_r, s_{h_r}}^{(n)}}{m_r} + p_{X_{h_r} \mid \mathbf{X}_{\mathbf{bc}}^{(n)}, \mathbf{X}_{\mathbf{m} \in pa(h_r)}^{(n)}, X_r}^{(n-1)} \cdot E^{(n-1)}[\mathbf{s}_{\mathbf{bc}}^{(n)}, \mathbf{s}_{\mathbf{m} \in pa(h_r)}^{(n)}]}{E^{(n)}[\mathbf{s}_{\mathbf{bc}}^{(n)}, \mathbf{s}_{\mathbf{m} \in pa(h_r)}^{(n)}]}, \end{aligned} \quad (\text{A.2})$$

where $m_{s_r, s_{h_r}}^{(n)}$ is the number of receptors of type r in zone s_r that experience hazard X_{h_r} in state s_{h_r} in simulation n . Further, we used $p_{X_{h_r} \mid \mathbf{X}_{\mathbf{bc}}^{(n)}, \mathbf{X}_{\mathbf{m} \in pa(h_r)}^{(n)}, X_r}^{(n-1)}$ to denote $P r^{(n-1)}(X_{h_r} = s_{h_r} \mid \mathbf{X}_{\mathbf{bc}} = \mathbf{s}_{\mathbf{bc}}^{(n)}, \mathbf{X}_{\mathbf{m} \in pa(h_r)} = \mathbf{s}_{\mathbf{m} \in pa(h_r)}^{(n)}, X_r = s_r)$. Thus, according to equation A.2, the CPT entries represent average proportions of receptors of type r subjected to s_{h_r} over all simulations with identical states for boundary conditions and DRR measures.

B

BN FOR WELLS-NEXT-THE-SEA IN GENIE

As graphical user interface for the BNs we used GeNIe, which is available free of charge for academic research and teaching use from BayesFusion, LLC, <http://www.bayesfusion.com/>. Because the state names of the nodes in GeNIe must be strings with no spaces and no symbols, we resorted to special characters to represent positive or negative values, decimals and ranges (Table B.1).

Figure B.1 shows the quantified BNs with its prior distributions and Figure B.2 shows the quantified BN conditioned on the states corresponding to a 1 in 100 year return storm (water level =4.41m; wave height =2.17m).

Special character	Meaning
p	+
m	−
–	·
—	to

Table B.1: Special characters used in state names in GeNIe

B




Figure B.1: Bayesian Network for Wells-next-the-Sea with prior distributions




Figure B.2: Bayesian Network for Wells-next-the-Sea showing a 100 year return period event

C

SAMPLING ALGORITHMS FOR VINE-COPULAS

This appendix presents the sampling algorithms for the three vine-copulas that have been used in this article. Conditional copulas are denoted by the following *h-function*:

$$h(u, v, \theta) = \frac{\partial C_{u,v}(u, v, \theta)}{\partial v}, \quad (\text{C.1})$$

where $u, v, \in [0, 1]$ and θ denotes the set of parameters for the copula of the joint distribution function of u and v . Thus, the second parameter of $h(\cdot)$ always corresponds to the conditioning variable. The inverse of the h -function with respect to u is denoted by h^{-1} and corresponds to the inverse conditional copula. We recall that n is the number of variables and k is the order of the Markov property.

Algorithm 1 (modified COPAR for $n = 2$ and $k = 1$)

1. Simulate $2T$ random variates $s_1, \dots, s_{2T}$ from $U(0, 1)$.
2. Set $x_1^{(1)} = s_1$
3. Set $x_1^{(2)} = h^{-1}(s_1, x_1^{(1)}, \theta_{X_t^{(1)}, X_t^{(2)}})$.

For $i = 2, \dots, T$:

4. Set $x_i^{(1)} = h^{-1}(h^{-1}(s_{2i-1}, h(x_{i-1}^{(2)}, x_{i-1}^{(1)}, \theta_{X_t^{(1)}, X_{t+1}^{(1)} | X_t^{(1)}}), x_{i-1}^{(1)}, \theta_{X_t^{(1)}, X_t^{(2)}})$
5. Set $x_i^{(2)} = h^{-1}(h^{-1}(h^{-1}(s_{2i}, h(x_{i-1}^{(2)}, x_{i-1}^{(1)}, \theta_{X_t^{(1)}, X_t^{(2)}}), h(x_i^{(1)}, x_{i-1}^{(1)}, \theta_{X_t^{(1)}, X_{t+1}^{(1)}}), \theta_{X_t^{(2)}, X_{t+1}^{(2)} | X_t^{(1)}, X_{t+1}^{(1)}}) \theta_{X_t^{(1)}, X_{t+1}^{(1)}}), x_{i-1}^{(1)}, \theta_{X_{t+1}^{(2)}, X_t^{(1)}})$

Algorithm 2 (Alternating D-Vine for $n = 2$ and $k = 1$)

1. Simulate $2T$ random variates $s_1, \dots, s_{2T}$ from $U(0, 1)$.
2. Set $x_1^{(1)} = s_1$.
3. Set $x_1^{(2)} = h^{-1}(s_2, x_1^{(1)}, \theta_{X_t^{(1)}, X_t^{(2)}})$.

For $i = 2, \dots, T$:

4. Set $x_i^{(1)} = h^{-1}(h^{-1}(s_{2i-1}, h(x_{i-1}^{(1)}, x_{i-1}^{(2)}, \theta_{X_t^{(1)}, X_t^{(2)}}), \theta_{X_{t+1}^{(1)}, X_t^{(1)} | X_t^{(2)}}), x_{i-1}^{(2)}, \theta_{X_{t+1}^{(1)}, X_t^{(2)}})$
5. Set $x_i^{(2)} = h^{-1}(h^{-1}(h^{-1}(s_{2i-2}, h(h(x_{i-1}^{(1)}, x_{i-1}^{(2)}, \theta_{X_t^{(1)}, X_t^{(2)}}), h(x_i^{(1)}, x_{i-1}^{(2)}, \theta_{X_{t+1}^{(1)}, X_t^{(2)}}), \theta_{X_t^{(1)}, X_{t+1}^{(2)} | X_t^{(2)}, X_{t+1}^{(1)}}), h(x_{i-1}^{(2)}, x_i^{(1)}, \theta_{X_t^{(2)}, X_{t+1}^{(1)}}), \theta_{X_t^{(1)}, X_{t+1}^{(2)} | X_t^{(1)}}), x_i^{(1)}, \theta_{X_t^{(2)}, X_t^{(1)}})$

Algorithm 3 (Branching D-vine for $n = 2$ and $k = 1$)

C

1. Simulate $2T$ random variates $s_1, \dots, s_{2T}$ from $U(0, 1)$.
2. Set $x_1^{(1)} = s_1$.
3. Set $x_1^{(2)} = h^{-1}(s_2, x_1^{(1)}, \theta_{X_t^{(1)}, X_t^{(2)}})$
4. Set $x_2^{(1)} = h^{-1}(h^{-1}(s_3, h(x_1^{(2)}, x_1^{(1)}, \theta_{X_t^{(1)}, X_t^{(2)}}), \theta_{X_{t+1}^{(1)}, X_t^{(2)} | X_t^{(1)}}), x_1^{(1)}, \theta_{X_t^{(1)}, X_{t+1}^{(1)}})$
5. Set $x_2^{(2)} = h^{-1}(h^{-1}(h^{-1}(s_4, h(h(x_2^{(1)}, x_1^{(1)}, \theta_{X_{t+1}^{(1)}, X_t^{(1)}}), h(x_1^{(2)}, x_1^{(1)}, \theta_{X_t^{(2)}, X_t^{(1)}}), \theta_{X_{t+1}^{(1)}, X_t^{(2)} | X_t^{(1)}}), \theta_{X_{t+1}^{(1)}, X_{t+1}^{(2)} | X_t^{(1)}, X_t^{(2)}}), h(x_1^{(1)}, x_1^{(2)}, \theta_{X_t^{(1)}, X_t^{(2)}}), \theta_{X_{t+1}^{(1)}, X_{t+1}^{(2)} | X_t^{(2)}}), x_1^{(2)}, \theta_{X_{t+1}^{(2)}, X_t^{(2)}})$

For $i = 3, \dots, T$:

6. Set $x_i^{(1)} = h^{-1}(s_{2i-1}, x_{i-1}^{(1)}, \theta_{X_t^{(1)}, X_{t+1}^{(1)}})$
7. Set $x_i^{(2)} = h^{-1}(h^{-1}(s_{2i}, h(x_i^{(1)}, x_{i-1}^{(1)}, \theta_{X_{t+1}^{(1)}, X_t^{(1)}}), \theta_{X_{t+1}^{(1)}, X_{t+1}^{(2)} | X_t^{(1)}, X_t^{(2)}}), x_{i-1}^{(2)}, \theta_{X_{t+1}^{(2)}, X_t^{(2)}})$

D

BACKGROUND ON SELECTED COPULA FAMILIES

D.1. TAWN

The Tawn copula [165] is an asymmetric extension of the Gumbel copula and its cdf is given by

$$C(u_1, u_2) = \exp\{\log(u_1 u_2) A\left(\frac{\log(u_1)}{\log(u_1 u_2)}\right)\}, \quad u_1, u_2 \in [0, 1], \quad (\text{D.1})$$

where $A(t) = 1 - \beta + (\beta - \alpha)t + \{\alpha^r t^r + \beta^r (1 - t)^r\}^{\frac{1}{r}}$, $0 \leq \alpha, \beta \leq 1$, $1 \leq r < \infty$, $t \in [0, 1]$. Two simplified versions of the Tawn copula with two parameters each are implemented in R's VineCopula package [148]. They are called "Tawn type 1" and "Tawn type 2" and have one of the asymmetry parameters fixed to 1 such that the dependence is either left- or right-skewed.

D.2. GAMMA 1-FACTOR MODEL

A bivariate gamma 1-factor model has the form

$$X_j = Z_0 + Z_j, \quad j = 1, 2, \quad (\text{D.2})$$

where Z_0 , Z_1 and Z_2 are independent gamma distributed variables with shape parameters θ_1 , θ_2 and θ_3 and scale parameters all equal to 1. The copula cdf is [85, p.211]

$$C_{\theta_1, \theta_2, \theta_3}(u_1, u_2) = F(F_{\theta_0 + \theta_1}^{-1}(u_1), F_{\theta_0 + \theta_2}^{-1}(u_2)), \quad u_1, u_2 \in [0, 1]. \quad (\text{D.3})$$

Here F is the joint cdf of X_1 and X_2 , while $F_{\theta_0 + \theta_1}$ and $F_{\theta_0 + \theta_2}$ are the marginal distribution functions of X_1 and X_2 obtained by convoluting the marginal distributions functions of Z_0 , Z_1 and Z_2 . This model is implemented in the R's CopulaModel package [86].

D.3. SKEW-T

The skew-t copula derives from the skew-t distribution, which is a normal mean variance mixture. While several formulations exist, only the formulation by Azzalini and Capitanio [10] is considered here. For this formulation an estimation procedure based

on maximum likelihood has been developed by Yoshiba [191], who also published the corresponding R code. The copula cdf is given by

$$C_{\rho, \delta_1, \delta_2, \nu}(u_1, u_2) = St(St_1^{-1}(u_1; 0, 1, \delta_1, \nu), St_2^{-1}(u_2; 0, 1, \delta_2, \nu); 0, \rho, \alpha, \nu), \quad u_1, u_2 \in [0, 1], \quad (\text{D.4})$$

where St is the multivariate skew- t distribution with correlation parameter ρ , a transformed skewness vector α and degrees of freedom ν . St_1 and St_2 are the univariate margins of this distribution and δ_1 and δ_2 are the respective skewness parameters.

REFERENCES

- [1] Aas, K., Czado, C., Frigessi, A., and Bakken, H. (2009). Pair-copula constructions of multiple dependence. *Insurance: Mathematics and economics*, 44(2):182–198.
- [2] Acar, E. F., Genest, C., and Nešlehová, J. (2012). Beyond simplified pair-copula constructions. *Journal of Multivariate Analysis*, 110:74–90.
- [3] Acero, W. G., Li, L., Gao, Z., and Moan, T. (2016). Methodology for assessment of the operational limits and operability of marine operations. *Ocean Engineering*, 125:308–327.
- [4] Ahern, M., Kovats, R. S., Wilkinson, P., Few, R., and Matthies, F. (2005). Global health impacts of floods: epidemiologic evidence. *Epidemiologic Reviews*, 27(1):36–46.
- [5] Ale, B. (1991). Risk analysis and risk policy in the Netherlands and the EEC. *Journal of Loss Prevention in the Process Industries*, 4(1):58–64.
- [6] Ale, B. (2005). Tolerable or acceptable: a comparison of risk regulation in the United Kingdom and in the Netherlands. *Risk analysis*, 25(2):231–241.
- [7] Alfieri, L., Salamon, P., Pappenberger, F., Wetterhall, F., and Thielen, J. (2012). Operational early warning systems for water-related hazards in Europe. *Environmental Science & Policy*, 21:35–49.
- [8] Angnuureng, D. B., Almar, R., Senechal, N., Castelle, B., Addo, K. A., Marieu, V., and Ranasinghe, R. (2017). Shoreline resilience to individual storms and storm clusters on a meso-macrotidal barred beach. *Geomorphology*, 290:265–276.
- [9] Athanassoulis, G. and Stefanakos, C. N. (1995). A nonstationary stochastic model for long-term time series of significant wave height. *Journal of Geophysical Research: Oceans*, 100(C8):16149–16162.
- [10] Azzalini, A. and Capitanio, A. (2003). Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):367–389.
- [11] Barquet, K., Dickin, S. K., Meijer, J. J., and Dastgheib, A. (2018). Testing RISC-KIT’s integrated approach for assessing disaster risk reduction measures on the coast of Kristianstad, Sweden. *Coastal Engineering*, 134:203 – 211.
- [12] Bates, P. D., Dawson, R. J., Hall, J. W., Horritt, M. S., Nicholls, R. J., Wicks, J., and Hassan, M. A. A. M. (2005). Simplified two-dimensional numerical modelling of coastal flooding and example applications. *Coastal Engineering*, 52(9):793–810.

- [13] Bates, P. D. and De Roo, A. (2000). A simple raster-based model for flood inundation simulation. *Journal of Hydrology*, 236(1):54–77.
- [14] Baxter, P. J. (2005). The east coast Big Flood, 31 January - 1 February 1953: a summary of the human disaster. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 363(1831):1293–1312.
- [15] Bedford, T. and Cooke, R. M. (2001). *Mathematical tools for probabilistic risk analysis*. Cambridge University Press.
- [16] Bedford, T. and Cooke, R. M. (2002). Vines: a new graphical model for dependent random variables. *Annals of Statistics*, pages 1031–1068.
- [17] Bevacqua, E., Maraun, D., Haff, I. H., Widmann, M., and Vrac, M. (2017). Multivariate statistical modelling of compound events via pair-copula constructions: analysis of floods in Ravenna (Italy). *Hydrology and Earth System Sciences*, 21(6):2701.
- [18] Blaser, L., Ohrnberger, M., Riggelsen, C., Babeyko, A., and Scherbaum, F. (2011). Bayesian networks for tsunami early warning. *Geophysical Journal International*, 185(3):1431–1443.
- [19] Boccotti, P. (2000). *Wave mechanics for ocean engineering*. Elsevier Oceanography Series. Elsevier Science.
- [20] Bogaard, T., De Kleermaeker, S., Jaeger, W. S., and van Dongeren, A. (2016). Development of generic tools for coastal early warning and decision support. In *E3S Web of Conferences*, volume 7, page 18017. EDP Sciences.
- [21] Bolle, A., das Neves, L., Smets, S., Mollaert, J., and Buitrago, S. (2018). An impact-oriented early warning and Bayesian-based decision support system for flood risks in Zeebrugge harbour. *Coastal Engineering*, 134:191 – 202.
- [22] Booij, N., Haagsma, I., Holthuijsen, L., Kieftenburg, A., Ris, R., van der Westhuysen, A., and Zijlema, M. (2014). *SWAN user manual, SWAN cycle III version 41.01*. Delft University of Technology.
- [23] Booij, N., Holthuijsen, L., and Ris, R. (1996). The “SWAN” wave model for shallow water. In *25th International Conference on Coastal Engineering*, pages 668–676.
- [24] Box, G. E. and Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 211–252.
- [25] Box, G. E., Jenkins, G. M., Reinsel, G. C., and Ljung, G. M. (2015). *Time series analysis: Forecasting and control*. John Wiley & Sons.
- [26] Brechmann, E. C. and Czado, C. (2015). COPAR - multivariate time series modeling using the copula autoregressive model. *Applied Stochastic Models in Business and Industry*, 31(4):495–514.
- [27] Brockwell, P. J. and Davis, R. A. (2016). *Introduction to time series and forecasting*. springer.

- [28] Brouwer, R. and Van Ek, R. (2004). Integrated ecological, economic and social impact assessment of alternative flood control policies in the Netherlands. *Ecological Economics*, 50(1):1–21.
- [29] Burzel, A., Dassanayake, D. R., Naulin, M., Kortenhaus, A., Oumeraci, H., Wahl, T., Mudersbach, C., Jensen, J., Gönner, G., Sossidi, K., et al. (2010). Integrated flood risk analysis for extreme storm surges (XtremRisk). In *Coastal Engineering Proceedings*, 32.
- [30] Cai, Y., Gouldby, B., Hawkes, P., and Dunning, P. (2008). Statistical simulation of flood variables: Incorporating short-term sequencing. *Journal of Flood Risk Management*, 1(1):3–12.
- [31] Callaghan, D., Nielsen, P., Short, A., and Ranasinghe, R. (2008). Statistical simulation of wave climate and extreme beach erosion. *Coastal Engineering*, 55(5):375–390.
- [32] Centraal Bureau voor de Statistiek (2017). Bevolking, huishoudens en bevolkingsontwikkeling; vanaf 1899.
- [33] Christie, E. K., Spencer, T., Owen, D., McIvor, A., Möller, I., and Viavattene, C. (2018). Regional coastal flood risk assessment for a tidally dominant, natural coastal setting: North Norfolk, southern North Sea. *Coastal Engineering*, 134:177–190.
- [34] Coles, S. (2001). *An introduction to statistical modeling of extreme values*, volume 208. Springer.
- [35] Cooke, R. M. (1991). *Experts in uncertainty: opinion and subjective probability in science*. Oxford University Press.
- [36] Cooke, R. M. (1997). Markov and entropy properties of tree- and vine-dependent variables. In *Proceedings of the ASA Section of Bayesian Statistical Science*, volume 27.
- [37] Corbella, S. and Stretch, D. D. (2012). Predicting coastal erosion trends using non-stationary statistics and process-based models. *Coastal Engineering*, 70:40–49.
- [38] Corbella, S. and Stretch, D. D. (2013). Simulating a multivariate sea storm using Archimedean copulas. *Coastal Engineering*, 76:68–78.
- [39] CoreTeam, R. (2015). R: A language and environment for statistical Computing. Vienna, Austria: R Foundation for Statistical Computing; 2015.
- [40] Cruz, A. M. and Krausmann, E. (2008). Damage to offshore oil and gas facilities following Hurricanes Katrina and Rita: an overview. *Journal of Loss Prevention in the Process Industries*, 21(6):620–626.
- [41] Cumiskey, L., Priest, S., Valchev, N., Costas, S., and Clarke, J. (2018). A framework to include the (inter)dependencies of disaster risk reduction measures in coastal risk assessment. *Coastal Engineering*, 143:81 – 92.
- [42] Cunha, C. and Soares, C. G. (1999). On the choice of data transformation for modelling time series of significant wave height. *Ocean Engineering*, 26(6):489–506.

- [43] Davies, G., Callaghan, D. P., Gravois, U., Jiang, W., Hanslow, D., Nichol, S., and Bal-dock, T. (2017). Improved treatment of non-stationary conditions and uncertainties in probabilistic models of storm wave climate. *Coastal Engineering*, 127:1–19.
- [44] Dawson, R. J., Dickson, M. E., Nicholls, R. J., Hall, J. W., Walkden, M. J., Stansby, P. K., Mokrech, M., Richards, J., Zhou, J., Milligan, J., et al. (2009). Integrated analysis of risks of coastal flooding and cliff erosion under scenarios of long term change. *Climatic Change*, 95(1-2):249–288.
- [45] De Michele, C., Salvadori, G., Passoni, G., and Vezzoli, R. (2007a). A multivariate model of sea storms using copulas. *Coastal Engineering*, 54(10):734–751.
- [46] De Michele, C., Salvadori, G., Passoni, G., and Vezzoli, R. (2007b). A multivariate model of sea storms using copulas. *Coastal Engineering*.
- [47] Delta Committee (2008). Working together with water - findings of the Deltacom-missie 2008. *Delta Committee The Netherlands*.
- [48] Den Heijer, C. (2013). *The role of bathymetry, wave obliquity and coastal curva-ture in dune erosion prediction*. Doctoral dissertation, Delft University of Technology, Delft.
- [49] Dissmann, J., Brechmann, E. C., Czado, C., and Kurowicka, D. (2013). Selecting and estimating regular vine copulae and application to financial returns. *Computational Statistics & Data Analysis*, 59:52–69.
- [50] Environment Agency (2007). Public response to flood warnings. Technical Report, SC020116.
- [51] Erhardt, T. M. and Czado, C. (2017). Standardized drought indices: a novel uni-variate and multivariate approach. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*.
- [52] Ernsteins, R. (2010). Participation and integration are key to coastal management. *Coastal Management*, 52(12):636–645.
- [53] Evans, E., Ashley, R., Hall, J., Penning-Roswell, E., Saul, A., Sayers, P., Thorne, C., and Watkinson, A. (2004). *Foresight: Future flooding. Scientific summary: Volume I, Future risks and their drivers*. Office of Science and Technology, London.
- [54] Faltinsen, O. (1993). *Sea loads on ships and offshore structures*, volume 1. Cambridge University Press.
- [55] Feagin, R. A., Sherman, D. J., and Grant, W. E. (2005). Coastal erosion, global sea-level rise, and the loss of sand dune plant habitats. *Frontiers in Ecology and the Envi-ronment*, 3(7):359–364.
- [56] Ferreira, O., Viavattene, C., Jiménez, J., Bolle, A., Das Neves, L., Plomaritis, T., Mc-Call, R., and Van Dongeren, A. (2018). Storm-induced risk assessment: evaluation of two tools at the regional and hotspot scale. *Coastal Engineering*, 134:241 – 243.

- [57] Fienen, M. N. and Plant, N. G. (2015). A cross-validation package driving Netica with Python. *Environmental Modelling & Software*, 63:14–23.
- [58] Fiksel, J. (2006). Sustainability and resilience: toward a systems approach. *Sustainability: Science, Practice and Policy*, 2(2):14–21.
- [59] Friedman, N. and Goldszmidt, M. (1996). Discretizing continuous attributes while learning Bayesian networks. In *ICML*, pages 157–165.
- [60] Garrote, L., Molina, M., and Mediero, L. (2006). Probabilistic forecasts using Bayesian networks calibrated with deterministic rainfall-runoff models. In *Extreme Hydrological Events: New Concepts for Security*, pages 173–183. Springer.
- [61] Genest, C. and Favre, A.-C. (2007). Everything you always wanted to know about copula modeling but were afraid to ask. *Journal of Hydrologic Engineering*, 12(4):347–368.
- [62] Gouldby, B. and Samuels, P. (2005). Language of risk-project definitions. FLOODsite Project Report, T32-04-01.
- [63] Guanche, Y., Mínguez, R., and Méndez, F. (2013). Climate-based Monte Carlo simulation of trivariate sea states. *Coastal Engineering*, 80:107–121.
- [64] Guedes Soares, C. and Cunha, C. (2000). Bivariate autoregressive models for the time series of significant wave height and mean period. *Coastal Engineering*, 40(4):297–311.
- [65] Guedes Soares, C. and Ferreira, A. (1996). Representation of non-stationary time series of significant wave height with autoregressive models. *Probabilistic Engineering Mechanics*, 11(3):139–148.
- [66] Guedes Soares, C., Ferreira, A., and Cunha, C. (1996). Linear models of the time series of significant wave height on the Southwest Coast of Portugal. *Coastal Engineering*, 29(1-2):149–167.
- [67] Gutierrez, B. T., Plant, N. G., and Thieler, E. R. (2011). A Bayesian network to predict coastal vulnerability to sea level rise. *Journal of Geophysical Research*, 116(F2).
- [68] Haff, I. H., Aas, K., and Frigessi, A. (2010). On the simplified pair-copula construction - simply useful or too simplistic? *Journal of Multivariate Analysis*, 101(5):1296–1310.
- [69] Hajat, S., Ebi, K., Kovats, R., Menne, B., Edwards, S., and Haines, A. (2005). The human health consequences of flooding in Europe: a review. In *Extreme weather events and public health responses*, pages 185–196. Springer.
- [70] Hallegatte, S., Green, C., Nicholls, R. J., and Corfee-Morlot, J. (2013). Future flood losses in major coastal cities. *Nature Climate Change*, 3(9):802–806.
- [71] Hapke, C. and Plant, N. (2010). Predicting coastal cliff erosion using a Bayesian probabilistic model. *Marine geology*, 278(1):140–149.

- [72] Henriksen, H. J., Rasmussen, P., Brandt, G., von Bülow, D., and Jensen, F. V. (2007). Public participation modelling using Bayesian networks in management of ground-water contamination. *Environmental Modelling & Software*, 22(8):1101–1113.
- [73] Hervouet, J.-M. (2000). TELEMAC modelling system: an overview. *Hydrological Processes*, 14(13):2209–2210.
- [74] Hinkel, J., Lincke, D., Vafeidis, A. T., Perrette, M., Nicholls, R. J., Tol, R. S., Marzeion, B., Fettweis, X., Ionescu, C., and Levermann, A. (2014). Coastal flood damage and adaptation costs under 21st century sea-level rise. *Proceedings of the National Academy of Sciences*, 111(9):3292–3297.
- [75] Holdgate, M. W. (1980). *A perspective of environmental pollution*. Cambridge University Press Archive.
- [76] Holman, R. (1986). Extreme value statistics for wave run-up on a natural beach. *Coastal Engineering*, 9(6):527–544.
- [77] Holthuijsen, L. H. (2010). *Waves in oceanic and coastal waters*. Cambridge University Press.
- [78] Jäger, W. S., Christie, E., Hanea, A., den Heijer, K., and Spencer, T. (2018). A Bayesian network approach for coastal risk analysis and decision making. *Coastal Engineering*, 134:48 – 61.
- [79] Jäger, W. S. and Morales Nápoles, O. (2015). Sampling joint time series of significant wave heights and periods in the North Sea. In Podofillini, L., Sudret, B., Stojadinović, B., Zio, E., and Kröger, W., editors, *Safety and Reliability of Complex Engineered Systems*.
- [80] Jäger, W. S. and Morales Nápoles, O. (2017). A vine-copula model for time series of significant wave heights and mean zero-crossing periods in the North Sea. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, 3(4):04017014.
- [81] Jäger, W. S., Nagler, T., Czado, C., and McCall, R. T. (n.d.). A statistical simulation model for joint time series of non-stationary hourly wave parameters. (*submitted for publication*).
- [82] JBA Consulting (2012). Establishing the cost effectiveness of property flood protection. Technical Report, FD2657.
- [83] Jensen, F. (1996). *An introduction to Bayesian networks*. University College London Press, London.
- [84] Joe, H. (1996). Families of m-variate distributions with given margins and m (m-1)/2 bivariate dependence parameters. *Lecture Notes-Monograph Series*, pages 120–141.
- [85] Joe, H. (2014). *Dependence modeling with copulas*. CRC Press.

- [86] Joe, H. and Krupskii, P. (2014). *CopulaModel: dependence modeling with copulas*. R package version 0.6.
- [87] Jongejan, R. and Maaskant, B. (2015). Quantifying flood risks in the Netherlands. *Risk Analysis*, 35(2):252–264.
- [88] Jonkman, S. and Kelman, I. (2005). Deaths during the 1953 North Sea storm surge.
- [89] Jonkman, S., Kok, M., Van Ledden, M., and Vrijling, J. (2009a). Risk-based design of flood defence systems: a preliminary analysis of the optimal protection level for the New Orleans metropolitan area. *Journal of Flood Risk Management*, 2(3):170–181.
- [90] Jonkman, S., Kok, M., and Vrijling, J. (2003). Economic optimisation as a basis for the choice of flood protection strategies in the Netherlands. In *International association for XXX IAHR Congress*, pages 19–26.
- [91] Jonkman, S. N., Bočkarjova, M., Kok, M., and Bernardini, P. (2008a). Integrated hydrodynamic and economic modelling of flood damage in the Netherlands. *Ecological Economics*, 66(1):77–90.
- [92] Jonkman, S. N., Kok, M., and Vrijling, J. K. (2008b). Flood risk assessment in the Netherlands: A case study for dike ring South Holland. *Risk Analysis*, 28(5):1357–1374.
- [93] Jonkman, S. N., Maaskant, B., Boyd, E., and Levitan, M. L. (2009b). Loss of life caused by the flooding of New Orleans after Hurricane Katrina: analysis of the relationship between flood characteristics and mortality. *Risk Analysis*, 29(5):676–698.
- [94] Jonkman, S. N., Vrijling, J. K., and Vrouwenvelder, A. C. W. M. (2008c). Methods for the estimation of loss of life due to floods: a literature review and a proposal for a new method. *Natural Hazards*, 46(3):353–389.
- [95] Karunaratna, H., Pender, D., Ranasinghe, R., Short, A. D., and Reeve, D. E. (2014). The effects of storm clustering on beach profile variability. *Marine Geology*, 348:103–112.
- [96] Knapp, R., Rhome, J., and Brown, D. (2005). Tropical Cyclone Report. Hurricane Katrina.
- [97] Koller, D. and Friedman, N. (2009). *Probabilistic graphical models: principles and techniques*. MIT press.
- [98] Kozlov, A. V. and Koller, D. (1997). Nonuniform dynamic discretization in hybrid networks. In *Proceedings of the Thirteenth conference on Uncertainty in artificial intelligence*, pages 314–325. Morgan Kaufmann Publishers Inc.
- [99] Kurgan, L. and Cios, K. J. (2001). Discretization algorithm that uses class-attribute interdependence maximization. In *Proceedings of the 2001 International Conference on Artificial Intelligence (IC-AI 2001)*, pages 980–987.
- [100] Kurowicka, D. and Cooke, R. M. (2006). *Uncertainty analysis with high dimensional dependence modelling*. John Wiley & Sons.

- [101] Lauritzen, S. L. and Spiegelhalter, D. J. (1988). Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 157–224.
- [102] Leontaris, G., Morales-Nápoles, O., and Wolfert, A. R. (2016). Probabilistic scheduling of offshore operations using copula based environmental time series - an application for cable installation management for offshore wind farms. *Ocean Engineering*, 125:328–341.
- [103] Li, F., van Gelder, P., Ranasinghe, R., Callaghan, D., and Jongejan, R. (2014a). Probabilistic modelling of extreme storms along the Dutch coast. *Coastal Engineering*, 86:1–13.
- [104] Li, F., van Gelder, P., Vrijling, J., Callaghan, D., Jongejan, R., and Ranasinghe, R. (2014b). Probabilistic estimation of coastal dune erosion and recession by statistical simulation of storm events. *Applied Ocean Research*, 47:53–62.
- [105] Lin, N. and Emanuel, K. (2016). Grey swan tropical cyclones. *Nature Climate Change*, 6(1):106–111.
- [106] Lin-Ye, J., Garcia-Leon, M., Gracia, V., and Sanchez-Arcilla, A. (2016). A multivariate statistical model of extreme events: an application to the Catalan coast. *Coastal Engineering*, 117:138–156.
- [107] Lins, H. F. and Cohn, T. A. (2011). Stationarity: Wanted dead or alive? *JAWRA Journal of the American Water Resources Association*, 47(3):475–480.
- [108] Marchand, M., Sanchez-Arcilla, A., Ferreira, M., Gault, J., Jiménez, J. A., Markovic, M., Mulder, J., van Rijn, L., Stănică, A., Sulisz, W., et al. (2011). Concepts and science for coastal erosion management - An introduction to the Conscience framework. *Ocean & coastal management*, 54(12):859–866.
- [109] Maris, A., De Blocq van Kuffeler, V., Harmsen, W., Jansen, P., Nijhoff, G., Thijsse, J., Van der Wal, L., et al. (1961). Rapport Deltacommissie. Deel 1. Eindverslag en interimadviezen.
- [110] Martín-Hidalgo, M., Martín-Soldevilla, M. J., Negro, V., Aberturas, P., and López-Gutiérrez, J. (2014). Storm evolution characterization for analysing stone armour damage progression. *Coastal Engineering*, 85:1–11.
- [111] Masina, M., Lamberti, A., and Archetti, R. (2015). Coastal flooding: a copula based approach for estimating the joint probability of water levels and waves. *Coastal Engineering*, 97:37–52.
- [112] McCall, R. T., De Vries, J. V. T., Plant, N., Van Dongeren, A., Roelvink, J., Thompson, D., and Reniers, A. (2010). Two-dimensional time dependent Hurricane overwash and erosion modeling at Santa Rosa Island. *Coastal Engineering*, 57(7):668–683.

- [113] McCann, R. K., Marcot, B. G., and Ellis, R. (2006). Bayesian belief networks: applications in ecology and natural resource management. *Canadian Journal of Forest Research*, 36(12):3053–3062.
- [114] Méndez, F. J., Menéndez, M., Luceño, A., and Losada, I. J. (2006). Estimation of the long-term variability of extreme significant wave height using a time-dependent peak over threshold (POT) model. *Journal of Geophysical Research: Oceans*, 111(C7).
- [115] Méndez, F. J., Menéndez, M., Luceño, A., and Losada, I. J. (2007). Analyzing monthly extreme sea levels with a time-dependent GEV model. *Journal of Atmospheric and Oceanic Technology*, 24(5):894–911.
- [116] Merz, B., Kreibich, H., Schwarze, R., and Thieken, A. (2010). Review article “Assessment of economic flood damage”. *Natural Hazards and Earth System Science*, 10(8):1697–1724.
- [117] Michel, W. H. (1999). Sea spectra revisited. *Marine Technology*, 36(4):211–227.
- [118] Milly, P. C., Betancourt, J., Falkenmark, M., Hirsch, R. M., Kundzewicz, Z. W., Lettenmaier, D. P., and Stouffer, R. J. (2008). Stationarity is dead: Whither water management? *Science*, 319(5863):573–574.
- [119] Möller, I. (2006). Quantifying saltmarsh vegetation and its effect on wave height dissipation: results from a UK East coast saltmarsh. *Estuarine, Coastal and Shelf Science*, 69(3):337–351.
- [120] Möller, I., Kudella, M., Rupprecht, F., Spencer, T., Paul, M., van Wesenbeeck, B. K., Wolters, G., Jensen, K., Bouma, T. J., Miranda-Lange, M., and Schimmels, S. (2014). Wave attenuation over coastal salt marshes under storm surge conditions. *Nature Geoscience*, 7(10):727–731.
- [121] Möller, I., Spencer, T., French, J., Leggett, D., and Dixon, M. (1999). Wave transformation over salt marshes: a field and numerical modelling study from North Norfolk, England. *Estuarine, Coastal and Shelf Science*, 49(3):411–426.
- [122] Montes-Iturrizaga, R. and Heredia-Zavoni, E. (2016). Multivariate environmental contours using c-vine copulas. *Ocean Engineering*, 118:68–82.
- [123] Morales-Nápoles, O. (2011). Counting vines. *Dependence Modeling: Vine Copula Handbook*, pages 189–218.
- [124] Morales-Nápoles, O., Cooke, R. M., and Kurowicka, D. (2010). About the number of vines and regular vines on n nodes. (working paper).
- [125] Morton, I. and Bowers, J. (1996). Extreme value analysis in a multivariate offshore environment. *Applied Ocean Research*, 18(6):303–317.
- [126] Nai Ruscone, M. and Osmetti, S. A. (2017). Modelling the dependence in multivariate longitudinal data by pair copula decomposition. In *Soft Methods for Data Science*, pages 373–380. Springer.

- [127] Narayan, S., Hanson, S., Nicholls, R., Clarke, D., Willems, P., Ntegeka, V., and Monbaliu, J. (2012). A holistic model for coastal flooding using system diagrams and the source-pathway-receptor (SPR) concept. *Natural Hazards and Earth System Science*, 12(5):1431–1439.
- [128] National Disaster Risk Reduction and Management Council (2013). Effects of Typhoon "Yolanda" (Haiyan).
- [129] Nelsen, R. B. (2013). *An introduction to copulas*, volume 139. Springer Science & Business Media.
- [130] Oumeraci, H., Kortenhaus, A., Burzel, A., Naulin, M., Dassanayake, D., Jensen, J., Wahl, T., Mudersbach, C., Gönnert, G., Gerkenmeier, B., et al. (2015). XtremRisk - integrated flood risk analysis for extreme storm surges at open coasts and in estuaries: methodology, key results and lessons learned. *Coastal Engineering Journal*, 57(01).
- [131] Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann.
- [132] Penning-Rowsell, E., Priest, S., Parker, D., Morris, J., Tunstall, S., Viavattene, C., Chatterton, J., and Owen, D. (2014). *Flood and coastal erosion risk management: a manual for economic appraisal*. Routledge.
- [133] Phillips, M. R. and Jones, A. L. (2006). Erosion and tourism infrastructure in the coastal zone: problems, consequences and management. *Tourism Management*, 27(3):517–524.
- [134] Plant, N. G. and Holland, K. T. (2011). Prediction and assimilation of surf-zone processes using a Bayesian network: Part I: Forward models. *Coastal Engineering*, 58(1):119–130.
- [135] Plomaritis, T. A., Costas, S., and Ferreira, Ó. (2018). Use of a Bayesian Network for coastal hazards, impact and disaster risk reduction assessment at a coastal barrier (Ria Formosa, Portugal). *Coastal Engineering*, 134:134 – 147.
- [136] Poelhekke, L., Jäger, W. S., van Dongeren, A., Plomaritis, T. A., McCall, R., and Ferreira, Ó. (2016). Predicting coastal hazards for sandy coasts with a Bayesian network. *Coastal Engineering*, 118:21–34.
- [137] Priest, S. J., Wilson, T., Tapsell, S. M., Penning-Rowsell, E. C., Viavattene, C., and Fernandez-Bilbao, A. (2007). Building a model to estimate risk to life for European flood events. Project Report.
- [138] Pullen, T., Allsop, N., Bruce, T., Kortenhaus, A., Schöter, H., and Van der Meer, J. (2007). *Wave overtopping of sea defences and related structure. Assessment Manual*.
- [139] Quartel, S., Ruessink, B., and Kroon, A. (2007). Daily to seasonal cross-shore behaviour of quasi-persistent intertidal beach morphology. *Earth Surface Processes and Landforms*, 32(9):1293–1307.

- [140] Repko, A., Van Gelder, P., Voortman, H., and Vrijling, J. (2004). Bivariate description of offshore wave conditions with physics-based extreme value statistics. *Applied Ocean Research*, 26(3):162–170.
- [141] Riggelsen, C. (2008). Learning Bayesian networks: a MAP criterion for joint selection of model structure and parameter. In *Data Mining, 2008. ICDM'08. Eighth IEEE International Conference on*, pages 522–529. IEEE.
- [142] Roelvink, D., Reniers, A., Van Dongeren, A., de Vries, J. v. T., McCall, R., and Lescinski, J. (2009). Modelling storm impacts on beaches, dunes and barrier islands. *Coastal engineering*, 56(11):1133–1152.
- [143] Rueda, A., Camus, P., Tomás, A., Vitousek, S., and Méndez, F. (2016). A multivariate extreme wave and storm surge climate emulator based on weather patterns. *Ocean Modelling*, 104:242–251.
- [144] Salvadori, G. and De Michele, C. (2006). Statistical characterization of temporal structure of storms. *Advances in Water Resources*, 29(6):827–842.
- [145] Salvadori, G., Tomasicchio, G., and D'Alessandro, F. (2014). Practical guidelines for multivariate analysis and design in coastal and off-shore engineering. *Coastal Engineering*, 88:1–14.
- [146] Sanuy, M., Duo, E., Jäger, W. S., Ciavola, P., and Jiménez, J. A. (2017). Linking source with consequences of coastal storm impacts for climate change and risk reduction scenarios for mediterranean sandy beaches. *Natural Hazards and Earth System Sciences Discussions*, 2017:1–57.
- [147] Sayers, P., Hall, J., and Meadowcroft, I. (2002). Towards risk-based flood hazard management in the UK. In *Proceedings of the Institution of Civil Engineers: Civil Engineering*, volume 150, pages 36–42.
- [148] Schepsmeier, U., Stoeber, J., Brechmann, E. C., Graeler, B., Nagler, T., Erhardt, T., Almeida, C., Min, A., Czado, C., Hofmann, M., et al. (2017). Package ‘VineCopula’.
- [149] Scotto, M. and Soares, C. G. (2000). Modelling the long-term time series of significant wave height with non-linear threshold models. *Coastal Engineering*, 40(4):313–327.
- [150] Serafin, K. A. and Ruggiero, P. (2014). Simulating extreme total water levels using a time-dependent, extreme value approach. *Journal of Geophysical Research: Oceans*, 119(9):6305–6329.
- [151] Serinaldi, F. and Grimaldi, S. (2007). Fully nested 3-copula: procedure and application on hydrological data. *Journal of Hydrologic Engineering*, 12(4):420–430.
- [152] Shumway, R. H. and Stoffer, D. S. (2017). *Time series analysis and its applications: with R examples*. Springer Science & Business Media.

- [153] Sills, G., Vroman, N., Wahl, R., and Schwanz, N. (2008). Overview of New Orleans levee failures: lessons learned and their impact on national levee design and assessment. *Journal of Geotechnical and Geoenvironmental Engineering*, 134(5):556–565.
- [154] Sklar, M. (1959). *Fonctions de répartition à n dimensions et leurs marges*. Publications de l'Institut Statistique de Université de Paris 8.
- [155] Smith, M., Min, A., Almeida, C., and Czado, C. (2012). Modeling longitudinal data using a pair-copula decomposition of serial dependence. *Journal of the American Statistical Association*.
- [156] Smith, M. S. (2015). Copula modelling of dependence in multivariate time series. *International Journal of Forecasting*, 31(3):815–833.
- [157] Solari, S. and Losada, M. (2011). Non-stationary wave height climate modeling and simulation. *Journal of Geophysical Research: Oceans*, 116(C9).
- [158] Solari, S. and van Gelder, P. (2011). On the use of vector autoregressive (VAR) and regime switching VAR models for the simulation of sea and wind state parameters. *Marine Technology and Engineering*, pages 217–230.
- [159] Soldevilla, M. J. M., Martín-Hidalgo, M., Negro, V., López-Gutiérrez, J., and Aberuras, P. (2015). Improvement of theoretical storm characterization for different climate conditions. *Coastal Engineering*, 96:71–80.
- [160] Spencer, T., Brooks, S. M., Evans, B. R., Tempest, J. A., and Möller, I. (2015). Southern North Sea storm surge event of 5 December 2013: water levels, waves and coastal impacts. *Earth-Science Reviews*, 146:120–145.
- [161] Steers, J., Stoddart, D., Bayliss-Smith, T., Spencer, T., and Durbidge, P. (1979). The storm surge of 11 January 1978 on the east coast of England. *Geographical Journal*, pages 192–205.
- [162] Stefanakos, C. N., Athanassoulis, G. A., and Barstow, S. F. (2006). Time series modeling of significant wave height in multiple scales, combining various sources of data. *Journal of Geophysical Research: Oceans*, 111(C10).
- [163] Stive, M. J., de Schipper, M. A., Luijendijk, A. P., Aarninkhof, S. G., van Gelder-Maas, C., van Thiel de Vries, J. S., de Vries, S., Henriquez, M., Marx, S., and Ranasinghe, R. (2013). A new alternative to saving our beaches from sea-level rise: the sand engine. *Journal of Coastal Research*, 29(5):1001–1008.
- [164] Stöber, J. and Czado, C. (2014). Regime switches in the dependence structure of multidimensional financial data. *Computational Statistics & Data Analysis*, 76:672–686.
- [165] Tawn, J. A. (1988). Bivariate extreme value theory: models and estimation. *Biometrika*, 75(3):397–415.

- [166] Thurston, N., Finlinson, B., Breakspear, R., Williams, N., Shaw, J., and Chatterton, J. (2008). Developing the evidence base for flood resistance and resilience. R&D Technical Report. FD2607/TR1.
- [167] Tolo, S., Patelli, E., and Beer, M. (2016). Risk assessment of spent nuclear fuel facilities considering climate change. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, page G4016003.
- [168] United Nations (2009). UNISDR terminology on disaster risk reduction.
- [169] Valchev, N., Eftimova, P., and Andreeva, N. (2018). Implementation and validation of a multi-domain coastal hazard forecasting system in an open bay. *Coastal Engineering*, 134:212 – 228.
- [170] van Dongeren, A., Ciavola, P., Martinez, G., Viavattene, C., Bogaard, T., Higgins, R., and McCall, R. (2018). Introduction to RISC-KIT: Resilience-increasing strategies for coasts. *Coastal Engineering*, 134:2 – 9.
- [171] van Gent, M., van Thiel de Vries, J., Coeveld, E., De Vroeg, J., and Van de Graaff, J. (2008). Large-scale dune erosion tests to study the influence of wave periods. *Coastal Engineering*, 55(12):1041–1051.
- [172] van Verseveld, H., van Dongeren, A., Plant, N., Jäger, W., and den Heijer, C. (2015). Modelling multi-hazard hurricane damages on an urbanized coast with a Bayesian network approach. *Coastal Engineering*, 103:1–14.
- [173] Vanem, E. (2016). Joint statistical models for significant wave height and wave period in a changing climate. *Marine Structures*, 49:180–205.
- [174] Vellinga, P. (1982). Beach and dune erosion during storm surges. *Coastal Engineering*, 6(4):361–387.
- [175] Viavattene, C., Micou, A. P., Owen, D., Priest, S., and Parker, D. (2015). Coastal vulnerability indicator library. RISC-KIT Deliverable D2.2.
- [176] Visser, P. J. (1998). *Breach growth in sand-dikes*. PhD thesis, Delft University of Technology, Delft.
- [177] Vitousek, S., Barnard, P. L., Fletcher, C. H., Frazer, N., Erikson, L., and Storlazzi, C. D. (2017a). Doubling of coastal flooding frequency within decades due to sea-level rise. *Scientific reports*, 7(1):1399.
- [178] Vitousek, S., Barnard, P. L., Limber, P., Erikson, L., and Cole, B. (2017b). A model integrating longshore and cross-shore processes for predicting long-term shoreline response to climate change. *Journal of Geophysical Research: Earth Surface*, 122(4):782–806.
- [179] Vousdoukas, M. I., Voukouvalas, E., Annunziato, A., Giardino, A., and Feyen, L. (2016). Projections of extreme storm surge levels along Europe. *Climate Dynamics*, 47(9):3171–3190.

- [180] Vrijling, J. (2001). Probabilistic design of water defense systems in The Netherlands. *Reliability engineering & system safety*, 74(3):337–344.
- [181] Wagenaar, D., De Bruijn, K., Bouwer, L., and De Moel, H. (2016). Uncertainty in flood damage estimates and its potential effect on investment decisions. *Natural Hazards and Earth System Sciences*, 16(1):1.
- [182] Wahl, T., Mudersbach, C., and Jensen, J. (2011). Assessing the hydrodynamic boundary conditions for risk analyses in coastal areas: a stochastic storm surge model. *Natural Hazards and Earth System Sciences*, 11(11):2925–2939.
- [183] Wahl, T., Mudersbach, C., and Jensen, J. (2012). Assessing the hydrodynamic boundary conditions for risk analyses in coastal areas: a multivariate statistical approach based on copula functions. *Natural Hazards and Earth System Science*, 12(2):495–510.
- [184] Wahl, T., Plant, N. G., and Long, J. W. (2016). Probabilistic assessment of erosion and flooding risk in the northern Gulf of Mexico. *Journal of Geophysical Research: Oceans*, 121(5):3029–3043.
- [185] Warren, I. and Bach, H. (1992). MIKE 21: a modelling system for estuaries, coastal waters and seas. *Environmental Software*, 7(4):229–240.
- [186] Werner, C., Bedford, T., Cooke, R. M., Hanea, A. M., and Morales-Nápoles, O. (2016). Expert judgement for dependence in probabilistic modelling: a systematic literature review and future research directions. *European Journal of Operational Research*, 258(3):801–819.
- [187] Whitehouse, R. J., Harris, J. M., Sutherland, J., and Rees, J. (2011). The nature of scour development and scour protection at offshore windfarm foundations. *Marine Pollution Bulletin*, 62(1):73–88.
- [188] Wirsching, P. H. and Chen, Y.-N. (1988). Considerations of probability-based fatigue design for marine structures. *Marine Structures*, 1(1):23–45.
- [189] Woodroffe, C. D. (2002). *Coasts: form, process and evolution*. Cambridge University Press.
- [190] Woodward, M., Kapelan, Z., and Gouldby, B. (2014). Adaptive flood risk management under climate change uncertainty using real options and optimization. *Risk Analysis*, 34(1):75–92.
- [191] Yoshida, T. (2015). Maximum likelihood estimation of skew t-copulas, with its applications to stock returns. *The Institute of Statistical Mathematics Research Memorandum*, 1195.
- [192] Zorrilla, P., Carmona García, G., Hera, A. d. l., Varela Ortega, C., Martínez Santos, P., Bromley, J., and Henriksen, H. J. (2010). Evaluation of Bayesian networks in participatory water resources management, Upper Guadiana Basin, Spain. *Ecology and Society*, 15(3):12.

- [193] Zuur, A. F., Ieno, E. N., Walker, N. J., Saveliev, A. A., and Smith, G. M. (2009). Zero-truncated and zero-inflated models for count data. In *Mixed effects models and extensions in ecology with R*, pages 261–293. Springer.

ACKNOWLEDGEMENTS

I would like to acknowledge and thank the following people for their contribution to this thesis and their support over the past four and a half years.

To my promoters and mentors: Marcel Stive, thank you for the great freedom you have given me for my research activities and your unconditional trust. Anca Hanea, as far as I can think back in Delft-time, you have been there for me, always, be it day or night, in person or on Skype. Your commitment and warmth never wavered. I am deeply grateful. Oswaldo Morales Nápoles, you suddenly joining the group in my second year came as a much needed surprise. Your enthusiasm is contagious and you have convinced me that all will be possible. Thank you.

To the independent members of my committee: Ap van Dongeren, Jose Jiménez, Bas Jonkman and Pieter van Gelder, thank you for being part of my doctoral committee, for your interest in my work and your time.

To my advisors - in the broader sense of the word: Ap van Dongeren, thank you for your guidance and support, you went well above and beyond your official role as RISC-KIT's project coordinator. Kees den Heijer, thank you for taking on a role as mentor; your perspective on the PhD process has been very valuable to me. Jaap van Thiel de Vries, thank you for introducing me to the world of coastal engineering. Our first 'informal meeting' has turned out to be a defining moment for my career. Robert McCall, Bas Jonkman, Pieter van Gelder, Harry Joe, Nathaniel Plant and Thomas Wahl, thank you for the thoughtful feedback that you have given at numerous occasions. Thanks also to the anonymous referees whose comments have been insightful and helpful.

To my colleagues and friends abroad: RISC-kids, I very much enjoyed the project meetings, workshops, the winter school, and of course collaborating with and learning from all of you. I treasure my RISC-KIT memories. The funding I received through this project also enabled me to travel across the world and build a unique network of colleagues and friends. I am very thankful for having had this opportunity. Nathaniel Plant, thank you for the warm welcome at the USGS in St. Petersburg, Florida, and the continuous support throughout my PhD. All of CEBRA, thank you for making me feel at home in Melbourne. It has been wonderful. Thanks also to Mark Burgman, who, as head of CEBRA, has supported my first visit financially. Claudia Czado, I am very grateful for the several opportunities I had to visit your group at the TUM. It has been a great pleasure to work with you and Thomas Nagler. Special thanks also to Dominik Müller and Nicole Barthel for taking care of me so well during my visits. I must also acknowledge the "Global Challenges for Women in Math Sciences" initiative for funding two of the research stays in Munich.

To my colleagues and friends at the TU Delft: You were the foundation of the pleasant working atmosphere. You have also contributed to this thesis by shaping my ideas, be it through formal colloquium talks or coffee machine chat. Huub van Verseveld and Lau-

rens Poelhekke, as Master students you have inspired, challenged and motivated me. Thank you. Tina Nane, thank you for your heart-felt encouragements. Alex Kosgodagan, I am glad we shared the first conference experiences. Special thanks also to Anais Couasnon, Kaveri Iychettira, and Antonia Sebastian for acting as sounding boards and providing incredible moral support. Marcel Stive, Myron van Damme and Maartje Godfroy, thank you for helping with the Dutch translation of my summary. Inge van Rooij, Judith Schooneveld-Oosterling, Agnes Groenestein and Diana Keijzer, thank you for your continuous assistance in all kinds of matters. Fedor Baart, thank you for several support sessions in front of the shell. Phaedra Oikonomopoulou, thank you for your awesome drawing and photographs. Telesilla Bristogianni and Tommaso Venturini, thank you for helping with perfecting the drawing's digital version.

Above all, I thank my friends and family, for everything.

CURRICULUM VITÆ

Wiebke Solvieg JÄGER

20-12-1989 Born in Hamburg, Germany.

EDUCATION

2003 - 2007 Abitur (Landesschule Pforta, Germany)

2007 - 2010 B.Sc. in Mechanical Engineering (ETH Zürich, Switzerland)

2010 - 2013 M.Sc. in Engineering and Policy Analysis (TU Delft, The Netherlands)

2011 - 2013 M.Sc. in Applied Mathematics (TU Delft, The Netherlands)

2013 - 2018 Ph.D. in Coastal Engineering (TU Delft, The Netherlands)

PUBLICATIONS

5. **W.S. Jäger**, T. Nagler, C. Czado, R.T. McCall *A statistical simulation model for joint time series of non-stationary hourly wave parameters*, submitted.
4. **W.S. Jäger**, E. Christie, A.M. Hanea, K. den Heijer, T. Spencer, *A Bayesian network approach for coastal risk analysis and decision making*, Coastal Engineering, 134, (2018).
3. **W.S. Jäger** and O. Morales-Nápoles, *A vine-copula model for time series of significant wave height and mean zero-crossing period in the North Sea*, ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering, 3(4), (2017).
2. L. Poelhekke, **W.S. Jäger**, A.R. van Dongeren, T.A. Plomaritis, R. McCall, Ó. Ferreira, *Predicting coastal hazards for sandy coasts with a Bayesian Network*, Coastal Engineering, 118, (2016).
1. H.C.W. van Verseveld, A.R. van Dongeren, N.G. Plant, **W.S. Jäger**, C. den Heijer, *Modelling multi-hazard hurricane damages on an urbanized coast with a Bayesian Network approach*, Coastal Engineering, 103, (2015).

CONFERENCES AND LECTURES

9. **W.S. Jäger**, *Modelling time series of wave parameters*, Oral Presentation, EVAN 2017, Third Conference on Advances in Extreme Value Analysis and Application to Natural Hazards, Southampton, United Kingdom, 5-7 September 2017.
8. **W.S. Jäger**, *Introduction to Bayesian networks*, Lecture, RISC-kit Summer School II, Faro, Portugal, 16 January 2017.
7. T. Bogaard, S. De Kleermaeker, **W.S. Jäger**, A.R. van Dongeren, *Development of generic tools for coastal early warning and decision support*, FLOODrisk 2016, 3rd European Conference on Flood Risk Management, Lyon, France, 17-21 October 2016.
6. **W.S. Jäger**, O.Morales-Nápoles, *Multivariate stochastic characterization of sea states*, Oral Presentation, Dependence Modeling in Finance, Insurance and Environmental Science, Munich, Germany, 17-19 May 2016.
5. **W.S. Jäger**, O.Morales-Nápoles, *How do waves and surges join forces in the Dutch North Sea? A multivariate analysis*, Oral Presentation, International Coastal Symposium 2016 , Sydney, Australia, 6-11 March 2016.
4. **W.S. Jäger**, O.Morales-Nápoles, *Sampling joint time series of significant wave heights and periods in the North Sea*, Oral Presentation, ESREL 2015: the 25th ESRA Conference , Zürich, Switzerland, 7-10 September 2015.
3. **W.S. Jäger**, C. den Heijer, A. Bolle, A.M. Hanea, *A Bayesian network approach to coastal storm impact modeling*, Oral presentation, ICASP12: 12th International Conference on Applications of Statistics and Probability in Civil Engineering, Vancouver, Canada, 12-15 July 2015.
2. S. De Kleermaeker, **W.S. Jäger**, A.R van Dongeren, *Development of Coastal-FEWS: Early warning system tool development*, IAHR 2015: 36th IAHR World Congress, Delft - The Hague, The Netherlands, 28 June - 3 July 2015.
1. **W.S. Jäger**, *The classical model of expert judgment. Opinion and subjective probability in science*, Lecture, Delft University of Technology, Delft, The Netherlands, 3 June 2014.