DTU Library

From soft sensing to anomaly detection in combined sewer systems

Palmitessa, Rocco

Publication date:
2021

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

DTU Environment
Nanyang Technological University

From soft sensing to anomaly detection in combined sewer systems

Rocco Palmitessa

PhD Thesis
May 2021

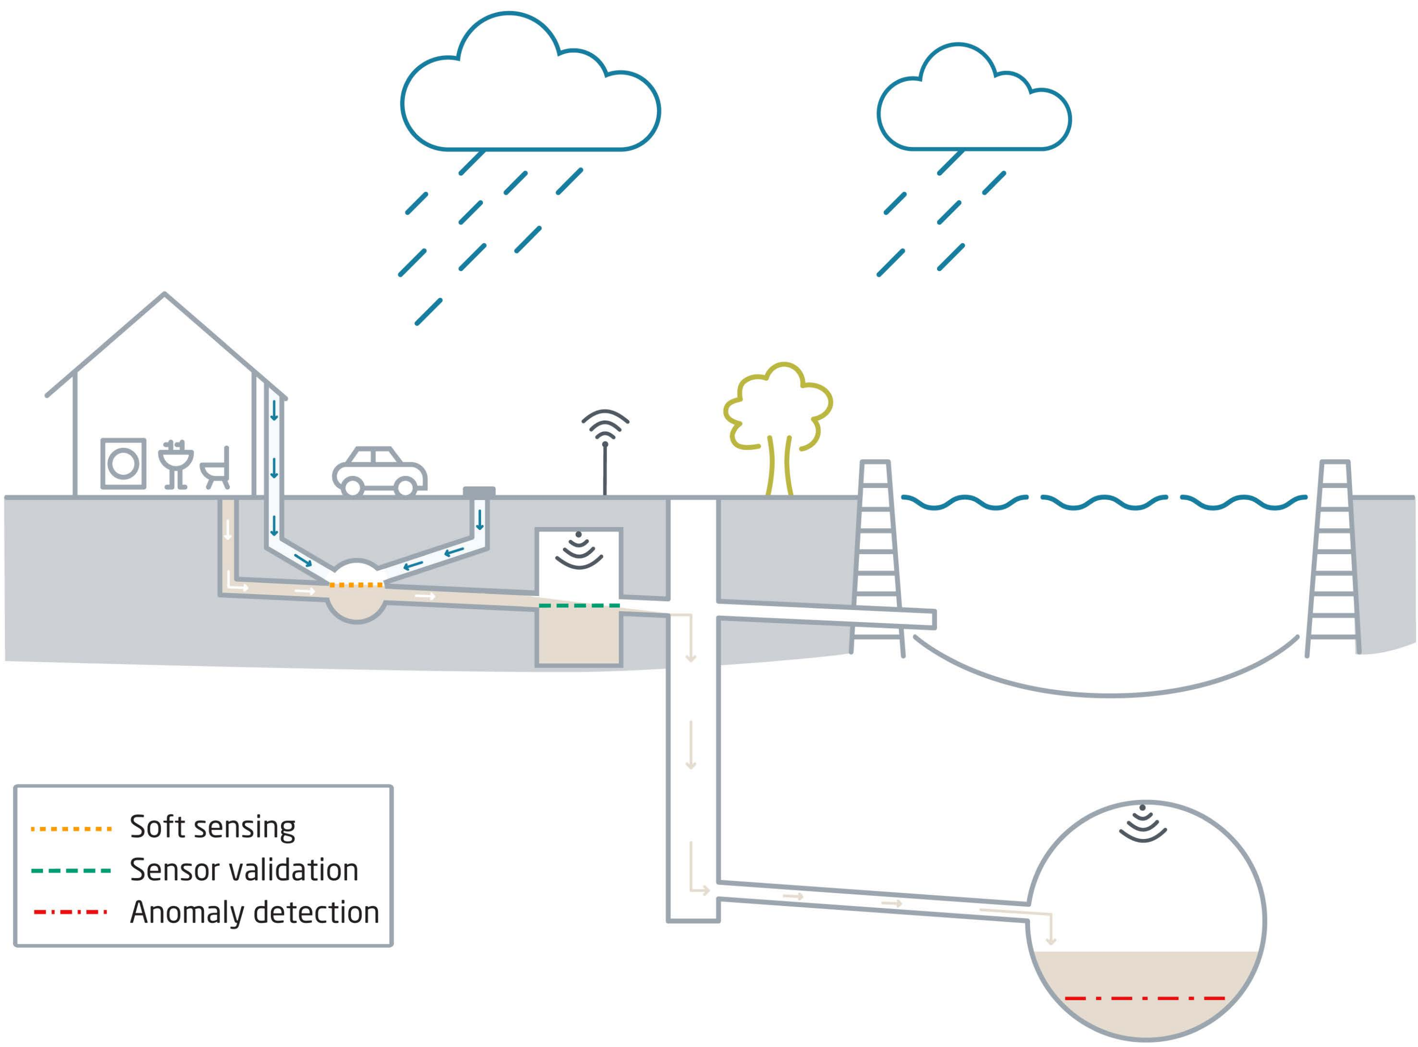

From soft sensing to anomaly detection in combined sewer systems

Rocco Palmitessa

PhD Thesis
May 2021

Department of Environmental Engineering
Technical University of Denmark

and

School of Civil and Environmental Engineering
Nanyang Technological University

From soft sensing to anomaly detection in combined sewer systems

Rocco Palmitessa

PhD Thesis, May 2021

The synopsis part of this thesis is available as a pdf-file for download from the DTU research database ORBIT: http://www.orbit.dtu.dk.

Address:	DTU Environment Department of Environmental Engineering Technical University of Denmark Bygningstorvet, Building 115 2800 Kgs. Lyngby Denmark
Phone reception:	+45 4525 1600
Fax:	+45 4593 2850
Homepage:	http://www.env.dtu.dk
E-mail:	reception@env.dtu.dk

Preface

This PhD thesis is submitted in partial fulfillment of the requirements for the framework agreement of Joint Doctorate Degree between the Technical University of Denmark (DTU) and the Nanyang Technological University (NTU), Singapore. The work was conducted in the period from December 2017 to March 2021 and guided by main supervisor Prof. Peter Steen Mikkelsen (DTU), co-supervisor Assoc. Prof. Morten Borup (DTU) and co-supervisor Prof. Adrian W.K. Law (NTU). The work was carried out mainly at DTU with two research stays at NTU in the periods February to July 2019 and January to July 2020.

The thesis is organized in two parts: the first part is a synopsis that puts into context the findings of the PhD in an introductory review; the second part consists of the papers listed below. These will be referred to in the text by their paper number written with the Roman numerals **I-III**.

I **Palmitessa, R.**, Mikkelsen, P.S., Law, A.W.K., Borup, M. (2020). Data assimilation in hydrodynamic models for system-wide soft sensing and sensor validation for urban drainage tunnels. *Journal of Hydroinformatics*. (In press) https://doi.org/10.2166/hydro.2020.074

II **Palmitessa, R.**, Mikkelsen, P.S., Borup, M., Law, A.W.K. (2021). Soft sensing of water depth in combined sewers using LSTM neural networks with missing observations. *Journal of Hydro-environment Research*. (In press) https://doi.org/10.1016/j.jher.2021.01.006

III **Palmitessa, R.**, Pedersen, A.N., Borup, M., Sørensen, L., Law, A.W.K., Clemmensen, L.K.H., Mikkelsen, P.S. (2021). Anomaly detection in water depth observations from combined sewers using LSTM neural networks. (Manuscript)

In this online version of the thesis, paper **I-III** are not included but can be obtained from electronic article databases e.g. via www.orbit.dtu.dk or on request from DTU Environment, Technical University of Denmark, Miljoevej, Building 113, 2800 Kgs. Lyngby, Denmark, info@env.dtu.dk.

In addition, the following publications, not included in this thesis, were also concluded during the PhD study:

Palmitessa, R., Borup, M., Mikkelsen, P.S. (2018). Urban tunnel systems for conveyance and storage of storm- and wastewater: features, classification, and modelling. In *11th International Conference on Urban Drainage Modelling*, 23-26 September, Palermo, Italy, pp. 251-254 (Extended abstract).

Palmitessa, R., Mikkelsen, P.S., Law, A.W.K., Borup, M. (2019). A data assimilation scheme tailor-made for real-time modelling of urban drainage tunnels. In *17th International Computing & Control for the Water Industry Conference*, 1-4 September, Exeter, United Kingdom, 2 pp (Extended abstract).

Palmitessa, R., Brink-Kjær, A., Clemmensen, L.K.H., Borup, M., Mikkelsen, P.S. (2021). Machine learning for anomaly detection in combined sewer systems. In *14th annual Water Research Conference: Danish Water Forum*, 2 March, Online, p. 40 (Abstract).

Palmitessa, R., Borup, M., Mikkelsen, P.S., Law, A.W.K. (2021). Soft sensing of water depth in combined sewers using LSTM neural networks. In *Singapore International Water Week 2021*, 22 June - 2 July, Online, 3 pp (Extended abstract).

Palmitessa, R., Chew, A.W.Z., Rieger, L., Borup, M., Mikkelsen, P.S., Law, A.W.K. (2021). Accuracy of water depth predictions in combined sewers using a recurrent neural network with gaps in the observed data. In *15th International Conference on Urban Drainage*, September, Melbourne, Australia, 3 pp (Extended abstract, postponed due to the Covi-19 pandemic).

Acknowledgements

My deepest gratitude goes to my team of supervisors at the Technical University of Denmark and the Nanyang Technological University for conceiving and supervising the PhD project that led to this thesis. I thank Professor Peter Steen Mikkelsen for his kind leadership and endless patience, Associate Professor Morten Borup for his joyful curiosity and reassuring honesty, and Professor Adrian W.K. Law for his warm hospitality and helpful pragmatism. They all contributed to this project both as researchers, proposing ideas and correcting mistakes, and mentors, providing guidance and support in the good and bad times. Like all endeavors, this PhD has been a bumpy ride, but my team of supervisors helped me make it to the end safe and sound.

All the results in this thesis are built on data and models generously provided by external collaborators. Mark Ridler and Henrik Madsen from DHI shared access to a beta version of MIKE URBAN and helped me setting up the data assimilation framework; Anders Breinholt and Lone Bo Jørgensen from the Greater Copenhagen Utility (HOFOR) shared models and data for the Damhus tunnel system, which were used for the first two publications; and Annette Brink-Kjær, Agnethe N. Pedersen and Lasse Sørensen from VandCenter Syd contributed to the third manuscript with their data, experience and opinions from an end-user perspective. I also need to thank Associate Professor Line Clemmensen and Laura Rieger from DTU Compute, as well as Alvin W.Z. Chew from NTU, for assisting me in the exploration of machine learning.

This research project has developed in parallel to a journey of personal growth and cultural exploration. I thank all those who have crossed my way and enriched my journey: the fellow PhD students at DTU Environment welcomed me as part of the gang and shared the joys and frustrations of the PhD life; the permanent staff at DTU taught me the benefits of a flat work hierarchy; Anne, the PhD secretary at DTU Environment, made my life a bit easier when needed; the students I assisted in the courses at DTU Environment inspired me with their curiosity and motivation; the friends and colleagues in Singapore introduced me to a whole new world and never let me feel lost.

Also, this final goal could not be met if it wasn't for all those who stood by my side: my family coped with the distance that separated us, wishing me a prosperous future; my long-time friends patiently waited through my long silences and made their best to make me feel their presence; my talented friend Karin kept up with tradition and supported my vanity for an elegant and communicative thesis cover; my beloved Marc made this whole journey more enjoyable by simply being himself.

Summary

Urban drainage systems are a strategic component of cities' infrastructure, as they safeguard citizens' health and properties. Most cities in Europe and North America are served by combined sewer systems built to meet the service levels of decades ago. Climate change, intensified urbanization and stricter environmental regulations are straining the existing infrastructure, and large investments are needed for futureproofing. The digitalization of urban drainage offers cost-effective solutions to upgrade the existing system without building new structures or adding new equipment. With a combination of models and data, utility companies can maximize the use of the available capacity and reduce the risk of flooding and pollution.

The hydraulic state of combined sewer systems is typically monitored with a network of sensors. These are expensive to install and maintain, therefore they are placed only in strategic locations, leaving large parts of the system unobserved. Moreover, sensor observations are vulnerable to instrument faults, communication errors and cyberattacks. Models of the system can supplement the information from the sensors with a system-wide picture of the hydraulic state. Both physics-based and data-driven models exist capable of reproducing the behavior of sewer systems. If the prediction accuracy is sufficiently high, models can also act as soft(ware) sensors, working alongside hardware sensors or replacing them for periods of time. Model predictions can also be used to validate available observations and detect anomalous behavior.

The potential of using models for real-time soft sensing and anomaly detection in combined sewer systems has not been fully exploited yet. Hydrodynamic models describe in detail the spatial and temporal distribution of the hydraulic variables but are computationally expensive and tend to drift off reality. Assimilating observations in real-time allows to update hydrodynamic models to the current state of the system, thus increasing their prediction accuracy. This is usually done by running an ensemble of model instances in parallel, which further increases the computational cost. A data assimilation scheme was developed and tested which limits the assimilation to a sub-system of the larger sewer system. This approach yielded accurate system-wide predictions of water depth, thus effectively enabling the soft sensing capabilities of the hydrodynamic model. The prediction was also used to validate an independent

sensor located 3.5 km upstream, revealing an issue of false echo in one of the analysed events.

Purely data-driven models offer an alternative to physics-based models. They learn the behavior of the system from a series of historical observations and predict accordingly its response to given inputs. Artificial neural networks have long been used to predict water depths in combined sewer systems, and recent advancements in terms of algorithms and hardware have unlocked new possibilities. Long Short-Term Memory (LSTM) neural networks have gained popularity in natural language processing but are also particularly suited for time series forecasting. However, little is known about their efficacy with water depth observations from combined sewers.

An LSTM neural network was trained with several months of water depth observations at 1-min resolution from different locations. The key network settings were calibrated to find an optimal setup that could work across different locations. The prediction accuracy was compared between scenarios with different gaps in the antecedent observations to simulate missing data. It was proven that the model could compensate the missing information on the antecedent state of the system with the other sources of information, namely the observed rainfall and the time of the day. This demonstrated the robustness of LSTM networks and their potential as soft sensing tools. In a separate test, the LSTM prediction was used as basis for anomaly detection. The observations were flagged as anomalous when they deviated significantly from the expected behavior of the system. However, the detection efficacy was dependent on the quality and quantity of the training data and changed across locations.

Updated hydrodynamic models can generate accurate system-wide predictions but require detailed physical information. For these reasons, they are suited for soft sensing and anomaly detection in strategic structures within combined sewer systems. On the other hand, LSTM neural networks are trained on observations alone but have point-wise validity. Therefore, they can easily be deployed as a screening tool for large networks of sensors. By investigating and testing updated hydrodynamic models and LSTM neural networks, this thesis demonstrates their concrete potential for soft sensing and anomaly detection applications and promotes their integration in the monitoring and control workflows of modern utility companies.

Dansk sammenfatning

Afløbssystemer er en vigtig komponent i byers infrastruktur, da de beskytter borgernes sundhed og ejendom. De fleste byer i Europa og Nordamerika betjenes af fælleskloakerede systemer, der er bygget til at opfylde de serviceniveauer, som blev opsat for årtier siden. Klimaændringer, intensiveret urbanisering og strengere miljølovgivning belaster den eksisterende infrastruktur, og der er behov for store investeringer til fremtidssikring. Digitaliseringen af afløbssystemer tilbyder omkostningseffektive løsninger til at opgradere det eksisterende system uden behov for at opføre nye bygværker eller tilføje nyt udstyr. Med en kombination af modeller og data kan forsyningsselskaber maksimere brugen af den tilgængelige systemkapacitet og dermed reducere risikoen for oversvømmelse og forurening.

Den hydrauliske tilstand af fælleskloakerede systemer overvåges typisk med et netværk af sensorer. Disse er dyre at installere og vedligeholde, og derfor placeres de kun strategiske steder, hvilket betyder at store dele af systemet ikke observeres direkte. Desuden er sensorobservationer sårbare over for instrumentfejl, kommunikationsfejl og cyberangreb. Modeller af systemet kan supplere informationen fra en sensor med et billede af den hydrauliske tilstand i hele systemet. Der findes både fysisk baserede og datadrevne modeller, der er i stand til at reproducere kloaksystemers opførsel. Hvis forudsigelsesnøjagtigheden er tilstrækkelig høj, kan modeller også fungere som soft(ware) sensorer, der arbejder sammen med hardwaresensorer eller erstatter dem i perioder. Modelforudsigelser kan også bruges til at validere tilgængelige observationer og opdage uregelmæssig adfærd.

Potentialet ved at bruge modeller til soft-observering i realtid samt anomali-detektion i fælleskloakerede systemer er endnu ikke fuldt udnyttet. Hydrodynamiske modeller giver en detaljeret beskrivelse af den rumlige og tidsmæssige fordeling af de hydrauliske variable, men er beregningsmæssigt dyre og har tendens til at drive væk fra virkeligheden. Assimilering af observationer i realtid gør det muligt at opdatere hydrodynamiske modeller til systemets aktuelle tilstand og dermed øge deres nøjagtighed. Dette gøres normalt ved at køre et ensemble af modeller parallelt, hvilket yderligere øger beregningsomkostningerne. En dataassimileringsmetode, som begrænser assimileringen til et undersystem af det større kloaksystem, er her blevet udviklet og testet. Denne fremgangsmåde gav nøjagtige forudsigelser af vanddybde i hele systemet, hvilket muliggør den hydrodynamiske models brug til soft-observering. Forudsigelsen blev også brugt til at validere data fra en uafhængig sensor placeret 3,5 km opstrøms, hvilket afslørede et problem med falske ekkoer i en af de analyserede begivenheder.

Rent datadrevne modeller tilbyder et alternativ til fysisk baserede modeller. De lærer systemets opførsel ud fra historiske observationer og forudsiger derefter dets reaktion på kendte input. Kunstige neurale netværk har længe været brugt til at forudsige vanddybder i fælleskloakerede systemer, men nylige fremskridt indenfor algoritmer og hardware har givet nye muligheder. Long Short-Term Memory (LSTM) neurale netværk har vundet popularitet inden for behandling af naturlige sprog, men er også særligt velegnede til tidsrækkeforudsigelser. Der findes dog ikke meget viden om deres brugbarhed til observationer af vanddybde fra fælleskloakker.

Et LSTM neuralt netværk blev trænet med flere måneders observationer af vanddybde ved 1-minuts opløsning fra forskellige lokationer. De vigtigste netværksindstillinger blev kalibreret til at finde en optimal opsætning, der kunne fungere på tværs af forskellige lokationer. Forudsigelsesnøjagtigheden blev sammenlignet mellem scenarier med forskellige huller i de foregående observationer for at simulere manglende data. Det blev bevist, at modellen kunne kompensere for den manglende information om systemets forudgående tilstand med andre informationskilder i form af observeret nedbør og tidspunktet på dagen. Dette demonstrerede LSTM-netværks robusthed og deres potentiale som software observeringsværktøjer. I en separat test blev LSTM-forudsigelsen anvendt som basis for opdagelse af anomalier. Observationerne blev markeret som anomalier, hvis de afveg markant fra systemets forventede opførsel. Opdagelseseffektiviteten var dog afhængig af kvaliteten og kvantiteten af træningsdataene og var forskellig fra sted til sted.

Opdaterede hydrodynamiske modeller kan generere nøjagtige systemforudsigelser, men kræver detaljerede fysiske oplysninger. Af disse grunde er de velegnede til soft-observering og anomali-detektion i udvalgte, vigtige bygværker i fælleskloakerede systemer. På den anden side trænes neurale LSTM-netværk kun på observationer og er gyldige for specifikke punkter. Derfor kan de let implementeres som et screeningsværktøj til store sensornetværk. Ved at undersøge og teste opdaterede hydrodynamiske modeller og LSTM neurale netværk demonstrerer denne afhandling deres potentiale for soft-observering og anomali-detektion samt fremmer deres integration i overvågnings- og kontrolarbejdsprocesserne i moderne forsyningsselskaber.

Table of Contents

Abbreviations

ANN	Artificial Neural Network
CFD	Computational Fluid Dynamics
CSO	Combined Sewer Overflow
DA	Data Assimilation
DUDM	Distributed Urban Drainage Model
EnKF	Ensemble Kalman Filter
LSTM	Long Short-Term Memory
RMSE	Root Mean Square Error
UDS	Urban Drainage System

1 Introduction

1.1 Background

Urban drainage systems (UDSs) are critical for the livability of urban areas, as they divert used and excess water to treatment and disposal. Most cities in Europe and North America are served by combined sewer systems built over decades. These receive used water from both households and non-residential facilities, mixed with precipitation runoff. Under normal conditions, the combined sewage is conveyed to a treatment facility. However, during intense precipitation, the excess water needs to be stored or released into the environment. The occurrence of intense events is expected to increase as a consequence of climate change (Arnbjerg-Nielsen et al., 2013). At the same time, environmentally conscious administrations are imposing stricter regulations on the release of untreated water from sewers. Thus, large investments are needed to ensure that these aging infrastructures can meet the ever-increasing service demand.

Responses to this problem typically fall within grey, green or digital solutions. Grey solutions aim at expanding the capacity of combined sewer systems by enlarging the existing structures or building auxiliary storage structures. As an alternative to volume expansion, a number of solutions aim at reducing or delaying runoff at source, often integrating vegetation and artificial structures. Thus, these solutions are termed green infrastructure, although a plethora of terms are in use (Fletcher et al., 2015). Grey and green strategies can be seen as opposed, but often the optimal solution relies on a combination of the two according to the local political, environmental, technical, economic and social factors (Dolowitz et al., 2018).

A third option is offered by digital solutions. The advancements in information and communication technologies have enabled new approaches for optimizing the management of urban drainage systems (Eggimann et al., 2017). Data and models are becoming essential for both monitoring and control purposes. Digitalization and automation allow information to be transmitted and processed in real time to inform prompt control decisions (García et al., 2015). This can help maximizing the use of the available storage capacity and thus reduce the need for additional physical structures (Kerkez et al., 2016). However, as data is poised to influence all stages of UDSs, from planning to monitoring, new ethical and technical challenges arise (Makropoulos & Savíc, 2019).

1.2 Data and models

Data is used to represent all aspects of combined sewer systems. Information about assets, including geometry and materials, are typically digitized and updated as the network evolves. Inputs to the system are also quantifiable, although with a margin of uncertainty. Wastewater inflows to the system are derived from the distribution of users and their water consumption, and generally follow repetitive patterns. Precipitation runoff, instead, is of aleatory nature and needs to be observed. Despite the emerging use of radar sensors, precipitation observations remain uncertain, mostly due to their heterogenous spatial distribution (Campisano et al., 2013). Infiltration from the surface and the groundwater can contribute significantly to the total inflow in dry weather, and generally requires long series of data to be quantified (de Ville et al., 2017).

Most importantly, the state of the sewer system can be characterized numerically to track the quantity and quality of water in the different components of the system. This data is needed to inform control decisions and ensure that the service demand is met. Water level measurements are the most common type of observations from UDSs. Water level sensors are generally very accurate but need to operate in harsh environments with limited access. Low-power wireless communication techniques allow cost-effective collection of spatially distributed, real-time observations with battery-powered sensors (Ebi et al., 2019). However, as the number of sensors deployed increases, so do maintenance costs. Therefore, water level sensors are typically placed in strategic structures, e.g storage basins and overflow chambers, leaving large parts of the system without direct observations.

Models, on the other hand, are capable of generating a system-wide estimate of the state of the system. To replicate the behavior of the sewer system, models need to incorporate our knowledge of the driving physical processes or to be exposed to long series of observations. Several types of models are in use with different levels of complexity. The choice between them depends on the purpose at hand, e.g planning, design, monitoring, forecasting or control.

Physics-based or white-box models carry a mathematical formulation of the hydrological and hydraulic processes dominating the behavior of urban drainage systems. Their level of detail generally relates to the required computational effort. Computational Fluid Dynamics (CFD) models have the highest level of detail and are mostly reserved for individual structures, such as overflow chambers (Isel et al., 2014). 1D hydrodynamic models are best suited for representing the entire sewer system, as they reproduce with “high fidelity" the temporal and spatial distribution of water across the system with reasonable computational resources. On the other hand, simplified and spatially lumped models are faster to run and, therefore, more suitable for control purposes (Vanrolleghem et al., 2005). At all levels of detail, models generate uncertain predictions that reflect the imperfect knowledge of inputs and the inherent simplifications of reality (Deletic et al., 2012). Observations can improve the prediction accuracy if used to calibrate the model parameters or update the model states dynamically (Borup et al., 2011), a process known as data assimilation (DA).

In contrast, data-driven or black-box models derive the behavior of the system from the available data. Grey-box models combine deterministic and data-driven elements and preserve some knowledge of the physical processes (Breinholt et al., 2011). A purely data-driven model is built to replicate the relationship between input and output without any domain knowledge. Generally, this type of models requires much less computational resources than a physics-based model and adapts easily to new information. Among black-box models, Artificial Neural Networks (ANN) have since long been used to model the complex, non-linear behavior of urban drainage systems (Loke et al., 1997). ANNs learn recurrent patterns in data in resemblance to biological neural systems (Schmidhuber, 2015).

Both physics-based and data driven models can act as soft (software) sensors to supplement or replace the hard (hardware) sensor, provided they are sufficiently accurate. This requires that the model is calibrated and validated against a set of observations. A soft sensor can provide an estimate of states of the system that are not directly observed or fill gaps of information on directly observed states. This could be the case if the hard sensor is faulty or under maintenance. Depending on the type of model, the soft sensor can be localized in space or cover large portions of the UDS. Examples of both types are reported in the literature, including a CFD-derived model for the monitoring of a combined sewer overflow (Ahm et al., 2016) and a system-wide digital twin of the urban drainage system for model-based real-time monitoring of the whole network (Pedersen et al., 2021).

In terms of water quantity in combined sewer systems, the major challenge for soft sensors is to replicate the complex, non-linear relationship between the rainfall input and the hydraulic state (water level/depth or discharge). An example of this relationship is shown in Figure 1. In dry weather, the water depth follows a repetitive diurnal pattern in addition to the infiltration contribution. In wet weather, instead, the response to the rainfall intensity is nonlinear and location-specific, as it is affected by the geometry of the system and the presence of control devices, such as gates, weirs, pumps, etc.

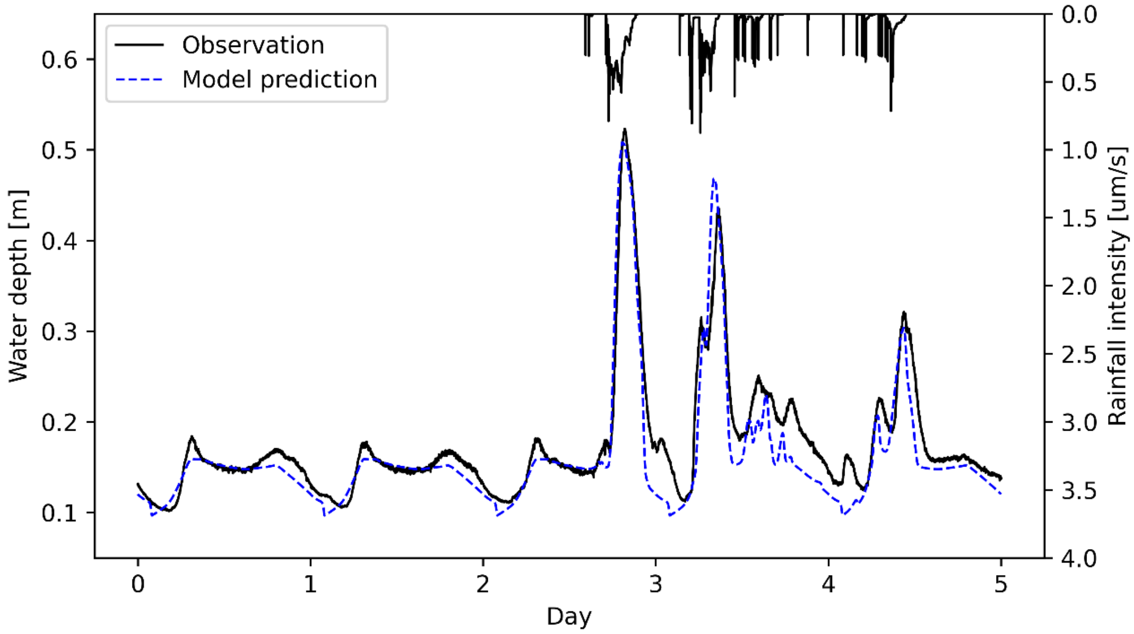

Figure 1. Example of water depth observations and model predictions in dry and wet weather. Adapted from Paper II.

1.3 Cyber-physical security

When the digital sensing and control become tightly intertwined with the physical assets, urban drainage systems can be regarded as cyber-physical systems (Sun et al., 2020). The operation of cyber-physical systems relies heavily on the quantity and quality of data available (Rajkumar et al., 2010). Thus, cyber-physical systems are more vulnerable than conventional physical systems to failure and attacks (Ding et al., 2018), and it is paramount that the data streams can be trusted with very high reliability (Blumensaat et al., 2019). Given the scale and critical function of UDSs, this increased vulnerability can potentially lead to serious consequences. One notable example is the breach that occurred in 2000 at the Maroochy Water Services in Queensland, Australia, in which a former contractor maliciously gained control of 150 pumping stations (Slay & Miller, 2007). By the time the attack was detected three months later, one million liters of untreated sewage had been discharged to the local waterways. While the digital transformation exposes new vulnerabilities of urban drainage systems, technological advancements at the same time offer solutions to help reducing risks (Moy De Vitry et al., 2019).

To ensure the security of cyber-physical systems, conventional network security should be backed up by an intrusion detection system, capable of identifying a breach that has already occurred (Ramotsoela et al., 2018). Anomaly detection is a subset of intrusion detection which aims at identifying deviations from the normal behavior of the system (Kwon et al., 2019). As such, it is also effective in detecting non-malicious malfunctions of the monitoring and control equipment. For highly dynamic systems like urban drainage, the normal behavior cannot be estimated solely based on the statistical distribution of the historical data, but also requires a model capable of understanding the driving processes and estimating the present state of the system. A model with a similar capability also fits the requirements of a soft sensor.

1.4 Objectives

The overall aim of this thesis is to bridge the gap towards the field implementation of soft sensing and anomaly detection in combined sewer systems, with an exclusive focus on water depth observations. The aim is pursued by investigating the following research questions:

i) How is the terminology related to soft sensing and anomaly detection used in the context of urban drainage systems?

ii) Can 1D hydrodynamic models be optimized for soft sensing and anomaly detection in combined sewer systems by means of data assimilation?

iii) Can Long Short-Term Memory neural networks be trained to replicate the behavior of combined sewer systems with sufficient accuracy for soft sensing and anomaly detection purposes?

iv) For which type of applications are the two investigated methods suitable?

These research questions reflect the objectives of the three attached papers (Table 1). Paper I presents a data assimilation scheme tailor-made for urban drainage tunnels that aims at enabling 1D hydrodynamic models with soft sensing and sensor validation capabilities. Paper II and III investigate the use of Long Short-Term Memory neural networks in predicting the behavior of combined sewer systems. Their focus is soft sensing and anomaly detection, respectively. Both Paper I and II use data from the Damhus overflow tunnel system in Copenhagen, Denmark, while Paper III uses data from combined sewer overflow chambers in Odense, Denmark.

Table 1. Overview of scientific papers included in this thesis: objectives, case study and model used.

Paper	Objectives	Case study	Model
I	• System-wide soft sensing • Cross-validation of sensors (anomaly detection) • Computational cost minimization	Damhus overflow tunnel system, Copenhagen, DK	Updated 1D hydrodynamic
II	• Point-wise soft sensing • High prediction accuracy • Reproducibility		Long Short-Term Memory neural network
III	• Anomaly detection (sensor validation) • Robustness • Minimal human intervention	Combined sewer overflow chambers, Odense, DK	

1.5 Thesis structure

This thesis is structured as follows. Chapter 2 discusses the main terminology and concepts related to soft sensing and anomaly detection in combined sewer systems. This first of the two investigated approaches, data assimilation in 1D hydrodynamic models, is presented in Chapter 3, together with the key findings from Paper I. Chapter 4, instead, focuses on the second approach, Long Short-Term Neural networks and summarizes the results from Paper II and III. Chapter 5 draws a comparison between the two investigated approaches and discusses advantages and disadvantages of both. The last two chapters present the conclusions and the future perspectives, respectively.

2 Soft sensing and anomaly detection

2.1 Terminology

The term “soft(ware) sensor” is commonly used in industrial environments to refer to mathematical models of a system designed to estimate relevant system variables (Fortuna et al., 2007). A soft sensing model can be used to supplement or replace hard(ware) sensors, which often have high installation and maintenance costs. This is especially the case in urban drainage systems, which are difficult to access and require specialized personnel. Generally speaking, all models of combined sewer systems return estimates of relevant variables at different forecast horizons. However, to qualify as soft sensor, a model needs to be able to run in real-time and predict the current value of the target variable at multiple relevant locations.

The terms “soft(ware) sensing” and “soft(ware) sensor” are not widely used in the context of urban drainage, compared to other research areas (Table 2). One recurrent use of the term refers to a model that predicts an unobserved quantity starting from an observed one. This is the case of flows derived from water levels (Carstensen et al., 1996; Leonhardt et al., 2012; Ahm et al., 2016). Alternative terms to soft sensing used with a similar meaning are “data reconstruction” (Benedetti et al., 2008) and “gap filling” (Langeveld et al., 2017). However, the majority of studies on soft sensing in combined sewer systems are focused on water quality, e.g. (Pedersen et al., 2020).

Chandola et al. (2009) define “anomaly detection” as *the problem of finding patterns in data that do not conform to expected behavior.* In this thesis, the analysed data are water depth observations from combined sewer systems and the expected behavior is computed with a model. The anomalies could have different origins. A system anomaly occurs when the water depth is correctly observed by the sensor, but it behaves differently from the norm, e.g. in case of a blockage or failure of the control equipment. A sensor anomaly, instead, refers to the scenario in which the system behaves as usual, but the sensor records a wrong observation, which could be caused, for example, by miscalibration or obstruction of the sensor. Finally, a transmission anomaly happens when the observation is altered maliciously or unintentionally while being sent from the data source to the data repository. Classifying anomalies according to their type would certainly add value to the detection mechanism but is not the purpose of this thesis. Therefore, in the following “anomaly” refers to any of the types described above.

Table 2. Number of hits on scopus.com for terms related to soft sensing and anomaly detection in a generic context or in associations with urban drainage ("urban drainage" OR "drainage network" OR "sewer system" OR "sewer network" OR "combined sewer" OR "separate sewer"). Accessed 24th February 2021.

Term	Generic context	Urban drainage context
Soft(ware) sensing/sensor	13880	94
Data reconstruction	7088	18
Gap filling	9470	34
Anomaly detection	65048	73
Data validation	10385	111

The use of “anomaly detection” in the context of combined sewer systems is rare (Table 2). Branisavljević et al. (2011) tested various anomaly detection techniques on water depth, velocity and conductivity observations from combined sewers and discussed the benefit of data pre-classification. Frequently, the term is used with the meaning of fault or defect detection in the inspection of drainage pipes (Hassan et al., 2019; Moradi et al., 2020). Techniques for the detection of anomalous observations from sewer systems are more commonly termed “data validation” (Branisavljevic et al., 2010; Mourad & Bertrand-Krajewski, 2002). The term “sensor validation” as it is used in Paper I, can be considered a synonym of data validation, or rather an abbreviation of “sensor data validation”.

2.2 Conceptual framework

Anomaly detection can be dealt with as a classification or regression problem. In the first case the observations are classified as normal or anomalous based on a set or criteria or using an intelligent algorithm. For example, several machine learning algorithms exist that are capable of automatically classifying data after being trained on a labelled dataset (Russo et al., 2020) With the regression approach, a model carrying information on the relationship among system variables is used to predict the behavior of the system in response to a given input. Then, the deviations of the observed state from the expected one are flagged as anomalies, according to some predefined criteria. The hydrological and hydraulic processes governing combined sewer systems are well understood but lead to highly dynamic and non-linear behavior. Also, water depth observations are arranged in sequences with temporal correlation. For these reasons, the regression approach was considered the most promising and thus received focus of this thesis.

A regression model for anomaly detection also fulfils the requirements of a soft sensing model, which allows to couple the two processes in the same workflow. This is conceptualized in Figure 2. The initial step is the selection and training/calibration of the regression model. Training and calibration are terms commonly used with data-driven and physics-based models, respectively. But, in principle, they both consist in optimizing a set of parameters to ensure a better fit between the model output and a series of historical observations. At each step in time, the trained model is used to predict the expected state or behavior of the system, given a set of known inputs. In this thesis the predicted state is limited to the water depth at the analysed locations within the combined sewer system.

If an observation from a hardware sensor is available at the current timestep, it is compared with the model prediction. A significant deviation between the two signals a potential anomaly. Instead, if no observation is available for the specific point in time and space, the prediction itself is used in replacement. Thus, the model acts as soft sensor. The algorithm then shifts forward in time in preparation for a new iteration as time advances. After a predefined period, or when a sufficient amount of new data has been collected, the model is updated with the new information to keep track of changes in the system behavior. The update can affect the model parameters or only the initial and boundary conditions. This framework is also applied when soft sensing is performed without detecting anomalies, as in Paper II. In that case the anomaly detection step in bypassed and all the observations are implicitly assumed to be valid.

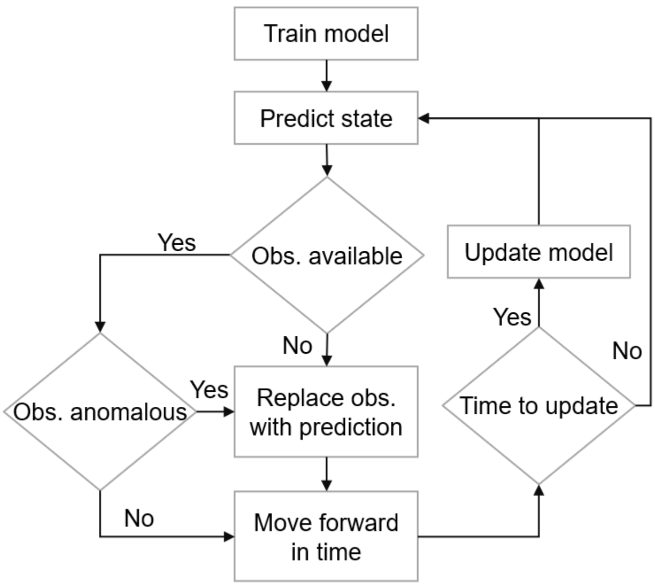

Figure 2. Conceptual workflow of soft sensing coupled with anomaly detection. From Paper II.

2.3 Urban drainage tunnels

Urban drainage tunnels are being widely adopted worldwide in response to the increased pressures on the sewer systems. These strategic structures often have massive scales and require considerable investments of taxpayers money. At the same time, urban drainage tunnels are usually deeper than the rest of the network to operate by gravity, thus limiting their accessibility for maintenance purposes. Monitoring and controlling theses structures with sensors alone can expose the operator to serious risks, including malicious attacks. This is why urban drainage tunnels make for an interesting case study for both soft sensing and anomaly detection.

A few large tunnels have been integrated in the combined sewer system of the city of Copenhagen, Denmark, and others are planned or under construction. They are a key component of the city's climate adaptation plan, an ambitious project aimed at securing the urban area from the potential damages of pluvial flooding. Some tunnels are designed to rapidly discharge excess stormwater directly at sea. Others intercept the existing overflow structures and withhold the combined sewage that would have otherwise reached the natural recipient untreated. In combined overflow tunnels the sewage is stored until the end of the precipitation event and then pumped back into the network for treatment. A different typology of drainage tunnels has been adopted by the city-state of Singapore. Located close to the equator, Singapore experiences a tropical climate with abundant precipitation throughout the year. Stormwater is therefore drained separately from wastewater and partly directed at large reservoirs for treatment and reuse. Used water from the nearly 6 million inhabitants is instead collected in a massive tunnel system crossing the island for a total length of about 80 km. This centralized structure allows to optimize the management of the sewer network and reduce the need for pumping stations at the surface level, but the operation of the entire city depends on its correct functioning.

Regardless of the type of water they are designed for (stormwater, wastewater, or combined sewage) urban drainage tunnels share a common defining feature: store and convey large quantities of water (Palmitessa et al., 2018). Despite the large dimensions, they are enclosed structures with a limited volume. Advanced monitoring techniques like soft sensing, can help optimize the available volume and inform smart control decisions. At the same time, it is greatly beneficial to be able to validate the incoming observations and detect anomalous data before they are given as input to the control algorithms.

2.4 Combined sewer overflows

Combined sewer systems are usually equipped with overflow structures to release excess water during wet weather and prevent the combined sewage from flooding streets and buildings. Originally, these structures discharged the overflow and its pollutants load directly into the natural recipient. In recent times, stricter environmental regulations mandate the reduction or elimination of overflow events. Nonetheless, many overflow chambers still exist in the combined sewers of modern cities. These chambers fundamentally consist of two components: a storage volume, determined by its geometry, and a regulation mechanism for the diversion of the overflow. The regulation can be either active, if controllable, or passive, if fixed. A passive overflow chamber fills up during a rain event until the water has reached a crest level. If the water depth raises further, the excess combined sewage overflows outside the storage chamber and is conveyed to the recipient (Figure 3).

To monitor the occurrence and magnitude of the overflow events, combined sewer overflows are often equipped with water level sensors. These are typically ultrasonic or pressure sensors and provide essential information for the monitoring and control of combined sewers. In dry weather, the water depth in the chamber follows the same repetitive pattern as the sewer pipes. During precipitation events, the water level increases in relation to the stored volume, which depends on the difference between inflow and outflow. When the storage capacity is met, the water starts overflowing as a function of the crest design. Across the three stages, the relationship between the precipitation and observed water depth changes.

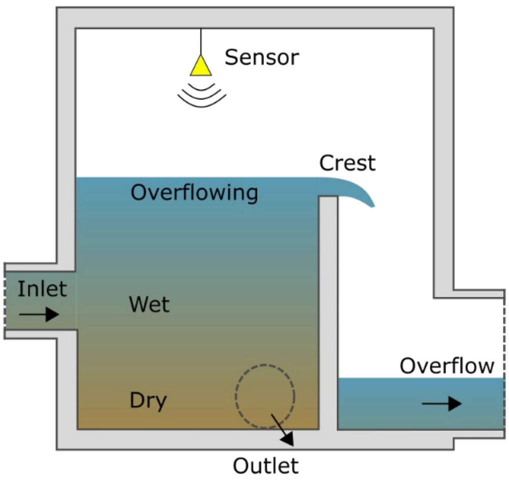

Figure 3. Schematics of a typical passive overflow chamber with water depth stages (dry, wet, and overflowing). Adapted from Paper II.

2.5 Inner and outer model

Combined sewer systems include in their most simple form two main physical components: a catchment, which collects wastewater and precipitation runoff from their sources, and a drainage network that conveys the combined sewage to treatment and disposal. A model capable of replicating the behavior of the system needs therefore to capture the basic mechanism of the various water sources being combined and conveyed across the system over time. As the water travels downstream, it mixes with the runoff and the sewage from larger areas of the system, leading to increasingly more complex behavior. Therefore, it may be beneficial to used models with a higher level of detail and complexity in the most downstream areas of the network.

Also, some structures have strategic roles in the operation of combined sewer systems, as is the case with drainage tunnels. Isolating a model of the selected structure can help optimizing the use of computational resources or increase the amount of information contained in the model. In this thesis, an "inner model" is focused on a downstream or strategic subsystem of the combined drainage system, whereas the "outer model" encompasses the upstream or peripheral parts of the system, including the catchment. The boundary between the outer and inner model is selected in a way that ensures the correct transfer of information from one to the other.

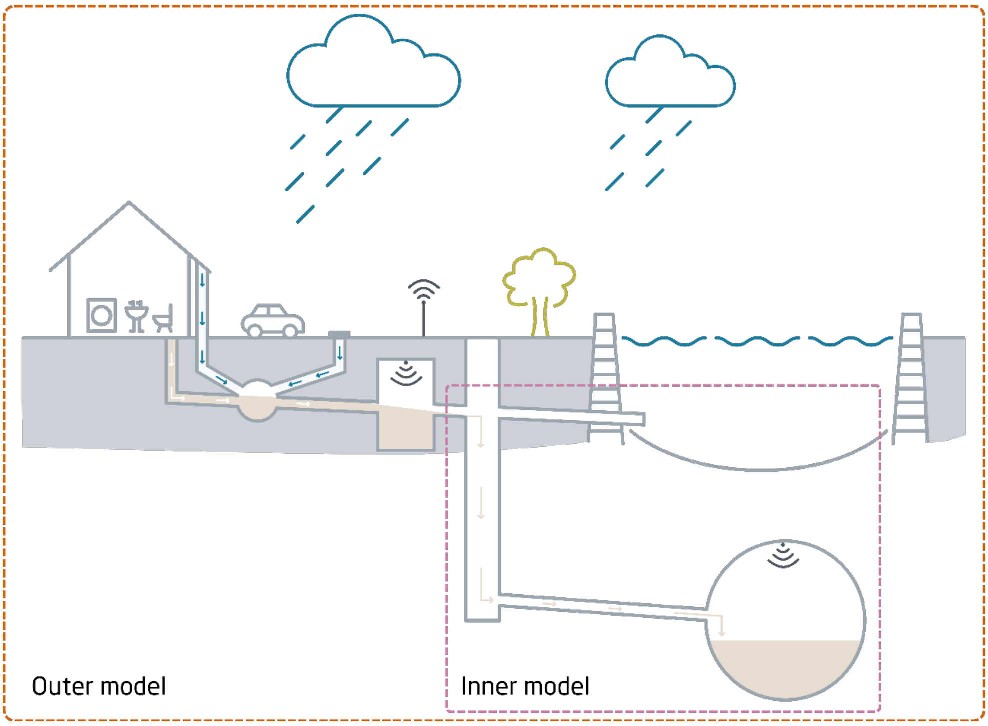

Figure 4. Conceptual overview of a combined sewer system equipped with an overflow tunnel, and boundaries of inner and outer model. Adapted from https://tideway.london/.

3 Updated hydrodynamic models

3.1 1D hydrodynamic models

Distributed urban drainage models (DUDMs) are physics-based models that simulate the spatial and temporal distribution of water across urban drainage systems. DUDMs typically include a model of the catchment, that converts precipitation and discharges in inflows to the network, and a hydrodynamic model of the drainage network, including pipes, basins, and control actuators. The catchment model describes the relative amounts of precipitation runoff, infiltration, and evapotranspiration, and connects the water sources to their point of entry into the sewer. Given the complexity of the surface flow and the large number of sources, the catchment component is usually represented in a lumped or conceptual fashion.

The hydrodynamic component receives as input the location, amount, and timing of the inflows at the inlets and simulates the transport of water to the outlets, where the water leaves the system. The hydrodynamic simulation relies on large amounts of data about the assets, e.g. location and geometry of the pipes, and has a higher level of detail than the catchment component. However, to optimize the computational resources needed, the hydraulic variables are computed by solving the continuity and momentum equations in 1 dimension along the flow direction. This 1D hydrodynamic module is discretized in space at time, so that the hydraulic variables are computed only at specific time and space intervals. Nonetheless, models of large drainage systems typically include thousands of computational nodes for both water depth and flow, that need to be computed at each simulation step.

3.2 Ensemble-based data assimilation

Data assimilation (DA) refers to a class of techniques aimed at improving the accuracy of a model by dynamically adjusting the model states and/or parameters using external data. In the case of urban drainage models, the assimilated data can be observations from the sensors. The assimilation or update thus ensures that the model prediction stays close to reality. This enhances the soft sensing capabilities of the model and provides a reliable initial condition for model forecasts. The update can be restricted to the location where observations are available or be propagated to unobserved locations given the spatial correlation between model variables, which is particularly useful for soft sensing purposes.

In ensemble-based data assimilation, the uncertainty of the model prediction is compared with that of the observation and the model is updated with a weighted average of the two (Figure 4). While the uncertainty of the observation is defined by the user, the uncertainty of the model prediction is derived by running an ensemble of model instances. Each member of the ensemble is an exact copy of the model but is forced with a slightly different input and/or boundary condition. Ensemble-based DA was originally formulated by Evensen (1994) as Ensemble Kalman Filter (EnKF). EnKF is more suitable than non-ensemble assimilation for urban drainage models, provided that the computational resources are not a limiting factor (Borup et al., 2014) To overcome the need of perturbing the observations, Sakov and Oke (2008) introduced a deterministic variation of the EnKF, which is the technique applied to 1D hydrodynamic models in this thesis. A detailed mathematical description of ensemble-based data assimilation is given in Paper I, Section 2.

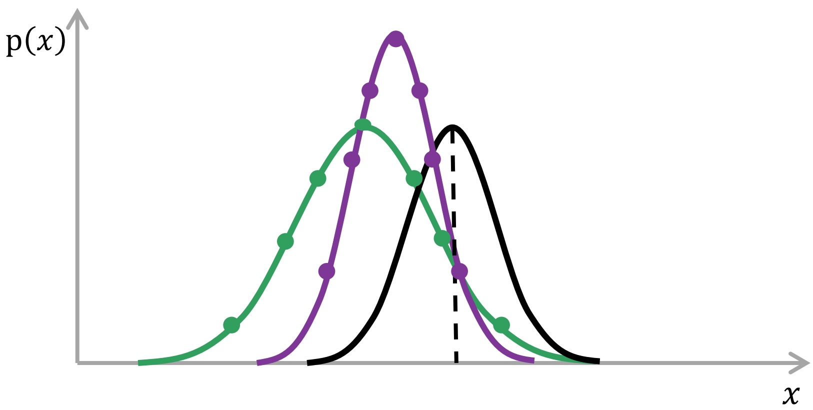

Figure 4. Probability distribution of observation (black), perturbed model ensemble (green) and updated model ensemble (purple).

3.3 Assimilation scheme

A scheme was developed and tested for reducing the computational cost of ensemble-based assimilation in hydrodynamic models of sewers, while retaining the full spatial resolution. This is achieved by running a single instance of the outer model and using its perturbed output as boundary condition to the inner model. The inner model retains the same level of detail of the original model but is limited to a selected subsystem. The boundary between the two models needs to be drawn strategically and should possibly coincide with the location of the sensors. To account for the uncertainty of the model, the hydraulic state at the boundaries is perturbed with a multiplication factor and a time displacement. The perturbation is different for each member of the ensemble.

Since the perturbations are defined arbitrarily, the inner model can be exposed to unrealistic inputs at the boundaries. By assimilating the observed water depth at some of the boundaries, the hydrodynamic simulation becomes more stable. Furthermore, the observations from within the inner model are assimilated and the update propagates through the entire sub-system proportionally to the correlation between states. The scheme is thus composed of three stages: i) running the outer model, ii) perturbing the boundaries of the inner model, iii) assimilating the observations in an ensemble of the inner model (Figure 5). In an offline setting, the stages are executed sequentially for the entire simulation, while in an online scenario the cycle is repeated at each time step. If observations are missing, the assimilation step can be bypassed, and the non-updated ensemble would still represent the uncertainty of the prediction.

This scheme was tailor-made for combined sewer tunnels. These types of structures are connected to the combined sewer via its inlets and outlets. The inlets are typically overflow chambers, which are often equipped with water level sensors and are ideal candidates to act as boundary. Since tunnels are deeper than the surrounding sewer, there are generally no backwater effects towards the outer system. Therefore, if the water level at the outlet is known, the inner and outer model can be effectively run independently. Also, drainage tunnels mostly consist of a single stretch of large pipes. If observations from within the tunnel are assimilated within the inner model, the update would propagate both upstream and downstream, increasing the accuracy of the model and enhancing its soft sensing and sensor validation capabilities.

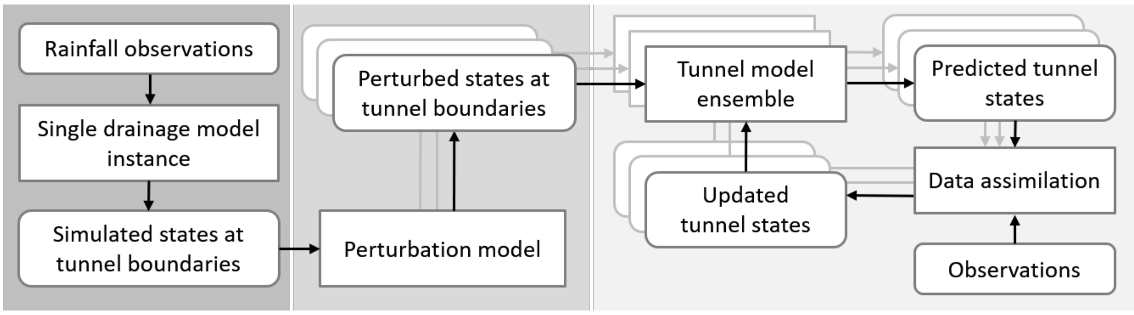

Figure 5. Data flow of a time step in the proposed data assimilation scheme: i) a single instance of the urban drainage model is forced with rainfall observations and the states at the tunnel boundaries are read from its output; ii) the boundary states are perturbed to account for the model uncertainty; iii) the states predicted by the tunnel model ensemble are compared to the available observations to compute the update, which serves as initial condition for the following prediction. From Paper I.

3.4 Damhus drainage tunnel

The Damhus combined sewer overflow tunnel in Copenhagen, Denmark, was used to test the proposed assimilation scheme. The surrounding sewer system was simulated with a highly detailed hydrodynamic model consisting of about 800 km of pipes. The scale of the model makes the direct application of ensemble-based assimilation impractical and justifies the use of the proposed scheme. The outer model receives as input the rainfall observations from 10 gauges distributed on the catchment and returns as output the water depth at the tunnel inlets and outlet. The tunnel itself is about 3.5 km long, with a constant diameter of 3 m and max depth of about 16 m below ground (Figure 6). Several dropshafts connect the tunnel with the shallow system. A large portion of the tunnel inflow is discharged from four overflow chambers connected to the most upstream dropshaft, while a system of pumps located in the most downstream dropshaft is used to empty the tunnel. Water depth sensors are located at three of the upstream overflows, at the tunnel outlet, and in both the upstream and downstream dropshafts. Both rainfall and water depth observations at 1-min resolution were made available for selected rain events when the tunnel was partially or completely filled.

Water depth observations from five sensors were assimilated in the inner model, while the observation from the upstream dropshaft was only used for cross-validation. All boundary states were perturbed with a time displacement uniformly distributed between -30 min and +30 min, and a multiplication factor normally distributed with standard deviation equal to 4%. The uncertainty of the assimilated observation was calibrated to optimize the prediction accuracy at the downstream dropshaft. Ensemble of various sizes were tested and an ensemble of 100 members was found to be a good compromise between accuracy and computational effort.

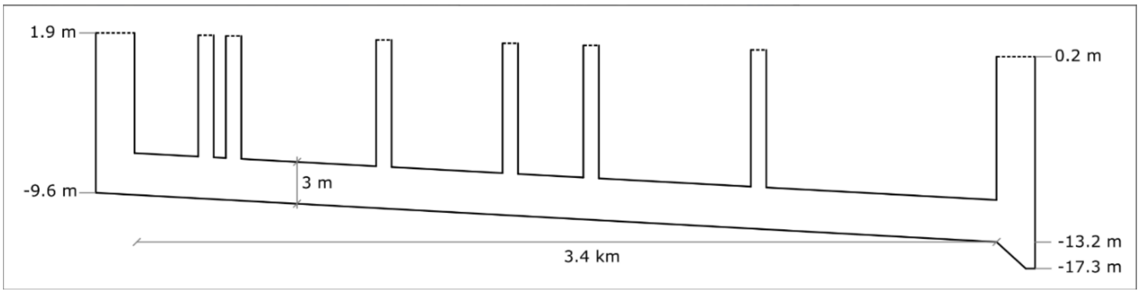

Figure 6. Longitudinal profile of Damhus tunnel system with main dimensions (depths are reported as relative to the Danish Vertical Reference 1990). Adapted from Paper I.

3.5 Sensor cross-validation

Three rainfall events of different magnitude were analysed (A, B and C). Each event covered a cycle of filling, emptying, and flushing of the tunnel. During event A the tunnel was completely filled until it overflowed to the nearby stream. At full capacity, the same water depth was observed at the upstream and downstream dropshaft, and the tunnel effectively behaved as a storage basin. Event A was used for testing the scheme and calibrating its parameters, including the perturbation coefficients and the standard deviation of the assimilated observations. The scheme was validated with event B, which was of moderate intensity and only led to a partial filling of the tunnel. In this case the tunnel behaved as a pipe and the water levels upstream and downstream were different. Nonetheless, the downstream update was propagated through the length of the tunnel and significantly improved the accuracy of the ensemble model prediction at the upstream dropshaft (Figure 7).

Meanwhile event C had a moderate peak intensity, but the sensor observations suggested a different behavior of the tunnel. While the downstream dropshaft was only partially filled, the water level at the upstream dropshaft was seen nearly reaching the crest level. After the assimilation, the model ensemble predicteded with confidence that the upstream water level was close to the invert, in disagreement with the sensor. Upon further investigation, it was discovered that the upstream sensor was affected by false echo when the inflow to the tunnel intersected the field of vision of the sensor. This application demonstrates that the updated hydrodynamic model can both be used for system-wide soft sensing, as the update propagates upstream, and validation of the non-assimilated observations.

For both soft sensing and anomaly detection purposes, it is important to quantify the degree of confidence in the model predictions. Classic accuracy metrics, e.g. Root Mean Square Error (RMSE) and Nash-Sutcliff Efficiency, do not fully capture the probabilistic information contained in the ensemble prediction. Other metrics were tested and compared, which are more suitable for probabilistic forecasts (e.g. the Continuous Ranked Probability Score). To validate the upstream observations the standardized residual of the ensemble prediction was computed at each time step. This was obtained by dividing the residual error of the ensemble mean by the expected standard deviation, which accounts for both the uncertainty of the sensor and the mode. All observations corresponding to a standardized residual error outside the [-2,2] interval were flagged as anomalous (Figure 7).

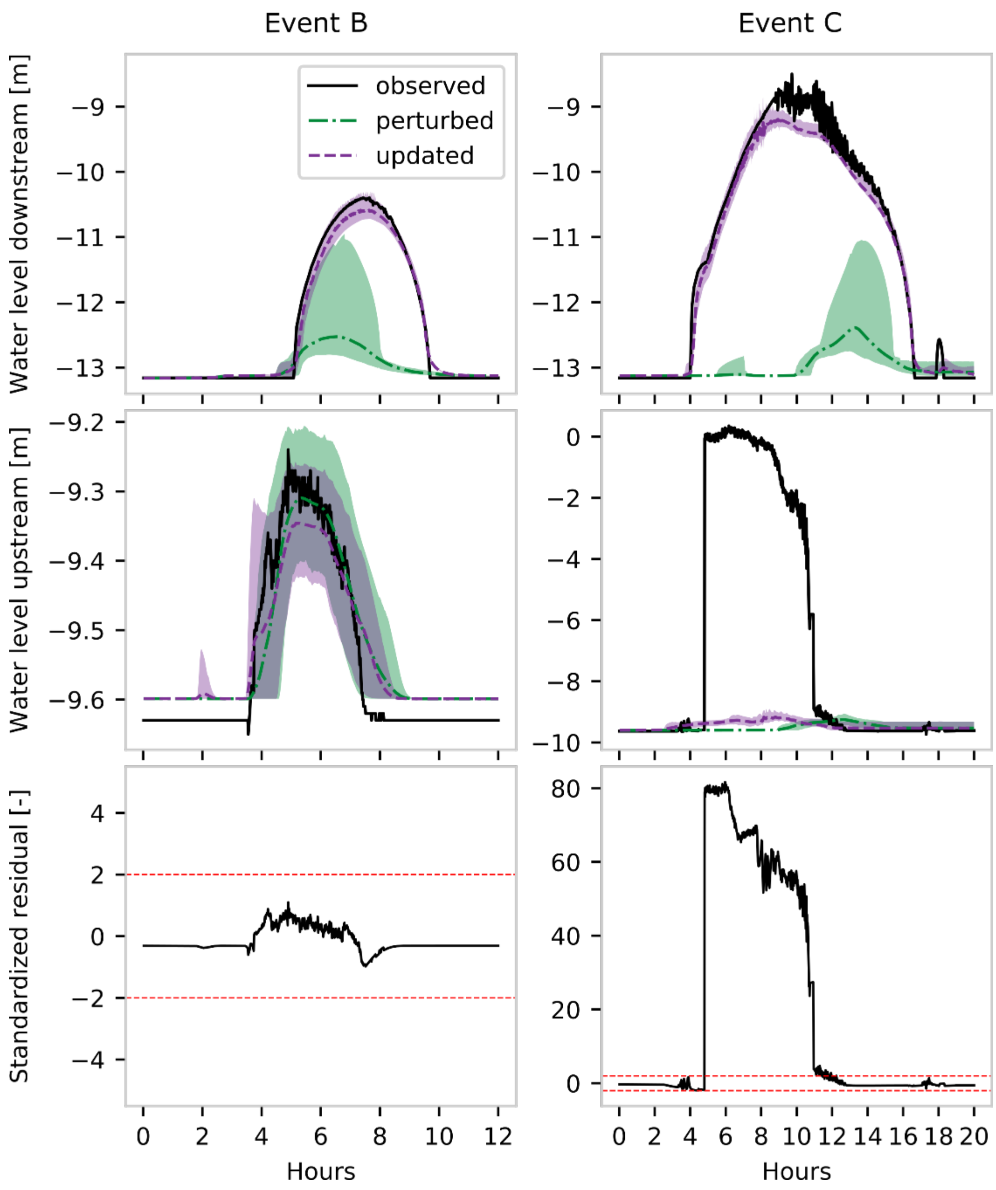

Figure 7. Water levels for Events B and C at the downstream (top row) and upstream dropshaft (middle row): observations in solid black line, and model ensembles (bands delimited by 5% and 95% quantiles, mean shown as dashed lines) without (green) and with (purple) data assimilation; Standardized residual (bottom row) of updated water level at upstream dropshaft. Adapted from Paper I.

4 Long Short-Term Memory networks

4.1 Artificial Neural networks

Artificial Neural Networks (ANN) mimic the learning processes of a human brain. The input signals are transformed and distributed between the artificial neurons via connections resembling artificial synapses (Figure 8). Repetitive patterns in the input signal strengthen some paths, thus embedding knowledge in the network. In practice, a weight is assigned to each connection and the optimal set of weights is calibrated in the training process by minimizing the prediction error of the output. If trained with a sufficiently representative dataset, the network is capable of predicting the expected behavior of the system in response to an independent set of inputs. This property can be exploited in modelling the complex non-linear behavior of combined sewer systems without any knowledge of their geometric and hydraulic properties.

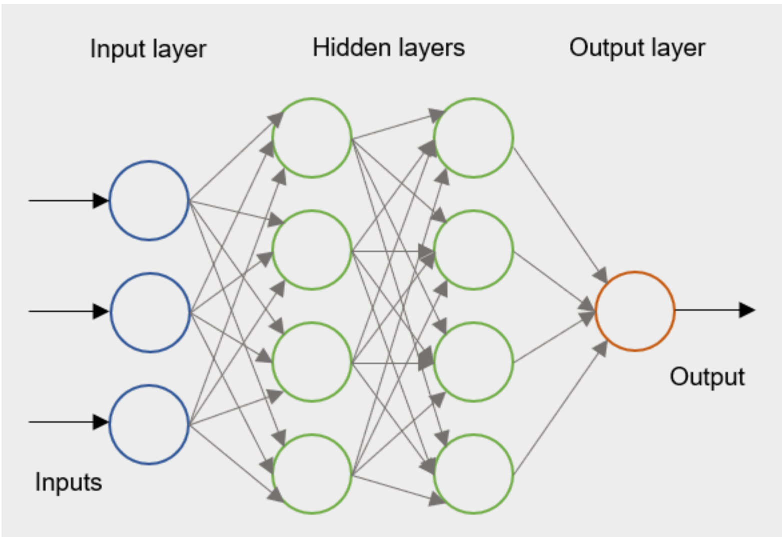

Figure 8. Internal structure of an Artificial Neural Network with neurons (circles) and synapses (grey arrows).

ANNs have been used extensively for forecasting the behavior of combined sewer systems, but only a few applications explicitly addressed soft sensing and anomaly detection. For example, Bailey et al. (2016) successfully used an ANN for the early detection of blockages. Higher prediction accuracies can be achieved with the recurrent type of neural networks, which learn time dependencies in the training data and are thus ideal for series of observations. Long Short-Term Memory (LSTM) neural networks are a variant of recurrent ANN which can store long-term information with a system of gates that further manipulate the signal being transmitted through the network (Hochreiter &

Schmidhuber, 1997). LSTM networks have been proven to outperform ANN in predicting both flows (Sufi Karimi et al., 2019) and water levels (Zhang et al., 2018) in combined sewers.

4.2 Setting up an LSTM network

LSTMs are very flexible tools and modern software libraries offer a myriad of options to customize the network to the problem at hand. The key settings to be defined are the number of layers and neurons, which can be seen as the depth and width of the network and determine the learning capacity. The input data is arranged in a 3-dimensional structure (Figure 9). A window of input data is associated to each target observation. The size of the window is given by the number and length of the input series (features). Gaps can be introduced in the input window to simulate scenarios of missing observations. The number of windows is determined by the length of the period for which learning data is available. The input signal is transformed at various steps of the learning process according to user-defined functions, which also affect the accuracy of the predictions. All together, these settings constitute the non-trainable parameters of the neural network and are also termed hyperparameters to differentiate them from the trainable parameters (internal weights).

The choice of optimal hyperparameters mainly depends on the characteristics of the output to be predicted and the number and type of input features. In this thesis, LSTM networks were used exclusively to predict water depth from combined sewer overflow chambers. To limit the required size of the network and consequently the computational effort, the set of inputs was reduced to only three features: minute of the day, for the recognition of the dry weather pattern; rainfall intensity, for the prediction of the wet weather response; and information on the antecedent hydraulic state. While the first two features are generally robust and reliable sources of information, the antecedent observations can potentially be missing or erroneous. This can undermine the usefulness of the prediction. The network learns that the antecedent observations have the highest correlation to the current observation among the three features and assigns them a higher weight. The prediction accuracy thus becomes largely dependent on the availability and validity of the antecedent observations, which is counterproductive when the aim of the prediction is to detect anomalous observations. To overcome this dependency, the network can be trained to operate without knowing the antecedent state or only a surrogate version, e.g. the baseline value of the day before.

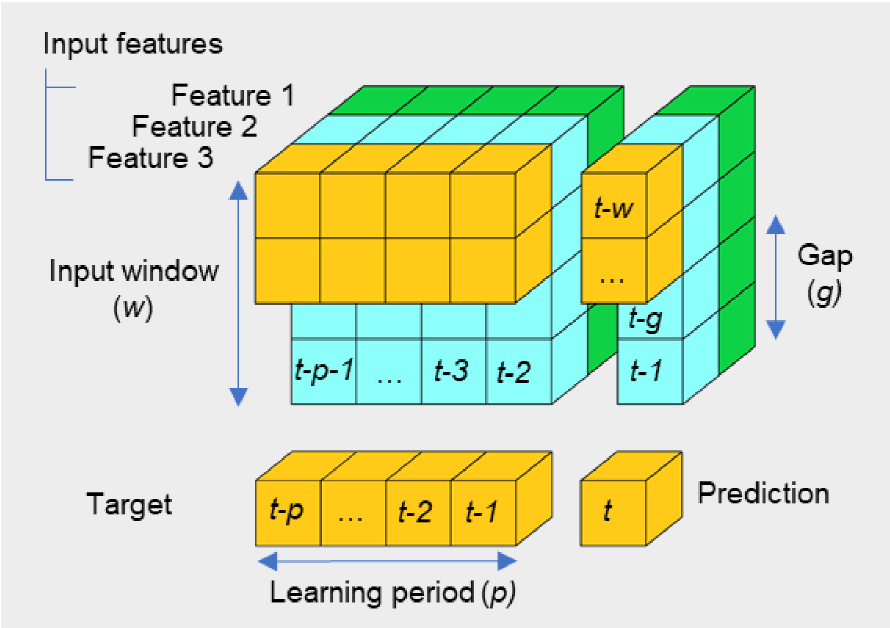

Figure 9. LSTM input features and target, with input window length (w), gap length (g) and learning period (p). Adapted from Paper II.

4.3 Prediction accuracy

Water depth observations from four different locations were used for testing the LSTM network and assessing the prediction accuracy. The first series was from one of the boundary structures of the Damhus tunnel system described in Section 3.4, and includes 11 months of observations at 1-min resolution (Paper II). The other three series were from different locations in the combined sewer system of Odense, Denmark, and cover periods ranging from 1.5 to nearly 3 years also at 1-min resolution (Paper III). All series were split into training, validation and testing subsets. The validation subset was used to guide the training process, which explains why a third independent subset was needed for testing the prediction accuracy.

For both Damhus and Odense cases, rainfall intensity observations from different gauges were combined in a single average series, thus limiting the number of features and the size of the input window. The average series showed in all cases a higher correlation to the observed water depth than the individual series. The correlation was highest for delays between 30 and 60 min and decreased significantly after 2 hours. This is a consequence of the nature of the hydrological processes occurring in a combined sewer network, as the inflows from the sub-catchments have different travel time to the observation location. Also, using the average series as input ensures that if one of the gauges becomes faulty, it does not compromise the network prediction.

As for the antecedent water depth observations, various approaches were tested to determine how much the knowledge of the recent past influences the prediction accuracy. Two networks with the same set of hyperparameters were trained with (scenario A) and without (scenario B) the antecedent observation as input. By comparing the RMSE of the network prediction in the two scenarios for different lengths of the input window it became evident that i) the network generally performed better if the antecedent observations were known, and ii) if the antecedent observations were unknown, a longer window was needed to capture the delayed response of the water depth to the rainfall input (Figure 10).

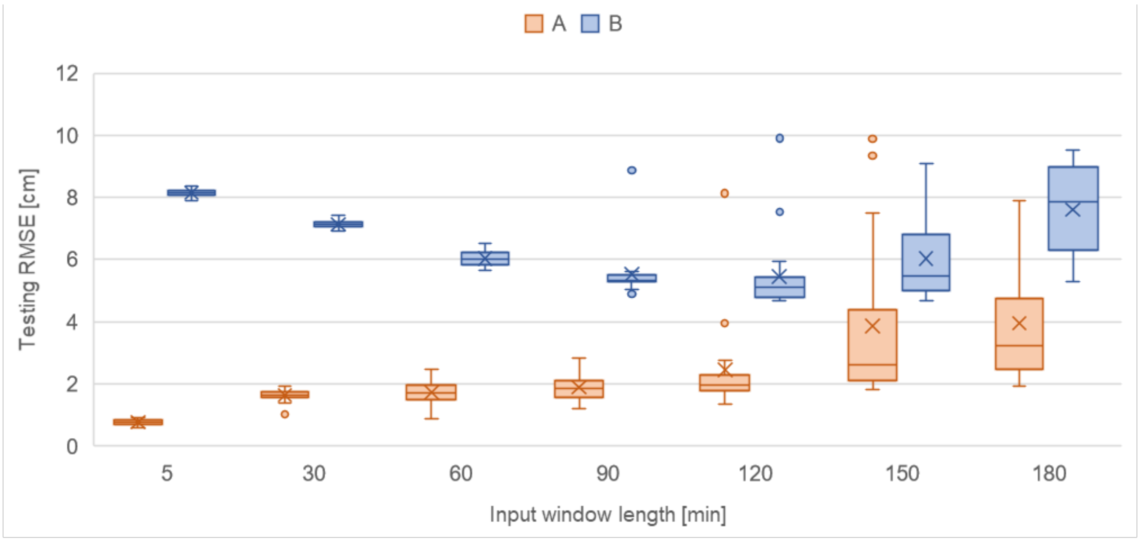

Figure 10. Root Mean Square Error (RMSE) of 20 LSTM predictions in the testing dataset from the Damhus case with different input window lengths. Networks trained with (A) or without (B) the antecedent observations. Results presented are quartiles (box) and mean (cross). The whiskers extend to the minimum and maximum value or 1.5 times the inter quantile range. Values outside the whiskers range are marked as outliers (circles). Adapted from Paper II.

Intermediate scenarios were also tested where a gap of constant length was introduced in the antecedent observations. This forced the network to compensate the missing information with the other input features and yielded intermediate accuracy compared to the other two scenarios. However, the LSTM best performed in a scenario if trained on the same scenario. This also applied to specific lengths of the gap. Therefore, in an operational setting several instances of the network should be maintained and applied at the occurrence of the specific scenario, which is a considerable computational overhead. On the other hand, the network trained without antecedent observations (scenario B) performed sufficiently well for soft sensing purposes while being independent from the availability and quality of observations.

Results from the Damhus case had, however, limited validity due to the short period of time used for testing (about 2 months), during which seasonal variations were negligible. Since the dry weather pattern was only described by the minute of the day, the same average behavior was predicted for all dry days of the year. This is often not the case due to seasonal variations in the amount of infiltrating water and in the response of the catchment to the rainfall. The validity of the approach for longer series of data was tested with observations from the Odense sewer system. It became evident that a large error was introduced by averaging the dry weather behavior over the year. Several combinations of input features were tested to identify one that could capture the seasonal variations without being vulnerable to sudden errors in the observation series. The 5% quantile of the previous day's observations proved the best option as it helped predicting consistently the water depth baseline (Figure 4 from Paper III) and it was only affected by long-lasting anomalies or gaps. In the worst case, the observed baseline could be replaced with the predicted one and would still probably yield a sufficient prediction accuracy, since the baseline normally changes gradually over days and weeks.

4.4 Anomaly detection indicators

To upgrade an LSTM from soft sensor to anomaly detection tool, the deviation between the prediction and the observation needs to be quantified. In the simplest case, the residual error of the LSTM prediction can be used as anomaly indicator. However, the error is generally larger in wet weather, due to the uncertain response to the rainfall, and lower in dry weather, when the observations fall within a relatively narrow range. To account for the state-dependency of the error, the absolute value of the residual was divided by the predicted value to obtain the relative residual (Figure 11). A small constant was added at the denominator to prevent the indicator from spiking when the prediction neared the invert level. Due to delays between predictions and observations, spikes of residual error were also observed at the beginning and the end of the wet weather events. These were downplayed by averaging the residual error over a day-long moving window. Using this indicator, only the long-lasting anomalies were highlighted, e.g. the blockage shown in Figure 11.

A detection threshold was computed for both indicators as the mean plus twice the standard deviation over the validation period. The threshold represented the level of confidence in the LSTM prediction and changed across locations. In theory, only the statistically significant events should exceed the threshold and be flagged as anomalous. However, the residual error exceeded the threshold on several occasions when the system was behaving as expected, potentially leading to numerous false alarms. This advocates for the use of separate indicators to discriminate between levels of urgency of the alarm. For example, a yellow warning could be raised when observations are filtered out because they are physically unrealistic; orange warnings could correspond to relative residual errors above the detection threshold; red warnings could be raised only when the daily average residual error threshold is exceeded.

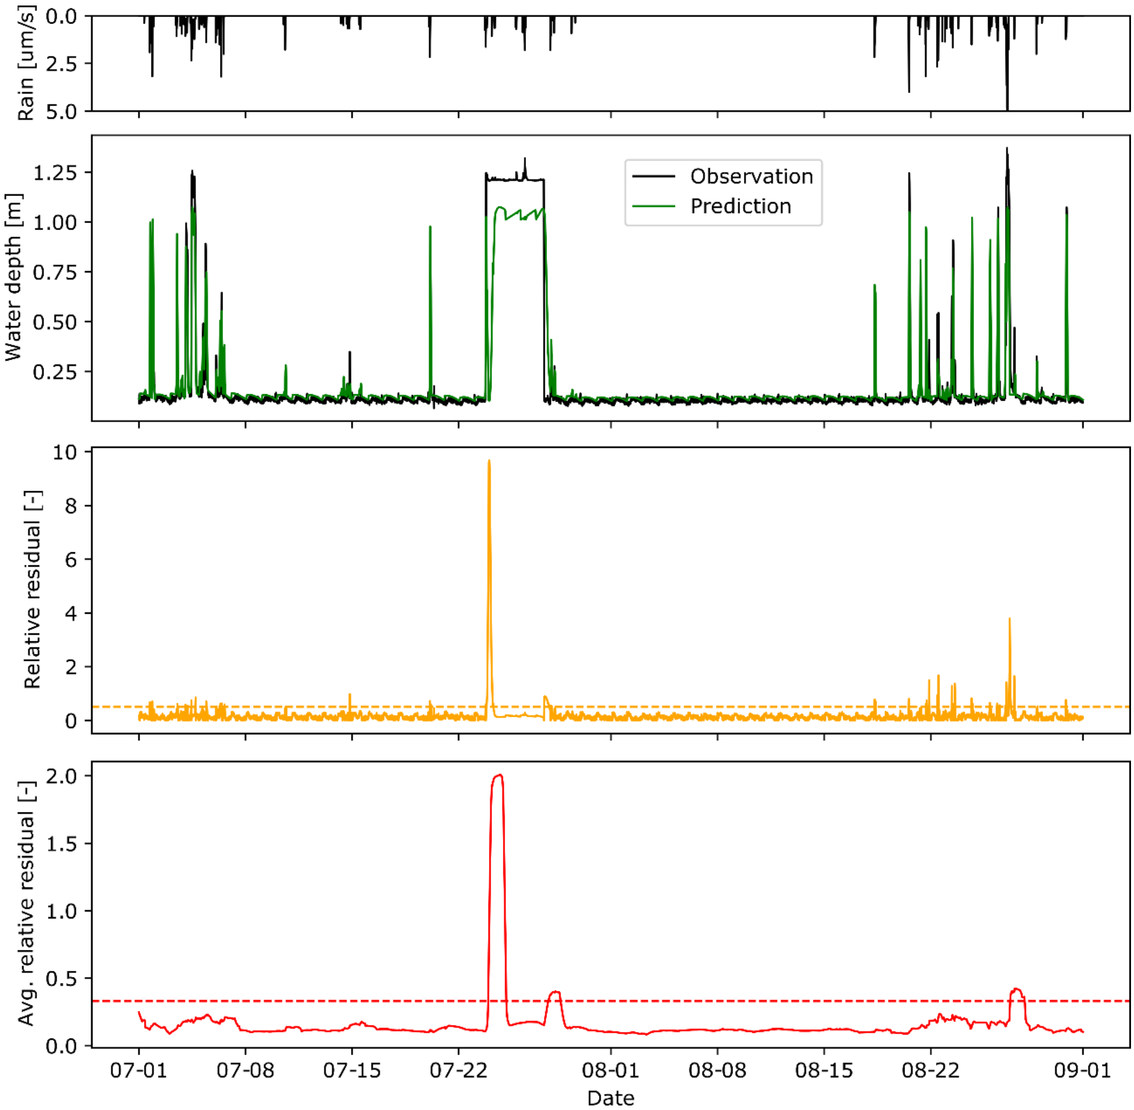

Figure 11. From top to bottom: rainfall intensity observations; water depth observation from a combined sewer overflow in Odense and LSTM prediction; relative residual error with anomaly threshold; average daily residual error with anomaly threshold. Thresholds are calculated as mean plus twice the standard deviation of the indicator over the validation period. Adapted from Paper III.

5 Discussion

This thesis investigated the use of two very different models for soft sensing and anomaly detection in combined sewer systems. 1D hydrodynamic models represent the state of the art in urban drainage modelling and are the result of decades of iterative refinements. They achieve a fine balance between level of detail and computational cost and cover all the main hydrological and hydraulic processes occurring in a sewer system. The model prediction is generally accurate enough for planning and monitoring purposes, but it lacks an assessment of the uncertainty. The main advantage of this type of models is that the prediction has a fine spatial and temporal resolution and is extended to the entire system.

Data assimilation is capable of enhancing the features of 1D hydrodynamic models, by dynamically adjusting the system states to better fit the available observations. Thus, an updated model can be regarded as hybrid physics-based and data-driven model. The model updates can significantly improve the prediction accuracy, while adding a probabilistic dimension. With ensemble-based methods, the uncertainty of the model prediction is computed from the spread of an ensemble of model instances. Running a large enough ensemble is a resource-demanding task and has been one of the limiting factors in the adoption of data assimilation techniques in the field of urban drainage. This thesis presented a scheme for localized data assimilation that enables soft sensing and anomaly detection capabilities in 1D hydrodynamic models at feasible computational costs.

LSTM neural networks can learn repetitive patterns in water depth observations from combined sewers and predict the behavior of the system in response to a given set of inputs. As such, they lack any knowledge on the physical processes occurring in the system and are purely data-driven. For this reason, their spatial and temporal validity is limited to that of the training dataset. If trained with a series of observation from a specific location, the LSTM prediction has point-wise validity as opposed to the system-wide prediction of hydrodynamic models. This limitation could be overcome by training a single network on several observation series at once, but this approach was not investigated in this thesis. Depending on the complexity of the network architecture and the quantity of training data, an LSTM neural network can be run at a fraction of the computational cost required by physics-based models. By default, the LSTM prediction is deterministic, but various possibilities to add a probabilistic dimension are discussed in Paper III.

The key features of the investigated models are summarized in Table 3. It should be noted that the computational cost and prediction accuracy assessments are qualitative and comparative in nature and based on the general experience of the author. From the comparison, no clear winner emerges. Generally, the computational cost increases with the prediction accuracy. Therefore, the resource-demanding updated hydrodynamic models could be reserved for soft sensing and anomaly detection in selected, strategic sub-systems of the larger combined sewer, e.g. large tunnels. Instead, LSTM networks can easily be deployed to several individual locations and could be suitable to supplement and validate large networks of sensors.

Table 3. Key features of investigated models for soft sensing and anomaly detection in combined sewer systems

Model	**1D Hydrodynamic**	**Updated 1D hydrodynamic**	**LSTM neural network**
Mechanism	Physics-driven	Hybrid physics- and data-driven	Data-driven
Spatial extent	System-wide	System-wide (inner model)	Point-wise
Computational cost	●●	●●●	●
Prediction accuracy	●●	●●●	●
Type of prediction	Deterministic	Probabilistic	Deterministic or probabilistic

6 Conclusions

The following conclusions were drawn in response to the research questions formulated in Section 1.4.

i) A limited number of examples can be found in the scientific literature of the use of “soft sensing” and “anomaly detection” in connection to urban drainage systems, despite their popularity in other fields of research. The concept of using a model prediction to extend and validate the information from the hardware sensors is well established, and both physics-based and data-driven models are used for the purpose. These applications are sometimes referred to as “gap filling”, “data reconstruction” and “data validation”. In this thesis, the terms soft sensing and anomaly detection imply that the underlying model is capable of running in real time and are preferred for their broad meaning.

ii) 1D hydrodynamic models represent an optimal compromise between level of detail and computational effort required. They can potentially qualify for soft sensing and anomaly detection already in their deterministic formulation. Ensemble-based data assimilation was capable of significantly improving the prediction accuracy of the 1D hydrodynamic model of an urban drainage tunnel and added a probabilistic dimension to its output. As a consequence, the updated model could be trusted to simulate the behavior of unobserved locations and observations could be flagged as anomalous if they fell outside the confidence interval of the ensemble prediction. The only drawback is that running a sufficiently large ensemble of models in parallels requires prohibitive computational resources for large models. A solution was proposed that restricts the ensemble application to a selected area of the combined sewer, thus optimizing the computational resources for strategic locations. However, the efficacy of the assimilation relies on the availability and quality of the assimilated observations, which should in principle be independently validated to prevent them from corrupting the update.

iii) A Long Short-Term Memory neural network was optimized to predict water depth observations from combined sewer overflows and tested with several scenarios of data availability. The prediction accuracy was highest when the antecedent observations were used as input. In this scenario, though, the network prediction was vulnerable to missing and anomalous observations. However, if appropriately trained, the LSTM was capable of compensating the missing information on the past state of the system with the other available input features, namely minute of the day and rainfall intensity. This yielded a more robust soft sensing model that consistently predicted the water depth with sufficient accuracy. The LSTM prediction was also used as basis for the automatic detection of anomalies in incoming observations from the hardware sensor. A detection method was proposed that uses simple indicators and thresholds representing the level of confidence in the model. Despite the indicators were tailored on the characteristics of water depth observations, the frequency of false alarms was still too high for practical applications. False alarms could be reduced either by further improving the prediction accuracy of the model or accounting for the large uncertainty in predicting the exact height and timing of the wet weather peaks.

iv) Both investigated models have their own advantages and disadvantages. Updated hydrodynamic models have higher computational cost but return more accurate and system-wide predictions. On the other hand, LSTM networks can be easily deployed to large sensor networks with minimal domain knowledge but have limited spatial validity. Therefore, as a general principle, the two approaches are better suited from the inner and outer model, respectively. A synergy between the two approaches remains to be explored.

7 Perspectives

Ensemble-based data assimilation was proven to significantly improve the prediction accuracy of 1D hydrodynamic models. The restriction of the update to a selected sub-system was justified by the need to contain the computational cost. However, as more tasks are assigned to computational servers and the cost to performance ratio of hardware decreases, the argument is bound to become less relevant in the future. At that point, ensembles of large hydrodynamic models could be run in parallel and observations from tens or hundreds of sensors could be assimilated at once. Also, information on the control actuators (e.g. the state of a gate or the flow through a pump) could be directly assimilated in the hydrodynamic model to yield an even better fit to the actual state of the system. The interplay between the different and potentially contrasting sources of information should than be investigated to prevent physically unrealistic phenomena to be introduced in the model. Finally, the use of updated 1D hydrodynamic models for probabilistic forecasting deserves further investigation, especially for the potential of informing risk-based control decisions.

Long Short-Term Memory neural networks were used to predict water depth from combined sewers with sufficient accuracy. The same network was used for both dry and wet weather periods with different degrees of success. LSTM predictions were generally less accurate in wet weather, when the response of the water depth to the rainfall input is highly non-linear due to the dynamic conditions of the catchment and the specific properties of overflow chambers. For example, when the water level exceeds a defined threshold the excess inflow is diverted. This threshold limits how much the water depth can increase in response to peaks of rainfall intensity and contributes to the non-linearity of the problem. In this thesis, it was chosen to apply the same network to all rainfall regimes (dry, wet and overflowing) to avoid the need of classifying the incoming observations, as it could introduce some form of bias. However, the classification could potentially be fully automated by means of fuzzy logic, as already done in hydrological applications (Talei et al., 2013). Also, probabilistic LSTM predictions have been proposed with an ensemble approach. Alternatively, a single neural network could be set to return as output the statistical properties of the observation (e.g. mean and standard deviation), provided that an adequate probabilistic loss function is used. A similar approach has already been described in the scientific literature (Gulshad et al., 2017) but no application can be found in the field of urban drainage so far.

Finally, there is untapped potential in the synergy between the two investigated techniques. As discussed, data assimilation relies on the assumption that the assimilated observations are valid. Therefore, a system is needed to independently assess the quality of the observations and cross-validation can only partially answer to this need. LSTM networks would serve the purpose well, since they are trained with historical observations from a single location and are independent from the nearby sensors. A probabilistic LSTM would even dynamically quantify the uncertainty of the assimilated observation, thus replacing the otherwise arbitrary parameters of the proposed assimilation scheme.

8 References

Ahm, M., Thorndahl, S., Nielsen, J. E., & Rasmussen, M. R. (2016). Estimation of combined sewer overflow discharge: a software sensor approach based on local water level measurements. *Water Science and Technology*, *74*(11), 2683–2696. https://doi.org/10.2166/wst.2016.361

Arnbjerg-Nielsen, K., Willems, P., Olsson, J., Beecham, S., Pathirana, A., Gregersen, I. B., Madsen, H., & Nguyen, V.-T.-V. (2013). Impacts of climate change on rainfall extremes and urban drainage systems: a review. *Water Science and Technology*, *68*(1), 16–28. https://doi.org/10.2166/wst.2013.251

Bailey, J., Harris, E., Keedwell, E., Djordjevic, S., & Kapelan, Z. (2016). The Use of Telemetry Data for the Identification of Issues at Combined Sewer Overflows. *Procedia Engineering*, *154*, 1201–1208. https://doi.org/10.1016/j.proeng.2016.07.524

Benedetti, L., Bixio, D., Claeys, F., & Vanrolleghem, P. A. (2008). Tools to support a model-based methodology for emission/immission and benefit/cost/risk analysis of wastewater systems that considers uncertainty. *Environmental Modelling and Software*, *23*(8), 1082–1091. https://doi.org/10.1016/j.envsoft.2008.01.001

Blumensaat, F., Leitão, J. P., Ort, C., Rieckermann, J., Scheidegger, A., Vanrolleghem, P. A., & Villez, K. (2019). How Urban Storm- And Wastewater Management Prepares for Emerging Opportunities and Threats: Digital Transformation, Ubiquitous Sensing, New Data Sources, and beyond - A Horizon Scan. *Environmental Science and Technology*, *53*(15), 8488–8498. https://doi.org/10.1021/acs.est.8b06481

Borup, M., Grum, M., Madsen, H., & Mikkelsen, P. S. (2014). Updating Distributed Hydrodynamic Urban Drainage Models. *Proceedings of the 13th International Conference on Urban Drainage*, *September*, 1–8. https://doi.org/10.13140/2.1.1682.1766

Borup, M., Grum, M., & Mikkelsen, P. S. (2011). Real time adjustment of slow changing flow components in distributed urban runoff models. *12th International Conference on Urban Drainage*, *1*(September), 11–16. http://orbit.dtu.dk/ws/files/6373440/prod11321622323249.PAP005261.pdf

Branisavljević, N., Kapelan, Z., & Prodanović, D. (2011). Improved real-time data anomaly detection using context classification. *Journal of Hydroinformatics*, *13*(3), 307–323. https://doi.org/10.2166/hydro.2011.042

Branisavljević, N., Prodanović, D., Pavlović, D., (2010). Automatic, semi-automatic and manual validation of urban drainage data. *Water Science and Technology*, *62*(5), 1013–1021. https://doi.org/10.2166/wst.2010.350

Breinholt, A., Thordarson, F. Ö., Møller, J. K., Grum, M., Mikkelsen, P. S., & Madsen, H. (2011). Grey-box modelling of flow in sewer systems with state-dependent diffusion. *Environmetrics*, *22*(8), 946–961. https://doi.org/10.1002/env.1135

Campisano, A., Ple, J. C., Muschalla, D., Pleau, M., Vanrolleghem, P. A., & Cabot Ple, J. (2013). Potential and limitations of modern equipment for real time control of urban wastewater systems Potential and limitations of modern equipment for real time control of urban wastewater systems. *Urban Water Journal*, *10*, 300–311. https://doi.org/10.1080/1573062X.2013.763996

Carstensen, J., Harremoës, P., & Strube, R. (1996). Software sensors based on the grey-box modelling approach. *Water Science and Technology*, *33*(1), 117–126. https://doi.org/10.1016/0273-1223(96)00164-3

Chandola, V., Banerjee, A., Kumar, V., & Almanza-Garcia, S. (2009). Anomaly detection. *ACM Computing Surveys*, *14*(3), 1–22. https://doi.org/10.1145/1541880.1541882

de Ville, N., Le, H. M., Schmidt, L., & Verbanck, M. A. (2017). Data-mining analysis of in-sewer infiltration patterns: seasonal characteristics of clear water seepage into Brussels main sewers. *Urban Water Journal*, *14*(10), 1090–1096. https://doi.org/10.1080/1573062X.2017.1363252

Deletic, A., Dotto, C. B. S., McCarthy, D. T., Kleidorfer, M., Freni, G., Mannina, G., Uhl, M., Henrichs, M., Fletcher, T. D., Rauch, W., Bertrand-Krajewski, J. L., & Tait, S. (2012). Assessing uncertainties in urban drainage models. *Physics and Chemistry of the Earth*, *42–44*, 3–10. https://doi.org/10.1016/j.pce.2011.04.007

Ding, D., Han, Q. L., Xiang, Y., Ge, X., & Zhang, X. M. (2018). A survey on security control and attack detection for industrial cyber-physical systems. *Neurocomputing*, *275*, 1674–1683. https://doi.org/10.1016/j.neucom.2017.10.009

Dolowitz, D. P., Bell, S., & Keeley, M. (2018). Retrofitting urban drainage infrastructure: green or grey? *Urban Water Journal*, *15*(1), 83–91. https://doi.org/10.1080/1573062X.2017.1396352

Ebi, C., Schaltegger, F., Rust, A., & Blumensaat, F. (2019). Synchronous LoRa Mesh Network to Monitor Processes in Underground Infrastructure. *IEEE Access*, *7*, 57663–57677. https://doi.org/10.1109/ACCESS.2019.2913985

Eggimann, S., Mutzner, L., Wani, O., Schneider, M. Y., Spuhler, D., Moy De Vitry, M., Beutler, P., & Maurer, M. (2017). The Potential of Knowing More: A Review of Data-Driven Urban Water Management. *Environmental Science and Technology*, *51*(5), 2538–2553. https://doi.org/10.1021/acs.est.6b04267

Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, *99*(C5). https://doi.org/10.1029/94jc00572

Fletcher, T. D., Shuster, W., Hunt, W. F., Ashley, R., Butler, D., Arthur, S., Trowsdale, S., Barraud, S., Semadeni-Davies, A., Bertrand-Krajewski, J.-L., Mikkelsen, P. S., Rivard, G., Uhl, M., Dagenais, D., & Viklander, M. (2015). SUDS, LID, BMPs, WSUD and more – The evolution and application of terminology surrounding urban drainage. *Urban Water Journal*, *12*(7), 525–542. https://doi.org/10.1080/1573062X.2014.916314

García, L., Barreiro-Gomez, J., Escobar, E., Téllez, D., Quijano, N., & Ocampo-Martinez, C. (2015). Modeling and real-time control of urban drainage systems: A review. *Advances in Water Resources*, *85*, 120–132. https://doi.org/10.1016/j.advwatres.2015.08.007

L. Fortuna, L., Graziani, S., Rizzo, A., Xibilia, M.G. (2007). Soft Sensors for Monitoring and Control of Industrial Processes. *Advances in Industrial Control* (Springer).

Gulshad, S., Sigmund, D., Kim, J.H. (2017). Learning to reproduce stochastic time series using stochastic LSTM. *Proceedings of the International Joint Conference on Neural Networks* 2017-May, 859–866. https://doi.org/10.1109/IJCNN.2017.7965942

Hassan, S. I., Dang, L. M., Mehmood, I., Im, S., Choi, C., Kang, J., Park, Y. S., & Moon, H. (2019). Underground sewer pipe condition assessment based on convolutional neural networks. *Automation in Construction*, *106*. https://doi.org/10.1016/j.autcon.2019.102849

Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. *Neural Computation*, *9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Isel, S., Dufresne, M., Bardiaux, J. B., Fischer, M., & Vazquez, J. (2014). Computational fluid dynamics based assessment of discharge-water depth relationships for combined sewer overflows. *Urban Water Journal*, *11*(8), 631–640. https://doi.org/10.1080/1573062X.2013.806561

Kerkez, B., Gruden, C., Lewis, M., Montestruque, L., Quigley, M., Wong, B., Bedig, A., Kertesz, R., Braun, T., Cadwalader, O., Poresky, A., & Pak, C. (2016). Smarter Stormwater Systems. *Environmental Science and Technology, 50*(14), 7267−7273. https://doi.org/10.1021/acs.est.5b05870

Kwon, D., Kim, H., Kim, J., Suh, S. C., Kim, I., & Kim, K. J. (2019). A survey of deep learning-based network anomaly detection. *Cluster Computing*, *22*(1), 949–961. https://doi.org/10.1007/s10586-017-1117-8

Langeveld, J., van Daal, P., Schilperoort, R., Nopens, I., Flameling, T., & Weijers, S. (2017). Empirical sewer water quality model for generating influent data for WWTP modelling. *Water (Switzerland)*, *9*(7). https://doi.org/10.3390/w9070492

Leonhardt, G., Fach, S., Engelhard, C., Kinzel, H., & Rauch, W. (2012). *A software-based sensor for combined sewer overflows*. https://doi.org/10.2166/wst.2012.331

Loke, E., Warnaars, E. A., Jacobsen, P., Nelen, F., & Almeida, M. do C. (1997). Artificial neural networks as a tool in urban storm drainage. *Water Science and Technology*, *36*(8–9), 101–109. https://doi.org/10.2166/wst.1997.0651

Makropoulos, C., & Savíc, D. A. (2019). Urban hydroinformatics: Past, present and future. In *Water (Switzerland)* (Vol. 11, Issue 10). MDPI AG. https://doi.org/10.3390/w11101959

Moradi, S., Zayed, T., Nasiri, F., & Golkhoo, F. (2020). Automated Anomaly Detection and Localization in Sewer Inspection Videos Using Proportional Data Modeling and Deep Learning–Based Text Recognition. *Journal of Infrastructure Systems*, *26*(3), 04020018. https://doi.org/10.1061/(asce)is.1943-555x.0000553

Mourad, M., & Bertrand-Krajewski, J.-L. (2002). A method for automatic validation of long time series of data in urban hydrology. *Water Science and Technology, 45*(4-5), 263-270 http://iwaponline.com/wst/article-pdf/45/4-5/263/734567/263.pdf

Moy De Vitry, M., Schneider, M. Y., Wani, O., Manny, L., Leitao, J. P., & Eggimann, S. (2019). Smart urban water systems: What could possibly go wrong? In *Environmental Research Letters* (Vol. 14, Issue 8). Institute of Physics Publishing. https://doi.org/10.1088/1748-9326/ab3761

Palmitessa, R., Borup, M., & Mikkelsen, P. S. (2018). Urban tunnel systems for conveyance and storage of storm- and wastewater: features, classification, and modelling. In G. Mannina (Ed.), *Proceedings of the 11th International Conference on Urban Drainage Modelling* (pp. 251–254). http://orbit.dtu.dk/ws/files/154733729/UDM18_abstract_Palmitessa_2906.pdf

Pedersen, A. N., Borup, M., Brink-Kjaer, A., Christiansen, L. E., & Mikkelsen, P. S. (2021). Living and Prototyping Digital Twins for Urban Water Systems: Towards Multi-Purpose Value Creation Using Models and Sensors. *Water*, *13*(5), 592. https://doi.org/10.3390/w13050592

Pedersen, J. W., Larsen, L. H., Thirsing, C., & Vezzaro, L. (2020). Reconstruction of corrupted datasets from ammonium-ISE sensors at WRRFs through merging with daily composite samples. *Water Research*, *185*, 116227. https://doi.org/10.1016/j.watres.2020.116227

Rajkumar, R., Lee, I., Sha, L., & Stankovic, J. (2010). Cyber-physical systems: The next computing revolution. *Proceedings - Design Automation Conference*, 731–736. https://doi.org/10.1145/1837274.1837461

Ramotsoela, D., Abu-Mahfouz, A., & Hancke, G. (2018). A survey of anomaly detection in industrial wireless sensor networks with critical water system infrastructure as a case study. *Sensors (Switzerland)*, *18*(8), 1–24. https://doi.org/10.3390/s18082491

Russo, S., Lürig, M., Hao, W., Matthews, B., & Villez, K. (2020). Active Learning for Anomaly Detection in Environmental data. *Environmental Modelling & Software*, 104869. https://doi.org/10.1016/j.envsoft.2020.104869

Sakov, P., & Oke, P. R. (2008). A deterministic formulation of the ensemble Kalman filter: An alternative to ensemble square root filters. *Tellus, Series A: Dynamic Meteorology and Oceanography*, *60 A*(2), 361–371. https://doi.org/10.1111/j.1600-0870.2007.00299.x

Schmidhuber, J. (2015). Deep Learning in neural networks: An overview. In *Neural Networks* (Vol. 61, pp. 85–117). Elsevier Ltd. https://doi.org/10.1016/j.neunet.2014.09.003

Slay, J., & Miller, M. (2007). Lessons learned from the maroochy water breach. In *Critical Infrastructure Protection* (pp. 73–82). Springer.

Sufi Karimi, H., Natarajan, B., Ramsey, C. L., Henson, J., Tedder, J. L., & Kemper, E. (2019). Comparison of learning-based wastewater flow prediction methodologies for smart sewer management. *Journal of Hydrology*, *577*(April). https://doi.org/10.1016/j.jhydrol.2019.123977

Sun, C., Puig, V., & Cembrano, G. (2020). Real-Time Control of Urban Water Cycle under Cyber-Physical Systems Framework. *Water*, *12*(2), 406. https://doi.org/10.3390/w12020406

Talei, A., Chua, L.H.C., Quek, C., Jansson, P.E. (2013). Runoff forecasting using a Takagi-Sugeno neuro-fuzzy model with online learning. *Journal of Hydrology* 488, 17–32. https://doi.org/10.1016/j.jhydrol.2013.02.022

Vanrolleghem, P. A., Benedetti, L., & Meirlaen, J. (2005). Modelling and real-time control of the integrated urban wastewater system. *Environmental Modelling & Software*, *20*(4), 427–442. https://doi.org/10.1016/j.envsoft.2004.02.004

Zhang, D., Lindholm, G., & Ratnaweera, H. (2018). Use long short-term memory to enhance Internet of Things for combined sewer overflow monitoring. *Journal of Hydrology*, *556*, 409–418. https://doi.org/10.1016/j.jhydrol.2017.11.018

9 Papers

I **Palmitessa, R.**, Mikkelsen, P.S., Law, A.W.K., Borup, M. (2020). Data assimilation in hydrodynamic models for system-wide soft sensing and sensor validation for urban drainage tunnels. *Journal of Hydroinformatics*. (In press) https://doi.org/10.2166/hydro.2020.074

II **Palmitessa, R.**, Mikkelsen, P.S., Borup, M., Law, A.W.K. (2021). Soft sensing of water depth in combined sewers using LSTM neural networks with missing observations. *Journal of Hydro-environment Research*. (In press) https://doi.org/10.1016/j.jher.2021.01.006

III **Palmitessa, R.**, Pedersen, A.N., Borup, M., Sørensen, L., Law, A.W.K., Clemmensen, L.K.H., Mikkelsen, P.S. (2021). Anomaly detection in water depth observations from combined sewers using LSTM neural networks. (Manuscript)

In this online version of the thesis, **papers I-III** are not included but can be obtained from electronic article databases e.g. via www.orbit.dtu.dk or on request from:

DTU Environment
Technical University of Denmark
Bygningstorvet, Building 115
2800 Kgs. Lyngby
Denmark

info@env.dtu.dk

The Department of Environmental Engineering (DTU Environment) conducts science-based engineering research within three sections: Circularity & Environmental Impact, Climate & Monitoring, Water Technology & Processes.

The department dates back to 1865, when Ludvig August Colding, the founder of the department, gave the first lecture on sanitary engineering as response to the cholera epidemics in Copenhagen in the late 1800s.

Department of Environmental Engineering

Bygningstorvet, Building 115
2800 Kgs. Lyngby
Tlf. +45 4525 1600
Fax +45 4593 2850

www.env.dtu.dk