Local Climate, Urban Form and Residential Gas Consumption: Assessing the Added Value of Crowdsourced Temperature Data in The Hague

Luuk Jochems – 0592307
TU Delft – Utrecht University

May 2026
Geographical Information Management and Applications

Supervisor: Arjen Meijer
Responsible Professor: Peter van Oosterom

Cover illustration: Residential gas consumption per PC6 area, The Hague (2021)

Abstract

Locally measured temperature data is increasingly available through citizen sensing networks, yet its added value for explaining spatial patterns in residential gas consumption remains unclear. Most Dutch energy studies treat outdoor temperature as a spatial constant or rely on the regional meteorological data that does not capture intra-urban climate variation. This study analyses residential gas consumption in 9410 postcode 6 areas across The Hague. The conclusion is that locally measured Heating Degree Days from the Netatmo citizen sensing network add statistically significant explanatory power to a spatial regression model of gas consumption, while regional Heating Degree Days derived from KNMI add no explanatory power. The KNMI network provides an almost uniform temperature surface across The Hague, and this is too coarse for intra-urban analysis for a single municipality.

The findings demonstrate that building age and property value remain the dominant predictors of residential gas consumption throughout all models, accounting for the largest share of explained variance. The results show that the baseline model already explains 75.4% of variance and rises to 75.5% when including the locally measured temperature. Urban morphology variables, including vegetation cover and building height, contribute to explaining spatial gas consumption patterns alongside the locally measured temperature. The choice of interpolation methods matters: the Kriging coefficient is three times larger than the IDW equivalent, while using the same stations. This research argues that the added value of including locally measured data is in reducing the measurement error that the use of regional stations brings, when applying at the postcode level. Using the locally measured citizen data produces more reliable estimates of the relationship between the outdoor temperature and residential gas consumption.

Statement on the use of Generative AI

In the preparation of this thesis, Generative AI tools were used to assist with writing structure and R programming. Claude Sonnet was used for coding support with the R program and English writing support for this thesis. All research decisions, data analysis, interpretation and academic conclusions are the author’s own. The research content, findings and conclusions of this thesis are the responsibility of the author.

Acknowledgement

I would like to thank Arjen Meijer for the supervision during this thesis process. His help with shaping the research and responding with constructive feedback. Thank you for also being available and very flexible throughout the thesis process and providing extra help when I needed this. Thanks for the nice discussions and meetings at the TU Delft. I also want to thank Peter van Oosterom, who provided critical feedback and shared practical tips to better my figures.

I also want to thank Ken Arroyo Ohori for collecting and providing the Netatmo data from 2015-2023. This thesis would not have been possible without his querying of the Netatmo API.

Special thanks to Sabine Steyn, who helped me stay motivated throughout the thesis.

Contents

Abstract 2
Statement on the use of Generative AI 2
Acknowledgement 2
Contents 3
List of Abbreviations 5
1 Introduction 6
1.1 Background 6
1.2 Relevance and Research Gap 6
1.3 Research Objectives 7
1.4 Research Questions 8
1.5 Research Scope 8
2 Theoretical Framework 9
2.1 Introduction 9
2.2 Determinants of residential heating energy consumption 9
2.3 Building determinants 10
2.4 Socio-economic determinants 10
2.5 Urban Morphology, Form and Energy demand 11
2.6 Degree Days 12
2.7 KNMI data 13
2.8 Crowdsourcing and Netatmo data 13
2.9 Conceptual model 14
3 Methodology 15
3.1 Research design 15
3.2 Data preprocessing 15
3.3 Control variables and operationalisation 15
3.4 Urban morphology and Urban Heat Island characteristics 17
3.5 Netatmo data quality 18
3.6 Degree Days calculation 19
3.7 Spatial interpolation of climate 20
3.8 Statistical methods 21
3.9 Model specification 21
3.10 Model comparison 22
3.11 Software and tools 22
4 Study area and data 23
4.1 Study area: The Hague 23
4.2 Data 23
4.3 Temperature data 26
4.4 CBS data 28

4.5 BAG 28
5 Results 33
5.1 Spatial Patterns of Local Temperature 33
5.2 Comparison: KNMI and Netatmo 34
5.3 Ordinary Least Squares Regression results 36
5.4 Spatial Autocorrelation and model selection 37
5.5 Model comparison 38
6 Discussion 42
6.1 Main findings 42
6.2 KNMI derived Heating Degree Days 43
6.3 Locally measured Heating Degree Days 43
6.4 Spatial Autocorrelation 44
6.5 Urban morphology and the Urban Heat Island 44
6.6 Policy implications 45
6.7 Limitations 46
7 Conclusion 50
8 Recommendations 52
9 References 53
10 Appendices 58
10.1 Appendix A: Gas consumption table 58
10.2 Appendix B: Kriging cross-validation statistics 58
10.3 Appendix C: R code for data transformation and regression models 58
10.4 Appendix D: Correlation Matrix 60

List of Abbreviations

Abbreviation	Full term
AHN	Algemeen Hoogte Bestand
AIC	Akaike Information Criterion
BAG	Basisregistratie Adressen en Gebouwen
CBS	Centraal Bureau voor de Statistiek
CDD	Cooling Degree Days
GIS	Geographic Information System
HDD	Heating Degree Days
IDW	Inverse Distance Weighting
KNMI	Koninklijk Nederlands Meteorologisch Instituut
NDVI	Normalized Difference Vegetation Index
OLS	Ordinary Least Squares
PC6	Postcode 6
PDOK	Publieke Dienstverlening Op de Kaart
R^2	Coefficient of determination
RMSE	Root Mean Square Error
RVO	Rijksdienst voor Ondernemend Nederland
SEM	Spatial Error Model
SLM	Spatial Lag Model
UHI	Urban Heat Island
VIF	Variance Inflation Factor
WMO	World Meteorological Organization
WOZ	Waardering Onroerende Zaken

1 Introduction

1.1 Background

The Netherlands aims to phase out natural gas by 2050. About 36% of the Dutch natural gas is used for the built environment (Klimaatmonitor, 2023). The largest share of this demand is allocated to space heating of buildings. Fossil fuel consumption for heating and cooling represents 32.7% of total emissions in the Netherlands, with around 19% of these emissions coming from direct gas use for personal residential heating (Du et al., 2025). Due to the temperate Dutch climate, residential heating demands remains substantial during large parts of the year. Reducing the amount of energy consumption for heating is one of the largest challenges in the Netherlands but is essential in order to meet the national and EU targets.

1.2 Relevance and Research Gap

The phasing out of natural gas is one of the most important subjects in the Dutch energy transition policies. The measures that are taken are primarily formulated at a national scale. The centralised approach relies on energy models where local differences are not accounted for and fail to differentiate between various urban contexts. Strong targets that are set are all at the national level. Patterns of heating energy use differ significantly across cities and neighbourhoods (Mashhoodi, 2019). Research into the underlying drivers has traditionally focussed on non-spatial determinants. A large part of the variation in energy consumption can be explained by socio-economic factors and building characteristics (van den Brom et al., 2019). Socio-economic factors include occupant behaviour, age, income and private housing tenure while building characteristics include building age, energy label and floor area. All these factors are well studied and often considered by policymakers (Niessink et al., 2025). The role of urban morphology, including factors such as building density and vegetation, remains underrepresented despite its influence on local temperatures and heating demand (Roxon et al., 2020).

Differences in local climatic conditions within cities, often referred to as microclimates, are not only technical but they translate into a growing social problem, namely energy poverty. Figure 1.1 shows the geographical distribution of the average energy quota in the Netherlands. The energy quota is the percentage of the income that is spent on energy costs. An uneven distribution is seen where the average energy quota of the Randstad region of the Netherlands is lower compared to the rest of the Netherlands. In 2024, 6.1% of all households in the Netherlands were affected by energy poverty. This percentage represents households characterized by a combination of high energy costs, low income, and poor building insulation (Batenburg et al., 2024). Understanding the local drivers behind energy use is therefore important for ensuring that energy policies addresses both climate goals and social problems.

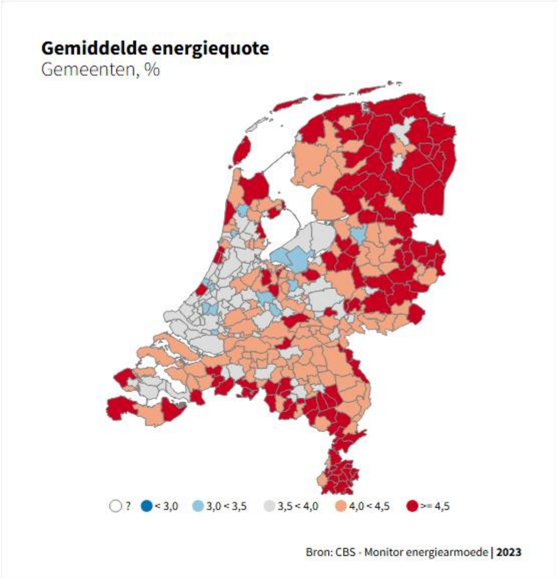

Figure 1.1: Average Energy quota in the Netherlands (CBS, 2023)

To address this gap, this research investigates how local climatic variation influences the residential gas consumption across The Hague, a coastal city with heterogenous urban landscape and the third biggest city in the country. Using temperature data from both the national Royal Netherlands Meteorological Institute network and a crowdsourced sensor network, this study assesses whether locally measured temperature data improves the explanation of spatial patterns in residential gas consumption compared to regional temperature data. Urban morphology is incorporated to capture the Urban Heat Island effects at the postcode level.

The residential sector relies heavily on gas heating, and phasing out natural gas is therefore named as a priority in Dutch heating transition plans (NPLW, 2023). By the end of 2026, municipalities need to provide a 10 year heat transition plan. A substantial share of these plans is focussed on the residential sector. This research could provide municipalities with insights into how local climatic variation influences residential heating demand at the postcode scale.

This research builds upon the previous work of Pijnenburg (2021) who assessed heating energy use at a PC6 level for The Hague and investigated the influence of the local climate on gas consumption. However, their research compares the difference in gas consumption and degree days over the span of 2 years. This study will extend on the research of Pijnenburg (2021) by spatially interpolating the Degree Days using two different methods. Furthermore, this research will test for spatial autocorrelation on the variables in the model. Lastly, this research adds urban morphology variables to incorporate the Urban Heat Island effect at the postcode level along with local temperature data.

1.3 Research Objectives

The objective of this research is to quantify the influence of local temperature variables on residential gas consumption at the postcode level in The Hague. The research combines a dense coverage of meteorological observation stations, based on a Netatmo dataset measured by Ohori (2022) with high-resolution gas consumption data from the Central Bureau of Statistics of the Netherlands. The research will develop spatial regression models that will assess how locally measured temperature, expressed as annual Heating Degree Days, and urban morphology variables contribute to explaining spatial differences in residential gas consumption. This research aims to assess whether incorporating local climatic data improves the model performance compared to models relying on regional interpolated temperature data from official meteorological stations.

1.4 Research Questions

The following research question has been set up:

To what extent does incorporating high-resolution, locally measured climatic data improve the explanation of spatial patterns in residential gas consumption, compared to using spatially coarse regional temperature data?

To help answer this question, the following sub-questions support the main research question:

1. How do Heating Degree Day patterns derived from KNMI and Netatmo datasets differ spatially within The Hague?
2. How do local temperature and urban morphology variables contribute to explaining spatial variation in residential gas consumption?
3. To what extent does including locally measured climatic data improve the performance of a spatial regression model of residential gas consumption, and how does the interpolation method affect this outcome?

1.5 Research Scope

To ensure that the thesis is feasible, the scope of this research will be on residential gas consumption. This excludes commercial and industrial gas consumption. Furthermore, this analysis is limited to the year 2021 as this was the most recent year for which both the Netatmo dataset and the CBS gas consumption are publicly available.

2 Theoretical Framework

2.1 Introduction

Residential heating accounts for a considerable share of the total energy consumption in many developed European countries in temperate and colder climates (Rafiee et al., 2019). Understanding the factors that influence the demand for space heating is therefore essential for this research. Traditional models rely on deterministic calculations to predict the consumption of energy. However, a significant discrepancy exists between these theoretical predictions and actual gas consumption, known as the 'energy performance gap' (Majcen, 2016).

This chapter will highlight previous research on the energy performance gap, and try to identify the most influential factors for space heating demand. Furthermore, energy demand models that include the influence of local climatic conditions and urban morphology on heating demand will be discussed. The following sections will also include the methods that are used in determining the demand for heating, urban climate studies and the role of local urban climate variations. Lastly, this chapter will discuss the use of local crowdsourced data and methods for handling local geographical data.

2.2 Determinants of residential heating energy consumption

Heating energy consumption can be estimated both directly and indirectly. The direct estimation involves obtaining data from local energy companies, investigating the physical aspects of a building and making use of questionnaires for costumers (Danylo et al., 2019). This direct approach requires a lot of effort and data collection, even on a smaller scale. The indirect approach is based on theoretical calculations that estimate heating demand based on the building energy performance. The building energy performance is dependent on four major groups of factors: Urban geometry, building design, systems efficiency and occupant behaviour (Ratti et al., 2005).

However, several studies have shown that there is often a discrepancy between theoretically calculated energy demand and the actual energy consumption (Majcen et al., 2013). This is generally referred to as the energy performance gap. Actual gas consumption differs from the theoretical gas consumption mainly due to differences in behaviour of occupants, but also due to different indoor preferences and the way dwellings are heated. For the Netherlands, the actual consumption is usually higher for dwellings with lower expected consumption, while dwellings with high expected gas consumption usually consume less (Sunnika-Blank & Galvin, 2012).

Van den Brom et al. (2019) found that for the Netherlands about 50% of the variation in residential heating consumption between households can be attributed to differences in occupants, and the other 50% to building characteristics and environmental parameters. Additionally, Namazkhan et al., (2020) used a decision tree approach and concluded that that household gas consumption was uniquely related to building characteristics, socio-demographics and psychological factors. Specifically, house size, residence type, and building age were the main building characteristics that are related to household gas use. They also noted that Income and employment status were the main socio-demographic variables related to gas use.

Despite the depth that both of these studies go in, they treat the environmental effect as a spatial constant, and use national averages or neglect it fully. However, outdoor temperature and local climate conditions directly influence heating demand (Mashhoodi et al., 2020). This research aims to account for this environmental effect by incorporating the local climatic conditions in the analysis. By doing so, it seeks to explain part of the variation in residential gas consumption that cannot be attributed to building characteristics or occupation behaviour.

The findings regarding the determinants are derived from studies in temperate climates in Northwestern Europe, similar to the Netherlands (Guerra Santin, 2010). In these climates with longer colder winters gas consumption is dominantly used for space heating

Based on the previous research in factors that explain heating energy demand, these factors can be split up into 3 groups; Building characteristics, socio-demographic factors and local climatic conditions (Guerra Santin et al., 2009; Namazkhan et al., 2020; van den Brom et al., 2018). All previous studies constantly show that the variation in residential heating demand is explained by the interaction of multiple factors. The following section will describe the specific group of factors and their influence on the heating demand.

2.3 Building determinants

The most important group of factors are the building characteristics. Around 42% of the energy demand of dwellings are affected by the building characteristics in the Netherlands (Guerra Santin et al., 2009). These characteristics influence the amount of heat required to reach comfortable indoor temperatures, and how efficiently the heat is retained in buildings. Van den Brom et al. (2018) and Guerra Santin et al. (2009) have quantified the contribution to the variance of gas consumption of these determinants, and the following will serve as control variables in the model for this study:

1. Dwelling size, with larger dwellings usually requiring more heat to reach the wanted indoor temperature, due to the larger volume of space that needs to be heated. This study will include the average floor area in m^2 in the model.
2. Building age, with 89% of neighbourhoods in the Netherlands having a positive relation with heating demand and building age (Mashhoodi et al., 2020). Brounen et al. (2012) found that gas consumption in dwellings built before 1980 is around 50% higher than dwellings built after 1980. This study will also use the building age as a proxy of energy efficiency, as buildings that are older lack the first major thermal regulations.
3. Insulation and Energy labels: Majcen et al. (2013) found that Energy labels can serve as a predictor for actual gas consumption. While building age can serve as a proxy for insulation standards, it does not account for post-construction adjustments e.g. additional wall insulation. Therefore, the energy label distribution will be included as an additional factor.
4. Dwelling type, where detached and single-family dwellings generally have higher heating requirements than apartments, due to greater exposure to outdoor temperatures (Voskamp et al., 2021). If proportions of apartments per postcode are high, heating demands can reduce due to better heat retention for example.

2.4 Socio-economic determinants

Occupant behaviour is methodologically difficult to measure, especially at a national scale. Therefore, occupancy characteristics or socio-economic characteristics are usually used a s a proxy for occupant behaviour (Guerra Santin et al., (2009). Occupant behaviour contribute around 4% to 15% to the variance of residential gas consumption in the Netherlands. While many factors can contribute, a selection was made based on the factors that deemed significant in studies that were conducted in the Netherlands:

1. Household size, both Mashhoodi et al. (2020) and Namazkhan et al. (2020) found that household size is an effective determinant of heating energy demand. Household composition influences the energy use, as families with children use heating more frequent and for a longer duration (Biesot & Noorman, 1999).

2. Income level: Households with higher income levels generally consume more gas. Namazkhan et al. (2020) found that households with an income above 12.000 euro use more gas than households with a lower income. Vringer (2005) found that in a study of 2200 Dutch households, a 1% increase in income resulted in a 0.63% increase in energy use. Other studies

that research gas consumption and income have similar findings (van den Brom et al., 2018; Mashhoodi, 2018).

Majcen et al. (2015) found that occupant characteristics in The Netherlands contributed a significant amount to the explained variance of gas consumption (7.5%). This study used extensive survey data and had access to preferred indoor temperatures and times of day when people were present in the dwelling. Furthermore, Namazkhan et al. (2020) also include psychological values in their research, as well as employment status (student, retired).

The selection of above socio-economic determinants is guided by data availability and theoretical importance, therefore aforementioned psychological factors and behavioural indicators are not used in this research. Together, these control variables capture the major non-climatic determinants of heating demand identified in previous research, allowing the remaining variation in gas consumption to be attributed to local climatic differences.

2.5 Urban Morphology, Form and Energy demand

While building characteristics and socio-economic factors establish the baseline for energy demand, the demand is further explained by local urban differences. Roughly 40-50% of the variance is explained by the building and socio-economic factors, leaving 50% of the unexplained variance.

One of the other factors that influences the heating demand in urban areas is the Urban Morphology. The physical structure of neighbourhoods, in this case postcodes, affects energy demand through the Surface-to-Volume ratio. Denser urban areas usually have more shared walls and are more compact. These areas benefit from this by having a reduced amount of heat loss compared to more rural layouts (Rode et al., 2014). They conclude that for Berlin, urban form explains 10-20% of variation after controlling for building and socio-economic characteristics. Ratti et al. (2005) used Digital Elevation Models (DEMs) to measure the urban surface-to-volume ratio and their influence on energy demand in three different European cities. The surface-to-volume ratio explained around 15% of heating demand differences between neighbourhoods.

Urban Heat Island effect

The Urban Heat Island effect (UHI) is a phenomenon that is well studied regarding energy use and temperature. The UHI effect is a phenomenon where urban areas experience significantly higher temperatures than their rural surroundings. UHI could lead to a median of 18.7% decrease in heating energy consumption (Li et al., 2009). The ‘Amsterwarm’ study by Van der Hoeven & Wandl (2015) provide a framework for understanding this effect in the Dutch context. They explain that the UHI does not occur constantly over a whole city, but is more like a fragmented spatial pattern that is influenced by the land use and the urban density.

In The Hague, similar to Amsterdam, dense urban centres and industrial zones accumulate heat on the surface while open green spaces and areas with water act as a cooling buffer. Van der Hoeven & Wandi (2018) again conducted research on UHI, this time in The Hague and found that during cold spells or heatwaves the temperatures can fluctuate by 8 – 13 degrees Celsius, with this effect being the highest in all of the measured Dutch cities (Klok et al., 2012). These differences only occur during heatwaves, with the average difference within the Heat Island being around 2.5 degrees Celsius.

For this research, an interesting dual effect is at place. The UHI can both lead to a reduction and an increase in heating demand. During the winter, the UHI acts as a thermal buffer and the retained heat leads to a reduction in space heating. Conversely, during the summer the UHIs do not cool down as much in the night causing more heat stress in these neighbourhoods. Findings by Rafiee et al. (2019) confirm this, as they found that dense Dutch urban neighbourhoods with less open space had lower heating demand. As the UHI is unevenly distributed over the city, the Netatmo data can capture the UHIs in The Hague.

Instead of defining Local Climate Zones, this study operationalizes the UHI using morphological factors that capture the UHI. In the case of The Hague these factors are the height of buildings, the density of buildings and the amount of vegetation in the area, captured by the NDVI and land use data van der Hoeven & Wandi (2018). The density of buildings in turn also represents the amount of sun that is reflected.

The factors named above all influence the demand for gas consumption as they directly influence the temperature. Indirect effects can also occur, such as the heat retention in buildings leading to less energy efficiency (Sharston & Singh, 2025). Rafiee et al. (2019) found for example that for Dutch cities every 1m increase in building height, they noticed a 111,3 m3 increase in gas consumption. A part of this is explained by that taller buildings have a lower volume-to-surface ratio. Wouters et al. (2017) research the impact of land cover in Belgian cities and found that land cover explains around 8 – 12% of local temperature variation.

2.6 Degree Days

To integrate the aforementioned climatic and morphological effects into a energy model, a meteorological metric is required. The concept of Degree Days is an established meteorological metric used globally to quantify the influence of ambient temperature on the demand for space heating (Heating Degree Days, HDD) or cooling (Cooling Degree Days, CDD) (Eto, 1988). The Degree Days serve as a normalization factor that isolates the temperature component of energy consumption.

The standard calculation of a Heating Degree Day (HDD) is a predetermined reference temperature minus the daily average temperature. This predetermined reference temperature in this case is 18 Degrees Celsius, the standard for Netherlands and other European countries and the needed external temperature to maintain the wanted 20 degrees Celsius internal temperature (Wever, 2008). The HDD are only measured when the daily average temperature is below the reference temperature. The formula for this calculation is defined as:

$$HDD = max(0, 18\ C\,° - Tmean\}$$

Traditionally, HDD used for energy models and historical analysis are derived from national meteorological stations. However, as Wolters et al. (2011) note, temperatures can differ up to 4 degrees when due to surface differences and urban heat. This makes it possible for the local temperature to be noticeably different over a distance of just a few kilometres. Figure 2.1 shows that differences in average regional temperature do certainly exist. More local climatic effects, like the influence of the sea on The Hague are not captured by the regional Degree Days. By calculating Heating Degree Days on a local scale, this research will try to capture the actual temperature at a local scale, incorporating UHI and urban form into the energy demand.

When the HDD are derived from the national KNMI network, the values represent the regional climate and are nearly uniform across The Hague due to the low station density. When the HDD are derived from the Netatmo network, the values represent the locally measured temperature and values differ across short distances capturing the intra-urban variations, including the UHI effect. Both HDD datasets represent a different scale of spatial climate data, and their relative performance in explaining gas consumption is tested.

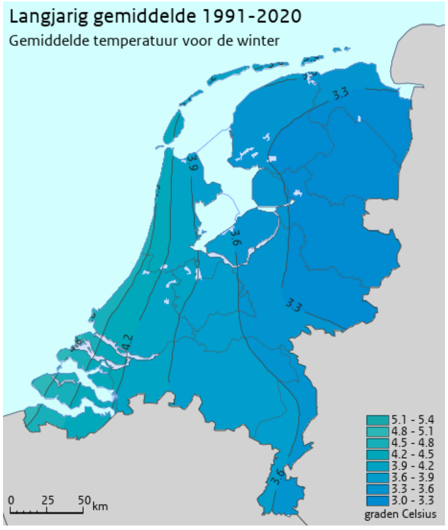

Figure 2.1: Average temperature in winter in Degree Celsius, 1991-2020, KNMI

2.7 KNMI data

The Royal Netherlands Meteorological Institute (KNMI) is the primary authority for weather and climate observations in the Netherlands. The national network consists of 35 high precision synoptic stations scattered around the Netherlands. The stations are designed to capture the large-scale atmospheric conditions, following the rules of the World Meteorological Organization (WMO).

KNMI purposefully locate their stations in open rural areas at a standard height of 1.5 meters (Wever, 2008). This is done to ensure consistency between stations. While this network provides highly calibrated and reliable longitude data, the spatial density is low. The stations usually represent a larger region and therefore have to stay ‘blind’ to microclimatic variations. Stations are usually in open areas such as fields or farmland to reduce the potential urban influences on the climate. Incorporating the regional climate using Heating Degree Days into the energy model can be done using the national network of the Netherlands, with consistent historical data.

2.8 Crowdsourcing and Netatmo data

Incorporating the regional climate using Heating Degree Days into the energy model can be done using the national network of the Netherlands, the KNMI (Wever, 2008). The shift to the implementation of more local climate data requires a higher spatial density of meteorological observations that the national network can provide.

Traditional meteorological networks are designed for synoptic observations and are often located outside urban areas to avoid urban heat influences (Muller et al., 2015). As a result, these networks are less suitable for intra-urban research. Crowdsourced data, such as Netatmo or WOW-NL, offers a solution by providing high-resolution citizen sensing data from within urban dense areas.

The cost of this high-density network of sensing data is that it lacks the professional calibration, storage capabilities and the accuracy that the KNMI network offers. In other studies that use citizen sensing data, it is viewed as a strategic trade-off: the absolute sensor accuracy is sacrificed in favor of increased spatial coverage (Chapman et al., 2016). In contrast of the KNMI network, where sensors are placed in open areas where few people live, the Netatmo sensors are located within the urban areas where people actually live (Demuzere et al., 2017). Moreover, because the sensors are located within urban neighbourhoods they are effective for capturing and quantifying the UHI effect within these neighbourhoods, as well as identifying local temperature variations (Voskamp et al., 2021; Rød & Maarse, 2021).

2.9 Conceptual model

Figure 2.2 represents the conceptual model of this research. It shows that residential gas consumption is influenced by building characteristics and the socio-economic characteristics of the occupant. The local climate is captured using citizen sensing temperature data. These climatic conditions are translated into Heating Degree Days which represent the climate-related component of gas consumption together with the influence of the urban morphology.

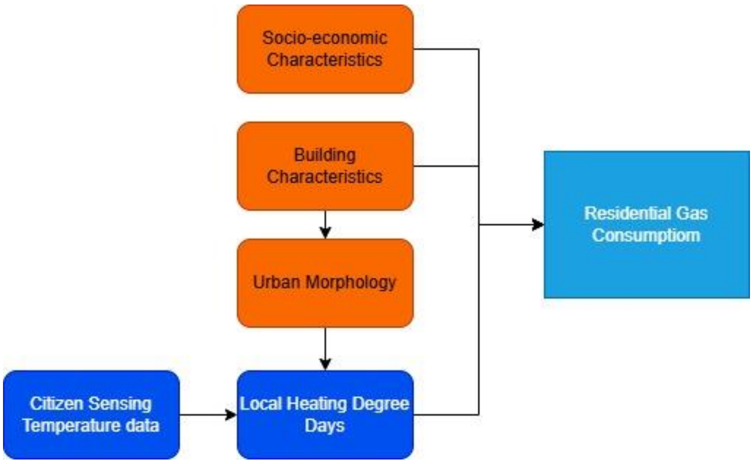

Figure 2.2: Conceptual model of theoretical framework

3 Methodology

3.1 Research design

This research will adopt a quantitative geospatial methodology to investigate how local climatic variables influence residential gas consumption at the postcode level.
The methodology will consist of multiple phases: Data collection, interpolation and aggregation of the data, creating Local Degree Days, developing the model and statistical analysis, and lastly mapping and visualizing results. After this visualization the results are interpreted for The Hague and the relevance for future policy for other municipalities in The Netherlands.

3.2 Data preprocessing

To ensure that the analysis goes correctly, all the datasets were investigated and aligned. The obtained datasets come in different spatial resolutions and temporal scales, as well as different data formats. While some datasets were GIS compatible, others needed to be transformed. The Postcode 6 boundary dataset from CBS Statline serves as the main dataset, as the spatial postcode boundaries are already defined. All other datasets are joined in this dataset, linking the explanatory variables to each postcode. Gas consumption data from CBS was obtained from the energy consumption dataset per PC6, and could be directly linked to the postcode boundary.

All variables are integrated and cleaned in R before model estimation. CBS suppression codes (-99997) are recoded as missing values, as -99997 indicated that there were privacy suppressed values, and not zero values. Postcodes with missing values on a core model variable were excluded. Joining BAG data with postcode data could not be done in R, and was performed in ArcGIS pro. Every building in The Hague was joined with the postcode the building has the most area in. Values were then transformed to return a mean per postcode. Exact processing steps for each variable can be found in chapter 4.

The meteorological data from KNMI consists of 2 sections. The first sections provide the station name, station id and geographical X and Y coordinates for each station. The second section provides the daily meteorological observations for each station id. The first section was split using Microsoft Excel and transformed to ArcGIS Pro, where the X and Y coordinates could be transformed to point data distributed across the Netherlands. The second section was first cleaned, and all stations without yearly coverage were deleted from the data. Subsequently, the daily temperature averages were then connected to each weather station. After calculating the Degree Days, the data was aggregated to yearly summaries to match the temporal resolution of the gas consumption data.

3.3 Control variables and operationalisation

Based on the sections in the theoretical framework (2.3 & 2.4) a selection was made of control variables that will account for the non-weather determinants of gas consumption. This section will explain on how these factors are represented in the model, and the data sources were identified. All variables were pre-processed and cleaned where needed.

The main source of the variables is the Centraal Bureau voor Statistiek (CBS) and the data they collected on the PC6 areas of the Netherlands in 2021. CBS excludes postcodes who have less than 5 in the categories population, households or dwellings. They also exclude postcodes that must remain secret. If that is the case, the data will be published with the number -99997. These numbers were transformed to missing data and excluded from this study, along with all the postcodes who have missing values on any of the variables.

Building characteristics

According to Brounen et al. (2012) buildings before 1980 consume more gas for heating due to less standards in building isolation compared to newer buildings, especially dwellings that were built after 2000. Rather than relying on the CBS for the data on building age classification, the variable was derived directly from the 3DBAG dataset (3D geoinformation research group, 2025). Each dwelling was classified as pre-1975 or post-1975 based on the building year (oorspronkelijkbouwjaar), and the share of pre-1975 buildings per PC6 was calculated. This approach ensures that the variable was complete instead of a variable that omitted values due to privacy.

As stated in multiple studies, the dwelling size affects residential gas consumption the most, as more space requires more gas use. The average usable floor area in m2 will be retrieved from BAG Verblijfsobjecten (Kadaster, 2021), where the 2021 data is extracted using QGIS and values were aggregated to the PC6 level using a spatial join to the PC6 polygons. To only represents the residential buildings, the BAG Verblijfsobjecten dataset was used instead of the 3DBAG. The 3DBAG dataset does not have a variable indicating what a building is used for, meaning that all structures within a postcode count towards the mean square meters. The mean area was calculated using all the qualifying residential units within each postcode. This represents the direct physical relation between the volume of space and the energy that is required to heat it.

For dwelling type the amount of multi family homes will be used, according to Voskamp (2021) apartments use less gas due to better heat retention from other apartments. This will be represented by the share of apartments or multi family homes. This was operationalized by using the BAG Verblijfsobjecten dataset and the 3DBAG dataset, where each BAG pand was classified as a multi family home if it contained more than 1 registered verblijfsobject. This reflected the presence of multiple dwellings within one building. The share of multi family homes was calculated by dividing it by the total amount of residential buildings within each postcode. This was again preferred over CBS data, as the CBS omitted around 40% (4605 areas) of all PC6 areas included in the analysis.

Additional data cleaning is done by calculating the percentage of vacant buildings. The CBS published the amount of homes that are occupied and the amount that is unoccupied (CBS Kerncijfers, 2021). Unoccupied homes do not consume gas, but the CBS calculates the average gas consumption by dividing it with the total amount of homes per postcode area. The share of vacancy was calculated by dividing the amount of vacant homes with the total amount of homes within each postcode. The CBS does not publish vacancy numbers for small postcodes due to privacy reasons. Areas were the vacancy rate was omitted for privacy reasons were still included in the analysis. All postcodes with an available vacancy rate higher than 10% were excluded.

Energy label distribution is represented by the share of dwellings with a high-efficiency label (A+++++, A++++, A+++, A++, A+ or B) per PC6. This was derived from the municipal dataset of energy labels from The Hague. This variable captures all post-construction insulation improvements that the building age alone can not account for, following Majcen et al. (2013). As energy label coverage is incomplete for some postcodes areas, the share of high efficiency labels (Label B and better) will serve as a complementary variable to building age. It is tested separately to avoid reducing the postcode sample.

Socio-Economic characteristics

The household size is represented by the average household size per PC6 directly published by the CBS, so no additional filtering is done. Income level will be represented using the average income of the households. Due to privacy regulations, the CBS will not display income levels of PC6 areas with less than 100 values. This conflicts with the PC6 data, where the small scale areas usually have less than 100 dwellings. Therefore, the WOZ-waarde is used as a proxy for income. While not the same as income, it measures the worth of the property and can serve as an indicator of people that are living comfortably. The WOZ waarde is also not suppressed like other CBS variables, as WOZ waarde is not personal but the worth of the property. The variable is calculated by calculating the mean WOZ-waarde per PC6, and this value is log-transformed in order to account for right-skewed property values and ensure a better fit (Arentze & Stevens, 2018). A limitation by using WOZ waarde is the potential mismatch between WOZ waarde and income, particularly for elderly homeowners in older neighbourhoods.

Although age was initially considered as a determinant of heating demand following Rafiee et al. (2019), the CBS variable for age classes was suppressed for 5797 of the 11440 PC6 areas, making it unsuitable for inclusion in the main model. Age is therefore excluded from the final analysis.

3.4 Urban morphology and Urban Heat Island characteristics

Following section 2.5 from the theoretical framework, the UHI effect is operationalized using measurable morphology variables. The amount of vegetation is used as a proxy for urban green spaces. The Normalized Difference Vegetation Index (NDVI) is a measurement on the presence of green surfaces and is captured by the Dutch National Institute for Public Health and the Environment (RIVM, 2022). Higher values of NDVI indicate more vegetation and therefore also weaker UHI intensity. The Urban density is operationalized using the building coverage ratio within each PC6 polygon, retrieved from the BAG verblijfsobjecten dataset (Kadaster, 2021). Using the geometry of building footprint divided by the total surface area, a coverage ratio per PC6 will be calculated. Higher building coverage suggest a more compact urban fabric, potentially reducing the amount of HDD required. Building height is obtained using the 3DBAG and mean height (at the 70th percentile) per postcode 6 area will be calculated. Taller buildings retain more heat during the night, as well as acting as wind shields for other buildings (Bansal & Quan, 2022).

Table 3.1: Overview of all datasets in the model along with the data source, spatial and temporal resolution.

Dataset	Description	Source	Spatial resolution	Temporal resolution	Usage
Residential Gas Consumption	Annual gas usage per postcode	CBS Statline	PC6	Annually (2021)	Dependant variable
KNMI weather observations	Official meteorological observations (temperature, wind, humidity, precipitation)	KNMI	Point data, 44 across Netherlands	Hourly/daily/yearly (2021)	Creating KNMI Interpolations for HDD
Netatmo (The Hague)	Crowdsourced microclimate observations (temperature)	Netatmo	Point data, The Hague	15 min/daily (2021)	Creating Netatmo Interpolations for HDD
Kerncijfers CBS	Population, dwellings, average income, building characteristics, etc.	CBS Statline	PC6, 13.483 PC6 areas The Hague	2021	WOZ-waarde, Household size
Postcode boundaries	Polygon boundaries for Dutch postcodes	PDOK	PC6	2021	Boundaries for interpolation and visualization

Dataset	Description	Source	Spatial resolution	Temporal resolution	Usage
Basisregistratie Adressen en Gebouwen	Data of all dwellings in the Netherlands	PDOK	Building level	2021	Building age, multi-family classification, building coverage, building height
Normalized Difference Vegetation Index (NDVI)	Index of the amount of green in the growth season	RIVM	10m x 10m	2021	Index of green spaces (0-1)
Energy Labels	Energy labels (A+++++-G)	Municipality of The Hague	Building level	2021	Known energy labels from A+++++ to G

3.5 Netatmo data quality

The Netatmo dataset is similar to the KNMI dataset, and also needs to be transformed in order to match the temporal resolution of the research. These meteorological observations are carried out by amateurs and therefore fall under citizen science. While not as accurate as the KNMI data, the coverage of these observations is more useful than the accuracy and the data may be used. To be as consistent as possible, the data is filtered after the Degree Day calculation.

As the Degree Days will be directly compared with the KNMI degree days, the number of days will need to be exactly the same. Therefore, stations without observations for the whole year are filtered out. As described before the dataset consists of .txt files of 15 minute intervals covering the whole of 2021. Each individual observation was grouped using the station id. Then, the 15 minute observations were transformed to daily averages. Now each station was filtered (<365) to ensure data quality before computing the Degree Days. The result is 116 stations that provided observations for each 15-minute interval.

The Netatmo data is filtered using a quality control scheme set up by Meier et al. (2017) and was developed to filter noise and outliers from crowdsourced Netatmo data without losing important urban climate effects. First, the data was checked for potential indoor stations. According to Napoly et al. (2018), indoor stations are one of the most common error sources in crowdsourced temperature data. To filter these indoor stations out of the data, the amplitude(Tmax-Tmin) is calculated. Stations that have an amplitude of <15 Celsius are filtered out, as a low amplitude means that the station was located indoors (or a different incorrect measurement). This resulted in deleting 3 invalid stations, with all stations having < 2 heating degree days. Total amount and percentage of deletion is visible in table 3.2.

Table 3.2: Filtering of Netatmo stations

Amount of stations	Type of quality improvement	Remaining stations
570	Removing of stations with <365 days of measurements	116
116	Removing stations with unrealistic amplitudes	113
113	Filtering of outliers using the IQR	105 (-81,6%)

For improved data quality, an outlier detection method will be used. The mean, the first quantile (Q1) and third quantile (Q3) of the heating degree days will be noted (see figure 3.3). To filter out the outliers, the Interquartile Range (IQR) was identified using the first quartile and the third quartile, and all outliers outside of the 1.5 x IQR range were then deleted from the analysis. According to the quality scheme, the filtered Netatmo data is now at quality level M2 (Meier et al., 2017). This resulted in the deletion of 8 stations, with the total remaining valid stations sitting at 105.

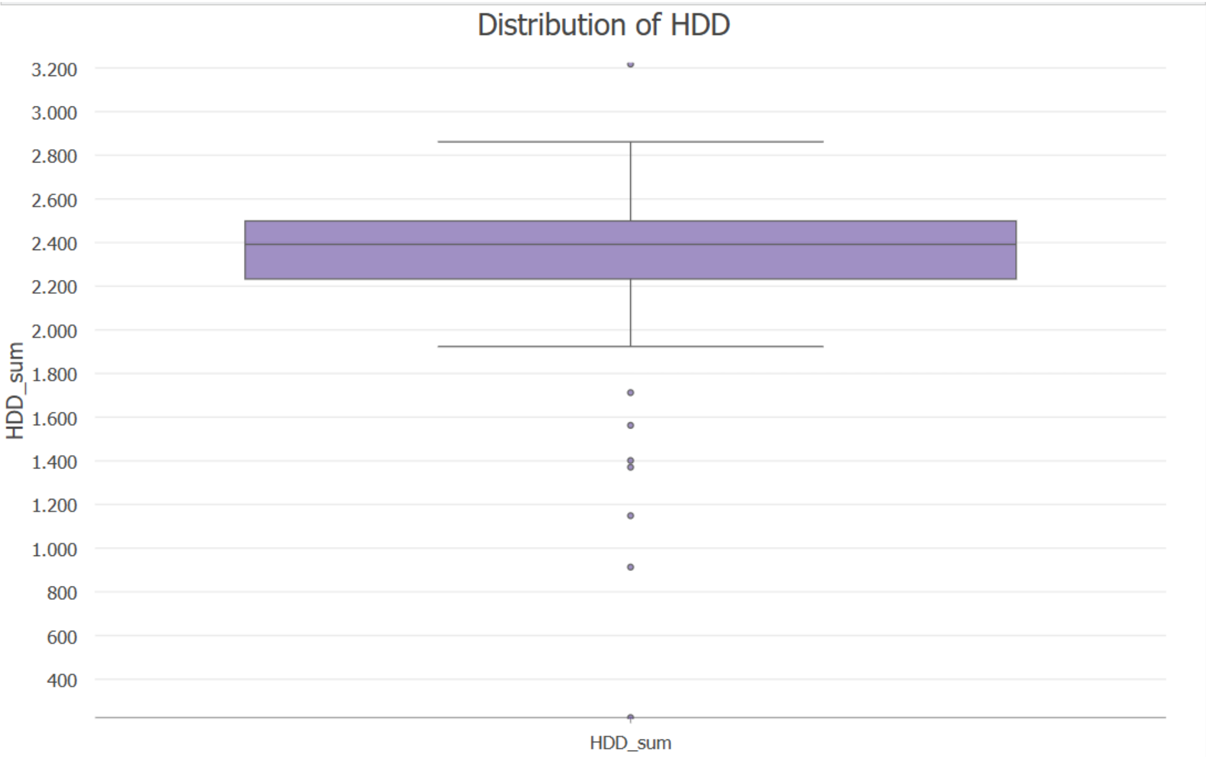

Figure 3.3: Box plot of Heating Degree Days derived from Netatmo stations via Ohori (2022)

All spatial datasets were projected to the Dutch coordinate system (RD New, EPSG:28992) to ensure that calculations will be consistent and that stations are projected at the correct corresponding location.

3.6 Degree Days calculation

To quantify the influence of the local temperature on heating demand, Heating Degree Days (HDD) were calculated. Degree days are commonly used in heating demand studies (Eto, 1988). The formula that this research uses is based on the formula from Mourshed (2012) who investigates the relationship between Degree Days and annual mean temperature. Similarly, Mourshed (2012) also uses the annual mean temperature from daily mean temperature observations to calculate HDD.

First, the HDD were calculated for each day, based on the mean daily temperature:

$$HDD_i = \max(0, T_{base} - T_{mean,i})$$

where:

- Each day = i
- T_{base}= 18°C
- $T_{mean,i}$= daily mean outdoor temperature
- HDD is set to zero when $T_{mean,i} > T_{base}$

Daily HDD values were summed to obtain annual Heating Degree Days per station:

$$HDD_{annual} = \sum_{i=1}^{365} HD\,D_i$$

Cooling Degree Days were not calculated for the stations, since Dutch residential gas consumption is predominantly based on space heating. Cooling Degree Days in the Netherlands are generally low, and demand for space cooling is lower compared to other European countries (Rovers et al., 2021). Additionally, space cooling is currently mostly absent in Dutch homes. This is changing however, with the increased amount of homes on a heat pump installation and increasing annual temperatures (de Wind et al., 2025).

3.7 Spatial interpolation of climate

Regional and local temperature data will be collected and will function as the main independent variables. The first dataset that will be used is the KNMI meteorological observations for the year 2021. 44 weather stations provide professional daily observations on temperature, humidity, wind speed and precipitation. This data will be spatially interpolated for nationwide coverage, as weather stations measure are static and measure from 1 set point**.** The raw temperature data from the 44 KNMI weather stations will be spatially interpolated to the centroid of every PC6 area (within The Hague). This approach of estimating local weather variables from the official network is a common and established methodology when using Dutch climate data (Sluiter, 2012). The result of the KNMI interpolation is that each postcode area gets assigned a regional HDD value derived from the KNMI measurements. The process described above is repeated for the Netatmo station measurements. This will create the Local Heating Degree Day variable for each PC6 area.

This spatial interpolation will be done using two different established methods, Inverse Distance Weighting (IDW) and Ordinary Kriging. These methods will be further explained in the following paragraphs. The results of the interpolation methods will both be compared to see which methods captures the climatic effects in The Hague the best. A limitation of the interpolation is the spatial coverage of the Netatmo and KNMI stations. The interpolated HDD values are extrapolated outside of the station network rather than only inside. This leads to some values being reliant on only a few distant stations, making the temperature estimates less reliable.

To interpolate the spatial distribution of climatic variables, several geostatistical approaches are commonly employed, including Inverse Distance Weighting (IDW), thin-plate spline regression and Kriging. These methods have been extensively used and compared on studies using temperature and precipitation (Hofstra et al., 2008; Stahl et al., 2006). This research will use Ordinary Kriging and IDW.

Ordinary Kriging

Ordinary Kriging, unlike other mentioned methods, is based on a stochastic framework and not just distance. This makes it suitable for interpolating the local climate, as temperature observations typically have spatial dependence. Kriging assigns weights based on the spatial covariance structure of the data. This allows the model to account for clustered observations and spatial trends in temperature measurements.

Kriging is a spatial interpolation method that predicts a value for a location based on the spatial autocorrelation of observed values (Goovaerts, 1997). Simply put, this means that two neighbouring points are more likely to have similar values than two points that are further apart. It predicts the target observation at an unknown location by taking a weighted average from nearby reference data where weights are based on the spatial variance of the input data with the location information (Wackernagel, 2003). A semivariogram model was used to define the relationship between spatial variance and physical distance. Furthermore, to account for the geographical orientation of The Hague, the model incorporated the anisotropy. This was done to reflect the coastline and of the North Sea into the model, as the spatial autocorrelation is higher parallel to the shore. Results of the model output are visible in table A2. To ensure more reliability, the interpolation values are at least based on the 12 closest neighbour observations.

Inverse Distance Weighting

Inverse Distance Weighting (IDW) is a deterministic interpolation method that estimates unknown values as a weighted average of nearby observations. In contrast to Kriging, IDW assumes that spatial similarity decreases with distance and does not explicitly model spatial autocorrelation (Van Ackere et al., 2015). Closer values will have a larger influence on the estimated values than sample points that are further away.

3.8 Statistical methods

The statistical analysis is performed in the statistical program R and follows a stepwise approach in order to determine the influence of the variables on gas consumption. In line with the methodology of Mashhoodi (2018) who seeks the local determinants of residential gas consumption in de Randstad, this process will follow two phases. A non-spatial estimation and a spatial estimation. A standard linear regression is first developed and will serve as the baseline.

The baseline model consists of the Ordinary Least Squares (OLS) model and includes only the socio-economic and building characteristics without accounting for the spatial location. Gas consumption in urban areas often experience spatial autocorrelation. Spatial autocorrelation is the tendency of neighbouring postcodes having similar values of gas consumption. This is checked by the Moran's I index on the model residuals. The Moran's I is a measure of spatial autocorrelation, and will give a value between -1 and +1 with values close to zero showing a random distribution. If this test indicates significant spatial dependency, the analysis will then select a spatial regression model. Based on the approach of Anselin (1988), the appropriate spatial model will be selected based on the results of the Rao's score test. Depending on this outcome, either the Spatial Error Model (SEM) or the Spatial Lag Model (SLM) will be selected. The SLM will be used when gas consumption in one postcode is influenced by gas consumption in neighbouring postcodes. The SEM will be selected if the dependency is caused by unobserved spatial factors, like neighbourhood types or building characteristics not taken into account for in the model.

Throughout this research, a significance threshold of $p < 0.05$ is applied as the standard for statistical significance. Values between 0.05 and 0.10 are reported as marginally significant. These results should be interpreted carefully since they do not meet the usual threshold but may still show an interesting trend.

3.9 Model specification

To test the added value of the local microclimate and morphological variables, the variables are divided into four different models. Each model builds upon the previous one:

- Model 1: Baseline. This model includes building characteristics (building age, dwelling size, dwelling type) and socio-economic factors (household size and the log-transformed WOZ-waarde).

- Model 2: Regional climate. Regional Heating Degree Days (HDD) derived from KNMI stations are added to the baseline.

- Model 3: Local microclimate & morphology. In this model, the KNMI data is replaced by local Netatmo HDD, alongside urban morphology variables (NDVI, building coverage, and building height).

- Model 4: Energy label test. As an additional step, the model is run including the energy label variable for the subset of postcodes where this data is available. This serves to check whether insulation improvements can further explain the effects originally attributed to building age.

3.10 Model comparison

To test the final research question, the KNMI and Netatmo models will be compared. Model performance for the OLS models is evaluated using the R^2 and Akaike Information Criterion (AIC) metrics, while the Spatial Error Models are evaluated using the Pseudo R^2 and AIC. A lower AIC indicates a better model fit

Several indicators will be used, such as: R^2 to measure the proportion of the variance explained and the AIC to compare the model quality. The spatial structure of residuals will be compared of both models, to check the extent of the inclusion of Degree Days and will allow the research to conclude whether the microclimate adds new information for explaining gas consumption.

3.11 Software and tools

This research makes use of Geographical Information Systems (GIS), combined with statistical tools to process, analyse and visualize the spatial and non-spatial data. QGIS software was used to obtain the 3DBAG dataset and extract the building data for The Hague. The main GIS environment that was used is ArcGIS Pro. ArcGIS Pro handled the spatial joins, and the interpolation of the point data. Lastly, it was used for the visualization of results.

The statistical analysis was performed in the R program. The datasets are imported from ArcGIS into RStudio to prepare the data, estimate OLS and spatial regression models, and compare model performance.

4 Study area and data

4.1 Study area: The Hague

The area of research will be the municipality of The Hague. The Hague is a city in the southwestern part of the Netherlands. The municipality consists of a range of different neighbourhoods with different characteristics that influence the climate on a micro scale. Both low-rise and high-rise urban neighbourhoods, green areas along the coast and lesser dense suburban neighbourhoods make up the municipality. Furthermore, with The Hague being a coastal city, sea influence and wind exposure may affect the climate on a micro level. The municipality also contains a large number of postcodes which makes it suitable for measuring fine scale differences, visible in figure 4.1. The urban density of The Hague also makes it that more Netatmo weather stations are active.

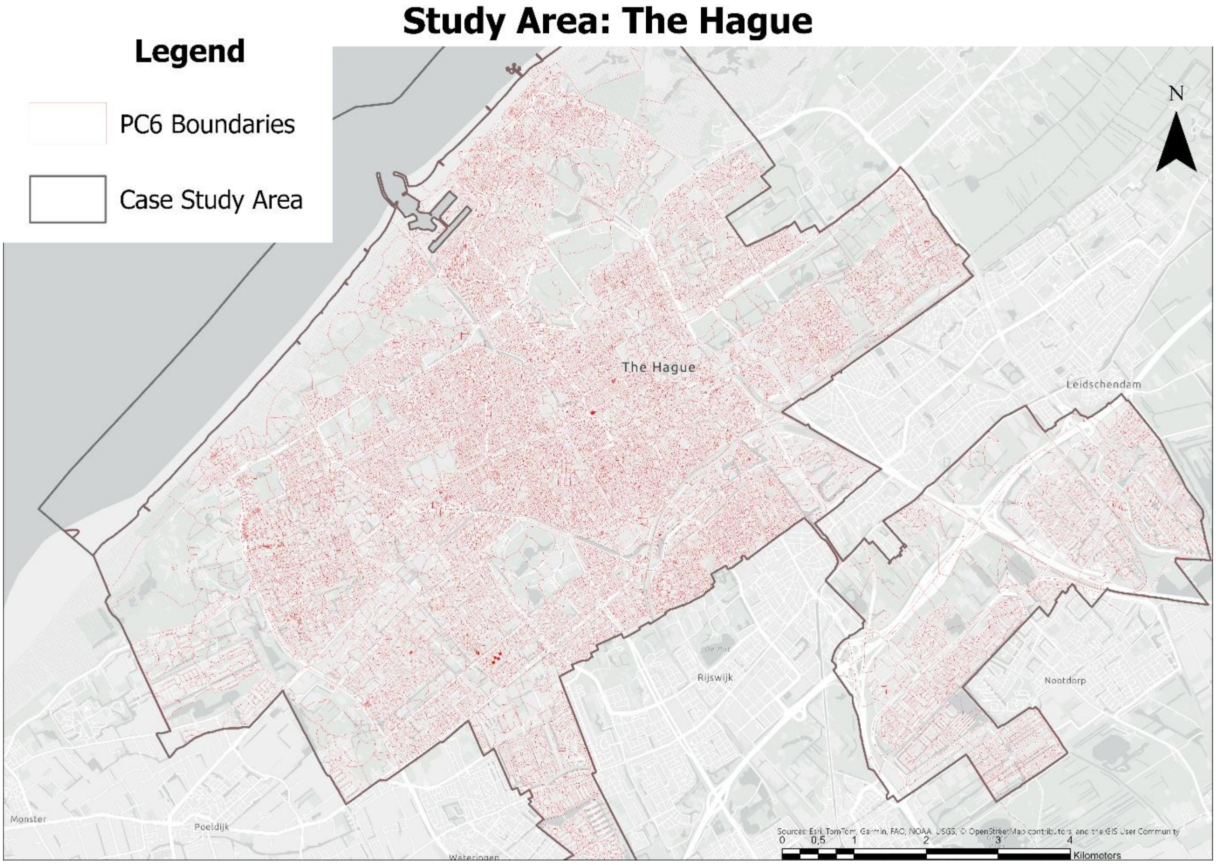

Figure 4.1: Study area and postcode 6 boundaries Source: CBS(2021);PDOK (2021).

4.2 Data

This section will consist of mainly describing on how the data is collected. Gas consumption data for The Hague will serve as the dependant variable. This dataset is the total yearly gas consumption at postcode 6 level in 2021. The year 2021 consisted of the last publicly available citizen sensing network data, as well as the most recent gas consumption data. Around 470.000 data entries exist in the dataset, with 13483 data entries for The Hague.

Gas consumption data

Residential gas consumption data is obtained from the Dutch Central Bureau of Statistics (CBS). The Statline data platform provides gas consumption at Postcode 6 level. The data is presented per postcode 6 area, but can be aggregated to both postcode 5 and postcode 4. The data is split up in consumption for homes and for businesses, including gas consumption and energy consumption. The numbers are calculated using data from the grid operators, making the source trustworthy. This split is based on sources from the Dutch Key Register of Addresses and Buildings, where the main function of the addresses located within the postcode decides the ‘residential’ or ‘business’’ group. For this research, the residential gas consumption will serve as the primary dependant variable. A preview of this table is presented in table A1

On average, each postcode 6 area includes 15-30 residential buildings and is mostly focussed on street level. Important to note is that this dataset includes the homes that are supported by district heating. Postcodes with district heating have low or zero gas consumption. Furthermore, areas with less than 5 dwellings are omitted because of privacy reasons. Furthermore, temperature corrected gas consumption is also included in the dataset, but will not be used as the corrected number removes the effect of the temperature on gas consumption.

The focus will be limited to analysing natural gas use for space heating. Gas use data on the postcode level however includes gas used for hot water and cooking, which are not sensitive to the external climate. Gas consumption of Dutch residential buildings is predominantly used for space heating, with cooking, hot water and lightning only accounting for ~20% share of the total consumption (RVO, 2024). This gas consumption is therefore a suitable proxy for heating energy demand on the postcode 6 scale.

Filtering of gas consumption data

The available data does not explicitly identify neighbourhoods connected to district heating. The values for gas consumption in these neighbourhoods will be low or 0. The abnormal values were filtered out for the model. The method of Mashhoodi (2018) was considered, where the author encountered the same problem when handling gas consumption data. Here, all incidents with z-value <= −2.5 or z-value >= +2.5) are identified as outlier and excluded. Figure 4.1 shows the distribution of residential gas consumption within The Hague. Due to the large amount of missing or 0 values, a different method was chosen.

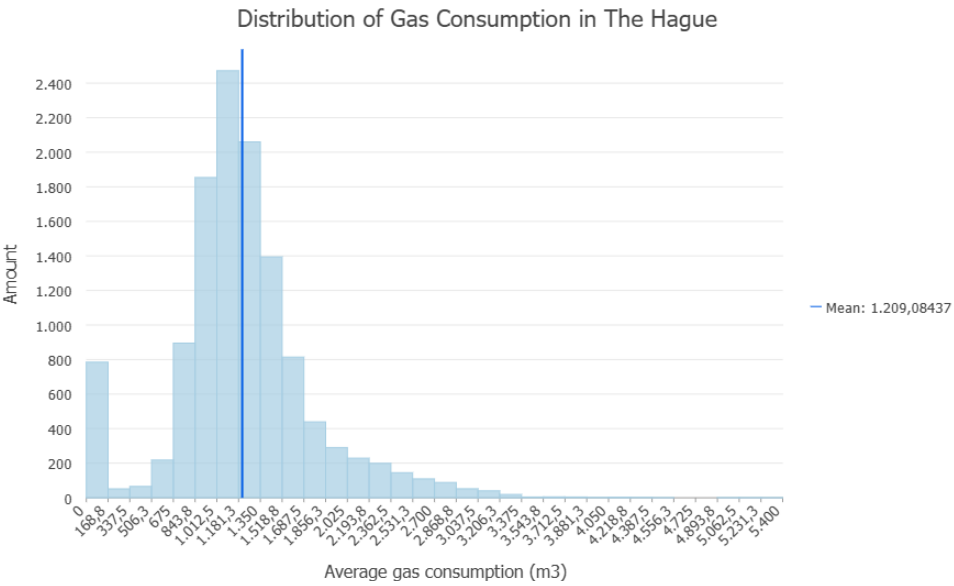

Figure 4.2: Histogram of Gas Consumption in The Hague. Source: CBS (2021)

Of the 13483 PC6 areas in The Hague, 1240 had a missing value. These missing values can be due different reasons: CBS suppresses values for postcodes with fewer than 5 dwellings for privacy reasons, postcodes with district heating have zero or low gas consumption, and postcodes with no registered addresses have no gas consumption. These missing values, along with zero or low values of gas consumption were filtered out from the study area as they are not representative of usual gas consumption numbers. 12243 areas are still left in the analysis with valid gas consumption values and are visible in figure 4.3.

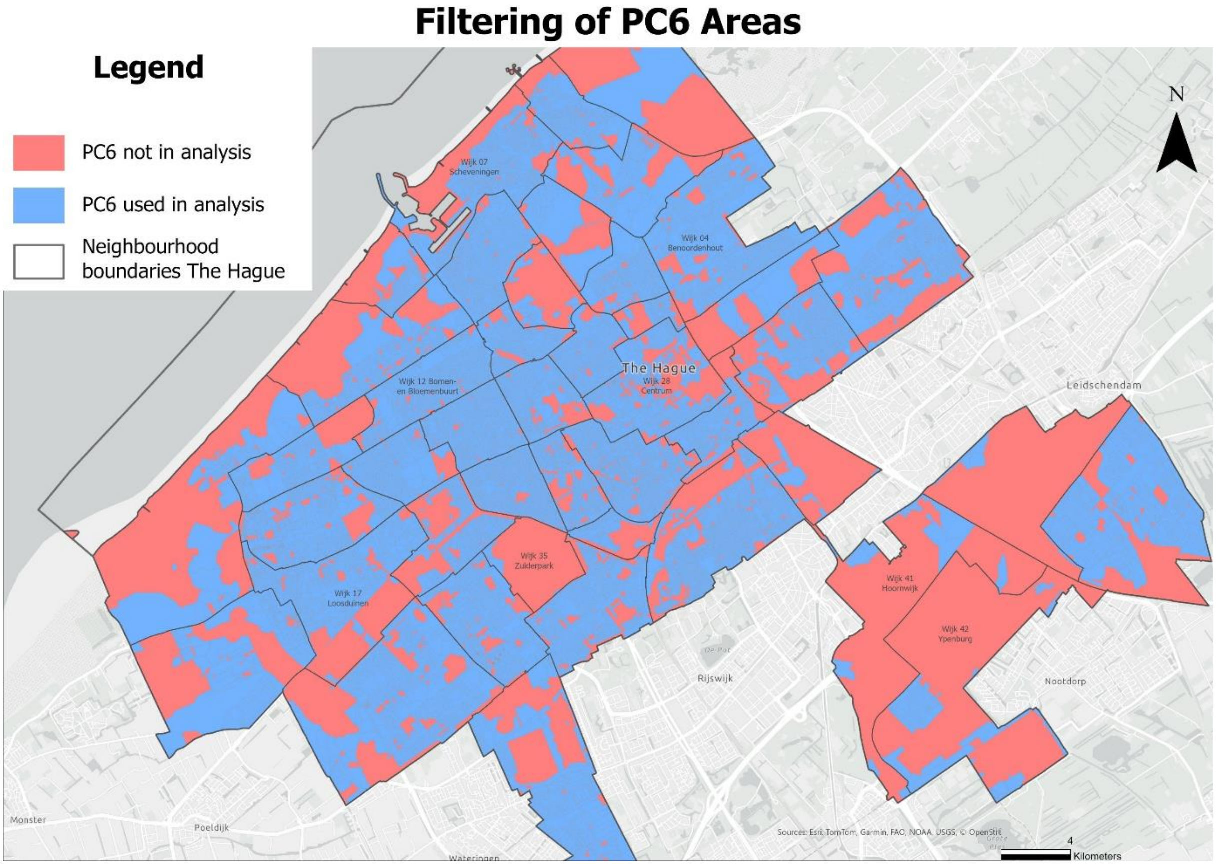

Figure 4.3: Excluded PC6 areas in analysis filtered by gas consumption. Source: CBS (2021); PDOK (2021)

Figure 4.4 shows the spatial distribution of average residential gas consumption per PC6 area. A pattern is visible, with higher concentration of gas consumption in the northern and western parts of the city, outside of the city centre. The lower average consumption values are found in the inner city and to the south of the city centre, in the Loosduinen and Morgenstond neighbourhoods. Large grey areas indicate that there either was no or low gas consumption or that the postcodes are connected to the district heating network.

Spatial Distribution of Residential Gas Consumption in The Hague (2021)

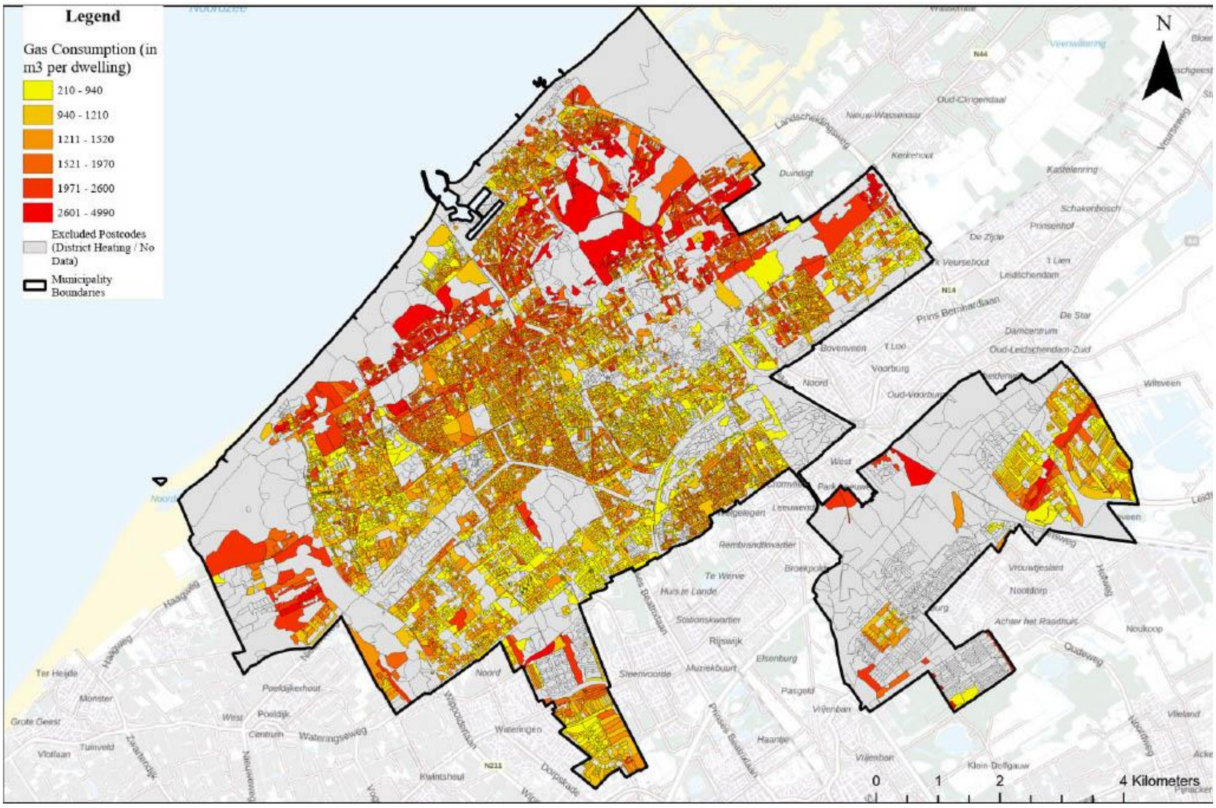

Figure 4.4: Spatial distribution of residential gas consumption. Source: CBS (2021);PDOK(2021).

4.3 Temperature data

KNMI data

Nationwide meteorological observations are obtained from the national weather station network of the KNMI. The KNMI is the national standard of climate observations and provide daily observations of key atmospheric variables, including temperature, wind speed, humidity and precipitation. The network is evenly distributed across the Netherlands, ensuring high measurement accuracy. On a daily basis, all observations are validated based on correctness and completeness. The validated data is then archived, ensuring data integrity (KNMI, 2000). Locations of weather stations are in open and standardized environments, minimizing the chance of inconsistencies and disturbances. The difference between station placement is visible in figure 4.5.

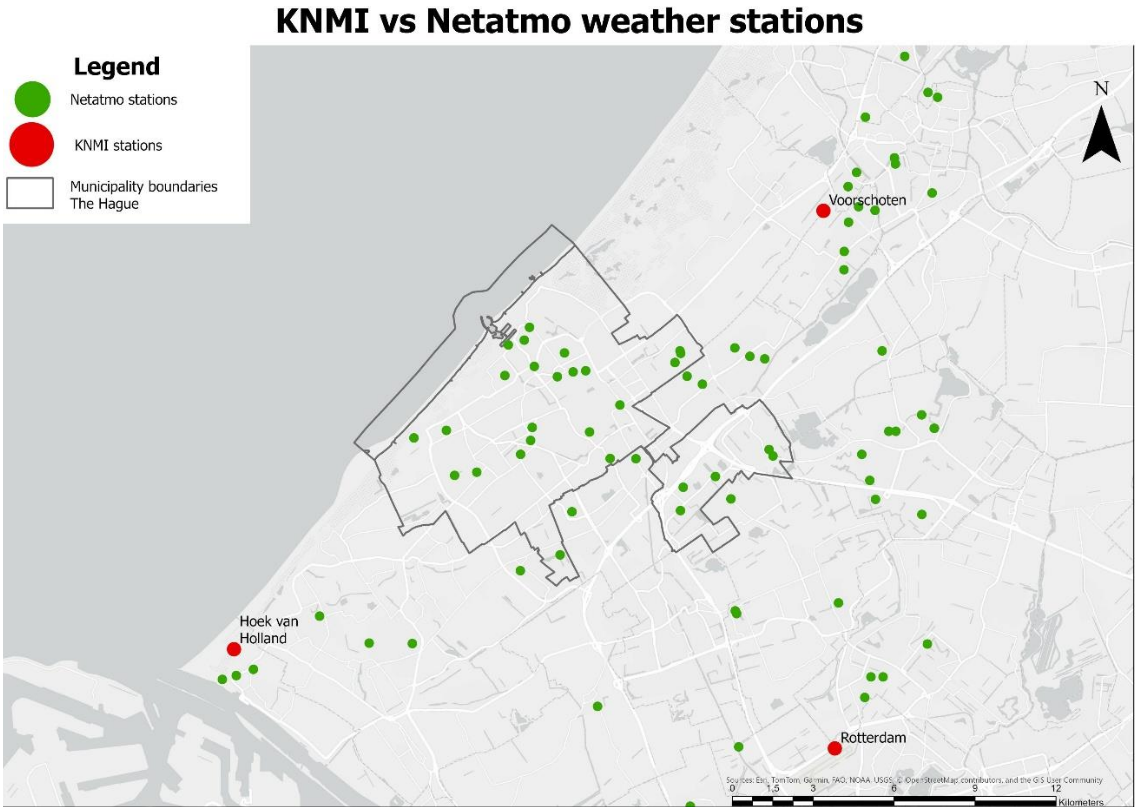

Figure 4.5: Distribution of KNMI stations and Netatmo stations in the western part of South Holland Source: KNMI (2021);Netatmo via Ohori (2022).

Netatmo data

The use of crowdsourced data in atmospheric research is relatively limited compared to other scientific studies, where citizen science is utilized more often (Muller et al., 2015). A dataset such as the Netatmo dataset is classified as citizen sensing, where amateur weather stations provide daily observations.

While official meteorological stations, like the KNMI, are primarily designed to offer data with high quality standards and quality control as the network. To ensure the accuracy and quality of the observations, stations are usually positioned in open areas outside urban areas. Moreover, the costs of setting up and maintaining a station is high, making official networks not suitable for high-density intra urban research (Meier et al., 2017).

Although crowdsourced observation networks lack the precision and standard quality assurance of professional networks, the use of this data is a strategic trade off. Sensor accuracy is sacrificed for dense coverage, making it useful for atmospheric research in urban areas. Chapman et al. (2016) highlight the use of Netatmo data for researching Urban Heat Island effects, where they use high resolution urban coverage of London and find similar results of heat islands as studies who used remote sensing techniques. By using the Netatmo sensing network in The Hague, this study aims to capture the intra urban temperature and climate variations, providing a more granular and realistic input for the model.

4.4 CBS data

The building characteristic data and socio-economic data are primarily derived from the CBS Kerncijfers Postcodegebieden 2021 dataset. This dataset provides aggregated data at the PostCode 6 (PC6) level. For this study, three variables are obtained directly from this dataset without additional spatial processing. Household size is represented as the average number of persons per household per PC6. The income level is represented by the WOZ-waarde, expressed in thousands of euros. The WOZ-waarde reflects property value rather than household income, but serves as a proxy. The CBS does not publish the income level of PC6 areas with less than 100 observations due to privacy regulations.

The WOZ-waarde is used as a proxy for income, but do not always correspond. Household with high property value do not always have high incomes. Household with younger residents in gentrified areas may also have higher WOZ but moderate incomes. The same mismatch may happen with long-term residents in appreciated neighbourhoods where the property value has been rising over the last few years. It is acknowledged as a limitation of using property value instead of income level.

4.5 BAG

Several building characteristics are derived from the BAG verblijfsobjecten dataset presented by the Publieke Dienstverlening Op de Kaart (PDOK). The building age will serve as a proxy for the insulation standards and is represented as Oorspronkelijkbouwjaar in the BAG database. This study differentiates between the share of buildings that were constructed before 1975 and the share that was constructed after. The threshold of 1975 corresponds with the time of no insulation standards (Brounen et al., 2012). The insulation standards have improved over time, and by including the building year this study tries to capture the structural differences more explicitly. Figure 4.6 shows the share of pre-1975 buildings across the Hague. The city mainly consist of older neighbourhoods and that is visible from this map. Some of the newer neighbourhoods are visible on this map as well, and are concentrated on the outskirts of the municipality. The energy label distribution is represented by the share of high efficiency labels (A and B) to further control for insulation standards that building age alone can not account for. This data is retrieved from Den Haag in Cijfers and RVO (2021).

Percentage of Residential Buildings Constructed before 1975 across The Hague

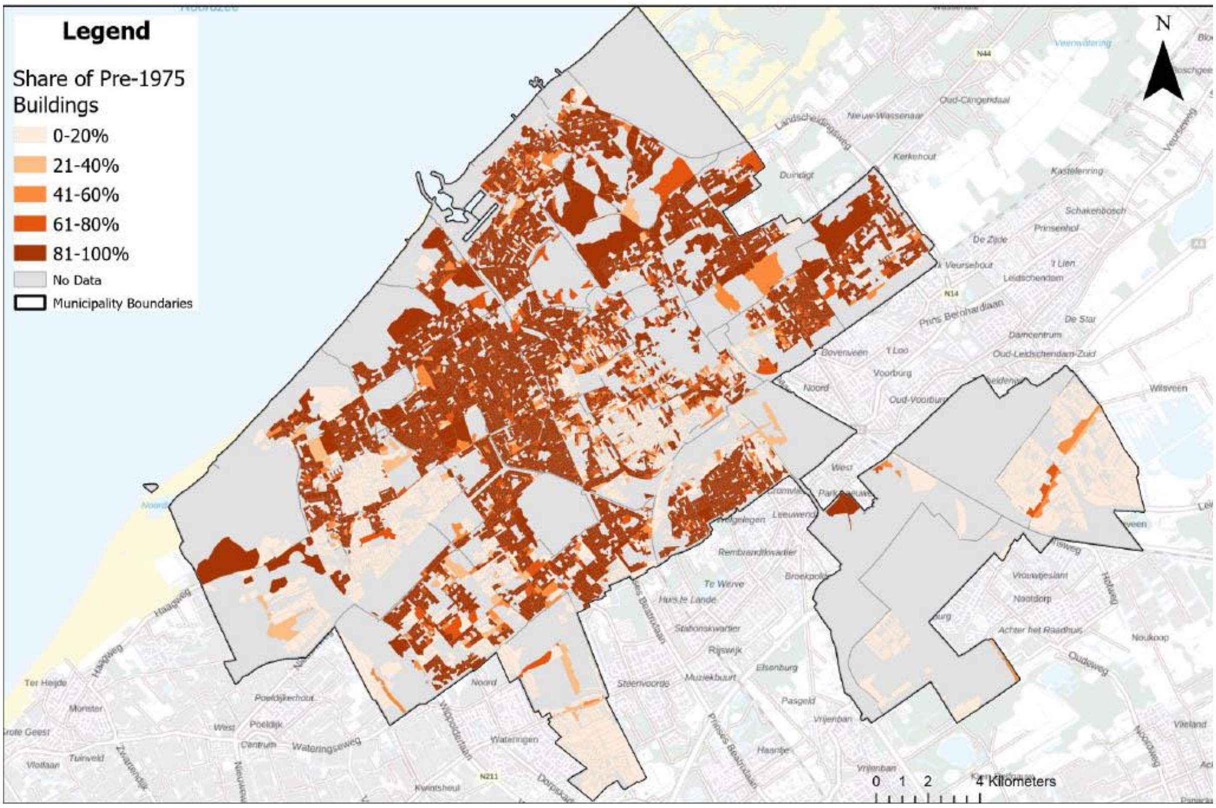

Figure 4.6: Share of Buildings constructed before 1975. Source: Kadaster (2021)

Dwelling type is represented by the share of multi-family homes per PC6. Apartments within multi-family homes benefit from shared walls and reduced heat loss compared to detached dwellings (Voskamp, 2021). For each PC6 area a share of multi-family homes was calculated. This was done by assigning a score of 0 or 1 to each building in the postcode. A score of 1 was given if more than 1 'verblijfsobject' was located within the building polygon. This was then transformed to a percentage for the whole postcode, from 0% to 100% multi-family homes.

The remaining building attributes and UHI variables are derived from the BAG Verblijfsobjecten and 3DBAG. The PDOK collects data at the individual dwelling level. A selection was made in the dataset where only residential units (gebruiksdoel = woonfunctie) in The Hague constructed before or in 2021 are selected. The floor area is aggregated to PC6 level by calculating the mean floor area per postcode, returning the average dwelling size in m2 per PC6.

Building height and urban density are both derived from the 3DBAG dataset. The 3DBAG integrates the footprint polygons with height information derived from the Algemeen Hoogte Bestand (AHN). The AHN provides building heights at the individual dwelling level. The dataset was downloaded using QGIS via PDOK and clipped to the The Hague area. Building height per PC6 is calculated using the mean of h_70 attribute across all buildings within each postcode. h_70 represents the 70^{th} percentile of roof height and is preferred over absolute maximum height as it is more robust to outliers (Roy et al., 2022). Urban density is operationalized as the building coverage ratio and is calculated by summing the footprint area of all buildings within each PC6 polygon area, then dividing this by the total area of the polygon. This will give the percentage of building area. This will capture the compactness of the built environment and will serve as the morphology proxy. Both variables are computed using a spatial join in ArcGIS Pro, where the individual building level attributes will be aggregated to the PC6 areas.

Average NDVI score per postcode in The Hague

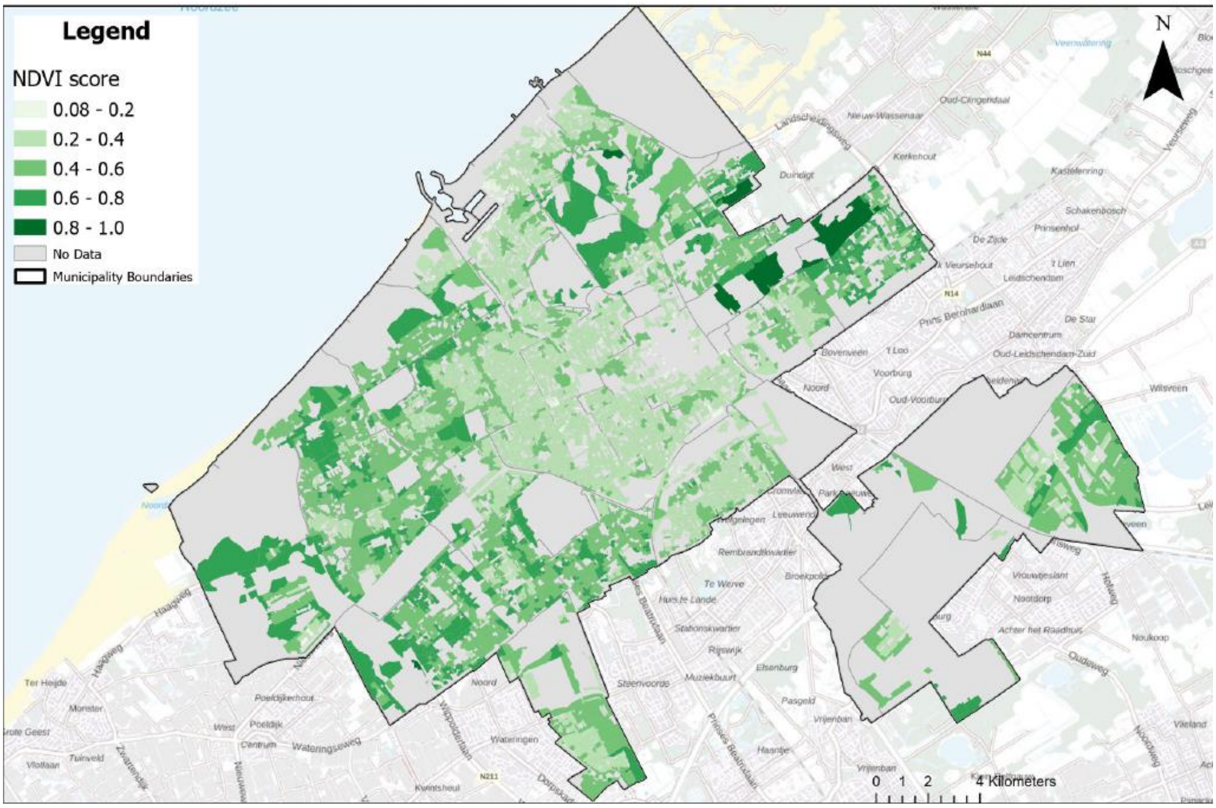

Figure 4.7: NDVI scores across postcode 6 areas in The Hague. Source: RIVM/DHIG (2022)

Figure 4.7 shows the average NDVI score per PC6 area and represents the intensity of vegetation cover across The Hague. Highest NDVI values are found in the northern part near Haagse Hout, and towards the southern postcodes near dune area or recreational area. The dense city centre is visible, with low amount of vegetation. It is expected that the areas with more green require more heating demand, as the temperatures in or near these areas would be lower compared to other neighbourhoods.

The spatial distribution of the two other variables that measure the UHI are visible in figure 4.8 and figure 4.9. Figure 4.8 shows the concentration of taller buildings within each postcode. Some clusters with taller buildings in their postcodes are visible, mainly in the city centre due to prominent skyscrapers that bump up the average for the postcodes. Interestingly, there is a pattern visible moving up from the city centre towards Scheveningen with a large cluster of postcodes with taller buildings. Some postcodes in the south also contain taller buildings, but these buildings are tall apartment blocks together with lower suburban building stock. Therefore, it is expected that the UHI has less effect in these suburban neighbourhoods.

Average building height per PC6

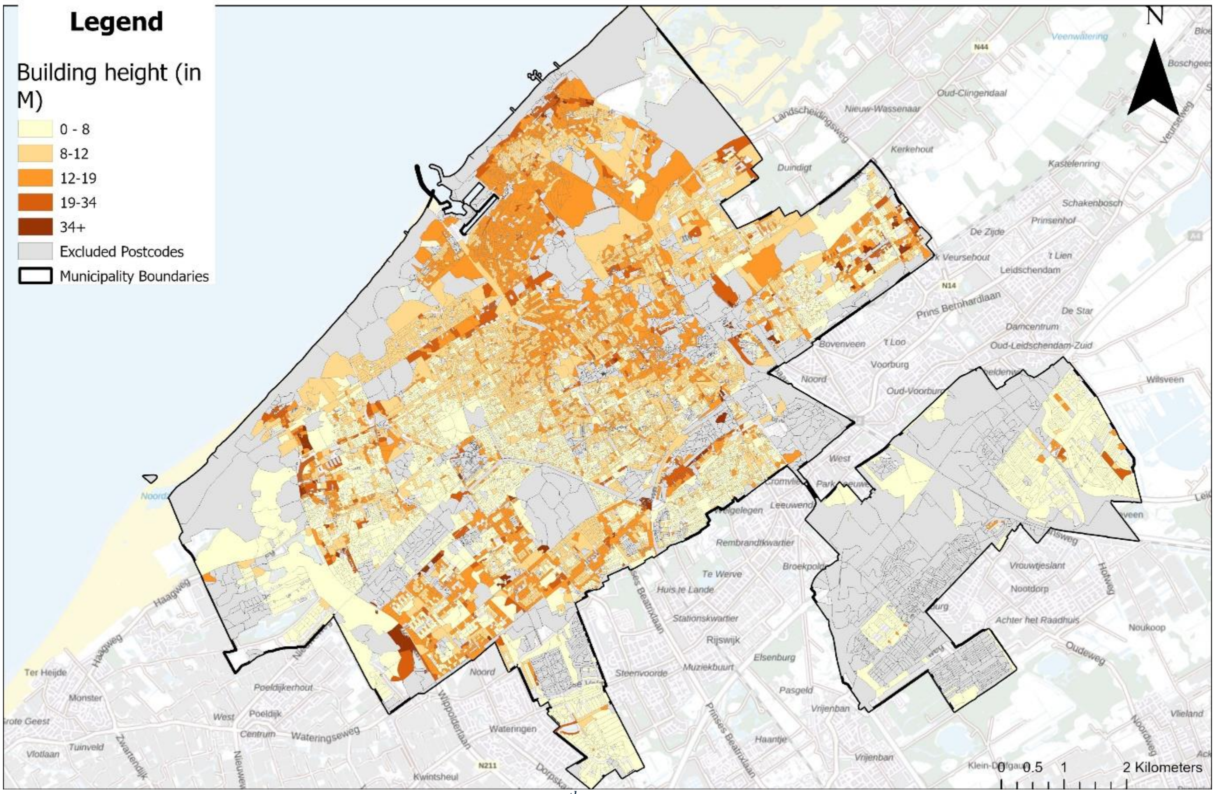

Figure 4.8: Average building height per postcode area (70th percentile). Source: 3D geoinformation group & 3DGI (2025); PDOK (2021).

Figure 4.9 shows the coverage rate and similar patterns are visible in the city centre and the Scheveningen area. The suburban neighbourhoods on the outskirts of the city with tall buildings show low coverage rates, confirming the presence of the tall apartment buildings.

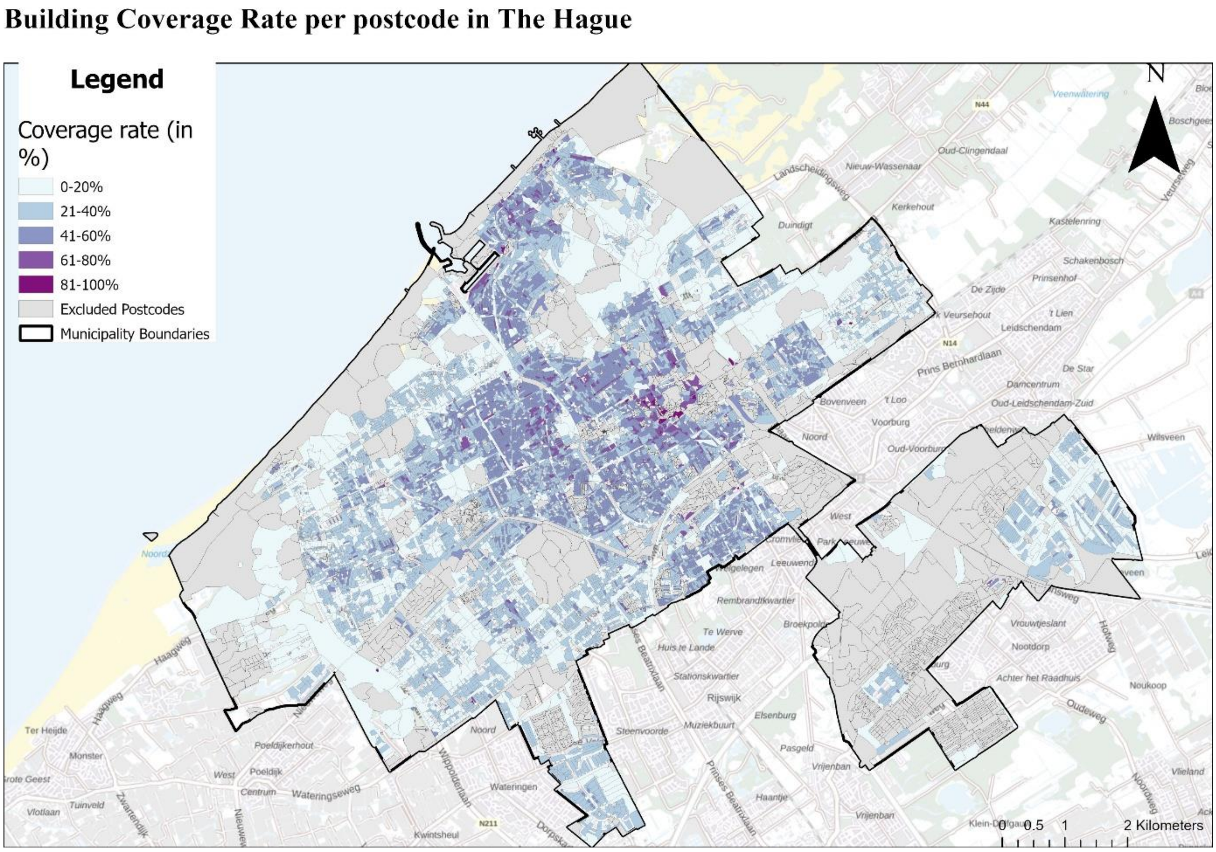

Figure 4.9: Coverage rate per postcode area. Source: Kadaster (2021); PDOK (2021).

Table 4.1: Mean, Median and Standard Deviation of variables for: N = 9,410 for Model 1 variables; N = 9,389 for urban morphology variables (Model 3); N = 8,955 for energy label variable (Model 4). HDD statistics based on N = 9,410. Statistics based on The Hague after filtering of unusable postcodes.

Variable	Description	Mean	Median	St. Dev.
Gas consumption	Annual gas use in m³	1285,17	1200,00	454,96
Building age	% built before 1975	74,24	100	39,66
Dwelling size	Avg. floor area in m2	97,80	83,68	50,29
Dwelling type	% multi-family buildings	51,91	56,95	35,88
Household size	Avg. person per household	2,06	2,00	0,57
WOZ-waarde	Log of property value x € 1.000	5,56	5,46	0,53
NDVI	Mean vegetation Index (0 to 1)	0,41	0,39	0,13
Urban Density	% building coverage area	35,65	36,54	14,95
Building height	Average height in meters (70th percentile)	10,83	9,79	6,77
HDD Netatmo (Kriging)	Annual HDD, Kriging interpolation	2339,09	2337,94	47,59
HDD KNMI (Kriging)	Annual HDD, Kriging interpolation	2679,27	2677,79	17,60
HDD Netatmo (IDW)	Annual HDD, IDW interpolation	2337,80	2336,91	112,14
HDD KNMI (IDW)	Annual HDD, IDW interpolation	2745,04	2746,57	20,01
Energy Label	% high performance (A and B)	39,65	33,33	38,21

.

5 Results

5.1 Spatial Patterns of Local Temperature

The Netatmo network was validated prior to analysis following the quality control scheme of Meier et al. (2017) and all stations with incomplete annual measurements or unrealistic temperatures were filtered out as described in table 3.2. This resulted in 105 stations remaining in the total network, with 31 valid stations in the The Hague area Netatmo network. The spatial distribution of annual Heating Degree Days derived from these stations is shown in figure 5.1.

Valid Netatmo station network in the The Hague area (2021)

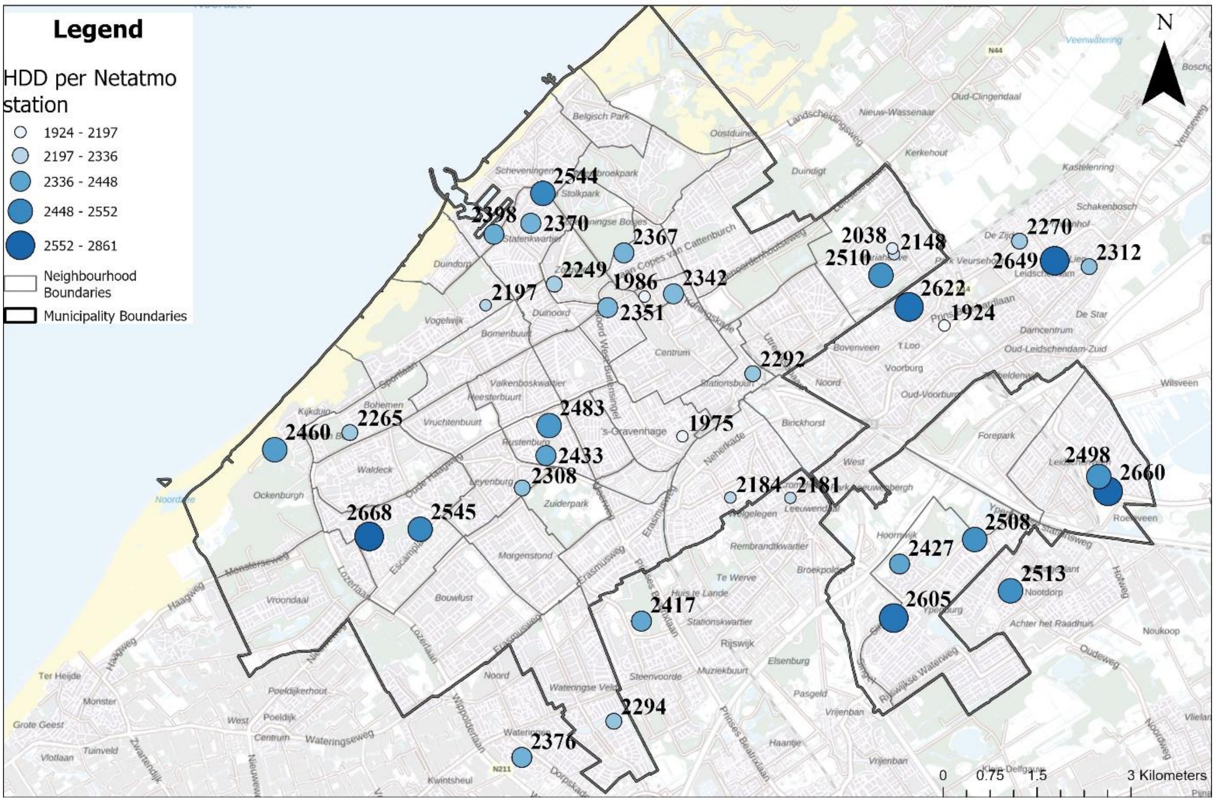

Figure 5.1: Heating Degree Days per station in The Hague. Source: Netatmo via Ohori (2022) (Base temperature = 18°C)

The spatial distribution of the Heating Degree Days per Netatmo station shows clearly that there is some intra-urban variation in The Hague. The inner city stations show a lower amount of Degree Days, indicating some UHI effects where the heat is retained and therefore reducing the heating demand. The outskirts or coastal stations have higher HDD values, indicating that temperatures are generally lower and that more heating is required in these areas. This spatial pattern confirms that microclimatic differences are measurable at the fine intra-urban scale using the crowdsourced observations and that regional climate data would fail to capture this variation.

Following quality control, annual HDD values were spatially interpolated to the centroid of each PC6 area using Ordinary Kriging and Inverse Distance Weighting. The kriging interpolation applied a spherical semi variogram model with anisotropy incorporated to reflect the North sea coastline. Cross-validation produced a Root Mean Square Standardized value of 0,94, indicating well-calibrated uncertainty estimates. The resulting interpolated surface is shown in figure 5.2, along with the IDW interpolation of the Netatmo HDD.

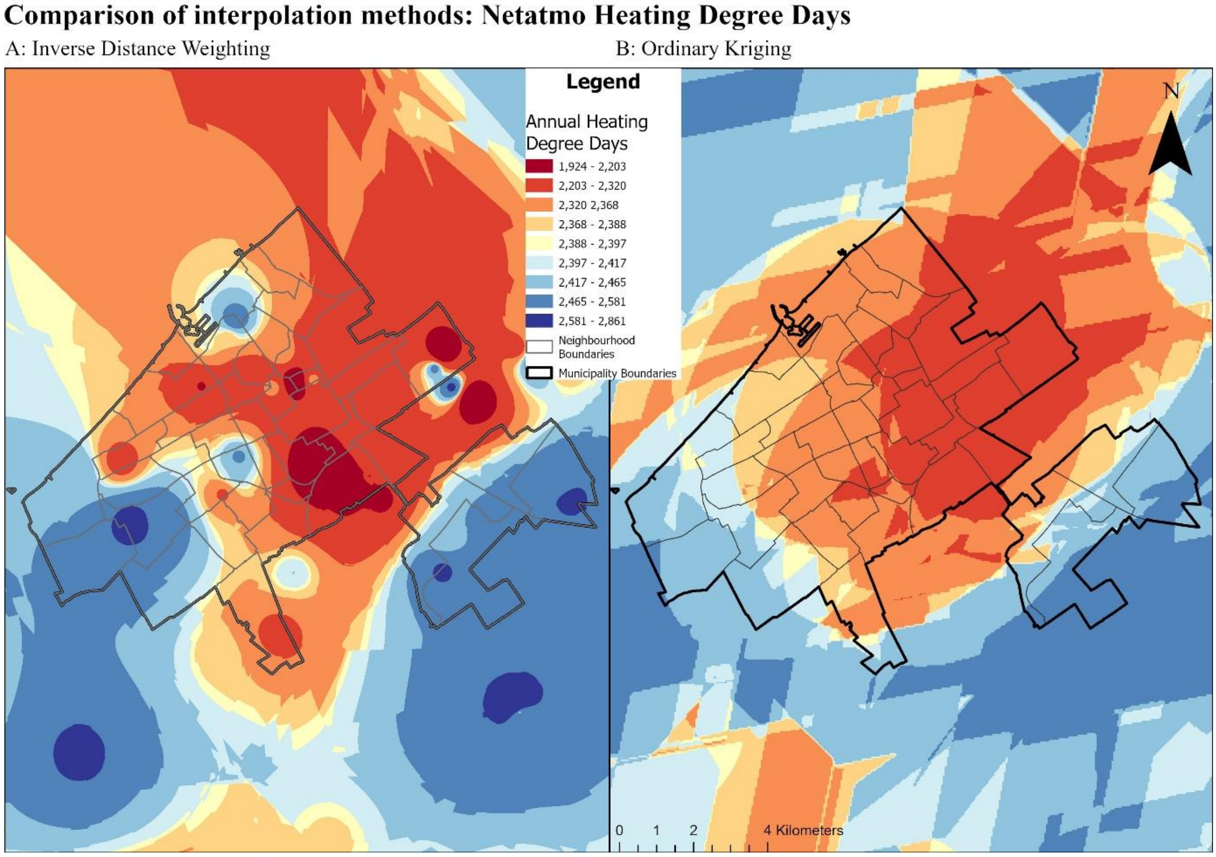

Figure 5.2: Comparison of Netatmo HDD interpolation surfaces, The Hague. Source: Netatmo via Ohori (2022); PDOK (2021).

Both the Inverse Distance Weighting (A) and the Ordinary Kriging surface (B) capture the Urban Heat Island effect in The Hague to some degree. The inner city shows lower HDD values in both maps while the coastal area and the suburbs of the city show higher values. The main distinction is that the IDW surface produces a more localized fragmented surface compared to the Kriging surface. This results in a clear visual difference, where the IDW surface produces 'bullseye' like patterns around each Netatmo station. In practice, this means that neighbouring postcodes can receive different HDD values depending on which station falls closer to their centre, rather than resembling real local temperature differences.

A useful test of interpolation quality would be to check if large green areas are visible and that these green zones consistently have lower HDD values. However, no active Netatmo stations are located within these parks and green areas, with the exception of a single station near the Uithof area. As a result, the HDD values assigned to postcodes covering these green areas are interpolated from surrounding urban stations rather than measured locally. This means the cooling effect of these green spaces within the urban area is likely underrepresented in both the IDW and Kriging surfaces.

The Kriging surface shows a smoother temperature gradient and provides a more realistic representation of The Hague, where outliers do not have much impact on the total surface. The cooling effect of the coastline is more visible in the Kriging surface as it is accounted for by the model. Both surfaces will be used as inputs for the regression models to compare model performance and see what method fits the best for The Hague.

5.2 Comparison: KNMI and Netatmo

Figure 5.3 presents the direct comparison of the interpolated Heating Degree Day surfaces derived from KNMI and Netatmo using Ordinary Kriging. The KNMI interpolation (A) is nearly a uniform surface across The Hague, with the HDD values ranging from 2631 to 2730. The total range of HDD for The Hague is only 99 HDD. The KNMI network covers the entire Netherlands, and calculating and interpolating HDD from only 3 points results in this uniform interpolation. No intra-urban differences are visible in temperature, with only a slight pattern from the sea being visible.

The Netatmo interpolation (B) shows a different picture. HDD values range from 2246 to 2541, a range of 295 Degree Days. The outer areas of The Hague show higher HDD values up to 2540 degrees, while the inner centre of the city shows values as low as 2246 degree days., consistent with the UHI and the findings of Van der Hoeven & Wandl (2018). The larger variation in the Netatmo HDD compared to the KNMI HDD, is the extra explanatory power that the KNMI HDD do not have. Clearly some intra-urban differences are visible, along with the effect of the coastline and the effect of the less urban dense areas near the outskirts of the city. The descriptive statistics confirm the contrast between Netatmo and KNMI: The standard deviation across 9410 postcodes for the KNMI HDD is only 17.6, while the Netatmo network standard deviation is 47.4, nearly three times larger.

Spatial variation in Heating Degree Days in The Hague (2021)

A: KNMI Kriging Interpolation B: Netatmo Kriging Interpolation

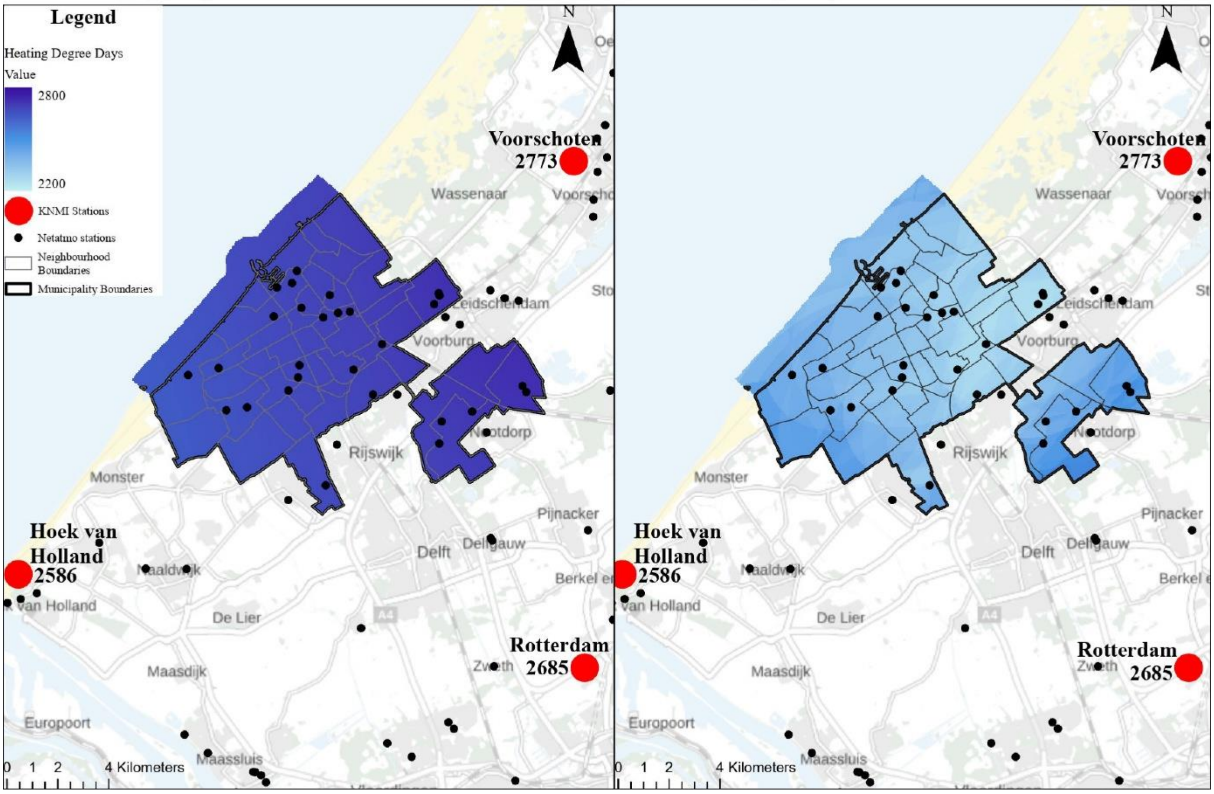

Figure 5.3: Comparison of KNMI and Netatmo HDD surfaces. Source: Netatmo via Ohori (2022); KNMI (2021); PDOK (2021).

It is important to note that the two networks also differ in absolute HDD level: the Netatmo mean of 2,339 is approximately 340 degree days lower than the KNMI mean of 2,679, with the whole range of Netatmo data also being lower than the lowest value of KNMI HDD. This difference reflects the placement of the weather stations. The Netatmo stations are placed in the city itself while the KNMI stations are located in rural open areas. For the purposes of this research, the absolute difference is less relevant than the spatial variation that each network captures. The variation is what determines the added value to the model.

The choice of interpolation method does not have an effect on the KNMI network. Only 3 valid stations affect The Hague and therefore this sparse input for both methods leads to an near-identical HDD surface. For the Netatmo network, Kriging produces a smoother surface while IDW preserves more of the local variation in HDD. The standard deviation of the IDW surface is 112.1, more then double than that of the Kriging surface (47.4). Further implications for the model are visible in section 5.5.

5.3 Ordinary Least Squares Regression results

Table 5.1: OLS coefficients (significant codes: . p < 0.10, * p < 0.05, ** p < 0.01, *** p < 0.001. Variables without a symbol are not statistically significant at the p < 0.10 level)

Variable	OLS M1	OLS M2 Kriging	OLS M2 IDW	OLS M3 Kriging	OLS M3 IDW	OLS M4 Labels
Intercept	-1731.95***	-77.66	-1012.71**	33.01	-1229.34***	-1601.79***
Building age (% pre-1975)	4.027***	4.004***	4.028***	3.937***	4.029***	3.333***
Dwelling size (m^2)	2.356***	2.351***	2.348***	2.298***	2.290***	2.415***
Multi-family (%)	-0.169.	-0.160.	-0.163.	-0.370***	-0.336***	0.006
Household size	94.99***	94.14***	94.08***	95.51***	89.47***	116.65***
Log WOZ-waarde	414.15***	417.83***	417.39***	412.51***	413.74***	400.85***
HDD KNMI Kriging	—	-0.624***	—	—	—	—
HDD KNMI IDW	—	—	-0.268.	—	—	—
HDD Netatmo Kriging	—	—	—	-0.762***	—	—
HDD Netatmo IDW	—	—	—	—	-0.218***	—
Building height (m)	—	—	—	-2.200***	-2.045***	—
Urban Density (%)	—	—	—	0.012	0.015	—
Mean NDVI	—	—	—	175.65***	139.72***	—
Energy label (% A/B)	—	—	—	—	—	-1.654***
R^2	0.673	0.673	0.673	0.680	0.677	0.691
AIC	131,390	131,376	131,388	130,887	130,969	124,500
N	9,410	9,410	9,410	9,388	9,388	8,954

Baseline

The baseline OLS model (Model 1) explains 67.3% of the variance in residential gas consumption across 9410 PC6 areas ($R^2 = 0.673$). All five control variables are statistically significant. Building age is the strongest predictor, with each percentage point increase in pre-1975 dwellings associated with an increase of approximately 4.04 m3 in annual gas consumption ($\beta = 4.035$, $p < 0.001$), consistent with the findings of Brounen et al. (2012). Dwelling size is the second strongest predictor, with each additional square metre of average floor area associated with 2.36 m^3 additional annual gas consumption ($\beta = 2.356$, $p < 0.001$). Log WOZ-waarde is also strongly significant ($\beta = 412$, $p < 0.001$), confirming the positive relationship between property value and energy consumption documented by Vringer (2005). Household size is positively associated with gas consumption ($\beta = 94.7$, $p < 0.001$), consistent with Mashhoodi et al. (2020). The share of multi-family dwellings is negatively associated with consumption, though the effect is marginally significant ($\beta = -0.169$, p =

0.064), but still confirming the heat retention effect of shared walls in apartment buildings (Voskamp et al., 2021).

Standardised coefficients confirm the relative importance of each variable. Log WOZ-waarde has the largest standardised effect (β = 0.481), followed by building age (β = 0.351) and dwelling size (β = 0.261). Household size has a moderate effect (β = 0.119), while multi-family share has a small and marginally significant effect (β = -0.014). This ranking confirms that socio-economic and structural building characteristics together dominate the spatial pattern of gas consumption, with WOZ-waarde and building age jointly accounting for the largest share of explained variance.

Model 2
Adding the KNMI Kriging HDD to the baseline model (Model 2 Kriging) produces a significant coefficient (β = -0.621, $p < 0.001$), indicating that postcodes with higher regional HDD values consume more gas as expected. The improvement in model fit is negligible however, with the R^2 increasing from 0.6727 to 0.6733 and AIC decreases by 14 points. The KNMI IDW Model performs similarly but a bit weaker with no AIC improvement (β = -0.263, $p = 0.057$). This confirms that the uniform KNMI HDD interpolation adds no explanatory power of intra-urban variation.

Model 3
Model 3, incorporating Netatmo Kriging HDD along with urban morphology variables, produces the strongest OLS performance ($R^2 = 0.680$). The Netatmo HDD coefficient is strongly significant (β = -0.775, $p < 0.001$) with a larger effect than KNMI. For the urban morphology variables, building height is negatively and significantly associated with gas consumption (β = -2.11, $p < 0.001$). This is consistent with the heat retention effect of taller buildings. Mean NDVI is positively associated with consumption (β = 165, $p < 0.001$). Reflecting that postcodes with a lot of green areas tend to have a higher heating demand. Building coverage is not significant in the model ($p = 0.474$), indicating that the raw building coverage adds no extra explanatory power. The Netatmo IDW model performs similarly but has a weaker R^2 (0.677).

Ordinary Kriging produces the best model fit for both the KNMI and Netatmo specifications and is therefore used as the primary interpolation method in the spatial models. The IDW results are retained in the analysis throughout to allow comparison between interpolation methods, and conclusions about both methods are drawn in section 5.5 and the discussion.

Model 4
Model 4, the sensitivity model including energy label share on a reduced sample of 8,955 postcodes, produces the highest OLS R^2 of all models ($R^2 = 0.691$). The energy label coefficient is strongly significant (β = -1.655, $p < 0.001$), indicating that each percentage point increase in high-efficiency labels is associated with 1.66 m^3 reduction in annual gas consumption. Interestingly, the building age coefficient decreases from 4.04 in Model 1 to 3.34 when energy labels are included, confirming that energy labels capture insulation improvements that building age alone cannot account for, as argued by Majcen et al. (2013). The actual insulation standards of a dwelling matters more than its construction year.

5.4 Spatial Autocorrelation and model selection

All OLS models showed significant positive spatial autocorrelation in their residuals with Moran's I values ranging from 0,29 to 0,32, all significant ($p<0.001$). This indicated that neighbouring postcodes share similar unexplained variation in gas consumption not explained by the variables that are included in the models. This violates the OLS assumption of independent errors, and therefore Spatial Error Models are performed.

The Lagrange Multiplier tests following Anselin (1988) were applied to select the appropriate spatial model. Both the LM error and the LM lag are significant across all models. However, the Spatial Error

Model is selected as the error statistics exceed the lag statistics (Model 1: adjRSerr = 1010, adjRSlag = 41; Model 3: adjRSerr = 916, adjRSlag = 49). The bigger error number indicates that residual spatial dependence comes from unobserved structural factors, such as neighbourhood typology. Variance Inflation Factors (VIF) were below 3 for all variables across all models, indicating that there is no problematic multicollinearity between variables.

The Lambda parameter of 0.55–0.56 across all models, with z-values around 50 and $p < 2.2e-16$, is significant. The AIC improvement from OLS to SEM ranges from approximately 1,900 to 2,020 points across specifications. The High Lambda values indicate that around 55% of the unexplained variation in gas consumption is spatially structured, driven by those unobserved values.

Table 5.2: SEM coefficients (. p < 0.10, * p < 0.05, ** p < 0.01, *** p < 0.001. Variables without a symbol are not statistically significant at the p < 0.10 level).

Variable	SEM M1	SEM M2 Kriging	SEM M2 IDW	SEM M3 Kriging	SEM M3 IDW	SEM M4 Labels
Intercept	-2029.29***	-424.36	-1053.00	-405.63	-1518.70***	-1947.51***
Building age (% pre-1975)	3.954***	3.946***	3.959***	3.912***	3.964***	3.457***
Dwelling size (m²)	1.626***	1.625***	1.622***	1.610***	1.600***	1.699***
Multi-family (%)	-0.394***	-0.390***	-0.390***	-0.470***	-0.461***	-0.301**
Household size	59.18***	58.87***	58.67***	58.86***	57.20***	70.73***
Log WOZ-waarde	497.02***	498.76***	499.39***	489.92***	491.23***	491.43***
HDD KNMI Kriging	—	-0.602*	—	—	—	—
HDD KNMI IDW	—	—	-0.360	—	—	—
HDD Netatmo Kriging	—	—	—	-0.701***	—	—
HDD Netatmo IDW	—	—	—	—	-0.225***	—
Building height (m)	—	—	—	-2.558***	-2.542***	—
Urban Density (%)	—	—	—	0.007	0.007	—
Mean NDVI	—	—	—	222.39***	203.67***	—
Energy label (% high efficiency)	—	—	—	—	—	-1.275***
Lambda (λ)	0.562***	0.561***	0.562***	0.551***	0.555***	0.542***
Pseudo R²	0.754	0.754	0.754	0.755	0.755	0.758
AIC	129,366	129,364	129,366	128,994	129,010	122,845
N	9,410	9,410	9,410	9,388	9,388	8,954

5.5 Model comparison

Adding in the KNMI HDD, regardless of interpolation method, adds no additional value. The AIC difference for Model 2 with IDW and Kriging is only 2 points, and the Pseudo R^2 are identical. The

Pseudo R^2 is also identical with that of the baseline model, indicating that the model almost stays the same with the addition of KNMI HDD. The KNMI Kriging HDD coefficient is marginally significant ($\beta = -0.602$, $p = 0.047$) while the KNMI IDW coefficient is not significant ($p = 0.179$). This again confirms that the KNMI interpolation adds no practical explanatory power for gas consumption variation.

Model 3 including the Netatmo Kriging interpolation and the urban morphology variables again performs the best (AIC 128994, $R^2 = 0.755$). It outperforms the IDW model by a small amount (ΔAIC = 16). Interestingly to note is that while both models are strongly significant, the coefficients differ a lot. The Netatmo Kriging coefficient ($\beta = -0.701$) is three times as big as the IDW coefficient ($\beta = -0.225$). This means that the interpolation method is not just a technical detail but has a substantial impact on how much the temperature-consumption effect is captured.

The Netatmo HDD coefficient remains strongly significant ($\beta = -0.701$, $p < 0.001$), confirming that locally measured temperature variation contributes independently to explaining spatial gas consumption patterns after accounting for residual spatial autocorrelation. Among the urban morphology variables, building height and mean NDVI remain significant ($\beta = -2.558$, $p < 0.001$, $\beta = 222.4$, $p < 0.001$) while the coverage ratio remains not significant ($p = 0.486$).

Model 4 including energy labels produces the highest model performance, with a Pseudo R^2 of 0.758 and AIC of 122,845, and the lowest RMSE of 223.5 m^3. The energy label coefficient remains strongly significant in the spatial framework ($\beta = -1.275$, $p < 0.001$). However, this model is not directly comparable to model 3 as it does not include the local HDD. The lower AIC also reflects a lower sample size rather than a better model fit. Model 4 is therefore only used as a sensitivity analysis showing that energy labels add independent explanatory value beyond the building age.

Table 5.3: Spatial Error Model Results and Model Comparison

Model	N	OLS R^2	OLS AIC	SEM Pseudo R^2	SEM AIC
M1: Baseline	9,410	0.673	131,390	0.754	129,366
M2: + KNMI Kriging HDD	9,410	0.673	131,376	0.754	129,364
M2: + KNMI IDW HDD	9,410	0.673	131,388	0.754	129,366
M3: + Netatmo Kriging + Morphology	9,388	0.680	130,887	0.755	128,994
M3: + Netatmo IDW + Morphology	9,388	0.677	130,969	0.755	129,010
M4: + Energy Labels	8,954	0.691	124,500	0.758	122,845

Table 5.3 presents an overview of model results. The transition from OLS to Spatial Error Models produces a substantial improvement in fit across all specifications. The AIC reduction from OLS to SEM ranges from approximately 1,600 points (Model 4) to 2,020 points (Model 1), confirming that accounting for spatial autocorrelation improves model quality. The RMSE decreases from 260.3 m^3 in OLS Model 1 to 226.1 m^3 in SEM 1, a reduction of 34.2 m^3 per postcode.

To assess prediction accuracy alongside AIC and pseudo R^2, the Root Mean Square Error (RMSE) was calculated for each model. RMSE measures the average difference between the model's predicted gas consumption and the actual observed values per PC6 area, expressed in the same unit as the dependent variable (m^3). A lower RMSE indicates that the model predictions are on average closer to the actual values. The Root Mean Standard Error (RMSE) of the baseline Spatial Error Model is 226.1 m^3, representing a reduction of 34.2 m^3 compared to the OLS baseline (RMSE = 260.3 m^3). The addition of Netatmo HDD and morphology variables reduces RMSE further to 225.5 m^3 in SEM 3, while the energy label model achieves the lowest prediction error at 223.5 m^3.

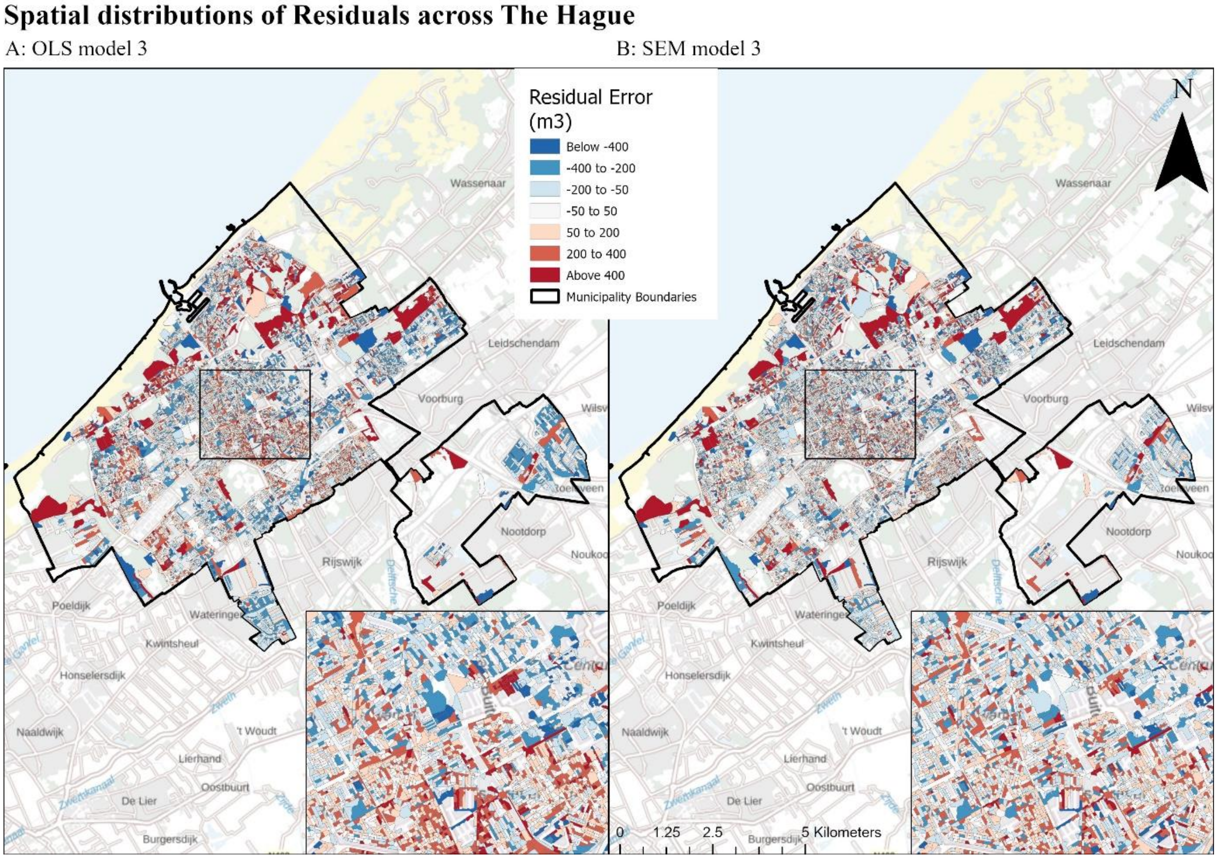

Figure 5.4: Comparison of spatial distribution of residuals between OLS and SEM models. Source: PDOK (2021).

Figure 5.4 presents the differences between the residuals between the OLS (A) and SEM (B) models across the Hague. The corner panel is zoomed in on the southern part of the city centre and the Schilderswijk neighbourhood. A positive residual means that the actual gas consumption is higher than the predicted value of the model, while a negative residual means that the actual gas consumption is lower than predicted.

The overall spatial pattern is similar between the models and is in line with the Lambda parameter that indicated that the autocorrelation was gradually reduced across the whole city rather than only a specific part or neighbourhood. The zoomed in map shows that the values in the Centrum are blue, and the model likely overestimates because property value is high and the buildings are old while in reality apartement sizes are small. De Schilderswijk shows more red postcodes, but those get reduced somewhat after accounting for autocorrelation. This pattern is likely due to poorer insulation standards than other pre-1975 buildings. Some red clusters are visible on the western side of both panels and those appear to be the wealthier villa areas like the Vogelwijk. Leidschenveen shows consistent overestimation of the model, likely that these new housing developments have other sources of heating, like a heat pump.

Spatial distribution of observed and predicted residential gas consumption across The Hague

A: Observed values B: Predicted values

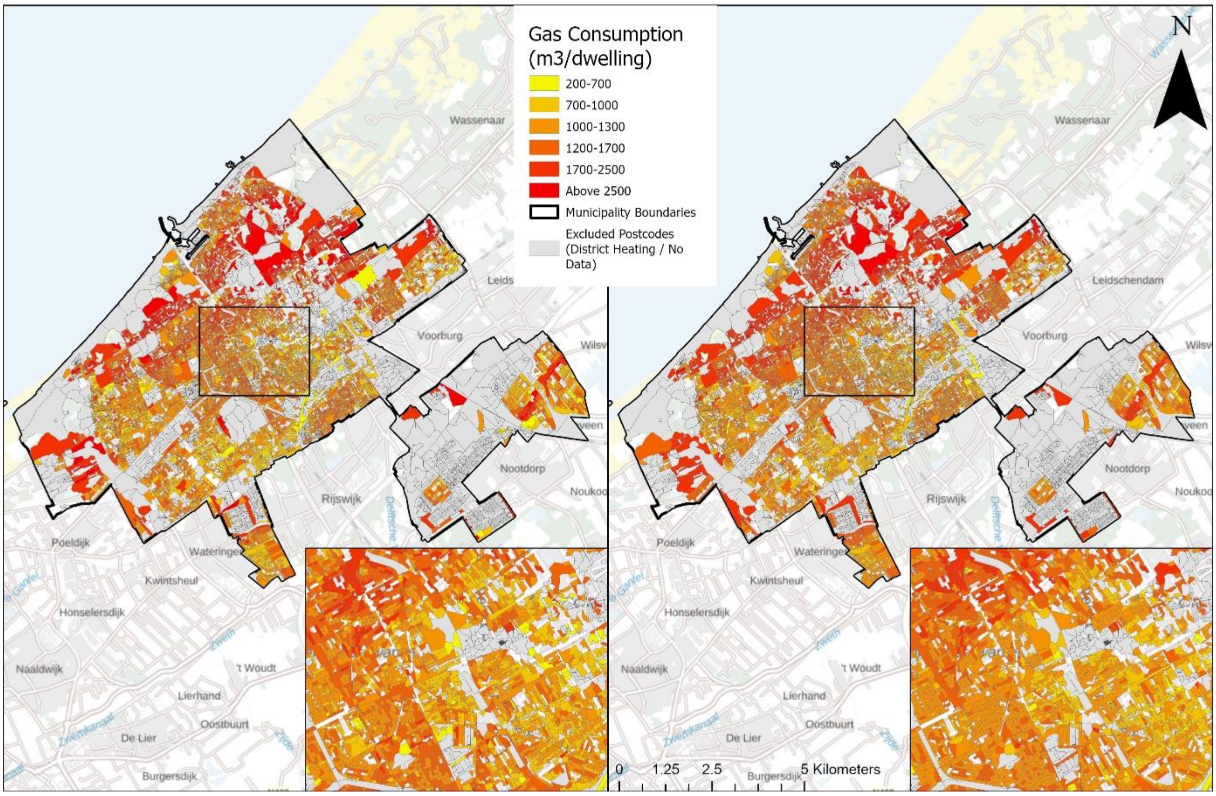

Figure 5.5: Actual gas consumption and predicted gas consumption. Source: CBS (2021); PDOK (2021).

Figure 5.5 presents the actual gas consumption alongside the fitted values from SEM model 3 Kriging. Both maps seem identical and this confirms the high explanatory power of the model ($R^2 = 0.755$). The differences that remain are visible in the residuals map in panel B of figure 5.4. These differences are not visible in this map due to the smoothing out of the classification.

6 Discussion

6.1 Main findings

This research investigated whether locally measured temperature data from a citizen sensing network improves the explanation of spatial patterns in residential gas consumption in The Hague, compared to regional KNMI data. The results show that the local temperature does improve the model performance, but only by a small magnitude. Most of the explanatory power comes from the building characteristics and socio-economic factors. Figure 6.1 visualize the standard coefficients and the domination of the building characteristics in explaining gas consumption is visible.

The baseline Spatial Error Model, including building age, dwelling area, dwelling type, property value and household size, already explains 75.4% of variance in residential gas consumption at the PC6 level across 9410 PC6 areas. This number is higher than comparable Dutch studies. Guerra Santin et al. (2009) found that building characteristics explain approximately 42% of variation in Dutch residential energy use, while Van den Brom et al. (2019) found that building characteristics and occupant factors together explain approximately 50%. The higher explanatory power of the baseline model in this study is likely due to several methodological differences. First, this study operates at the PC6 level, which aggregates individual household variation into postcode averages. Aggregation typically reduces unexplained variance because individual behaviour cancels out at the postcode scale, producing more predictable average patterns. Guerra Santin et al. (2009) worked at the individual household level for example, where occupant behaviour contributed to more noise in the model (4-15% variance).

Compared to the earlier work of Pijnenburg (2021), who also investigated gas consumption and local climate at the PC6 level in The Hague, this study makes use of several methodological upgrades. Pijnenburg compared average gas consumption and degree days between 2018 and 2019 and found the climatic difference between the two years too small to explain meaningful variation in gas consumption after degree day correction. His approach did not make use of multiple spatial regression modes or interpolation techniques. Another conclusion that was made in his thesis is that without local wind and rain data it was difficult to isolate the local climate effect. This study addresses this limitation by operationalising the Urban heat Island effect using urban morphological variables, and directly incorporating this in regression models. At the same time, the small improvement of the explained variance when adding the local climatic variables is consistent with the conclusion of Pijnenburg, who argues that local climate beyond temperature only plays a small role. Both studies therefore point to the same fundamental conclusion: the thermal quality of the building stock, not local climate variation, is the primary driver of gas consumption patterns in The Hague.

Building age was the strongest predictor for residential gas consumption across all model specifications, with each percentage point increase in pre-1975 dwellings per postcode was associated with around 3.95 – 4.04 m³ additional annual gas consumption. This finding was consistent with Brounen et al. (2012), who found that pre - 1980 dwellings consumed 50% more gas than newer buildings. The sensitivity model that included energy labels, decreased the building age coefficient from 3.95 to 3.46. This confirms that the addition of energy labels capture some insulation improvements to the buildings that were constructed before 1975. This also confirms the argument of Majcen et al. (2013), who states that energy labels provide additional value that building age alone does not give.

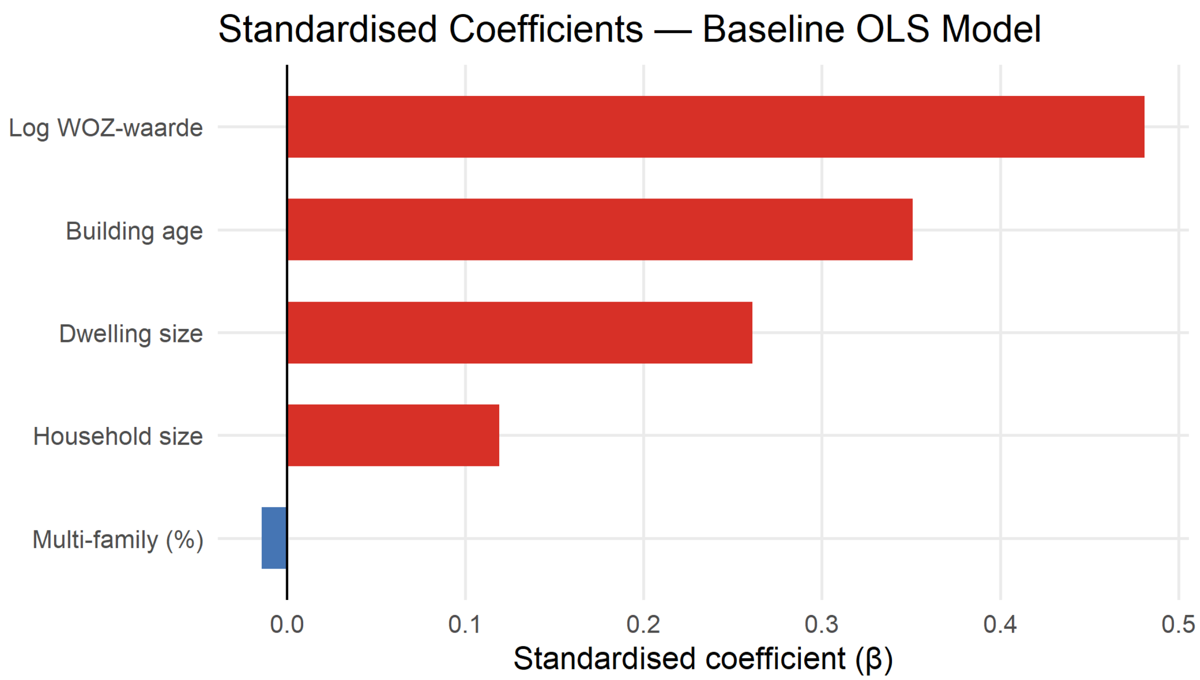

Figure 6.1: Standardised regression coefficients from the baseline OLS model

6.2 KNMI derived Heating Degree Days

The HDD derived from the KNMI network did not add any explanatory power to the models. This finding is not a finding about the irrelevance of temperature for gas consumption, but rather about the unsuitability of the KNMI network for intra-urban analysis. The KNMI network is designed as such that the weather stations capture the least urban influence on their climate measurements. Applying this to a postcode level, it means that all off the 9410 postcodes in The Hague receive a HDD value derived from either Hoek van Holland, Voorschoten or Rotterdam. None of these station locations are measuring the temperature conditions that the postcodes actually experience.

The result is not only a limitation of spatial resolution. It is a measurement error problem. The standard error on the KNMI Kriging coefficient is nearly three times larger than the Netatmo coefficient, and the KNMI IDW coefficient is not statistically significant at all. This means that the KNMI is not necessarily measuring the wrong thing, it measures the right thing but at the wrong location for this research purpose. A more practical consequence of this is that Dutch energy models that use KNMI-derived HDD as the control for temperature may be underestimating the temperature sensitivity. The coefficients in those models do not reflect the relationship between temperature and gas but rather the relationship between gas consumption and a temperature measurement taken in an open rural environment, rather than within the urban area itself.

6.3 Locally measured Heating Degree Days

The Netatmo network captures the effects that the KNMI fail to capture, because the locations of the stations are embedded within the urban area that is meant to be measured. The HDD surface based on the Netatmo network creates a similar spatial pattern that is visible in the study on the UHI I The Hague of Van der Hoeven & Wandl (2018), but based on a ground level citizen sensing network rather than satellite based observations. The temperature surface captures the temperature gradient for The Hague and ranges from 2246 to 2541 HDD across the distance of a few kilometres. This surface is able to capture the intra-urban temperature variation and therefore usable for postcode level energy analysis. The KNMI network, by design, can not reproduce a similar surface for analysis on the same level.

The regression results confirm this added value. The Netatmo Kriging HDD coefficient is strongly significant in the Spatial Error Model 3 ($\beta = -0.701$, $p < 0.001$), with a standard error of 0.111 compared to 0.303 for KNMI Kriging. The AIC improvement of 372 points over the baseline model, compared to only 2 points for KNMI Kriging, confirms that the locally measured temperature signal contributes independently to explaining spatial gas consumption patterns after all building, socio-economic and spatial autocorrelation effects are accounted for. Each degree day of additional heating demand is associated with around 0.70 m^3 additional annual gas consumption per postcode. This relationship could not be found by only the KNMI derived HDD.

The magnitude of this improvement should however be interpreted with care. The baseline model without climate already explains round 75.4% of variance. The model including Netatmo explains 75.5%. The HDD is significant and theoretically meaningful for the model, but the spatial pattern is mostly determined by the building stock and socio-economic variables.

The choice of interpolation method proved to be more than a technical detail. The Netatmo Kriging coefficient ($\beta = -0.701$) is approximately three times larger than the Netatmo IDW coefficient ($\beta = -0.225$), despite both using identical station observations. Kriging, assigns values to postcode that reflect the neighbourhood level temperature better, mostly by incorporating the direction of the North sea coastline. The IDW method preserved more local variation, and reflects the sensor specific conditions more than the Kriging method. The interpolation method therefore has implications for the estimated temperature sensitivity, and one should consider both options when modelling citizen science data.

6.4 Spatial Autocorrelation

The high Lambda parameter ($\lambda \approx 0.55$) across all of the spatial models confirms that around 55% of the unexplained variance in residential gas consumption is spatially structured. The AIC improvement from OLS to SEM is between 1900 and 2020 points, larger than the improvement from adding any climate variable. This means that accounting for the spatial structure of the data matters more for the model quality than KNMI or Netatmo data. Where this spatial dependence comes from is not clear. District heating is accounted for and all postcodes with less than 200 m^3 of gas consumption were not included in final analysis. This result indicates that there are neighbourhood effects in place that are not captured by the building stock or socio-economic variables. Similar occupant behaviour in neighbourhoods could be an explanation as the occupant behaviour is not totally accounted for in the models.

From a physical perspective, the high Lambda value is not surprising when the urban structure of The Hague is considered. Residential neighbourhoods in Dutch cities tend to be built in relatively homogeneous phases. A pre-war expansion area like Schilderswijk consists almost entirely of similar housing built within the same decade. Within such neighbourhoods, gas consumption patterns are highly similar between neighbouring postcodes because the buildings and thermal conditions are almost the same.

6.5 Urban morphology and the Urban Heat Island

The positive NDVI coefficient ($\beta = 222$) is consistent with the theoretical expectation from section 2.5.1 when the seasonal context is considered. Vegetation reduces UHI intensity through evaporative cooling and reduced heat retention, an effect that lowers summer temperatures in green areas. In the context of winter heating demand, this same mechanism means that greener postcodes experience cooler ambient temperatures and therefore require more space heating. The correlation matrix visible in Appendix D, shows that the correlations between NDVI and the other model variables are all near zero, confirming that the NDVI effect in the regression is largely independent of building age, dwelling size and socio-economic variables. The only slight correlation is between Netatmo HDD and NDVI ($r = 0.246$), which is the strongest correlation the HDD variable has with any other variables in the model. This is consistent with the theoretical expectation that greener areas experience weaker UHI intensity and therefore cooler temperatures as established by van der Hoeven & Wandl (2018).

The non-significance of building coverage ratio is also interpretable. The variable captures the footprint area of buildings but excludes roads, parking areas, and other asphalt surfaces that together with buildings create the urban heat island. In The Hague, the urban fabric consists of wide roads, squares and tram infrastructure. Including all of these in the variable would lead to a higher coverage percent per postcode, resulting in a better proxy of urban density, something that is now not represented well by building footprint alone.

Figure 6.2 presents the UHI index and the Netatmo Kriging HDD surface. The core of The Hague with low HDD values corresponds with a high UHI index value. The HDD values of the postcodes in the southern areas and along the coast also correspond with the UHI index value. While not a perfect match, this spatial pattern suggests that the urban morphology variables together represent some thermal conditions of an UHI, and this effect is also captured by the interpolated Netatmo HDD surface.

Comparison of UHI and HDD in The Hague

A: UHI Index (Building height, coverage rate and NDVI) B: Netatmo Kriging HDD surface

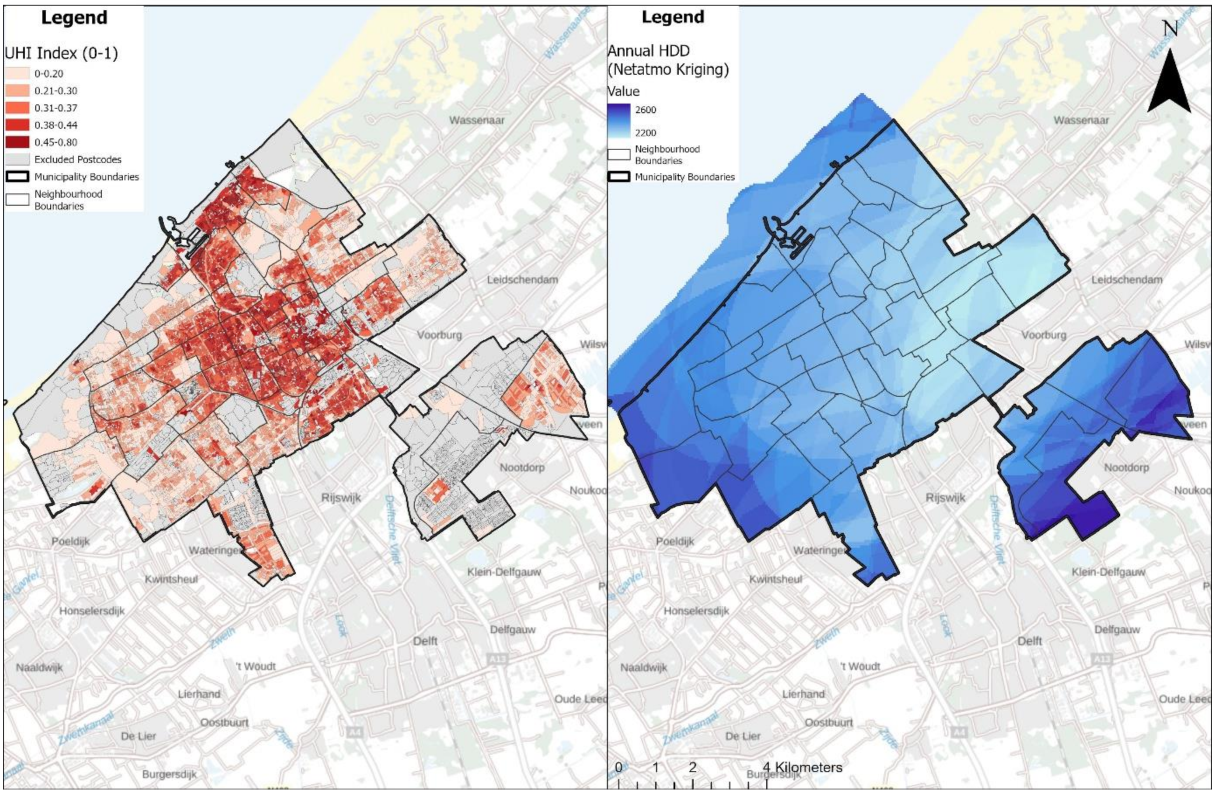

Figure 6.2: Urban Heat Island Index from measured UHI variables compared to the Heating Degree Day surface from Netatmo. Source: 3D geoinformation group & 3DGI (2025); RIVM (2022); Ohori (2022); PDOK (2021).

6.6 Policy implications

The findings of this research have direct relevance for Dutch municipalities that have to provide a ten year heat transition plan by the end of 2026. This section provides concrete policy implications from the result of this research. These implications are based on The Hague, but are applicable to other municipalities in the Netherlands.

The building stock is the highest priority. The results show that building age and dwelling size are the dominant predictors for residential gas consumption, even after controlling for socio-economic variables and the local temperature Each percentage point increase in the share of pre-1975 dwellings within a postcode is associated with approximately 4 m³ higher annual gas consumption per dwelling, compared to postcodes with newer building stock. Looking at the large concentrated amount of pre-1975 buildings in the inner city, visible in figure 4.5, this is where the challenge lies for The Hague. Under the current Dutch climate policies, the municipalities need to phase out natural gas by 2050 through localized neighbourhood execution plans (Wijkuitvoeringsplannen) (NPLW, 2023). In The Hague, the strategy for these historic urban neighbourhoods relies on collective district heating networks, because individual retrofits are often not feasible in the historic buildings.

The findings of this research, especially the drop in coefficient when including energy labels, show that post-construction insulation improvements do have an effect on the residential gas consumption. The policies can not rely entirely on changing the heat source, and must pair the district heating with insulation improvements, focussing on the postcodes with high percentages of pre-1975 buildings.

Local climate data can improve the spatial resolution of climate inputs to energy models. The Netatmo network demonstrates that intra-urban temperature variation of up to 295 HDD exists across The Hague. A range that translates to approximately 207 m^3 difference in predicted annual gas consumption between the warmest and coolest postcodes. This variation is invisible to KNMI-based energy models, which assign a near-uniform temperature to all postcodes in the city. Implementing citizen science could provide municipalities a cost-effective way to incorporate local climate inputs. Whether expanding the citizen sensing network is a task for the municipality is debatable. Netatmo stations are privately owned and their placement can not be fully controlled. A better way of improving local weather observations for the municipality is to install stations in less covered areas. This would improve interpolation quality and as a result would lead to more accurate local weather observations.

A more practical example could be the implementation of local climate in a tool the municipality of The Hague provides to its residents. On the municipality website they operate a address tool, where residents can enter their street name and home address. The tool returns the best heating alternative and transition plans for their neighbourhood currently in place. But is based on regional climate data, and this study shows that provide uniform temperature surfaces over the city. For an address level tool, the implementation of locally measured temperature data on the postcode level could provide better insights on recommendations on heating alternatives.

6.7 Limitations

The interpretation and generalizability of this research's findings are limited by several factors. These will be described below.

Scope

The temporal resolution and single year analysis of this study comes with some limitations. The gas consumption data and explanatory variables all refer to the year 2021. Aggregating these numbers to an annual number means that some anomalies are not captured. Especially for 2021, where gas consumption may be higher due to the implications of COVID-19 pandemic and the increasing amount of people that would work from home. A multi-year analysis would enable the assessment of whether the contribution of local temperature data varies across years with different temperature profiles.

For this research, gas consumption was used as proxy for space heating. The dependent variable that was used includes gas consumption for water and cooking next to space heating. While the RVO (2024) estimates that space heating accounts for approximately 80% of the residential gas use, the remaining 20% introduces noise into the dependent variable that is not responsive to temperature variation. The dependent variable also only measures natural gas drawn from the municipal grid. It does not account for secondary or alternative heating sources. Heat pumps consume electricity rather than gas for space heating, meaning postcodes with a high heat pump adoption rate will show structurally lower gas consumption than their building characteristics would predict. The share of heat pumps in Dutch homes grew substantially between 2019 and 2021 (De Wind et al., 2025), meaning some postcodes in newer developments such as Leidschenveen may already show suppressed gas consumption for reasons unrelated to building quality or local temperature.

Lastly, the Degree Days were calculated using a base temperature of 18 °C and represents the threshold below which space heating is required. But well insulated newer buildings with high energy labels retain heat much longer than older buildings with inefficient labels for example. Older buildings would require heating much sooner in this case. The threshold of 18 °C does not account for the fact that the effective base temperature varies with building quality.

Another limitation is the comparability of results across years in a potential multi-year analysis. The stations are required to have complete 365 days of measurements before they are included in the HDD calculation. In this study, applying this filter to 2021 reduced the dataset from 570 to 116 stations before further quality filtering. Adding a multi-year comparison would likely reduce valid stations even more. A snapshot comparison between station activity in January 2020 and January 2021 illustrates this problem: 570 stations were active in the 2021 snapshot while 767 were active in January 2020, with only 353 stations appearing in both snapshots. While this comparison reflects a single moment rather than full-year coverage, it confirms that the station network changes substantially between years due to stations being installed, moved or taken offline by residents

Another limitation is the geographic scope. This analysis is limited to the municipality of The Hague. The Hague was selected as a case study because of the dense Netatmo network and public available data. The coastal location and heterogenous urban morphology make it an interesting case for looking at the added value of local climate data. Smaller or less dense Dutch cities or towns are likely to not have dense Netatmo coverage or have less intra-urban differences. This makes this research not replicable in other cities, apart from the largest cities in the Netherlands. A comparison of two Dutch cities with different local climates, such as an inland city and a coastal city would provide more insights on the difference in their local climates. While this is currently possible with KNMI data, this study proved that KNMI data does not reflect the intra-urban climate well.

Netatmo data limitations

The usage of the Netatmo data also comes with some limitations. The first being that the network is not designed to be a scientific measurement. The resident could place the stations to their own choosing, resulting in that the network has no spatial presentiveness. The quality control procedure of Meier et al. (2017) removed all stations with unrealistic values e.g. indoor stations and outliers. Moreover, because the gas consumption was a yearly number, the temperature data needed to be aggregated to yearly values as well. This resulted in stations that were missing 1 measurement not being included in the analysis. Figure 4.4 visualizes the remaining stations and shows what areas of The Hague are covered. The raw dataset included more than 570 stations with measurements, and the final dataset used for the calculation of Degree Days contained only 105 stations, some not even covering The Hague but surrounding areas.

Figure 4.5 shows the spatial distribution of the 105 valid Netatmo stations across The Hague and surrounding area. While coverage is relatively dense in the residential neighbourhoods of the city centre and the western coastal strip, several areas have no or very few stations. The large green areas of Zuiderpark, Westbroekpark and the Haagse Bos contain no active stations, meaning HDD values assigned to postcodes covering these parks are entirely interpolated from surrounding urban stations rather than locally measured. More evidently, the area above Scheveningen along the coast lacks coverage, meaning that this less dense area got temperature interpolation values of urban dense area, therefore showing low HDD numbers. In reality, the cooling effect of the sea would have an effect, which is now not captured.

Even if the network was more complete and all 570 stations were included in the model, the whole surface of The Hague was not 100% accurately represented. Some areas would still have sparser station coverage while other areas would have a more dense coverage. Furthermore, even after validation measures the Netatmo network remains a citizen science network and is not professionally calibrated like the KNMI. The result of this is that stations could still be placed sub optimally near heat emitting surfaces or in the shade. Lastly, the HDD values are derived from a point measurement and while interpolation techniques try to capture the most realistic situation, the true local climate in every postcode can never be 100% represented.

The small increase in Pseudo R^2 from 0.754 to 0.755 when adding locally measured HDD could be interpreted as noise rather than a real improvement. Also because the analysis is performed for a single year. A multi-year analysis would make it possible to test whether this improvement is consistent across years with different temperature profiles, The AIC improvement when adding local HDD to the model suggest that adding locally measured temperature improves the model-fit.

CBS data suppression

The CBS Kerncijfers dataset suppresses values for PC6 areas below privacy thresholds, encoding these as negative placeholder values. This affected the age variable, which was suppressed for 5,797 of 11,440 postcodes and could therefore not be included in the main model. Rafiee et al. (2019) identified the 65+ age group as a significant determinant of heating demand due to higher daytime occupancy and preferred indoor temperatures. The exclusion of this variable may introduce upward bias in the WOZ-waarde coefficient in older inner-city neighbourhoods where elderly homeowners with high property values but low or moderate incomes live. Furthermore, the share of multi-family homes needed to be calculated using the BAG because of the same suppression issue.

This study includes only a limited set of building and occupant characteristics. Several theoretically relevant variables could not be included due to data availability or CBS suppression. The age distribution of residents for example, or the home ownership tenure identified by Leth-Petersen & Togeby (2001).

Climate data

This research operationalizes local climate through Heating Degree Days derived from temperature observations. Wind speed and precipitation are not included despite being theoretically relevant determinants of heating demand (Wever, 2008). While both variables being theoretically available in both KNMI and Netatmo, they were excluded from the spatial models for two reasons. First, the Netatmo dataset lacks sufficient and reliable spatial coverage for wind and precipitation. These specific sensors are more vulnerable to setup errors by the user. Underneath a roof or behind a wall for example. Furthermore, not every Netatmo sensor had a rain gauge. This would shrink the dataset even more, after validating the temperature measurements.

Second, because the dependent variable gas consumption is aggregated at an annual level. Constructing annual averages for wind speed or total precipitation smooths out the extreme spikes. These extreme spikes, like strong wind or rainstorms, likely drives the increase in heating demand. But as these extreme spikes usually only last for a couple of days, they will not be visible in an annual aggregation, making them weak predictors for the spatial variance.

To compensate for the limited representation of local climate, the effects of an Urban Heat Island was included to include an additional dimension next to the temperature measurements. Three morphological variables were used to operationalize the UHI, namely building height, NDVI and coverage ratio. While these variables all capture an element of the UHI, they do not represent the UHI as a whole. Building height affects gas consumption through multiple ways. It contributes to heat retention, but taller buildings also have a higher surface-to-volume ratio which increases heat loss. They are also more exposed to wind speed compared to other buildings. The negative coefficients in the model therefore could represent a different mechanism than a pure UHI proxy. The coverage ratio limitations were already explained in section 6.5, but this variable also does not represent the UHI in full. The NDVI is a vegetation index that does not differentiate in the types of vegetation. Ground level vegetation has a different effect on the UHI effect than tree cover for example. Furthermore, it represents growth during the growth season which is the summer, while the heat demand is highest during the winter. The UHI effects that are active during the summer may not be relevant during the winter.

7 Conclusion

This research investigated to what extent incorporating high-resolution, locally measured climatic data improves the explanation of spatial patterns in residential gas consumption, compared to using spatially coarse regional temperature data. This section will first answer all sub questions before answering the main research question.

Sub question 1: How do Heating Degree Day patterns derived from KNMI and Netatmo datasets differ spatially within The Hague?

The KNMI and Netatmo networks produce fundamentally different HDD surfaces across The Hague. The KNMI Kriging interpolation, based on three stations located in open rural environments near Hoek van Holland, Voorschoten and Rotterdam, produces a near-uniform surface with a standard deviation of only 17.6 HDD across 9410 postcodes. The total intra-urban range is 99 HDD and is not sufficient to differentiate between the dense city centre and the suburban neighbourhoods of The Hague. The Netatmo Kriging interpolation, based on 105 stations distributed within and around the urban area, produces a standard deviation of 47.6 HDD and captures a gradient ranging from 2,246 HDD in the warmest inner-city postcodes to 2541 HDD in the cooler coastal and suburban areas. A range of 295 Degree Days across a distance of only a few kilometres, is more representative of The Hague. The Netatmo network provides intra-urban temperature variation that the KNMI network cannot produce due to its design.

Sub question 2: How do local temperature and urban morphology variables contribute to explaining spatial variation in residential gas consumption?

Locally measured Netatmo HDD contributes a statistically significant and independent explanation of spatial gas consumption variation. In the Spatial Error Model, each degree day of additional heating demand is associated with approximately 0.70 m^3 additional annual gas consumption per postcode. Among the urban morphology variables, building height is negatively and significantly associated with gas consumption, consistent with the heat retention and wind shielding effects of taller buildings. Mean NDVI is positively associated with consumption. This is consistent with the seasonal interpretation of the UHI effect: vegetation reduces urban heat retention through evaporative cooling, resulting in cooler winter temperatures in greener postcodes and therefore higher heating demand.

However, these climate and morphological variables operate at a small margin that is determined by building stock quality and socio-economic composition. The explained variance without local HDD is 75.4%, and adding the local HDD and UHI variables increases it with 0.1% to 75.5%. Building age is the strongest predictor across all model specifications, with each percentage point increase in pre-1975 dwellings associated with approximately 3.95–4.03 m^3 additional annual gas consumption. Log-transformed property value has the largest standardised effect of all variables, followed by building age and dwelling size. These building and socio-economic variables together explain the large majority of spatial variance in gas consumption before any climate variable is introduced.

Sub question 3: To what extent does including locally measured climatic data improve the performance of a spatial regression model of residential gas consumption, and does the interpolation method affect this outcome?

Including Netatmo HDD and urban morphology variables improves the Spatial error Model by 372 AIC points relative to the baseline model, compared to only 2 points for the KNMI Kriging. The Netatmo Kriging coefficient is three times larger than the Netatmo IDW equivalent, confirming that the choice of interpolation method has implications for the estimated temperature sensitivity. Ordinary Kriging is the preferred method for modelling the intra-urban citizen sensing data in this case.

Together, these individual findings answer the main research question:

To what extent does incorporating high-resolution, locally measured climatic data improve the explanation of spatial patterns in residential gas consumption, compared to using spatially coarse regional temperature data?

Locally measured Netatmo temperature data improves the explanation of spatial patterns in residential gas consumption compared to the regional KNMI data, but the improvement is small. The baseline Spatial Error Model without climate data already explains 75.4% of the variance, and the addition of Netatmo HDD and urban morphology increases this variance to 75.5%. The building age, dwelling size and property value remain the main determinants of spatial gas consumption patterns in The Hague. The locally measured Netatmo data does not change the explanation of why some postcodes consume more gas than others. But the value of the Netatmo network lies in that it can produce a temperature surface that varies realistically across the city, and is something that can not be produced by using current KNMI data.

The primary novel contribution of this study is the first direct empirical comparison of KNMI and Netatmo-derived Heating Degree Days as competing climate inputs to a spatial regression model of residential gas consumption, demonstrating that the choice of climate data source. not just the inclusion of climate data, determines whether temperature adds meaningful explanatory power.

If citizen sensing networks become more available, have increased coverage and accuracy, integrating the locally measured temperature into municipal energy models is a logical step and is statistically justified by this research.

8 Recommendations

For future research, it is recommended to expand the research to a multi-year analysis. Evaluating multiple consecutive years would help determine if the small explanatory power of the local HDD observed in 2021 would remain or whether it fluctuates during years. Next to a different temporal resolution, the geographic scope could be extended. Including the local temperature of more cities could provide insights on the locally measured temperatures and their effect on gas consumption, and if their effect remains modest in other cities like Utrecht or Eindhoven.

For the usage of citizen sensing climate networks, it is recommended that future research looks past Netatmo data. The data is incomplete and not calibrated. Due to time and capability constraints, this analysis is limited to the year 2021. Ohori has made data available from 2015 to 2023, meaning a multi-year analysis is possible and recommended for future research. This dataset is however limited to only the The Hague area. While there are options to make use of the Netatmo data for other areas it would take a long time and is not cost effective to activate their API.

Future research could improve the operationalisation of the Urban Heat Island effect by selecting more variables that could represent the UHI. The coverage percentage could be improved upon by including roads, parking lots and other asphalt. Additionally, including the defined Local Climate zones by Stewart and Oke (2012) could lead to a better proxy of UHI. Moreover, to better capture the local climate, it would be recommended to look into wind speed and precipitation and how to include those in the annual aggregation. Or drop the annual aggregation and take monthly values, but then analysis at the same postcode scale would not be possible. Since neighbourhood effects were quite large in this study, looking at a slightly larger scale could be an improvement and would also lead to less suppressed CBS data.

Practical recommendations

For the municipality of The Hague, the findings of this research show that building characteristics are dominant in gas consumption variance, and that the local climate only adds 0.1% to the variance. Practically, this means that climate adaptation strategies, like urban greening to combat the UHI effect, will have a negligible impact on gas consumption. The municipality’s focus remains in improving the thermal conditions of the building stock or connecting neighbourhoods to district heating.

A suitable alternative would be the WOW-NL network, which is an UK portal for personal weather stations. Currently, 10000 stations are active in 220 countries. It is currently in progress of transferring to KNMI, as WOW is cancelling their platform. The KNMI could make it available for public, and as a result a citizen sensing network other than Netatmo could be available. It must be noted that it remains difficult to publicly share data and locations of personal weather stations due to privacy reasons, and that such a network with a dense coverage could never become available.

Whether a fully professional measurement network is needed depends on the purpose. For nationwide climate monitoring and regional weather forecasting the KNMI is well-suited and covers most of the Netherlands sufficiently. For intra-urban studies at the postcode or neighbourhood level, as this study demonstrates, the KNMI is insufficient due to the placement in open rural areas. However, the modest improvement when adding locally measured climatic data does not justify the costs that are required for setting up a dense professional station network at the municipal level.

9 References

3D geoinformation group & 3DGI. (2025). 3DBAG: Landelijke 3D gebouwmodellen van Nederland (Versie v2025.09.03) [Dataset]. https://3dbag.nl/

Anselin, L. (1988). *Spatial econometrics: methods and models* (Vol. 4). Springer Science & Business Media.

Arentze, T. A., & Stevens, R. P. (2018). Het belang van secundaire objectkenmerken woningen bij de bepaling van de WOZ-waarde: vooronderzoek WOZ Woningen REN.

Bansal, P., & Quan, S. J. (2022). Relationships between building characteristics, urban form and building energy use in different local climate zone contexts: An empirical study in Seoul. *Energy and Buildings*, *272*, 112335.

Batenburg, A., Hopman, B., Wijhuizen, E., Dalla Longa, F., & Mulder, P. (2024). Energiearmoede in Nederland 2019–2023. TNO innovation for life, Amsterdam, The Netherlands.

Biesiot, W., & Noorman, K. J. (1999). Energy requirements of household consumption: a case study of The Netherlands. *Ecological Economics*, *28*(3), 367-383.

Brounen, D., Kok, N., & Quigley, J. M. (2012). Residential energy use and conservation: Economics and demographics. *European Economic Review*, *56*(5), 931-945.

Centraal Bureau voor de Statistiek. (2021). Energielevering aan woningen en bedrijven naar postcode 2021 [Dataset]. https://www.cbs.nl/nl-nl/maatwerk/2023/46/energielevering-aan-woningen-en-bedrijven-naar-postcode

Centraal Bureau voor de Statistiek. (2021). Kerncijfers postcodegebieden 2021 [Dataset]. CBS Statline

Centraal Bureau voor de Statistiek. (2025, 24 juli). *Monitor Energiearmoede 2023*. Centraal Bureau Voor de Statistiek. https://www.cbs.nl/nl-nl/maatwerk/2025/30/monitor-energiearmoede-2023

Chapman, L., Bell, C., & Bell, S. (2016). 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(6), 3168-3180.

Danylo, O., Bun, R., See, L., & Charkovska, N. (2019). High-resolution spatial distribution of greenhouse gas emissions in the residential sector. *Mitigation and Adaptation Strategies for Global Change*, *24*(6), 941-967.

Demuzere, M., Harshan, S., Järvi, L., Roth, M., Grimmond, C. S. B., Masson, V., ... & Wouters, H. (2017). Impact of urban canopy models and external parameters on the modelled urban energy balance in a tropical city. *Quarterly Journal of the Royal Meteorological Society*, *143*(704), 1581-1596.

de Wind, J., Alpízar-Castillo, J., Visser, J., & Ramírez-Elizondo, L. (2025). An in-depth analysis of residential E-cooling demand in The Netherlands. *Case Studies in Thermal Engineering*, *73*, 106469.

Du, H., Loenders, S., & Han, Q. (2025). Navigating the built environment energy transition in the Netherlands: A tool for energy-oriented design and planning. *Energy and Buildings*, *346*, 116223.

Eto, J. H. (1988). On using degree-days to account for the effects of weather on annual energy use in office buildings. Energy and buildings, 12(2), 113-127.

Gemeente Den Haag. (2022). Transitievisie Warmte Den Haag 2021-2030. Gemeente Den Haag. https://denhaag.raadsinformatie.nl/document/10731340/1?connection_type=16&connection_id=706160

Goovaerts, P. (1997). *Geostatistics for natural resources evaluation*. Oxford university press.

Guerra Santin, O. G., Itard, L., & Visscher, H. (2009). The effect of occupancy and building characteristics on energy use for space and water heating in Dutch residential stock. *Energy and buildings*, *41*(11), 1223-1232.

Guerra Santin, O. (2010). *Actual energy consumption in dwellings: the effect of energy performance regulations and occupant behaviour*. [Dissertation (TU Delft), Delft University of Technology]. IOS Press.

Hofstra, N., Haylock, M., New, M., Jones, P., & Frei, C. (2008). Comparison of six methods for the interpolation of daily, European climate data. Journal of Geophysical Research: Atmospheres, 113(D21).

Kadaster. (2021). *Basisregistratie Adressen en Gebouwen (BAG)* [Dataset]. PDOK.

Klimaatmonitor. (2023). Aardgasverbruik gebouwde omgeving. Rijkswaterstaat. https://klimaatmonitor.databank.nl

Klok, E. J., Schaminée, S., Duyzer, J., & Steeneveld, G. J. (2012). De stedelijke hitte-eilanden van Nederland in kaart gebracht met satellietbeelden. Tno report, 6.

Koninklijk Nederlands Meteorologisch Instituut. (2000). *Handboek waarnemingen*. Geraadpleegd van KNMI Handboek Waarnemingen [1]

Koninklijk Nederlands Meteorologisch Instituut [KNMI]. (2021). *Daggegevens van het weer in Nederland* [Dataset]. https://daggegevens.knmi.nl/

Koninklijk Nederlands Meteorologisch Instituut. (2021). Klimaatatlas Nederland [Dataset]. https://www.knmi.nl/klimaat-viewer/kaarten

Leth-Petersen, S., & Togeby, M. (2001). Demand for space heating in apartment blocks: measuring effects of policy measures aiming at reducing energy consumption. *Energy Economics*, *23*(4), 387-403.

Li, X., Zhou, Y., Yu, S., Jia, G., Li, H., & Li, W. (2019). Urban heat island impacts on building energy consumption: A review of approaches and findings. *Energy*, *174*, 407-419.

Majcen, D., Itard, L., & Visscher, H. (2013). Actual and theoretical gas consumption in Dutch dwellings: What causes the differences?. *Energy Policy*, *61*, 460-471.

Majcen, D., Itard, L., & Visscher, H. (2015). Statistical model of the heating prediction gap in Dutch dwellings: Relative importance of building, household and behavioural characteristics. *Energy and Buildings*, *105*, 43-59.

Majcen, D. (2016). Predicting energy consumption and savings in the housing stock: A performance gap analysis in the Netherlands. *A+ be| Architecture and the Built Environment*, (4), 1-224.

Mashhoodi, B. (2018). Spatial dynamics of household energy consumption and local drivers in Randstad, Netherlands. *Applied Geography*, *91*, 123-130.

Mashhoodi, B., Stead, D., & van Timmeren, A. (2019). Spatial homogeneity and heterogeneity of energy poverty: a neglected dimension. *Annals of GIS*, *25*(1), 19-31.

Mashhoodi, B., Stead, D., & van Timmeren, A. (2020). Local and national determinants of household energy consumption in the Netherlands. GeoJournal, 85(2), 393-406.

Meier, F., Fenner, D., Grassmann, T., Otto, M., & Scherer, D. (2017). Crowdsourcing air temperature from citizen weather stations for urban climate research. *Urban Climate*, *19*, 170-191.

Mourshed, M. (2012). Relationship between annual mean temperature and degree-days. *Energy and buildings*, *54*, 418-425.

Muller, C. L., Chapman, L., Johnston, S., Kidd, C., Illingworth, S., Foody, G., ... & Leigh, R. R. (2015). Crowdsourcing for climate and atmospheric sciences: current status and future potential. *International Journal of Climatology*, *35*(11), 3185-3203.

Namazkhan, M., Albers, C., & Steg, L. (2020). A decision tree method for explaining household gas consumption: The role of building characteristics, socio-demographic variables, psychological factors and household behaviour. *Renewable and Sustainable Energy Reviews*, *119*, 109542.

Napoly, A., Grassmann, T., Meier, F., & Fenner, D. (2018). Development and application of a statistically-based quality control for crowdsourced air temperature data. *Frontiers in Earth Science*, *6*, 118.

Nationaal Programma Lokale Warmtetransitie. (2023). Proces totstandkoming uitvoeringsplan warmtetransitie. In *Proces Totstandkoming Uitvoeringsplan Warmtetransitie*. https://www.nplw.nl/uploads/files/Uitvoeringsplan/Proces-totstandkoming-uitvoeringsplan-warmtetransitie-NPLW.pdf

Niessink, R., Obbink, Y., Van Leeuwen, M., & TNO Publiek. (2025). Energieverbruik en energiebesparing in de bestaande woningbouw en utiliteitsbouw. In Rijksdienst voor Ondernemend Nederland (RVO) & TNO, *Achtergrondrapport Ten Behoeve van Monitor Verduurzaming Gebouwde Omgeving 2024* [Report]. TNO.

Ohori, K. (2022). Weather data for Netatmo personal weather stations around the Hague (2021) [Dataset]. In *4TU.ResearchData*. https://doi.org/10.4121/18065285

Pijnenburg, B. (2021). Assesment of heating energy use at postal code level 6: influence of the local climate, Master's thesis, GIMA (TUD, UT-ITC, UU, WUR), pp. 91, 2021

Publieke Dienstverlening Op de Kaart. (2021). Postcode 6 gebieden [Dataset]. PDOK. https://www.pdok.nl

Rød, J. K., & Maarse, M. J. (2021). Using citizen sensing to identify heat-exposed neighbourhoods. Urban Science, 5(1), 14.

Rafiee, A., Dias, E., & Koomen, E. (2019). Analysing the impact of spatial context on the heat consumption of individual households. *Renewable and Sustainable Energy Reviews*, *112*, 461-470.

Ratti, C., Baker, N., & Steemers, K. (2005). Energy consumption and urban texture. Energy and buildings, 37(7), 762-776.

Rijksdienst voor Ondernemend Nederland. (2021). Energielabels woningen gemeente Den Haag [Dataset]. Den Haag in Cijfers / RVO. https://denhaag.incijfers.nl

Rijksdienst voor Ondernemend Nederland. (2024). *Monitor Verduurzaming Gebouwde Omgeving 2024*. RVO.

Rijksinstituut voor Volksgezondheid en Milieu. (2022). *Conditie - Groenkaart van Nederland 5 m resolutie* [Map]. Atlas Natuurlijk Kapitaal. https://www.atlasnatuurlijkkapitaal.nl

Rode, P., Keim, C., Robazza, G., Viejo, P., & Schofield, J. (2014). Cities and energy: urban morphology and residential heat-energy demand. Environment and Planning B: Planning and Design, 41(1), 138-162.

Rovers, V., Niessink, R., Loonen, P., van der Wal, A., & Matthijssen, E. (2021). Energievraag van ruimtekoeling in woningen. TNO.

Roxon, J., Ulm, F. J., & Pellenq, R. M. (2020). Urban heat island impact on state residential energy cost and CO2 emissions in the United States. *Urban Climate*, *31*, 100546.

Roy, E. I., Pronk, M., Agugiaro, G., & Ledoux, H. (2022). Inferring the number of floors for residential buildings.International Journal of Geographical Information Science (online), 37(4), 938-962. https://doi.org/10.1080/13658816.2022.2160454

Sharston, R., & Singh, M. (2025). Urban morphology, urban heat island (UHI) and building energy consumption: A critical review of methods and relationships among influential parameters. Building Services Engineering Research & Technology, 46(4), 561-584.

Stahl, K., Moore, R. D., Floyer, J. A., Asplin, M. G., & McKendry, I. G. (2006). Comparison of approaches for spatial interpolation of daily air temperature in a large region with complex topography and highly variable station density. Agricultural and forest meteorology, 139(3-4), 224-236.

Stewart, I. D., & Oke, T. R. (2012). Local Climate Zones for Urban Temperature Studies. *Bulletin of the American Meteorological Society*, *93*(12), 1879-1900. https://doi.org/10.1175/BAMS-D-11-00019.1

Sluiter, R. (2012). *Interpolation methods for the climate atlas* (Vol. 705). De Bilt, The Netherlands: KNMI.

Sunikka-Blank, M., & Galvin, R. (2012). Introducing the prebound effect: the gap between performance and actual energy consumption. Building Research & Information, 40(3), 260-273.

Van Ackere, S., Van Eetvelde, G., Schillebeeckx, D., Papa, E., Van Wyngene, K., & Vandevelde, L. (2015). Wind resource mapping using landscape roughness and spatial interpolation methods. *Energies*, *8*(8), 8682-8703.

van den Brom, P., Meijer, A., & Visscher, H. (2018). Performance gaps in energy consumption: household groups and building characteristics. Building Research & Information, 46(1), 54-70.

van den Brom, P., Hansen, A. R., Gram-Hanssen, K., Meijer, A., & Visscher, H. (2019). Variances in residential heating consumption–Importance of building characteristics and occupants analysed by movers and stayers. *Applied Energy*, *250*, 713-728.

Van Der Hoeven, F., & Wandl, A. (2015). Amsterwarm: Mapping the landuse, health and energy-efficiency implications of the Amsterdam urban heat island. Building Services Engineering Research and Technology, 36(1), 67-88.

van der Hoeven, F., & Wandl, A. (2018). Haagse Hitte: Het Haagse warmte-eiland in kaart gebracht. https://doi.org/10.7480/isbn.9789463660037

Voskamp, I. M., Visscher, M. N., Vreugdenhil, C., Van Lammeren, R. J., & Sutton, N. B. (2021). Spatial, infrastructural and consumer characteristics underlying spatial variability in residential energy and water consumption in Amsterdam. *Sustainable Cities and Society*, *72*, 102977.

Vringer, K. (2005). Analysis of the energy requirement for household consumption. Utrecht University.

Wackernagel, H. (2003). Ordinary kriging. In Multivariate geostatistics: an introduction with applications (pp. 79-88). Berlin, Heidelberg: Springer Berlin Heidelberg.

Wever, N. (2008). Effectieve temperatuur en graaddagen. Klimatologie en klimaatscenario's. KNMI publicatie, 219.

Wolters, D., Homan, C., & Bessembinder, J. (2011). Ruimtelijke klimatologische verschillen binnen Nederland. Koninklijk Nederlands meteorologisch instituut, De Bilt.

Wouters, H., De Ridder, K., Poelmans, L., Willems, P., Brouwers, J., Hosseinzadehtalaei, P., Tabari, T. Vanden Broucke, S., Van Lipzig, N. & Demuzere, M. (2017). Heat stress increase under climate change twice as large in cities as in rural areas: A study for a densely populated midlatitude maritime region. Geophysical Research Letters, 44(17), 8997-9007.

10 Appendices

10.1 Appendix A: Gas consumption table

Table A1: Preview of gas consumption table delivered by the CBS (2021)

Postcode6	*GC homes*	*GC Homes corrected*	*Electricity Homes*	*GC Business*	*Electricity Business*
2171BM	1390	1410	2470	4410	22600
2171BN	1400	1420	3060	3290	18920
2171BP	1080	1100	2630	2030	24660
2171BR	900	910	1860	1640	9630

10.2 Appendix B: Kriging cross-validation statistics

Table A2: Ordinary Kriging model output

Statistic	KNMI	Netatmo
Mean Error	2,30	-3,01
Root Mean Square	64,87	168,15
Root Mean Square Standardized	1,00	0,90
Average Standard Error	65,79	187,01

10.3 Appendix C: R code for data transformation and regression models

C1: Data cleaning and variable construction

```
library(dplyr); library(spdep); library(spatialreg)

# Load CBS data - negative values are privacy suppression codes
pc6 <- read.csv2("PC6_final.csv")
pc6_clean <- pc6 |> mutate(across(where(is.numeric),
                          ~ ifelse(. < 0, NA, .)))

# Log-transform WOZ-waarde to reduce right skew in property values
pc6_clean$log_woz <- log(pc6_clean$gemiddelde_woz_waarde_woning)

# Building age from BAG rather than CBS (CBS suppresses ~40% of postcodes)
bag <- read.csv2("bag_bouwjaar_pc6.csv")
pc6_clean <- left_join(pc6_clean,
                    bag |> mutate(pct_pre_1975_bag = pre_1975_float * 100),
                    by = "postcode6")

# Filter: remove district heating and incomplete cases
pc6_m1 <- pc6_clean |>
  filter(gas > 200,
         is.na(pct_leegstand) | pct_leegstand <= 10,
         complete.cases(pick(gas, pct_pre_1975_bag, gem_opp_woning,
                          pct_meergezins_bag, hh_grootte, log_woz)))
```

C2 Spatial weights and autocorrelation diagnostics

```
# k-nearest neighbour weights (k=8), row-standardised
make_listw <- function(df) {
  coords <- cbind(df$POINT_X, df$POINT_Y)
```

```
  nb    <- knn2nb(knearneigh(coords, k = 8))
  nb2listw(nb, style = "W", zero.policy = TRUE)
}

lw_m1  <- make_listw(pc6_m1)
lw_m3k <- make_listw(pc6_m3_kriging)

# Moran's I confirms spatial autocorrelation in OLS residuals
moran.test(residuals(model1), lw_m1, zero.policy = TRUE)

# LM tests determine model type: adjRSerr >> adjRSlag -> Spatial Error Model
lm.LMtests(model1, lw_m1,
           test = c("RLMerr", "RLMlag"))
```

C3 Spatial Error Model estimation

```
# Baseline SEM (building stock and socio-economic variables)
sem1 <- errorsarlm(
  gas ~ pct_pre_1975_bag + gem_opp_woning + pct_meergezins_bag +
        hh_grootte + log_woz,
  data = pc6_m1, listw = lw_m1, zero.policy = TRUE)

# Model 3 adds locally measured HDD and urban morphology
sem3_kriging <- errorsarlm(
  gas ~ pct_pre_1975_bag + gem_opp_woning + pct_meergezins_bag +
        hh_grootte + log_woz + Netatmo_Kriging +
        b3_h_70p + coverage_percent + Mean_NDVI,
  data = pc6_m3_kriging, listw = lw_m3k, zero.policy = TRUE)

# AIC comparison - lower is better, difference > 10 is meaningful
cat("Baseline AIC:", AIC(sem1), "\n")
cat("Netatmo AIC: ", AIC(sem3_kriging), "\n")
cat("Improvement: ", AIC(sem1) - AIC(sem3_kriging), "points\n")
```

10.4 Appendix D: Correlation Matrix

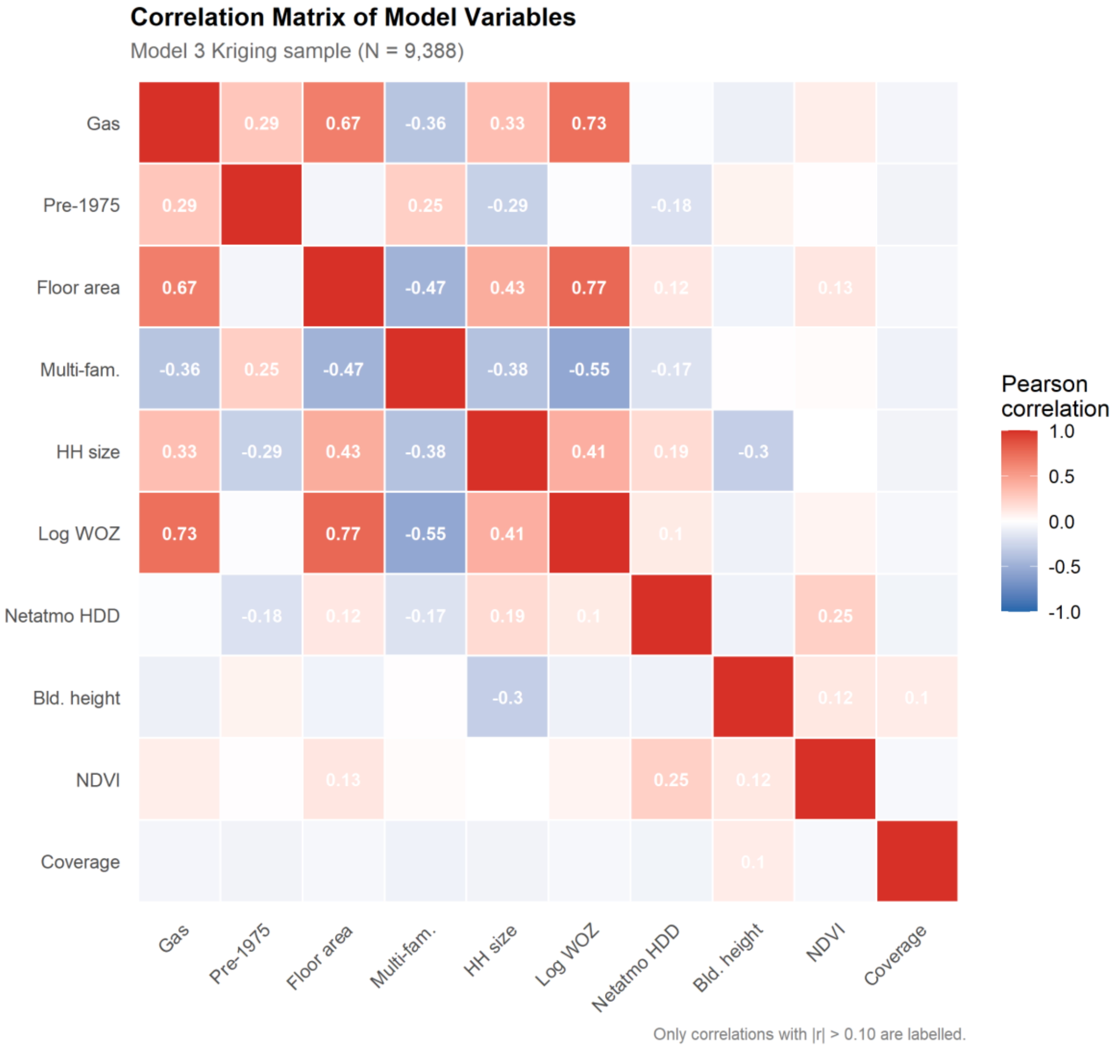

Figure D1: Correlation matrix of model 3 variables