

Understanding the Urban Atmosphere through Conceptual Modelling and Opportunistic Sensing

Arjan Droste

Understanding the Urban Atmosphere through Conceptual Modelling and Opportunistic Sensing

Arjan Droste



Propositions

- 1) Opportunistic sensing has unprecedented potential for urban meteorological observations.
(this thesis)
- 2) Urban wind is as important for human health as urban heat.
(this thesis)
- 3) To encourage interdisciplinary collaboration and understanding, scientific articles would benefit from a more poetic writing style.
- 4) The scientific system rewards ambition and productivity, not necessarily progress
- 5) Meteorology is the mother of all environmental sciences
- 6) Scientific results cannot be objective if the underlying statistics are not provided
- 7) Involvement with people with mental disabilities is an enriching experience, which can bring to light the true important things in life
- 8) A sense of humour is a helpful quality for any scientist

Propositions belonging to the thesis entitled

Understanding the urban atmosphere through conceptual modelling and opportunistic sensing

Arjan Marten Droste

Wageningen, 4 November 2020

Understanding the urban atmosphere through conceptual modelling and opportunistic sensing

Arjan Marten Droste

Thesis committee

Promotor:

Prof. dr A.A.M. Holtslag
Professor of Meteorology
Wageningen University & Research

Co-promotor:

Dr G.J. Steeneveld
Associate Professor, Meteorology and Air Quality
Wageningen University & Research

Other members:

Prof. dr D. Tuia, Wageningen University & Research
Prof. dr J. Barlow, University of Reading
Prof. dr D. Scherer, Technische Universität Berlin
Prof. dr B. Van de Wiel, Delft University of Technology

This research was conducted under the auspices of the Graduate School for Socio-Economic and Natural Sciences of the Environment (SENSE)

Understanding the urban atmosphere through conceptual modelling and opportunistic sensing

Arjan Marten Droste

Thesis

submitted in fulfilment of the requirements for the degree of doctor
at Wageningen University
by the authority of the Rector Magnificus
Prof. dr A.P.J. Mol,
in the presence of the
Thesis Committee appointed by the Academic Board
to be defended in public
on Wednesday 4 November 2020
at 1:30 p.m. in the Aula.

Arjan M. Droste

Understanding the urban atmosphere through conceptual modelling and opportunistic sensing
xii + 161 pages.

PhD thesis, Wageningen University, Wageningen, the Netherlands (2020)
With references, with summary in English

ISBN 978-94-6395-533-1

DOI <https://doi.org/10.18174/530485>

Just as we yoke our own destinies to the seasons,
to give our lives the impression of circularity and renewal,
rather than just an uninterrupted march towards death,
so we look to weather for meaning and metaphor,
to avoid the recognition of life's calamitous randomness.

Zoe Williams

Summary

This thesis explores the urban climate, with a focus on urban wind speed, from two distinct approaches. The first approach is the use of crowdsourcing, or opportunistic sensing techniques, to collect and process vast quantities of urban meteorological data. The second approach is to use (conceptual) physical modelling to gain a better understanding of wind differences between city and countryside.

Chapter 2 expands upon a previously developed technique to derive urban air temperatures from smartphone battery temperatures, studying finer spatial and temporal scales than ever before. Over 10 million smartphone temperature records for the city of São Paulo, Brazil, are combined to derive daily and even hourly averaged urban air temperatures. Optimal results are achieved for 700 or more retrievals aggregated into daily or hourly temperature values. Daily temperature estimates are good (coefficient of determination of 86 %), and temperature differences between Local Climate Zones (LCZs) can be distinguished at this scale. Hourly estimations of air temperature require a correction through a diurnally varying parameter in the used heat transfer model. The results show the value of smartphones as a measuring platform when routine observations are lacking.

Chapter 3 makes use of Personal Weather Station (PWS) data to investigate their use for urban wind research. While the potential of PWS data for rain and temperature has been established, PWS wind data remained unused because of its high risk of error. We use crowdsourced wind speed observations from 60 Netatmo-brand PWSs in Amsterdam, the Netherlands, to analyse wind speed distributions of different LCZs. In a field test against a reference station, the Netatmo PWS anemometer appears to systematically underestimate the wind speed, and episodes with rain or a high relative humidity deteriorate the measurement quality. Therefore, we developed a quality assurance (QA) protocol to correct PWS measurements for these errors. The applied QA protocol strongly improves PWS data to a point where they can be used to infer the probability density distribution of wind speed of a city or a neighbourhood. The PWS performs poorly for periods with very low wind-speed; however, results for a year-long climatology of wind speed are satisfactory.

In **Chapter 4** the aforementioned crowdsourcing techniques are combined, along with Commercial Microwave Link (CML) data, to study meteorological phenomena. During a 17-day summer period in Amsterdam the passage of a cold front, followed by a warm episode, is monitored using the opportunistic sensing techniques. Measurements of temperature (from smartphone and PWSs); wind speed (PWSs); precipitation (PWSs and CML); air pressure (smartphone and PWS); radiation (smartphones) and humidity

(PWSs) are compared to an established reference network. While the opportunistic sensing data have large uncertainties, both the cold front and the Urban Heat Island (UHI) during the hot period are successfully captured by these innovative techniques, showcasing their combined potential for urban (hydro-)meteorological monitoring.

Chapter 5 introduces the Urban Wind Island (UWI): a positive wind difference of around 0.5 m/s between city and countryside. We research urban–rural wind differences with mixed-layer model, a bulk representation of the urban and rural boundary layers. In the model, these two distinctly surfaces are not connected, but experience the same large-scale atmospheric influences. The UWI appears to be caused by a combination of differences in boundary-layer dynamics between city and countryside. Oscillation of the wind around the geostrophic equilibrium can cause these UWI episodes. Sensitivity of the UWI to urban morphology is researched by implementing the 10 urban LCZs. The ideal circumstances for UWI formation seem to be moderate wind speeds (around 5 m/s), low building heights (up to 12 metres) and a deeper initial urban boundary layer.

The UWI exploration is expanded upon in **Chapter 6**, which uses the mesoscale model WRF at 500m resolution to investigate the characteristics of the UWI over 2 consecutive summers in Amsterdam, in a more realistic model setting than the previous chapter. Large scale influences that could influence the wind field, such as frontal passages and strong precipitation events, are filtered out to focus on surface-induced wind changes. In this more realistic modelling environment, the data show that the UWI is present roughly half of the time, with a very similar magnitude as previously found (in the order of 0.5 m/s). The formation of the daytime UWI proceeds very similar to that of the theoretical concept, but a nocturnal UWI also forms, as a result of a delayed transitions towards a stable nocturnal urban boundary layer. While the relation between UWI and the initial conditions of the boundary layer (depth, geostrophic wind speed and temperature) are less clear, the UWI is consistently present in the data with a similar magnitude and timing as the conceptual model.

The main findings of this thesis are that opportunistic sensing and crowdsourced data form a valuable addition to urban meteorological measurements, though care needs to be taken to account for the large errors present in these data and associated techniques. Secondly, that average wind in urban areas can be faster than over the countryside under certain conditions: this UWI can form both during the day, through different boundary-layer growth dynamics; and during night-time, related to the delayed collapse of the convective boundary layer. Both unconventional observation techniques as well as conceptual modelling can contribute to understanding the urban atmosphere.

Voorwoord

Een promotie in tijden van corona. Het voelt raar, dit stuk te typen vandaag. Het is 13 Mei: de oorspronkelijke datum van mijn promotie. Maar ik sta niet stijf van de spanning in de Aula, gekleed in mij oneigen nette kledij. In plaats daarvan zit ik met sloffen aan, in pyjama achter m'n bureautje deze eerste paar zinnen te typen; mijn promotie uitgesteld tot een onbepaalde tijd. Waarom? Had ik wat harder door moeten werken? Nee, dit boek is al lang en breed goedgekeurd, ik heb mezelf het zweet in gewerkt om alles op tijd af te krijgen. De wereld is in de greep van COVID-19, het coronavirus, dat ons allemaal binnenhoudt en de wereld gegijzeld heeft. Maar ondertussen gaat de wereld door: de universiteit is open, al het onderwijs massaal online gegoooid. Samen met wat collega's assisteer ik om een veldwerk vak volledig online te geven: een enorme inspanning voor student en docent, maar het gebeurt wel. We leren onszelf in rap tempo nieuwe skills aan: het maken en bewerken van instructiefilmpjes, omgaan met virtuele leeromgevingen, de beste houding om in te loungen tijdens eindeloze videovergaderingen...

Die blijken we de maanden erna ook nodig te hebben: het virus gaat nergens heen, en we proberen ons dagelijks leven weer op te pakken, terwijl we ook de veiligheid in acht (trachten te) nemen. Het academisch jaar is inmiddels weer van start gegaan, en we mogen mondjesmaat weer de campus op: de tijden van elke dag naar de WUR fietsen lijken eeuwen geleden. We worden het thuiswerken wel zat, de sleur van je eigen huis, ik begon zelfs de dubieuze koffie op kantoor te missen. Vreemd genoeg blijkt eindeloos thuiszitten best stressvol te zijn. Maar tegelijkertijd leren we er ook wel wat uit: we zijn enorm flexibel, maar ook ontzettend kwetsbaar als ons dagelijkse ritme zo wild verstoord wordt. Mocht, over hoe lang dan ook, deze situatie weer voorbij zijn, is het zaak om terug te kijken naar hoe we als samenleving ons hier doorheen hebben geslagen, en niet deze periode maar gewoon te vergeten als een nare herinnering.

U, lezer, heeft dit boekje nu in de hand. Mijn nieuwe promotiedatum is inmiddels al lang bekend en nadert rap. Hopelijk treffen we elkaar in de Aula, en anders via een scherm. Hopelijk heeft u plezier aan het doornemen van dit werk, gemaakt in een andere wereld, toen we nog bij elkaar op kantoor mochten zitten en elkaar liederlijk de hand konden schudden.

Stel je eens voor.

Contents

	Page
Summary	vii
Voorwoord	ix
Chapter 1 Introduction	1
Chapter 2 Crowdsourcing urban air temperatures through smartphone battery temperatures in São Paulo, Brazil	15
Chapter 3 Crowdsourcing for urban wind research	37
Chapter 4 Hydrometeorological monitoring using opportunistic sensing networks in the Amsterdam metropolitan area	59
Chapter 5 Introducing the Urban Wind Island Effect	81
Chapter 6 The urban wind island from a three-dimensional perspective	103
Chapter 7 Synthesis	123
References	135
Authorship Statement	153
Acknowledgements	155
List of Publications	159
Graduate School Certificate	161

The image features a large, white, serif capital letter 'I' centered on a background of abstract watercolor splatters. The splatters are in shades of red, orange, and brown, with some darker, almost black, spots. The overall effect is a high-contrast, artistic composition. The letter 'I' is simple and bold, standing out prominently against the textured, organic background.

I

Chapter I

Introduction

1.1 The roots of meteorology

Humanity has always been fascinated by the weather, and trying to make sense of the influence the heavens have on our every-day lives has been one of the oldest endeavours of our kind. Indeed, the number one hobby in the Netherlands still has to be talking (or complaining) about the weather. Over the centuries, what we now call “meteorology” – based on the ancient Greek μετέωρος: “high in the sky” – has been practised by countless people for a myriad of purposes, and their views and visions determined how they went about interpreting the weather. The current state of meteorology, or atmospheric sciences to encompass the whole field, is highly advanced, and new research is published on a daily basis to further develop our understanding of the atmosphere and its interactions with the solid and liquid Earth.

This thesis is but another addition to the ever growing compendium of human knowledge, building greatly on those who have observed, described and theorised the beauty and complexity of the Earth. Making use of the work of one’s predecessors is one of the primary principles of science: it is not for nothing that the phrase “[...] *nos esse quasi nanos gigantium humeris incidentes*” (we are like dwarves sitting on the shoulders of giants) has been attributed to a number of eminent scientists, chief of which Sir Isaac Newton. In this light it makes sense to dwell upon the historical roots of meteorology, before I provide you, the reader, with yet another of these scientific works, from my seat on the giants’ shoulders.

In its most basic form, the relation between weather and human wellbeing is in agriculture: knowing about periods of rain or drought or the optimal time for sowing are crucial for a bountiful harvest. It is within this frame of mind that the first meteorology was practised, to predict rain from annual cycles. In the ancient world, there was therefore no real distinction between astrology and meteorology: the heavens move with the seasons, and the skies change correspondingly. Babylonian tablets contain thousands of omens related to the sky and the heavens: there were interpreted by advisers to the Assyrian kings for matters political and military. During the time of the Greek and Roman civilisations the interpretation shifted: the celestial sphere had a direct influence on the atmosphere, and the relation between the stars and meteorological phenomena are detailed in *paraepgmata*: essentially farmers’ almanacs or lunar calendars, linking celestial movements to meteorological phenomena (Taub, 2003).

One of the earliest examples of this practical information, agricultural and commercial advice based on the seasons, is Hesiod’s *Work and Days*, which among others also describes how to identify safe days for sailing, linking the importance of meteorology to naval travel. This work is not only purely practical: its encompassing theme is how to improve one’s life in a difficult world, by working to escape the inherent evil in the world, and studying the heavens in this sense also serves an ethical purpose (Bowker, 2011). So not only is meteorology practical, it can improve one’s character!

With regards to that lofty goal, it is perhaps not strange to find meteorological advice



not just in the form of purely factual writings, but rather in the form of epic and didactic poetry. In Homer's epic poetry, the *Iliad* and *Odyssey*, weather plays a prominent role to thwart and aid its heroes. The meteorological phenomena are then works of the gods: Zeus is the cloud gatherer, hurling lightning bolts to his enemies; Athena is mentioned to create fog, and Poseidon shakes the Earth and creates tidal waves. Aeolus, keeper of the winds, gives Odysseus a bag containing the four major winds to safely guide him home to Ithaca, but man's inevitable greed makes his men open the bag in search for riches, generating a devastating storm blowing his ship to the far reaches of the Mediterranean... (Taub, 2003)

This link between the divine and the natural was at some point severed: the Roman poet Lucretius used rational explanations of meteorology to prove that the basis of the universe is natural rather than divine, to banish the fear of the divine in mortals. Similarly, Seneca in his *Natural Questions* discusses meteorological work by his predecessors, with the aim to uncover the secrets of the universe and pass on this knowledge. However, the old epic poets were still regaled as sources of useful information, and poetry was an important medium to disseminate information to a wider public. For instance the *Phaenomena* by Aratus of Soli, a didactic poem that starts with an ode to Zeus for his creation of the order of the heavens which allow for the reading of signs to be interpreted into a meteorological phenomenon, was very popular, if not always accurate by contemporary standards. The popularity of this work ensured its survival up through the medieval period, where it was picked up by the Church and adopted into a Catholic practice of weather forecasting (Lawrence-Mathers, 2019)

While most treatises do not endeavour to explain the causes of meteorological phenomena, Aristotle's *Meteorologica*, where we get the name of our discipline from, provides a theoretical framework for the natural world. The world consists of four elements in 4 states (dry-moist; hot-cold), and the elements changing states govern the physical processes on Earth. An important practice of his is the review of his predecessor's work, reviewing their theories and refuting or complementing them with his own to develop improved theories: demonstrating the value of the scientific community.

Another remark of his, and one that is shared by many classical authors writing on the subject of the natural world, is the admission that not everything can be sufficiently explained (*Meteorologica* 339a: 2-3):

We consider that we have given a sufficiently rational explanation of things inaccessible to observation by our senses if we have produced a theory that is possible: and the following seems, on the evidence available, to be the explanation of the phenomena now under consideration.

Pliny the Elder, famous for his *Natural History* which remained relevant for centuries after his death in 79 BCE, notes the difficulty with observations being made at different locations, and that these cannot always be translated and generalised to a different area. He shows a particular interest in winds, distinguishing between steady winds that blow with a constant speed and direction, and erratic gusts of wind: a first mention of turbu-

lence! The Greeks already made use of a wind-rose, with 12 cardinal directions, similar to the current 16 directions. The Greek author Theophrastus wrote a treaty *On Winds*, where he regards wind as the movement of air, to restore balance upset by the sun's heat (Taub, 2003). The sun and the earth generate winds, and since the moon acts as a weaker version of the sun, so are winds stronger at night – the first mention of the low-level jet?

The works of Aristotle and Pliny the Elder would remain relevant for centuries, with their practical information for agriculture and their framework to explain the phenomena of the atmosphere. As late as the 18th century, Pliny was used as an authority by English agricultural writers (Lawrence-Mathers, 2019). And even now in our current age, atmospheric science still shares some of its characteristics already present millennia ago. After all, Pliny already bemoans the lack of interest in his contemporaries to study the winds, lamenting that they only study to get rich quick and not for the knowledge itself... Surprisingly sharp still to this day!

Ultimately, meteorology is built on tradition: first through incorporating practical traditions on predicting the weather through signs and omens, then through a traditional framework of the Aristotelean elements explaining the natural world. Now, in the modern era, the WMO guidelines are our 'tradition' for correct measurements, and our numerical models based on parametrisations, concepts and techniques designed by our predecessors. But even these traditions need to be challenged in our changing world: observations are concentrated in rural areas free of human influence, and operational models typically do not differentiate the urban land-use (at least not on the larger scales). Nowadays the centre of our activity is the heterogeneous urban area, for which rural measurements are hardly of any use, and models require special modules to deal with the strong local character of turbulence and the altered energy balance. The next section will deal with this "urban problem" that needs to be tackled for meteorology to advance.

1.2 The Urban Problem

Advancing from the Classical world to current day, it is obvious that our living environment is vastly different. Human population has exploded in those 2,500 years, with currently over 7 billion people alive (Max Roser and Ortiz-Ospina, 2020). In this time humanity has created its own unique living environment: the city. Classical Athens may have been an impressive metropolis with 300,000 inhabitants; it pales in comparison to current-day Athens, steadily growing with over 3 million people living in its extensive urban area. Such overwhelming human activity has distinct effects on its immediate environment: pollution springs to mind, but also the atmosphere feels the effects of concentrated human presence. Modern urban areas struggle with air pollution and heat (Pascal et al., 2013), all exacerbated by the effects of buildings and anthropogenic influence on the atmosphere: an urban problem indeed. This urban problem extends to the very nature of doing measurements: can meteorological observations in cities be representative?

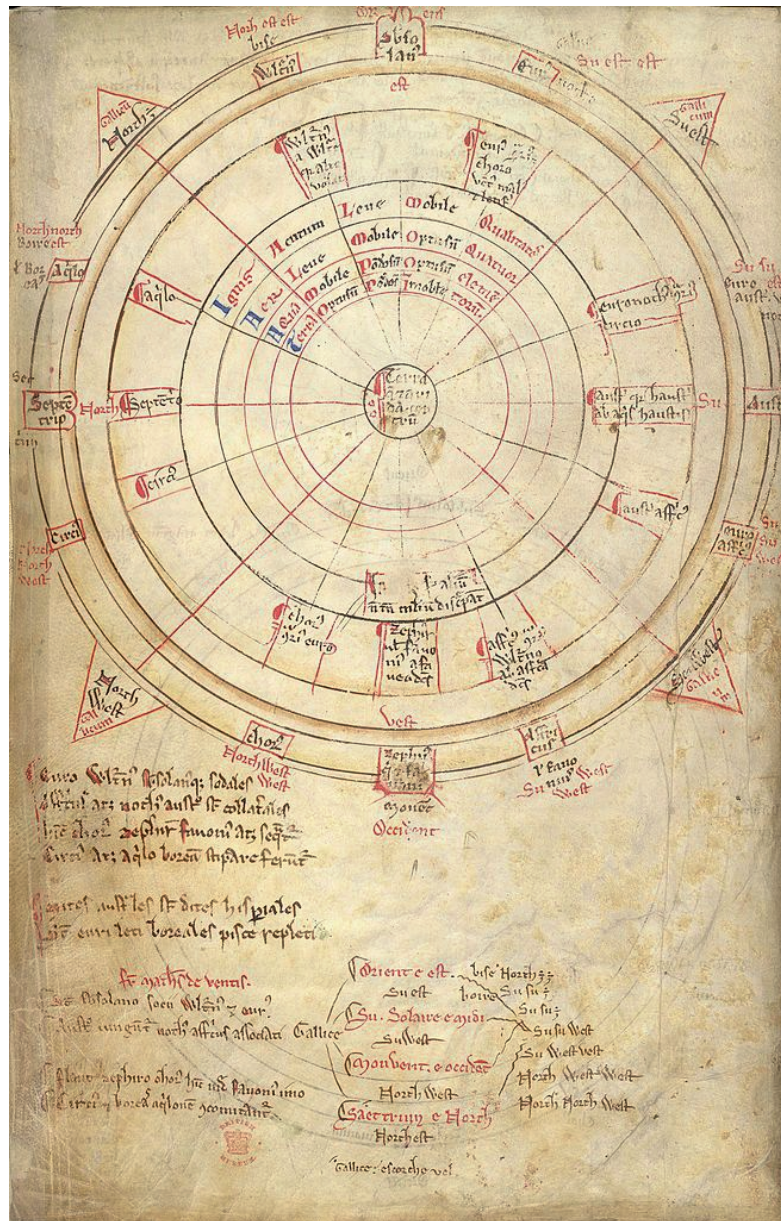


Figure 1.1: Sketch of the classical (12-point) wind-rose imposed on the modern 16-point wind-rise. Sketch from *Liber Additamentorum*, c. 1250, Matthew Paris, manuscript held by the British Library, London, UK.

We have already seen that Pliny the Elder noted that it is difficult to translate observations made in one place to those at another if their geography is different. To this end, we now have the guidelines of the WMO (WMO, 2018), with a complete list of what a proper meteorological observational site must adhere to. In this way, when we compare data of two sites with each other, we know that the observations are taking place over relatively smooth rural (grass)land with no obstructions close to the measurement site. It is based on long-term and field campaign data over these types of land that modern meteorological theories are based, e.g. the Monin-Obukhov similarity based on the Kansas dataset (Businger et al., 1971). There is a lot of study regarding the effect of heterogeneity in (rural) areas on micrometeorology (e.g. Mahrt, 1996), but these are all as exceptions to the general case of the homogeneous grassland.

In reality, the world is not as nice and smooth as our theories would have it: heterogeneity is the norm. It is in strongly heterogeneous terrain that many conventional theories do not hold: over mountains, over areas with rapidly varying land-use, and of course over cities, which can wildly vary in structure and building density over the course of a single neighbourhood (Barlow et al., 2017). This makes cities a unique landscape, to the point where the boundary layer over a city is distinctly different from that over the rural terrain: the Urban Boundary Layer (UBL) is a distinct entity from the Rural Boundary Layer.

Taking a step back to the world of Classical meteorology as described in the previous section, it has been a long-known fact that the Sun is a primal influence on the Earth. The Sun's energy is the main driver of many Earth-system processes, in particular through the surface energy balance. The available short-wave radiation reaching the Earth's surface can be used to heat up the air and develop the boundary layer: the sensible heat flux. Energy can also be allotted to the evaporation of water: the latent heat flux. Any energy that is not used for these purposes is stored in the ground: the ground heat flux. Since roughly the same solar energy would reach an urban and a rural area, how can their atmosphere be so different? The key to the differences in boundary-layer development lies within the distribution of energy (Figure 1.2).

Rural areas are typically highly vegetated, so most of the energy is partitioned towards the latent heat flux, leaving less to heat up the air or store in the ground, since heat storage of natural areas tends to be low. As a result, the boundary layer heats up quickly in the morning, but reaches a relatively lower depth. The urban area contains large amounts of impervious material with high heat storage capacity: roads, buildings, concrete; there is little vegetated area. As such, the latent heat flux is low, and the ground heat flux is an important component, storing heat during the day to release it later during the night. The ground heat flux in cities is such a typical phenomenon for the urban area that it's referred to as the storage heat flux in urban meteorology. The UBL development is typically slower in the morning, as most energy is absorbed into the storage heat flux, but since there is nearly no vegetation to evapotranspire water, the UBL grows deeper due to the large sensible heat flux.

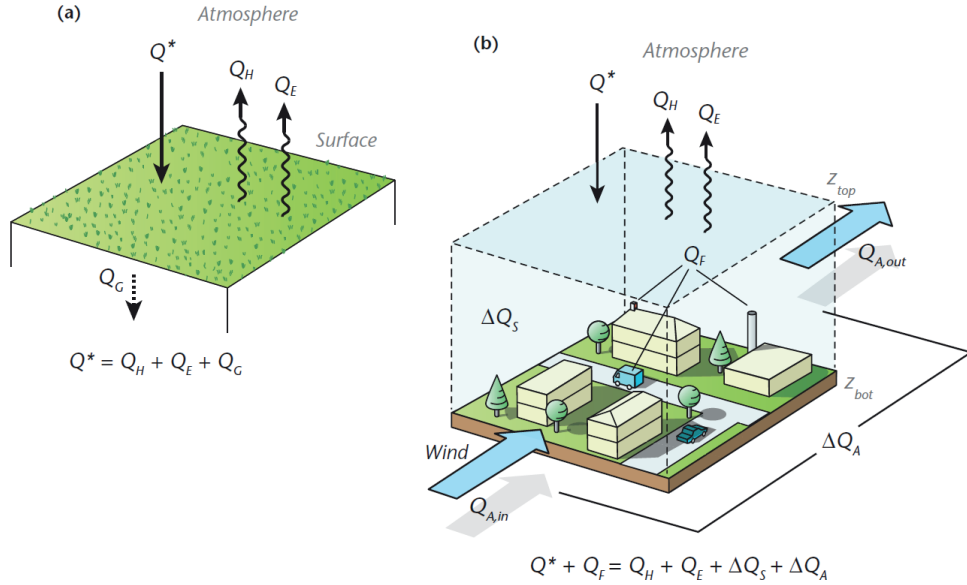


Figure 1.2: Schematic overview of surface energy balance for a rural (a) and an urban (b) surface. Q_H represents the sensible heat flux; Q_E the latent heat flux; Q^* the net radiation; Q_G the (rural) ground heat flux; Q_F the anthropogenic heat flux; ΔQ_A the heat flux through advection; ΔQ_S the (urban) storage heat flux. Source: Figure 6.2 from Oke et al. (2017), Chapter 6.

At night, when there is no solar energy input, the rural boundary layer collapses and becomes stable, but the UBL still receives energy from below: the stored heat from the day is released by the building materials to heat up the air. The nocturnal UBL stays neutral or weakly unstable for several hours during the night (Barlow et al., 2015) and several degrees warmer than its rural counterpart: this is when the Urban Heat Island, the temperature contrast between city and countryside, is the strongest. An additional contribution to the energy balance is the Anthropogenic Heat Flux: energy released through human activity such as car engines, industry or air conditioning. This energy flux is dependent on the population density of a city and highly variable between cities, but can be a strong contributor to the overall energy budget. Typical values are daily averages of 40 W/m^2 but can be as high as 400 W/m^2 in dense cities such as Tokyo (Ichinose et al., 1999).

The structure of the UBL is also different from the rural boundary layer (see Figure 1.3). The roughness elements of the city (the buildings, trees and other obstacles) create a roughness sub-layer above the building height (typically 2–5 times building height, Barlow (2014)); the air within the street canyons and between the buildings is the Urban Canopy Layer, which is directly influenced by the transfer of energy between the ground, road, roof and wall surfaces. Turbulent flows dominate the UBL up to the roughness sublayer, above which the UBL behaves more like the classical (rural) mixed layer. The




Figure 1.3: Schematic overview of the urban boundary layer during daytime. Source: Figure 2.13a from Oke et al. (2017), Chapter 2.

dominance of turbulence within the lower layers of the UBL makes doing representative measurements notoriously difficult: anything measured in one street canyon is not representative for the street adjacent to it, so to speak. In addition, it is expensive to maintain an urban measurement site, and regulations can make it difficult to do certain obnoxious forms of measurements such as SODAR, not to mention the increased risk of vandalism or construction works blocking or removing equipment altogether. Modelling urban meteorological processes is equally laborious: conventional similarity theory does not work in the lower layers of the UBL, so measurements of wind, temperature and moisture can not be extrapolated towards vertical profiles above rooftop height (Kent et al., 2018).

There have been some solutions to tackle this urban problem. In terms of measurements, while long-term measurements are often not feasible, there have been many measurement campaigns with either fixed (Rotach et al., 2005; Barlow et al., 2011; Bohnenstengel et al., 2015) or mobile (Heusinkveld et al., 2014; Shi et al., 2018) intensive measurements for shorter periods of time. Fixed measurements can gain insight into physical processes at the cost of representativeness; mobile measurements are limited in what they can sample (e.g. no high-resolution fluxes) but can cover the wide range of urban landscapes, from the packed city centre to the more spacious outskirts.

When modelling the urban environment, or incorporating it into a larger-scale model for weather forecasting, the scale of the model decides how the urban area is implemented. A common implementation in NWP models is the Urban Canopy Model (UCM, Kusaka



et al., 2001), which is a schematic representation of the urban canopy, modelled as an infinite street with buildings on either side. Building height, road width and street orientation can be altered, as well as the amount of vegetated area based on the building density. The scheme translates modelled values of heat, wind etc. down from the first model level into an urban canyon value when coupled to the actual atmospheric model (Tsiringakis et al., 2019). Since the UCM is only one layer, it circumvents the erratic vertical behaviour of turbulent quantities in the urban canyon. Other more sophisticated schemes are multi-layered (Salamanca et al., 2010), or take building energy interactions with the UBL into account (Martilli et al., 2002). If the model scale is finer, e.g. for Large-Eddy Simulation (LES) models, many of the parameterised processes can be solved explicitly, such as flow around walls and obstacles (Maronga et al., 2019). This comes at the cost of requiring high computational power, and therefore LES models tend to take limited model domains into account, or just the flow around single buildings (Toparlar et al., 2017).

In summary, the strong heterogeneity of cities is the core issue which scientists must take into account when modelling or measuring the urban atmosphere. Making simplifications or capture the larger effect of the city as a whole or limiting oneself to a single measuring location or modelled street are ways to deal with this urban problem.

1.3 Innovative solutions

At the start of this Introduction, we have noted that meteorology is an old science, and was already rooted in tradition in Antiquity, taking care to incorporate practical experience and frameworks from predecessors into new, refined theories and practice. Current practice focuses on increasing computational power to construct finer model grids, to explicitly resolve processes that are difficult to parameterise, such as convection or cumulus cloud formation (Gentine et al., 2018). On the observational side, new instruments are capable of measuring fluxes at high-frequencies, and development focuses on high-precision observations, or difficult-to-measure quantities such as isotopes (Werner et al., 2012). But is this focus on higher quality observations and models the only way forward? After all, the Greeks already knew that some meteorological phenomena are hard to study simply because the quantity of observations is too low, which is exactly the issue urban meteorology suffers from. As explained before, there have been some thorough urban field studies and some ongoing networks, but none of these are permanent, and when looking for data on a specific city there often is none. Similarly, model studies with either Numerical Weather Prediction (NWP) models or LES models are not always the most suitable for getting a grasp on the underlying mechanics and physics of the studied process, regardless of how accurate they are. Sensitivity studies are time consuming with these sophisticated models, so a more simplified model based on the core physics of the atmosphere can provide valuable insight can could be overlooked by the myriad of interacting processes in a high-resolution model.

In this thesis I make use of some more unconventional techniques to study the urban

atmosphere, both with observations and with models, focusing on wind. On the observational side that is crowdsourcing, or opportunistic sensing: making use of a large amount of low-quality data, often generated by devices not intended for (scientific) measuring purposes. On the modelling side, a conceptual slab-model of the boundary layer is used to gain insight in core physical differences between a rural and an urban area.

1.3.1 Crowdsourcing and Opportunistic Sensing

The word crowdsourcing, a portmanteau of ‘crowd’ and ‘outsourcing’, is relatively young: it was first coined as a phrase in 2006, as a way to describe how companies outsource their work to the Internet. In broader terms, it can be defined as “*the practice to obtain information from a large group of people, typically over the Internet*” (Oxford English Dictionary). The practice itself goes much further back: a nice meteorological example would be the United States’ National Weather Service, which started from an 1849 project to make weather maps based on volunteers’ data sent in by telegraph (NWS, 2013). And Aristotle himself already noted that a crowd, even made up of “lesser men”, knows more than its individuals. Nowadays, crowdsourcing in meteorology can be used to obtain weather data from weather stations owned by citizens: so-called Private Weather Stations (PWSs).

Several brands offer affordable weather stations that can be placed in a garden, or on a balcony or roof. Most of these stations have the capability to connect to an online platform to upload their data to form a live weather map. Examples of these platforms are the Weather Underground (<https://www.wunderground.com/wundermap>), which allows data from any brand PWS; or the Netatmo Weather map (<https://weathermap.netatmo.com>) which shows data from their own Netatmo brand PWS. There is a wide variation in measurement capability, going from simple thermometers to rain intensity measurements, as well as in quality, with the cheaper stations typically having lower measurement accuracy or resolution (Bell et al., 2013, 2015). While individual stations can be used to e.g. research many different areas at once (e.g. Steeneveld et al., 2011), the power of crowdsourcing lies in their number. Since individual quality tends to be poor, depending on the type of station, many observations pooled together can provide a good estimate of the actual signal. Examples are Fenner et al. (2017); De Vos et al. (2017) and Chapters 3 and 4 of this thesis. However, Quality Control is also necessary to filter out erroneous stations or measurements, since the PWSs are usually installed by hobbyists without knowledge of correct sensor setup.

In some cases, data such as described in the above section is obtained not from a large number of volunteers, but from a third party or a small group, yet still used in the same manner. In that sense, the broader term “opportunistic sensing” is also used to describe data originating from devices that were not meant to measure a (meteorological) variable but can be used for this purpose. Examples include the use of Commercial Microwave Links to estimate rainfall (Overeem et al., 2013a; Uijlenhoet et al., 2018; de Vos et al., 2018) or smartphone observations of air pressure (Mass and Madaus, 2014; Madaus and Mass, 2017; Hintz et al., 2019) and temperature (Overeem et al. (2013b); Chau (2019) and Chapter 2 of this thesis). These data tend to require more preprocessing and quality



control, and the resulting product is of lesser quality than a PWS observation would be, but their spatial coverage can be much wider, especially in the case of smartphone-derived observations. With the ubiquity of smartphones and other smart devices, which obviously have the highest density in cities, their usage for urban meteorology is a logical step, and could be used for local forecasts or heat stress maps, for instance.

1.3.2 Conceptual Modelling

A conceptual model only takes the core elements of a physical problem into account: enough to describe its basic state. Unlike crowdsourcing or opportunistic sensing, the use of conceptual models in meteorology is nothing new. Conceptual meteorological models include those to study cloud formation in satellite data (EUMETSAT, 2009), and the global circulation of wind that forms the basis of many synoptic meteorological theories is in itself a conceptual model.

The conceptual mixed-layer bulk model of the boundary layer, describing it as a single slab governed by exchange with the surface below and the free troposphere above, has seen much usage in the past (e.g. Driedonks and Tennekes, 1984; Byun and Arya, 1986) as well as in current PBL research (e.g. Schröter et al., 2013; Wouters et al., 2019). This model relies on a set of equations describing the rate of change of state variables of the boundary layer (wind, temperature, humidity) depending on surface and entrainment fluxes (these equations are given in Chapter 5). Its main usage in the past has been for the rural boundary layer, researching cloud formation, vegetation interaction or wind shear but fewer applications for the UBL (Byun and Arya (1990); Theeuwes et al. (2015), and Chapter 5 of this thesis). Most urban model studies are specifically interested in the Urban Canopy Layer and its vertical interactions with the air aloft and the surface (e.g. Tsiringakis et al., 2019), rather than the UBL as a whole. The bulk model is however very suitable to research how the different urban surface influences the entire boundary layer on a larger scale: the effect of the entire city, as it were. Microscale effects are then ignored in favour of the broader influence the city itself has, such as the Urban Heat Island through the high heat capacity, or the effect of the increased roughness and turbulence of the surface.

Findings made in such conceptual models can afterwards be translated to a more “real-world” setting by incorporating it into a NWP or LES model, connecting the conceptual realm to the reality of forecasters and operational meteorologists.

1.4 Scope of this thesis

This Introduction connected the early history of meteorology to current practice, where we noticed that the split between modelling or theory, and practice or observations is nothing new. With a discipline as old and steeped in tradition as meteorology, innovative new approaches to tackle old problems can be met with resistance. The overarching theme of this broad thesis is to introduce the usage of opportunistic sensing and conceptual mod-

elling to investigate the urban atmosphere, and to show that these new techniques can be complementary to the established approach, to improve our knowledge.

This thesis is roughly divided in two themes: the first 3 chapters centre around crowdsourcing and opportunistic sensing, whereas the subsequent 2 chapters deal with conceptual and 3D modelling. The two themes are connected through the focus on wind: a highly turbulent variable that's difficult to capture in both modelling and observations, and therefore fairly under-studied in the urban domain, where urban heat tends to be the main focus point.

Urban heat is the focus point of chapter 2: the usage of crowdsourced smartphone battery temperatures to derive urban air temperature. The area of interest for this chapter is the Brazilian mega-city São Paulo. The city has two official measurements sites, in the form of an Urban Fluxnet tower <https://ibis.geog.ubc.ca/urbanflux/> and the WMO airport station surrounded by urban area thanks to the enormous urbanisation. In contrast, nearly 10 million smartphone observations are available for the study period of 2 years, spanning the entire city. This shows the strength of crowdsourced data especially in developing countries or megacities. These data are used to answer the following question: **to what extent are crowdsourced temperature data from smartphones useful to study urban air temperature at the hourly scale? (RQ 1)**

With chapter 3 starts this thesis' focus on wind: researching the quality of PWS wind speed observations. The area of focus is Amsterdam, the Netherlands, which has a dedicated meteorological observation network run by Wageningen University, and has been the focus of several urban meteorological campaigns. While rain and temperature have been successfully researched using PWS data, wind has remained unstudied. This chapter uses the combined presence of PWS and established network data to assess **whether PWS wind data can be applied for urban wind research, and how it can be quality controlled to improve the data (RQ 2).**

Chapter 4 synthesises the preceding two chapters, applying the developed techniques to study two meteorological events in June 2017, again in Amsterdam. PWS, smartphone and Commercial Microwave Link data are combined to study the passage of a cold front, followed by an extended hot period with a high UHI. The combination of several opportunistic sensing techniques uses their respective strengths while compensating for weaknesses. **The question is whether a combination of opportunistic sensing techniques can be used to study meteorological phenomena in sufficient detail. (RQ 3)**

Chapter 5 moves from opportunistic sensing to the realm of conceptual modelling, using the slab mixed-layer model to research differences in wind behaviour between the urban and rural boundary layer. The same large-scale forcings are imposed on different surfaces, for a range of initial conditions, **to understand how wind behaviour in urban and rural areas responds to similar large-scale influences (RQ4).** Under certain circumstances, this gives rise to periods in time where the urban wind is faster

than the rural wind: the Urban Wind Island effect (UWI).

Chapter 6 takes this concept of the UWI and researches it in a more realistic model setup, using the 3D NWP model WRF (Weather Research and Forecasting). Two summers of model data from the Summers in the City project are studied to investigate whether the UWI mechanisms found in the conceptual setting are still the main drivers in the 3D model, and whether there are different drivers when transitioning into the stable boundary layer at night. **Are the drivers behind the UWI in a three-dimensional model setting similar to the conceptual setup previously used? (RQ 5)**

Chapter 7 closes with a synthesis and discussion of the thesis, highlighting some of the main issues encountered with the used methods, and what the future holds for the studied phenomena and the used techniques.



The image features a large, white, serif-style number '2' centered on a red, textured background. The background has a mottled, painterly appearance with various shades of red and some darker, almost black, spots. The entire composition is set within a white border that has a subtle, faint grid pattern. The number '2' is the primary focus, rendered in a clean, classic font that contrasts sharply with the busy, organic texture of the red background.

2

Chapter 2

Crowdsourcing urban air temperatures through smartphone battery temperatures in São Paulo, Brazil

Crowdsourcing as a method to obtain and apply vast datasets is rapidly becoming prominent in meteorology, especially for urban areas where traditional measurements are scarce. Earlier studies showed that smartphone battery temperature readings allow for estimating the daily and city-wide air temperature via a straightforward heat transfer model. This study advances these model estimations by studying spatially and temporally smaller scales. The accuracy of temperature retrievals as a function of the number of battery readings is also studied. An extensive dataset of over 10 million battery temperature readings is available for São Paulo (Brazil), for estimating hourly and daily air temperatures. The air temperature estimates are validated with air temperature measurements from a WMO station, an Urban Fluxnet site, and crowdsourced data from 7 hobby meteorologists' private weather stations. On a daily basis temperature estimates are good, and we show they improve by optimising model parameters for neighbourhood scales as categorised in Local Climate Zones (LCZ). Temperature differences between LCZs can be distinguished from smartphone battery temperatures. When validating the model for hourly temperature estimates, initial results are poor, but are vastly improved by using a diurnally varying parameter function in the heat transfer model rather than one fixed value for the entire day. The obtained results show the potential of large crowdsourced datasets in meteorological studies, and the value of smartphones as a measuring platform when routine observations are lacking.

This chapter is published as:

Droste, A. M., J. J. Pape, A. Overeem, H. Leijnse, G. J. Steeneveld, A. J. Van Delden, and R. Uijlenhoet, 2017: Crowdsourcing Urban Air Temperatures through Smartphone Battery Temperatures in São Paulo, Brazil. *Journal of Atmospheric and Oceanic Technology*, **34** (9), 1853–1866, doi: 10.1175/JTECH-D-16-0150.1

2.1 Introduction

The need for high-resolution meteorological measurements is ever increasing. Numerical weather prediction models are improving continuously, which requires more precise measurements in both time and space for data assimilation (Ronda et al., 2017). Ongoing and projected global urbanisation (United Nations, 2012) make a thorough understanding of the urban atmosphere vital for urban planning, as well as for reliable forecasts of air quality, energy demand and heat stress. The Urban Heat Island (UHI) is widely studied, i.e. the difference in air temperature (T_{air}) between the rural background and the urban core (e.g. Oke, 1982; Arnfield, 2003; Steeneveld et al., 2011; Heusinkveld et al., 2014). Cities experience enhanced radiation uptake during the day as a result of their lower albedo and high heat storage capacity. Because of the slow nocturnal heat release from the urban fabric to the atmosphere, cities cool down more slowly than their surroundings, which creates the UHI. This may amount to 8K on hot and calm summer days (Oke, 1982). The UHI effect can exacerbate the degree of heat stress experienced by citizens ((Reid et al., 2009), which is projected to increase in the future by the combination of climate change and global urbanisation (United Nations, 2012; Miralles et al., 2014). Moreover, Hajat and Kosatky (2010) show that mortality increases with 2% per 1 ° C increase in high temperature. Oleson et al. (2015); Molenaar et al. (2016) project a drastic increase of future heat stress days caused by climate change for the USA and Canada, and the Netherlands, respectively. Both studies show that future heat stress is amplified in urban areas, underlining the need for knowledge of the temperature within the urban fabric. To understand these developments, we require urban temperature observations.

Traditional measurements in urban areas are scarce, and usually organised as intensive measurement campaigns (Heusinkveld et al., 2014) or cover just a small area (Kotthaus et al., 2012). This data scarcity can be harmful for e.g. megacities in developing countries, where knowledge on the mitigation of urban heat can be vital for citizen's health. Part of this data scarcity can be overcome by crowdsourcing: utilising data that is routinely collected by citizens or public sensors, and transferred over the Internet (the "Internet of Things" Muller et al., 2015; Warren et al., 2016), most notably by smartphones. Examples in the atmospheric sciences include the mPING app where users can share information about precipitation (Elmore et al., 2014); the iSPEX phone add-on which allows a user to measure optical thickness with a smartphone (Snik et al., 2014); estimating rain employing microwave links from cellular telecommunication networks (Overeem et al., 2013a); mapping forest fires by using voluntary observations sent by smartphones (Dalyot and Sosko, 2015); or making use of the built-in pressure sensor in many smartphone models to improve surface pressure forecasts (Mass and Madaus, 2014). Also, crowdsourced data from citizen weather stations uploaded to, for example, Weather Underground (Wunderground) and the Weather Observation Website (WOW) project (MetOffice, 2011), has proven to be valuable in urban research (e.g. Steeneveld et al., 2011; Bell et al., 2013; Warren et al., 2016; Meier et al., 2017). Using these stations asks for strict control of the quality of station and data, in terms of site-setup, measurement accuracy and data gaps. A thorough overview of crowdsourcing projects in the atmospheric sciences is given by Muller et al. (2015).

An innovative way of estimating urban air temperatures from smartphones was presented by Overeem et al. (2013b), henceforth *O13*. Using the OpenSignal application, *O13* employ 6-month datasets of smartphone battery temperature readings from 8 cities (including São Paulo, Brazil), with on average 844 selected battery temperature readings per city per day (1383 per day for São Paulo alone). They use a straightforward heat transfer model between phone, human

body and T_{air} (Figure 2.1, left) to translate the temperatures of smartphone batteries into daily average, city-averaged T_{air} . These daily temperature estimates (T_{est}) are shown to correspond well with measurements taken in the respective cities.

This paper builds upon the study of *O13*. Here, we employ a much longer (2 years) and denser (12×10^3 readings per day) dataset for São Paulo, Brazil. This study explores the potential of the *O13* heat transfer model at refined spatial and temporal scales. This study explores the potential of the *O13* heat transfer model at refined spatial and temporal scales. The current extended availability of temperature readings per day facilitates hourly air temperature estimates that have not been possible in previous research. We also investigate whether the model performance for daily averaged temperatures improves when applied only to selected neighbourhoods with their characteristic morphology, that is, the so-called Local Climate Zones (LCZs) (Stewart and Oke, 2012). By using an extensive set of validation data obtained from both certified sources (WMO, Urban Fluxnet) and crowdsourced weather stations (Wunderground and Netatmo), we provide a more robust representation of the actual city temperature dynamics on both daily and hourly time scales. Section 2.3 deals with the background of the heat transfer model and the Urban Heat Island; data and methodology are discussed in Section 2.2; results in Section 2.4; the discussion follows in Section 2.5 and we end with conclusions and perspectives in Section 2.6.

2.2 Background

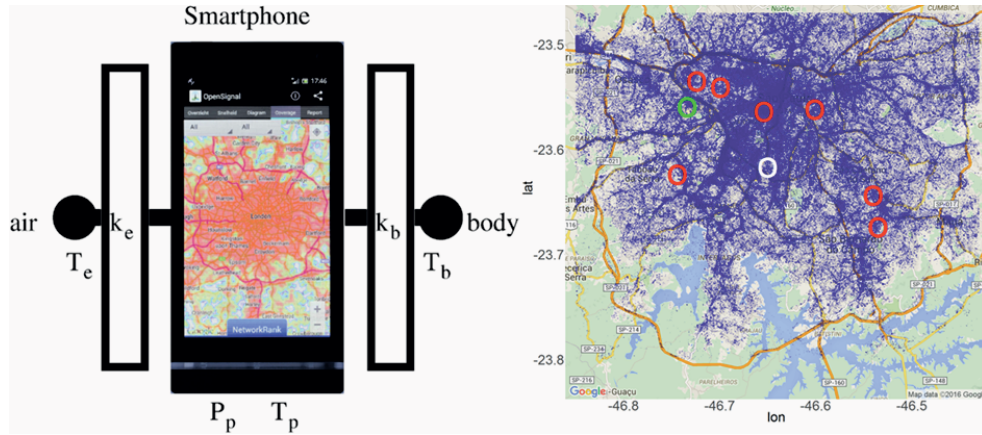


Figure 2.1: (left) Conceptual diagram of the heat transfer model (from Overeem et al. (2013b)). The $T_{e,p,b}$ is the temperature of the environment (e), phone (p), and body (b); P_p represents the thermal energy produced by the phone (or power); $k_{e,b}$ represents heat transfer rate between phone and environment (e) and body and phone (b). (right) Readings of smartphone battery temperature (in 2013, after selections) taken in São Paulo. Symbols show locations of the hobby meteorology stations (red circles), the FluxNet station (green circle), and WMO station Congonhas (white circle). For the exact coordinates of the stations consult Table 2.1.

2.2.1 Heat transfer model

O13 showed that smartphone battery temperatures can be used to obtain a daily average urban T_{air} , using a linear heat transfer model (Figure 2.1, left). In their model, the phone battery temperature T_p [° C] is regulated by the environmental air temperature T_e [° C], the body temperature T_b [° C] and the thermal energy generated by the phone P_p [W] (Equation S-2 in *O13*),

$$T_e = (1 + \frac{k_b}{k_e})T_p - (\frac{k_b}{k_e}T_b + \frac{P_p}{k_e}), \quad (2.1)$$

where the coefficients k_b and k_e [W ° C⁻¹] are determined by the thermal insulation between body and phone, and between phone and environment, respectively (Equation S-2 in *O13*). Assuming independence among values of T_p , k_e , k_b , P_p , and T_b over the set of measurements, and equilibrium between P_p and the heat flow to the body and the environment, leads to (Equation S-8 in *O13*)

$$T_e = m_j(T_p - T_0) + T_0 + \epsilon, \quad (2.2)$$

Here m_j is the average of $(1 + k_b / k_e)$ for a set of observations for a city j , ϵ is a random error, and T_0 is interpreted as the human body temperature (T_b in Equation 2.1), plus a constant, under the assumption that the heat transfer from phone to the environment is approximately zero. *O13* found $T_0 = 39$ ° C for their 8 cities, and calibrates m_j separately per city. Both constants are calibrated over the entire dataset (i.e., a constant value for both m_j and T_0). Note that the heat transfer model initially uses daily (and hourly later in this paper) and spatially averaged battery temperatures, rather than instantaneous battery readings. For the full derivation, we refer to the supporting information of *O13*.

2.2.2 Urban Heat Island and Local Climate Zones

This study defines the UHI as the difference in canopy air temperatures between urban and rural sites. Urban areas differ from their rural surroundings by the high prevalence of impervious surface and buildings, and little vegetation. Building materials have a high heat capacity, storing radiative energy during the day for subsequent slow nocturnal release. Additionally, the low sky-view factor induces efficient heat trapping inside the urban canopy (Oke, 1982). These combined effects cause the city to cool more slowly at night than the countryside, where energy is released much faster by virtue of the high sky-view factor and low heat capacity of vegetation. This creates the UHI, which peaks a few hours after sunset, when rural air temperatures have dropped and urban air temperatures can still be high.

Defining the UHI can be highly subjective: a clear definition of ‘urban’ and ‘rural’ is lacking (Stewart and Oke, 2012). Many UHI studies lack proper metadata, making comparisons between cases difficult (Stewart, 2011; Stewart and Oke, 2012). Defining the UHI as a temperature difference between LCZs can increase objectivity. The LCZ framework classifies land-use into 10 urban and 7 rural zones, each with its distinct surface (e.g. impervious fraction, vegetation cover) and building properties (e.g. building height, aspect ratio). The UHI can thereby also be defined as the difference in temperature between a rural and an urban LCZ, or even between two urban LCZs. In this study we define the UHI as a difference in canopy air temperature between two urban LCZs (section 2.4.2).

2.3 Data & Methodology

2.3.1 Smartphone battery temperature data

The study region is São Paulo (Brazil), which is located just south of the Tropic of Capricorn, at roughly 23.55° S and 46.63° W, at ~ 760 m MSL. The city is characterised by a subtropical maritime climate with mild dry winters and humid summers. The study area is confined to a rectangle around the city centre, between 23.47 and 23.80° S, and 46.43 and 46.85° W (Figure 2.1, right).

The battery temperatures are obtained from the OpenSignal app, a smartphone application which measures network signal strength from available providers. This app also logs T_p from the temperature sensor present in smartphone batteries. A T_p reading is taken when 1): the phone is being plugged in or out of the power source, and 2): when the phone is turned on or off. The selection procedure in this study follows that of *O13*. Only those readings made when 1): the phone is just being plugged into the power source, or 2): when the phone is turned on or off and the battery is discharging, are considered. To avoid spurious data in the analyses, an additional selection removes those battery temperatures outside the range between 10 and 47° C, since these readings are likely to be erroneous (because of, e.g., battery charging or intensive processor use). Battery temperature values are typically around 30° C (*O13*), so a large deviation to the positive or negative side would likely be a faulty reading, or severely influenced by battery charging. The battery temperature dataset covers January 1st 2013 up to December 31st 2014. During this period an average of 16×10^3 battery readings per day are left after filtering, though this number is significantly lower ($\sim 10^3$ per day) at the start of 2013, and rises to as much as 40×10^3 per day for several months in 2014.

2.3.2 Weather station air temperature data

Three main sources of weather station T_{air} data are employed for calibration and validation, i.e. WMO station Congonhas (WMO no. 837800), the 17-m.-tall Urban Fluxnet tower of the Micrometeorology group of the University of São Paulo (IAG-USP), and a set of 7 citizen weather stations. Congonhas is located at an airport, south of the city centre, in the middle of built-up environment (see Table 2.1 for station metadata, including classification into LCZs). The WMO data fully cover 2013 and 2014, with very few hours missing (less than 3 hours per month) and with 1.5-m T_{air} ($^\circ$ C) measured at the full hour, available as rounded integers. The Fluxnet T_{air} measurements are taken every 5 minutes and subsequently averaged into hourly values around the hour, at 0.1° C accuracy.

The data from the citizen stations are freely available for download from the Netatmo and Wunderground platforms (<http://www.netatmo.com> and <http://www.wunderground.com>, respectively), where weather enthusiasts can share their station data. At the start of this study, around 20 stations were active in the São Paulo area. First we selected only stations with fewer than 100 missing days per year. A day is considered as missing if it contains fewer than 21 hours of data. Very few stations meet these criteria in 2013, but 7 stations remain in 2014, one of which also has a sufficient record length in 2013 (see Table 2.1). Measurement accuracy is variable between brands of weather stations, since the more expensive stations tend to measure at a higher degree of accuracy, as a result of better radiation shielding and sensor quality (Bell et al., 2015). Typically, the better citizen stations have temperature measurement errors during

daytime of around 0.5°C (Steenneveld et al., 2011; Bell et al., 2015).

Since the instrument placement and setup of these hobby stations are not bound to strict rules, we have applied a series of filters, to ensure quality, accounting for the recommendations made by Stewart (2011). Data entries with sudden large temperature jumps ($>2^{\circ}\text{C}$ increments between 2 consecutive measurements) that are not confirmed in either the WMO or Fluxnet are removed. The temporal resolution of the measurements varies between the stations, but lies mainly between 5- and 10-minute intervals. For comparison to the other stations we have averaged the measurements to an hourly mean temperature. Past studies (e.g. Steenneveld et al., 2011; Bell et al., 2013, 2015; Meier et al., 2017) have demonstrated the value of these citizen data to good effect.

2.3.3 Calibration and validation

To create independent calibration and validation sets, data from 2013 are designated to the calibration set, and data from 2014 to the validation set. Both the WMO station and the Urban Fluxnet station are fully active throughout these years and are used for both calibration and validation (Table 2.1). However, the majority of the citizen stations only has a sufficient number of measurement days in 2014, and will therefore be used for validation purposes only. Hence, the model is calibrated and validated against the best possible representation of the average urban T_{air} , rather than just one fixed station, in order to ensure the most robust results. The original validation method in *O13* may suffer from autocorrelation between calibration and validation datasets, since the authors alternately assign days to the calibration and validation sets. The longer set of battery temperature readings in this study allows for a statistically independent calibration and validation set, to ensure that positive model outcomes are not due to autocorrelation effects.

The number of selected battery readings for 2013 averages at $\sim 8 \times 10^3$ per day, though this number is significantly lower in the first few months (below $\sim 1 \times 10^3$ per day up to May), totalling nearly 3 million readings in 2013. The battery dataset for 2014 is much larger, on average around 24×10^3 selected readings per day for 8.8 million readings in total. Battery readings are averaged into hourly and daily values. Days with fewer than at least 200 readings per day are excluded from the analysis. This results in the loss of 6 days in 2013, and none in 2014. For the hourly analysis, July 2013 and July 2014 are set as calibration and validation datasets, respectively. All days in July 2013 and 2014 have more than 200 measurements per day.

The T_0 parameter has been determined by *O13* for 8 cities across the world, rather than separately for each city under consideration. In this work the T_0 parameter is optimised for São Paulo using a least-squares approach, based on the one-year calibration (2013) dataset of battery readings (section 2.4.1). The value of m_j (Equation 2.2) is likewise determined, separately for the daily and the hourly calibration datasets. Parameter T_0 can be interpreted as the approximate human body temperature, which is not expected to fluctuate, whereas m_j represents a ratio of insulation coefficients. The factors influencing these coefficients (such as clothing) will be more variable over time. Therefore, m_j is calibrated separately for the analysis of hourly temperatures, since diurnal variation in heat isolation is expected to differ from the seasonal variation in the daily averaged temperature. This results in two calibration datasets: one for the daily dataset and one for the hourly dataset, which are used to train the model. Validation is done using the 2014 battery readings in the subsequent analyses.

Since the dataset used for analysis is extensive, we can determine the minimum number of battery readings \bar{N} needed for a stable model result. Employing random sampling, \bar{N} measurements are picked for every day and averaged into one daily value. This procedure is repeated 100 times per chosen value of \bar{N} to capture the mean battery temperature as accurately as possible, so every day has 100 mean battery temperatures. Each of these temperatures are validated against the city average air temperature (section 2.4.1).

2.3.4 Data air temperature modelling for a single neighbourhood

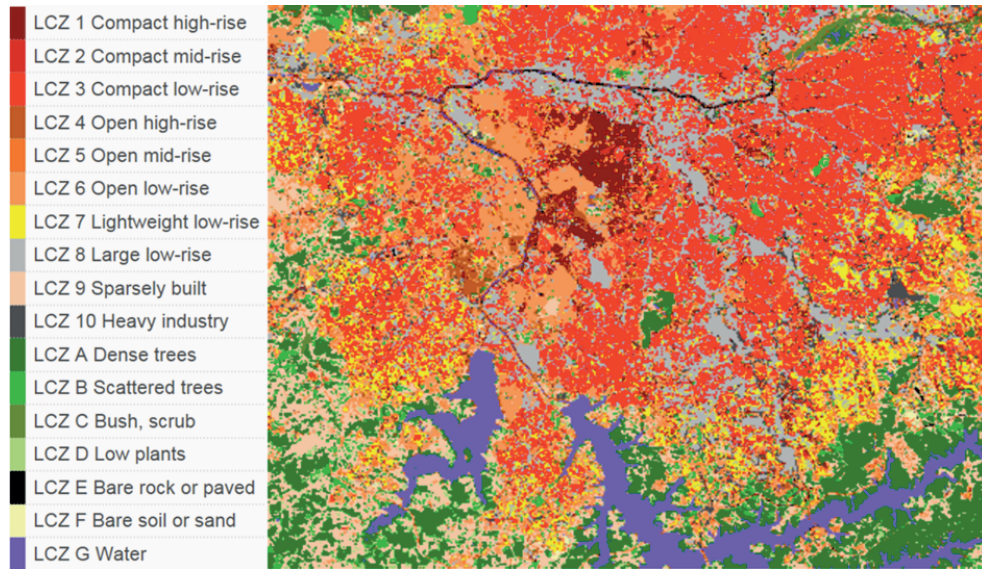


Figure 2.2: LCZ in the research area (São Paulo) at a resolution of 120m². The city is mainly made up of: LCZ3, compact low-rise buildings (22%); LCZ6, open low-rise buildings (15%); and LCZ8, large low-rise buildings (14%). [Image and legend are obtained from Geopedia, the World Urban Database and Access Portal Tools (WUDAPT) for visualizing LCZ data (<http://geopedia.world/>)

A point of interest is the role of the environment on the T_p reading. An analysis of model performance as function of the distance between battery reading and validation station yielded no significant relation (not shown). Instead, we study the influence of the urban fabric on the battery temperature estimates. Muller et al. (2015) write that “*the utility of smart phones for higher resolution UHI analysis [...] is still to be explored*”. To this end we utilise the LCZ classification for São Paulo, which was constructed using a GIS algorithm (Mills et al., 2015) and is available on the WUDAPT portal (www.wudapt.org). The location of each battery reading is coupled to the corresponding location on the LCZ map (Figure 2.2); the battery readings are subsequently grouped by LCZ and are used to validate the heat transfer model per LCZ. São Paulo mainly consists of low-rise buildings: LCZ3 in the centre (compact low-rise); a wide spread




Figure 2.3: Number of available battery readings after selection in July 2013 (circles) and July 2014 (squares) per hour as function of time of day. São Paulo is in timezone UTC - 3 h (summertime)

of LCZ6 (open low-rise) closer to the city border and several clusters of LCZ8 (large low-rise) (Figure 2.2).

2.3.5 Hourly air temperature estimation

For determining hourly T_{est} , we use July 2013 as calibration data and July 2014 as validation data. Model parameter m_j is calibrated to the diurnal course of temperature in July. Parameter T_0 is set at the optimal value for the entire year, determined using the methods described in section 2.3.3. For this analysis, July is the preferred month due to high data availability, and because July is one of the driest and cooler months, reducing possible effects of data distortion as a result of weather conditions (e.g., more people staying inside during heat waves or rain). Additional data selections, such as selections on LCZ and smartphone series, are not feasible with the hourly averaged data due to the strong reduction in available measurements, especially during night-time. Figure 2.3 shows the availability of smartphone readings against time of day (UTC). Around 8:00 UTC (5:00 LT) the number of measurements is at its minimum, at less than 10% of the daytime data density. Removing these data will strongly reduce the applicability of the dataset; however, excluding the night-time hours will lead to an unreliable calibration of m_j , and to missing hours in the resulting validation. In addition, we will explore the effect of using 24 hourly m_j constants, to better capture T_{air} variation. By this methodology the average diurnal variability of human behaviour (different clothing, being inside/outside, etc.) in July will be accounted for through m_j .

2.4 Results

2.4.1 Estimation of daily air temperature




Figure 2.4: Average daily air temperature in São Paulo during 2014 from T_{est} (blue line) validated against city average (WMO, FluxNet, and hobby stations) air temperature observations (black line) and daily averaged battery temperatures used as input (orange line)

This section presents the results for daily T_{est} from smartphone battery temperature readings. Figure 2.4 shows a validation of time series of daily T_{est} estimates against observed average city T_{air} , computed as the average of the various temperature measurements available (WMO, Fluxnet and the citizen weather stations). Shown statistics are coefficient of determination (ρ^2); Mean Error (ME); Mean Absolute Error (MAE) and Root Mean Square Error (RMSE). In general T_{est} compares very well with the observed air temperature, as was also concluded by *O13*. The analysis uses optimised values of m_j and T_0 : optimising T_0 for São Paulo only slightly changes its value in comparison to the standard value in *O13* (from 39 to 39.8 °C). The ρ^2 is 0.86, with a ME (or bias) of -0.53 °C. This bias is largest in January and February 2014. Mean T_{est} is consistently up to 2 °C lower than the actual measured temperature. We hypothesise that this is related to the number of measurements available in the calibration data. The number of measurements per day in the period of January – May is roughly 12 times lower than in the rest of the year ($\sim 1 \times 10^3$ versus $\sim 12 \times 10^3$). This could affect the calibration, since the months with the highest temperature peaks are under-represented in the model calibration (fewer measurements are available). The model seems to perform well for temperatures in the middle of the range; however, for temperatures close to the upper and lower limits, the model response underestimates the amplitude. When solely WMO data are used for calibration, results deteriorate due to the coarser resolution (1 °C) of the WMO data. Interestingly, calibrating m_j for separate seasons does not improve the performance (not shown), which indicates that variability in the (daily-averaged) heat transfer is not very strong over the year. Though São Paulo experiences seasonal variation in temperature, daily-averaged temperature variability is




Figure 2.5: Relation of the mean \bar{N} used per day (horizontal axis) against model performance: (a) ρ^2 and (b) RMSE. Average of 100 random samples of \bar{N} readings (large black dots) and the 100 samples (small black dots) to indicate the spread. Calibration of the data is performed on the entire dataset, using FluxNet as an air temperature reference. Validation is done against the average city air temperature (all stations combined).

weaker than, for instance, that of continental climates.

Figure 2.5 shows the model performance (ρ^2 and RMSE) as a function of the number of measurements \bar{N} used per day. It appears that after roughly 700 measurements the performance quality does not significantly increase and converges to a constant value. Apparently adding more data does not raise the quality beyond a certain threshold, but rather opens up more options for detailed analyses. This justifies stratifying the large dataset at our disposal into subsets for individual LCZs and even into hourly time intervals. The number of measurements left in these selections should still produce reliable results.

2.4.2 Daily temperature estimates per LCZ

The next step is to explore the model potential for the selected LCZ data, and whether spatial temperature differences can be identified and quantified by the smartphones. Whereas the city surface cover is mainly LCZ3 and LCZ6 (Figure 2.2), a disproportionately large number of measurements (19%) originates from LCZ1 (3% of surface cover): “compact high-rise”, which is typically found in the city centre (Figure 2.6). Datasets of battery readings from the LCZs with the most measurements (LCZ1; LCZ3; LCZ6; LCZ8) are used as model input. The resulting T_{est} values for these LCZs are compared to each other to study whether the urban fabric discernibly influences T_{est} . A daily average UHI per LCZ is calculated by subtracting the resulting temperatures from the daily averaged background temperature, taken from the WMO station. Note that this station is surrounded by built-up area and cannot be considered as an ideal rural station, though T_{air} differences between LCZs will still be visible using this approach. From this analysis a surprisingly high daily mean UHI arises for LCZ8 (~ 0.9 K), and a moderate UHI




Figure 2.6: Histogram of the distribution of battery temperature readings in 2013 over the urban LCZs in São Paulo. Fraction is calculated over the total of the 17 LCZs; the 7 rural LCZs (A–G) have not been plotted; and the LCZ corresponding to each number is explained in Figure 2.2.

for LCZ3 (~ 0.3 K; Figure 2.7a). Standard error in the mean for all LCZs is ~ 0.09 ° C. Using LCZ-specific battery temperatures does not strongly affect the model output: that is, the main difference with the full dataset is the change in sign of the ME for LCZ8 (from -0.48 ° C to 0.37 ° C; Figures 2.7b and c). Where the original model output underestimated the urban T_{air} , for LCZ8 the bias is positive, suggesting higher model temperatures as is indeed seen in the large positive UHI (Figure 2.7a). For LCZ3 the bias as compared to the full set remains negative but decreases (to -0.27 ° C). Since the statistical distribution of the data is unknown, the significance of the UHI effect in these two LCZs is investigated using the non-parametric Kruskal-Wallis test for two independent samples. Test results (not shown) confirm that the UHI magnitude between LCZs is significantly different, and the UHI magnitude between LCZs is significantly different. Hence, there is a discernable difference in T_{est} between these LCZs, which shows that the UHI can indeed be observed with this method.

2.4.3 Estimation of hourly air temperatures

In general the results for the daily T_{est} are satisfactory, as appears from the good model statistics in Figure 2.4. Next, we explore whether the method can also correctly estimate hourly averaged temperatures, despite the significantly reduced number of measurements available (Figure 2.3). The hourly T_{est} shows a relatively poor result (ρ^2 of 0.35) with a large spread (RMSE of 3.2 ° C) and an ME of roughly 0.9 ° C (Figure 2.8a). Upon a more careful examination, a delay in the timing of the peaks in the smartphone estimates as compared to the reference measurements is found (Figure 2.8a). That is, for the period around July 9th the peak of T_{est} occurs several hours after the peak of the measured temperature. Furthermore, the cooling rate in the evening is more rapid in the measured temperature, whereas T_{est} lags behind, cooling later and at a slower rate. This may be due to the heat capacity of the system (the phone itself, and the insulating layers




Figure 2.7: (a) UHI [K] derived from T_{est} per LCZ, using Congonhas as reference background station. The mean standard errors (whiskers) are shown. Scatterplots of observations (horizontal axis) against T_{est} (vertical axis) of the full (b) dataset (Figure 2.4) and (c) LCZ8. The 1:1 lines indicate where measurements would equal model results (straight black lines).

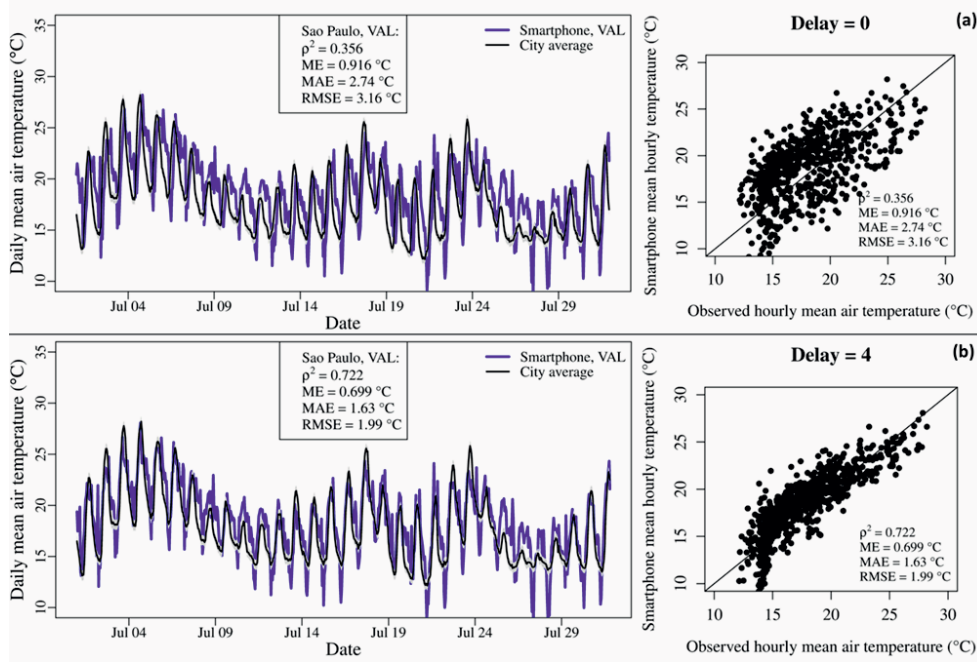


Figure 2.8: (left) Average hourly air temperature of São Paulo during July 2014 estimated from smartphone battery temperature readings (blue line) validated against city average air temperature observations (black line), and (right) scatterplots of observations (horizontal axis) against T_{est} (vertical axis); the 1:1 lines indicate where measurements would equal model results (straight black lines). (top) Results of the validation without a delay correction in the smartphone data and (bottom) with a delay-corrected dataset with delay $H = 4$ h

between phone & air, and phone & body) causing a delay in response. To explore whether results might improve, a delay is introduced to Equation 2:

$$T_{est}(t) = m_j[T_p(t + H) - T_0] + T_0 + \epsilon, \quad (2.3)$$

Here $T(t)$ is the temperature (in °C) at hour t (in hours UTC), and H is the delay in whole hours ($H=1,2,\dots,h$). Residual analysis of hourly T_{est} against the city average air temperature yields the best match between smartphone estimates and temperature measurements at $H=4$ hours, although already at $H=1$ results improve significantly over the uncorrected set. The ρ^2 doubles (from 0.35 to 0.72) and the large MAE and RMSE are reduced with over 1° C each, to 1.63 °C and 1.99 °C, respectively (Figure 2.8b). While the magnitude of the peaks (positive and negative) is still much larger than the measurements indicate, the timing of the estimated temperatures now corresponds much better to the observations. Analysis of the daily peaks in temperature reveals that on average the delay during the day is roughly 2 hours between model and observations, whereas at night the delay can be longer, on average up to 3 or 4 hours.

In search of a physical explanation for the delay, we formulate a simplified differential equation for the change in temperature of the phone as caused by the differences in between the phone

temperature, and the air- and body temperatures:

$$\frac{\delta T_p}{\delta t} = \frac{k_e(T_{air} - T_p)}{m * c} + \frac{k_b(T_b - T_p)}{m * c}, \quad (2.4)$$

Here m is the mass of the phone, taken as 0.13 kg, and c is the specific heat of the phone, taken as $600 \text{ JK}^{-1}\text{kg}^{-1}$ (based on specific heat of glass and sand, for simplicity). From the supporting information of *O13* it follows that k is conductivity times the area divided by the insulating material thickness. We take the typical dimensions of the phone as $10 \times 4 \text{ cm}$, and the conductivity of the clothing layer between phone and body as $0.037 \text{ Km}^2 \text{ W}^{-1}$ (0.24 clo, resembling a pair of trousers, ASHRAE (2010)). Furthermore, in our results $m_j = (1 + k_b/k_e) \approx 2$, so $k_b \approx k_e = 0.11 \text{ WK}^{-1}$. Using these typical values, T_b at 37° C and a linearly cooling of the atmosphere with 1 K per hour , we can simulate the cooling (heating) rate of the phone. This simple analysis indicates the phone arrives at a steady cooling rate after approximately 1 to 2 hours, depending on the exact initial values of $T_{air} - T_p$ and the specific heat and mass. The data seem to suggest a larger delay time (up to 4 hours): in reality, the heat capacity of the phone will be larger than assumed, by including the heat capacity of the bag or clothes in which it is being carried. The inside T_{air} for those readings taken indoors will influence the calibration: inside T_{air} reacts to outside T_{air} , with a lower amplitude and another delay factor, thus increasing the response time of the total smartphone system.

A second, implicit way to correct for the delay is by using 24 hourly m_j values, rather than a single fixed m_j for the entire dataset. By determining one m_j value per hour, the variations in heat transfer efficiency over the day are taken into account, since m_j is the ratio of the thermal insulation k -values (Equation 2.2). Possible variations in human behaviour (e.g., clothing) and the available measurements per hour can also be implicitly accounted for with this method. According to theory, the heat flow between phone and environment should decrease when the difference in temperature decreases. One would expect this to happen during the day when air temperature is higher and therefore closer to the smartphone battery temperature ($\sim 30^\circ \text{ C}$). At night, the temperature difference is larger and the rate of heat exchange would increase. When implementing an hourly variation of m_j the results (not shown) are very similar to the results for the delay-corrected series (Figure 2.8b). There are no appreciable differences between the two sets (i.e., hourly m_j versus delay-corrected), indicating that m_j implicitly corrects for the delay in the battery response. Values for m_j vary between 1.4 and 2.2 throughout the day, with the higher values occurring during night-time. A high value indicates that the thermal insulation between phone and environment (k_e) increases, or that the thermal insulation between phone and body (k_b) decreases. A higher k_e indicates a larger temperature difference between phone and environment (*O13*), as does indeed happen during the night (if the reading takes place outdoors).

The compensating effect of m_j on the delay in the battery temperature (Figure 2.8b) is confirmed when calibrating hourly m_j values to the explicitly delay-corrected set (Equation 3). When calibrating hourly m_j values to the explicitly corrected set, the range of m_j is halved (ranging between 1.7 and 2.1), though the diurnal pattern (lower m_j during the day) persists: no constant m_j is obtained for the delay-corrected set. Results do not notably improve: a 0.08° C reduction in the MAE while ρ^2 and RMSE remain equal in comparison to the hourly m_j calibration on the uncorrected dataset. This means that the observed delay in the smartphone battery estimates can be corrected for by either explicitly accounting for the delay (as in Equation 2.3), or taking 24 hourly m_j values rather than a single fixed value for the entire day.

2.5 Discussion

2.5.1 Relation to other studies

Our study extends *O13* by employing a more extensive dataset for just one city, and independent calibration and validation datasets. The São Paulo results of *O13* show nearly the same ρ^2 based on two periods of 3 months (ρ^2 of 0.65 and 0.85 for winter and spring 2012, respectively), as our results provide a ρ^2 of 0.86. Their MAE for São Paulo is only slightly higher than our values (1.2 ° C in *O13* and ~ 1.1 ° C here). This indicates that even with a smaller dataset (*O13* used on average 1383 measurements per day for São Paulo, whereas this study has roughly 10 times more), the daily averaged temperature on a city-wide scale can be captured well. Though our study has only focused on one city, the *O13* study was carried out for 8 different cities in vastly different climate zones, with different temperature seasonality and extremes. Their sound results indicate that the method is valid across a wide variation of climates, rather than only for São Paulo. The specific calibration constants of this work are optimised for São Paulo and are statistically not valid for any other city. However, this is not a fundamental limitation of the proposed method, since for other regions the model can be recalibrated using region-specific data.

Considering the data availability, we find that results deteriorate below a minimum number of battery measurements, even on the daily scale. As an illustration, Overeem et al. (2014) have applied the same method to Rotterdam and Amsterdam (the Netherlands), but their model statistics are less satisfactory (ρ^2 of 0.77 and 0.67; MAE of 1.22 and 1.40 ° C, respectively). The daily data availability was much lower for these comparatively small cities (382 and 116 per day for Rotterdam and Amsterdam, respectively). Similarly, with 203 readings per day Muller et al. (2015) report an even lower ρ^2 of 0.52, and a higher MAE of 1.71 ° C for Birmingham (UK). Overeem et al. (2014) provide a relation between the data availability and model ρ^2 , showing that the results become inaccurate for fewer than 100 measurements per day. Also, >350 daily measurements are preferable for accurate results: the optimal number of measurements is 700+ (Figure 2.5). *O13* fulfils this requirement, but the data availability reported in Overeem et al. (2014) and Muller et al. (2015) lies below this threshold. Note that the ρ^2 found by Muller et al. (2015) is lower than the lowest ρ^2 in Figure 2.5.

In our study 12×10^3 battery readings are available on average per day (section 2.3.1), and the ρ^2 value is 0.86 for the daily analysis (section 2.4.1 & Figure 2.4), which clearly illustrates the necessity of having enough data. Overall, our results tend to be of equal or better skill compared to earlier studies with the same heat transfer model. The high data availability provides possibilities for studying the method at an hourly scale or making selections for separate city areas (LCZs).

2.5.2 Data quality and additional filtering

A notable issue is the uncertainty in location of the smartphone, which is often in the order of tens of meters. This may introduce an uncertainty in coupling a battery reading to a location in the LCZ map (Figure 2.2). However, this map is based on satellite imagery with a resolution of 120m²; the uncertainty in smartphone location should fall within this range. In addition, the phone's GPS tracking is not always turned on (*O13*), so it remains difficult to discern between indoor and outdoor readings. However, the applied data selection (see *O13* and section 2.3.1) aims to minimise the uncertainty. In addition, the calibration process will account for this effect

as well if relevant. Importantly, the apparent time-lag between temperature changes in phone and environment (section 2.4.3) suggests that readings taken inside may still have been affected by the outdoor temperature.

Moreover, our approach assumes the phone is carried in a pocket, which allows for an exchange of heat between body and phone. In practice this assumption may be violated, e.g., because the phone is carried in a bag or elsewhere, and on an hourly scale the system might not be in equilibrium. However, calibrated m_j values appeared to be close to a-priori estimated m_j values from clothing properties (*O13*), which supports confidence in the followed approach. Additionally, we assume m_j to be constant over time, whereas clothing thickness (insulation) will obviously undergo a diurnal and seasonal cycle. Possibly, using the light sensor that many smartphone types now possess, a distinction can be made between indoor and outdoor measurements, if these data are available. A follow-up study that improves the heat transfer model by reducing the assumptions made could be very valuable for further research with these data. The weather can also influence human behaviour. On very hot days or days with extreme precipitation, people are more likely to stay indoors, meaning that readings taken during those periods will not reflect the outside air temperature, but rather the indoor environment. For instance, in May 2014 several hail events occurred in São Paulo, during which the error between T_{est} and the observations was relatively high (up to 3° C on May 19th) compared to clear days. A sensitivity test where all days with rainfall higher than 0 mm were excluded (leaving 238 dry days for calibration and 254 for validation) did not significantly improve model results enough to justify losing several weeks of data completely in the rainy months. Less strict filters (1, 2 and 5 mm) made the results nearly identical to the results without any filtering for precipitation. Therefore we decided to not pursue this sensitivity aspect further.

These issues are an inherent drawback of using smartphones for air temperature data, but by averaging a large amount of battery data in space (on the scale of a city or a LCZ) and in time (daily or hourly), the errors will be filtered out to a certain extent, as can be concluded from the favourable results of this study. For a thorough analysis of the reaction of the smartphone battery to changes in air temperature over the course of the day, a controlled trial should be set up with a conventional temperature sensor and several smartphones logging battery temperature. Because of several limitations (such as an inability to do continuous battery logging) we could not perform such a trial, but we strongly recommend it for any future research.

Two additional analyses did not yield any improvement: an attempt was made to use only smartphones with similar hardware (Samsung-GTI series), which in theory will have similar k coefficients, similar heat capacities, similar battery temperature sensors, and a similar thermal energy generated by the battery (P). However, this was no improvement over using all the smartphone data available. Additionally, we constructed an extra filter of the raw battery temperatures by Gaussian mixture modelling (Reynolds, 2009). This statistical technique assumes that the dataset consists of several sub-distributions or data-clusters, each with its own mean and standard deviation. A data-cluster that is e.g. characterised by high temperatures could be influenced by the battery charging or extensive use of the phone. These faulty data could be filtered out to improve results. However, the resulting data clusters of the mixture modelling had mean temperature differences smaller than 1° C with standard deviations around 6° C, which hampers distinguishing of clusters.

2.5.3 T_0 calibration

Initially, results with the optimised T_0 (section 2.4.1) were worse than the results using the reference T_0 from *O13*. The T_0 was found to be as high as 49°C and the MAE of the validation results significantly increased (by $\sim 0.3^\circ\text{C}$ compared to the set using $T_0 = 39^\circ\text{C}$). This large T_0 value cannot be realistically interpreted as the approximate body temperature, which is ideally near 37°C , plus a constant (Equation S-9 in *O13*). A more physically sound value for T_0 (39.8°C , used in the analyses) was obtained from repeating the optimisation procedure for incrementally increasing random samples of battery temperature used for calibration. For an increasing number (\bar{N}) of measurements, the T_0 and RMSE values decrease until $\bar{N} = 3 \times 10^3$ (or 229 days of battery readings used). Beyond this point T_0 remains constant at 39.8°C and RMSE does not appreciably decrease any more (value $\sim 1.49^\circ\text{C}$).

2.5.4 Weather station measurement data

With three different sources of measurement data, each with their own measurement accuracy, resolution, and location which determines the fetch of the station, knowing which station represents ‘the truth’ is nearly impossible. What the ‘true’ city temperature is remains open for debate. Since the city is heterogeneous by default, the temperature is increasingly influenced by local characteristics when moving from the boundary layer top to the surface layer (Barlow, 2014). What is measured by the smartphones is the urban canyon temperature, influenced by the local microclimate. Using a city-averaged T_{air} constructed from all these measurements seems to be the most robust option, to represent the urban air temperature as accurately as possible. However, this approach will not always yield the best model statistics. Particularly with the LCZ analysis, calibration and validation of the separate LCZs would ideally be performed with a station located in the same LCZ. Calibration of the model using specific LCZ air temperature data would make it better suited to detect differences in smartphone response between LCZs. The number of suitable citizen stations to use was scarce, however. For a city with more stations to choose from, a more thorough selection procedure (based on e.g. measurement height, meta-data or neighbourhood) could be performed according to the principles in Stewart (2011) and Bell et al. (2015). This might also reduce the high uncertainty at night, which can among others be caused by the high variability in measured minimum temperature (Brandsma and van der Meulen, 2008), in combination with the low number of battery readings available during those hours. For estimating the absolute value of the UHI with smartphones, a robust rural background station is essential, but one was unavailable in this study.

2.5.5 Applicability

Though this article primarily functions as a proof of principle, smartphone-derived air temperatures can have various applications to complement conventional data. For instance, in developing countries where weather stations are scarce, but smartphone ownership is high, they can add valuable information about the urban temperature. This knowledge can be vital during for instance heat waves, where knowledge of which neighbourhoods are most prone to the UHI can potentially save lives. Additionally, whereas a traditional urban measurement network is very expensive to set up and maintain, and will be prone to vandalism, a smartphone network will not be hindered by these limitations, providing valuable data virtually for free. This will be particularly valuable for those cities where funds for urban research are limited.

Alternatively, data assimilation in NWP models can be beneficial for NWP since model resolution is steadily increasing to an extent that the influence of cities will be felt (ECMWF, 2016). An urban scheme is often lacking within these NWP models, so data assimilation of the urban meteorological data will be crucial for reliable forecasts. Given the scarcity of urban data, even relatively coarse data such as the smartphone-derived temperatures can make a contribution to the forecasts. In broader terms, the developed methodology of this study may also be useful for algorithms which are being developed for application to other types of crowdsourced data. A preliminary test in which smartphone based temperature data were assimilated within the WRF modelling system for São Paulo, revealed that maximum temperatures were forecasted by about 0.5-1K better for the studied week (not shown).

Table 2.1: Metadata of the stations used in S\~ao Paulo. LCZs, distinction between balcony and garden stations, and the street aspect ratio (Height/Width) have been visually determined by using Google Streetview. Greenness fraction has been determined in a radius of 250m around the stations using Google Earth. Note that begin and end of the record indicate the length of the record at our disposal: “active” means the station is still running as of 31 Nov 2015. Citizen stations have missing days indicated for only 2014, since they are not used in the calibration process. Netatmo is the name of stations and is the brand name, i.e., hardware name (Netatmo).

Station code / name	Begin / End	no. missing days in 2014	Lat ($^{\circ}$ S); Lon ($^{\circ}$ W)	LCZ	Citizen station hardware	H/W; balcony / garden station	Greenness (%)
Congonhas (WMO nr. 837800)	1-1-2013 / Active	0 (0 in 2013)	23.617; 46.650	3		N/A (Airfield)	10
Fluxnet station USP	1-1-2013 / 31-12-2014	3 (13 in 2013)	23.560; 46.734	6		Roof top. H/W ~ 0.3	40
ISPSOPAU3	2-9-2013 / 20-5-2015	4	23.624; 46.744	4	DIY Arduino Funnel input/output (Fio), BMP085, DHT22	Balcony. H/W ~ 1 to 2	14
ISPSANTO2	5-8-2013 / Active	91	23.674; 46.535	3	Ambient Weather WS2080	Garden. H/W ~ 0.5	8
ISOPAULO61	3-3-2014 / Active	63	23.562; 46.601	2	AcuRite Pro Weather Center	Garden. H/W ~ 1 to 1.5	4
ISAOPAU11	1-1-2013 / Active	0	23.541; 46.699	4	Davis VT Pro 2	Balcony. H/W up to 3	12
ISOPAULO5	1-1-2013 / 5-12-2014	44	23.564; 46.654	4	HOBO U30	Balcony. H/W up to 3	4
Alfenas	7-4-2014 / 31-3-2015	96	23.644; 46.540	3	Netatmo	Garden. H/W ~ 0.5 to 1	6
Casinha	14-9-2013 / 31-3-2015	0	23.535; 46.724	3		Garden. H/W ~ 0.5 to 1	12

2.6 Conclusions

This study utilises a heat transfer model to translate smartphone battery temperature readings into city-wide air temperatures, on both a daily and an hourly scale. This work extends earlier work by Overeem et al. (2013b) by using an extensive dataset spanning 2 years of over 10 million battery readings taken in São Paulo, Brazil. We use a wide variety of multiple measurement stations spread across the city for calibration, thereby better capturing the average urban air temperature than using a single WMO station. The extensive data availability allows for a division of the dataset per Local Climate Zone (LCZ) to investigate spatial differences in temperature, as well as zooming in to the hourly temperature variations as captured by the smartphones. The consistent division into a separate calibration (the year 2013) and validation period (the year 2014) for both daily and hourly temperatures ensures that all results are statistically robust, and not subject to autocorrelation.

Estimated daily averaged air temperatures are good and can even be used to calculate temperatures of specific LCZs. A daily averaged Urban Heat Island (UHI) can be found in LCZ8 (large low-rise) and LCZ3 (compact low-rise): these LCZs have a significant difference in temperature in comparison to the official WMO airport station. However, insufficient battery temperature data are available to estimate hourly UHI. This would also need a proper rural background station: the airport is fully surrounded by built-up area (LCZ3).

On the hourly scale, initial results for temperature were poor, but vastly improve after correcting for a seemingly delayed response of the battery temperatures to changes in air temperatures. An analogous improvement can be obtained by using 24 hourly calibration (m_j) constants rather than one average value for all hours. The incorrect magnitude of especially the night-time peaks remains an unsolved issue, possibly as a result of the low number of battery temperature readings taken at night. A larger set of battery temperatures, especially when taken at night, is required to reduce the negative night-time bias. Making use of an urban testbed like Rotterdam (Heusinkveld et al., 2014) or Birmingham (Muller et al., 2015; Warren et al., 2016), could aid with this issue.

From a large number of smartphones reading an accurate air temperature estimate for the daily and even hourly scale of a city can be obtained, which underlines the strength of crowd-sourced data. With newer smartphone models regularly carrying temperature, moisture or pressure sensors, as well as applications such as mPing or WeatherSignal, there is no denying that measurements from smartphones may hold a lot of potential for future (urban) meteorological studies given their interconnectivity and every-day use in great numbers

Acknowledgements

We would like to thank Prof. Amauri Oliveira (University of São Paulo), for providing their Fluxnet station data. Dr. Gerald Mills (University College Dublin), Michael Foley (UCD) and Maria de Fatima Andrade (USP) for creating and supplying the LCZ WUDAPT data for São Paulo, and Dr. James Robinson and the OpenSignal company for supplying the battery temperature dataset. We extend our thanks to the hobby meteorologists for maintaining and uploading the data from their personal weather stations. We thank Prof. Berthold Horn (Massachusetts Institute of Technology), for sharing the idea of Gaussian mixture modelling. Gert-Jan Steen-

evel and Arjan Droste acknowledge funding from the Netherlands Organisation for Scientific Research (NWO)-VIDI grant “The Windy City” (File number 864.14.007) and NWO – eScience project ‘ERA-URBAN’ (grant 027.014.203)

3

Chapter 3

Crowdsourcing for urban wind research

The use of crowdsourcing – obtaining large quantities of data through the internet – has been of great value in urban meteorology. Crowdsourcing has been used to obtain urban air temperature, air pressure, and precipitation data from sources such as mobile phones or personal weather stations (PWS), but so far wind data has not been researched. Urban wind behaviour is highly variable and challenging to measure, since observations strongly depend on the location and instrumental setup. Crowdsourcing can provide a dense network of wind observations and may give insight into the spatial pattern of urban wind. We evaluate the skill of the popular “Netatmo” PWS anemometer against a reference for a rural and an urban site. Subsequently, we use crowdsourced wind speed observations from 60 PWSs in Amsterdam, the Netherlands, to analyse wind speed distributions of different Local Climate Zones (LCZs). The Netatmo PWS anemometer appears to systematically underestimate the wind speed, and episodes with rain or high relative humidity deteriorate the measurement quality. Therefore, we developed a quality assurance (QA) protocol to correct PWS measurements for these errors. The applied QA protocol strongly improves PWS data to a point where they can be used to infer the probability density distribution of wind speed of a city or neighbourhood. This density distribution consists of a mixture of two Weibull distributions, rather than the typical single Weibull distribution used for rural wind speed observations. The limited capability of the Netatmo PWS anemometer to measure near-zero wind speed causes the QA protocol to perform poorly for periods with very low wind speeds. However, results for a year-long wind speed climatology of the wind speed are satisfactory, as well as for a shorter period with higher wind speeds.

This chapter is accepted for publication after minor revisions as:

Droste, A. M., D. Fenner, B. Heusinkveld, and G. J. Steeneveld, 2020: Assessing the potential and application of crowdsourced urban wind data. *Quarterly Journal of the Royal Meteorological Society*, 1–18, doi: 10.1002/qj.3811

3.1 Introduction

The urban climate is a subject of increasing interest in science and society. With ongoing climate change and urbanisation, the need for accurate urban weather information has never been more urgent. In order to combat the effects of heat waves, air pollution or urban flooding, knowledge of urban weather can assist in identifying risk prone areas, to which urban planners can find a solution. The urban climate is of particular importance to human thermal comfort and air quality, and their associated health risks (Moonen et al., 2012).

The lack of urban weather observations is a major challenge in characterising the urban climate. Several cities have dedicated observational networks, e.g. Birmingham (England, Warren et al. (2016)), Berlin (Germany; Fenner et al. (2014)), Novi Sad (Serbia; Savić et al. (2019)), Ghent (Belgium; http://www.observatory.ugent.be/index_eng.html) or Amsterdam (the Netherlands; Ronda et al. (2017)), but the majority of cities across the world lacks such a detailed network. WMO regulations prevent official weather stations to be located in cities, since they require a relatively open field. Though obstructions characterise the urban climate, their heterogeneity complicates taking representative measurements of the city as a whole. From street to street, vast differences can occur in especially wind speed, but also temperature and, to a lesser extent, humidity (Heusinkveld et al., 2014).

To compensate for this lack of data, the urban meteorological research community has embraced the use of crowdsourcing. Muller et al. (2015) define crowdsourcing as “*Obtaining data or information by enlisting the services of a (potentially large) number of people and/or from a range of public sensors, typically connected via the Internet.*” An increasing number of crowdsourcing studies has been conducted recently, mainly to study the urban heat island effect (Steenekveld et al., 2011; Chapman et al., 2017; Fenner et al., 2017), urban rainfall (De Vos et al., 2017) and air pressure (Mass and Madaus, 2014). Urban air temperature and rainfall are well-captured by so-called Personal Weather Stations (PWSs): small weather stations designed for use by citizens, that can be installed on balconies, in gardens or on roofs.

Wind has so far not been researched through crowdsourcing. Urban wind speed and direction are hard to quantify due to the strong turbulent nature of wind. Observational studies have usually confined to single streets, where canyon profiles of wind speed and direction are measured with masts (Rotach et al., 2005; Eliasson et al., 2006). Urban wind studies often rely on wind tunnel experiments or computational fluid dynamics models to study wind loads on buildings, at pedestrian level, and for urban pollutant dispersion (Carpentieri and Robins, 2015; Ramponi et al., 2015; Toparlar et al., 2017). Knowledge of the urban wind is important for topics such as air pollution dispersion (Pascal et al., 2013), mechanical wind loads on buildings (Ramponi et al., 2015), human thermal comfort (Hsieh and Huang, 2016; Heusinkveld et al., 2017), and urban wind energy potential (Kent et al., 2017).

Crowdsourcing might be useful to investigate the urban wind climate. Crowdsourced data are available in great quantities, but quality is often relatively low, so serious scrutiny of data is required. The station setup and site representativeness are generally less well-known (Muller et al., 2015), which impacts data interpretation. In this study, we aim to learn whether crowdsourced wind speed data from “Netatmo” PWSs are suitable for analysing urban wind speed. We perform this research in Amsterdam (the Netherlands), where the Meteorology & Air Quality group of Wageningen University operates an urban network of automated weather stations with

high-quality wind measurements (Ronda et al., 2017), to serve as reference against which the crowdsourced stations can be tested. First, the Netatmo wind module is compared to reference sonic anemometers records in the field, in both a rural and an urban setting. Subsequently, from these field tests a bias correction and Quality Assurance (QA) protocol is established, which is then applied to the crowdsourced urban wind speed measurements. Finally, these data are used to analyse the wind speed characteristics of different Local Climate Zones (LCZs, Stewart and Oke (2012)) in Amsterdam, and compare the results to the reference network.

Below, Section 3.2 introduces the data (crowdsourced and reference stations), as well as the QA protocol used to filter the crowdsourced data; section 3.3 shows the results, which are discussed in further detail in section 3.4, before final conclusions are drawn in section 3.5.



3.2 Data & Methodology

3.2.1 The Netatmo wind module and data gathering

We focus on the Netatmo brand PWS, because of its popularity as PWS brand in Europe. As an illustration, in large cities such as Berlin or Paris, hundreds to thousands of Netatmo PWSs are set up (Meier et al., 2017), but even smaller cities such as Amsterdam or Toulouse appear to be equipped with hundreds of stations measuring urban weather (De Vos et al., 2017; Napoly et al., 2018). All Netatmo PWSs contain equal hardware, which limits discrepancies between stations, and allows development of a uniform bias-correction and QA procedure. Meier et al. (2017) have developed such a QA procedure for PWS air temperature data, and we will follow their example to develop such a system for PWS wind speed data, using reference measurements during QA.

The Netatmo company started distributing their wind sensor midway through 2015, as the latest addition to their weather station. It is a cylindrical sonic anemometer of 11 cm tall and 8.5 cm in diameter, using 4 nodes in an opening in the middle of the cylinder to measure the zonal and meridional wind components (Figure 3.2). Measurements are made every 6 seconds and aggregated to mean and maximum output values every 5 minutes. Accuracy of the wind speed measurements is 0.5 m/s, and 5 °for wind direction (Netatmo, 2019).

The data of this study are obtained through the Netatmo online API (Application Programming Interface) method *getstationdata*. This method provides wind data at roughly 5-minute resolution (variable timeframe) in rounded integer km/h for wind speed and degrees for wind direction. The API requires station and module identifiers, that are requested from the *getpublicdata* API method. This method outputs a list of station identifiers and their corresponding weather modules (outdoor module, wind, or rain meters) which can be used in *getstationdata*. For Amsterdam, 60 PWSs measuring wind speed were present within the period January 2016 – July 2018. Not all stations were active for this whole period: at most 52 stations actively measured in a single day, but we see a general increase of the amount of stations measuring over time (Figure 3.1). At the beginning of 2016, few PWS owners will have had the new wind module, though the PWS itself has been gaining popularity over the years as well.




Figure 3.1: Number of actively measuring Netatmo wind stations over the course of the study period in Amsterdam.

3.2.2 Data evaluation

To evaluate the PWS wind speed measurements against a known reference, without the complexity of an urban environment, one Netatmo station was installed at the experimental rural weather field in Wageningen, The Netherlands (51.981°N; 5.622°E; 5.0 m.a.s.l.). This weather field is a well-watered flat grass field that conforms to WMO regulations for weather observations. The Netatmo wind module was installed at 2 m height (Figure 3.2, left panel), allowing for direct comparison with the Gill/Campbell Scientific CSAT3 3D sonic anemometer (measurement rate 10 Hz; resolution: 0.001 m/s; 2% accuracy), also installed at 2 m, distanced roughly 10 m away from the Netatmo sensor. Rain and relative humidity are also measured at the weather field and used in the development of the QA procedure (see section 3.2.3). The field comparison at the weather field ran from April 2018 to December 2018.

As a second reference, located in an urban setting, we utilise observations of three Netatmo anemometers (Figure 3.2, right panel) installed on the rooftop of the Chair of Climatology building of the Technische Universität Berlin, Germany (52.457°N, 13.316°E). Here the reference sonic is a Gill Windmaster Ultrasonic Anemometer installed on a pole, measurement height (middle of path): 9.74 m above ground level (3.74 m above roof level). Wind speed range: 0–45 m/s, resolution: 0.001 m/s, sampling at 10 Hz to give 1-minute data. The three Netatmo sensors were installed on a boom 0.55 m below the reference sensor, in a north-south configuration, sensors each 0.25 m apart (Figure 3.2, right panel). These comparison measurements ran from June 2018 to March 2019.

The reference to which we compare the urban PWS observations is the Amsterdam Atmospheric Meteorological Supersite (AAMS), which consists of 25 stations covering the city centre and suburbs, measuring wind, temperature, and relative humidity. The air temperature and humidity sensor (Decagon VP-3, U.S.A.) is mounted inside a 184 mm aspirated radiation shield (Davis, U.S.A.). The ventilation fan is powered by two small solar panels mounted on top of the shield. The fans work at global radiation levels $>100 \text{ W/m}^2$. The radiation screens are mounted

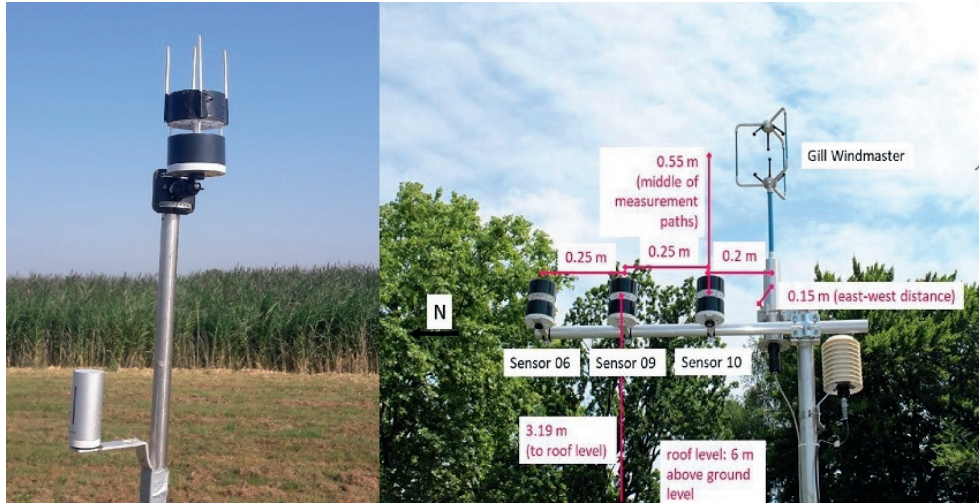


Figure 3.2: Left: Netatmo wind module (top sensor) at the Wageningen weather field, outfitted with spikes to prevent birds resting and disturbing the position. The silver cylinder is the outdoor temperature and humidity module, which is not used in this study. Right: Netatmo setup at TU Berlin. The sensor numbers refer to the identifiers of the Netatmo wind modules.

onto lantern posts, 0.46 m away from the edge of the lantern post, 4.0 m above ground level. The ultrasonic anemometer (Decagon DS-2, U.S.A.) has an accuracy of 0.30 m/s or 3% (whichever is larger). The anemometer is mounted above the radiation screen 0.50 m away from the lantern post edge and at a height of 4.30 m above ground level.

Rain and humidity observations from the WMO station at Amsterdam airport (Schiphol, WMO 06240, situated 10 km to the southwest of the city centre) are used in the bias correction for the PWS wind data. By using the WMO data as input for the bias correction (rather than the AAMS reference network), the correction protocol can be used for any city with a nearby WMO station, and does not require an extensive urban network. The WMO wind speed data is not used as a reference. The LCZ framework allows for an objective division of a city and its surroundings into zones with equal morphological properties, such as building heights, vegetation fraction, and building material. The LCZ map (Figure 3.3) we constructed for Amsterdam, following the WUDAPT guidelines (<http://www.wudapt.org/>) shows the location of the AAMS and PWSs, their respective LCZs, as well as the location of the WMO station. Table S1 contains the full table of AAMS station locations and LCZs.

Most PWSs are concentrated around the city centre (LCZ2 and LCZ5; compact / open midrise, and LCZ6, open low-rise), with some very close to the river and canals (LCZG, water), and three stations in sparsely built areas (LCZ9), near farmlands. We use LCZs as an indicator for urban morphology, which has a strong impact on wind speed (especially the ratio between building height and street width), so comparing stations with similar LCZs is required. We assume that morphology strongly determines a certain wind speed distribution, and by pooling the individual station data into one overall distribution per LCZ we are more likely to sample

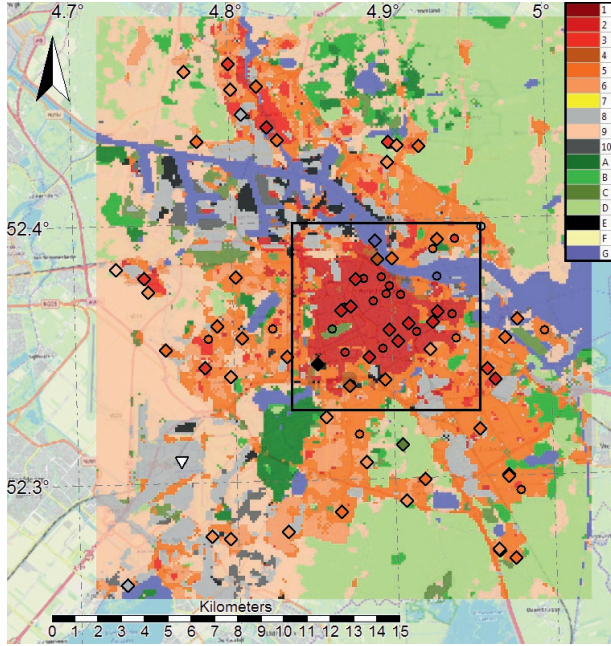


Figure 3.3: Local Climate Zone map of Amsterdam. The diamonds indicate the PWSs; the circles the AAMS stations, and the white triangle the WMO Amsterdam Airport station. The black rectangle indicates the “city centre” stations used for analyses.

the ‘true’ wind speed distribution for a given LCZ, and not the microscale wind climate of one particular station.

In section 3.3.4, the PWS data will be compared to the AAMS data over the period January 2017 to June 2018, which has good data availability for both AAMS and PWS. Only stations roughly in the city centre (between 52.33 and 52.4°N, and 4.837 and 4.95 °E) are used. Using the entire PWS network would give a biased image of the fit to the AAMS reference, since a large amount of PWSs are outside the city centre.

The applicability of the data in calm or windy conditions is examined in sections 3.3.4 and 3.3.4, for a relatively calm period (August – September 2017) and a relatively windy period (February – April 2018) respectively. Both periods contain a relatively high data density of the PWS network. Both sections also address the effect of morphology on the wind speed distribution, analysing stations in LCZ2 and LCZ5.

3.2.3 Quality Assurance protocol

To set up a QA protocol to improve the quality of the crowdsourced wind observations we follow Meier et al. (2017), who developed a rigorous QA procedure for air temperature measurements from crowdsourced data. We adapt and extend their QA protocol to be suitable for wind data, as follows:

- A Location requirement & morphology (QA A1 in Meier et al. (2017)). This criterion is based on the provided location as present in the PWS data (latitude and longitude). Stations with equal latitude and longitude are excluded. Additionally, we did a visual assessment using Google Earth to filter out unusual locations (such as a station in the middle of a canal). The initial is constructed from these criteria, leaving 60 PWSs. The stations are classified into LCZs, using the LCZ map of Figure 3.3.
- B **Data averaging & filtering.** The PWS data is provided in integer km/h, at roughly 5-minute resolution. This step aggregates all data (PWS, WMO, test field, AAMS) into hourly means. According to the Netatmo website, the minimum wind speed measurement is 0 m/s, with an accuracy of 0.5 m/s (1.8 km/h). However, having placed the wind module indoors for a period of time, we found the minimum measurement tended to be 1 or 2 km/h rather than 0 km/h, meaning that very low wind speed or calm conditions are probably not well captured by the sensor. This is also often reported by users at the official Netatmo forum (<https://forum.netatmo.com/>). The histograms of the raw PWS data indeed show a peak at precisely 1.0 km/h, much more than seen in reference sonic anemometer data (not shown). The crowdsourced urban data especially suffer from the large uncertainty at low wind speeds, which comprise a significant part (up to 20 % for some locations) of the wind distribution. To eliminate the large bias in wind distribution, all hourly means below 1.0 km/h are excluded from the analysis. A peak at 2.0 km/h remains visible but is less pronounced.
- C **Filtering for meteorological conditions.** From the field experiments we determine whether meteorological circumstances, such as rain or humidity, negatively influence the measurements. Netatmo users report that rain disturbs the measurements, and that the stations are prone to collecting moisture inside the sonic module. We investigate any significant influence of rainy and humid conditions, and whether filtering for these circumstances can improve data quality.
- D **Systematic bias correction.** Any systematic deviation from the actual wind speed as measured during the comparison measurements at the experimental sites will be corrected for. The bias correction based on the experimental setup will be applied to the (filtered) crowdsourced data from Amsterdam.

3.2.4 Wind statistics

A direct comparison between the crowdsourced data and the professional AAMS data, as was performed for Wageningen and Berlin, is complicated by the urban heterogeneity and the contrasting setup between PWSs and the AAMS stations. Whereas the AAMS stations are installed on lampposts within the street canyon (public space), the PWSs are installed in private space. The exact PWS setup is unknown, which adds uncertainty. Hence, the wind statistics, rather than instantaneous values, are compared with each other. Under idealised, undisturbed conditions, wind speed follows a 2-parameter Weibull distribution (Justus et al., 1978; Takle and Brown, 1978; Conradsen et al., 1984). This distribution is invalid below 0, has a peak at low values, and a long tail. The distribution is determined by a shape (a) and a scale (b) factor (Equation 3.1):

$$f(x) = \frac{a}{b} \left(\frac{x}{b}\right)^{a-1} e^{-(x/b)^a} \quad (3.1)$$

However, the observations do not always match the Weibull distribution for sites with disturbances, for instance where the wind speed distribution shows a bimodal pattern, or where there is a high probability of null (near-zero) wind speeds. This may occur in mountainous regions, but also in complex environments such as cities. This especially holds for PWSs, which have a relatively low accuracy, and are therefore very likely to have high peaks at the lower end of the wind speed distribution. Carta et al. (2009) have investigated several statistical distributions to capture a variety of wind regimes at the Canary Islands, including stations in mountainous regions. They found that a mixture Weibull distribution can represent a wind speed regime with a large probability of null winds, which is what we would expect in a city. Such a distribution combines two Weibull distributions into one overall mixture distribution: one representing the peak, and one representing the tail end of the distribution. Equation 3.1 then turns into Equation 3.2:

$$f(x) = \omega_1 \frac{a_1(x)^{a_1-1}}{b_1^{a_1}} \exp\left[-\left(\frac{x}{b_1}\right)^{a_1}\right] + \omega_2 \frac{a_2(x)^{a_2-1}}{b_2^{a_2}} \exp\left[-\left(\frac{x}{b_2}\right)^{a_2}\right] \quad (3.2)$$

Here $a_{1,2}$ and $b_{1,2}$ are shape and scale parameters, respectively, for the first and second mixture components. $\omega_{1,2}$ is the proportionality of the two mixture components, and their total sums to 1. In this case, as for equation 1, x represents the measured wind speed [km/h], and $f(x)$ is the probability density function.

We use the R-package *mixdist* to fit a mixture distribution to the PWS and AAMS data (Macdonald and Du, 2018). This method iteratively fits the shape and scale parameters, and estimates the proportionality of the two distributions, through a maximum likelihood procedure. This requires an initial estimate of the first order statistical moments and proportions of the two distributions. These are estimated from the data share below 3.0 km/h, which also provides an estimate of the proportionality.

We assess the performance of the PWS data against the AAMS reference data through the resulting probability density distribution (PDD). We do this graphically, and numerically using the coefficient of determination (R^2) between the two PDDs, as is common practice in the wind energy field (Garcia et al., 1998; Celik, 2004; Carta et al., 2009). In order to compare the wind speed distributions between the AAMS and PWSs, mixture Weibull PDDs are constructed for both data sources for similar LCZs, for an equal time period. All shown histograms of the measurements are made with 0.5 km/h bins: the PDDs are constructed from the fitted distributions parameters using 0.1 km/h intervals, up to the maximum range of the data (often 20 km/h).

3.3 Results

3.3.1 Evaluating the mixture Weibull distribution

To examine whether the mixture Weibull distribution indeed outperforms a regular Weibull distribution for representing urban wind speed, we fit various distributions to a single AAMS station, located in the city centre (station 2194, Spuiplein, see Table S1.). This site is a square, surrounded by midrise buildings (LCZ2). Since the site is not in a narrow street canyon, the influence of turbulent flows will likely be smaller than at a sheltered PWS site. Data were filtered for rain and humidity (QA step C), but not for low wind speeds, since the AAMS sonic

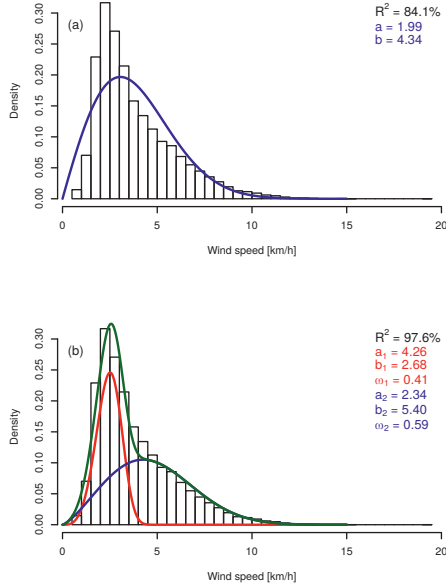


Figure 3.4: Regular (panel a) and mixture (panel b) Weibull distributions fit to the data of AAMS station 2194. Shape (a), scale (b) and proportionality (ω) parameters are given for all (mixture) components. In panel b, the green line indicates the resulting mixture distribution made from its two mixture component distributions (red and blue line).

anemometer is well-capable of measuring very low wind speeds. Fitting a normal 2-component Weibull distribution to the observations (Figure 3.4a) offers a fairly poor result: the peak at the low wind speeds is poorly captured whereas the right tail tends to overestimate the frequency of higher wind speed. The fitted mixture distribution on the other hand (Figure 3.4b) captures the peak and the tail end, though a slight underrepresentation of the transition between the peak and tail seems to occur (at ≈ 4 km/h). While graphically the mixture Weibull distribution fits well to the observations, the R^2 confirms that the mixture distribution ($R^2 = 98\%$) outperforms the regular Weibull ($R^2 = 84\%$). Note that R^2 values of Weibull distribution fits tend to be relatively high, since the 0 and the furthest tail generally match (always very close to 0) because of the shape of the distribution. Alternatively, we can evaluate the bulk of the wind speed, below 5 km/h, while excluding 0.0 km/h. R^2 for the mixture Weibull distribution for this lower end still amounts to 93%, but for the regular Weibull distribution drops down to 66%, confirming the superiority of the mixture Weibull distribution.

While these results show that in general the mixture Weibull distribution represents the disturbed urban wind environment well (other AAMS stations show similar results, not shown), for some stations the single Weibull distribution performs equally well, for instance in relatively open environments. However, the mixture Weibull distribution is a viable tool to use for the typically urban, sheltered PWSs.

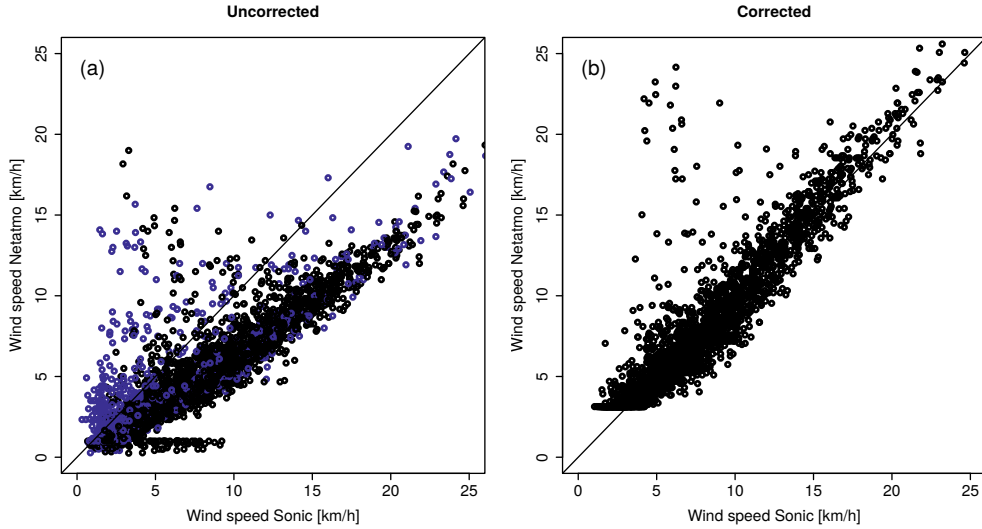


Figure 3.5: Scatterplots of hourly averaged wind speeds as observed by the sonic reference instrument (x axis) and the Netatmo station (y axis) at the Wageningen weather field. (a) shows the unfiltered, uncorrected data, with rain and/or humidity events marked in blue. (b) shows the resulting data after filtering and bias correction.

3.3.2 Bias analysis through comparison measurements at Wageningen weather field

For the Wageningen comparison measurements the unfiltered PWS data (Figure 3.5a) show a systematic underestimation of the wind speed which increases with the wind speed. Also, the PWS frequently measures 1.0 km/h when the actual wind speed is higher. Thus, hourly mean wind speeds of 1.0 km/h and lower are excluded from all crowdsourced datasets (QA step B).

Moisture can collect inside the device, which is not completely watertight, and which can influence the measurements (a common issue according to the users' forums). This problem appeared after 3 months when our installed PWS stopped measuring, at which point we dismantled the module, cleaned and dried it, and re-installed it in the field. Sonic anemometer measurements are known to be disturbed by rain and water droplets, which can affect the path of the sonic and the instrument itself (Campbell Scientific, 2017). To investigate the effect of humidity and rain, rain events are classified as hours with more than 0.1 mm accumulated rainfall, measured by the Wageningen pluviometer, as well as two hours afterwards, to also take possible collection of droplets onto the anemometer path into account. In a similar way, humidity events are hours with $>95\%$ relative humidity, measured by the reference. High humidity events mainly coincide with positive wind bias by the Netatmo station in the lower wind speed regions, whereas rain events are distributed over the entire wind speed distribution (Figure 3.5a).

After rain and humidity filtering, and removing the 1.0 km/h data, 64% of the data remain for analysis (2294 hours from 3570). We correct for the systematic bias based on the wind speed measured by the PWS, not on the 'ground truth' of the reference sonic. Hence, the correction is

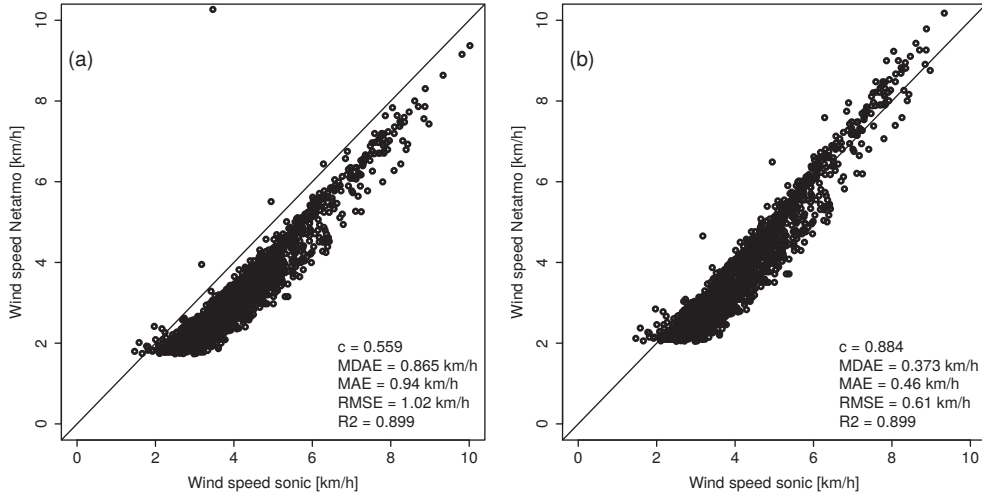


Figure 3.6: Hourly averaged wind speeds of the three Berlin Netatmo station (y-axis) against reference sonic observations (x-axis). Netatmo data have been filtered and corrected. (a) shows the correction using $c = 0.559$ obtained from the Wageningen dataset; (b) shows the corrections using $c = 0.884$, obtained from optimising on the Berlin data.

independent from reference data and can be used universally. The data is corrected with a linear regression model, optimised for the median absolute error (MDAE) of resulting corrected wind speed, to give high outliers less weight. Other regression models, including higher polynomial models and multiple linear regression models including other variables such as humidity or rain, have been tested, but a linear regression model explained most variance whilst maintaining model simplicity.

$$v = N(1 + c) \quad (3.3)$$

In Equation 3.3, v is the resulting, corrected wind speed [km/h], N the uncorrected (but filtered) wind speed measured by the Netatmo anemometer, and c is the regression coefficient ($c = 0.559$ in this case). The majority of the corrected data follows the 1:1 line (Figure 3.5b), though a portion of positive outliers remains. MDAE amounts to 0.78 km/h, down from 2.5 km/h in the uncorrected set. However, the RMSE of the corrected data is still 1.94 km/h, compared to 3.5 km/h in the uncorrected dataset, indicating the spread still visible in the corrected dataset.

3.3.3 Bias correction of comparison measurements in Berlin

To investigate the robustness of the correction coefficient obtained at the Wageningen site, it is applied to correct the Berlin Netatmo data, after QA step B. Data for rain and humidity are obtained from the WMO station Dahlem (WMO 10381), close to the measurement location (52.454 °N, 13.302 °E, 51 m.a.s.l.).

The Berlin wind speed data (Figure 3.6) are notably lower than the Wageningen data, with hourly averages only reaching 8 km/h, and some rare outliers to 14 km/h. This is partly due to the lower wind speed in urban areas in general but can also be caused by the trees surrounding

the building, sheltering the station. In this sense, the Berlin data represent what we can expect in terms of Netatmo setup: sheltered environments in gardens or on balconies. Consequently, around 40 % of the hourly values are 1.0 km/h, which are filtered out in QA level B. Filtering for rain and humidity events in QA level C (of which rain is the main contributor, humidity hardly ever reaches 90% for this site), brings the total data reduction to 45.1%. This indicates the main expected contributor to error is the underperformance of the hardware at low wind speeds.

The correction factor obtained from the Wageningen dataset ($c = 0.559$) is initially applied to the (filtered) Berlin dataset (Figure 3.6a). Results are good, with the output statistics showing a better response than for Wageningen: MDAE is only 0.55 km/h, RMSE is 0.64 km/h and R^2 is 83%. However, some underestimation still seems to be present, albeit not as prominent as in the initial data. Therefore, we repeated the optimisation procedure for the Berlin dataset to derive a new correction constant. This new correction constant is higher ($c = 0.884$), indicating that the underestimation of the Netatmo is stronger in Berlin compared to Wageningen. Especially at relatively low wind speeds, the fit to the reference observations is better, and model statistics improve, though not strongly (Figure 3.6b). At the upper end of the measurements, some overestimation appears due to the higher correction factor which weighs heavier on higher wind speeds (see Equation 3.3). Since this new correction factor is tuned towards a wind speed distribution characterised by lower (urban) wind speeds, this might be a preferred tool to correct the urban PWS measurements than the correction coefficient derived from the rural Wageningen data, which covers a much wider wind speed spectrum. In Section 3.3.4, both will be tested to see which one results in a better fit to the reference network. A separate calculation of the correction coefficient on the Wageningen data, using only wind speeds below 10 km/h, did not significantly change the previously found value for c .

3.3.4 Application of QA to Amsterdam data

The previous sections showed that the Netatmo anemometer is capable of measuring wind speeds, but requires a substantial correction for rain/humidity events and portrays a systematic negative bias. The following sections apply the QA protocol from Section 3.2.3 to the urban PWS data, comparing the data after each QA step to the rain and humidity filtered AAMS data as the reference, with the R^2 of the PDDs as measure of goodness of fit. In Sections 3.3.4 and 3.3.4 the data are separated into periods of relatively low and high wind speeds to analyse the dependency of the QA protocol on mean wind speed.

QA effect on the entire data set (January 2017 – July 2018)

The AAMS dataset contains 99950 records over 17 stations, after filtering for rain and humidity (which removed 34.7% of the originally available data). This filtered dataset serves as the reference against which the PWS data is tested across the various QA steps. For the PWS data, the total unfiltered data available (prior to any QA) is 164267 records over 17 stations (Figure 3.7a). At QA level B, which removes the 1.0 km/h values, the data set is reduced by 21.6 % to 128655 records (Figure 3.7b). At QA level C, filtering for rain and humidity, total data reduction is 45.9 %, or 88897 records (Figure 3.7c). The disturbing effect of the 1.0 km/h measurements can be seen in Figure 3.7a, where the peak is prominent and strongly influences the PDD of the PWS data.

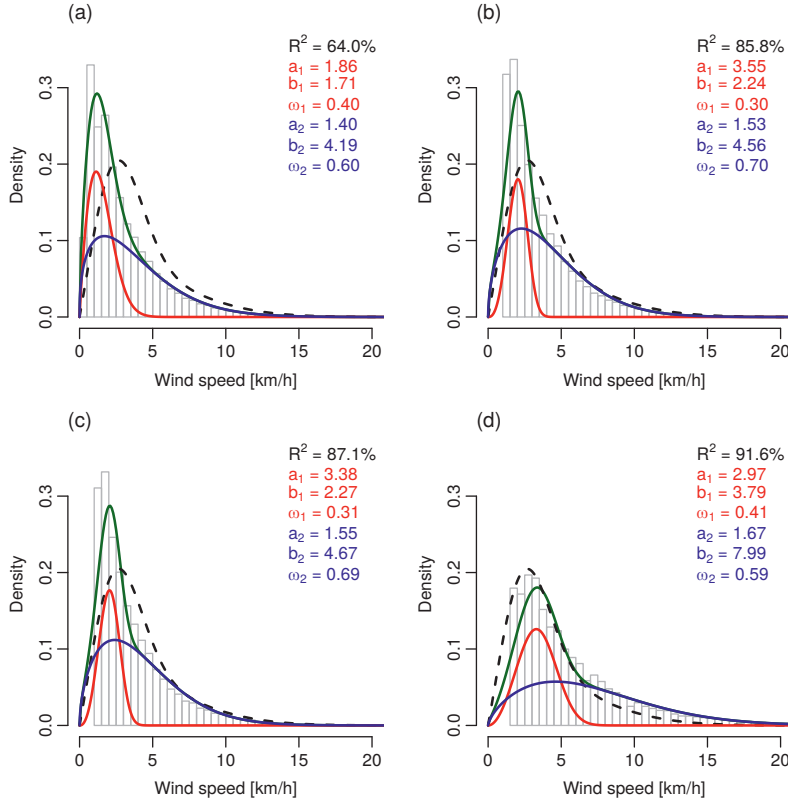


Figure 3.7: QA protocol applied to all city centre PWSs. Bars indicate the PWS data in 0.5 km/h bins; red and blue lines are the two mixture Weibull components that make up the mixture Weibull distribution (green). Black dashed line is the rain and humidity filtered AAMS reference probability density distribution (PDD); $a_{1,2}$, $b_{1,2}$ and $\omega_{1,2}$ are shape, scale, and proportionality parameters of the components, respectively, as in Equation 3.2. R^2 is the squared correlation between the PWS PDD and the AAMS PDD. Panel (a) is the unfiltered, hourly averaged data; (b) the data with 1.0 km/h values removed (QA level B); (c) the data filtered for rain and humidity (QA level C), and (d) the bias corrected data using $c = 0.559$ in Equation 3.3 (QA level D).

Fit to the reference AAMS data is poor ($R^2 = 64\%$), since the centre of the PWS PDD is much lower than the reference. Removing these 1.0 km/h values strongly improves the fit (to $R^2 = 85.8\%$, Figure 3.7b) but the peak at low wind speeds is still prominent. Filtering for rain and humidity (Figure 3.7c) does little to improve the peak values, since rain events tend to coincide with relatively high wind speeds, and as such the main data reduction occurs at the tail end of the PDD. The fit only marginally improves with respect to the previous QA step (to $R^2 = 87.1\%$). Applying the bias correction (using $c = 0.559$) results in a strong improvement ($R^2 = 91.6\%$), which eliminates the high peak value, and the PWS PDD fits that of the AAMS data much better. However, due to the linear nature of the correction, some overestimation of the

higher wind speeds is introduced. Using $c = 0.884$ deteriorates the result since low wind speeds are underrepresented in that case, and R^2 is only 77.7% (not shown).

Calm period (August – September 2017)

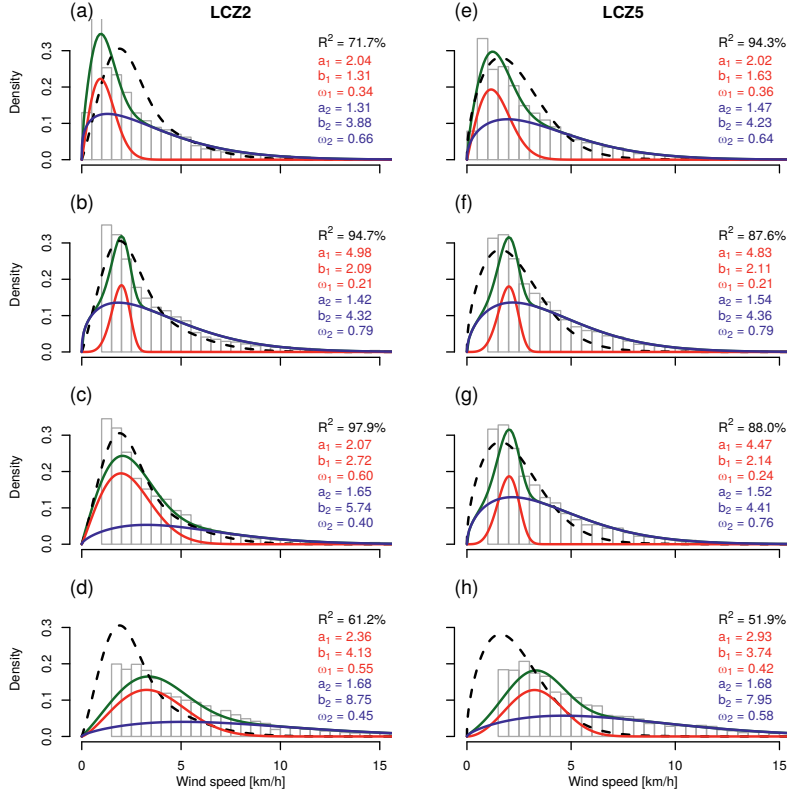


Figure 3.8: QA applied to the calm wind period, for LCZ2 (left, panels a to d), based on 9 PWSs (bars) and 12 AAMS reference stations (dashed black lines); and LCZ5 (right, panels e to h) based on 18 PWSs (bars) and 6 AAMS reference stations (dashed black lines). Setup similar as in Figure 7. In panels (d) and (h), QA level D, the value for c is 0.559.

August and September 2017 were characterised by generally calm conditions in Amsterdam: 78% of hourly wind speed values measured at Amsterdam airport are below 5 m/s. During the period 9 PWSs were operational in LCZ2, compared to 12 AAMS stations. Prior to any filtering (QA level A), the PWS data contains 10243 observations across all 9 stations. The 12 AAMS stations only have 2 missing hours in this period, but the rain and humidity filters remove 36.4 % of the data (leaving 11182 records).

The histograms of the wind data showcase the calm conditions in this period (Figure 3.8): both the unfiltered PWS data and the AAMS data (Figure 3.8a) peak at very low wind speeds, and the 1.0 km/h peak of the PWS data makes up for over 20 % of the histogram. Indeed, at QA level B (Figure 3.8b), 27.6 % of the data is filtered out, leaving 7418 records. This shifts the

centre of the distribution to almost exactly the value of the peak in the AAMS distribution, resulting in a very good fit to the reference even at QA level B ($R^2 = 94.7\%$).

The tail of the PWS PDD slightly overestimates the occurrence of high wind speeds, but by filtering for rain and humidity (Figure 3.8c) this is corrected, and the fit to the AAMS reference is further improved ($R^2 = 97.9\%$). When applying the bias correction (with $c = 0.559$), results become worse (R^2 drops to 61.2% , Figure 3.8d), since low wind speeds are overcorrected towards higher values. The bias correction does not seem to perform well during these circumstances with very low wind speeds, and just data filtering is enough to obtain a good fit to the reference network. The PWS's tendency to underestimate wind speed seems to only become an issue when wind speeds are not low (median wind ≈ 2 km/h), so QA step D for these situations is not recommended.

To ensure this result is not only valid for LCZ2, stations from LCZ5 are evaluated for the same period (Figure 3.8e-h). Here we have 6 AAMS stations with the same relative data reduction after rain and humidity filtering (5576 observations left), and 18 PWSs. There are a total of 18515 hourly PWS observations across the 18 stations prior to any filtering. The low winds seem even more prominent in LCZ5: the AAMS observations peak at 2.0 km/h (Figure 3.8e). Removing the 1.0 km/h measurements (QA level B) reduces PWS data by 20.6% (14700 measurements left). The rain and humidity filter (QA level C) brings the total data reduction to 49.6% (9329 measurements left), comparable to LCZ2. This mainly filters the higher wind speed observations at the right tail, though some overrepresentation of high wind speed remains. Oddly enough, here the best results have been made by just the hourly, unfiltered data (QA step A, $R^2 = 94.3\%$), since the AAMS wind speed values themselves are very low. Understandably, applying the bias correction (QA level D) deteriorates the fit to the AAMS values further since the distribution is shifted to the right (Figure 3.8h).

Some differences in the characteristics of the Weibull distribution appear between the two LCZs. Especially the scale (b) parameter is lower for both mixture components in LCZ5, indicating lower wind speeds altogether (the scale parameter scales with mean wind speed). The shape (a) parameter is only different for the first mixture component (red lines in Figure 3.8): at 4.47 in Figure 3.8f it shows a very narrow distribution with a clear peak, related to the narrow shape of the overall distribution, indicating low wind speed values and low spread. Regardless of the LCZ, the QA protocol is not fully able to improve wind speed estimations from PWSs under calm conditions, only up to a certain point (QA level C, and arguable not even that for LCZ5), and bias correction deteriorates results. Applying QA step D should therefore be dependent on the mean wind speed. However, applying a different type of correction model, consisting of two separate corrections (for wind speed either above or below 3 km/h) did not significantly improve the results compared to the model used here (results not shown). The underperformance is therefore likely due to the strong relative error introduced by the integer data resolution and the measurement accuracy of 1.8 km/h (0.5 m/s).

Windy period (February — April 2018)

The period of February — April 2018 was relatively windy, with 54% of the hourly wind speeds at Amsterdam airport above 5 m/s. 15497 AAMS records are available in LCZ2 after filtering, out of 10 AAMS stations, of which one reported no data for 1 month. 16945 unfiltered hourly PWS records are available for LCZ2, over 9 PWSs. Data reduction is 17.9% at QA level B

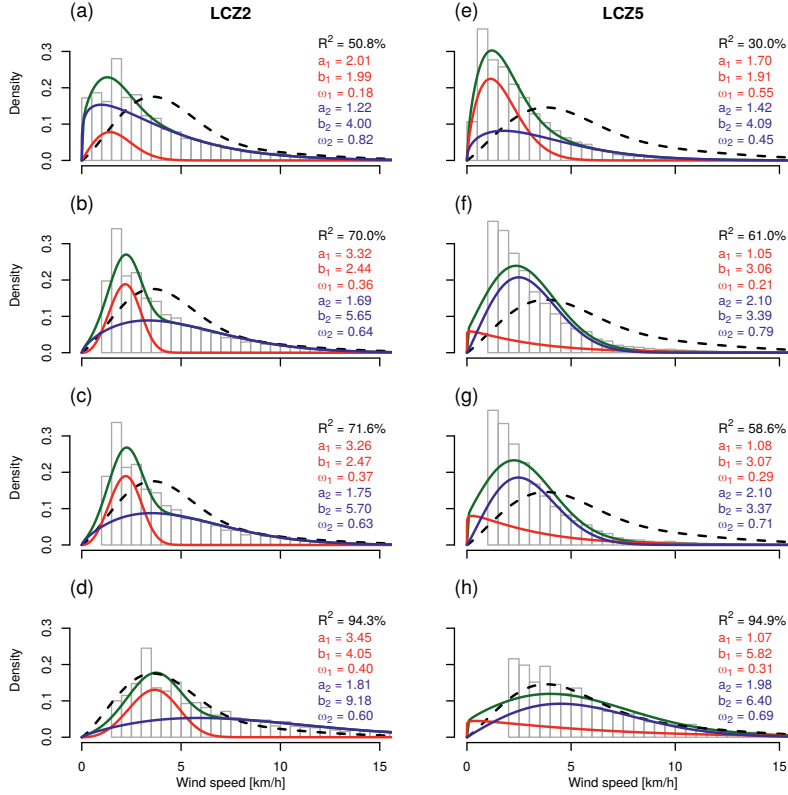


Figure 3.9: QA applied to the windy case, for LCZ2 (left, panels a to d), based on 9 PWSs (bars) and 10 AAMS reference stations (dashed black lines); and LCZ5 (right, panels e to h) based on 18 PWSs (bars) and 6 AAMS reference stations (dashed black lines). Setup similar as in Figure 7. In panels (d) and (h), QA level D, the value for $c = 0.559$.

(13910 records left); 38.0 % for QA level C (10507 records left). Compared to Figure 3.8, Figure 3.9 shows a wider distribution of wind speed, and the peak of the data is not at 1.0 km/h this time, but at 2.0 km/h (Figure 3.9b). The AAMS data is centred around 5 km/h, indicating that the average wind speed measured by the reference network is higher than in the previous case. Here, the bias correction is clearly valuable, resulting in a very good fit at the centre of the distribution (Figure 3.9d, $R^2 = 94.3\%$). Comparing the proportionality parameters of the tail distribution (ω_2) between Figures 3.8d and 3.9d reveals that the tail end of the distribution contributes relatively more than the peak: ω_2 always exceeds 0.5.

LCZ5 contains 6407 reference observations across 6 AAMS stations, and 30413 PWS observations after excluding missing data, across 18 PWSs. 23.4 % of the data is removed at QA level B (23290 records left); 43.5 % for QA level C (17176 records left). The AAMS data is similar to LCZ2, but with a longer tail end (observed wind speeds also exceed 15 km/h). Initial unfiltered PWS data is still focused on the strong peak at 1 km/h (Figure 3.9e) resulting in a poor fit ($R^2 = 30\%$) with the higher-tailed reference data. Removing the 1.0 km/h peak improves the fit

by shifting the distribution to the right, but still suffers from a large number of low wind speed observations around 2 km/h (Figure 3.9f and 3.9g). Applying the systematic bias correction strongly improves the fit to the AAMS data ($R^2 = 95.4\%$ with $c = 0.559$; $R^2 = 94.9\%$ with $c = 0.884$). The shape of the first Weibull component is rather curious owing to its low shape (a) parameter. A shape parameter very close to 1 approaches the exponential distribution (where $a = 1$) which carries a strong weight at the very low end of the distribution. The wider tail of LCZ2 is the most remarkable difference in wind speed distribution between the two LCZs, suggesting a higher likelihood of higher wind speeds, which could be caused by tunnelling through street canyons (Macdonald, 2000).

For LCZ5 we conclude that, as for LCZ2, the QA protocol strongly improves the raw PWS data to the point where it can provide an estimate of the wind speed distribution for a windy case. The poor result of Section 3.3.4 can therefore be attributed to the very low wind speed which makes up the bulk of the distribution, at which the Netatmo anemometer is not capable of measuring due to hardware limitations and the coarse output in integer km/h.

3.4 Discussion

3.4.1 Use of PWS data in other studies

The use of citizen science is not new in meteorology, and has even aided in the birth of the field (Eden, 2009). The emergence of PWSs has contributed to various studies, especially in urban areas lacking traditional measurements (Steeneveld et al., 2011; Wolters and Brandsma, 2012; Chapman et al., 2017; De Vos et al., 2017; Fenner et al., 2017; Chapman and Bell, 2018), but also at the level of national weather authorities (Krennert et al., 2018). While PWS data provides valuable insights into the under-sampled urban regions, the data remains of relatively low quality compared to WMO standards. The nature of the technique means there is little information regarding the setup of the station, which will not have been performed by an expert, and can thereby lead to substantial errors in the data. Additionally, the hardware itself can be a cause of error, such as unventilated screens heating up during the day to strongly overestimate outside air temperature (Bell et al., 2015; Meier et al., 2017). Our study shows that hardware can play a significant role in wind speed errors as well, here related to the effect of humidity and rain on the anemometer. For meaningful results, extensive quality checks need to be made, and have been constructed for PWS air temperature data (Meier et al., 2017; Napoly et al., 2018) and rain (De Vos et al., 2017). These account for the different error sources associated with crowdsourced PWS data by statistical checks and/or using data from surrounding (reference) stations to determine and filter out potentially erroneous data. For wind speed this is not applicable due to the microscale character of the measurements, but grouping stations into LCZs is an appropriate solution, reducing the impact of potentially wrong individual station data.

Using crowdsourcing or non-traditional data sources for wind measurements is comparatively rare: (Agüera-Pérez et al., 2014) utilise several data sources to construct a global wind field for Andalusia (Spain). For instance, the Spanish Meteoclimatic platform (<https://www.meteoclimatic.net/>) offers data, though without QA. Furthermore, the density of the network is scarce in comparison to our work (1 station every $\approx 450 \text{ km}^2$), and only some metadata such as measurement height is known. An interesting effort has been made to use mobile phones to measure wind, using an add-on device to create crowdsourced Handheld Wind Observations (Hintz et al., 2017), but the dataset is limited and seems focused towards the coast (mainly used

by windsurfers).

3.4.2 Causes of error in PWS wind

The greatest issue with crowdsourcing data is its lack of metadata (Muller et al., 2015). The aforementioned scarcity of studies exploring crowdsourced or otherwise unorthodox measurements can be ascribed to the large uncertainties inherent to (urban) wind measurements. Wind is strongly variable in time and space, more so than air temperature and rainfall. Especially the station setup is a crucial factor: air temperature measurements are affected by radiation from nearby walls or direct sunlight in the case of unshielded thermometers; rain is sensitive to the orientation and level of shelter of the station (which can cause underestimation of rainfall), but wind speed is affected not only by orientation (the sonic anemometer needs to be level), and strongly by shelter and obstacles (which induce turbulence and block the flow), but also the height of the measurement, which is unknown. Only the most basic location information is given by Netatmo, which does not contain any information regarding setup, calibration, or height of the measurements. Citizens have the option to provide more information about their station on the Wunderground platform (<https://www.wunderground.com/>), but this is entirely voluntary and is only provided by a few owners. PWS owners on rare occasions have their own weather website, which often provides very detailed information about the station, but these tend to be the higher-end stations such as Davis Vantage stations (Steeneveld et al., 2011; Wolters and Brandsma, 2012), not the cheaper Netatmo devices.

A large uncertainty lies within the Netatmo PWS itself: our bias analysis and experimental setup shows that the station has a tendency to underestimate relatively high wind speeds, and substantially underestimates very low (<2 km/h) wind speeds due to the coarse output resolution. We know that the measurement frequency is roughly 0.16 Hz and that these measurements are then aggregated into the ≈ 5 -minute output obtained through the API, so potentially the raw unprocessed data could provide a solution to the low wind speed errors. Other PWS brands might not suffer from the issues at low wind speeds combined with coarse output, which could make QA level B unnecessary.

Wind speed in the open field follows a logarithmic profile increasing with height: in the urban area the usual boundary-layer similarity theory is not valid, but wind still increases with height (Rotach, 1995; Castro, 2017; Kent et al., 2018). We expect the PWSs to be usually mounted at some height above the ground, at best 2m, but also on balconies, window sills, or wherever it is most suitable for the respective owner. A more ambitious weather hobbyist might install the anemometer on a pole or on top of a shed, for example, to better measure the actual wind. Stations installed on balconies or rooftops will give a completely different signal in time than a station in a sheltered garden. When installing the anemometer, the software offers the opportunity to report the station's height above ground. However, this information cannot be extracted from the Netatmo data obtained through the API, so we cannot check whether any correction towards a standard value (e.g. 10m) is performed prior to data storage.

3.4.3 Application & perspectives

Due to the uncertainty of station setup, the data obtained from individual PWSs cannot give an accurate representation of the wind speed climate: the station might be located on a balcony, a shed, near a wall, etc. A 1-to-1 time series for individual PWSs compared to their closest

matching AAMS stations displayed large deviations around the mean (not shown), even for daily averages, indicating the strong microscale character of the measured wind speed. The mixture Weibull distribution successfully captures the variability of wind within a city, and provides insight into the wind speed differences between neighbourhoods and LCZs. For a more quantitative temporal assessment of wind speed, PWSs appear not to be the right tool, and specialised measurement networks setup by professionals are still necessary. For instance, the effect of wind speed on thermal comfort during heat waves is better researched using stations situated in street canyons with a higher accuracy (such as the AAMS network): the effect of wind on thermal comfort is substantial, and strongly dependent on the local conditions (Heusinkveld et al., 2017). While a technique like Generalised Extreme Value statistics of thermal comfort like in Steeneveld et al. (2011) would be possible using the substantial length of the crowdsourced dataset, this is usually not applied by policy makers, but could be in the future with increasing number of PWS in cities worldwide.

The use of a probability density distribution is very common in the field of wind energy, where these functions are used to assess the wind energy potential of given areas (Celik, 2004; Carta et al., 2009; Drew et al., 2013). It would therefore be interesting to research the value of the crowdsourced wind data for identification of possible urban wind energy generation given the difficulties in estimating urban wind resource (Walker, 2011). A potential issue here could again be the unknown height of the stations: a (Weibull) PDD is typically constructed at a certain level and transposed to the hub-height of the expected wind turbine. A possible way to circumvent this issue is to assume a likely range of heights (between 2 and 5 m seems the most logical given the residential character of the PWS locations, though balcony stations can be much higher) and construct the transposed PDDs for the minimum and maximum height to see the spread in the results. GIS data on building height could also provide some information for rooftop and balcony stations.

The sensitivity of the station to setup errors (i.e. tilting, sheltering, measurement height) needs to be investigated in a systematic way, as Bell et al. (2015) did for air temperature. Wind tunnel experiments using several Netatmo anemometers could investigate the influence of the angle of tilt of the station on the reported wind speed values, and the threshold wind speed when the measurements are of sufficient quality. A long-term urban experimental setup of the Netatmo anemometer next to a known reference station could give some insight into expected sensor drift and the cause of errors for the low urban wind regimes.

3.5 Conclusions

This study makes use of wind speed data of 60 Netatmo Personal Weather Stations (PWSs) collected between January 2016 and June 2018 in the area of Amsterdam, the Netherlands, as well as data from an urban reference network and two experimental setups. From these data, we have established a Quality Assurance (QA) protocol to filter incorrect data, remove the effects of rain and humidity, and correct for a systematic underestimation of wind speed measured by PWSs. The quality-controlled PWS wind speed data can be used to construct a mixture Weibull probability density distribution (PDD) which conforms to the reference network. The wind distribution of different Local Climate Zones (LCZs) can be investigated by aggregating the data from the stations in those LCZs. While we conclude that for extended periods with very low (<2 km/h) wind speeds, the QA protocol does not improve the raw data, and the bias

correction even deteriorates results, we acknowledge that other PWS devices might not suffer from hard- and software issues associated with the Netatmo PWS anemometer, which needs to be investigated in further studies.

Based on the results obtained, we conclude that Netatmo PWS wind speed data is useful under the following conditions:

- The record is of sufficient length ($>\approx 2$ months) to have a large amount of data and document meaningful probability density distributions.
- The mean wind speed in this period is not low (>2 km/h): inherent issues with the Netatmo hardware induces substantial errors at low wind speeds, and the output of the stations in integer km/h only increases the relative error made.
- External WMO data of rain and humidity is available to apply the QA protocol, which filters out rain and high relative humidity ($RH > 95\%$) events. Humidity impacts the sonic anemometer and reduces its quality. Humidity and rain data could also be collected from (QA controlled) PWSs.
- The research in question is interested in the distribution of wind, rather than the wind at one given moment in time or space.

Acknowledgements

The authors wish to acknowledge all contributors to the Netatmo web platform, without whom the vast dataset this article uses would not have existed. We thank Lotte de Vos (KNMI) for her help with the Netatmo API setup. We acknowledge the Amsterdam Institute for Metropolitan Solutions for their financial support of the AAMS observational network. Arjan Droste and Gert-Jan Steeneveld acknowledge funding from the Netherlands Organization for Scientific Research (NWO) VIDI Grant “The Windy City” (file 864.14.007). Daniel Fenner acknowledges funding from the Deutsche Forschungsgemeinschaft (DFG) under Grant No. SCHE 750/15-1. Hartmut Küster, Ingo Suchland, Fred Meier, and Carl Benz are thanked for support with the comparison measurements at the Chair of Climatology at the Technische Universität Berlin.

Supporting Information

Table 3.1: Used AAMS reference network station details.

Station code	latitude	longitude	LCZ
2194	52.3687	4.88883	2
2195	52.3965	4.9579	5
2198	52.371	4.90615	2
2199	52.3173	4.87797	5
2221	52.3538	4.94017	6
2223	52.3591	4.82495	5
2225	52.3667	4.87053	2
2226	52.3558	4.99615	5
2227	52.378	4.89422	2
2228	52.3922	4.94083	3
2229	52.3715	4.89677	2
2230	52.3585	4.8621	2
2231	52.3632	4.93825	2
2235	52.3567	4.91572	2
2236	52.3505	4.89412	2
2237	52.3493	4.86987	2
2238	52.3777	4.8832	2
2239	52.3554	4.78377	6
2240	52.3367	4.87368	4
2241	52.3779	4.92937	G
2245	52.3745	4.89887	2
2246	52.3884	4.92737	5
2247	52.2942	4.97872	5





4

Chapter 4

Hydrometeorological monitoring using opportunistic sensing networks in the Amsterdam metropolitan area

The ongoing urbanisation and climate change urges further understanding and monitoring of weather in cities. Two case studies during a 17-day period over the Amsterdam metropolitan area, the Netherlands, are used to illustrate the potential and limitations of hydrometeorological monitoring using non-traditional and opportunistic sensors. We employ three types of opportunistic sensing networks to monitor six important environmental variables: (1) air temperature estimates from smartphone batteries and personal weather stations; (2) rainfall from commercial microwave links and personal weather stations; (3) solar radiation from smartphones; (4) wind speed from personal weather stations; (5) air pressure from smartphones and personal weather stations; (6) humidity from personal weather stations. These observations are compared to dedicated, traditional observations where possible, although such networks are typically sparse in urban areas. First we show that the passage of a front can be successfully monitored using data from several types of non-traditional sensors in a complementary fashion. Also we demonstrate the added value of opportunistic measurements in quantifying the Urban Heat Island (UHI) effect during a hot episode. The UHI can be clearly determined from personal weather stations, though UHI values tend to be high compared to records from a traditional network. Overall, this study illustrates the enormous potential for hydrometeorological monitoring in urban areas using non-traditional and opportunistic sensing networks.

This chapter is published as:

¹ de Vos, L., A. Droste, M. Zander, A. Overeem, H. Leijnse, B. Heusinkveld, G. Steeneveld, and R. Uijlenhoet, 2019a: Hydrometeorological monitoring using opportunistic sensing networks in the Amsterdam metropolitan area. *Bulletin of the American Meteorological Society*, BAMS-D-19-0091.1, doi: 10.1175/BAMS-D-19-0091.1

¹ The first two authors contributed equally to this work and share co-first authorship

4.1 Introduction

Traditionally, hydrologists and meteorologists, scientists and practitioners alike, have relied on dedicated measurement equipment in their research and operations. Such instruments are typically owned and operated by governmental agencies. Installed and maintained according to (inter)national standards, they offer accurate and reliable information about the state of environment we study, monitor and manage. Standard instruments are often based on novel measurement techniques that originate in the research community and have been tested extensively during dedicated field campaigns.

Unfortunately, the operational measurement networks available to the hydrometeorological community today often lack the required spatial and/or temporal density for high-resolution monitoring or forecasting of rapidly responding environmental systems. Apart from the high installation and maintenance costs of such dedicated networks, it can be challenging, if not impossible, to install meteorological monitoring instruments according to the official requirements in urban areas (Oke, 2006).

Yet, sensors are omnipresent in our environment nowadays, often related to the rapid development in wireless communication networks (e.g. McCabe et al., 2017; Balsamo et al., 2018; Tauro et al., 2018; Zheng et al., 2018). To make use of such opportunistic sensors could be greatly beneficial to (meteorological) science and environmental monitoring and management operations. Opportunistic sensors are devices that were not installed with the intention to generate large-scale observations, but can be used as such. They may not be as accurate or reliable as the dedicated equipment we are used to, let alone meet official international standards. However, they are typically available in large numbers and are often readily accessible online. Hence, combined with smart retrieval algorithms and statistical treatment, opportunistic sensors may provide a valuable complementary source of information regarding the state of our environment.

This article surveys recent opportunistic sensing techniques in meteorology, from (1) rainfall monitoring using commercial microwave links (CML) from cellular communication networks, via (2) crowdsourcing urban air temperature, pressure and solar radiation using smartphones to (3) high-resolution urban monitoring of air temperature, pressure, humidity, wind speed, and rainfall using personal weather stations (PWS). Other opportunistic sensing examples are: using security cameras as rainfall indicators (Allamano et al., 2015), rainfall information from sensors in driving cars (Rabiei et al., 2013), deriving the UHI from measurements of gradients of shallow groundwater (Buik et al., 2004), using fiber-optic cables (Bense et al., 2016), using aeroplanes to measure upper-air wind and temperature (de Haan, 2011), using hot-air balloons to measure boundary-layer winds (de Bruijn et al., 2016), smartphone anemometers (Hintz et al., 2017), or using networks of solar panels for radiation monitoring. Muller et al. (2015); Zheng et al. (2018) provide excellent overviews of past and ongoing projects making use of opportunistic sensing techniques, and USAID (2013) showcases practical applications of crowdsourcing projects for agricultural purposes in Africa. We limit ourselves to the presented techniques since these are relatively established even in developing countries, discussed in detail in literature, and observe near the Earth's surface.

We present a 17-day analysis for the Amsterdam metropolitan area, the Netherlands, where these opportunistic sensors are employed in a complementary fashion, in particular to provide detailed monitoring (both time series and spatially) of the passage of a front, as well as to

demonstrate the potential of opportunistic sensors to quantify the Urban Heat Island (UHI) effect. This study aims to showcase the availability of several opportunistic sensing techniques and their ability to capture meteorological events.

4.2 Sampling techniques

4.2.1 Traditional sensing methods

We use three traditional data sources as reference for the opportunistic sensing observations: a gauge-adjusted radar product; the WMO station at Amsterdam airport; and the Amsterdam Atmospheric Monitoring Supersite (AAMS (Ronda et al., 2017)) urban network.

4.2.2 Gauge-adjusted radar dataset

The Royal Netherlands Meteorological Institute (KNMI) operates two C-band Doppler weather radars. The 5-min reflectivity data from these radars are combined into one composite using a weighing factor as a function of distance from the radar. Beekhuis and Mathijssen (2018) provide detailed characteristics on the radars and the processing of their data. Reflectivity factors Z ($\text{mm}^6 \text{m}^{-3}$) are converted to rainfall intensities R (mm h^{-1}) with a fixed Z - R relationship (Marshall et al., 1955), $Z = 200R^{1.6}$, and, subsequently, accumulated to rainfall depths for different durations. The two KNMI rain gauge networks are employed to adjust the radar-based accumulated rainfall depths: an automatic network with 1-h rainfall depths for each hour (~ 1 station per 1000 km^2) and a manual network with 24-h 08:00–08:00 UTC rainfall depths (~ 1 station per 100 km^2). A daily spatial adjustment utilizing the manual gauge data is combined with an hourly mean-field bias adjustment employing the automatic gauge data. The resulting gauge-adjusted radar rainfall dataset has a spatial resolution of 0.9 km^2 , with no missing data for the study period. Overeem et al. (2011) provide a more detailed description of this radar dataset, which largely uses the methodology developed by Overeem et al. (2009a,b). Finally, 15-min path-averaged rainfall intensities are derived from the radar pixels covering each link path of the CML dataset (described in Section 3.b.2). The gauge-adjusted radar rainfall dataset is used as a reference to validate rainfall estimates from CMLs and PWSs.

4.2.3 WMO station Amsterdam airport

The WMO station Amsterdam airport, WMO 06240 (4.78°E , 52.32°N ; Figure 4.1a) provides hourly air temperature and cloud cover observations. This surface synoptic station is operated by KNMI, situated in a polder (4.18 m below MSL) and surrounded by meadows, arable land, and buildings as well as infrastructure from Amsterdam airport. Air temperature is observed at 1.5-m height above short mowed grass. The sensor is covered by a radiation screen and well ventilated. Cloud cover aloft is obtained from a LD40 ceilometer, which uses LIDAR to detect the height and concentration of particles, such as cloud droplets. KNMI (2000) provides more information on the temperature observation.

4.2.4 Amsterdam Atmospheric Monitoring Supersite

As an urban reference network we utilise the observations from the Amsterdam Atmospheric Monitoring Supersite (AAMS; Ronda et al. (2017)), which consists of 30 weather stations across

the city. The network consists of temperature and humidity sensors (Decagon VP-3, U.S.A.) mounted inside a 184 mm aspirated radiation shield (Davis, U.S.A.). The ventilation fan is powered by 2 small solar panels mounted on top of the shield. The fans work at global radiation levels $>100 \text{ W m}^{-2}$. The radiation screens are mounted onto lantern posts using a boom to mount the centre of the radiation screen 0.46 m away from the edge of the lantern post at a height of 4.0 m above ground level. The sonic anemometer (Decagon DS-2, U.S.A.) has an accuracy of 0.30 m s^{-1} or 3% (whichever is larger). The anemometers were mounted above the radiation screens 0.50 m away from the lantern post edges and at heights of 4.30 m (from ground level to centre of the anemometer).

4.2.5 Opportunistic sensing methods

Smartphone data

Smartphones contain many sensors to support their functionality, including sensors for light levels to adjust screen brightness, pressure sensors to complement the GPS for an accurate (vertical) location estimation, and thermometers for the battery to avoid damage from overheating. Readings from such sensors can be used for opportunistic environmental sensing by collecting them through mobile applications ('apps'). These apps sample the sensor readings with a certain frequency, along with the last stored GPS coordinates. Examples of apps that collect and store smartphone sensor readings include Pressurenet (<http://www.cumulonimbus.ca/>) (Mass and Madaus, 2014; Madaus and Mass, 2017), OpenSignal (<https://opensignal.com/>), and Atmos (Niforatos et al., 2014, 2017).

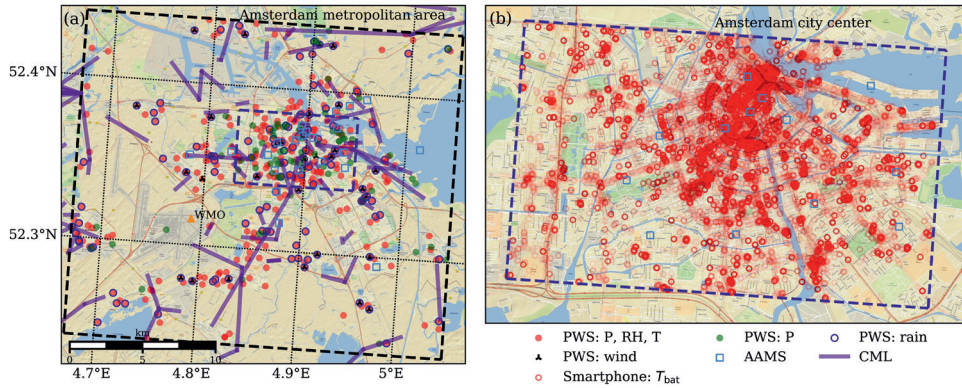


Figure 4.1: Map of Amsterdam metropolitan area and city centre with locations of all sensor networks: Personal Weather Stations (PWS) and Commercial Microwave Links (CML) (a), and of smartphone battery temperature readings and Amsterdam Atmospheric Monitoring Supersite (AAMS) stations (b).

Mass and Madaus (2014); Madaus and Mass (2017); McNicholas and Mass (2018) show that assimilating smartphone pressure data into NWP models improves representation of convective events. Likewise, Hintz et al. (2019) show for a case in Denmark that assimilating smartphone pressure observations decreased the surface pressure bias in a NWP model. Different quality

control methods were applied: for Madaus and Mass (2017) the raw smartphone pressure readings were filtered to only include one value per smartphone per assimilation time step, and were also corrected for the terrain elevation and checked for spatial and statistical consistency. Hintz et al. (2019) additionally use a consistency check with synoptic observations. In McNicholas and Mass (2018) a machine learning algorithm was used to remove outliers. Niforatos et al. (2017) compared smartphone light sensor readings with manually reported classifications of weather, which showed light readings to be indicative of present weather conditions.

City-wide air temperatures can be estimated from smartphone battery temperature readings, as has been shown for eight major cities (Overeem et al., 2013b), for the city of Birmingham (Muller et al., 2015) for daily temperatures, and for São Paulo for hourly and daily temperatures (Droste et al., 2017). Statistical training with independent temperature measurements was performed based on a steady-state heat transfer model: a smartphone is typically carried close to the user's body. The thermal energy generated by the smartphone must be balanced by heat exchange to the body and the environment. The conductive heat flow between two adjacent systems is assumed to be proportional to their temperature difference, and depends on the thermal insulation between smartphone and environment, and between smartphone and body. This principle allows us to estimate hourly-averaged air temperatures from hourly-averaged battery temperatures (Overeem et al., 2013b):

$$\bar{T}_{e,j,h}^{A,\text{hour}} = m_j^h (\bar{T}_{\text{bat},j,h}^{A,\text{hour}} - T_0) + T_0 + \epsilon_{j,h}, \quad (4.1)$$

where $\bar{T}_{e,j,h}^{A,\text{hour}}$ is the hourly mean urban air temperature, $\bar{T}_{\text{bat},j,h}^{A,\text{hour}}$ is the hourly-averaged battery temperature (both in space A and time), and T_0 a constant equilibrium temperature. m_j^h is a coefficient, $\epsilon_{j,h}$ is a random disturbance, and h denotes the hour.

In this study we build upon a large dataset of observations obtained from the Android application OpenSignal, which crowdsources data relevant to wireless connectivity along with the aforementioned sensor readings. Compared to the previously mentioned studies, readings were obtained at a far higher frequency, i.e. 15-s intervals whenever the smartphone screen is active, not requiring the app to be opened by the user. A total of 3.14 million smartphone observations are available for the entire study period for the Amsterdam metropolitan region (larger domain in Figure 4.1a).

The OpenSignal dataset includes self-reported accuracy scores (1, 2 or 3) of the light and pressure readings, as determined by the sensor management software in the smartphones (Android, 2019). Only readings with the highest possible accuracy were included in our analysis. All smartphone pressure sensor readings below 950 hPa are excluded, based on the lowest recorded pressure in the Netherlands, 954.2 hPa (de Haij, 2009), which results in a dataset of 2.06 million pressure readings. Light sensor readings above 0 lux are taken into account, leaving 2.32 million light readings in the whole study period. We only include battery temperature readings between 10 – 47°C when the smartphone is not charging: 0.4 million temperature readings within the city centre. Hourly battery temperature readings are averaged spatially over the city centre domain (Figure 4.1b), light and pressure are averaged over the entire region for each hour (Figure 4.1a).

Ambient air temperatures are estimated from battery temperature (Equation 4.2.5); the value of equilibrium temperature (T_0) as optimized by Overeem et al. (2013b), 39°C, is used. Figure 4.1b shows the positions of the underlying battery temperature readings. Two different datasets

are derived: one without and one with optimising the coefficients of the heat transfer model for the available dataset. The first dataset uses a fixed value of m_j^h for all hours, 2.4, as found for a summer period in London based on daily averages (Overeem et al., 2013b). These results, without further model calibration, are presented in Figure 4.2b, which also shows the 25th and 75th percentile. For the second dataset, records from 1 June 00:00 UTC – 15 June 00:00 UTC are employed to calibrate a value of m_j^h for each clock-hour (24 in total, ranging between 2.0 to 2.6). These optimised values, found using a least squares regression, are applied to the validation dataset from 16 June 00:00 UTC – 23 June 00:00 UTC.

A smartphone light sensor measures illumination in lux (lumen m^{-2}), i.e. irradiance weighted for the visible part of the electromagnetic spectrum, so a measure for the perceived brightness for the human eye. To estimate the equivalent solar radiation, we use an empirical factor of 0.0079 lux per W m^{-2} , based on the spectral distribution of sunlight (Chua, 2009). By applying this transformation, the readings are treated as if they were measurements of solar radiation. This is a fairly strong assumption, as we expect that most readings will not be made in a representative manner: with the smartphone perpendicular towards the Sun and in direct sunlight. User behavior plays a large role (e.g. indoor versus outdoor measurements), so one may expect that most light readings will underestimate the solar radiation, resulting in a skewed distribution. A light sensor in a smartphone has a limited view angle ($<180^\circ$) and has a relatively poor cosine response. Additionally, the sensor can over-saturate at high light intensities (the sensor limit is typically around 200 W m^{-2}). Therefore it is desirable to have many readings to increase the probability of observations taken in favorable conditions (unshaded and perpendicular to direct sunlight).

Because smartphone measurements are taken when the smartphone is used, most data is available for those times where people are active. Since hundreds of smartphone measurements are required to obtain a good signal of air temperature (Droste et al., 2017), spatial detail is limited to Local Climate Zone (LCZ, (Stewart and Oke, 2012)) scale at best, and temporal resolution to roughly hourly. The data at this availability is useful to get a broad overview of urban temperature, but not for (spatially) detailed studies.

Commercial microwave links

Cell phone communication relies on a telecommunication link network that consists of transmitting and receiving antennas, typically several km apart, between which radio signals propagate. Telecom operators commonly use signal frequencies that are sensitive to hydrometeors. This causes attenuation of the microwave link signals when liquid precipitation occurs between the antennas. Upton et al. (2005) first suggested to use signal attenuation in CML networks, which is typically monitored for quality control purposes, to determine rainfall. Soon after, this was shown to be successful with actual CML data (Messer et al., 2006; Leijnse et al., 2007). This was promising as microwave link networks are widespread, also in areas of the world with limited to no traditional rainfall sensors. Subsequent research has focused on improving the techniques to obtain accurate rainfall estimates from these datasets, (e.g. Leijnse et al., 2008; Zinevich et al., 2008; Overeem et al., 2011; Chwala et al., 2012) and produce rainfall maps (Overeem et al., 2013a, 2016b) with real-time applicability (Chwala et al., 2016; Andersson et al., 2017; Chwala et al., 2018). Comprehensive overviews of literature on this technique were provided by Messer and Sendik (2015), Uijlenhoet et al. (2018), and Chwala and Kunstmann (2019). Several tools have

been developed, documented and made (freely) available for users to construct rainfall observations with CML data: ‘Rcmlrain’ (<https://github.com/fenclmar/Rcmlrain>), ‘RAINLINK’ (<https://github.com/overeem11/RAINLINK>), ‘pycomlink’ (<https://github.com/pycomlink/pycomlink>), and ‘pySNMPdaq’ (<https://github.com/cchwala/pySNMPdaq>).

The relation between rainfall attenuation and rainfall intensity can be described with a power law between path-averaged specific signal attenuation (k in dB km⁻¹) and link path-averaged rainfall intensity (R in mm h⁻¹) (Atlas and Ulbrich, 1977):

$$R = ak^b \quad (4.2)$$

where

$$A = \text{TSL} - \text{RSL}; \quad k = \frac{A_{\text{wet}} - A_{\text{dry}} - A_a}{L} \quad (4.3)$$

Coefficients a (in mm h⁻¹ dB^{-b} km^b) and b (-) are dependent on signal frequency and polarisation (Olsen et al., 1978; Jameson, 1991). TSL and RSL are the transmitted and received signal level (dB) respectively, A_a is the attenuation due to wet antennas (dB) assumed as a fixed value, A_{wet} and A_{dry} are the attenuation under wet and dry weather conditions respectively (dB) and L is the length of the link path (km). The specific attenuation due to rainfall is what remains when the attenuation due to other causes (i.e. dry weather conditions and wet antennas) are subtracted.

The time series shown in Figure 4.2d originate from the T-Mobile CML network visualised in Figure 4.1a. Between 6 June 00:00 UTC and 10 June 14:00 UTC, 74 links were operational in the study area. Power levels were instantaneously sampled every 15 min. Due to data transfer issues, no power levels were available at the end of the study period. Rainfall time-series for each link were constructed with the open source package RAINLINK (Overeem et al., 2016a), using the approach and optimised parameters from (de Vos et al., 2019c). The wet antenna attenuation makes up a larger fraction of the total attenuation for short links, meaning that a small error in A_a , a constant, will result in a relatively large error in k for short links, and the effect on the estimated value of R would subsequently be larger than for long links given the same error in A_a .

Crowdsourced personal weather stations

PWSs allow anyone to measure weather variables in their direct environment. Many automatic PWSs can upload their measurements directly to online platforms where they can be visualized and shared. Weather Underground (<https://www.wunderground.com/wundermap>), WOW-NL (<https://wow.knmi.nl/>) and the Netatmo Weathermap (<https://weathermap.netatmo.com/>) are examples of platforms where weather observations are visualized in real time. Ideally, weather variables can be crowdsourced from such platforms in far higher spatial and temporal resolution than from traditional sensor networks.

The devices are often low-cost with a lower expected measurement accuracy than typical sensors from meteorological institutes. The PWSs are installed by citizens without expert knowledge on sensor placement requirements and/or lacking available measurement site without interference from surroundings. Hence we expect that many of the PWSs generate compromised measurements. For tipping bucket rain gauges, obstructions (e.g. insects, twigs) and the device not being

completely level with the ground, could hinder the tipping mechanism. A shielded location will also lead to underestimation of rainfall. Overestimation of rainfall can result from PWS owners cleaning or handling the device, resulting in tipping bucket tips, creating measurements of artificial rain. PWS wind measurements are also largely affected by their position in relation to obstacles and the shielding effect of buildings. Furthermore, PWSs with a sonic anemometer are sensitive to rain blocking the path of the sound waves, so data quality might be compromised during rain events. Urban wind is highly variable in space, and is often measured as profile using e.g. LIDAR (Drew et al., 2013), so spatial averaging of PWS wind measurements is needed to obtain useful data. Temperature readings are highly affected by direct radiation: the lack of a proper radiation screen in most PWSs can result in overestimation of temperature by several degrees when positioned in direct sunlight (Bell et al., 2015; Chapman et al., 2017). Finally, the updates of measurements to the platform can be infrequent, and connectivity problems will result in large gaps in the time series.

Only a few studies compared PWSs with high-end sensors; temperature, relative humidity, radiation, pressure, rainfall, wind speed and direction: Jenkins (2014); Bell et al. (2015), temperature: Meier et al. (2015), rainfall: De Vos et al. (2017). Other studies have benefited from available PWS temperature records in cities. The UHI is then defined as the difference between PWS temperatures and a rural reference station (Meier et al., 2017; Chapman et al., 2017; Fenner et al., 2017; Golroudbary et al., 2018; Napoly et al., 2018). Preliminary work has been performed on crowdsourced wind (Droste et al., 2018a) and rainfall measurements (De Vos et al., 2017; Golroudbary et al., 2018; Chen et al., 2018) (and explored with simulated PWS rainfall measurements by de Vos et al. (2018) as well). In other studies code has been developed and made available to apply quality control on crowdsourced PWS data (the CrowdQC R-package for PWS temperature observations <https://depositonce.tu-berlin.de/handle/11303/7520.3> and TITAN <https://github.com/metno/TITAN/>, and code to filter crowdsourced rainfall observations PWSQC <https://github.com/LottedeVos/PWSQC>).

Measurements from all personal weather stations from the brand Netatmo in the Amsterdam study area (Figure 4.1a) are evaluated. All devices measure temperature, pressure and humidity. Additionally, rain and/or wind are measured in case those optional modules are installed for that PWS. In order to standardize the variable time intervals, all measurements are attributed to the timestamp of the 5-min interval in which it occurred. If multiple measurements occurred within the 5-min interval they are averaged (or accumulated in case of rainfall). The measurements over the study period are shown in Figure 4.2 (panels (a), (c), (f) and (h)), where panel (i) indicates the dewpoint depression (*DPD*) as calculated from the temperature and humidity measurements from the PWS. No QC treatment is applied on the PWS data to showcase the raw potential. *DPD* is here preferred over dewpoint temperature itself to identify the frontal passage.

4.3 Case selection & study area

We selected Amsterdam (capital of the Netherlands) and its surroundings and the period between 6 June 2017 00:00 UTC and 23 June 2017 00:00 UTC as case study period (local time is UTC+2 hours). This period contains both sufficient data from opportunistic sensing techniques, and interesting meteorological events to illustrate the potential of the opportunistic sensing techniques. The selected region is bound by 4.67–5.05°E & 52.24–52.44°N (26 km × 22 km). To be able to distinguish between the inner city and suburbs, the study area was divided into two

parts, i.e. the urban centre dataset: 4.83–4.95°E. & 52.34–52.38°N and the suburban dataset 4.67–5.05°E. & 52.24–52.44°N, excluding the urban centre area (Figure 4.1a).

The Netherlands has a temperate maritime climate (Köppen Cfb). With a mean temperature of 18.0°C and 50.5 mm of rainfall June 2017 was about 2.5°C warmer and 10.5 mm drier than the climatological mean (based on the past 30 years of observations at station WMO 06240 Amsterdam airport, henceforth referred to as "Amsterdam airport"). The month had eight summer days and two tropical days (max. temp. above 25 & 30°C respectively).

On June 6, a small low-pressure system developed over the North Sea off the coast of the Netherlands and passed over the country, resulting in a substantial pressure drop to 992 hPa, an hourly maximum wind speed of 54 km h⁻¹ (7 Bft) and 12 mm of rainfall measured at Amsterdam airport. In the morning of June 9, an active cold front brought in relatively cold air which resulted in 27 mm of rainfall. A clear-sky episode occurred 9–11 June, while another cold front passed in the early morning of June 12 (Figure 4.2i). In the following period, no rainfall occurred, and temperatures were mild (daily maximum temperatures below 25°C), followed by a warm episode between June 16 and June 19. On June 19 the maximum air temperature reached 29.8°C at Amsterdam airport. This warm episode ended with the passage of a cold front and associated rainfall and thunderstorms on June 22. For the remainder of the paper we will focus on two cases, i.e. case A, describing the passing front and resulting rainfall at the start of the study period, and case B, containing the hot summer period, with a focus on UHI detection.

For this study, the UHI is defined as the instantaneous urban air temperature difference between the city and the countryside (Stewart, 2011). The UHI develops as a result of the relatively low albedo of cities, high heat capacity of the urban fabric, thermal radiation trapping, and low surface evapotranspiration. The UHI is favored by weather conditions with high solar insolation (low cloud cover) and low wind speeds (Oke, 1982; Theeuwes et al., 2017). Earlier crowdsourcing observations indicated that Dutch urban areas experiences a mean daily maximum UHI of 2.3°C and the 95th percentile amounts to 5.3°C (Steenefeld et al., 2011). Ronda et al. (2017) found a mean evening UHI of $\sim 1^\circ\text{C}$, and a maximum of 4.5°C in Amsterdam for the summer of 2015 as a whole.

4.4 Results

4.4.1 Case A: Weather front

First we focus on the passage of a cold front over the study area on June 9. At 6:00 UTC the operational model analysis provided by KNMI locates the frontal zone to the west of Amsterdam (not shown), and by 12:00 UTC the front has passed the city. Prior to the frontal passage itself, an upper air disturbance passed over Amsterdam between 3:00 and 4:00 UTC, bringing strong convection and rainfall. Such frontal zones cause distinctive behavior in various meteorological variables, which we expect to be distinguishable in the crowdsourced data (Figure 4.2).

The passage of the front is clearly visible in the observed *DPD* and the wind speed (Figure 4.3a). The *DPD* steadily drops during the approach and passage of the cold front, reaching a minimum of 1.4°C at 9:00 UTC. Between 10:00 and 11:00 UTC, when the front has passed, the *DPD* increases again up to 6.8°C, indicating the cold and dry air mass brought in by the cold




Figure 4.2: Time series of opportunistic measurements of weather variables. PWS observations of temperature (a), cumulative rainfall (c), wind (f), pressure (g) and dewpoint depression calculated from humidity and temperature (i), smartphone battery derived air temperature (b), light (e) and pressure (g). The colored areas indicate the interquartile range (IQR is mean 25–75 percentile) of all observations at that time; lines show the median values, except for (b) where the line shows mean temperature. Shaded areas indicate night-time. Pearson correlation (r), standard deviation of the difference (SD) and absolute bias (bias) are calculated based on hourly values compared with WMO observations at Schiphol.

front. Crowdsourced and reference wind speed steadily increase as the front passes (from 2 to over 4 km h⁻¹), before reaching its maximum (5 km h⁻¹) directly after the passage. The convection associated with the upper air disturbance at around 3:00–4:00 UTC generates a strong peak in the wind speed. Despite the unknown measurement setup of the PWS anemometers, the average signal of all PWSs corresponds well to that of the quality-controlled reference AAMS network (mean bias of 0.4 km h⁻¹), which shows the same behavior for the upper air disturbance and the front passing. However, the AAMS signal indicates a delayed onset of the wind speed increase (at around 9:00 UTC) and takes longer to reach a higher maximum wind speed.




Figure 4.3: (a) Hourly average wind speed measured by PWS (solid line) and AAMS (dotted line), as well as the PWS dewpoint depression (red line, right y-axis). (b) Hourly averaged air pressure measured by PWS (solid line) and smartphone (dashed line) on June 9. The vertical lines indicate the boundaries of the frontal passage. At 6 UTC the front is located to the west of Amsterdam; at 12 UTC the front has passed over the city.

The ambient air pressure (Figure 4.3b), measured by PWSs and smartphones, starts increasing at the moment the front passes (8:00 UTC). Typically, air pressure decreases before a cold front, rapidly increases during the passage, and increases at a slower rate afterwards. The expected drop prior to the frontal passage is not very pronounced in the measurements: there is a slight decrease in pressure between 0:00 and 2:00 UTC (1.7 hPa decrease for PWS; 3.5 hPa for smartphone). The latter is more likely associated with the upper air disturbance. After the

frontal passage at 8:00 UTC, the pressure rises, from 1006–1008 hPa (PWS–smartphone) up to a maximum of 1013–1016 hPa at midnight. The pressure tendency remains roughly 1 hPa hour^{-1} after the front has passed.

The light intensity as measured by smartphones shows a distinct diurnal pattern in Figure 4.2e, following the course of solar radiation. The measured data are strongly skewed, so the median light intensity values are low (Figure 4.2e). Figure 4.4 shows the 99th percentile of light readings to capture the readings made in the most favorable light conditions (see Section 2.b.1). The sky on June 9 is overcast (8 octas) until 11:00 UTC, at which time the front has passed over Amsterdam and the sky clears up to scattered cloudiness (Figure 4.4). The light intensity is also very low until 10:00 UTC, even though this is well within daylight hours. Compared to June 18 (a clear day) the light intensity is roughly halved, and the shape of the line is not as symmetrical (as we would expect from the diurnal cycle of global radiation). The green lines in Figure 4.4 indicate the other days over the study period, showing the strong variability in the daily course of light intensity. The light intensity measured by smartphones not only depends on incoming radiation, but also strongly on user behavior (indoors vs outdoors, the angle of the phone) and the type of light sensor in the smartphone, which can differ between brands. The light sensor may also be oversaturated during high light intensities, resulting in flattened peak values.




Figure 4.4: Hourly 99th percentile of smart phone radiation (green lines) with Amsterdam airport cloud cover in octas on June 9 (blue, circles) and June 18 (orange, diamonds). June 9 (blue solid line) and June 18 (orange solid line) are a cloudy and cloud-free day, respectively.

The light intensity peak at 4:00 UTC coincides with the the upper air disturbance seen in Figure 4.3, but is actually an artefact of the low number of observations. The number of available observations is higher during the day than during night and early morning, since it is related to

user activity whether the smartphone logs an observation (as detailed in Section 2.b.1). At 4:00 UTC there are only 502 smartphone observations, compared to 10,373 at 15:00 UTC (17:00 local time, the typical end of the working day), so the data is more sensitive to outliers.

The upper air disturbance, and subsequent frontal passage, of June 9 results in 27 mm rainfall as measured by the gauge-adjusted radar reference. Figures 4.2c and 4.2d show that the peak of rainfall occurs after sunrise, coinciding with the timing of the frontal passage. Figures 4.5a and 4.5b depict the cumulative rain over June 9, measured by CML and PWS, against the reference. Total amounts differ between the two methods, but both show the same time response. The relatively short links (< 2 km) overestimate rainfall, with the majority reporting > 30 mm rainfall (relative bias is 87%). The longer links (≥ 2 km) also tend to overestimate, but much less extreme (relative bias is 12%). Although the expected uncertainty in rainfall estimates is higher for short links, the larger systematic bias (54% relative bias, or 0.13 mm absolute bias, for all links) indicates that the methods to derive rainfall (RAINLINK) were not ideal for this rainfall event, especially for short links.

PWS measurements tend to underestimate the rain as measured by the reference, with some occurrences of large reported rainfall values that are not otherwise captured (Figure 4.5b). Nevertheless, the majority of PWSs seem to agree overall with the reference (Figures 4.5b and d). The spatial distribution of rainfall (Figure 4.6) measured by PWS and CML corresponds to that of the gauge-adjusted radar reference. We find that areas with high rainfall in the reference also yield high accumulations in the CML and PWS data in these areas. The overestimation by short links up to 8 mm is visible to the northwest of the band with high rainfall. The rainfall observations by PWSs correspond well to the spatial pattern of rain, although a number measure little (< 1 mm) rain during the hour represented in Figure 4.6. These stations are mainly clustered in the city centre. The large amount of obstructions inside the city centre could reduce the rainfall received by the stations, which may partly explain the underestimation tendency already seen in Figures 4.5b and d.

4.4.2 Case B: Urban Heat Island

The last days of the study period are characterized by high temperatures and generally clear, sunny weather, leading to higher urban temperatures (PWS median air temperature up to 30°C on June 19, Figure 4.7). Air temperature is measured by PWS, and derived through the smartphone battery temperature using the second, calibrated dataset (Section 2.b.1). The AAMS network serves as urban reference, and the Amsterdam airport measurements are used as rural reference for the UHI (Figure 4.7b). The smartphone-derived air temperature differs clearly from the PWS and AAMS measurements, with more erratic behavior and strong minimum values at night and early morning (as low as 7°C when the AAMS values are above 16°C). Figure 4.7c showcases this larger spread, also indicated by the large standard deviation (2.82°C compared to 0.66 and 1.08°C for the PWS). During daytime the smartphone-derived temperatures correspond better with the PWS and AAMS measurements than at night. The diurnal cycle is clearly visible: the low values at night are most likely due to a low number of measurements available, increasing the sensitivity to outliers. Despite these occasional large deviations, the bias amounts to -0.6°C compared to AAMS (Figure 4.7c), which is relatively small. A large positive bias (2.0°C) is found when a fixed literature value for m_j^h is used, for the time series shown in Figure 4.2b, whereas the other model statistics are mostly uninfluenced by optimizing m_j^h .

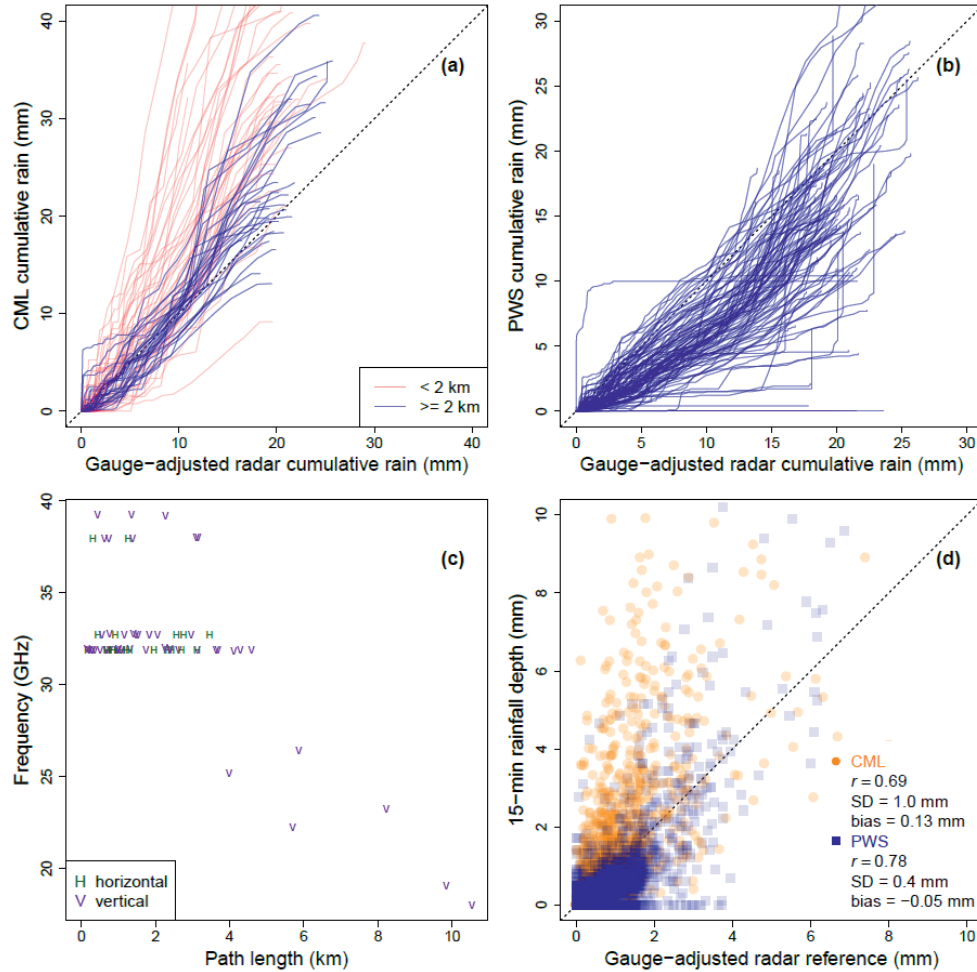


Figure 4.5: Double mass plots of commercial microwave link derived rainfall observations (a) and PWS rainfall observations (b) on 9 June against the reference of respectively the path-averaged and overlying pixel gauge-adjusted radar rainfall observations. Polarisation, path length and frequency of the CML network is given in (c) and (d) shows scatter plots of both together with their Pearson correlation (r), standard deviation of the difference (SD) and absolute bias (bias) at 15 min time steps.




Figure 4.6: Map of 60-min rainfall depths over the Amsterdam metropolitan area based on gauge-adjusted radar data (pixels; 100% availability), CML data (paths; only CMLs with 100% availability are shown), and PWS data (circles; only PWSs with at least 83.3% availability are shown).

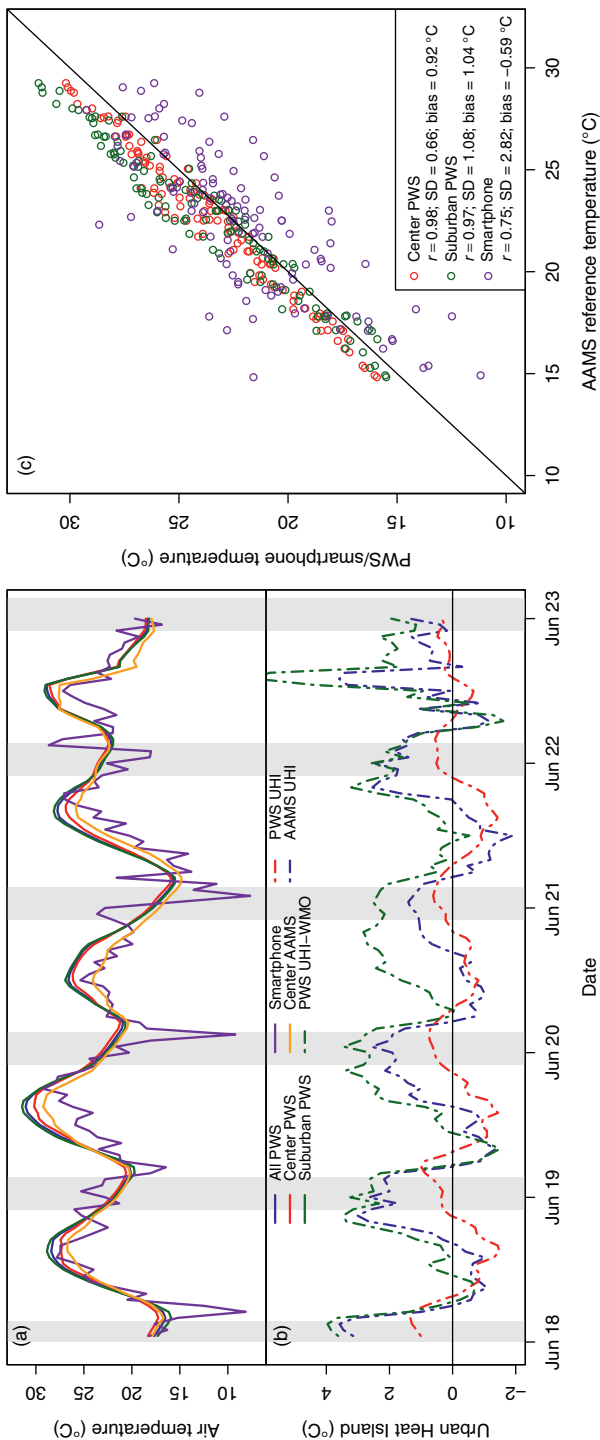


Figure 4.7: (a) Time series of temperature measurement according to the median of all (blue), centre (red) and suburban (green) PWSs, as well as the median of the AAMS network (orange) and the mean smartphone derived air temperatures (purple). (b) The difference between the median PWS centre and suburban temperatures (red, dashed); the AAMS and WMO (blue, dashed); and the PWS centre and WMO (green, dashed). Shaded areas indicate night-time. (c) Scatter plot of hourly median PWS and smartphone temperatures against median AAMS station data, with Pearson correlation (r), standard deviation of the difference (SD) and absolute bias (bias).

The diurnal pattern of air temperature between centre and suburban PWSs is similar, although the centre stations tend to be warmer at night, and colder during the day (Figure 4.7a). The suburban stations contain a higher spread and bias than the centre stations, though both show good agreement to the reference (Figure 4.7c). The AAMS air temperature is typically about 2 to 3°C lower during the day: this could partially be caused by the unknown setup of the Netatmo station, which is likely to be exposed to direct sunlight or close to walls, making it sensitive to radiation errors. Figure 4.7b depicts the UHI estimated by subtracting the centre and suburban PWS (red dashed line). This particular PWS–UHI shows spatial variability within the PWS data, which is most pronounced during daytime, where the difference can be up to -1.5°C (i.e. the centre is 1.5°C colder than the suburban area). Higher urban shading in the morning, and the faster heating rate of the relatively thin rural boundary layer compared to the deeper urban boundary layer cause this urban cool island in the morning (Theeuwes et al., 2015).

The other two UHI estimates are constructed using Amsterdam airport as rural background, showing that the city centre is indeed much warmer at night than the rural surroundings. Urban cool islands typically form in the morning, persisting for several hours before the city heats up more. A remarkable 6°C UHI peak is visible on June 22, in the afternoon (13:00–14:00 UTC). This seems to be mainly caused by the Amsterdam airport temperature, since the PWS–UHI (which has no true rural reference) shows a value close to 0°C at that time. This is visible in Figure 4.7a, where temperatures rapidly decrease in the course of a few hours on June 22 afternoon. Thunderstorms were reported on this day, and several mm of rain were measured at Amsterdam airport (according to radar) between 14:00 and 15:00 UTC. The UHI in this case is likely caused by the sudden cooling of the rural reference, rather than strong urban heating.

Figure 4.8 presents the spatial variability in the AAMS and PWS temperature recordings between 2:00 and 3:00 UTC on June 18, when the UHI is typically largest. The cluster of stations in the centre yields higher values than the suburban stations, although in both areas many stations deviate from this trend. The centre PWS report an average UHI of 4.0°C, the AAMS UHI is 3.6°C, whereas the suburban areas have an average UHI of 2.7°C. Variability between measurement sites is high: some stations report a temperature difference of up to 12.4°C, and even a few with negative UHI (up to -0.6°C).

4.5 Discussion and conclusions

4.5.1 General

We have shown that even though each technique has considerable limitations regarding accuracy, the data from opportunistic sources can be used to monitor meteorological phenomena. The potential of these techniques lies in the high spatial density of such observations, especially in urban areas.

We explicitly consider observations that can be obtained near-directly from the opportunistic sensors, without applying many correction schemes, to illustrate their inherent potential: validation using the available quality assurance schemes was not the aim of this research. We use temperature from smartphone batteries and personal weather stations (PWS), rainfall from commercial microwave links (CML) and PWS, solar radiation from smartphones, wind speed from PWS, air pressure from smartphones and PWS, and humidity from PWS. Two case studies in a 17-day period over the city of Amsterdam, the Netherlands, are explored. In the first case



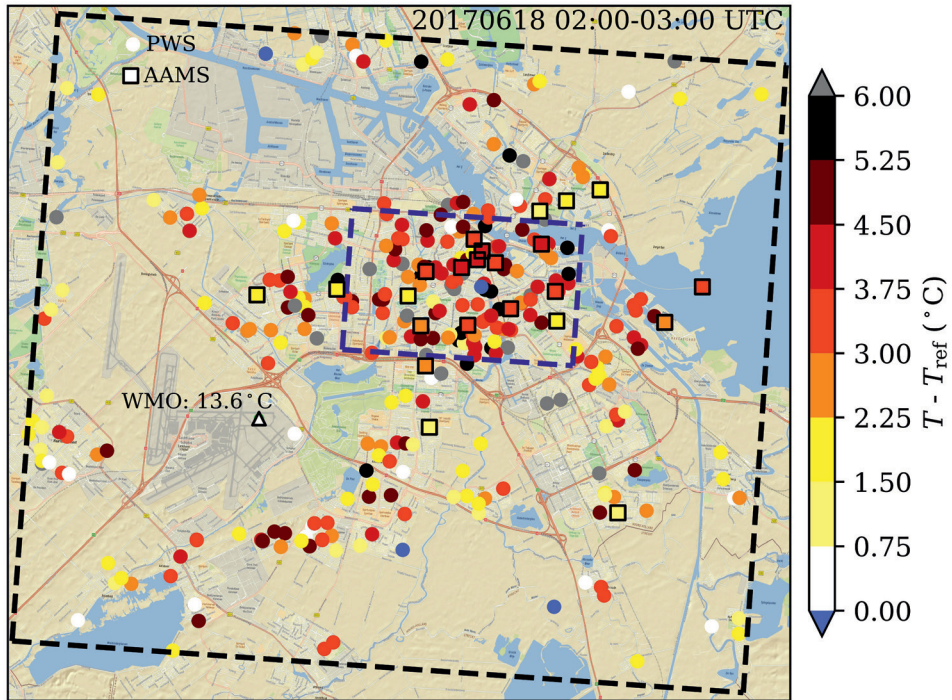


Figure 4.8: Urban Heat Island map for the Amsterdam metropolitan area (black rectangle) showing difference between hourly averaged air temperature for AAMS network (squares) and PWSs (circles) with respect to 1.5-m air temperature at WMO station Amsterdam airport observed at 3 UTC (triangle). The blue rectangle represents Amsterdam city centre, the remainder of the metropolitan area is suburban. Only stations with at least 80% availability are shown. Of the 309 PWSs only 4 are colder than WMO, at most 0.6°C, and 24 are at least 6.0°C warmer than WMO, at most 12.4°C.

study we show how the passage of a front is apparent from many of the data sources. The second case study shows that these measurements can be valuable in monitoring the Urban Heat Island (UHI) effect, especially given the fact that WMO stations in urban areas are very rare.

The passage of a cold front is visible in all of the studied opportunistic sensing data sources. The dynamics of the temperature (especially from PWS, less so from smartphones), rainfall, solar radiation, wind speed, air pressure, and humidity all show the passage of the front. However, not every aspect of the weather events is sufficiently captured by the data: techniques using smartphone observations can only estimate a variable as a spatial average over the city and cannot be used to describe detailed spatial variability. Also, the PWS wind observations were too noisy to describe spatial patterns in the city with confidence.

A rural reference station is needed to quantify the UHI, for which we use the WMO station Amsterdam Airport. Even the PWS locations outside the city centre of Amsterdam (suburban, see Figures 4.1 and 4.8) are mostly in built-up areas, and are hence expected to experience the UHI, although less severely. This is supported by Figure 4.7b, where the temperature difference between the city centre and the suburban areas is much less pronounced than when the WMO station is used as rural reference. Using a single rural reference to quantify the UHI, instead of multiple background stations, is a good practice when the main interest is in the intra-urban variability of temperature, as in this work and e.g. Fenner et al. (2017). However, a limitation of this practice is evident from the artificially high UHI we see in Figure 4.7 on June 22, which was caused by local cooling at the rural site. Finally, we note that the distinction between centre and suburban in this study was made rather crudely. In future UHI studies we recommend a more sophisticated partitioning of the stations into different classes (such as Local Climate Zones (Stewart and Oke, 2012)).



4.5.2 Temperature

The PWSs are suitable for monitoring the UHI. When compared to the AAMS urban reference network, these PWS show an UHI of the same order of magnitude (2–4°C), especially during the night. We observe urban cool islands during the period between sunrise and local noon. Air temperatures derived from smartphone battery temperatures exhibit much more noise than PWS temperatures, which limits their use for UHI measurement. We note that the PWS thermometers are not shielded from solar radiation or ventilated, whereas the AAMS are. This is clearly visible in Figure 4.7, where the PWS temperature (and derived UHI) increases much more quickly than the AAMS temperature. This corresponds to the findings of Bell et al. (2015).

4.5.3 Pressure

Both smartphones and PWSs provide good estimations of pressure. Pressure fields are relatively constant in time and space, and both opportunistic sensing techniques show ability to describe them.

4.5.4 Light

Light estimations derived from smartphones are highly variable in time. The indirect nature of the measurement and the typical suboptimal conditions during sampling result in merely an

indicative observation of light. Such measurements should only be considered in the absence of dedicated sensor observations and considered with caution.

4.5.5 Wind

Average wind speed from the PWSs are very low compared to what would be expected during the passage of a front. This may partly be due to how the PWS anemometers are installed. However, the carefully installed AAMS stations show average wind velocities of the same order of magnitude (around 5 km h^{-1}), indicating that the placement of the anemometers does not play a large role here. The reasons for these wind speeds to be lower than expected lies in the fact that the urban wind measurements are made at a lower level, and the urban fabric greatly reduces wind speeds at street level (Macdonald, 2000). This also means that wind speeds are expected to be highly variable across the city, which is clear from Figure 4.2f. Spatial averaging over the city is therefore needed in order to see clear signals in the wind speed. Spatial averages of wind speed show the same behavior between PWS and AAMS, indicating their use to measure the urban wind as a whole. Note that the wind sensor on the Netatmo PWS is a sonic anemometer, which are negatively affected by precipitation, hence wind observations during rainfall can be less reliable. This illustrates the need of a quality-control procedure which could improve overall data quality by filtering out precipitation events (Droste et al., 2018a).

4.5.6 Rainfall

Data from PWS and CML are shown to provide useful information on both rainfall amount and space-time variation. Their ability to show detailed variations in space and time makes them useful for qualitative use in rainfall monitoring. The CML network overestimated rainfall in case A (Figure 4.5a), although the relative bias of long links ($\geq 2 \text{ km}$) was 71% smaller than that of short links ($< 2 \text{ km}$). This is likely related to the larger error contribution wet antennas have for shorter links, and that the correction was calibrated on a different dataset, possibly with more long links (de Vos et al., 2019c).

The PWS also show a good agreement with the reference, although most stations underestimate rainfall (Figure 4.5b). This may be due to the higher wind speed above the urban fabric which could cause buildings to act as a shield for the PWS rain gauges. It is also apparent from Figure 4.5d that some PWSs report either zero rainfall when there is clearly rain or large amounts of rainfall where there was none. Such, and other errors could be corrected by using automated filters (de Vos et al., 2019b). The difference in accumulations between the city-averaged CML and PWS rainfall data (see Figure 4.2) is partly caused by the overestimation by CML. However, differences may also be due to the spatial variation of rain and the respective locations of the PWS and CML. Figure 4.6 shows that for the examined hour the CML are more abundant in high-rainfall areas, whereas the PWS are more clustered in the city centre, where less rainfall was observed.

The method used to derive rainfall estimates from CML data (RAINLINK) is one of many possible methods (see Section 2.b.2). Our dataset consists of instantaneously sampled CML data, which is more prone to errors than CML data obtained with other sampling strategies and/or more frequently than every 15 min (de Vos et al., 2019c).

4.5.7 Outlook

Our study shows that the research opportunistic sensing techniques all yield meaningful results. However, without quality control procedures, PWS data performs better than smartphone or CML measurements. The PWS sensors are designed to measure hydrometeorological variables and are less reliant on quality control than the indirect CML or smartphone observations. A thorough procedure which removes error sources will therefore be most effective for the CML and smartphone data, which can strongly improve with regards to the unfiltered signal. This may change in the future when the expected measurement density of smartphones increases and their hardware (sensor capability) improves. The observations contain large errors, as found by the larger spread in the data than would be explained by spatial or temporal variability. However, the opportunistic sensors provide information in time scales and areas that cannot be achieved with traditional sensing techniques.

Many PWSs are found in densely populated areas, where also many smartphones are operational. This is mainly true for urban areas in parts of the world where people have funds to invest in these devices (although smartphones are considered so important that they are essentially ubiquitous, independent of living standards). CML networks differ as well, in sampling strategy and frequency (which affects the accuracy of rainfall estimates) and in network density (depending on replacement by fiber optic technology). The availability of opportunistic sensing observations should be explored in order to judge their usefulness, especially as their accuracy heavily relies on the quantity of observations. Because traditional meteorological measurements are generally absent in urban areas, these new data provide a welcome addition. This is particularly important for monitoring the UHI, and wind and rainfall at street level. We therefore urge the scientific community to keep investigating new sources of data, and to study the uncertainties therein. In combination with reference networks of meteorological measurements or stand-alone, these new sources will provide much needed hydrometeorological information for citizens and scientists, in any part of the world.

Acknowledgements

Gert-Jan Steeneveld and Arjan Droste acknowledge funding from the Netherlands Organisation for Scientific Research (NWO) project 864.14.007. Lotte de Vos is financially supported by KNMI MSO project 2013.09. The reference stations of the Amsterdam Atmospheric Monitoring Supersite have been financially supported by the Amsterdam Institute for Advanced Metropolitan Solutions (AMS), which has also provided co-funding (project VIR17006). We thank James Robinson from OpenSignal (London, UK) for providing the smartphone dataset. We gratefully acknowledge Ronald Kloeg and Ralph Koppelaar from T-Mobile NL for providing the cellular telecommunication link data. All PWS owners who have contributed to the Netatmo dataset are greatly appreciated, as well as the Netatmo company for making the measurements available to us. See <https://dev.netatmo.com/resources/technical/reference/weatherapi/getpublicdata> for information on how to obtain Netatmo data. Datasets retrieved from third parties (CML and OpenSignal) are available upon request. The gauge-adjusted radar rainfall dataset is freely available in netCDF4 format, 'Radar precipitation climatology' via <http://climate4impact.eu> or in HDF5 format at the KNMI Data Centre https://data.knmi.nl/datasets/rad_n125_rac_mfbs_em_5min/2.0?q=radar. The AAMS data are available upon request (contact: bert.heusinkveld@wur.nl).



5

Chapter 5

Introducing the Urban Wind Island Effect

Wind is a key component of the urban climate due to its relevance for ventilation of air pollution and urban heat, wind nuisance, as well as for urban wind energy engineering. These winds are governed by the dynamics of the atmosphere closest to the surface, the atmospheric boundary layer (ABL). Making use of a conceptual bulk model of the ABL, we find that for certain atmospheric conditions the boundary-layer mean wind speed in a city can surprisingly be higher than its rural counterpart, despite the higher roughness of cities. This Urban Wind Island effect (UWI) prevails in the afternoon, and appears to be caused by a combination of differences in ABL growth, surface roughness and the ageostrophic wind, between city and countryside. Enhanced turbulence in the urban area deepens the ABL, and effectively mixes momentum into the ABL from aloft. Furthermore, the oscillation of the wind around the geostrophic equilibrium, caused by the rotation of the Earth, can create episodes where the urban boundary-layer mean wind speed is higher than the rural wind. By altering the surface properties within the bulk model, the sensitivity of the UWI to urban morphology is studied for the 10 urban Local Climate Zones (LCZs). These LCZs classify neighbourhoods in terms of building height, vegetation cover etc. and represent urban morphology regardless of culture or location. The ideal circumstances for the UWI to occur are a deeper initial urban boundary-layer than in the countryside, low-rise buildings (up to 12 metres) and a moderate geostrophic wind ($8 \text{ to } 15 \text{ ms}^{-1}$). The UWI phenomenon challenges the commonly held perception that urban wind is usually reduced due to drag processes. Understanding the UWI can become vital to accurately model urban air pollution, quantify urban wind energy potential or create accurate background conditions for urban Computational Fluid Dynamics models.

This chapter is published as:

Droste, A. M., G. J. Steeneveld, and A. A. M. Holtslag, 2018b: Introducing the urban wind island effect. *Environmental Research Letters*, **13** (9), 094 007, doi: 10.1088/1748-9326/aad8ef

5.1 Introduction

The heterogeneous landscape of (large) cities causes a complex micro-climate, which can vary from street to street. The combination of ongoing urban expansion and climate change underlines the need to understand the dynamics behind this microclimate. The urban microclimate directly impacts citizen health through additional heat Oke (1982), air pollution Pascal et al. (2013), and influences energy demand Dhakal (2009). While urban heat has been widely studied (Oke, 1982; Arnfield, 2003; Steeneveld et al., 2011; Heusinkveld et al., 2014), knowledge about urban wind and its variability remains scarce Barlow (2014); Wise et al. (2018). These winds are governed by the dynamics of the atmospheric layer closest to the surface, i.e. the atmospheric boundary layer (ABL, Figure 5.1).

This study aims to quantify the difference in wind dynamics between city and countryside, using the conceptual Mixed-Layer Model for the daytime ABL Driedonks and Tennekes (1984); Tennekes and Driedonks (1981); Byun and Arya (1986); Conzemius and Fedorovich (2006a). Urban wind has hitherto mainly been studied for individual buildings or street canyons Carpentieri and Robins (2015), or over limited areas using Computational Fluid Dynamics (CFD) models Kondo et al. (2015); O'Neill et al. (2016); Wise et al. (2018), though CFD models are increasingly able to model larger urban areas Toparlar et al. (2017). However, knowledge regarding mean wind behaviour at the scale of the ABL can be valuable for more general aspects, such as the mean wind load on buildings, wind potential for energy production Walker (2011) or wind nuisance in urban planning.

Differences between urban and rural wind dynamics can be caused by increased roughness, by enhanced surface heating due to the Urban Heat Island (UHI) and changed ABL evolution. We aim to identify and understand these differences in wind behaviour through simulating contrasting wind regimes, by altering the surface parameters (roughness, displacement height, built fraction) and the initial conditions forcing the model (wind profile, geostrophic wind speed, boundary-layer depth).

Theeuwes et al. (2015) have used the Mixed-Layer Model to show that for certain conditions, boundary-layer dynamics can explain the Urban Cool Island: a period during the day where the city is colder than the countryside. Being aware of the impact of ABL dynamics, we hypothesise that an Urban Wind Island (UWI), where the urban wind speed is larger than the rural wind speed, might form under favourable conditions.

5.2 Methodology

We use the conceptual Mixed-Layer Model with a separated rural and urban column, as in Theeuwes et al. (2015). We expand their model set-up with the bulk equations for the zonal (U) and meridional (V) wind components (Figure 5.1).

5.2.1 Model Description

The Mixed-Layer Model is a slab model describing the mean ABL state. The model predicts well-mixed vertical profiles of turbulent quantities (heat, moisture, momentum), topped with a sharp

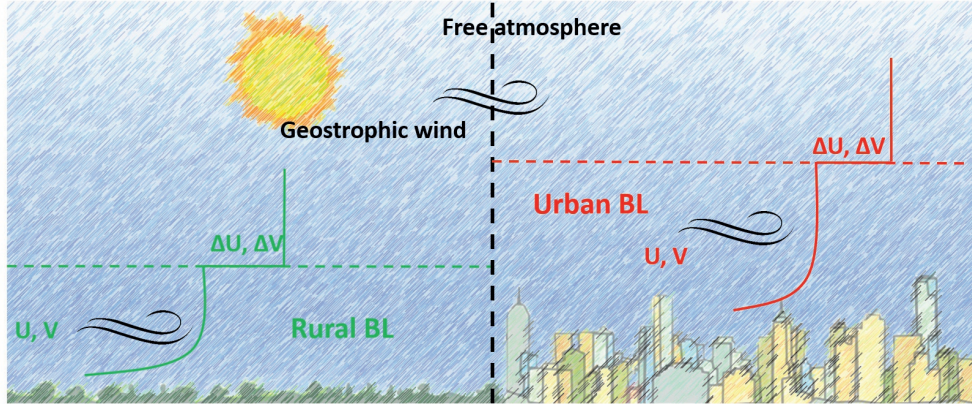


Figure 5.1: Schematic overview of the model. The rural column is indicated in green; the urban column in red. The dashed line represents the boundary-layer depth; U , V and ΔU , ΔV refer to the zonal and meridional wind, and their respective jumps to the geostrophic wind.

jump (ΔU , $\Delta \Theta$, etc.) to the free tropospheric profile, as well as the evolution of the ABL-depth (h). These conditions represent the convective daytime ABL, which is relevant when studying wind in the urban atmosphere. The tendency of the boundary-layer quantities is governed by their respective surface flux, and the entrainment flux, which mixes quantities down into the ABL from the troposphere and vice versa. Neither advection nor horizontal heterogeneity is considered. The simplified representation of the convective boundary layer makes these models very suitable for a wide range of conceptual studies (Tennekes (1973); Tennekes and Driedonks (1981); Schröter et al. (2013); Conzemius and Fedorovich (2006a); Byun and Arya (1986); Mahrt and Lenschow (1976)).

The mixed-layer equations governing the wind budget are:

$$\frac{dU}{dt} = f(V - V_{geo}) + \frac{1}{h}(\overline{u'w'_s} - \overline{u'w'_h}) \quad (5.1)$$

$$\frac{dV}{dt} = -f(U - U_{geo}) + \frac{1}{h}(\overline{v'w'_s} - \overline{v'w'_h}) \quad (5.2)$$

Here $U_{(geo)}$ is the zonal boundary-layer (geostrophic) wind; $V_{(geo)}$ the meridional boundary-layer (geostrophic) wind (all in ms^{-1}); $\overline{u'w'_s}$ and $\overline{u'w'_h}$ are the surface and entrainment momentum fluxes [m^2s^{-2}], respectively; h is the ABL depth (m); f is the Coriolis-parameter taken at 10^{-4} s^{-1} , representing mid-latitudes. A full set of governing equations is presented in the supporting material of this paper.

5.2.2 Surface model

The model consists of two, non-communicating columns: an urban column and a rural column (Figure 5.1). The columns are uncoupled to capture the behaviour of the urban effect on its own, rather than looking at rural-to-urban interactions (advection). The urban column represents a large metropolitan area, uninfluenced by the rural surrounding on the diurnal time-scale we are interested in. To distinguish between the urban and rural parts, underlying surface model parameters vary (e.g. the displacement height d is 0 m in the rural area, but significant in

the urban area). The surface model contains a well-validated land-surface parametrisation valid for clear-sky, daytime cases Theeuwes et al. (2015). The urban surface includes the effect of energy storage in impervious surface by using the Objective Hysteresis Model (Grimmond and Oke, 2002) which calculates the storage heat flux from the net radiation and the urban building fractions.

5.2.3 Model Validation

The rural part of the model is validated against observations from the Cabauw research tower, which measures the profile of wind speed among other quantities up to a height of 180 m. The Cabauw measurement site is a grassy field with a roughness length z_0 of approximately 0.2 m, and few other roughness elements in its direct neighbourhood, which represents a typical rural field area Beljaars and Bosveld (1997). Since the mixed-layer model is only valid for clear-sky, convective conditions, a set of clear days was chosen from the EUCAARIE campaign conducted in Cabauw in May 2008 Kulmala et al. (2011) for model evaluation. We evaluate the modelled wind against the vertically averaged, layer-weighted wind observations between 10 m and 180 m conform the bulk character of the model results. Geostrophic wind forcing is determined from surface pressure observations in a radius of 100 km around Cabauw.

The selected validation case is May 4 2008, a clear-sky day with moderate wind speeds (geostrophic wind components U_{geo} and V_{geo} were on average -5 and +4.5 ms^{-1} , respectively). The model simulation starts at 6:00 UTC (about 2 hours after sunrise) and ends at 18:00 UTC (LT is UTC +2), to capture the full extent of the convective boundary-layer regime. The model captures the mean wind behaviour to a sufficient degree (Figure 5.2a), though small-scale variability in the wind signal cannot be captured by the model due to its simplified physics and forcing. It seems the model produces relatively high friction in the first 2 hours compared to the observations, visible by the sharp decline in the U and V wind components, which leads to an initial underestimation of the wind speed. Afterwards the model reproduces the measurements very well, though some low extremes in the V -component of the wind cannot be reproduced. Overall, the RMSE in the mean wind amounts to 0.8 ms^{-1} which is a robust result, given the simplifications within the model. Air temperature is modelled equally well, with an RMSE of 0.6 °C. The modelled surface energy balance (not shown) corresponds to observations, though the model initially underestimates the latent heat flux (by $\sim 40 \text{ Wm}^{-2}$), and slightly overestimates the sensible heat flux ($\sim 15 \text{ Wm}^{-2}$) near the end of the model simulation.

The urban part of the model is validated separately, against measurements taken at King's College, London, at July 23rd 2012. Surface parameters are taken from Kotthaus and Grimmond (2014), and forcing parameters (geostrophic wind, initial profiles) are taken from the SUBLIME case description for urban model intercomparison Steeneveld et al. (2017). Observations are made on top of a building at 49 m above ground level (2.2 times mean building height), in a compact mid-rise neighbourhood. The chosen day is a clear-sky, convective day, though with substantial advection of momentum and a turning of the geostrophic wind with time, which we apply from SUBLIME.

The model performs satisfactorily (Figure 5.2b), with an RMSE of 0.78 ms^{-1} for the mean wind speed. The model does not simulate the sudden jumps in wind speed seen in the observations, which can be attributed to the relatively low height of the observations, which still induces some turbulent behaviour from the street canyons below. The overall result is good, and

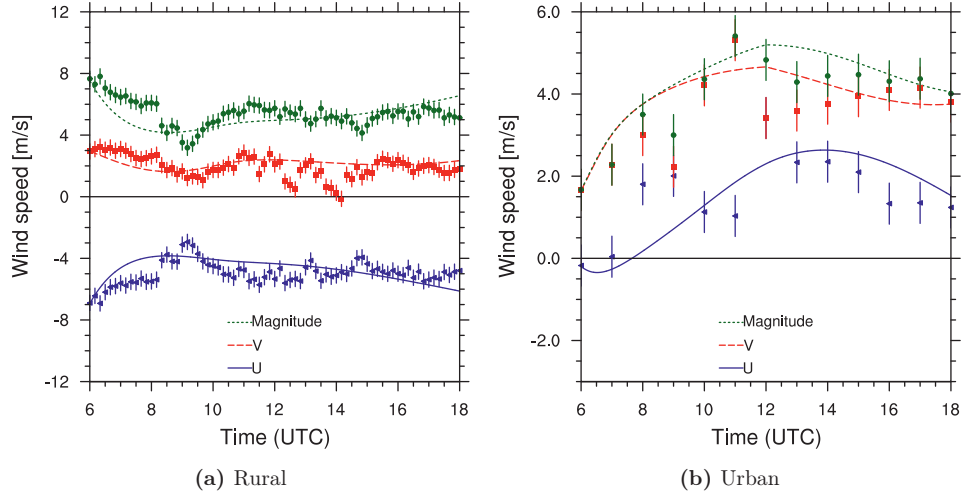


Figure 5.2: Validation of the model, for (a) May 4th 2008 at the Cabauw observational site and (b) July 23rd 2012 at the London King's College mast. Markers represent observations with error bars (10 min. at Cabauw, 30 min. at London); lines are the model outcomes. Green represents the wind magnitude (speed); in red and blue the meridional and zonal wind velocity components, respectively.

the model can therefore be used with confidence to model both urban and rural mean wind speed.

5.2.4 Experiments

Having validated the model, we follow Theeuwes et al. (2015) for the model initialization, using initial values from the BUBBLE campaign in Basel (Switzerland) (Rotach et al., 2005). During this campaign 3 masts measured wind speed in and outside of Basel for an intensive observation period of one month. Since most of the urban measurements actually took place inside the urban canopy and roughness sublayer (below ~ 2 times rooftop level (Barlow, 2014)) which causes a lot of disturbance, direct validation of the urban wind is challenging. Hence we use the BUBBLE measurements to provide a typical urban - rural contrast at the start of the model simulation. This set-up is then used as a base run, against which further setups of the model are compared.

Considering the budget equations for momentum (Equations 5.1 and 5.2), the modelled wind is influenced by 3 processes. The first process, $f(V - V_{geo})$, is the ageostrophic term, which redistributes momentum between the U and V components of the wind, due to the Earth's rotation. It follows that if $(V - V_{geo})$ (or $(U - U_{geo})$) is large, $\frac{dU}{dt}$ has to increase as well. The second and third processes are in the momentum flux divergence term, which contains a surface effect ($\overline{u'w'_s}$) and an entrainment effect (in $\overline{u'w'_h}$). When the fluxes are of equal sign, the momentum distribution term is small and the ageostrophic term will dominate the momentum budget. When the fluxes are of opposite sign, the momentum distribution term increases and

contributes more to the momentum budget. To determine which of these 3 effects (surface, entrainment, ageostrophic) is dominant in determining urban-rural wind contrasts we perform a set of experiments.

1. In the first experiment we eliminate the effect of the different urban surface. By setting the urban roughness lengths for heat and momentum to those of the rural surface, any influence caused by the buildings in a city is removed.
2. In the second experiment the entrainment rate is set to zero, to simulate a boundary-layer evolution without entrainment or detrainment from the free troposphere. The momentum budget is thereby only governed by the surface flux and the ageostrophic term.
3. The third experiment is to set the Coriolis parameter f to 0, which eliminates the ageostrophic term from Equation 5.1, so only the momentum divergence plays a role. Though $(f * V_{geo})$ is not 0 in this case, but just leaves the horizontal pressure gradient, we assume this to be negligible in order to investigate the importance of the ageostrophic term as a whole.

Finally, the 10 urban Local Climate Zones (LCZs, Stewart and Oke (2012)) are implemented in the urban surface model, to study the influence of the urban fabric on the wind behaviour. Building height (and thereby displacement height and roughness length), urban and vegetation fractions are varied between these LCZs to simulate varying degrees of urbanisation and urban morphology, and their respective effects on the wind.

5.3 Results

5.3.1 Reference case study

The urban model set-up resembles the city of Basel to represent the conditions of the BUBBLE campaign (Rotach et al., 2005). The geostrophic wind is set at 5 and 1 ms^{-1} for U_{geo} and V_{geo} , respectively, based on a climatology of the wind at sounding station Payerne (WMO code 06610) for June 2002 (the period of the BUBBLE campaign). Initial values of relevant model variables are given in Table 5.1.

Table 5.1: Initial values of the model variables. $z0m$ is the roughness length for momentum; θ is the potential temperature.

Parameter	Urban	Rural
$z0m$	1.5 m	0.2 m
h	400 m	100 m
$U(0)$	2 ms^{-1}	3 ms^{-1}
$V(0)$	0 ms^{-1}	0 ms^{-1}
$\theta(0)$	288 K	287 K
$\Delta U(0)$	3 ms^{-1}	2 ms^{-1}
$\Delta V(0)$	1 ms^{-1}	1 ms^{-1}
$\Delta \theta(0)$	4 K	5 K

The model results for the reference case are depicted in Figure 5.3. While the rural wind has a higher initial value than the urban wind speed, the urban wind speed accelerates for a longer




Figure 5.3: Model results for the reference case. In green are the rural model outcomes; urban in red. The asterisk in (b) indicates the the geostrophic equilibrium wind, respectively; dots demarcate the hours (starting from 6 UTC). The numbers indicate the start of the model run (at 0) and the direction in which the hodograph should be read.

amount of time after an initial drop at the start of the model run. This longer acceleration phase causes the urban wind to ultimately surpass the rural wind speed, creating a UWI at around 11 UTC, which reaches its maximum of 0.4 ms^{-1} at 14 UTC.

An interesting feature of this model run is visible in the hodograph (Figure 5.3b). Both the urban and the rural wind show a clear inertial oscillation in time, moving clockwise around the geostrophic equilibrium wind vector. Such oscillations are typically associated with the stable nocturnal boundary layer (Blackadar, 1957; Van de Wiel et al., 2010; Shapiro and Fedorovich, 2010), though several modelling and observational studies have also found these oscillations in the convective boundary layer (e.g. (Schröter et al., 2013; Byun and Arya, 1986)), where they significantly influence the wind, despite the effect of surface friction.

Figure 5.3 shows that the rural oscillation seems to be more dampened than the urban oscillation, which describes a wider circle around the geostrophic equilibrium. This would suggest that either the rural part has a higher internal friction dampening the oscillation, or that the ageostrophic wind at the start of the model causes a larger amplitude of the urban oscillation. The larger amplitude of the urban oscillation causes the urban wind to accelerate for a longer period of time, thereby forming the UWI.

5.3.2 Experiments

The experiments described in section 5.2.4 will focus on the UWI, and how the boundary-layer processes influence its modelled formation. Results of the first experiment, the equal roughness

(Figure 5.4b), are not very different from the reference run, though the urban wind values increase slightly. The urban surface momentum fluxes show a decrease, indicating lower friction generated by the surface, which causes the wind speeds to increase and enhances the UWI to $\sim 1 \text{ ms}^{-1}$. Hence, roughness affects the UWI, but does not seem essential in its formation.

The second experiment, without the entrainment effect (Figure 5.4c), shows remarkable differences between the urban and the rural response. Both the urban and rural wind become nearly stationary after 4 hours, but the urban wind increases more rapidly, to $\sim 4.1 \text{ ms}^{-1}$, whereas the rural wind initially decreases before slowly returning to its initial value ($\sim 3 \text{ ms}^{-1}$). The UWI is much larger than in the reference case ($\sim 1.5 \text{ ms}^{-1}$ versus $\sim 0.4 \text{ ms}^{-1}$) due to the stationary rural wind. An analysis of the two components of equation 5.1 shows that the ageostrophic term and the momentum divergence term are nearly equal but of opposite sign after 4 hours, effectively cancelling each other out. In the reference case the momentum divergence term is much weaker: this suggests that the entrainment counteracts against the surface flux, and that entrainment is an important source of momentum. To study whether the result of Figure 5.4c is attributable to the entrainment, and not to the difference in initial boundary-layer height between urban and rural, we repeat the experiment with equal boundary-layer depths (shown in the Supporting Information). In this case, a UWI is not formed, indicating that the increased boundary-layer growth of the city plays a crucial role in reducing the impact of friction over the whole of the boundary-layer.

Wind evolution is nearly stationary in the third experiment, ($f = 0$, Figure 5.4d). Neither urban nor rural wind change much over the course of the model run, though the urban wind increases more than the rural wind. This indicates that the inertial oscillation is a key driver of the UWI, since a strong ageostrophic wind ($f(V - V_{geo})$) causes a stronger acceleration of the wind, and this ageostrophic wind will be higher for the cases where the urban wind is initially lower than the rural wind. The influence of the initial conditions of the model seems apparent here: with urban and rural wind speeds not evolving as dramatically as in the reference case, the initial conditions determine whether the UWI occurs or not (Figure 5.4d).

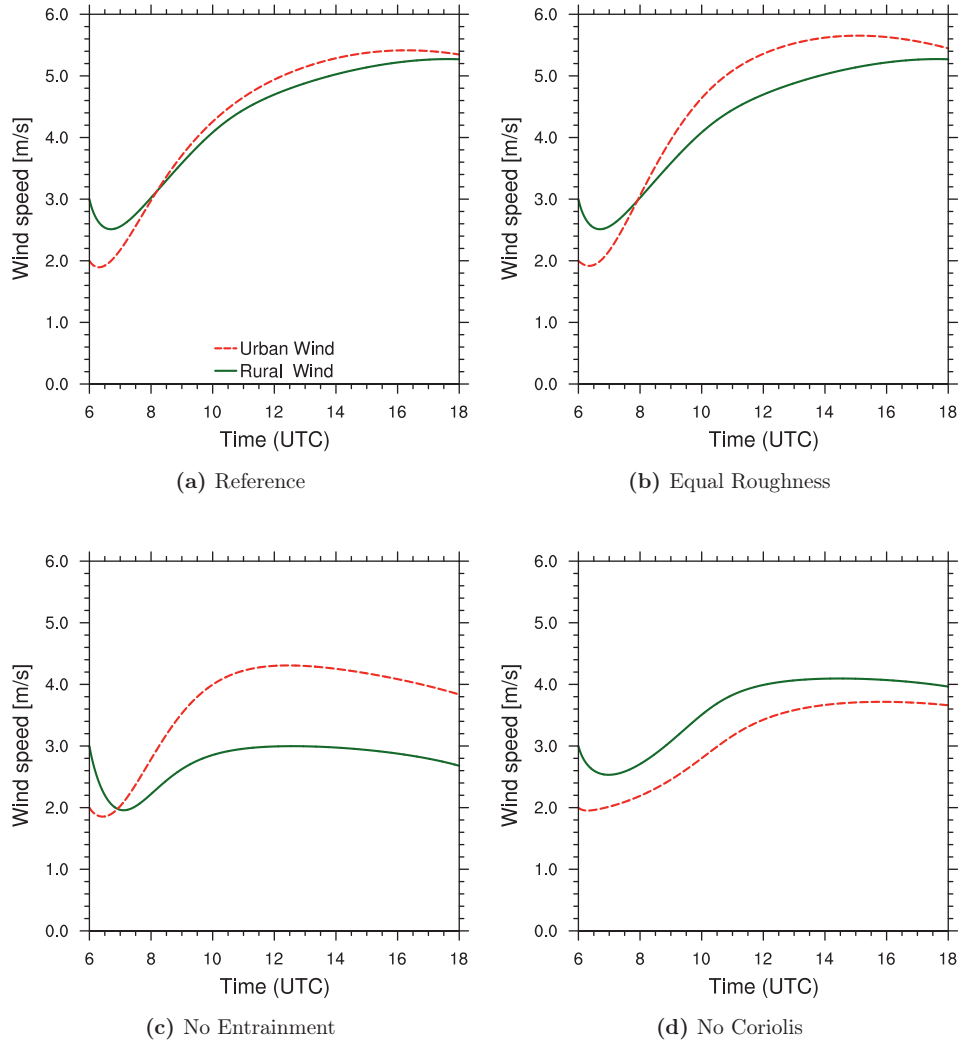


Figure 5.4: Modelled wind speed magnitude [ms^{-1}] for (a) the reference case; (b) the equal roughness case; (c) the no entrainment case and (d) the no Coriolis force case. Urban wind is given in red; rural wind in green.

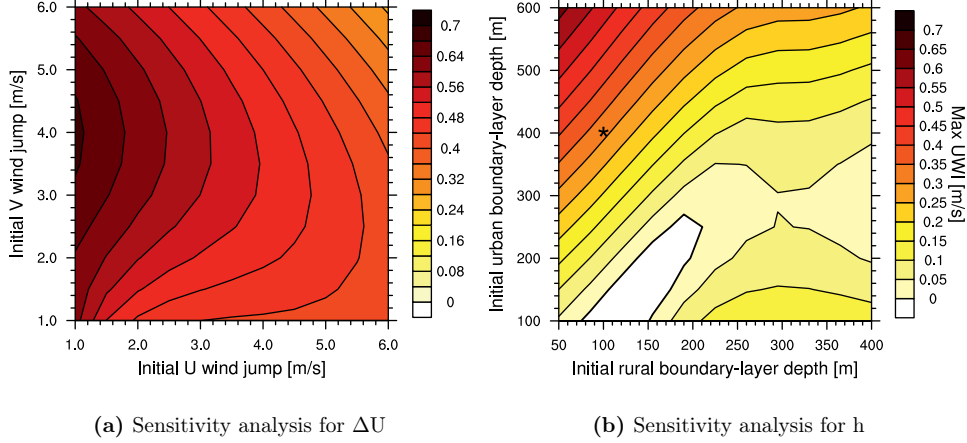


Figure 5.5: Sensitivity of the modelled UWI for the initial ΔU & ΔV (a) and boundary-layer depth (b). Contours indicate the value of the maximum (positive) UWI found, in ms^{-1} . Model initialisation for the other parameters is as in Table 5.1. In (b) the asterisk denotes the values for the control case. Note that since $U(0)$, $V(0)$ are constant, the change in $\Delta U, \Delta V$ changes U_{geo} and V_{geo} .

5.3.3 Sensitivity to initial values and LCZs

The previous section revealed that the initial conditions are important contributors to the formation of the UWI. We examine the sensitivity for initial boundary-layer depth, wind speed, and geostrophic wind speed. In addition, we run the reference case for all 10 urban Local Climate Zones (LCZs) (Stewart and Oke, 2012) to test how surface parameters such as roughness affect the model outcomes. For each sensitivity analysis, the model is run with the same settings as in Table 5.1, with a range of values for the sensitivity parameter of interest. The sensitivity analysis is based on 11 values per parameter, linearly increasing between minimum and maximum values, for a total of 121 model runs per analysis. The wind-jumps range between 0 and 6.0 ms^{-1} , and initial ABL depth from 50 to 400 m (rural) and 100 to 600 m (urban).

The modelled UWI formation appears to depend on the initial ageostrophic wind (Figure 5.5a). When the ageostrophic wind of the urban area is large, the amplitude of wind oscillation (as seen in Section 5.3.2) becomes larger and allows for a stronger UWI formation. Figure 5.5b shows the sensitivity of the modelled UWI to the initial boundary-layer depth (h_0) in the city and countryside. The maximum UWI is found for those cases where the initial urban boundary layer is several hundred metres higher than its rural counterpart. This is likely linked to the boundary-layer dynamics as seen in section 5.3.2. The higher urban boundary layer has a more efficient mixing, and smears out the surface friction over a thicker boundary layer, decreasing $\frac{dU}{dt}$ and $\frac{dV}{dt}$ (Equation 5.1).

The effect of the urban surface on the UWI is further explored by implementing the surface characteristics of the 10 urban LCZs. Each LCZ has a distinct combination of properties such as building height and impervious fraction that alter the boundary-layer dynamics. The model

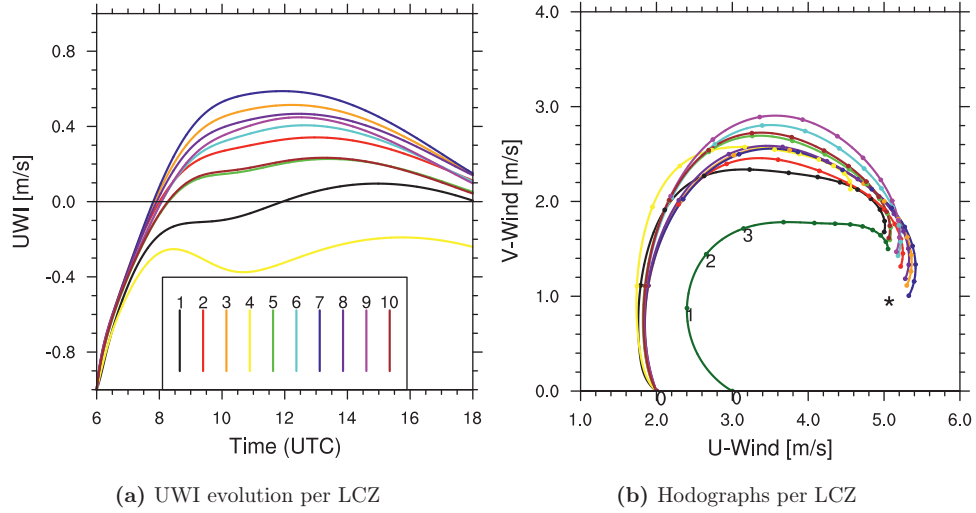


Figure 5.6: Time evolution of the UWI for the 10 urban LCZs (a), and hodographs for the 10 LCZs and the rural area (b). The hodograph is similarly annotated as Figure 5.3b. LCZs 1 and 4 are high-rise, 2 and 5 are mid-rise, and 3 and 6 are low-rise classes.

is initialised with the reference case settings (given in Table 5.1): the LCZ surface parameters are given in the Supporting Material, and more details of the LCZs in Stewart and Oke (2012).

All but one LCZ show the UWI formation over the course of the model run, with the peak commonly around 12 UTC (Figure 5.6a). LCZ 4 (open high-rise) does not show any UWI at all, though the similar LCZ in terms of building height, LCZ 1 (compact high-rise), shows a small UWI near the end of the model simulation (15 UTC). The building height for these LCZs is high (40m), which in turn means that the displacement height is large. The increased friction causes a strong sink of momentum at the surface, which reduces the urban wind speed. The UWI formation in LCZ 1 can be attributed to the high urban fraction, which enhances the sensible heat flux, allowing for greater boundary-layer growth. This growth can enhance entrainment of momentum, and offset the strong friction at the surface by diluting the friction effect over a thicker boundary layer compared to LCZ 4, which is more open and therefore contains less high heat-storage material. The hodograph (Figure 5.6b) shows the familiar oscillation for all LCZs, with few differences between LCZs. The weakest oscillation is seen in LCZ 1, where the wind is dampened by the strong surface friction.

5.4 Discussion

5.4.1 UWI in previous studies

While there is a range of published urban wind studies, those that look at urban and rural wind differences are scarce, due to the difficulties in observing and modelling urban wind. Studies that

take the broader scale of wind differences into account hint at the UWI phenomenon. Bornstein and Johnson (1977) measure wind speed in and around the city of New York. For low wind conditions they find the urban wind, downwind of the city centre, can be increased with respect to the rural upwind values. They attribute this to the influence of the UHI, which causes enhanced turbulent flow. Other studies confirm these results, though the threshold below which the UWI varies per city and season. Chandler (1965) and Lee (1979) find a UWI for the city of London (England); Fortuniak et al. (2006) for Łódź in Poland; and Shreffler (1979) for St Louis (United States).

The results of the first experiment in section 5.3.2 resemble the effect of the UHI, where enhanced mixing in the urban boundary-layer can accelerate the wind speed. Using a two-dimensional mesoscale model, Byun and Arya (1990) also model the differences between a rural and an urban site. Their results correspond to that of Chandler (1965); Bornstein and Johnson (1977) and Lee (1979), as they also find a region of increased wind speed downwind of the city centre, attributed to the intensity of the UHI, but they do not directly relate its magnitude to that of the rural wind velocity. While our results are focused on the daytime wind behaviour with a conceptual model, both observations Bornstein and Johnson (1977); Chandler (1965); Lee (1979); Fortuniak et al. (2006); Shreffler (1979) and a more advanced mesoscale model Byun and Arya (1990) detect the UWI, and its formation due to the enhanced urban boundary-layer mixing. The strongest urban acceleration (UWI) occurs at night in most of these studies: however, Bornstein and Johnson (1977) and Shreffler (1979) also find a UWI during the weak daytime UHI with a similar magnitude as our results ($\sim 0.5 \text{ ms}^{-1}$).

5.4.2 Model choice

The model used in this study is a conceptual bulk model, which makes several assumptions to investigate a wide range of idealised wind conditions. However, an independent run of the single-column WRF model Skamarock and Klemp (2008) (not shown) also displays a UWI with similar magnitude and timing, confirming the robustness of our applied model. The Mixed-Layer Model only takes the mean boundary-layer wind into account, hence no vertical profiles of the wind can be modelled. In reality the boundary layer can experience vertical wind shear during neutral and weakly unstable conditions, especially in the complex terrain of cities (Rotach, 1995; Macdonald, 2000; Castro, 2017), and knowledge of the vertical structure of wind speed is important for applications such as urban planning (e.g. through mechanical loads on buildings). Current models cannot easily simulate the flow field of wind inside an entire city: most limit themselves to either single buildings (Blocken et al., 2012), homogenized canyons under various wind regimes (Rotach, 1995) or rely on resource-intensive computation for an entire urban area Toparlar et al. (2017), and be used to select interesting circumstances which a more sophisticated model can then study in more detail, in a limited area setting.

5.4.3 UWI

The UWI that we find seems to be caused by urban boundary-layer dynamics, but it stands to reason that wind differences can also be caused by specific rural boundary-layer dynamics. An example is the transition towards the stable boundary layer at the end of the afternoon, when surface heating dies down and buoyancy suppresses turbulence. This transition in the rural boundary layer will happen earlier than in the urban boundary layer, due to the UHI, caused by

heat storage in the built-up surface Oke (1982); Grimmond and Oke (2002).

The nocturnal urban boundary layer often becomes neutral, where weak turbulence still produces (surface) winds, whereas the nocturnal rural boundary layer becomes stably stratified. The wind dying down in the rural area (before low-level jet formation) when there is still turbulence in the urban boundary layer, could potentially form a UWI. Observational studies (Chandler (1965); Lee (1979); Bornstein and Johnson (1977); Shreffler (1979); Fortuniak et al. (2006)) suggest an interplay between the UHI and the UWI, where a strong nocturnal UHI can fortify downward mixing of momentum into the urban boundary layer to further enhance the UWI. Whether this nocturnal UWI shares the same mechanisms with the afternoon-UWI that we find is a natural follow-up to this research. Rural-to-urban interactions of wind, such as urban breezes, are also not taken into account in this model, whereas they might play an important role in the horizontal distribution of momentum. We assume a decoupling between the urban and rural model parts, which in reality happens at distances of over ~ 50 km, where the advection is small compared to the other components of the momentum budget. Using a 3D weather model could provide valuable insight in the daily cycle of the UWI, and how momentum is transported to create the UWI. Observations of the UWI are difficult to make, due to the strong influence of the urban canopy on the wind field, but at sufficient height (2-5 times building height Barlow (2014), tower measurements can be very useful for UWI detection.

5.5 Conclusions

Using a conceptual mixed-layer model of the urban and rural atmospheric boundary layer we show that the mean wind in cities can exceed the wind in the rural surrounding. The model is validated against observations of the Cabauw tower facility in the Netherlands Beljaars and Bosveld (1997). The model initial values and forcing are based on a climatology of the typical morning values measured in the BUBBLE campaign Rotach et al. (2005). This Urban Wind Island effect (UWI) occurs primarily at the beginning of the afternoon, and has a magnitude typically around 0.5 ms^{-1} . The UWI is caused by a combination of effects:

1. Enhanced mixing in the urban boundary layer which facilitates entrainment of free tropospheric air;
2. A deeper urban boundary layer that dilutes the increased urban roughness;
3. The imbalance between boundary-layer wind and the geostrophic wind, which accelerates the wind.

The initial model conditions also determine the magnitude of the UWI. Optimal conditions for the UWI are:

1. A deeper initial urban boundary layer than its rural counterpart;
2. A moderate initial wind speed and geostrophic wind speed;
3. Relatively low building heights (around 12 m).

An analysis of UWI magnitude for all 10 urban Local Climate Zones (LCZs) Stewart and Oke (2012) reveals that the urban roughness from buildings decreases the UWI, and no UWI will for high-rise LCZs. Insight in the UWI can be used to determine whether a city has potential for urban wind farming, and to provide background knowledge for more detailed studies, for instance in air quality modelling.

Acknowledgements

The authors acknowledge the CESAR database (www.cesar-database.nl) from which the Cabaauw data was obtained, and would like to thank Fred Bosveld for supplying the geostrophic wind data. The authors thank Natalie Theeuwes for her help with the model development; Simone Kotthaus for supplying the urban data for London, and Aristofanes Tsiringakis for his help with the SUBLIME boundary values. Arjan Droste and Gert-Jan Steeneveld acknowledge funding from the Netherlands Organization for Scientific Research (NWO) VIDI Grant 'The Windy City' (File no. 864.14.007).

Supporting Material

Model formulation

In addition to the budget equations of momentum given in the main text, the following equations govern the wind in the Mixed-Layer Model:

$$\frac{d\Delta U}{dt} = w_e * \gamma_u - \frac{1}{h} (\overline{u'w'_s} - \overline{u'w'_h}) + f\Delta U \quad (5.3)$$

$$\frac{d\Delta V}{dt} = w_e * \gamma_v - \frac{1}{h} (\overline{v'w'_s} - \overline{v'w'_h}) + f\Delta V \quad (5.4)$$

$$w_e = \frac{dh}{dt} = \frac{T_0 V_*^3 (C_F - C_D \frac{hN}{V_*})}{g\Delta\Theta h + C_T T_0 V_*^2} \quad (5.5)$$

Here U is the zonal boundary-layer wind; V the meridional boundary-layer wind; w_e is the entrainment velocity (all in ms^{-1}); γ_u and γ_v represent the tropospheric lapse rate of the wind speed components (s^{-1}); $\overline{u'w'_s}$ and $\overline{u'w'_h}$ are the surface and entrainment momentum fluxes [$\text{m}^2 \text{s}^{-2}$], respectively; h is the boundary-layer depth (m); Δ indicates the jump of a quantity between boundary-layer and free troposphere; f is the Coriolis-parameter taken at 10^{-4} s^{-1} . For the entrainment rate equation, T_0 is a reference temperature (300 K); V_* is approximately the convective velocity scale w_* [ms^{-1}] under convective conditions (Conzemius and Fedorovich, 2006a,b); C_F , C_T and C_D are constants with values taken from Conzemius and Fedorovich (2006b) and Byun and Arya (1986) ($C_F = 0.5$, $C_T = 0.5$, $C_D = 0.024$); N is the Brunt-Väisälä frequency [s^{-1}]; $\Delta\Theta$ is the potential temperature jump [K] and g is the gravitational acceleration (9.81 ms^{-2}). In Equation 5.5, $T_0 V_*^3 C_F / gh$ represents the surface heat flux; $C_T T_0 V_*^2$ represents the non-stationarity in the turbulent kinetic energy (TKE) budget Byun and Arya (1986); and $-C_D \frac{hN}{V_*}$ accounts for the dissipation of TKE Conzemius and Fedorovich (2006b).

The surface and entrainment fluxes of momentum are calculated as follows:

$$\overline{u'w'_e} = -\Delta U w_e \quad (5.6)$$

$$\overline{v'w'_e} = -\Delta V w_e \quad (5.7)$$

$$\overline{u'w'_s} = -C_d U <U> \quad (5.8)$$

$$\overline{v'w'_s} = -C_d V <U> \quad (5.9)$$

$$C_d = \kappa^2 [\ln((z+d)/z_0) - \Psi_M((z+d)/L)] \quad (5.10)$$

Where $<U>$ is the absolute wind speed in ms^{-1} (i.e. the vector-sum of U and V), and C_d is the bulk transfer coefficient for momentum (Stull, 1988). In equation (5.10), κ is the Von Kàrmàn constant (≈ 0.4); d is the displacement height (0 m in rural conditions, and 4/5 of the building height in the urban surface); z and z_0 are height (taken as 10 % of the boundary-layer height) and the roughness-length of momentum (in m), respectively; and Ψ_M is the stability correction function for momentum (Stull, 1988).

These equations govern the wind and momentum in the Mixed-Layer Model: in addition to these we also calculate atmospheric potential temperature and moisture, and the interactions with the surface with a simple surface model. Since these are not the focus of this paper we omit writing these out: we refer to Theeuwes et al. (2015) and its supplementary material for the equations and the surface model formulation.




Figure 5.5.7: Model results for the idealised initial case. Ent refers to the entrainment flux; sfc refers to the surface flux of momentum. (a) shows the value of the urban momentum fluxes; (b) shows the tendency of the U and V components of the wind.

Supplementary results

Figure 5.5.8 depicts an additional validation case for Cabauw, May 5 2008. Model setup is similar as in the rural validation described in the main manuscript. The geostrophic wind components are set at -8 and 4 ms^{-1} for U and V , respectively. Surface parameters remain the same as in the original validation case, and initial values of the turbulent quantities (wind, moisture, temperature) have been taken from the observations. The model performs well in the wind speed, though overestimates the V component of the wind during mid-day (11 to 12 UTC), where it cannot reproduce the short-term fluctuations in the wind speed. Overall, the result is good: the RMSE is 0.91 ms^{-1} in the wind magnitude.

Momentum entrainment

For the reference case (Figure 3 of the main text), a striking difference between urban and rural is the entrainment of V -momentum, which shows a large peak at around 11 UTC (Figure 5.5.7). This coincides with the extended acceleration of the urban wind speed which ultimately forms the UWI. Slower momentum is entrained downwards (V is higher than V_{geo} at the time of the peak, so entrainment slows down the boundary-layer wind), which causes a positive value for the entrainment flux. The entrainment flux is opposite in sign to the surface flux, decreasing the momentum divergence term. The budget of urban V wind is thereby strongly governed by the ageostrophic term. While the U -components of both the urban and rural wind show a similar tendency (Figure 5.5.7b), the urban V -component tendency is larger where the rural tendency is close to zero. This matches the behaviour of the entrainment flux of V -momentum, as seen in figure 5.5.7a.




Figure 5.5.8: Validation of the model, for May 5th 2008 at the Cabauw observational site. Dots represent observations (every 10 min.); lines are the model outcomes. Green represents the wind magnitude (speed); in red and blue the meridional and zonal wind velocity components, respectively.

Experiments

Figure 5.5.9 depicts the value of the U and V components of the wind of the 4 cases described in Section 3.2 of the main text. In all cases, including the reference, the wind magnitude is mainly governed by the U component, though for the zero-entrainment case (Fig 5.5.9c) the V -component for both urban and rural wind becomes dominant, which can be attributed to the absence of entrainment of geostrophic V momentum which slows the boundary-layer V wind down.

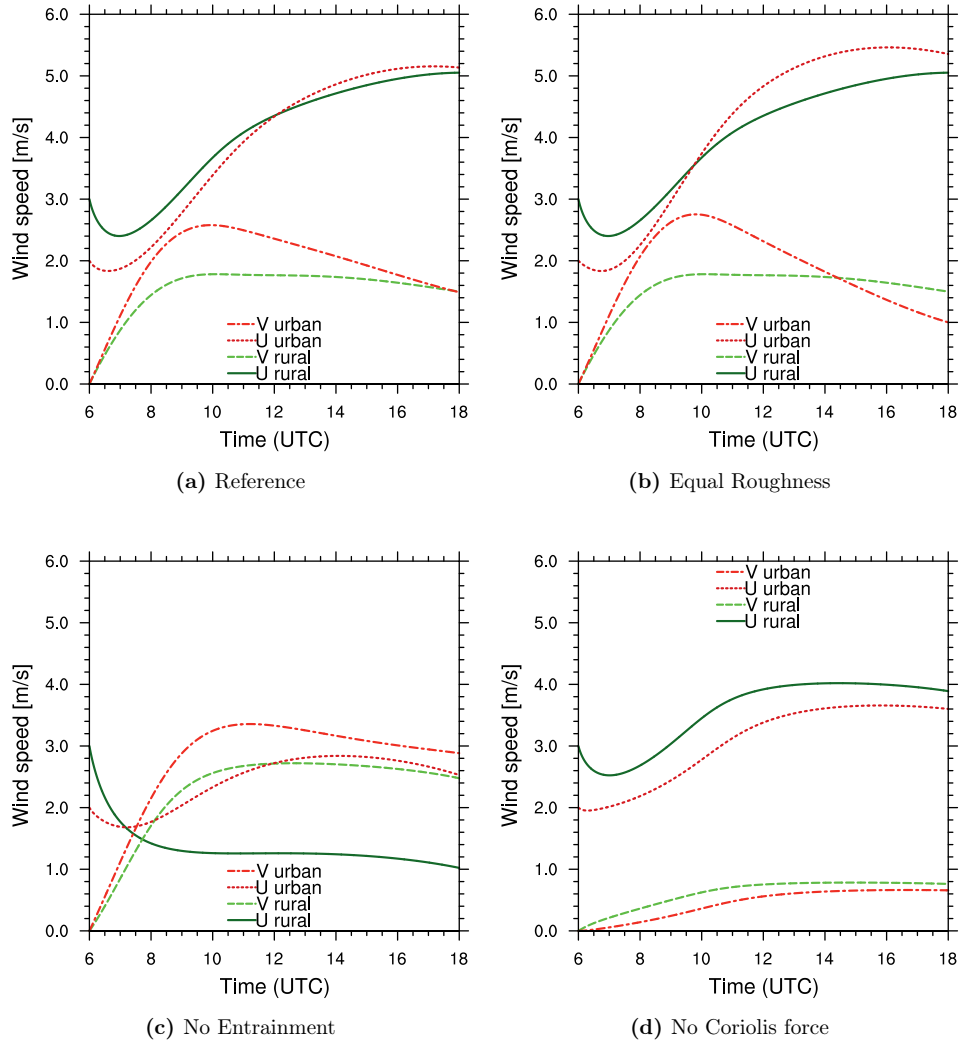


Figure 5.5.9: Modelled U and V components of the urban and rural wind [ms^{-1}] for (a) the reference case; (b) the equal roughness case; (c) the no entrainment case and (d) the no Coriolis force case.




Figure 5.5.10: Additional entrainment case experiment where initial ABL depth for both urban and rural parts is set at 400 m, in addition to the entrainment being turned off. Compare with Figure 4c in the main article.

Figure 5.5.10 provides the results for the additional entrainment experiment described in section 3.2 of the main article. In addition to entrainment being turned off, the initial boundary-layer depth for both urban and rural parts of the model is now set at 400 m. In this case, the rural and urban wind show very similar responses, though the urban wind accelerates earlier than its rural counterpart. A UWI does not form in this situation, highlighting the importance of the increased boundary-layer growth through entrainment, which is turned off in this experiment, on the UWI formation.

Local Climate Zone details

The values of various parameters of the urban surface model for each urban Local Climate Zone are provided in table 5.5.2. The values were taken from the data tables in Stewart and Oke (2012). For the analysis in section 3.3 of the main text, these parameters were changed in the surface model to alter the surface energy balance and surface roughness, to study the UWI formation under different urban morphological compositions.

Table 5.5.2: Surface model parameters for the 10 urban Local Climate Zones. ϵ_{surf} refers to the surface emissivity [-]; Alb to the surface albedo [-]; f_{gr} , f_{imp} and f_{build} refer to fractions of green, impervious and built area, respectively; z_{0m} is the roughness length for momentum [m] and z_h is the building height [m].

LCZ	ϵ_{surf}	Alb	f_{gr}	f_{imp}	f_{build}	z_{0m}	z_h
1	0.91	0.15	0.05	0.48	0.48	6.75	40
2	0.91	0.15	0.1	0.4	0.5	1.5	20
3	0.91	0.15	0.15	0.35	0.5	0.4	10
4	0.91	0.2	0.35	0.35	0.3	5.25	40
5	0.91	0.2	0.3	0.4	0.3	1.25	20
6	0.91	0.2	0.4	0.35	0.25	0.5	10
7	0.28	0.25	0.15	0.1	0.75	0.2	3
8	0.91	0.2	0.1	0.45	0.45	0.55	7
9	0.91	0.2	0.7	0.15	0.15	0.35	7
10	0.91	0.15	0.45	0.3	0.25	1.1	10

6

Chapter 6

The urban wind island from a three-dimensional perspective

The Urban Wind Island (UWI), a small but persistent positive wind anomaly over the city as a whole, has previously been revealed using a simplified conceptual model of the convective atmospheric boundary layer. However, the urban boundary layer is strongly heterogeneous and complex, and many interactions with surrounding land-use are not taken into account with the conceptual model used. Additionally, the transition to a stable or neutral nocturnal boundary layer substantially influences wind speed, for instance leading to nocturnal jets, which could also lead to UWI formation. This study extends the UWI research into less idealised cases by using the 3D WRF mesoscale model for Amsterdam (the Netherlands) and its surroundings, at 500m resolution. Two summers of forecast results for in total 173 days are used to identify whether the UWI persists in a 3-dimensional modelling environment, and which conditions are optimal for its formation and persistence. In order to focus only on wind modified by surface processes, large-scale influences which modify wind speed, such as frontal passages, are identified and eliminated from the dataset. We find that a positive UWI is present roughly half the time, with an order of magnitude that is similar to the previous work (103 0.5 m/s). In addition we find an evening UWI that is caused by the delayed onset of the transition from an unstable to a stable or a neutral boundary layer in the urban area, while the rural area is already stable and calm.

6.1 Introduction

At present, much remains unknown regarding the complex microclimate of cities and urban areas. Within a neighbourhood or even a single street canyon, strong contrasts in temperature, wind or radiation are the norm, rather than the exception. Due to this strong heterogeneity, taking representative urban meteorological measurements is extremely challenging and a costly affair. Though human activity concentrates in cities, routine meteorological measurement sites are exclusively situated in rural areas, that are preferably flat and homogeneous to capture the undisturbed larger-scale meteorological background required for accurate large-scale weather forecasting. Some weather stations have even been moved further away from encroaching urban sprawl (Yang et al., 2013; Bassett et al., 2019), lest it “contaminates” the rural background with the enhanced heat or roughness so typical of the urban environment.

The advances in computational processing power have improved numerical weather prediction models up to the point where cities should no longer be merely considered a different land-use class, but actually resolve several of the local-scale meteorological processes that strongly influence the urban area and its direct surroundings (Best and Grimmond, 2015). The urban heat island (UHI) is the most well-known consequence: the larger heating effect of the urban area caused by the relatively efficient absorption of solar radiation in the urban fabric, and longer retention of long-wave radiation. Abundance of building materials such as concrete, metal and asphalt, combined with a relatively low vegetation fraction and low average albedo shift a relatively high latent heat flux at a typical rural site, into a high ground heat flux, or storage heat flux as the term goes for urban meteorology. This heat is stored during the day, where the vertical structure of (large) cities captures much of the incoming radiation into roads, roofs and walls (Oke, 1988). This heat is released later at night, delaying the onset of the stable boundary layer, rather keeping a neutral or weakly unstable urban nocturnal boundary layer, which behaves very differently from its rural counterpart during the night.

These contrasting boundary-layer dynamics give rise to other local-scale phenomena, such as the urban wind island (UWI): a positive difference in boundary-layer averaged wind speed between urban and rural areas. (Droste et al., 2018b) (from here on referred to as *D18*) showed this UWI in a conceptual model of the daytime urban and rural boundary layers (see Figure 6.1.1). This mixed-layer model is a bulk-representation of the boundary-layer, with influences from the underlying surface and the free troposphere above. In *D18*, two separate columns are modelled, with the same free atmosphere, but distinct surface properties: urban and rural. A positive wind difference between the urban and rural column can be found under certain circumstances (clear-sky days with moderate wind speeds): the UWI.

Its two driving factors appeared to be the ageostrophic wind speed magnitude (the difference between the large-scale geostrophic wind speed, free of surface influence, and the mean boundary-layer wind speed), and the difference in the boundary-layer depth between urban and rural areas at the start of the day. The UWI is a fairly small effect, with values up to 0.5 m/s found in *D18*. Hence, it stands to reason that under the idealised circumstances of the conceptual model used, the UWI might not persist if other local processes are dominant. On the other hand, also observational evidence of the UWI has been provided (Fortuniak et al., 2006), meaning that there is some indication that the UWI is a real phenomenon, distinct from the sharp microscale differences in wind speeds often observed near buildings (canyon tunnelling or wind lulls at the leeward side of obstacles).

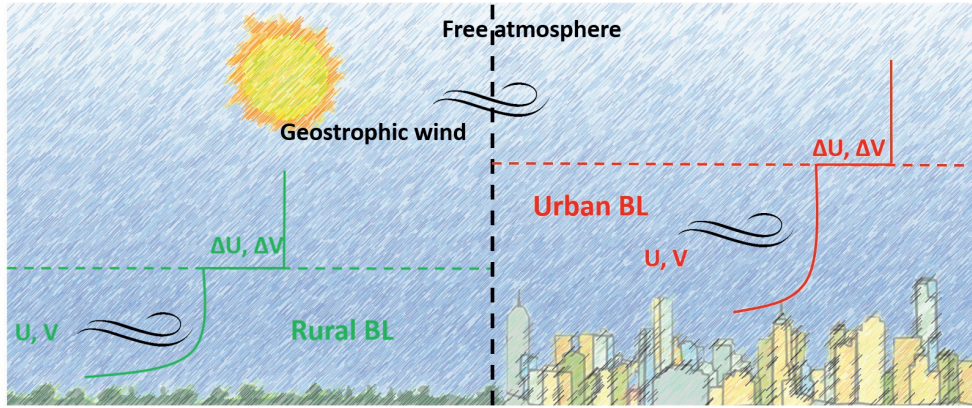


Figure 6.1.1: Schematic overview of the conceptual model used in (Droste et al., 2018b). The urban and rural part of the model do not communicate, but are influenced by the same geostrophic wind speed.

In *D18*, the most favourable conditions for UWI formations appear to be a moderate geostrophic wind speed (~ 5 m/s), a positive difference in boundary-layer height between urban and rural at the at sunrise, and cloudless conditions which encourage strong boundary-layer development dominated by local effects. Note that outside of an idealised where the boundary conditions can be controlled, frontal effects and other mesoscale phenomena may also influence the wind fields in both urban and rural areas, which can cause wind differences between the two that should not be considered an UWI. Thus, we narrow our definition of an UWI to be “a positive wind difference between an urban and nearby rural area, caused by differences in local dynamics”, and therefore need to remove large-scale effects that could suggest an UWI where there is none.

The objective of this study is therefore to establish whether the UWI phenomenon manifests itself in a 3D-context. Since urban wind observations are challenging, especially of a boundary-layer averaged phenomenon such as the UWI, we use mesoscale model results (Ronda et al., 2017). These are based on the WRF model (Weather Research & Forecasting, (Powers et al., 2017)) at very high resolution (100m at the inner domain) with improved urban physical schemes and land-use properties for Amsterdam, the Netherlands. Using the data of 2 summers, 2017 and 2018, we research whether the UWI occurs in these model results, which conditions are optimal to its formation, and what similarities or differences these have with respect to the theory of *D18*.

This paper has been organised as follows: Section 6.2 introduces the model setup, study area and research strategy; Section 6.3 provides the results, looking at UWI statistics under various conditions, its time of formation, and role of initial conditions; Section 6.4 discusses the implications of the found results and answers the research questions.

6.2 Methodology and Data

This section introduces the utilised WRF-model results and its setup for the study area of Amsterdam. Additionally, a single day is picked out that is expected to showcase a UWI formation in section 6.2.2, based on the conditions as found by *D18*. The data analysis strategy is laid out in section 6.2.3, outlining potential causes of the UWI and outside disturbances, and how to deal with those.

6.2.1 WRF-model and study area

The area of interest for our study is the city of Amsterdam, the Netherlands, and its rural surroundings, located at roughly 52.37° N , 4.89° E , in the north-west of the Netherlands. The region is characterised by a temperate maritime climate (Köppen Cfb climate) with the IJssel Lake in the vicinity ($\sim 8 \text{ km}$) to the east and the North Sea ($\sim 30 \text{ km}$) to the west.

The WRF model has been run for the study area with a customised land-surface and urban physics scheme as in (Ronda et al., 2017), which forms the core dataset of this work. The model setup works with 4 one-way nested domains, each increasing with a factor 5 in resolution, from 100m of the inner (4th) domain up to 12.5 km for the outer domain, roughly comprising north-western Europe. The forecasts have a lead time of 48 hours, with the first 24 hours used as spin-up. Soil- and urban fabric temperatures are initialised from the previous forecasts, and water temperatures for the two inner domains are taken from measurements by Rijkswaterstaat (Directorate-General for Public Works and Water Management): the governmental water body and infrastructure managing agency. (Ronda et al., 2017) focus their analysis on the inner, 100m resolution domain which is mainly the urbanised area of Amsterdam. Since this domain does not comprise of a rural area to serve as reference, but only the urban extent, we uniquely use the 500m resolution 3rd domain, centred on Amsterdam, which encompasses a large amount of the western Netherlands, including a swath of rural land to the south and north of Amsterdam (Figure 6.2.1).

To identify a robust urban wind signal, we use the grid-average of 5×5 grid-cells ($2.5 \text{ km} \times 2.5 \text{ km}$ in total) in the city centre, which is distant from the city borders and large water bodies. Selecting the rural reference location is more challenging, since to the east and west of Amsterdam there are no large undisturbed rural areas, with various cities to the west and the IJssel-lake to the east. Ideally, the rural reference wind speed would be upwind of the city, so no plume effects or momentum changes are advected from the relatively rough urban area to the rural reference, essentially “polluting” the rural signal which needs to be clear of urban influence.

The area north of Amsterdam has some urban structures, and the south is bordering a large nature reserve with well-watered meadows. Since the southern reference is quite far from the city, it is likely that any urban influence will already be dispersed by the time air parcels reach the location in case of northerly winds. Hence, the UWI could be independent of wind direction, using a fixed reference location instead, since the only other option would be the northern area, and northerly winds occur infrequently in the Dutch climate (Sluiter, 2011). In section 6.3.1 we first show that indeed we can use this fixed reference even during northerly winds. A grid-average of 5×5 cells is also used for the rural wind values. Since we use a grid-average instead of a single grid point, the exact location of the 5×5 grid is not crucial. When the grid is moved 1 or 2 grid cells, the resulting averaged wind speed hardly changes, indicating its robustness.




Figure 6.2.1: Land-use map of the WRF model, centred on Amsterdam. The black dots in the north and south mark areas designated as potential rural reference points, while the dot in the urban centre denotes the urban grid used

6.2.2 Convective case example

In *D18*, the UWI is most prominent on clear-sky, convective summer days with moderate wind speed: days that also tend to have higher UHI values. For example, August 18, 2018, was a hot day that met these criteria. Using WRF model results for UWI (Figure 6.2.2), the typical evolution of the UWI is similar as by the conceptual model of *D18*. It first emerges during the morning and peaks during the afternoon (with a value of ~ 0.6 m/s, consistent with earlier model results), after which it decreases as it becomes night-time.

In addition, a second peak is visible during the evening (20:00 UTC, or 22:00 LT) around sunset. This hints at a separate mechanism for the UWI, unrelated to convective boundary-layer growth. Rather, we hypothesise it is the delayed transition towards the urban stable boundary layer: due to the high heat storage of the urban fabric, more energy to drive turbulence is available even when insolation is low or 0 (after sunset). This energy reserve is the main driver of the UHI, and also causes a delayed transition towards a more stable (neutral) boundary layer. Thus, a timeframe exists where the rural boundary layer is already stable while the urban boundary layer is not, which could generate a UWI until the urban boundary layer becomes stably stratified as well. Section 6.3.3 in the Results will explore the influence of atmospheric stability on the formation of a UWI during the evening and early night.




Figure 6.2.2: Hourly PBL-averaged UWI over Amsterdam (The Netherlands) from WRF output on August 18th, 2018. Time is in UTC (LT = UTC+2). The southern point (Figure 1) is used as rural reference to calculate the UWI.

6.2.3 Data analysis strategy

In the introduction, we made the distinction between locally created wind differences (the UWI) and larger-scale influences which could cause an apparent UWI. Frontal effects, rain causing local downbursts, or storms can all modify the wind at either the urban or rural area (or both), without the surface playing any role. To study the pure UWI, these circumstances need to be identified and removed, by utilising some of the parameters present in the WRF model results. Hence, the presence of clouds and their impact on surface radiation, the boundary-layer depth, boundary-layer stability, and the geostrophic wind speed are all expected to contribute to UWI formation, and are extracted or calculated from the WRF output variables for selection purposes:

1. Rain is accounted for by the modelled accumulated grid precipitation, which tracks how much precipitation reaches the ground for a given grid cell over the course of that model run. If this is non-zero for the analysed day (the second day in the 48h forecasts), it has rained, and the entire day is flagged as such. We subdivide the rain flag into 3 classes: a 0 for no or very little rain (≤ 0.1 mm), a 1 for rain between 0.1 and 1.0 mm, and 2 for anything above 1 mm rain. WRF typically struggles with the timing and spatial distribution of rain, and this model setup was not validated for precipitation. However,

we are not concerned with the accuracy of the modelled rain, rather how it impedes the modelled wind speed. Therefore, whether the modelled rain values are correct is not very important: rather that we can filter out days with (heavy) rain.

2. Clouds are identified from 3 modelled cloud fraction variables: the cloud fraction at each model level, cloud liquid water, and cloud ice water contents. Since the level and composition of a cloud determines its physical impact (e.g. low, warm clouds impact surface radiation to a larger degree than high, cold clouds), we divide the cloud data into low, middle and high clouds based on approximate altitudes (the model layer closest to the mentioned altitude). Low between 0 and 2 km; middle between 2 and 7 km, and high above 7 km up to 12 km. Radiation, which governs the surface heat fluxes and therefore the strength of boundary-layer development, is expressed through the incoming shortwave radiation.
3. The geostrophic wind speed is not a direct output variable in WRF. Instead, we calculate it from the geopotential following Wallace and Hobbs (2006), at various levels. Since the grid spacing is relatively small, our estimation is liable to induce numerical instability if we do not average over a large field. Therefore, we use the geopotential gradient over 30 grid cells to estimate the geostrophic wind positioned over the city centre. The geostrophic wind can be calculated from the geopotential gradient at various levels: for analysis purposes, we use the geostrophic wind calculated from the 900 hPa geopotential field.
4. Thermodynamic stability of the surface layer is approximated using the bulk Richardson number between the surface and the first model level (approx. 30 m), which can be calculated from WRF output variables in a straightforward manner (equation 5.6.3 in Stull (1988)). We approximate static stability by the sign of the bulk Richardson number: determining the layer as stable when Rib is positive, and unstable when it is negative. We also use this bulk Richardson number to determine the boundary-layer height, using the parcel method (Vogelezang and Holtslag, 1996), with the critical value of the Richardson number set at 0.25. At night, we approximate the PBL depth by $700 \times u^*$ (Koracin and Berkowicz, 1988; Vogelezang and Holtslag, 1996; Steeneveld et al., 2007), while during the day the maximum of the two methods is selected as the returned PBL height. WRF does provide a PBL estimation itself but especially at night this leads to unrealistically low PBL depths, presumably because of the parcel method not being optimal during stable conditions. During convective daytime, the values from the two methods corresponds rather closely (mean relative error of $\sim 6\%$).
5. The magnitude of the total wind speed is the vector sum of the WRF outputted u and v components of the wind speed at all model levels inside the boundary layer. The UWI is then the difference between the PBL-averaged wind speeds of the urban and the rural model locations.

Aided by these identifiers, we construct histograms of the modelled UWI for various conditions, to identify under which circumstances the UWI is highest, how long it persists (in terms of a continuous positive UWI episode) and during what time of the day it occurs. We also examine whether the influence of the (initial) PBL depth is similar as in *D18*. Additionally, we explore the relation with the geostrophic wind (components) to the value of the maximum daily UWI values.

6.3 Results

This section presents the results of our data analysis, starting with a justification for the selected rural reference location in the model. Following that, the mean UWI is analysed with special attention to the effects of rain and stability forming the nocturnal UWI. Finally, the effect of the initial boundary-layer conditions on the UWI (PBL depth, UHI and geostrophic wind) are investigated.

6.3.1 Rural reference location

We first test whether a fixed rural reference location can be used, even for situations where the site is downwind from the city, and therefore potentially contaminated by the urban plume. The period between 20 and 28 July 2018, had predominantly northerly winds, varying between north-west and north-east. For that episode, 185 hours are available for analysis, after removing some hours where a sea-breeze front was present. The modelled UWI is calculated for both the northern (upwind) and southern (downwind) rural reference grids, as in Figure ref:fig:1. If the difference in UWI between north and south sites is not statistically significant (using a 2-sided t-test of the difference), both sites are equally usable under these conditions, even when a site is downwind from the city. An insignificant result would suggest that this urban influence is minimal at this distance (20 km) and can therefore be ignored, and the southern site used as reference regardless of wind direction.

The results of the statistical test are strongly insignificant: with $p = 0.769$ ($t = -0.294$; $df = 338.18$) we cannot reject the null hypothesis at $\alpha = 0.05$. Ergo, the mean difference between the results for the two sites is likely zero, and there is no distinguishable urban plume effect on the downwind wind speed at the southern site. Based on this result we will use the southern site as rural reference for any results from here onwards.

6.3.2 Mean UWI under various conditions

A histogram of the hourly UWI values over both years (no filters) shows that the distribution is roughly normally distributed, centred around zero, with both positive and negative outliers, favouring generally lower positive values (Figure 6.3.1). Some extreme values lie beyond the plotted range: these consist of only 17 hours out of a total 4152, and coincide with frontal passages. The positive UWI generally is within the range found in *D18*, $\sim 0.5 - 1$ m/s, with some distinct outliers going up to 2 m/s. The negative side of the distribution shows a more gradual decline, but is in values roughly the same. Negative UWI seems to occur mostly at night-time (also indicated in Figure 6.2.2), where the rural area might experience a low-level jet, or the high roughness of the urban area has a strong effect on the boundary-layer wind speed, since the boundary-layer depth is much lower than during the day.

When subdividing the dataset into the 3 rain categories as described in section 6.2.3, it initially seems that daily maximum UWI increases with the accumulated precipitation for that day. Convective precipitation events can include downbursts, i.e. modifications of the wind field due to non-local phenomena. These events are therefore better disregarded, as they do not represent the surface-induced UWI. Rain flag 1 however contains little precipitation (0.1 up to 1 mm/d) which is typically not associated with heavy wind activity. Possibly warm fronts, accompanied by drizzle, affect the wind field by altering the wind direction. For the rest of this study, days with




Figure 6.3.1: Histogram of the modelled hourly UWI over Amsterdam over 2017 and 2018. Each bin is 0.2 m/s, totalling 4135 hours.

> 1 mm accumulated precipitation are disregarded, leaving a total of 106 days up for analysis (61 % of the original dataset).

Figure 6.3.2a shows the averaged diurnal UWI timeseries for those days without rain (69 days in total, 40 % of the original dataset; 1656 hours). The red line in Figure 4a is the conceptual model of *D18*, run for the average forcings of the 2 years of WRF 3D data, resulting in a very similar course of UWI between the bulk model and the 3D WRF model. The daytime values of the WRF-UWI (blue line) show a strong similarity to the course of the UWI in *D18*, which also reaches its maximum at the end of the afternoon. A difference is that the UWI in this data already starts positive a little after sunrise (which is around 4 UTC, 6 LT), whereas the UWI in *D18* didn't become positive before 10 LT (8 UTC). The value of this mean UWI is rather low (0.2 m/s), but realize this contains all days with the right rain flag, rather than an idealized case. The control case of *D18* also showed UWI values around 0.3 m/s, and only under more optimized circumstances did that value increase. The similarity between the WRF-3D data and the mixed-layer model of *D18* suggests that on average, the UWI is formed by the core processes described in *D18*, but individual days see a strong spread in UWI formation and values (as seen in the histogram of Figure 6.3.1).

The UWI is a boundary-layer averaged phenomenon, and though wind is generally a well-mixed quantity in the convective boundary layer, in practice there will be differences with height, e.g. related to shear. Figure 6.3.2b also includes the UWI calculated from 3 different model levels, aside from the boundary-layer averaged UWI also shown in panel a. Model level 1 is close to the surface at 32m; model level 3 is moderate, at 120m; and model level 6 higher at 300m,

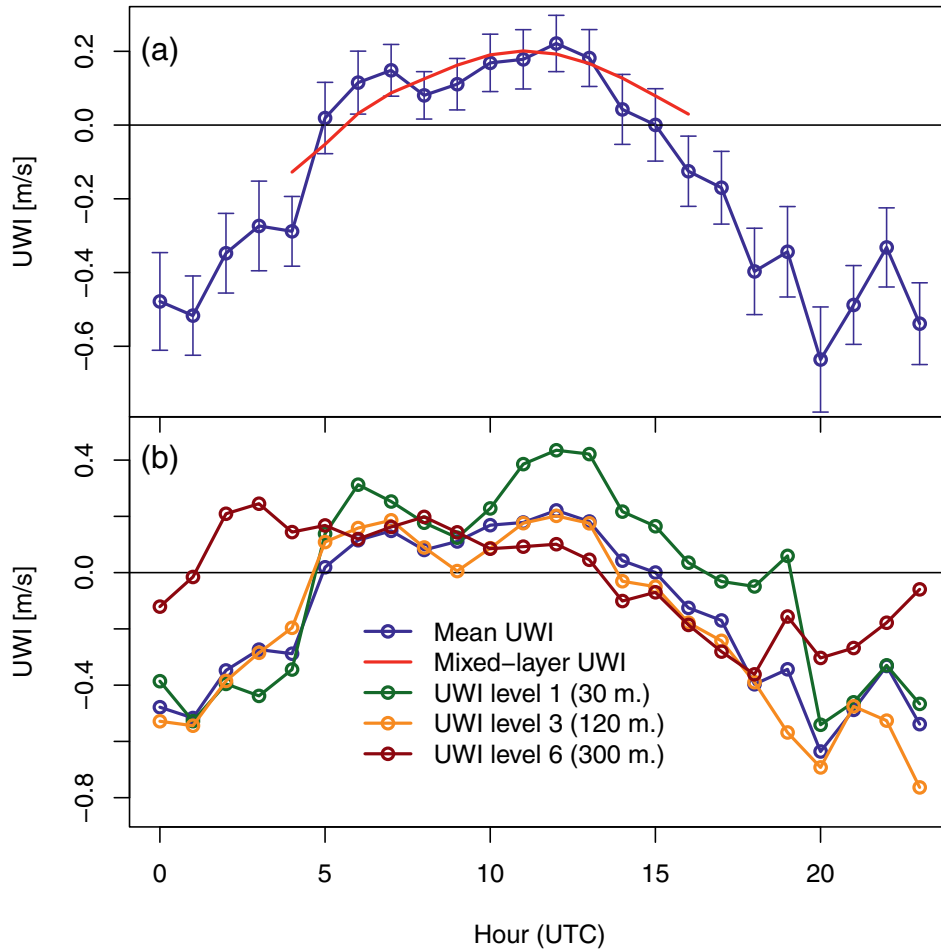


Figure 6.3.2: Mean hourly WRF-modelled UWI for Amsterdam calculated over days without rain (a; blue line) and the mixed-layer model UWI run with averaged forcings over those days (red line). Vertical crossbars in (a) indicate the standard error. Panel (b) shows the mean hourly UWI at various model levels, as well as the boundary-layer average which is equal to that of (a).

meaning it will occasionally be above the boundary-level height during early morning or late at night. Especially the surface-level UWI is more variable during the day, likely due to turbulence in the urban area generating higher wind speeds. At night, the UWI at 300m is positive which could be caused by low-level jets. The medium level UWI shows very similar behaviour to the boundary-layer weighted UWI, which is less directly reliant on the surface processes but still within the boundary-layer during the night.

The hodograph of the average wind time series data (Figure 6.3.3) separates the PBL-averaged wind speed in its U and V component for the urban and rural location. We indeed find a typical inertial oscillation of the wind speed over the day, turning clockwise. This is precisely the behaviour found in *D18* and also in other studies of the inertial oscillation (Blackadar, 1957; Byun and Arya, 1986), now also extending towards night-time rather than only the convective boundary layer. Note that this is the average wind over 2 summers of data: individual days have quite some variability between them and may not portray an oscillation at all or only for a certain portion of the day. Similarly, the diurnal UWI evolution also shows a lot of spread over the 2 years [see histogram, Figure 6.3.1].




Figure 6.3.3: Hodograph (with dots representing time rather than height) of urban (red) and rural (green) mean hourly modelled wind speed for days without rain (selection similar to figure 4). Each point represents one hour, starting at the middle at 00 UTC and going counter-clockwise, ending at 23 UTC.

6.3.3 Stability and night-time UWI

As seen in the case of Figure 6.2.2 and various other days, the UWI can be positive shortly after sunset, which the conceptual study in *D18* does not cover. During this time of the day, the boundary-layer transitions from unstable to stable stratification, which suppresses buoyancy-driven turbulence and generally lowers wind speed. The majority of positive UWI values does not occur during this time but rather during the day (as the mean diurnal timeseries in Figure 6.3.2 shows), so it seems that the inertial oscillation mechanic is not the (main) driver of this evening-transition UWI.

The frequency distribution for the evening hours where the urban boundary layer is unstable and the rural boundary layer is already stable (Figure 6.3.4), shows that the UWI for these transitional cases leans more towards the negative side compared to the convective daytime values. The positive side of the distribution looks very similar however, though any strongly positive values are absent. The total number of hours where this transition happens is also limited, 173 in total after filtering for high rain days (rain flag 2), but even when including those days, the shape of the distribution changes only slightly, centring more around 0 and containing larger outliers. Nearly all days show this transition, which typically occurs between 18:00 and 20:00 UTC, i.e. around or a little after sunset. The UWI during this time shows no distinct difference from that of the entire dataset averaged. This means that the evening UWI is not strongly different in magnitude from the UWI formed during (convective) daytime.

However, when looking at the days where the maximum daily UWI happens during this evening time window, a pattern seems to emerge (12 days in total). The UWI over the course of the day is generally positive or very close to zero, and at the end of the afternoon ($\sim 16:00$ UTC) starts to increase to reach its maximum during the transition phase (Figure 6.3.5a). This maximum UWI tends to be fairly high, over 2 m/s in half of the cases, and then shows a reduction as the night progresses and the urban boundary layer becomes stable. When taking the transition window somewhat wider, considering the daily maximum UWI between 18:00 and 23:00 UTC (23 days in total), the time evolution is similar, but the values are somewhat lower, with a more gradual decrease during night-time (Figure 6.3.5b).

6.3.4 Role of initial conditions

The relation between initial PBL-depth and the strength of the UWI as found in *D18* seems to be much less strong in this dataset. Figure 6.3.6 relates PBL-depth of urban and rural boundary layers at 4:00 UTC (6:00 LT, just after sunrise) to the maximum UWI during that day, but a clear pattern is not present. The higher UWI values (> 2 m/s) all lie above the 1:1 line, so for situations with a higher initial urban than rural PBL depth, but there are also plenty of low-UWI days in this area of the plot, particularly for the very low values. The size of the points indicates the duration of the UWI episode of that day (i.e. consecutive hours with positive UWI values).

A higher urban PBL depth occasionally coincides with higher daily maximum UWI values, but there is no convincing relation: a statistical (non-parametric) test of the difference between the two data clusters above and below the 1:1 line yielded no significant difference. This is different from the *D18* study which showed a dependence on the initial boundary-layer depth, which was especially strong for days with a high initial difference between rural and urban boundary-layer depths. Such days do not seem to exist within this dataset, and might therefore not be




Figure 6.3.4: Frequency distributions of the hourly modelled UWI where both urban and rural boundary-layers are unstable (upper plot) and where the urban PBL is still unstable where the rural PBL is already stable (lower plot).

very realistic for Amsterdam. A statistical extrapolation based on these data (not shown) does show a similar relation between UWI and PBL depth as in *D18*, but the required combination for PBL depths does not show up in this dataset, so that extrapolation cannot be reliably used.

In a similar fashion, the value of the geostrophic wind speed is a less dominant driver for the UWI formation than was apparent in *D18*. The relation with the initial ageostrophic wind




Figure 6.3.5: Hourly UWI for days that have the maximum daily UWI occur between 18:00 UTC and 20:00 UTC (a) or 18:00 UTC and 23:00 UTC (b). Vertical black lines denote the timeframe where the maximum daily UWI occurs, which coincides roughly with the evening transition of the PBL.

($U_{geo} - U_{PBL}$) at sunrise is weakly positive, where a larger difference between the boundary-layer wind speed and the geostrophic wind speed seems to lead to larger UWI values (results not shown). Exploring the daily mean geostrophic wind speed (Figure 6.3.7), the relation between geostrophic wind and the daily maximum UWI is again not as clear-cut as for the conceptual model, though generally moderate geostrophic wind speeds coincide with higher UWI values. Interestingly, higher geostrophic wind correlates with the number of hours that a positive UWI does occur on a given day: especially within the 10–15 m/s geostrophic wind range. Why exactly this happens is unclear: potentially the higher geostrophic wind speed incites stronger boundary-layer flows that get accelerated through the enhanced urban turbulence.




Figure 6.3.6: Modelled boundary-layer depth (rural and urban on x and y axis, respectively) at 4:00 UTC (6:00 LT) related to the modelled maximum daily UWI (colours). The size of the points is an indication of the duration of a positive UWI episode of that day: consecutive positive UWI values.

In addition, we do not find a significant relation between the UWI and the amount of solar radiation received at the surface over a day, nor does the type or height of cloud presence during the day seem to have a relation to the UWI occurrence for that day (results not shown). Both hourly values or daily sums of radiation show no clear signal, indicating that cloudiness or radiation do not influence the UWI formation or magnitude, apart from the relation to intense rainfall events that induce downbursts, which we have intentionally removed from the dataset. However, there is a weak relation to the temperature difference (UHI) at just after sunrise (5:00 UTC, Figure 10). A higher UHI seems to lead to higher UWI values during the day, though for UHI values higher than 1 °C the effect lessens. A higher UHI at the start of the day means more energy available for the urban boundary-layer to grow and develop turbulence, which could in turn lead to a higher UWI over the course of the day.

6.4 Discussion and conclusion

In this study we have looked at the Urban Wind Island (UWI), the positive urban wind anomaly, in the setting of the 3D mesoscale meteorological model WRF. Model data from 2 consecutive summers have shown that the UWI is present in, and that aside from a daytime UWI there is also an evening UWI, induced by the delayed transition to the stable nocturnal boundary layer.




Figure 6.3.7: Modelled daily mean geostrophic wind versus the number of positive UWI hours for a given day (with no to little rain). The coloured points indicate the value of the daily maximum UWI for that day.

The UWI is convincingly present in the modelled data, and not just in very specific circumstances, but even those that are not “typical” nice summer days. The value of the positive UWI is not very high, in the order of 0.5 m/s, which is similar order of magnitude as the conceptual model predictions of *D18*. The average diurnal profile of the UWI, as well as the time-averaged hodograph in figures 4 and 5 correspond very well to the theoretical model of the UWI, but individual days can strongly diverge from the average pattern, and reach much higher values, up to 3 m/s. The average time evolution of the UWI corresponds to that of the conceptual model, with an acceleration phase and positive UWI’s during morning/early afternoon, and then negative values. The UWI timing on average is very similar to the earlier work, with the onset a little earlier (few hours after sunrise), but the maximum UWI is also found in the afternoon.

The absence of a relation with cloudiness or radiation suggests that the clear-sky convective situation as modelled in *D18* is not a requirement for the UWI to form. However, on days with frontal passages or strong rainfall events, an apparent UWI can form through downbursts or time delays in frontal passage over the urban or rural area. These are separate from the boundary-layer phenomenon and have been filtered out in this study.

In situations where the rural boundary layer reaches stable conditions earlier than the urban area, a UWI can form through the suppression of turbulence in the rural area while it persists in the city, typically a few hours around sunset (18:00-20:00 UTC). In these situations, the UWI increases at the end of the afternoon, peaking at around 19:00 UTC before reducing in value again as the night progresses and the urban boundary layer becomes (weakly) stable or neutral.




Figure 6.3.8: Modelled initial UHI at 5:00 UTC versus modelled daily maximum UWI for days with little or no rain.

This is a separate effect from the UWI caused by differences in the inertial oscillation of wind during the day, and typically leads to a weaker UWI that persists for a shorter period of time (~ 2 -3 hours).

Overall the UWI's conceptual framework from *D18* broadly fits the UWI formation as seen in the more realistic WRF setting, but the importance of the difference in initial urban and rural PBL height seems to play a lesser role. There is some relation between the PBL depth at the start of the day (4:00 UTC, corresponding to 6:00 LT as also used in *D18*), with a higher urban boundary layer typically leading to a higher maximum UWI, but not in all situations. The signal is somewhat noisier and not as clear-cut as the conceptual study. The magnitude of the geostrophic wind speed is seemingly an important component of the UWI formation in the conceptual study, but this relation is much less apparent and much noisier without any real pattern in the 3D data. While a general relation is still visible, and similar to the conceptual study, we find quite some spread, but the number of positive UWI hours in a day seems to depend on the mean geostrophic wind speed.

Since we have deliberately chosen two model locations well in the middle of their respective land-use categories, what happens on the interface between urban and rural land has been held out of the scope of this research. The urban breeze phenomenon, analogous to sea-breeze effects, could also have some influence on the UWI, but this would be a locally driven circulation, which is a separate, though not to be ignored, effect on wind speeds.

We have chosen for a grid-averaged UWI, rather than a spatial analysis, mainly because the

model infrastructure is too coarse to do detailed spatial analysis of this phenomenon. The definition of a UWI (or indeed a UHI) requires a background value, which necessarily is either a point or grid average, otherwise a meaningful UWI cannot be constructed. While we could construct a map of the UWI in Amsterdam, not much of value would be gained, since the length scales of urban heterogeneity are much smaller than the grid size of 500m, so spatial variability would not be very valid, and could suggest spatial patterns that are perhaps not valid in a realistic urban setting.

Amsterdam is, by international standards, a fairly small city. The conceptual model in *D18* assumed separate urban and rural columns, with boundary-layer properties fully formed by these respective surfaces. Since Amsterdam is quite small, the urban boundary layer might well have some rural advection changing its composition. Similarly, the land surrounding Amsterdam are also of mixed land-use, with various small towns interspersed with the natural areas. Repeating this study for larger cities with distinct urban atmospheres (e.g. New York) might yield firmer results that could be closer to the conceptual model work.

While we have taken care to identify cases where larger scale processes influence the wind pattern, it is possible that not all of these situations have been properly captured when constructing the longer-period statistics. Amsterdam is relatively close to the North Sea (~30 km), and can experience sea breeze fronts in situations where the land-sea contrast is very high. These fronts can reach either the rural or urban area before the other, resulting in large positive or negative wind differences which would suggest a UWI. A workaround could be to look at sudden changes in wind direction at either area, but this is no guarantee that a sea breeze front is actually passing, or that some other process, such as onset of stability (which is a local process we are interested in) is modifying wind direction.

Validation of the model data is primarily done in Ronda et al. (2017). However, they don't validate wind speed in their model setup. Urban wind measurements are notably difficult and non-representative, effectively making an urban point measurement and the 500m grid scale wind entirely different quantities which should not directly be compared to each other. There is a measurement network in Amsterdam which measured urban canyon wind speed, which is again an entirely separate quantity, which can be translated to boundary-layer wind, with a notable degree of error. Wind direction, rather than speed, could also be an influence: if the speed is correct but its direction is not, the inertial oscillation might start in a different way, and therefore generate a different UWI time evolution.

In summary, the UWI is present in a realistic 3D modelling setting. Its diurnal evolution corresponds to the conceptual framework presented in *D18*, with an extra mechanism generating a UWI around sunset. This evening UWI is caused by the difference in stability between urban and rural, where the onset of the stable rural boundary layer suppresses turbulence and subsequently reduces wind speed. Unlike the conceptual framework, a clear-sky convective day with low wind speeds that would be beneficial for the UHI is not required for the UWI, since there is no clear relation to radiation or the magnitude of the geostrophic wind, and days with little rain still generate a positive UWI.

A longer dataset (e.g. the ERA-URBAN urban meteorology reanalysis project, van Haren et al. (2018)), also encompassing multiple cities in different geographical situations, would solidify the statistics of this article. Amsterdam is surrounded by urban agglomerations and water bodies,

which makes a good rural reference hard to identify. A more theoretical case study in WRF, by creating an “academic city” (e.g. Theeuwes et al., 2013), could help further research the influence of the forcings such as the geostrophic equilibrium and the boundary-layer depth, especially since a statistical extrapolation does suggest a dependence on the initial boundary-layer depth, but the data does not show the combination of PBL depth where this occurs.

Acknowledgements

The authors would like to thank Hendrik Wouters for his input and help in conceiving the idea behind this paper; and Aristofanis Tsiringakis for his help in calculating the geostrophic wind speed from WRF. Arjan Droste and Gert-Jan Steeneveld acknowledge funding from the Netherlands Organisation for Scientific Research (NWO) VIDI Grant ‘The Windy City’ (file no. 864.14.007).

7

Chapter 7

Synthesis

7.1 This work

This thesis has covered a broad topic: the urban atmosphere, analysed from different perspectives. While a typical scientific study typically demarcates a small area of interest within this broader topic, I have chosen to approach the field of urban meteorology from two angles: the observational and the modelling side. This Synthesis will be similarly broad: looking at the potential and pitfalls of opportunistic sensing, whether we can use the Urban Wind Island (Chapter 5) to our advantage or not, and what the future holds for the subjects covered by this thesis, and (urban) meteorology in general: both the observational and the modelling sides.

In Chapter 1 we have noted that the connection between these sides, and their occasional rivalry, stems from Antiquity, where the practical versus theoretical debate was as fierce as it is today. Even the potential of crowdsourcing was already acknowledged by Aristotle in his *Politics*, remarking that the many can outperform the individual (Waldron, 1995): what we now call the "Wisdom of the Crowd".

On the observational side of the spectrum, Chapter 2 answered the question whether crowd-sourced temperature data from smartphones could be useful to study urban air temperature at the hourly scale. In this Chapter we further refined an existing method to extract urban air temperatures from smartphone battery temperature data, to the point where Local Climate Zone differences in temperature were distinguishable, as well as the diurnal cycle of hourly urban air temperature. **This shows that even data obtained through such indirect sensors can yield valuable results at the hourly scale, albeit at lower quality than our usual high-end meteorological equipment would.**

In Chapter 3 the strong potential of Personal Weather Stations (PWSs) is explored further by their use in wind research, having knowledge of their usefulness in temperature and rain sensing. The Netatmo brand PWS allows citizens to setup an anemometer to measure wind speed where they choose to, typically a garden or balcony. Wind is turbulent by nature, and it seems that the crowdsourced wind observations have a high degree of uncertainty, caused by a combination of hardware limitations and an unknown station configuration. Regardless, the same principle of Chapter 2, the Wisdom of the Crowd, applies: many lower-quality observations can yield an average of sufficient quality, which is very similar to that of an established reference network. To answer the second research question of this thesis: **After correcting for bias and and filtering for precipitation, the quality-controlled PWS wind speed obtained over a longer duration and over several stations rather than one is usable for urban wind research.**

However, (urban) meteorology typically deals with day-to-day variations or even sub-daily processes, not the statistics gathered over several months. Chapter 4 showcased the capability of opportunistic sensing techniques to actually measure typical meteorological phenomena such as a cold front passage and a hot spell. The combination of PWS, smartphone and Commercial Microwave Link (CML) data to measure temperature, wind, rain and radiation showed the strengths of these techniques when combined with one another. The answer to RQ 3: **While the data shows differences when compared to the established reference, and more subtle variations in time or space are not captured, the overall picture can be painted very well with opportunistic sensing data.**

Chapters 5 and 6 deal with the other side of the spectrum: Chapter 5 uses a theoretical

model framework, the Mixed-Layer Model, capturing only the very essence of the urban and rural boundary layer, to investigate wind differences between these two surfaces under the same large-scale conditions. The positive urban wind anomaly, the Urban Wind Island (UWI), that arises from this theoretical model is small (~ 0.5 m/s) but persistent through the afternoon. That this phenomenon is not an artefact generated by the idealised model assumptions is proven through the use of the more sophisticated WRF 3D mesoscale model, where analysis of 2 summers of data shows the UWI to be a common phenomenon. These chapters answer RQs 4 and 5: **The UWI is generated through different boundary-layer growth dynamics between city and countryside in a conceptual model (RQ 4), and is shown in a three-dimensional modelling setting to also form during the evening through delayed transition towards the stable boundary layer, though its formation seems less reliant on the initial boundary-layer conditions. (RQ 5)**

Having answered these questions, some new ones have risen to the surface, that will be addressed in this synthesis. Firstly; given the shown potential of opportunistic sensing data: are they on their way to replace the status quo of meteorological observations, or are there other disadvantages to these techniques? Secondly; while the UWI persists even in a realistic modelling environment, is a 0.5 m/s boundary-layer averaged wind difference really that noteworthy? Is it relevant for practical purposes, or overshadowed by the turbulent nature of the urban surface? Is the UWI a benefit, a detriment, or just simply something we can ignore? The following sections will shed some light on these fresh questions, placing the subjects of this thesis in the broader scientific and societal context. Finally, I will wrap up with a Vision of the Future: a vision of crowdsourcing and citizen science; of the meteorology of the future, and our role as scientists in an increasingly sceptic world.

7.2 Crowdsourcing: sensing opportunities

The “big data” approach, turning millions of seemingly worthless data into a usable product, would seem to be the polar opposite of usual meteorological data gathering. For forecasting purposes, the nature of meteorological data needs to be highly accurate, which in turn means data is sparse. About 11,000 registered WMO surface monitoring stations measure weather across the entire world (https://cpdb.wmo.int/volume_a_observing_stations/list_stations), typically concentrated in developed countries. Compare that to the literal billions of smartphones present all over the globe (GSMA, 2019); the 250,000+ stations connected to the Weather Underground network (<https://www.wunderground.com/pws/overview>); or the strongly increasing amount of other smart devices such as smartwatches: the gap is quickly seen. Of course, comparing numbers is hardly fair when WMO stations are specifically built to measure climate as accurately as possible for decades, whereas a smartphone typically does not even contain an outside air thermometer nowadays and is quickly traded in for a newer model within years.

Regardless, while this obvious discrepancy exists, it is not a good reason to disregard the usage of unconventional devices outright. Opportunistic sensing data should always be complementary to existing, high-quality networks: where these networks show shortcomings or absence in general is where innovative sensing techniques can show their worth (Kamel Boulos et al., 2011; Muller et al., 2015; Chapman et al., 2017). Especially in areas which have very limited coverage by traditional measurement sources, opportunistic sensing can provide a basic level of knowledge of meteorological circumstances. Ways of obtaining precipitation data through CML



have been refined since 2005 (Upton et al., 2005), and have been very valuable in remote areas where rain radar does not penetrate (Uijlenhoet et al., 2018). Similarly, the very nature of opportunistic sensing data means that data is concentrated where people are, which is generally the urban areas, which the usual measurements are ill-equipped for.

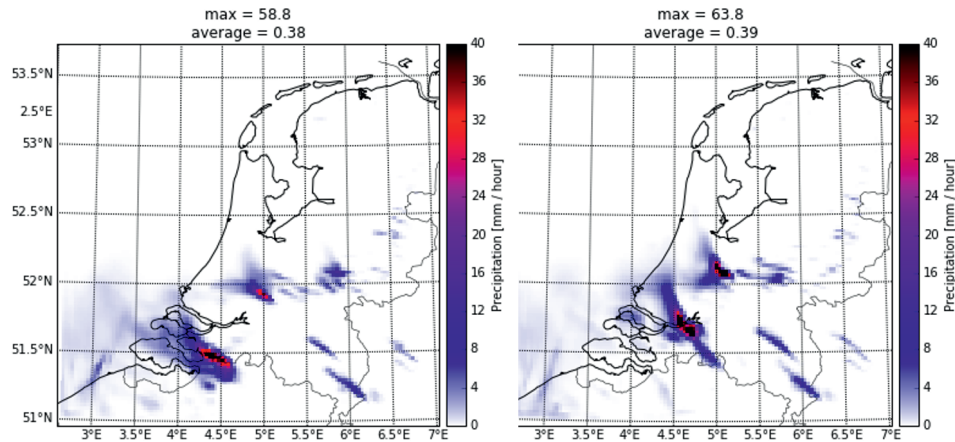


Figure 7.2.1: Impact of data-assimilation of PWS rainfall data (right-hand side), on a forecast of heavy precipitation over the Netherlands, 28 July 2014 11 UTC, using the WRF model. Model forecast is 7 hours ahead; run started at 4:00 UTC at 2.5 km resolution. The original forecast without data assimilation is on the left. Image courtesy of Sytse Koopmans, Gert-Jan Steeneveld and Ronald van Haren.

Modelling efforts can also benefit through opportunistic sensing data: for instance by data assimilation of PWS data. Figure 7.2.1 shows preliminary results of assimilating data of 68 PWSs into WRF, for a heavy precipitation event in the Netherlands, July 28 2014. Temperature, dew point temperature and air pressure data from one specific type of PWS with low measurement errors (Davis Vantage Pro; Bell et al. (2015)) are assimilated into the WRF model to nudge the forecasts. The shown forecast is 7 hours ahead, and by ingesting PWS data the location and intensity of the precipitation event correspond better to radar observations, though especially location is still somewhat inaccurate. Other applications of opportunistic sensing data include the use of pressure data obtained through smartphones to improve location of frontal zones (Mass and Madaus, 2014; Madaus and Mass, 2017; Hintz et al., 2019), or using PWS measurements for the validation of local processes such as deep convection (Mandement and Caumont, 2019).

Increasingly, crowdsourced data are finding their way into the established meteorological institutes: the ECMWF acknowledges the value of crowdsourced data, which can "[capture] rapid temporal/spatial variations in weather parameters; enable new quality control methods (e.g. artificial intelligence) which depend on the availability of large amounts of data; verify small-scale features and high impact weather, allowing the highlighting of situation-dependent model limi-

tations; enable more robust upscaling of weather parameters for verification purposes" (Haiden et al. (2018), p. 5). They list a number of requirements such data need to adhere to, mainly real-time availability and accurate metadata in order to get an understanding of measurement uncertainty and develop quality assurance. Such recommendations are given in many scientific articles working with crowdsourced data (e.g. Fenner et al., 2017; Gharesifard et al., 2017; de Vos et al., 2018; Chapman and Bell, 2018), including chapters 2–4 in this thesis. National weather services could provide a central hub in order to impose some regulations on a widely unregulated source of data, and efforts are being made such as the WOW-network for PWS data (<https://wow.metoffice.gov.uk/>), run by the English MetOffice, to which 10.000 stations in 220 countries are now connected (though data this is not yet used in data assimilation, rather for verification purposes).

Krennert et al. (2018) provide an inventory of crowdsourced data used at national weather services, such as the Norwegian weather service improving its forecasts using purchased Netatmo data. This last sentence highlights the potential danger of using crowdsourced data: large amounts are in commercial hands. In some cases these data are "leftovers" generated by companies: used for their own quality-control, or essentially just an artefact of generating their actual end-product. The telecom companies running the CML towers use the attenuation data as quality-control rather than derive a product from it. Similarly, the OpenSignal company behind the app used to derive air temperature from smartphone battery thermometers (Ch. 2) is interested in the communication signal the phone receives, rather than its battery. In these cases, companies would be more willing to share this data for free (or a nominal fee for processing the data), though data of this type tend not to comply to the real-time availability the ECMWF recommends. However, when the meteorological data is the prime objective of the company behind the devices generating that data, the main advantage of crowdsourcing might quickly fade. The Netatmo company has sold datasets to the Norwegian weather service, for their use short-term forecasts (nowcasting) and correction of meteorological products (Nipen et al., 2019).

Acquisition of datasets is in itself not a problem: we could see purchasing instruments in a similar light, where we pay to create data; data from measurement campaigns is often not free for use by other parties either. The issue lies with the continuity of the data platform: if the Netatmo company shifts its business model away from weather stations, this whole extensive network could collapse if there is no alternative. The other way around is also a potential problem: the Weather Underground platform shut down its previously free API, allowing access to its data, and replaced it with a paid format. The data might be crowdsourced, with thousands of people contributing to one goal: but the platform where all of this data collected is ultimately in private hands, a single company with economical profits as their goal.

This is one of the issues of relying on opportunistic sensing techniques: science moves slowly, whereas industry moves fast. By the time a proper quality-assurance procedure is set up for a new type of measurement to make it ready to be accepted into mainstream scientific use, the industry might be miles away, having already discarded that instrument in favour of an upgraded version. For instance, the Netatmo wind sensor, subject of Chapter 3, was released in 2015, and this Chapter is the first work to even use the data and provide a first quality assessment, 4 years later. Technological advancement can also have unintended side-effects for science: the move towards a 5G network, operating at higher frequencies, has the unfortunate side-effect of potentially disrupting satellite measurements of water vapour (Witze, 2019). While a stronger mobile network enables a wider range of opportunistic sensing opportunities, it can threaten existing established



networks in this way, or even existing opportunistic sensing networks: CML towers upgrading to 5G may no longer be able to sense rainfall in the same way, requiring new retrieval techniques.

Since technological advancements of smart devices are not driven by scientific questions, but by profits, potentially interesting scientific uses may be scrapped in later releases. Such was the case of meteorological sensors on smartphones: the Samsung S4 model was equipped with an actual air thermometer and hygrometer, but subsequent series have removed some of the environmental sensors (though a barometer and light sensor are typically present in all current models). Such external sensors seem to be pricey and could impede the design of the phone, without providing a direct essential function for the phone usage. Rather, several apps have been developed to turn the phone into a makeshift thermometer, often based on the battery temperature conversion. Data from a direct smartphone air temperature sensor could have been very valuable, even if it suffers from the human behaviour problems as described in (Overeem et al., 2013b; Droste et al., 2017). Preliminary tests of smartphone-measured air temperature (direct measurements as opposed to the battery-derived temperatures from Chapters 2 and 4) show that the smartphone temperature data is negatively biased and strongly influenced by radiation, since the thermometer is not shielded (Noyola Cabrera (2018), Figure 7.2.2).

Windspeed/Radiation	<0 W/m ²	0 - 100 W/m ²	100 - 200 W/m ²	200 - 300 W/m ²	300 - 400 W/m ²	400 - 500 W/m ²	500 - 600 W/m ²	600 - 700 W/m ²	700 - 800 W/m ²	800 - 900 W/m ²	> 900 W/m ²
<0.5 m/s	-4.2	-3.1	-2.5	-2.3	-2.2				13.5		
0.5 - 2 m/s	-3.9	-3.2	-1.72	-1.6	-1.4	-0.5	3.35	7.9	11.3		
2 - 3.5 m/s	-3.7	-3.1	-1.7	-0.3	0.9	0.3	2.7	5.7	8.5	8	
3.5 - 5 m/s						2.5	6.9	8.2	8.5	9.7	
5 - 6.5 m/s										9	

Figure 7.2.2: Mean error of smartphone temperature ($^{\circ}$ C) readings as a function of wind speed and global radiation (direct + diffuse). Image courtesy of Aly Noyola Cabrera (Figure 18 in his MSc thesis).

Ultimately, human behaviour might be the key factor causing uncertainty for a large number of opportunistic sensing techniques (Lin et al., 2016). For smartphone-derived data, the whereabouts and routine of the person using the phone is crucial to understand the data they produce. Indoor or outdoor; whether the phone is in a pocket, on a surface or in someone's hand; whether the phone is out in the sun or shade, in an air-conditioned building or not, at the ground floor or several floors above ground level, etc. The possible circumstances influencing the raw opportunistic sensing data are endless, and so the causes of error are endless. In a lab setting some manner of correction is possible for these error sources, but they are difficult to translate to actual crowdsourced-obtained data (Chau, 2019). We can try to correct for average behaviour, as we invoke the Wisdom of the Crowd once again, but we cannot ethically correct for individual behaviour. When error courses are known and can be corrected for, the potential of these measurements is enormous: bias-corrected smartphone temperature measurements carried out in Wageningen by Noyola Cabrera (2018), show very good agreement with transect measurements of verified measurement equipment. The smartphone measurements are able to capture spatial and temporal temperature differences to a good degree (Figure 7.2.3), which could translate very well to the heterogeneous urban microclimate.

The rise of big data has not been a wholly positive one: scandals such as the Cambridge

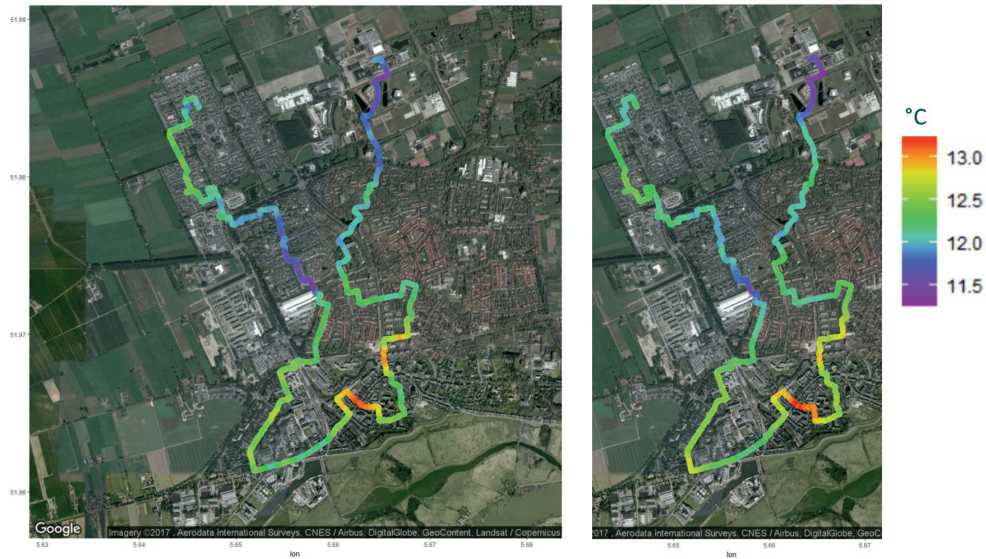


Figure 7.2.3: Left: air temperature readings a mobile measurement platform of known quality. Right: air temperature readings from the smartphone after bias correction. Transect measurements made in March 2018, Wageningen, the Netherlands. Image courtesy of Aly Noyola Cabrera (Figure 26 in his MSc thesis).

Analytica misuse of private data have tainted big data for the general public, as we collectively have come to realise how much valuable information can be distilled from the data we generate through our daily routines (Isaak and Hanna, 2018). Privacy laws have been made stricter with regards to the use of personal data, and the kind of information needed to correct environmental monitoring for human behaviour steers dangerously close to very private information. While we may (profess to) have a pure goal in mind, e.g. improving smartphone-derived measurements by correcting for measurements taken in vehicles, a less savoury scientist might use that information to deduce where a single smartphone owner works and lives by analysing their movement pattern over a larger amount of time. As such, we as scientists using opportunistic sensing data are split between the scientific curiosity, and a moral obligation to respect privacy, and we will have to make find clever ways to use the data we have to its best purpose.

7.3 The Urban Wind Island

The first parts of this thesis have been largely focused on unconventional observation practices, by aiming to turn opportunistic sensing observations into meteorological output of sufficient quality. On the modelling side however, the usage of conceptual models is not exactly new or ground-breaking for meteorology. One could say modern meteorology is built around weather models (fed by high-quality observations) which all started in such a simple, conceptual manner, and the real new things in that regard are the ever-increasing resolution and accuracy of models. Such advances bring their own unique challenges to meteorology (Baklanov et al., 2011), especially when scales are being pushed far enough to solve processes that were previously parametrised.



For instance, the ECMWF aims to model the entire globe at a staggering 1 km resolution as a trial, to explicitly solve deep convection (ECMWF, 2020). Such process understanding could make great advances in the skill of numerical weather prediction.

Rather than the methodology, let us then consider the "new" phenomenon this thesis has introduced through model usage: the Urban Wind Island effect. The UWI has been shown to exist both in a conceptual modelling framework which uses the essential physics of the boundary layer, as well as in the 3-dimensional modelling environment of the WRF mesoscale model. It is important to once again note that the UWI is a boundary-layer scale phenomenon: separate from the hyper-localised wind gusts present around buildings within the urban canyon. This is the realm of very fine Large Eddy Simulation or Direct Numerical Simulation models, as opposed to the (relative) coarse-resolution mesoscale models used for weather forecasting (such as WRF). These wind gusts can present a real danger to pedestrians (Bottema, 2000; Blocken and Carmeliet, 2004), and knowledge of wind load on building fronts is essential for a safe built environment.

In contrast, the UWI is a small effect, a positive urban wind anomaly in the order of 0.5 m/s, that manifests itself mostly during the afternoon and early night on nice summer days. We picked its name to deliberately correspond to the well-known Urban Heat Island, the positive temperature anomaly that can cause elevated levels of heat stress during hot days and nights. The UHI is an obvious threat to human health: since it is strongest during night-time, sleeping patterns can be disrupted for urban dwellers since they are less capable of cooling down, which increases the risk of cardiovascular diseases, especially among the elderly or the infirm (Kovats and Hajat, 2008).

What about the UWI then? The dangers of wind are instinctively more associated with high wind speeds: storms, disruptive wind tunnels in streets (Bottema, 1999), or cold stress: e.g., the Wind Chill thermal comfort index (Steadman, 1971). It is not difficult to imagine that a ~ 1 m/s difference is not going to matter much when average wind speeds are already fairly high. Wind speeds above ~ 5 m/s are broadly categorised as "disrupting" (Penwarden, 1973), but such high wind speeds are mostly seen in autumn or winter, associated with storm activity rather than boundary-layer mechanisms. The UWI would therefore not influence these disruptive winds much: this is disregarding that the UWI might be much lower during winter since the boundary layer is less convective, so urban-rural contrasts will be lower. More broadly: cold stress is easier to deal with than heat. You can dress for the cold, but there is a strict (legal) limit to how much you can dress down for the heat.

However, wind speed is an important component for heat related stress as well: hot days can be made much more bearable if there is even the slightest gust of wind to cool the body down and promote evaporation of sweat. The effect of wind speed on thermal comfort is strongly non-linear and generally at its maximum for the very low, or null, wind speeds. Because of this non-linearity, the way thermal comfort indices account for the effect of wind speed is highly variable: some simple indices use a simple linear relation, while more complex energy-balance based indices account for its non-linearity (Blażejczyk et al., 2012; Coccolo et al., 2016). Because of the strong influence of wind speed at low wind speed, even a 0.5 m/s increase through the UWI could greatly improve thermal comfort (Coccolo et al., 2017). This even holds for inside temperatures, by making use of natural ventilation to drive down indoor temperature (Aynsley, 1999; Lien and Ahme, 2011). Not only heat-related issues can be mitigated by wind, but urban air quality

as well, by dispersing pollutants more effectively (Cheng et al., 2008). In the Gaussian plume dispersion model (Awasthi et al., 2006), concentration of a pollutant is inversely proportional to the wind speed, so especially at low mean wind speed, the UWI contribution could significantly reduce pollutant concentrations.

The geometry of a city (its density, height, aspect ratio) determine how the city reacts to the wind passing over it: as a whole construct analogous to the city in Chapter 5, or as individual building blocks disrupting flow (Buccolieri et al., 2010). A denser, more homogeneous city will be akin to a single block, with few winds from the inertial sub-layer penetrating down into the urban canyon: the "skimming" regime as described in (Oke, 1988). When variability in building height is larger, and the aspect ratio decreases, a turbulent wind regime forms with more vertical motions that move down into the urban canyon (Chapter 4 in Oke et al. (2017)) Related to the UWI, which should be seen as a phenomenon relating to the inertial sub-layer, this means that for the skimming regime, the UWI will not penetrate downward towards a pedestrian level, since the canopy layer is partially decoupled from the air above due to the density of the built environment. The rapid variations in wind speed we feel as pedestrians in the urban canopy layer are much more governed by the local obstacles and roughness, as opposed to a city-wide effect. For less dense urban structures, there is more vertical exchange of air, so it is possible that the UWI could translate downwards to a pedestrian effect, but this is highly speculative, and would require an in-depth study. In general terms, it is important to realise the difference between scales: the UWI operates on a boundary-layer scale, whereas individual human thermal comfort is at the local scale, where extremely large gradients can occur. On that individual scale, the UWI influence is small; however, when considering the health of the city as a whole, it can certainly contribute through a small but persistent effect.

7.4 Vision of the Future

I have started this thesis by noting that meteorology is an old discipline with firm roots in Antiquity. Its methods and practitioners have changed substantially over the past 2,500 years, but its goal has always been the same: understanding and predicting the weather for our own benefit. We have come a long way since then: from divine intervention and interpreting celestial patterns, to highly accurate computational fluid dynamics and models forecasting the weather over the entire globe at an incredible skill (Haiden et al., 2019). And yet, despite all of this scientific and technological advancement, there are still many more questions to be answered regarding the atmosphere, and how it influences our lives. Rapid urbanisation, a problem highlighted in nearly every single urban research paper, has caused us to realise that the human environment interacts with the atmosphere as much as the other way around. Climate change, the energy transition, global population growth: all these change how urban life will be in the future, and thus what role urban meteorology will play in our society. Already there is much more attention on a policy level for the "green city" of the future: a healthy urban space with low emissions and even self sufficient in food and/or energy sources (e.g. the Green Cities EU project <https://www.thegreencity.eu/>). Such a city would need to be designed to draw from the environment: cool spaces during summertime; integrated energy production through solar and wind power, while also keeping harmful cold stress out and being resilient against flooding and storm damages. Urban meteorological measurements and modelling can contribute to the planning, and maintenance, of these cities of the future (Baklanov et al., 2018).



7.4.1 Forging Connections through Citizen Science

Such a future green city is a hopeful, optimistic Vision of the Future. But, current signs indicate a darker tomorrow. Carbon emissions are not going down, despite years of overwhelming scientific evidence and pledges by companies and nations (Peters et al., 2020). Global economic inequality is increasing, and climate change is disproportionately affecting poorer countries (Diffenbaugh and Burke, 2019), suggesting that these green cities of the future will only be a reality for the more well-off regions and people. There seems to be a discrepancy between science and the general public: a form of mutual misunderstanding where scientists cannot understand why their (apparently) clear results are not being understood, while the public considers them disconnected from "real life". This is a worrying prospect that will only worsen if we, as the scientific community, dismiss this as an acute issue. Without interest and participation in science by the public, the very authority of science will diminish, if it cannot connect in a meaningful way to the inhabitants of the world it studies.

In this light, there is a place for one of the subjects of this thesis: crowdsourcing and citizen science, which connects activities of the general public to scientific work. For instance, the crowdsourcing platforms WOW-UK and WOW-NL (run by the English and Dutch national weather services, respectively) offer opportunities for schools to participate in their network by supplying them their own weather station, to teach children about weather and climate using their own measurements (<https://wow.metoffice.gov.uk/education>). In a world that is getting increasingly sceptic towards science and scientific institutions, letting citizens (children and adults alike) participate and contribute towards science through their own experience could create a more positive view of science and its place in society.

Campaigns where citizens can actively contribute to improve their own neighbourhood, e.g. by measuring particulate matter (Snik et al., 2014) or urban heat to make their street "climate proof" (e.g. through the "Measuring your City" project in Amersfoort, the Netherlands, Mureau et al. (2018)) have been very successful, with citizens more involved and positive about changes in their living environment (Paulos et al., 2009). Studies relying on citizen science are getting more common and accepted, though the scientific community is still slow to accept these new forms of data (Ottinger, 2010; Bonney et al., 2014): if the public is to embrace science, then science has to be more open to contributions from citizens. In that way, there are many more exciting opportunities to develop opportunistic sensing strategies for urban meteorology. With an increase in measuring devices, we could even use PWS data to estimate urban fluxes (following the Holtslag and Van Ulden (1983) method); identify air pollution and improve urban health through fitness app data (Sun and Mobasheri, 2017); and create a critical, science-minded society in the process.

7.4.2 Data-driven meteorology?

This increase in citizen science goes hand in hand with the increased possibilities of opportunistic sensing, since both have their origin in the increased global interconnectivity. It is easier than ever to record, share and search data of any possible kind through the Internet of Things. The increased level of technology in our everyday lives allows for much more participation in science,

and cross-contamination of ideas (i.e. science answering questions from the public domain). In a similar manner, the increased level of technology in the scientific field has not only given rise to more sophisticated models or higher resolution weather forecasting, but whole new techniques based on giant amounts of data, as this thesis has successfully shown. The use of machine learning to recognise patterns within large datasets is rising dramatically in all scientific fields; in meteorology it is being used to for example train climate models (Boukabara et al., 2019) or classify extreme weather patterns (Racah et al., 2017). Machine learning approaches are a double-edged blade for scientists. Its benefits are that it can be used to speed up traditionally expensive parts of complex models such as the radiation scheme (Krasnopolsky et al., 2008) or convection schemes (Gentine et al., 2018), or forecasting highly variable processes such as solar irradiance for renewable energy purposes (Voyant et al., 2017).

In contrast, there is the danger of losing physical understanding once machine learning is used to replace physical understanding in a model. While there are currently no forecast models based on neural network approaches that rival physics-based model in forecast skill (Dueben and Bauer, 2018), this could be a reality in the future. Another downside of machine learning models is that they are based on training data, which can make it difficult to extrapolate to entirely new situations, especially if the physical basis beneath the model is (partially) replaced by machine learning. That could make long-range climate forecasts, or forecasts of exceptionally rare phenomena for which little historical evidence exists, very uncertain. A synergy between physics-driven (traditional) modelling, and data-driven (machine-learning) approaches is therefore necessary, to draw upon the strengths of both methods, while ensuring the quality of the resulting outcomes (Reichstein et al., 2019).

7.4.3 Final Thoughts

To sum it all up, my vision of the future of meteorology is an integration of established techniques and process understanding, coupled with innovative new data and modelling approaches. Its goal has not changed, and will not change: to understand and forecast the fickle nature of the atmosphere for our own benefit. Whether that is to design a healthier urban living space, improve early warning systems, or forecast solar and wind energy resources at an even more precise level, meteorological research is always connected to an societal issue in the end. That does not mean fundamental research is pointless: increased understanding of the core processes driving our atmosphere leads to better forecasts further down the line. As much as society needs meteorology, so does meteorology need feedback from society, to be engaged in its research and to stimulate new exciting questions to be answered and fields of research to be discovered.

This thesis now forms a part of the impressive lexicon of meteorological knowledge, formed over 25 centuries. Whether it will withstand the jaws of time, and still be relevant 25 centuries hence, remains a question I can only guess at. That seems unlikely; but it is my hope that this thesis has added some novel perspectives to the current-day understanding of the urban atmosphere. At the very least, it has been fun to write.

References

- Agüera-Pérez, A., J. C. Palomares-Salas, J. J. G. de la Rosa, and J. M. Sierra-Fernández, 2014: Regional wind monitoring system based on multiple sensor networks: A crowdsourcing preliminary test. *Journal of Wind Engineering and Industrial Aerodynamics*, **127**, 51–58, doi: 10.1016/j.jweia.2014.02.006.
- Allamano, P., A. Croci, and F. Laio, 2015: Toward the camera rain gauge. *Water Resources Research*, **51** (3), 1744–1757, doi: 10.1002/2014WR016298.
- Andersson, J. C. M., P. Berg, J. Hansryd, A. Jacobsson, J. Olsson, and J. Wallin, 2017: Microwave links improve operational rainfall monitoring in Gothenburg Sweden. *Proceedings of the 1st International Conference on Environmental Science and Technology*, 1–4.
- Android, 2019: Android Developer Guide for Sensor Management. Android Developers, URL <http://developer.android.com/reference/android/hardware/SensorManager.html>.
- Aristotle, 1978: *Meteorologica*, Vol. 397. Loeb Classical Library, translated by H.D.P. Lee, 1978.
- Arnfield, A. J., 2003: Two decades of urban climate research: A review of turbulence, exchanges of energy and water, and the urban heat island. *International Journal of Climatology*, **23** (1), 1–26, doi: 10.1002/joc.859.
- ASHRAE, 2010: Thermal environmental conditions for human occupancy. Tech. rep., ANSI/ASHRAE Standard 55-2010, 39 pp.
- Atlas, D., and C. W. Ulbrich, 1977: Path-and area-integrated rainfall measurement by microwave attenuation in the 1–3 cm band. *Journal of Applied Meteorology*, **16** (12), 1322–1331.
- Awasthi, S., M. Khare, and P. Gargava, 2006: General plume dispersion model (GPDM) for point source emission. *Environmental Modeling & Assessment*, **11** (3), 267–276, doi: 10.1007/s10666-006-9041-y.
- Aynsley, R., 1999: Estimating summer wind driven natural ventilation potential for indoor thermal comfort. *Journal of Wind Engineering and Industrial Aerodynamics*, **83** (1-3), 515–525, doi: 10.1016/S0167-6105(99)00098-7.
- Baklanov, A., M. Alexander, and R. Sokhi, Eds., 2011: *Integrated Systems of Meso-Meteorological and Chemical Transport Models*. Springer Berlin Heidelberg, Berlin, Heidelberg, doi: 10.1007/978-3-642-13980-2.
- Baklanov, A., and Coauthors, 2018: From urban meteorology, climate and environment research to integrated city services. *Urban Climate*, **23**, doi: 10.1016/j.uclim.2017.05.004.

- Balsamo, G., and Coauthors, 2018: Satellite and In Situ Observations for Advancing Global Earth Surface Modelling: A Review. *Remote Sensing*, **10** (12), 2038.
- Barlow, J., and Coauthors, 2017: Developing a research strategy to better understand, observe and simulate urban atmospheric processes at kilometre to sub-kilometre scales. *Bulletin of the American Meteorological Society*, BAMS-D-17-0106.1, doi: 10.1175/BAMS-D-17-0106.1, URL <http://journals.ametsoc.org/doi/10.1175/BAMS-D-17-0106.1>.
- Barlow, J. F., 2014: Progress in observing and modelling the urban boundary layer. *Urban Climate*, **10** (P2), 216–240, doi: 10.1016/j.uclim.2014.03.011.
- Barlow, J. F., T. M. Dunbar, E. G. Nemitz, C. R. Wood, M. W. Gallagher, F. Davies, E. O'Connor, and R. M. Harrison, 2011: Boundary layer dynamics over London, UK, as observed using Doppler lidar during REPARTEE-II. *Atmospheric Chemistry and Physics*, **11** (5), 2111–2125, doi: 10.5194/acp-11-2111-2011.
- Barlow, J. F., C. H. Halios, S. E. Lane, and C. R. Wood, 2015: Observations of urban boundary layer structure during a strong urban heat island event. *Environmental Fluid Mechanics*, **15** (2), 373–398, doi: 10.1007/s10652-014-9335-6.
- Bassett, R., X. Cai, L. Chapman, C. Heaviside, and J. E. Thornes, 2019: Semi-idealized urban heat advection simulations using the Weather Research and Forecasting mesoscale model. *International Journal of Climatology*, **39** (3), 1345–1358, doi: 10.1002/joc.5885.
- Beekhuis, H., and T. Mathijssen, 2018: From pulse to product, Highlights of the upgrade project of the Dutch national weather radar network. *10th European Conference on Radar in Meteorology and Hydrology (ERAD 2018)*, 960–965, doi: 10.18174/454537.
- Beljaars, A. C. M., and F. C. Bosveld, 1997: Cabauw Data for the Validation of Land Surface Parameterization Schemes. *Journal of Climate*, **10** (6), 1172–1193, doi: 10.1175/1520-0442(1997)010<1172:CDFTVO>2.0.CO;2.
- Bell, S., D. Cornford, and L. Bastin, 2013: The state of automated amateur weather observations. *Weather*, **68** (2), 31–36, doi: 10.1002/wea.1980.
- Bell, S., D. Cornford, and L. Bastin, 2015: How good are citizen weather stations? Addressing a biased opinion. *Weather*, **70** (3), 75–84, doi: 10.1002/wea.2316.
- Bense, V. F., and Coauthors, 2016: Distributed Temperature Sensing as a downhole tool in hydrogeology. *Water Resources Research*, **52** (12), 9259–9273, doi: 10.1002/2016WR018869.
- Best, M. J., and C. S. B. Grimmond, 2015: Key Conclusions of the First International Urban Land Surface Model Comparison Project. *Bulletin of the American Meteorological Society*, **96** (5), 805–819, doi: 10.1175/BAMS-D-14-00122.1.
- Blackadar, A. K., 1957: Boundary Layer Wind Maxima and Their Significance for the Growth of Nocturnal Inversions. *Bulletin of the American Meteorological Society*, **38**, 283–290.
- Blazejczyk, K., Y. Epstein, G. Jendritzky, H. Staiger, and B. Tinz, 2012: Comparison of UTCI to selected thermal indices. *International Journal of Biometeorology*, **56** (3), 515–535, doi: 10.1007/s00484-011-0453-2.

- Blocken, B., and J. Carmeliet, 2004: Pedestrian Wind Environment around Buildings: Literature Review and Practical Examples. *Journal of Thermal Envelope and Building Science*, **28** (2), 107–159, doi: 10.1177/1097196304044396.
- Blocken, B., W. D. Janssen, and T. van Hooff, 2012: CFD simulation for pedestrian wind comfort and wind safety in urban areas: General decision framework and case study for the Eindhoven University campus. *Environmental Modelling & Software*, **30**, 15–34, doi: 10.1016/j.envsoft.2011.11.009.
- Bohnenstengel, S. I., and Coauthors, 2015: Meteorology, Air Quality, and Health in London: The ClearfLo Project. *Bulletin of the American Meteorological Society*, **96** (5), 779–804, doi: 10.1175/BAMS-D-12-00245.1.
- Bonney, R., J. L. Shirk, T. B. Phillips, A. Wiggins, H. L. Ballard, A. J. Miller-Rushing, and J. K. Parrish, 2014: Next Steps for Citizen Science. *Science*, **343** (6178), 1436–1437, doi: 10.1126/science.1251554.
- Bornstein, R. D., and D. S. Johnson, 1977: Urban-rural wind velocity differences. *Atmospheric Environment (1967)*, **11** (7), 597–604, doi: 10.1016/0004-6981(77)90112-3.
- Bottema, M., 1999: Towards rules of thumb for wind comfort and air quality. *Atmospheric Environment*, **33** (24–25), 4009–4017, doi: 10.1016/S1352-2310(99)00142-9.
- Bottema, M., 2000: A method for optimisation of wind discomfort criteria. *Building and Environment*, **35** (1), 1–18, doi: 10.1016/S0360-1323(98)00065-1.
- Boukabara, S.-A., V. Krasnopolsky, J. Q. Stewart, E. S. Maddy, N. Shahroudi, and R. N. Hoffman, 2019: Leveraging Modern Artificial Intelligence for Remote Sensing and NWP: Benefits and Challenges. *Bulletin of the American Meteorological Society*, **100**, 473–491, doi: 10.1175/bams-d-18-0324.1.
- Bowker, D., 2011: Meteorology and the ancient Greeks. *Weather*, **66** (9), 249–251, doi: 10.1002/wea.758.
- Brandsma, T., and J. P. van der Meulen, 2008: Thermometer screen intercomparison in De Bilt (the Netherlands)—Part II: description and modeling of mean temperature differences and extremes. *International Journal of Climatology*, **28** (3), 389–400, doi: 10.1002/joc.1524.
- Buccolieri, R., M. Sandberg, and S. Di Sabatino, 2010: City breathability and its link to pollutant concentration distribution within urban-like geometries. *Atmospheric Environment*, **44** (15), 1894–1903, doi: 10.1016/j.atmosenv.2010.02.022.
- Buik, N., P. Stolk, and G. Willemsen, 2004: Analyse van temperatuurmetingen in de Nederlandse ondergrond. *Stromingen*, **10** (4), 13–22.
- Businger, J. A., J. C. Wyngaard, Y. Izumi, and E. F. Bradley, 1971: Flux-Profile Relationships in the Atmospheric Surface Layer. *Journal of the Atmospheric Sciences*, **28** (2), 181–189, doi: 10.1175/1520-0469(1971)028<0181:FPRITA>2.0.CO;2.
- Byun, D. W., and S. P. S. Arya, 1986: A study of mixed-layer momentum evolution. *Atmospheric Environment (1967)*, **20**, 715–728, doi: 10.1016/0004-6981(86)90186-1.

- Byun, D. W., and S. P. S. Arya, 1990: A two-dimensional mesoscale numerical model of an urban mixed layer-I. Model formulation, surface energy budget, and mixed layer dynamics. *Atmospheric Environment Part A, General Topics*, **24**, 829–844, doi: 10.1016/0960-1686(90)90284-T.
- Campbell Scientific, 2017: CSAT3 Three Dimensional Sonic Anemometer Instruction Manual. URL <https://s.campbellsci.com/documents/us/manuals/csat3.pdf>, 74 pp.
- Carpentieri, M., and A. G. Robins, 2015: Influence of urban morphology on air flow over building arrays. *Journal of Wind Engineering and Industrial Aerodynamics*, **145**, 61–74, doi: 10.1016/j.jweia.2015.06.001.
- Carta, J. A., P. Ramírez, and S. Velázquez, 2009: A review of wind speed probability distributions used in wind energy analysis. Case studies in the Canary Islands. *Renewable and Sustainable Energy Reviews*, **13** (5), 933–955, doi: 10.1016/j.rser.2008.05.005.
- Castro, I. P., 2017: Are Urban-Canopy Velocity Profiles Exponential? *Boundary-Layer Meteorology*, **164** (3), 337–351, doi: 10.1007/s10546-017-0258-x.
- Celik, A. N., 2004: On the distributional parameters used in assessment of the suitability of wind speed probability density functions. *Energy Conversion and Management*, **45** (11–12), 1735–1747, doi: 10.1016/j.enconman.2003.09.027.
- Chandler, T., 1965: *The Climate of London*. Hutchinson & Co, London, 292 pp.
- Chapman, L., C. Bell, and S. Bell, 2017: Can the crowdsourcing data paradigm take atmospheric science to a new level? A case study of the urban heat island of London quantified using Netatmo weather stations. *International Journal of Climatology*, **37** (9), 3597–3605, doi: 10.1002/joc.4940.
- Chapman, L., and S. J. Bell, 2018: High-resolution monitoring of weather impacts on infrastructure networks using the internet of things. *Bulletin of the American Meteorological Society*, **99** (6), 1147–1154, doi: 10.1175/BAMS-D-17-0214.1.
- Chau, N. H., 2019: Estimation of air temperature using smartphones in different contexts. *Journal of Information and Telecommunication*, **3** (4), 1–14, doi: 10.1080/24751839.2019.1634869.
- Chen, A. B., M. Behl, and J. L. Goodall, 2018: Trust me, my neighbors say it’s raining outside: Ensuring data trustworthiness for crowdsourced weather stations. *Proceedings of the 5th Conference on Systems for Built Environments*, ACM, 25–28.
- Cheng, W., C.-H. Liu, and D. Y. Leung, 2008: Computational formulation for the evaluation of street canyon ventilation and pollutant removal performance. *Atmospheric Environment*, **42** (40), 9041–9051, doi: 10.1016/j.atmosenv.2008.09.045.
- Chua, S., 2009: Light vs. distance. UC Berkeley, URL http://bccp.berkeley.edu/o/Academy/workshop_09/pdfs/InverseSquareLawPresentation.pdf, University Lecture, Powerpoint Slides.

- Chwala, C., F. Keis, and H. Kunstmann, 2016: Real-time data acquisition of commercial microwave link networks for hydrometeorological applications. *Atmospheric Measurement Techniques*, **9** (3), 991–999.
- Chwala, C., and H. Kunstmann, 2019: Commercial microwave link networks for rainfall observation: Assessment of the current status and future challenges. Accepted.
- Chwala, C., G. Smiatek, and H. Kunstmann, 2018: Real-time rainfall from a country-wide network of commercial microwave links in Germany. *10th European Conference on Radar in Meteorology and Hydrology (ERAD 2018)*, 265, doi: 10.18174/454537.
- Chwala, C., and Coauthors, 2012: Precipitation observation using microwave backhaul links in the alpine and pre-alpine region of Southern Germany. *Hydrology and Earth System Sciences*, **16** (8), 2647–2661.
- Coccolo, S., J. Kämpf, J.-L. Scartezzini, and D. Pearlmutter, 2016: Outdoor human comfort and thermal stress: A comprehensive review on models and standards. *Urban Climate*, **18**, 33–57, doi: 10.1016/j.uclim.2016.08.004.
- Coccolo, S., D. Mauree, E. Naboni, J. Kaempf, and J.-L. Scartezzini, 2017: On the impact of the wind speed on the outdoor human comfort: a sensitivity analysis. *Energy Procedia*, **122**, 481–486, doi: 10.1016/j.egypro.2017.07.297.
- Conradsen, K., L. B. Nielsen, and L. P. Prahm, 1984: Review of Weibull Statistics for Estimation of Wind Speed Distributions. *Journal of Climate and Applied Meteorology*, **23** (8), 1173–1183, doi: 10.1175/1520-0450(1984)023<1173:ROWSFE>2.0.CO;2.
- Conzemius, R. J., and E. Fedorovich, 2006a: Dynamics of Sheared Convective Boundary Layer Entrainment. Part I: Methodological Background and Large-Eddy Simulations. *Journal of the Atmospheric Sciences*, **63** (4), 1151–1178, doi: 10.1175/JAS3691.1.
- Conzemius, R. J., and E. Fedorovich, 2006b: Dynamics of Sheared Convective Boundary Layer Entrainment. Part II: Evaluation of Bulk Model Predictions of Entrainment Flux. *Journal of the Atmospheric Sciences*, **63** (4), 1179–1199, doi: 10.1175/JAS3696.1.
- Dalyot, S., and S. Sosko, 2015: Towards the use of Crowdsourced Volunteered Meteorological Data for Forest Fire Monitoring. *GEOProcessing 2015, The Seventh International Conference on Advanced Geographic Information Systems, Applications, and Services*, C. P. Ruckemann, and Y. Doytsher, Eds., IARIA, Lisbon, Portugal, 127–132.
- de Bruijn, E. I. F., S. de Haan, F. C. Bosveld, B. W. Schreur, and A. A. M. Holtslag, 2016: Observing Boundary-Layer Winds from Hot-Air Balloon Flights. *Weather and Forecasting*, **31** (5), 1451–1463, doi: 10.1175/WAF-D-16-0028.1.
- de Haan, S., 2011: High-resolution wind and temperature observations from aircraft tracked by Mode-S air traffic control radar. *Journal of Geophysical Research: Atmospheres*, **116** (D10), doi: 10.1029/2010JD015264.
- de Haij, M., 2009: Automatische validatie van druk- en windwaarnemingen op het KNMI - een verkenning. Tech. rep., 84 pp. URL <http://bibliotheek.knmi.nl/knmipubIR/IR2009-03.pdf>, in Dutch.

- de Vos, L., A. Droste, M. Zander, A. Overeem, H. Leijnse, B. Heusinkveld, G. Steeneveld, and R. Uijlenhoet, 2019a: Hydrometeorological monitoring using opportunistic sensing networks in the Amsterdam metropolitan area. *Bulletin of the American Meteorological Society*, BAMS-D-19-0091.1, doi: 10.1175/BAMS-D-19-0091.1.
- De Vos, L., H. Leijnse, A. Overeem, and R. Uijlenhoet, 2017: The potential of urban rainfall monitoring with crowdsourced automatic weather stations in Amsterdam. *Hydrology and Earth System Sciences*, **21** (2), 765–777, doi: 10.5194/hess-21-765-2017.
- de Vos, L. W., H. Leijnse, A. Overeem, and R. Uijlenhoet, 2019b: Quality Control for Crowdsourced Personal Weather Stations to Enable Operational Rainfall Monitoring. *Geophysical Research Letters*, **46** (0), doi: 10.1029/2019GL083731.
- de Vos, L. W., A. Overeem, H. Leijnse, and R. Uijlenhoet, 2019c: Rainfall estimation accuracy of a nation-wide instantaneously sampling commercial microwave link network: error-dependency on known characteristics. *Journal of Atmospheric and Oceanic Technology*, **36**, 1267–1283, doi: 10.1175/JTECH-D-18-0197.1.
- de Vos, L. W., T. H. Raupach, H. Leijnse, A. Overeem, A. Berne, and R. Uijlenhoet, 2018: High-Resolution Simulation Study Exploring the Potential of Radars, Crowdsourced Personal Weather Stations, and Commercial Microwave Links to Monitor Small-Scale Urban Rainfall. *Water Resources Research*, **54** (12), 10,293–10,312, doi: 10.1029/2018WR023393.
- Dhakal, S., 2009: Urban energy use and carbon emissions from cities in China and policy implications. *Energy Policy*, **37** (11), 4208–4219, doi: 10.1016/j.enpol.2009.05.020.
- Diffenbaugh, N. S., and M. Burke, 2019: Global warming has increased global economic inequality. *Proceedings of the National Academy of Sciences*, **116** (20), 9808–9813, doi: 10.1073/pnas.1816020116.
- Drew, D. R., J. F. Barlow, and S. E. Lane, 2013: Observations of wind speed profiles over Greater London, UK, using a Doppler lidar. *Journal of Wind Engineering and Industrial Aerodynamics*, **121**, 98–105, doi: 10.1016/j.jweia.2013.07.019.
- Driedonks, A. G. M., and H. Tennekes, 1984: Entrainment effects in the well-mixed atmospheric boundary layer. *Boundary-Layer Meteorology*, **30** (1-4), 75–105, doi: 10.1007/BF00121950.
- Droste, A. M., D. Fenner, B. Heusinkveld, and G. J. Steeneveld, 2020: Assessing the potential and application of crowdsourced urban wind data. *Quarterly Journal of the Royal Meteorological Society*, 1–18, doi: 10.1002/qj.3811.
- Droste, A. M., B. G. Heusinkveld, D. Fenner, and G. J. Steeneveld, 2018a: Crowdsourcing the urban wind. *{EGU} General Assembly Conference Abstracts*, Vol. 21, 5671.
- Droste, A. M., J. J. Pape, A. Overeem, H. Leijnse, G. J. Steeneveld, A. J. Van Delden, and R. Uijlenhoet, 2017: Crowdsourcing Urban Air Temperatures through Smartphone Battery Temperatures in São Paulo, Brazil. *Journal of Atmospheric and Oceanic Technology*, **34** (9), 1853–1866, doi: 10.1175/JTECH-D-16-0150.1.
- Droste, A. M., G. J. Steeneveld, and A. A. M. Holtslag, 2018b: Introducing the urban wind island effect. *Environmental Research Letters*, **13** (9), 094007, doi: 10.1088/1748-9326/aad8ef.

- Dueben, P. D., and P. Bauer, 2018: Challenges and design choices for global weather and climate models based on machine learning. *Geoscientific Model Development*, **11** (10), 3999–4009, doi: 10.5194/gmd-11-3999-2018.
- ECMWF, 2016: ECMWF strategy 2016-2025: the strength of a common goal. Tech. rep., 27 pp. URL https://www.ecmwf.int/sites/default/files/ECMWF_Strategy_2016-2025.pdf.
- ECMWF, 2020: ECMWF scientists to simulate global weather at 1 km resolution. ECMWF, URL <https://www.ecmwf.int/en/about/media-centre/news/2020/ecmwf-scientists-simulate-global-weather-1-km-resolution>.
- Eden, P., 2009: Traditional weather observing in the UK: An historical overview. *Weather*, **64** (9), 239–245, doi: 10.1002/wea.469.
- Eliasson, I., B. Offerle, C. S. B. Grimmond, and S. Lindqvist, 2006: Wind fields and turbulence statistics in an urban street canyon. *Atmospheric Environment*, **40** (1), 1–16, doi: 10.1016/j.atmosenv.2005.03.031.
- Elmore, K. L., Z. L. Flamig, V. Lakshmanan, B. T. Kaney, V. Farmer, H. D. Reeves, and L. P. Rothfus, 2014: MPING: Crowd-sourcing weather reports for research. *Bulletin of the American Meteorological Society*, **95** (9), 1335–1342, doi: 10.1175/BAMS-D-13-00014.1.
- EUMETSAT, 2009: Introduction to Conceptual Models. URL https://www.eumetsat.int/website/home/Data/Training/TrainingLibrary/DAT_2043086.html.
- Fenner, D., F. Meier, B. Bechtel, M. Otto, and D. Scherer, 2017: Intra and inter ‘local climate zone’ variability of air temperature as observed by crowdsourced citizen weather stations in Berlin, Germany. *Meteorologische Zeitschrift*, **26** (5), 525–547, doi: 10.1127/metz/2017/0861.
- Fenner, D., F. Meier, D. Scherer, and A. Polze, 2014: Spatial and temporal air temperature variability in Berlin, Germany, during the years 2001-2010. *Urban Climate*, doi: 10.1016/j.uclim.2014.02.004.
- Fortuniak, K., K. Klysik, and J. Wibig, 2006: Urban–rural contrasts of meteorological parameters in Łódź. *Theoretical and Applied Climatology*, **84** (1-3), 91–101, doi: 10.1007/s00704-005-0147-y.
- Garcia, A., J. L. Torres, E. Prieto, and A. De Francisco, 1998: Fitting wind speed distributions: A case study. *Solar Energy*, **62** (2), 139–144, doi: 10.1016/S0038-092X(97)00116-3.
- Gentine, P., M. Pritchard, S. Rasp, G. Reinaudi, and G. Yacalis, 2018: Could Machine Learning Break the Convection Parameterization Deadlock? *Geophysical Research Letters*, **45** (11), 5742–5751, doi: 10.1029/2018GL078202.
- Ghariesifard, M., U. Wehn, and P. van der Zaag, 2017: Towards benchmarking citizen observatories: Features and functioning of online amateur weather networks. *Journal of Environmental Management*, **193**, 381–393, doi: 10.1016/j.jenvman.2017.02.003.
- Golroudbary, V. R., Y. Zeng, C. M. Mannaerts, and Z. B. Su, 2018: Urban impacts on air temperature and precipitation over The {N}etherlands. *Climate Research*, **75** (2), 95–109.

- Grimmond, C. S. B., and T. R. Oke, 2002: Turbulent Heat Fluxes in Urban Areas: Observations and a Local-Scale Urban Meteorological Parameterization Scheme (LUMPS). *Journal of Applied Meteorology*, **41** (7), 792–810, doi: 10.1175/1520-0450(2002)041<0792:THFIUA>2.0.CO;2.
- GSMA, 2019: The Mobile Economy - 2019. Tech. rep., The GSM Association, 56 pp. URL <https://www.gsmainelligence.com/research/?file=b9a6e6202ee1d5f787cfebb95d3639c5>.
- Haiden, T., M. Janousek, F. Vitart, L. Ferranti, and F. Prates, 2019: Evaluation of ECMWF forecasts, including the 2019 upgrade. *ECMWF Technical Memoranda*, **853**, doi: 10.21957/mlvapkke, URL <https://www.ecmwf.int/node/19277>.
- Haiden, T., and Coauthors, 2018: Use of in situ surface observations at ECMWF. *ECMWF Technical Memoranda*, **834**, doi: 10.21957/dj9lpy4wa.
- Hajat, S., and T. Kosatky, 2010: Heat-related mortality: a review and exploration of heterogeneity. *Journal of Epidemiology & Community Health*, **64** (9), 753–760, doi: 10.1136/jech.2009.087999.
- Heusinkveld, B. G., G. J. Steeneveld, L. W. A. Van Hove, C. M. J. Jacobs, and A. A. M. Holtslag, 2014: Spatial variability of the Rotterdam urban heat island as influenced by urban land use. *Journal of Geophysical Research : Atmospheres*, **119** (2), 677–692, doi: 10.1002/2012JD019399.
- Heusinkveld, B. G., and Coauthors, 2017: Smartphone App Brings Human Thermal Comfort Forecast in Your Hands. *Bulletin of the American Meteorological Society*, **98** (12), 2533–2541, doi: 10.1175/BAMS-D-16-0082.1.
- Hintz, K., H. Vedel, J. Muñoz-Gomez, and N. Woetmann, 2017: An Examination of the Quality of Wind Observations with Smartphones. *EGU General Assembly Conference Abstracts*, EGU General Assembly Conference Abstracts, Vol. 19, 19568, URL <http://adsabs.harvard.edu/abs/2017EGUGA..1919568H>.
- Hintz, K. S., H. Vedel, and E. Kaas, 2019: Collecting and processing of barometric data from smartphones for potential use in numerical weather prediction data assimilation. *Meteorological Applications*, **26** (4), 733–746, doi: 10.1002/met.1805.
- Holtslag, A. A. M., and A. P. Van Ulden, 1983: A Simple Scheme for Daytime Estimates of the Surface Fluxes from Routine Weather Data. *Journal of Climate and Applied Meteorology*, **22** (4), 517–529, doi: 10.1175/1520-0450(1983)022<0517:ASSFDE>2.0.CO;2.
- Hsieh, C. M., and H. C. Huang, 2016: Mitigating urban heat islands: A method to identify potential wind corridor for cooling and ventilation. *Computers, Environment and Urban Systems*, **57**, doi: 10.1016/j.compenurbsys.2016.02.005.
- Ichinose, T., K. Shimodozono, and K. Hanaki, 1999: Impact of anthropogenic heat on urban climate in Tokyo. *Atmospheric Environment*, **33** (24–25), 3897–3909, doi: 10.1016/S1352-2310(99)00132-6.
- Isaak, J., and M. J. Hanna, 2018: User Data Privacy: Facebook, Cambridge Analytica, and Privacy Protection. *Computer*, **51** (8), 56–59, doi: 10.1109/MC.2018.3191268.

- Jameson, A. R., 1991: A comparison of microwave techniques for measuring rainfall. *Journal of Applied Meteorology*, **30** (1), 32–54.
- Jenkins, G., 2014: A comparison between two types of widely used weather stations. *Weather*, **69** (4), 105–110.
- Justus, C. G., W. R. Hargraves, A. Mikhail, and D. Graber, 1978: Methods for Estimating Wind Speed Frequency Distributions. *Journal of Applied Meteorology*, **17** (3), 350–353, doi: 10.1175/1520-0450(1978)017<0350:MFEWSF>2.0.CO;2.
- Kamel Boulos, M. N., and Coauthors, 2011: Crowdsourcing, citizen sensing and sensor web technologies for public and environmental health surveillance and crisis management: trends, OGC standards and application examples. *International Journal of Health Geographics*, **10** (1), 67, doi: 10.1186/1476-072X-10-67.
- Kent, C. W., C. Grimmond, D. Gatey, and J. F. Barlow, 2018: Assessing methods to extrapolate the vertical wind-speed profile from surface observations in a city centre during strong winds. *Journal of Wind Engineering and Industrial Aerodynamics*, **173**, 100–111, doi: 10.1016/j.jweia.2017.09.007.
- Kent, C. W., S. Grimmond, J. Barlow, D. Gatey, S. Kotthaus, F. Lindberg, and C. H. Halios, 2017: Evaluation of Urban Local-Scale Aerodynamic Parameters: Implications for the Vertical Profile of Wind Speed and for Source Areas. *Boundary-Layer Meteorology*, doi: 10.1007/s10546-017-0248-z.
- KNMI, 2000: *Handbook for the Meteorological Observation*. Koninklijk Nederlands Meteorologisch Instituut, 91–110 pp., URL http://projects.knmi.nl/hawa/pdf/Handbook_H01_H06.pdf.
- Kondo, H., A. Inagaki, and M. Kanda, 2015: A New Parametrization of Mixing Length in an Urban Canopy Derived from a Large-Eddy Simulation Database for Tokyo. *Boundary-Layer Meteorology*, **156** (1), 131–144, doi: 10.1007/s10546-015-0019-7.
- Koracin, D., and R. Berkowicz, 1988: Nocturnal boundary-layer height: Observations by acoustic sounders and predictions in terms of surface-layer parameters. *Boundary-Layer Meteorology*, **43** (1-2), 65–83, doi: 10.1007/BF00153969.
- Kotthaus, S., and C. S. B. Grimmond, 2014: Urban Climate Energy exchange in a dense urban environment – Part I : Temporal variability of long-term observations in central London. *Urban Climate*, **10**, 261–280, doi: 10.1016/j.uclim.2013.10.002.
- Kotthaus, S., H. C. Ward, C. Grimmond, and J. G. Evans, 2012: Urban Flux Measurements in Southern England, by KCL and CEH. *FluxLetter*, **5** (1), 25–31.
- Kovats, R. S., and S. Hajat, 2008: Heat Stress and Public Health: A Critical Review. *Annual Review of Public Health*, **29** (1), 41–55, doi: 10.1146/annurev.publhealth.29.020907.090843.
- Krasnopolsky, V. M., M. S. Fox-Rabinovitz, and A. A. Belochitski, 2008: Decadal climate simulations using accurate and fast neural network emulation of full, longwave and shortwave, radiation. *Monthly Weather Review*, **136** (10), 3683–3695, doi: 10.1175/2008MWR2385.1.

- Krennert, T., G. Pistotnik, R. Kaltenberger, and C. Csekits, 2018: Crowdsourcing of weather observations at national meteorological and hydrological services in Europe. *Advances in Science and Research*, **15**, 71–76, doi: 10.5194/asr-15-71-2018.
- Kulmala, M., and Coauthors, 2011: General overview: European Integrated project on Aerosol Cloud Climate and Air Quality interactions (EUCAARI): integrating aerosol research from nano to global scales. *Atmospheric Chemistry and Physics*, **11** (24), 13 061–13 143, doi: 10.5194/acp-11-13061-2011.
- Kusaka, H., H. Kondo, Y. Kikegawa, and F. Kimura, 2001: A Simple Single-Layer Urban Canopy Model For Atmospheric Models: Comparison With Multi-Layer And Slab Models. *Boundary-Layer Meteorology*, **101** (3), 329–358, doi: 10.1023/A:1019207923078.
- Lawrence-Mathers, A., 2019: *Medieval Meteorology: Forecasting the Weather from Aristotle to the Almanac*. Cambridge University Press, Cambridge, 228 pp., doi: 10.1017/9781108289948.
- Lee, D., 1979: The influence of atmospheric stability and the urban heat island on urban-rural wind speed differences. *Atmospheric Environment (1967)*, **13** (8), 1175–1180, doi: 10.1016/0004-6981(79)90042-8.
- Leijnse, H., R. Uijlenhoet, and J. N. M. Stricker, 2007: Rainfall measurement using radio links from cellular communication networks. *Water Resources Research*, **43** (3), doi: 10.1029/2006WR005631.
- Leijnse, H., R. Uijlenhoet, and J. N. M. Stricker, 2008: Microwave link rainfall estimation: Effects of link length and frequency, temporal sampling, power resolution, and wet antenna attenuation. *Advances in Water Resources*, **31** (11), 1481–1493, doi: 10.1016/j.advwatres.2008.03.004.
- Lien, J., and N. Ahme, 2011: Wind Driven Ventilation for Enhanced Indoor Air Quality. *Chemistry, Emission Control, Radioactive Pollution and Indoor Air Quality*, InTech, doi: 10.5772/17059.
- Lin, Y. W., J. Bates, and P. Goodale, 2016: Co-observing the weather, co-predicting the climate: Human factors in building infrastructures for crowdsourced data. *Science and Technology Studies*, **29** (3), 10–27.
- Macdonald, P., and J. Du, 2018: mixdist: Finite Mixture Distribution Models. URL <https://cran.r-project.org/package=mixdist>.
- Macdonald, R. W., 2000: Modelling the mean velocity profile in the urban canopy layer. *Boundary-Layer Meteorology*, **97** (1), 25–45, doi: 10.1023/A:1002785830512.
- Madaus, L. E., and C. F. Mass, 2017: Evaluating Smartphone Pressure Observations for Mesoscale Analyses and Forecasts. *Weather and Forecasting*, **32** (2), 511–531, doi: 10.1175/WAF-D-16-0135.1.
- Mahrt, L., 1996: The bulk aerodynamic formulation over heterogeneous surfaces. *Boundary-Layer Meteorology*, **78** (1), 87–119, doi: 10.1007/BF00122488.

- Mahrt, L., and D. H. Lenschow, 1976: Growth Dynamics of the Convectively Mixed Layer. *Journal of the Atmospheric Sciences*, **33** (1), 41–51, doi: 10.1175/1520-0469(1976)033<0041:GDOTCM>2.0.CO;2.
- Mandement, M., and O. Caumont, 2019: Contribution of personal weather stations to the observation of deep-convection features near the ground. *Natural Hazards and Earth System Sciences Discussions*, doi: 10.5194/nhess-2019-229.
- Maronga, B., and Coauthors, 2019: Development of a new urban climate model based on the model PALM – Project overview, planned work, and first achievements. *Meteorologische Zeitschrift*, **28** (2), 105–119, doi: 10.1127/metz/2019/0909.
- Marshall, J. S., W. Hirschfeld, and K. L. S. Gunn, 1955: Advances in radar weather. *Advances in Geophysics*, Vol. 2, Academic Press, New York, 1–56.
- Martilli, A., A. Clappier, and M. W. Rotach, 2002: An Urban Surface Exchange Parameterisation for Mesoscale Models. *Boundary-Layer Meteorology*, **104** (2), 261–304, doi: 10.1023/A:1016099921195.
- Mass, C. F., and L. E. Madaus, 2014: Surface pressure observations from smartphones: A potential revolution for high-resolution weather prediction? *Bulletin of the American Meteorological Society*, **95** (9), 1343–1349, doi: 10.1175/BAMS-D-13-00188.1.
- Max Roser, H. R., and E. Ortiz-Ospina, 2020: World Population Growth. *Our World in Data*, URL <https://ourworldindata.org/world-population-growth>.
- McCabe, M. F., and Coauthors, 2017: The future of Earth observation in hydrology. *Hydrology and Earth System Sciences*, **21** (7), 3879–3914.
- McNicholas, C., and C. F. Mass, 2018: Smartphone Pressure Collection and Bias Correction Using Machine Learning. *Journal of Atmospheric and Oceanic Technology*, **35** (3), 523–540.
- Meier, F., D. Fenner, T. Grassmann, B. Jänicke, M. Otto, and D. Scherer, 2015: Challenges and benefits from crowd-sourced atmospheric data for urban climate research using Berlin, Germany, as testbed. *ICUC9 - 9th International Conference on Urban Climate jointly with 12th Symposium on the Urban Environment*.
- Meier, F., D. Fenner, T. Grassmann, M. Otto, and D. Scherer, 2017: Crowdsourcing air temperature from citizen weather stations for urban climate research. *Urban Climate*, **19**, 170–191, doi: 10.1016/j.uclim.2017.01.006.
- Messer, H., and O. Sendik, 2015: A new approach to precipitation monitoring: A critical survey of existing technologies and challenges. *IEEE Signal Processing Magazine*, **32** (3), 110–122, doi: 10.1109/MSP.2014.2309705.
- Messer, H., A. Zinevich, and P. Alpert, 2006: Environmental monitoring by wireless communication networks. *Science*, **312** (5774), 713, doi: 10.1126/science.1120034.
- MetOffice, 2011: Weather observations website. URL <https://www.metoffice.gov.uk/>.
- Mills, G., B. Bechtel, J. Ching, L. See, J. Feddema, M. Foley, P. Alexander, and M. O’Connor, 2015: An Introduction to the WUDAPT project. *Proceedings of the ICUC9. Meteo France*, 6, URL http://www.wudapt.org/wp-content/uploads/2015/05/Mills_et_al_ICUC9.pdf.

- Miralles, D. G., A. J. Teuling, C. C. van Heerwaarden, and J. Vilà-Guerau de Arellano, 2014: Mega-heatwave temperatures due to combined soil desiccation and atmospheric heat accumulation. *Nature Geoscience*, **7** (5), 345–349, doi: 10.1038/ngeo2141.
- Molenaar, R. E., B. G. Heusinkveld, and G. J. Steeneveld, 2016: Projection of rural and urban human thermal comfort in The Netherlands for 2050. *International Journal of Climatology*, **36** (4), 1708–1723, doi: 10.1002/joc.4453.
- Moonen, P., T. Defraeye, V. Dorer, B. Blocken, and J. Carmeliet, 2012: Urban Physics: Effect of the micro-climate on comfort, health and energy demand. *Frontiers of Architectural Research*, **1** (3), 197–228, doi: 10.1016/j.foar.2012.05.002.
- Muller, C. L., L. Chapman, S. Johnston, C. Kidd, S. Illingworth, G. Foody, A. Overeem, and R. R. Leigh, 2015: Crowdsourcing for climate and atmospheric sciences: Current status and future potential. *International Journal of Climatology*, **35** (11), 3185–3203, doi: 10.1002/joc.4210.
- Mureau, R., J. Wesseling, and H. Zipp, 2018: Meet je Stad. *Meteorologica*, **3**, 18–20, URL https://meetjestad.net/downloads/Belang_Minimumtemperatuur_Meteorologica.pdf, (*In Dutch*).
- Napoly, A., T. Grassmann, F. Meier, and D. Fenner, 2018: Development and Application of a Statistically-Based Quality Control for Crowdsourced Air Temperature Data. *Frontiers in Earth Science*, **6**, 1–16, doi: 10.3389/feart.2018.00118.
- Netatmo, 2019: Full Specifications for the Netatmo Weather Station. URL <https://www.netatmo.com/en-us/weather/weatherstation/specifications>.
- Niforatos, E., A. Vourvopoulos, and A. Langheinrich, 2017: Understanding the potential of human-machine crowdsourcing for weather data. *International Journal of Human-Computer Studies*, **102**, 54–68, doi: <https://doi.org/10.1016/j.ijhcs.2016.10.002>.
- Niforatos, E., A. Vourvopoulos, M. Langheinrich, P. Campos, and A. Doria, 2014: Atmos: a hybrid crowdsourcing approach to weather estimation. *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*, ACM, 135–138, doi: 10.1145/2638728.2638780.
- Nipen, T. N., I. A. Seierstad, C. Lussana, J. Kristiansen, and O. Hov, 2019: Adopting citizen observations in operational weather prediction. *Bulletin of the American Meteorological Society*, **0** (0), BAMS-D-18-0237.1, doi: 10.1175/BAMS-D-18-0237.1.
- Noyola Cabrera, A., 2018: Smartphone as a weather station for the urban environment: An exploratory analysis. Wageningen University, MSc thesis, 33 pp.
- NWS, 2013: History of the National Weather Service. URL <https://www.weather.gov/timeline>.
- Oke, T. R., 1982: The energetic basis of the urban heat island. *Quarterly Journal of the Royal Meteorological Society*, **108** (455), 1–24, doi: 10.1002/qj.49710845502.
- Oke, T. R., 1988: Street design and urban canopy layer climate. *Energy and Buildings*, **11** (1-3), 103–113, doi: 10.1016/0378-7788(88)90026-6.

- Oke, T. R., 2006: Initial guidance to obtain representative meteorological observations at urban sites. *World Meteorological Organization, Geneva, ION report*.
- Oke, T. R., G. Mills, A. Christen, and J. A. Voogt, 2017: *Urban Climates*. Cambridge University Press, Cambridge, 14–43 pp., doi: 10.1017/9781139016476.
- Oleson, K. W., A. Monaghan, O. Wilhelmi, M. Barlage, N. Brunzell, J. Feddema, L. Hu, and D. F. Steinhoff, 2015: Interactions between urbanization, heat stress, and climate change. *Climatic Change*, **129** (3–4), 525–541, doi: 10.1007/s10584-013-0936-8.
- Olsen, R. L., D. V. Rogers, and D. Hodge, 1978: The R^2 relation in the calculation of rain attenuation. *IEEE Transactions on Antennas and Propagation*, **26** (2), 318–329, doi: 10.1109/TAP.1978.1141845.
- O'Neill, J. J., X. M. Cai, and R. Kinnersley, 2016: Improvement of a stochastic backscatter model and application to large-eddy simulation of street canyon flow. *Quarterly Journal of the Royal Meteorological Society*, 1121–1132, doi: 10.1002/qj.2715.
- Ottinger, G., 2010: Buckets of Resistance: Standards and the Effectiveness of Citizen Science. *Science, Technology, & Human Values*, **35** (2), 244–270, doi: 10.1177/0162243909337121.
- Overeem, A., T. A. Buishand, and I. Holleman, 2009a: Extreme rainfall analysis and estimation of depth-duration-frequency curves using weather radar. *Water Resources Research*, **45**, doi: 10.1029/2009WR007869.
- Overeem, A., I. Holleman, and A. Buishand, 2009b: Derivation of a 10-year radar-based climatology of rainfall. *Journal of Applied Meteorology and Climatology*, **48**, 1448–1463, doi: 10.1175/2009JAMC1954.1.
- Overeem, A., H. Leijnse, and R. Uijlenhoet, 2011: Measuring urban rainfall using microwave links from commercial cellular communication networks. *Water Resources Research*, **47**, doi: 10.1029/2010WR010350.
- Overeem, A., H. Leijnse, and R. Uijlenhoet, 2013a: Country-wide rainfall maps from cellular communication networks. *Proceedings of the National Academy of Sciences*, **110** (8), 2741–2745, doi: 10.1073/pnas.1217961110.
- Overeem, A., H. Leijnse, and R. Uijlenhoet, 2016a: Retrieval algorithm for rainfall mapping from microwave links in a cellular communication network. *Atmospheric Measurement Techniques*, **9** (5), 2425–2444, doi: 10.5194/amt-9-2425-2016.
- Overeem, A., H. Leijnse, and R. Uijlenhoet, 2016b: Two and a half years of country-wide rainfall maps using radio links from commercial cellular telecommunication networks. *Water Resources Research*, **52** (10), 8039–8065, doi: 10.1002/2016WR019412.
- Overeem, A., J. C. R. Robinson, H. Leijnse, G. J. Steeneveld, B. K. P. Horn, and R. Uijlenhoet, 2013b: Crowdsourcing urban air temperatures from smartphone battery temperatures. *Geophysical Research Letters*, **40** (15), 4081–4085, doi: 10.1002/grl.50786.
- Overeem, A., J. C. R. Robinson, H. Leijnse, G. J. Steeneveld, B. K. P. Horn, and R. Uijlenhoet, 2014: Stadstemperatuur in beeld met Smartphones. *Meteorologica*, **23**, 16–18, (*In Dutch*).

- Pascal, M., and Coauthors, 2013: Assessing the public health impacts of urban air pollution in 25 European cities: Results of the Aphekom project. *Science of The Total Environment*, **449**, 390–400, doi: 10.1016/j.scitotenv.2013.01.077.
- Paulos, E., R. Honicky, and B. Hooker, 2009: Citizen Science. *Handbook of Research on Urban Informatics*, M. Foth, Ed., IGI Global, chap. 28, 414–436, doi: 10.4018/978-1-60566-152-0.ch028.
- Penwarden, A., 1973: Acceptable wind speeds in towns. *Building Science*, **8** (3), 259–267, doi: 10.1016/0007-3628(73)90008-X.
- Peters, G. P., R. M. Andrew, J. G. Canadell, P. Friedlingstein, R. B. Jackson, J. I. Korsbakken, C. Le Quéré, and A. Peregon, 2020: Carbon dioxide emissions continue to grow amidst slowly emerging climate policies. *Nature Climate Change*, **10** (1), 3–6, doi: 10.1038/s41558-019-0659-6.
- Powers, J. G., and Coauthors, 2017: The Weather Research and Forecasting Model: Overview, System Efforts, and Future Directions. *Bulletin of the American Meteorological Society*, **98** (8), 1717–1737, doi: 10.1175/BAMS-D-15-00308.1.
- Rabiei, E., U. Haberlandt, M. Sester, and D. Fitzner, 2013: Rainfall estimation using moving cars as rain gauges-laboratory experiments. *Hydrology and Earth System Sciences*, **17** (11), 4701–4712.
- Racah, E., C. Beckham, T. Maharaj, S. Ebrahimi Kahou, M. Prabhat, and C. Pal, 2017: ExtremeWeather: A large-scale climate dataset for semi-supervised detection, localization, and understanding of extreme weather events. *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., Curran Associates, Inc., 3402–3413.
- Ramponi, R., B. Blocken, L. B. de Coo, and W. D. Janssen, 2015: CFD simulation of outdoor ventilation of generic urban configurations with different urban densities and equal and unequal street widths. *Building and Environment*, **92**, 152–166, doi: 10.1016/j.buildenv.2015.04.018.
- Reichstein, M., G. Camps-Valls, B. Stevens, M. Jung, J. Denzler, N. Carvalhais, and Prabhat, 2019: Deep learning and process understanding for data-driven Earth system science. *Nature*, **566** (7743), 195–204, doi: 10.1038/s41586-019-0912-1.
- Reid, C. E., M. S. O'Neill, C. J. Gronlund, S. J. Brines, D. G. Brown, A. V. Diez-Roux, and J. Schwartz, 2009: Mapping Community Determinants of Heat Vulnerability. *Environmental Health Perspectives*, **117** (11), 1730–1736, doi: 10.1289/ehp.0900683.
- Reynolds, D., 2009: Gaussian Mixture Models. Springer US, Boston, MA, 659–663 pp., doi: 10.1007/978-0-387-73003-5_196.
- Ronda, R. J., G. J. Steeneveld, B. G. Heusinkveld, J. J. Attema, and A. A. Holtslag, 2017: Urban finescale forecasting reveals weather conditions with unprecedented detail. *Bulletin of the American Meteorological Society*, **98** (12), 2675–2688, doi: 10.1175/BAMS-D-16-0297.1.
- Rotach, M. W., 1995: Profiles of turbulence statistics in and above an urban street canyon. *Atmospheric Environment*, **29** (13), 1473–1486, doi: 10.1016/1352-2310(95)00084-C.

- Rotach, M. W., and Coauthors, 2005: BUBBLE - An urban boundary layer meteorology project. *Theoretical and Applied Climatology*, **81** (3-4), 231–261, doi: 10.1007/s00704-004-0117-9.
- Salamanca, F., A. Krpo, A. Martilli, and A. Clappier, 2010: A new building energy model coupled with an urban canopy parameterization for urban climate simulations—part I. formulation, verification, and sensitivity analysis of the model. *Theoretical and Applied Climatology*, **99** (3-4), 331–344, doi: 10.1007/s00704-009-0142-9.
- Savić, S., D. Arsenović, V. Marković, and D. Milošević, 2019: Temperature Risk Assessment in Urban Environments During Heat Wave Periods: A Case Study on the City of Novi Sad (Serbia). *Climate Change Adaptation in Eastern Europe: Managing Risks and Building Resilience to Climate Change*, W. Leal Filho, G. Trbic, and D. Filipovic, Eds., Springer International Publishing, Cham, 185–197, doi: 10.1007/978-3-030-03383-5_13.
- Schröter, J. S., A. F. Moene, and A. A. M. Holtslag, 2013: Convective boundary layer wind dynamics and inertial oscillations: The influence of surface stress. *Quarterly Journal of the Royal Meteorological Society*, **139** (676), 1694–1711, doi: 10.1002/qj.2069.
- Shapiro, A., and E. Fedorovich, 2010: Analytical description of a nocturnal low-level jet. *Quarterly Journal of the Royal Meteorological Society*, **136** (650), 1255–1262, doi: 10.1002/qj.628.
- Shi, Y., K. K.-L. Lau, C. Ren, and E. Ng, 2018: Evaluating the local climate zone classification in high-density heterogeneous urban environment using mobile measurement. *Urban Climate*, **25**, 167–186, doi: 10.1016/j.uclim.2018.07.001.
- Shreffler, J. H., 1979: Urban-Rural Differences in Tower-Measured Winds, St. Louis. *Journal of Applied Meteorology*, **18** (7), 829–835, doi: 10.1175/1520-0450(1979)018<0829:URDITM>2.0.CO;2.
- Skamarock, W. C., and J. B. Klemp, 2008: A time-split nonhydrostatic atmospheric model for weather research and forecasting applications. *Journal of Computational Physics*, **227** (7), 3465–3485, doi: 10.1016/j.jcp.2007.01.037.
- Sluiter, R., 2011: Klimaatatlas. KNMI, URL <http://www.klimaatatlas.nl/>.
- Snik, F., and Coauthors, 2014: Mapping atmospheric aerosols with a citizen science network of smartphone spectropolarimeters. *Geophysical Research Letters*, **41** (20), 7351–7358, doi: 10.1002/2014GL061462, 10.1175/2014GL061462, 10.1175/2014GL061462.
- Steadman, R. G., 1971: Indices of Windchill of Clothed Persons. *Journal of Applied Meteorology*, **10** (4), 674–683, doi: 10.1175/1520-0450(1971)010<0674:IOWOCP>2.0.CO;2.
- Steeneveld, G. J., S. Koopmans, B. G. Heusinkveld, L. W. A. Van Hove, and A. A. M. Holtslag, 2011: Quantifying urban heat island effects and human comfort for cities of variable size and urban morphology in the Netherlands. *Journal of Geophysical Research*, **116** (D20), D20 129, doi: 10.1029/2011JD015988.
- Steeneveld, G. J., B. J. H. Van de Wiel, and A. A. M. Holtslag, 2007: Diagnostic Equations for the Stable Boundary Layer Height: Evaluation and Dimensional Analysis. *Journal of Applied Meteorology and Climatology*, **46** (2), 212–225, doi: 10.1175/JAM2454.1.

- Steeneveld, G.-J., and Coauthors, 2017: Single-column Urban Boundary Layer Inter-comparison Modelling Experiment (SUBLIME): Call for participation. *Urban Climate News*, **66**, 21–26.
- Stewart, I. D., 2011: A systematic review and scientific critique of methodology in modern urban heat island literature. *International Journal of Climatology*, **31**, 200–217, doi: 10.1002/joc.2141.
- Stewart, I. D., and T. R. Oke, 2012: Local climate zones for urban temperature studies. *Bulletin of the American Meteorological Society*, **93** (12), 1879–1900, doi: 10.1175/BAMS-D-11-00019.1.
- Stull, R. B., 1988: *An Introduction to Boundary Layer Meteorology*. Springer Netherlands, Dordrecht, 670 pp., doi: 10.1007/978-94-009-3027-8, 9809069v1.
- Sun, Y., and A. Mobasheri, 2017: Utilizing Crowdsourced Data for Studies of Cycling and Air Pollution Exposure: A Case Study Using Strava Data. *International Journal of Environmental Research and Public Health*, **14** (3), 274, doi: 10.3390/ijerph14030274.
- Takle, E. S., and J. M. Brown, 1978: Note on the Use of Weibull Statistics to Characterize Wind-Speed Data. *Journal of Applied Meteorology*, **17** (4), 556–559, doi: 10.1175/1520-0450(1978)017<0556:NOTUOW>2.0.CO;2.
- Taub, L., 2003: *Ancient Meteorology*. Routledge, London, doi: 10.4324/9780203634288.
- Tauro, F., and Coauthors, 2018: Measurements and Observations in the {XXI} century ({MOXXI}): innovation and multi-disciplinarity to sense the hydrological cycle. *Hydrological Sciences Journal*, **63** (2), 169–196.
- Tennekes, H., 1973: The Logarithmic Wind Profile. *Journal of the Atmospheric Sciences*, **30** (2), 234–238, doi: 10.1175/1520-0469(1973)030<0234:TLWP>2.0.CO;2.
- Tennekes, H., and A. G. M. Driedonks, 1981: Basic entrainment equations for the atmospheric boundary layer. *Boundary-Layer Meteorology*, **20** (4), 515–531, doi: 10.1007/BF00122299.
- Theeuwes, N. E., A. Solcerová, and G. J. Steeneveld, 2013: Modeling the influence of open water surfaces on the summertime temperature and thermal comfort in the city. *Journal of Geophysical Research Atmospheres*, **118** (16), 8881–8896, doi: 10.1002/jgrd.50704.
- Theeuwes, N. E., G.-J. Steeneveld, R. J. Ronda, and A. A. M. Holtslag, 2017: A diagnostic equation for the daily maximum urban heat island effect for cities in northwestern Europe. *International Journal of Climatology*, **37** (1), 443–454, doi: 10.1002/joc.4717.
- Theeuwes, N. E., G.-J. Steeneveld, R. J. Ronda, M. W. Rotach, and A. A. M. Holtslag, 2015: Cool city mornings by urban heat. *Environmental Research Letters*, **10** (11), 114022, doi: 10.1088/1748-9326/10/11/114022.
- Toparlar, Y., B. Blocken, B. Maiheu, and G. van Heijst, 2017: A review on the CFD analysis of urban microclimate. *Renewable and Sustainable Energy Reviews*, **80** (September 2016), 1613–1640, doi: 10.1016/j.rser.2017.05.248.

- Tsiringakis, A., G. J. Steeneveld, A. A. Holtslag, S. Kotthaus, and S. Grimmond, 2019: On- and off-line evaluation of the single-layer urban canopy model in London summertime conditions. *Quarterly Journal of the Royal Meteorological Society*, **145** (721), 1474–1489, doi: 10.1002/qj.3505.
- Uijlenhoet, R., A. Overeem, and H. Leijnse, 2018: Opportunistic remote sensing of rainfall using microwave links from cellular communication networks. *Wiley Interdisciplinary Reviews: Water*, **5** (4), e1289, doi: 10.1002/wat2.1289.
- United Nations, 2012: World Urbanization Prospects: The 2005 revision. URL https://www.un.org/en/development/desa/population/publications/pdf/urbanization/WUP2011_Report.pdf, 302 pp.
- Upton, G. J., A. R. Holt, R. J. Cummings, A. R. Rahimi, and J. W. Goddard, 2005: Microwave links: The future for urban rainfall measurement? *Atmospheric Research*, **77** (1-4 SPEC. ISS.), 300–312, doi: 10.1016/j.atmosres.2004.10.009.
- USAID, 2013: Crowdsourcing Applications for Agricultural Development in Africa. URL https://www.agrilinks.org/sites/default/files/resource/files/Crowdsourcing_Applications_for_Agricultural_Development_in_Africa.pdf, 1–6 pp.
- Van de Wiel, B. J. H., A. F. Moene, G. J. Steeneveld, P. Baas, F. C. Bosveld, and A. A. M. Holtslag, 2010: A Conceptual View on Inertial Oscillations and Nocturnal Low-Level Jets. *Journal of the Atmospheric Sciences*, **67** (8), 2679–2689, doi: 10.1175/2010JAS3289.1.
- van Haren, R., S. Koopmans, G.-J. Steeneveld, N. Theeuwes, R. Uijlenhoet, and A. A. Holtslag, 2018: Weather Reanalysis on an Urban Scale using WRF. *2018 IEEE 14th International Conference on e-Science (e-Science)*, IEEE, 279–280, doi: 10.1109/eScience.2018.00049.
- Verdonck, M.-L., and Coauthors, 2019: The Human Influence Experiment (Part 2): Guidelines for Improved Mapping of Local Climate Zones Using a Supervised Classification. *Urban Science*, **3** (1), 27, doi: 10.3390/urbansci3010027, URL <https://www.mdpi.com/2413-8851/3/1/27>.
- Vogelezang, D. H. P., and A. A. M. Holtslag, 1996: Evaluation and model impacts of alternative boundary-layer height formulations. *Boundary-Layer Meteorology*, **81** (3-4), 245–269, doi: 10.1007/BF02430331.
- Voyant, C., G. Notton, S. Kalogirou, M. L. Nivet, C. Paoli, F. Motte, and A. Fouilloy, 2017: Machine learning methods for solar radiation forecasting: A review. *Renewable Energy*, **105**, 569–582, doi: 10.1016/j.renene.2016.12.095.
- Waldron, J., 1995: The Wisdom of the Multitude: Some Reflections on Book 3, Chapter 11 of Aristotle’s Politics. *Political Theory*, **23** (4), 563–584.
- Walker, S. L., 2011: Building mounted wind turbines and their suitability for the urban scale-A review of methods of estimating urban wind resource. *Energy and Buildings*, **43** (8), 1852–1862, doi: 10.1016/j.enbuild.2011.03.032.
- Wallace, J. M., and P. V. Hobbs, 2006: *Atmospheric Science*. Elsevier, doi: 10.1016/C2009-0-00034-8.

-
- Warren, E. L., D. T. Young, L. Chapman, C. Muller, C. Grimmond, and X.-M. Cai, 2016: The Birmingham Urban Climate Laboratory—A high density, urban meteorological dataset, from 2012–2014. *Scientific Data*, **3**, 160 038, doi: 10.1038/sdata.2016.38.
- Werner, C., and Coauthors, 2012: Progress and challenges in using stable isotopes to trace plant carbon and water relations across scales. *Biogeosciences*, **9** (8), 3083–3111, doi: 10.5194/bg-9-3083-2012.
- Wise, D. J., V. B. L. Boppana, K. W. Li, and H. J. Poh, 2018: Effects of minor changes in the mean inlet wind direction on urban flow simulations. *Sustainable Cities and Society*, **37** (August 2017), 492–500, doi: 10.1016/j.scs.2017.11.041.
- Witze, A., 2019: Global 5G wireless deal threatens weather forecasts. *Nature*, **575** (7784), 577–577, doi: 10.1038/d41586-019-03609-x.
- WMO, 2018: Guide to Meteorological Instruments and Methods of Observation. URL https://library.wmo.int/index.php?lvl=notice_display&id=12407.
- Wolters, D., and T. Brandsma, 2012: Estimating the urban heat Island in residential areas in the Netherlands using observations by weather amateurs. *Journal of Applied Meteorology and Climatology*, **51** (4), 711–721, doi: 10.1175/JAMC-D-11-0135.1.
- Wouters, H., I. Y. Petrova, C. C. van Heerwaarden, J. Vilà-Guerau de Arellano, A. J. Teuling, V. Meulenbergh, J. A. Santanello, and D. G. Miralles, 2019: Atmospheric boundary layer dynamics from balloon soundings worldwide: CLASS4GL v1.0. *Geoscientific Model Development*, **12** (5), 2139–2153, doi: 10.5194/gmd-12-2139-2019.
- Yang, Y.-J., and Coauthors, 2013: Impacts of Urbanization and Station-relocation on Surface Air Temperature Series in Anhui Province, China. *Pure and Applied Geophysics*, **170** (11), 1969–1983, doi: 10.1007/s00024-012-0619-9.
- Zheng, F., and Coauthors, 2018: Crowdsourcing Methods for Data Collection in Geophysics: State of the Art, Issues, and Future Directions. *Reviews of Geophysics*, **56** (4), 698–740.
- Zinevich, A., P. Alpert, and H. Messer, 2008: Estimation of rainfall fields using commercial microwave communication networks of variable density. *Advances in Water Resources*, **31** (11), 1470–1480, doi: 10.1016/j.advwatres.2008.03.003.

Authorship Statement

- AD = Arjan Droste (WUR)
- AO = Aart Overeem (KNMI/WUR)
- AvD = Aarnout van Delden (IMAU)
- BHe = Bert Heusinkveld (WUR)
- BHo = Bert Holtslag (WUR)
- DF = Daniel Fenner (TUB)
- GS = Gert-Jan Steeneveld (WUR)
- HL = Hidde Leijnse (KNMI)
- JP = Jan-Jaap Pape (IMAU)
- LV = Lotte de Vos (KNMI/WUR)
- MZ = Marjanne Zander (Deltares/WUR)
- RU = Remko Uijlenhoet (WUR)

Chapters 1 (Introduction) and 7 (Synthesis) have been solely written by AD.

The data of Chapter 2 was collected by AO and preprocessed by JP. Data analysis and interpretation of results was performed by AD with input from AO and HL. The article was drafted by AD, revised by AO, HL, GS, AvD and RU, and approved by all co-authors.

The PWS data in Chapter 3 was collected by AD; the Berlin reference data was obtained through DF; the Veenkampen NetAtmo setup was done by AD and BHe; BHe provided the AAMS data. AD performed all data analyses, interpretation of results was aided by DF, BHe and GS. AD drafted the article, which was revised with input from all co-authors, and approved for publication by all.

The idea for the work in Chapter 4 was proposed by RU, in consultation with LV, AD, AO and GS. The dataset collection was performed by LV (PWS data for rain, temperature, relative humidity and pressure), AO (CML data, smartphone data and radar reference), BHe (AAMS data) and AD (PWS data for wind and AWS reference). LV and AO were responsible for the rain results of the study. AD produced the UHI results assisted by AO, as well as the wind and smartphone pressure results. MZ was in charge of the smartphone results on light. GS, RU and HL assisted in the interpretation of the results. LV and AD were responsible for the final manuscript in which RU wrote the abstract and introduction, GS provided text on case selection and study area and information on AAMS, AO wrote sections on the traditional measurement techniques and smartphone data, LV provided text on CML and PWS, AD took the lead in writing the results sections assisted by LV and MZ, and HL wrote the conclusion. Critical revision was performed by LV, AD, AO, MZ, GS, RU and HL, and all (including BHe) approved the article for publication.

The model used in Chapter 5 is based upon Theeuwes et al. (2015), and further developed by AD with aid from GS and BHo. Interpretation of results was done by AD with input from GS

and BHo. The article was drafted by AD, revised and approved for publication by all co-authors.

The model results and methodology used to obtain the data in Chapter 6 were originally developed by Ronda et al. (2017). Postprocessing and analysis of model results was performed by AD with input from GS and BHo. The chapter was drafted by AD and revised with input from GS and BHo, and all co-authors approved it for inclusion in this thesis.

Acknowledgements

Het enige interessante deel van een proefschrift is het dankwoord. Dat was een (mogelijk terecht) verworpen stelling die ik bedacht had ergens in de afgelopen 4 jaar. Laten we eerlijk zijn: niemand leest een volledig proefschrift van kaft tot kaft. De schrijfstijl is saai, de materie te diepgaand en onbekend, en de conclusie is altijd dat er meer onderzoek nodig is en dat we niks zeker weten. Dus duiken we, als we weer zo'n boekje in de handen gedrukt krijgen, eerst eens lekker naar pagina 170 om eens te kijken wat voor menselijks er nou achter al dit wetenschappelijke geneuzel schuilt. En, tikje ijdell: misschien word je zelf wel genoemd als bron van inspiratie...? Aan al die licht ijdele collega's, vrienden en familie: welkom! Ook als je naam er niet bij staat ben ik je vast ergens dankbaar voor (maar bel me alsjeblieft niet om te vragen waarvoor). Laten we maar beginnen.

Het doen van een PhD is een rare ervaring. Het kost zowel meer als minder tijd dan gepland, meer en minder moeite dan vertevoren gedacht, en er komt een hoop meer bij kijken dan alleen maar hoge cijfers voor (Meteorologie) vakken. Aan het begin heb ik hard lopen twijfelen of ik er überhaupt wat mee kon, dit onderwerp. "*The Windy City*" is vaag en breed genoeg, en ik lijk er in geslaagd te zijn dat nog vager en breder te maken: de titel van dit proefschrift is niet voor niets zo lang. In wezen lijkt het nauwelijks op het voorstel dat ik al die tijd geleden heb ingediend bij de onderzoeksschool, maar ik ben dankbaar dat ik de vrijheid en ruimte had om zo mijn eigen draai aan het onderwerp te geven.

Dat ik die vrijheid had is vooral te danken aan mijn begeleiders. Er is me ooit een keer verteld dat een PhD staat of valt met de supervisor, en ik heb helaas zelf enkele PhDers inderdaad vroegtijdig zien stoppen door die reden.

Gert-Jan. Als mijn dagelijkse begeleider is het jouw "eer" om dit dankwoord aan te voeren, en die dank heb ik ook zeker. Ik kende je al een beetje voordat we aan dit project begonnen: eerst tijdens Atmospheric Modelling (waarin je subtiel je eigen werk propageerde als basis voor mijn case-study voor het Amsterdam UHI), later tijdens de bakfietsmetingen in Amsterdam, en natuurlijk de ICUC in Toulouse. Blijkt dat je toch een heel ander beeld van een docent hebt als je een week een hotelkamer deelt en een flink aantal roddels meekrijgt! Maar pas tijdens m'n PhD leerde ik je echt goed kennen en waarderen: het duurde even voor ik goed hoogte van je kon krijgen (en ik je durfde te bombarderen met mijn arsenaal flauwe grappen). Je droge humor, voorliefde voor belastingen, de KLM, en verjaardagen van voormalig Duitse staatshoofden kan ik altijd erg waarderen, evenals je impressies van sommige (ex-)collega's. Ook projectmatig konden we goed met elkaar overweg: je liet me vrolijk aanmodderen, kwam met nuttige opmerkingen en hulp waar nodig, en herinnerde me er vooral aan dat dingen ook afgemaakt en aan de wereld getoond moeten worden. De les die ik van jou heb geleerd is tenslotte: we onderzoeken dit alleen als we er geld uit kunnen halen, of een paper over kunnen publiceren! Je bevologenheid, enorme werklust en aanpak bewonder ik ten zeerste, en volgens mij weet je het wonderwel te combineren

met je vaderschap. Ik hoop dat we in de toekomst nog veel kunnen samenwerken aan interessante urban studies.

Bert. Je was wat minder intiem betrokken bij het dagelijks reilen en zeilen van mijn PhD, maar zeker niet onbelangrijk. Je kwam vaak op cruciale momenten met een nuttig idee, een handig naslagwerk om te raadplegen, of één van je talloze connecties om input aan te vragen, als ik ergens vastzat. Je interesse lag niet zozeer bij het crowdsourcen, maar meer bij het stedelijk windeiland wat we ontdekt hebben (de naam UWI hebben we aan jou te danken!), en daar kwam je eigen expertise met grenslaagmeteo ook zeker veelvuldig van pas. Ook was je vaak de kalmerende factor: waar Gert-Jan en ik nogal graag op alles wat los en vast zit zaten te vitten, bleef jij de rust zelve. Je geveugelde uitspraak, die me altijd zal bijblijven en ik graag herhaal: *er gaan toch zoveel dingen goed*. Moge dat in je pensioen ook door blijven gaan!

Of course, not only my supervisors made my PhD bearable: everyone at MAQ did. I've greatly enjoyed the time I've worked at our chair group, getting to know the colourful personalities and odd quirks of everyone. First, I'd like to thank my officemates, of Lumen A127: Aris, Super Ingrid, Auke, Ruben, Sjoerd, Sytse and Felipe. It was a lot of fun to be in the office with you all, and we enjoyed our share of Aris' culinary rage, regularly held the stress-ball Olympics, greatly improved the WUR-slogan, created start-up companies, and almost-but-not-quite published the greatest collaboration paper the ESG has ever known. You'll have to miss my jokes (and cakes) from now on, though if you insist, I will pop on by every so often from across the HWM-corridor.

Special thanks to Peter K and Imme: we've known each other for a pretty long while now (ever since BWA), and it was great to do a PhD with you. Peter: I'll miss our odd, philosophical discussions and ground-breaking experiments on Skittles and lunar escalators, and Imme: I'll miss your contagious laughter and skills in PhD-songwriting.

I'd also like to thank some (former) staff members that have been inspiring to me, or helped me out during my PhD: Bert Heus; Reinder and Natalie; Oscar and Arnold. Bert Heus: I really admire your incredible enthusiasm for your field, and your seemingly endless knowledge on all kinds of subjects, be they MAQ-related or wildly different areas. And double thanks for helping me out so much with the experimental side of things! Reinder and Nathalie: without you I might not even have considered doing a PhD. As my MSc supervisors, you showed me how much fun it was to do science, and to present at a conference (which is still one of the best I have ever attended). Oscar and Arnold: it's inspiring to see how much energy and effort you both commit to the courses you give (and I played a small part in). Such heart for our students is what makes MAQ education special, with no small thanks to you.

But I've saved perhaps the most important MAQ members for this last bit: Sandra and Caroline. Because, let's be honest, without you being there, the whole chair group would just collapse. Thank you for always being there to help with practical things, and also just to have a nice chat and a cup of tea when everyone else is too busy working like crazy.

Van MAQ verschuif ik nu naar HWM: in het bijzonder Lotte, Aart, Remko, Hidde en Marjanne. We hebben een hoop mooie dingen samen gedaan (2 van de hoofdstukken in dit boekwerk!) en het was altijd goed samenwerken (of in ieder geval goed samen zeiken tegen R, Lotte). Ik kijk uit naar de vervolging van die samenwerking als jullie nieuwe HWM collega!

Dan is er naast het werk ook nog een ander leven (je zou het bijna vergeten), waarvoor ik dankbaar ben dat ze me ervan weerhouden hebben om mezelf plat te werken. De hele Wagenin-gen lunch-club: Monique, Tijn, Eva, Sasja, Dinja, Nadine, Amanda, Djoline, Dolf, en alweer Marjanne: het is altijd een gezellige boel met jullie erbij, en onder het genot van idiote hoeveelheden thee & crumble hebben we een hoop lol gehad (en geklaagd over de WUR...). Duane: ik kan je droogheid altijd erg op prijs stellen, evenals onze gedeelde waarden voor woordgrappen en goede herinneringen aan alle anime en games die we tot ons genomen hebben. Bo: ook na 20+ jaar ben je me nog niet zat, en ik hoop dat dat zo blijft; je mafheid en Bo-wijsheid maakt het leven beter. Ryan: thanks for always being there at the end of the working day, to play games with me, share the crazy things going on in your neck of the woods, and show your genuine interest in my science. It's helped stay sane over these past few busy years.

Als laatste wil ik nog mijn ouders en zus bedanken. Jullie stonden altijd voor me klaar, om te helpen en advies te geven, en het is altijd prettig thuiskomen in Heerhugowaard. Zonder jullie steun had ik het een stuk zwaarder gehad.

List of Publications

Peer-reviewed Journal Publications:

Droste, A. M., J. J. Pape, A. Overeem, H. Leijnse, G. J. Steeneveld, A. J. Van Delden, and R. Uijlenhoet, 2017: Crowdsourcing Urban Air Temperatures through Smartphone Battery Temperatures in São Paulo, Brazil. *Journal of Atmospheric and Oceanic Technology*, **34** (9), 1853–1866, doi: 10.1175/JTECH-D-16-0150.1

Droste, A. M., G. J. Steeneveld, and A. A. M. Holtslag, 2018b: Introducing the urban wind island effect. *Environmental Research Letters*, **13** (9), 094007, doi: 10.1088/1748-9326/aad8ef

Verdonck, M.-l., and Coauthors, 2019: The Human Influence Experiment (Part 2): Guidelines for Improved Mapping of Local Climate Zones Using a Supervised Classification. *Urban Science*, **3** (1), 27, doi: 10.3390/urbansci3010027, URL <https://www.mdpi.com/2413-8851/3/1/27>

de Vos, L., A. Droste, M. Zander, A. Overeem, H. Leijnse, B. Heusinkveld, G. Steeneveld, and R. Uijlenhoet, 2019a: Hydrometeorological monitoring using opportunistic sensing networks in the Amsterdam metropolitan area. *Bulletin of the American Meteorological Society*, BAMS–D–19–0091.1, doi: 10.1175/BAMS-D-19-0091.1

Droste, A. M., D. Fenner, B. Heusinkveld, and G. J. Steeneveld, 2020: Assessing the potential and application of crowdsourced urban wind data. *Quarterly Journal of the Royal Meteorological Society*, 1–18, doi: 10.1002/qj.3811



*Netherlands Research School for the
Socio-Economic and Natural Sciences of the Environment*

D I P L O M A

for specialised PhD training

The Netherlands research school for the
Socio-Economic and Natural Sciences of the Environment
(SENSE) declares that

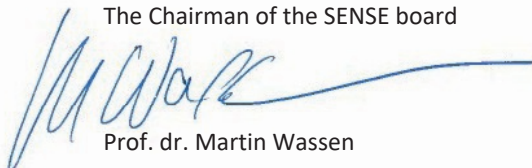
Arjan Marten Droste

born on 22 August 1992 in Heiloo, The Netherlands

has successfully fulfilled all requirements of the
educational PhD programme of SENSE.

Wageningen, 4 November 2020

The Chairman of the SENSE board



Prof. dr. Martin Wassen

the SENSE Director of Education



Dr. Ad van Dommelen

The SENSE Research School has been accredited by the Royal Netherlands Academy of Arts and Sciences (KNAW)



K O N I N K L I J K E N E D E R L A N D S E
A K A D E M I E V A N W E T E N S C H A P P E N



The SENSE Research School declares that **Arjan Marten Droste** has successfully fulfilled all requirements of the educational PhD programme of SENSE with a work load of 43.8 EC, including the following activities:

SENSE PhD Courses

- o Bayesian statistics (2017)
- o Environmental research in context (2017)
- o Research in context activity: 'Initiating and co-organizing a symposium on energy meteorology entitled 'From Model Field to Power Yield' (11 May 2017, Wageningen), and writing accessible summaries for a dedicated website and the Meteorologica popular science journal

Other PhD and Advanced MSc Courses

- o Summer school Urban Physics, University of Cyprus (2016)
- o Communication with the media and general public, Wageningen Graduate Schools (2016)
- o Sub-grid parametrisation in numerical weather models, ECMWF (2017)
- o Teaching and Supervising Thesis students, Wageningen Graduate Schools (2018)

Management and Didactic Skills Training

- o Supervising two MSc students with thesis (2017-2019)
- o Supervising three BSc student with thesis (2017-2019)
- o Teaching in the MSc course 'Urban Hydrometeorology' (2016-2019)
- o Teaching in the MSc course 'Environmental Data Collection & Analysis' (2016-2019)

Selection of Oral Presentations

- o *Introducing the Urban Wind Island Effect*, EGU General Assembly, 8-13 April 2018, Vienna, Austria
- o *Exploring the Potential of Smartphone Temperature Measurements*, 10th International Conference on the Urban Climate, 6-10 August 2018, New York, United States of America
- o **Invited talk:** *Het weer in je achtertuin - crowdsourcen in de stad*, NHV/NVBM symposium, 23 November 2018, Wageningen, The Netherlands
- o *Crowdsourcing the urban wind*, EGU General Assembly, 7-12 April 2019, Vienna, Austria

SENSE coordinator PhD education

Dr. ir. Peter Vermeulen

This research received funding from The Netherlands Organisation for Scientific Research (NWO), grant number 864.14.007.

Financial support from Wageningen University for printing this thesis is gratefully acknowledged.

The reference stations of the Amsterdam Atmospheric Monitoring Supersite have been financially supported by the Amsterdam Institute for Advanced Metropolitan Solutions (AMS).

Cover design by Arjan Droste, inspired by the art of *Shin Megami Tensei: Persona 5*.

