

Online Appendix

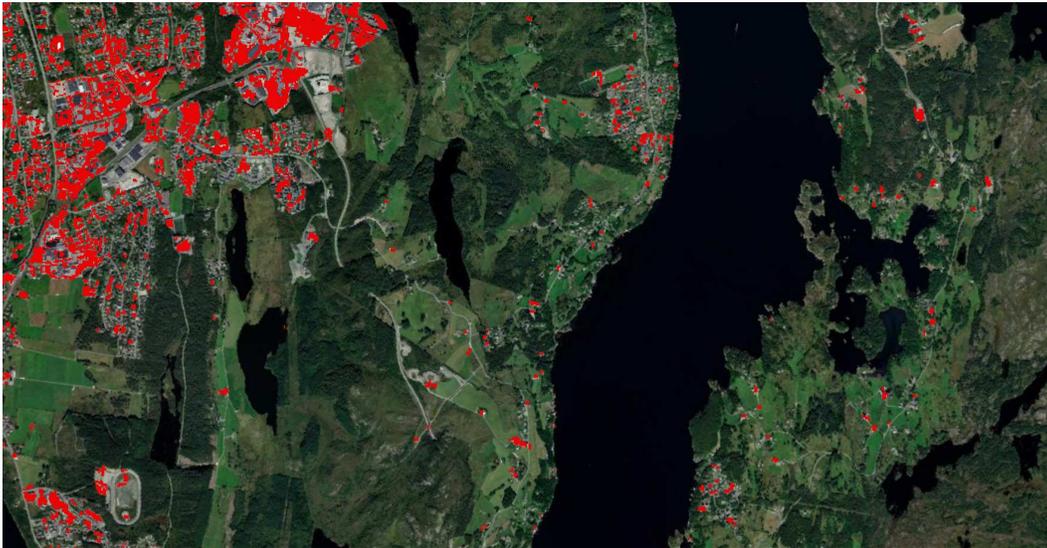
A Appendix

A.1 Selection of Urban Residential Built-Up Areas

Below, we delve into greater detail about the processes involved in defining our unit of observation. The area showcased in the satellite imagery provided by ESRI's ArcGIS World Imagery (2017) is the southeastern end of Haugesund, focusing on the neighborhood of Norheim, which we use to illustrate these processes.

Stage 1: Data Extraction from ESM

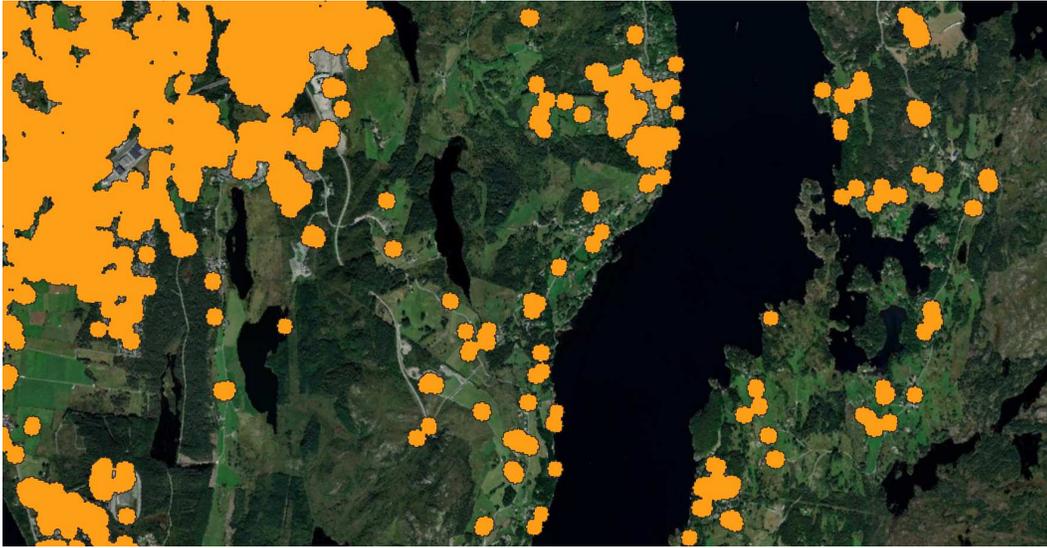
High-resolution remote sensing data (10m x 10m) from the European Settlement Map (ESM) of 2015 is used to indicate residential built-up areas. The ESM data focuses on residential zones, explicitly excluding areas of industrial build-up.



The image above highlights the initial stage with residential build-up in red excluding industrial build-up.

Stage 2: Buffer Calculation

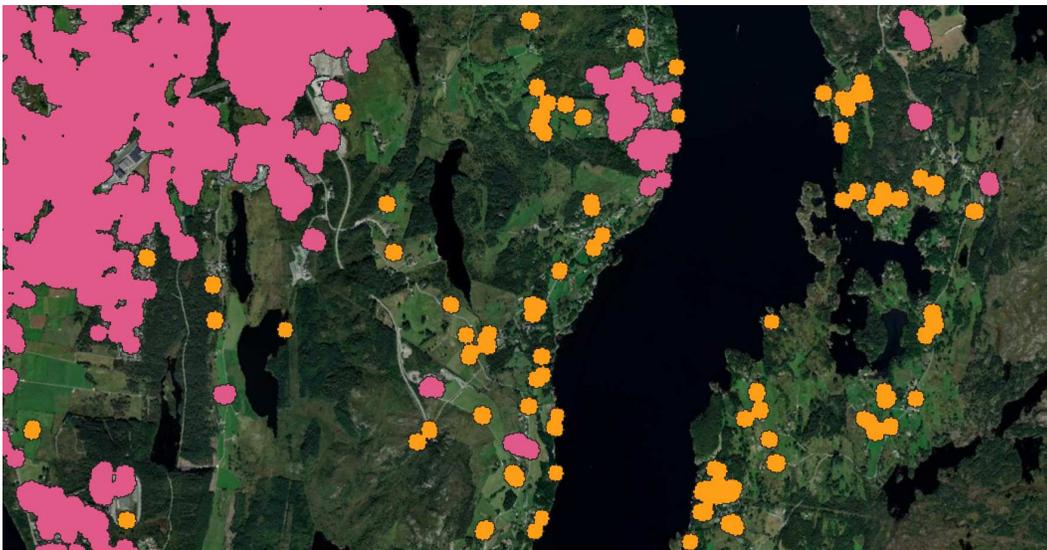
A buffer with a radius of 50 meters is created around all identified residential built-up areas to include the immediate surroundings.



Above, the 50-meter buffers around residential build-up are shown in orange.

Stage 3: Removal of Non-Contiguous Areas

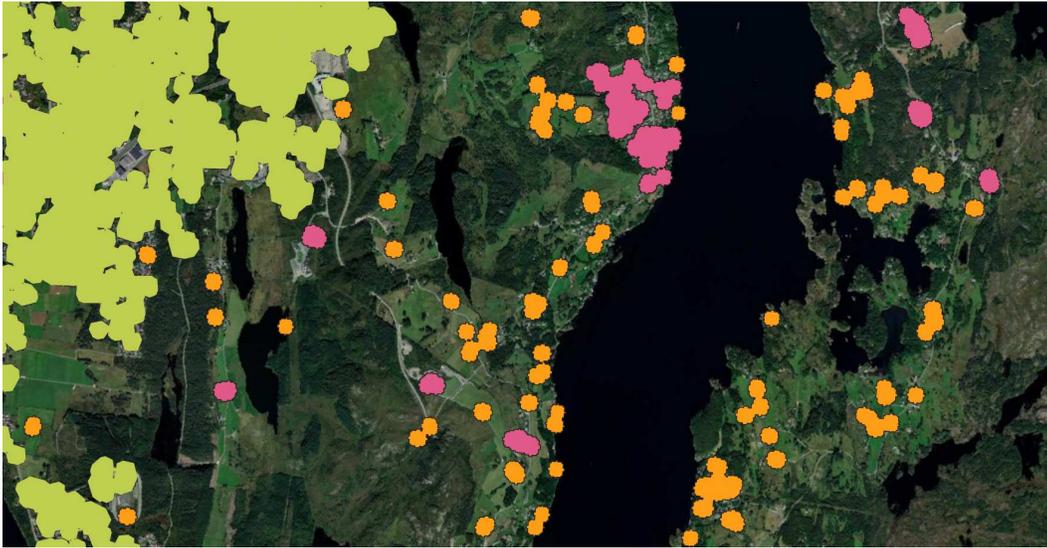
Non-contiguous areas, where the ratio of built-up area to open space is less than one to ten, are dropped. This step helps exclude small standalone housing settlements far from urban agglomerations, focusing on the integrated urban structure.



The image above shows contiguous areas pink, excluding the isolated orange ones in order to focus on the contiguous urban structure.

Stage 4: Matching with GHS-SMOD Data for Validation and Final Selection

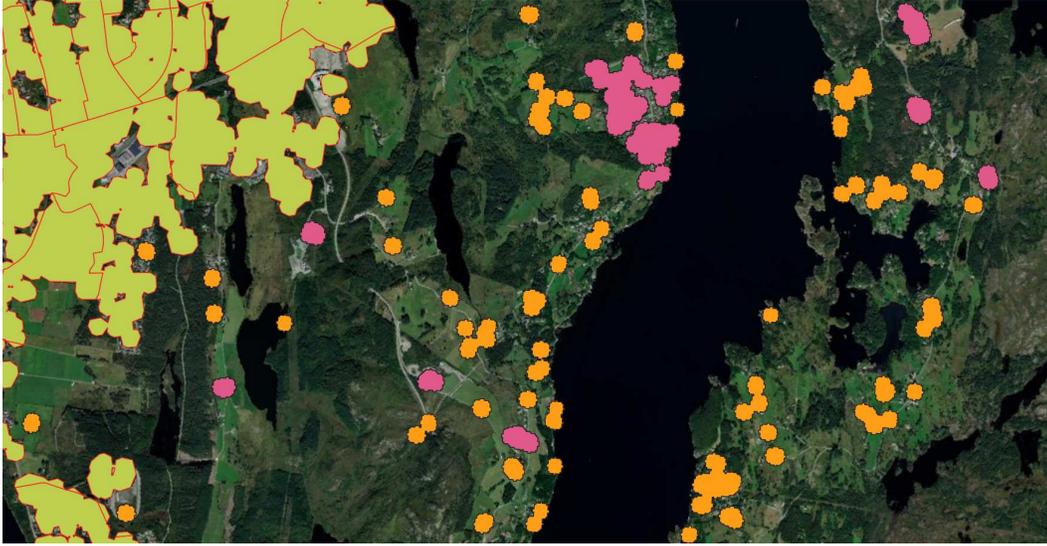
For validation, we match the residential built-up areas with with Global Human Settlement Model (GHS-SMOD) grid data from 2015, which provides the degree of urbanization on a 1 km x 1 km grid. We retain all residential built-up areas that are within or touch areas classified as 'urban' in the GHS-SMOD data (DN>20). This includes both urban core and peripheral areas like suburbs.



Displayed above, areas in light green have been matched with GHS-SMOD data, identifying the continuum of urban neighborhoods for our final sample.

Stage 5: Separation Using *Grunnkrets* Borders

The identified areas are then separated based on *grunnkrets* borders as provided by Kartverket. This step ensures that the neighborhoods fall within the official administrative boundaries for accurate analysis and reporting.



The *grunnkrets* borders in red delineate our final units of observation, the neighborhoods.

A.2 Distance to the CBD

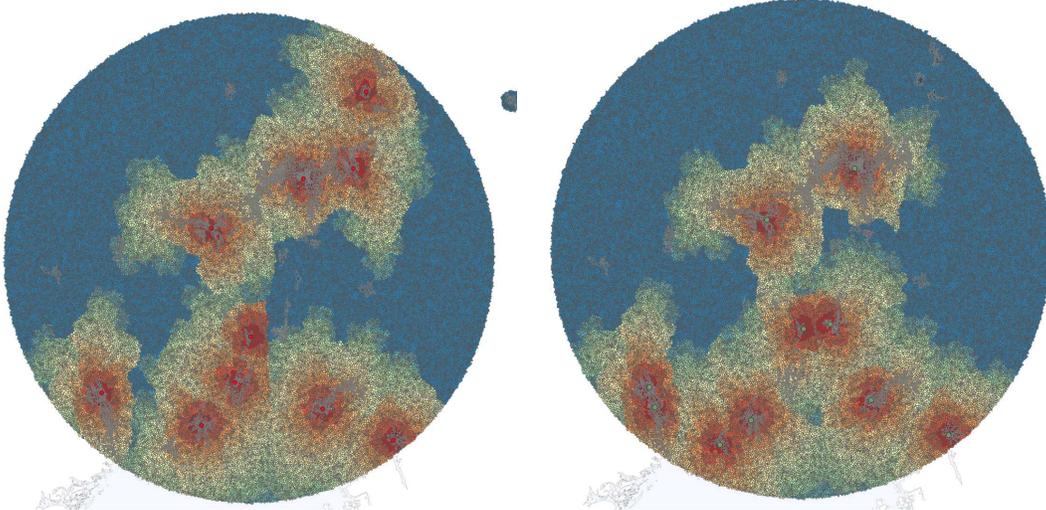
For the gradients of density and its components with respect to the city center, we need a definition of where the city center is located. There is no unambiguous approach in the empirical literature; often-used definitions work with job density or qualitative criteria (historical center, public transport hub etc), see [Liotta et al. \(2022\)](#). We follow these lines of thought and work with the density of cafés recorded in the Open Street Map data as a proxy for the CBD. Assuming that where people work they have to consume food and beverages, implies that a high density of cafés signals high levels of business activity. This approach has recently been used by [Ahlfeldt et al. \(2022\)](#), who relate the density of Starbucks franchises to prime business locations within world cities. In addition, recent work has highlighted the role of cafés and restaurants as endogenous amenities, in particular in the city center ([Aguar and Bils, 2015](#), [Baum-Snow and Hartley, 2020](#)). Using the Open Street Map data on the location of cafés, we define the gravitational center of consecutive areas that are larger than half a km^2 and have a café density of more than 5 cafés per km^2 to be a city center. This definition allows us to define at least one city center in all except three of our urban areas. In downtown Oslo, we merged the city centers that had less than 5km distance to one another. In this way, we obtain a total of 25 city centers in all urban areas in Norway

in our final sample. Most urban areas only have one city center, but some have more (such as Oslo).

As a robustness check, we create an alternative measure for the city center based on ports. In Norway, ports are natural harbors, and in an economy strongly driven by fishing, sea trade and - more recently - oil, they correlate strongly with historical city centers (Helle et al., 2006). Because of strong path dependence, these tend to overlap with modern centers in many cities (unless external shocks to the urban structure have occurred which does not seem to be the case here), see Ahlfeldt et al. (2022). We use data on the size of ports from the World Port Index. As the coordinates of the ports reported in the World Port Index are in some cases on land and in others on water, we unify locations by hand using daylight satellite images. Moreover, we compare pre-industrial-revolution maps of Norway with the location of ports in urban areas to prove that they are highly correlated. Hence, the location of ports captures historical - and still modern-day - city centers. Based on the port location, we obtain 19 city centers for all urban areas in Norway in our final sample. Most urban areas only have one city center, but some have more (such as Oslo). Figure A-1 compares the two definitions of the city center for Oslo.

For distance calculations, we work with the shortest path through the terrain. We do so for various reasons. First, there is a certain amount of unevenness in the harsh Norwegian terrain (including mountains, islands and waterways) which might impact Euclidean distance ('how the crow flies'). On the other hand, roads are man-made, so road distance potentially entails an endogeneity issue that we would like to avoid. Our way out of this is to work with the shortest path through the terrain. Also, we assume that transport costs are equal to the incline of the terrain and that traveling over water has a cost equal to a 10 degree incline in a 100m \times 100m raster. Comparing actual road data and shortest paths reveals that overground roads are often very close to shortest paths. Larger deviations are often associated with the location of tunnels. To ensure that our choice of distance measurement does not drive the results, we conduct a robustness check in Figure A-7, where we measure distance with Euclidean distance rather than the shortest path given the terrain. The pattern of the density gradients is very similar to the one based on the shortest path through the terrain.

Figure A-1: Oslo City Centers



Note: The figure shows neighborhoods within the circumference of the metropolitan area of Oslo. Color from red to blue indicates in increasing order the distance to the city centers measured by the shortest path. On the left, city centers are defined by café density; on the right based on port locations. Gray borders indicate neighborhoods with urban development.

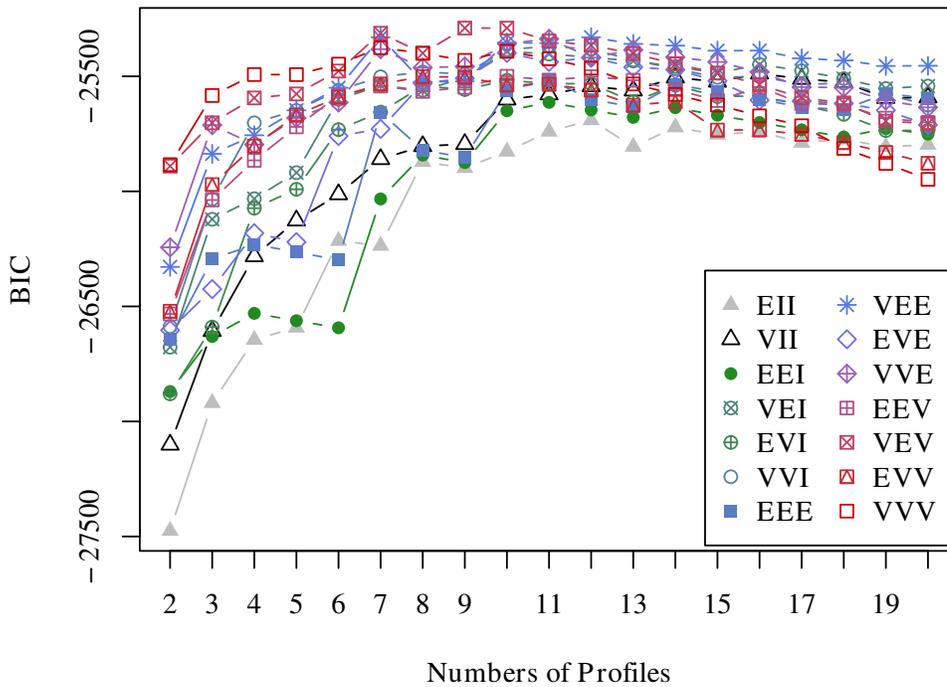
A.3 Latent Profile Analysis

Here we provide more detail on the latent profile analysis (LPA), which we use to algorithmically categorize neighborhoods according to their density components. It assumes that there are unobserved latent profiles (i.e. clusters of individual observations) that generate patterns of outcomes on various variables, here the three density components. Using (quasi-)maximum likelihood techniques, LPA tries to uncover these latent profiles as clusters (Hagenaars and McCutcheon, 2002, Hancock and Samuelsen, 2008).

LPA belongs to the branch of finite Gaussian mixture models (GMM). Latent Profile Analysis is conceptually very similar to Latent Class Analysis (LCA), with the only difference being that the outcome variable - here the magnitude of the density components - is continuous rather than binary. Also, LPA with GMMs can be considered a generalization of the often-used K-means clustering. K-means assumes spherical profiles due to its reliance on Euclidean distance. GMMs, however, model each profile with a Gaussian distribution, which includes parameters for both the mean (similar to the centroid in K-means) and the covariance (determining the spread and orientation of the cluster). They are more flexible in supporting other shapes than spherical, as well as varying orientations.

We run numerous different GMMs with a varying number of profiles (from 2 to 20), using the R routine *mclust* (Scrucca et al., 2016). All models are fitted by the Expectation-Maximization (EM) algorithm based on the outcomes of the logarithm of the three density components. We compare the Bayesian Information Criteria (BIC) of both the different models and the number of profiles to see which perform best. The highest score is reached by VEV models with 7, 9, and 10 profiles. In Figure A-2, we display the Bayesian Information Criteria (BIC) of the various models.

Figure A-2: BIC of Tested Gaussian Mixture Models

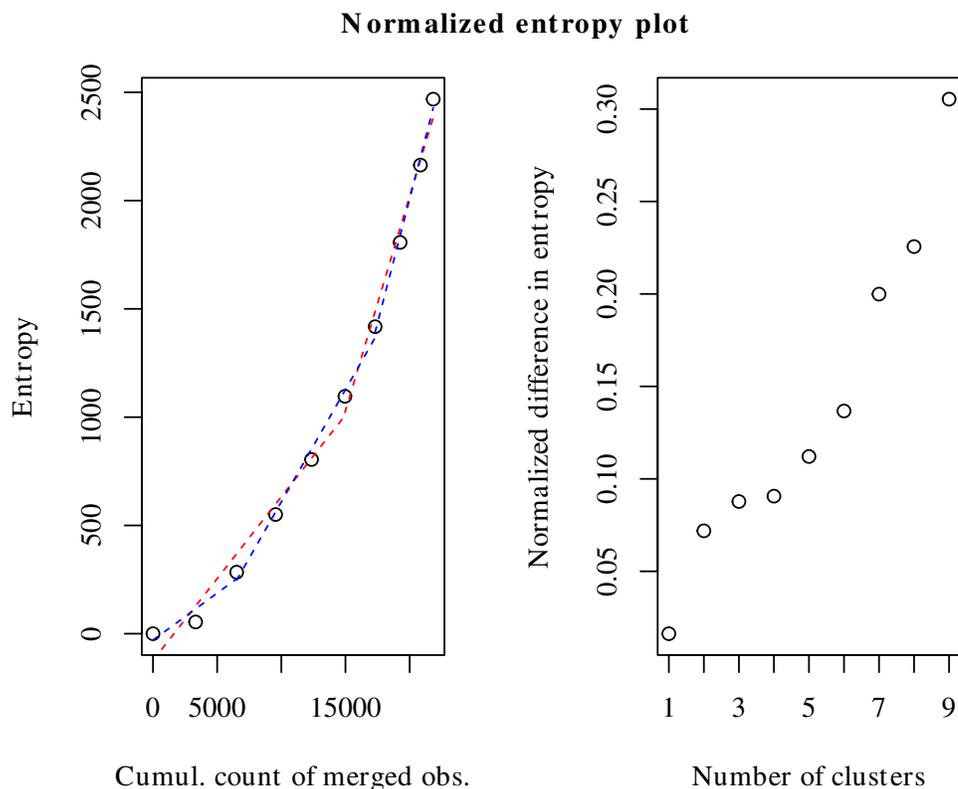


Note: The figure shows the performance of all Gaussian mixture models run by the *mclust* routine in R, using all possible numbers of profiles. All models are fitted with the EM-algorithm based on the neighborhood-level logarithms of building height, crowding, and residential coverage.

The VEV type of Gaussian mixture models allows for **V**arying volume, **E**qual shape (here ellipsoidal) and **V**arying orientation of the latent profiles. The next step is to choose between the models with the 7, 9, and 10 profiles which all perform similarly in terms of BIC. We argue in favor of parsimony, but also consider that the BIC tends to overestimate the number of profiles in case of non-normal variables (such as building

height). For robustness, we analyze the change in entropy associated with varying numbers of profiles. The results in Figure A-3 suggest notable change points at 4 and 7 profiles. Additionally, we conduct a Bootstrap Likelihood Ratio Test and find that increasing the number of profiles beyond 7 does not significantly increase the model fit, see Table A-2. We consequently settle on 7 profiles.

Figure A-3: Entropy of Tested VEV models



Note: The left plot shows the entropy associated with the number of profiles. The red line corresponds to the estimate of a piece-wise linear regression with one change point. A kink at 6 profiles is discernible. The blue line corresponds to the estimate of a piece-wise linear regression with two change points. One kink at 4 profiles and one at 7 can be seen. The right panel shows the normalised change in entropy when going down step-wise from 10 to 2 profiles. A notable jump occurs at 7 profiles.

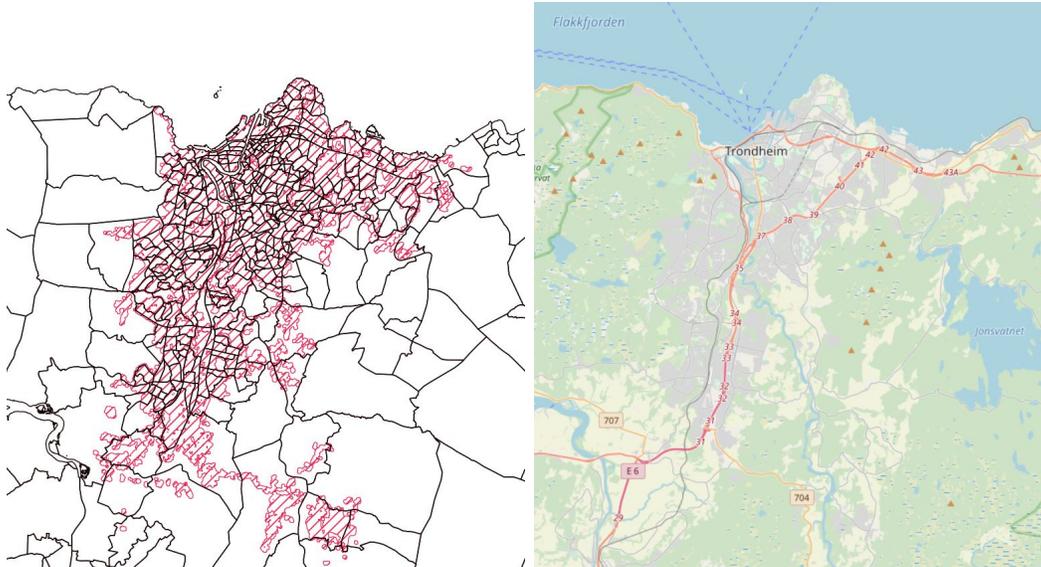
Table A-2: Bootstrap Sequential Likelihood Ratio Test for the Number of Latent Profiles

	LRTS	bootstrap p-value
1 vs 2	2009.16966	0.001
2 vs 3	253.28442	0.001
3 vs 4	172.84890	0.001
4 vs 5	82.36948	0.001
5 vs 6	164.38859	0.001
6 vs 7	229.25799	0.001
7 vs 8	-21.83383	0.674

Note: The model used is VEV and the number of replications is 999. We see the model fit increase, when the number of profiles increases until 7, but then a decrease when 8 profiles are used.

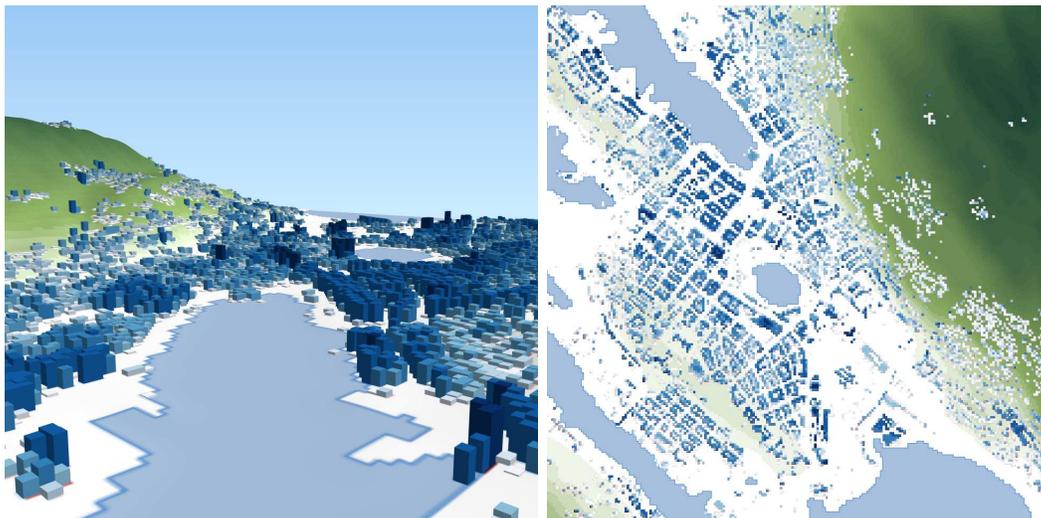
A.4 Supplementary Tables and Figures

Figure A-4: Neighborhoods in Trondheim



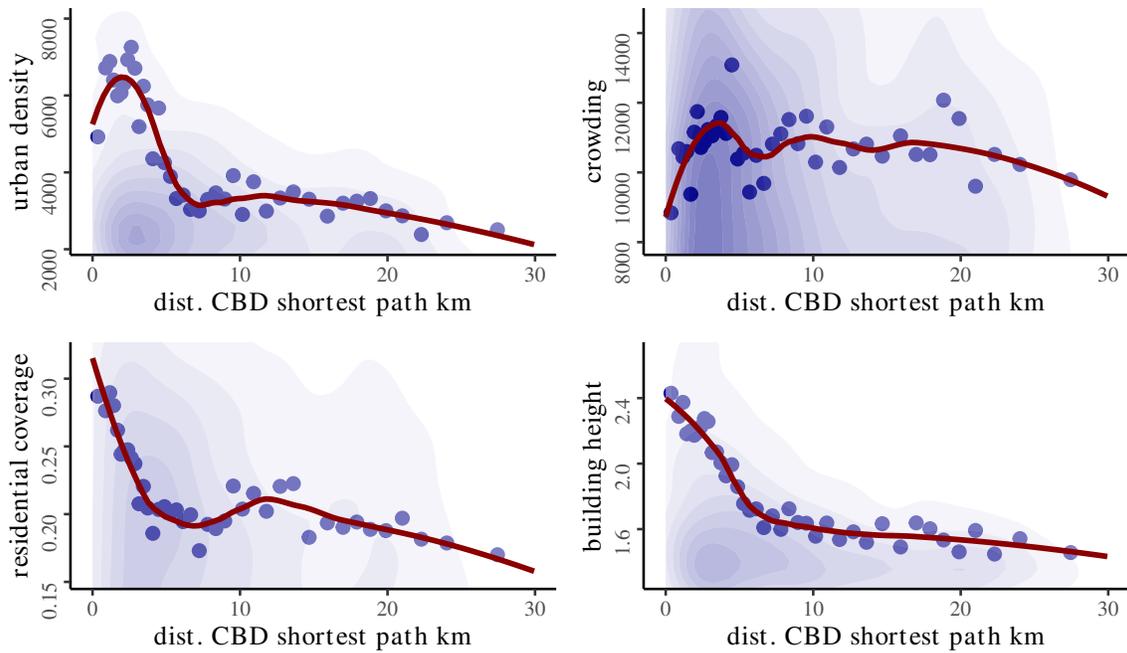
Note: The figure shows the *grunnkrets* borders in black, and the urban residential built-up areas in red (on the left), compared to the area of Trondheim in the OpenStreetMap project (on the right).

Figure A-5: Building Height and Footprint in Bergen



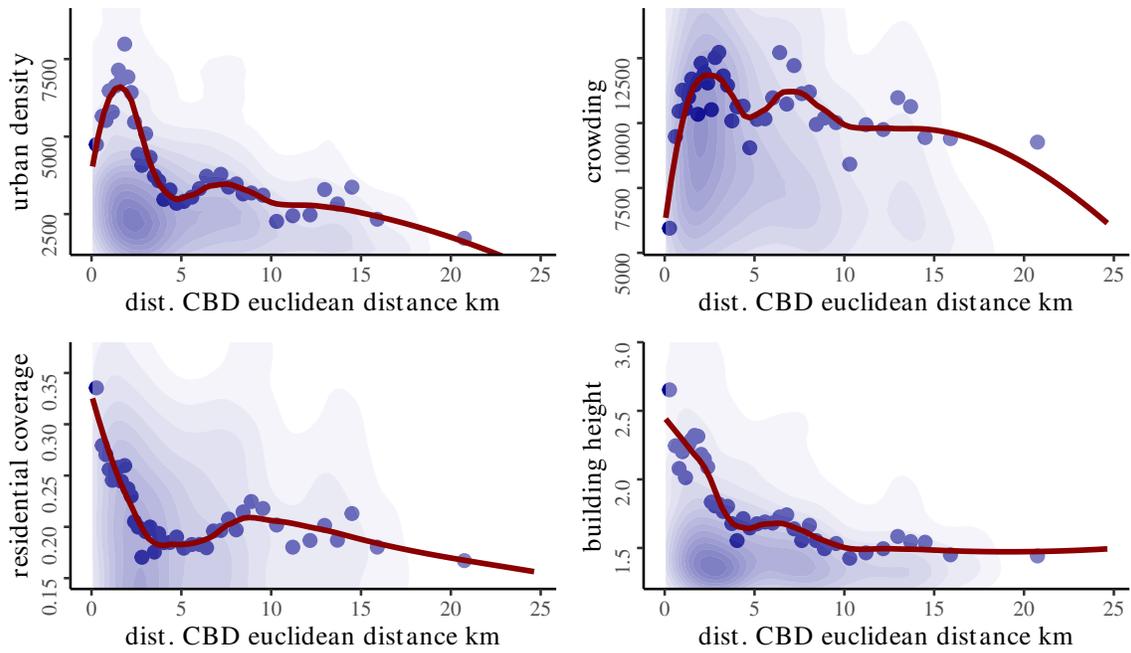
Note: The figure illustrates our approach of constructing our data set involving building heights and footprints. It shows the 3D view of the old port of Bergen (Brugten) from the sea (on the left), and a “bird’s-eye view” of the city center (on the right). In both figures, blue indicates developed areas, with a darker blue indicating higher buildings.

Figure A-6: Distance Gradients (with City Centers Based on Ports)



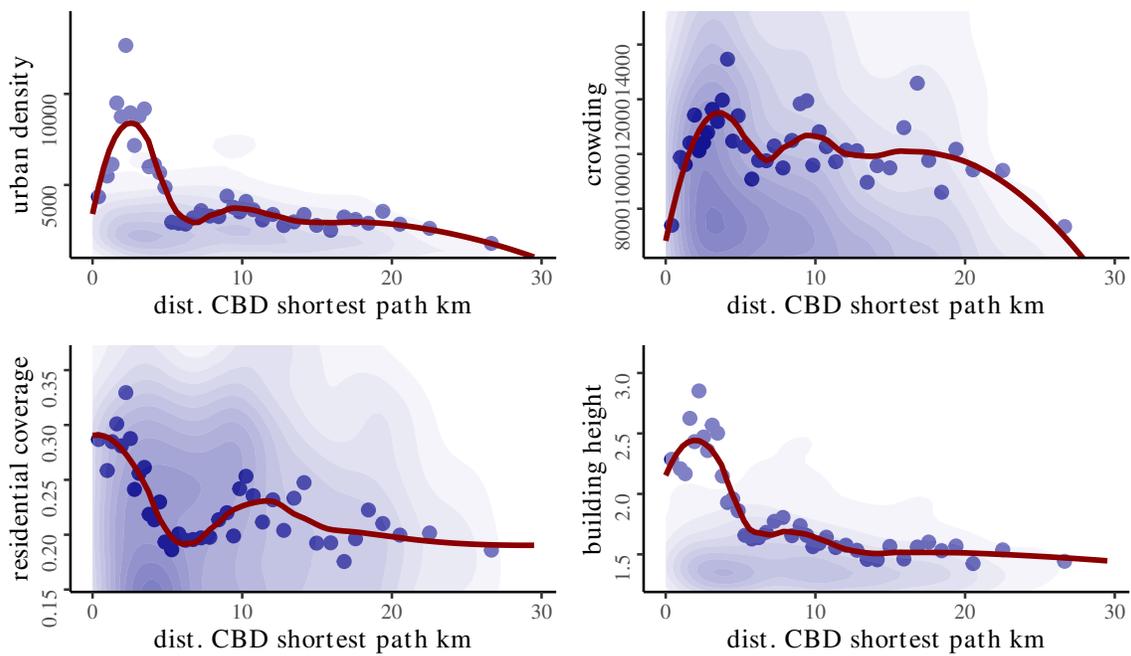
Note: The figure displays a scatter plot with bins of 2.5% (blue dots), a locally estimated smoothed scatter plot with a 50% bandwidth (red line), and a 2D kernel density estimation plot (blue shades), all of which represent urban density and its components plotted against the distance to city center, measured by the shortest path given the terrain. The city center is defined based on port locations.

Figure A-7: Distance Gradients (with Euclidean Distance)



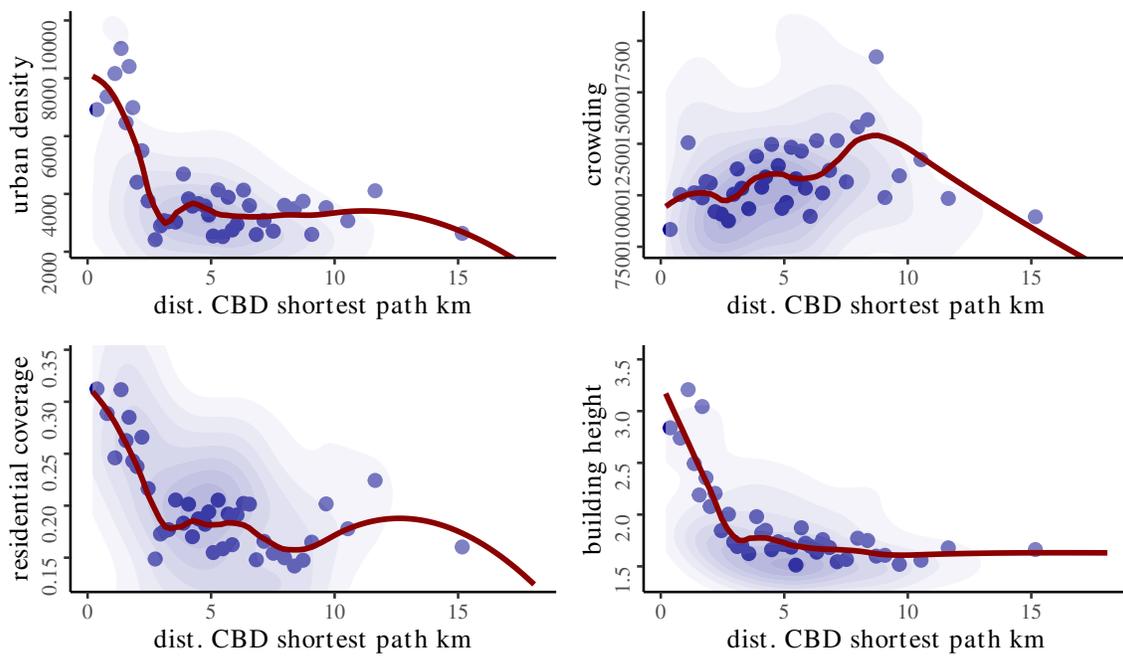
Note: The figure displays a scatter plot with bins of 2.5% (blue dots), a locally estimated smoothed scatter plot with a 50% bandwidth (red line), and a 2D kernel density estimation plot (blue shades), all of which represent urban density and its components plotted against distance to the city center (measured as Euclidean distance).

Figure A-8: Distance Gradients for all Neighborhoods in Oslo



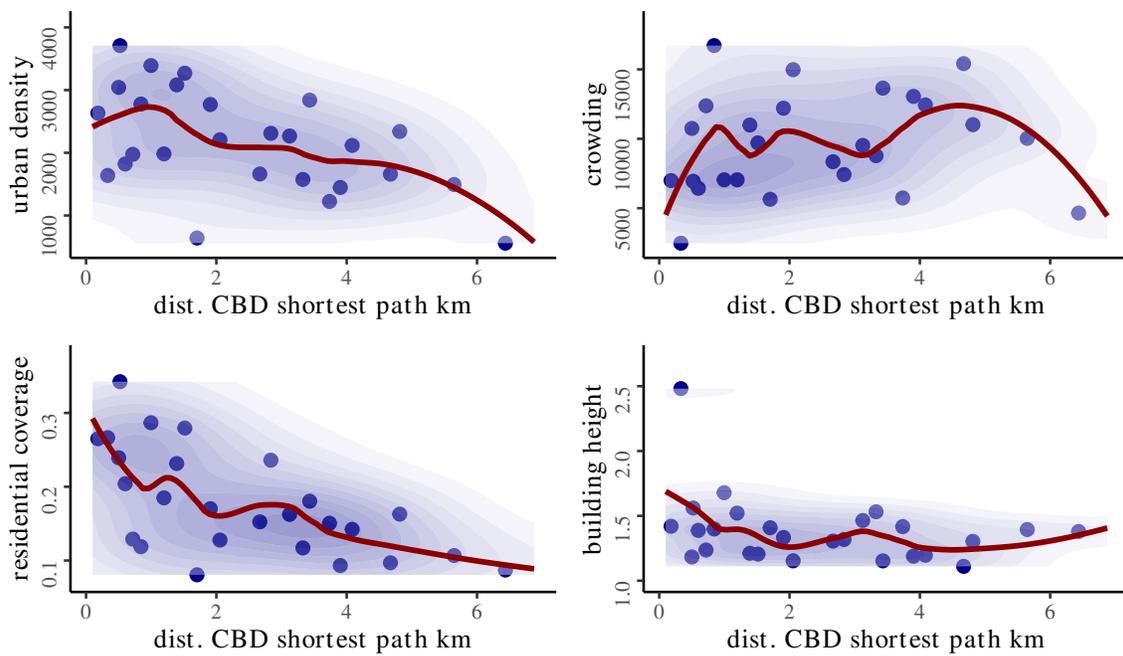
Note: The figure displays a scatter plot with bins of 2.5% (blue dots), a locally estimated smoothed scatter plot with a 50% bandwidth (red line), and a 2D kernel density estimation plot (blue shades), all of which represent urban density and its components plotted against the distance to the city center of Oslo, measured by the shortest path given the terrain.

Figure A-9: Distance Gradients for all Neighborhoods in Trondheim



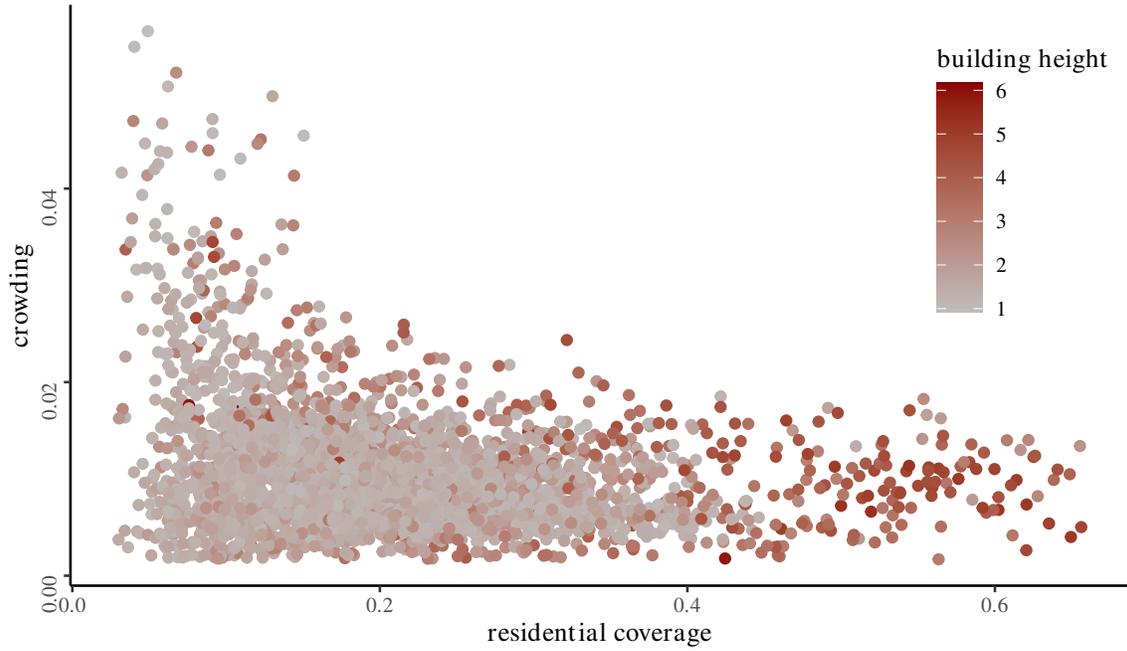
Note: The figure displays a scatter plot with bins of 2.5% (blue dots), a locally estimated smoothed scatter plot with a 50% bandwidth (red line), and a 2D kernel density estimation plot (blue shades), all of which represent urban density and its components plotted against the distance to the city center of Trondheim, measured by the shortest path given the terrain.

Figure A-10: Distance Gradients for all Neighborhoods in Molde



Note: The figure displays a scatter plot with bins of 2.5% (blue dots), a locally estimated smoothed scatter plot with a 50% bandwidth (red line), and a 2D kernel density estimation plot (blue shades), all of which represent urban density and its components plotted against the distance to the city center of Molde, measured by the shortest path given the terrain.

Figure A-11: Scatter Plot of the Three Components



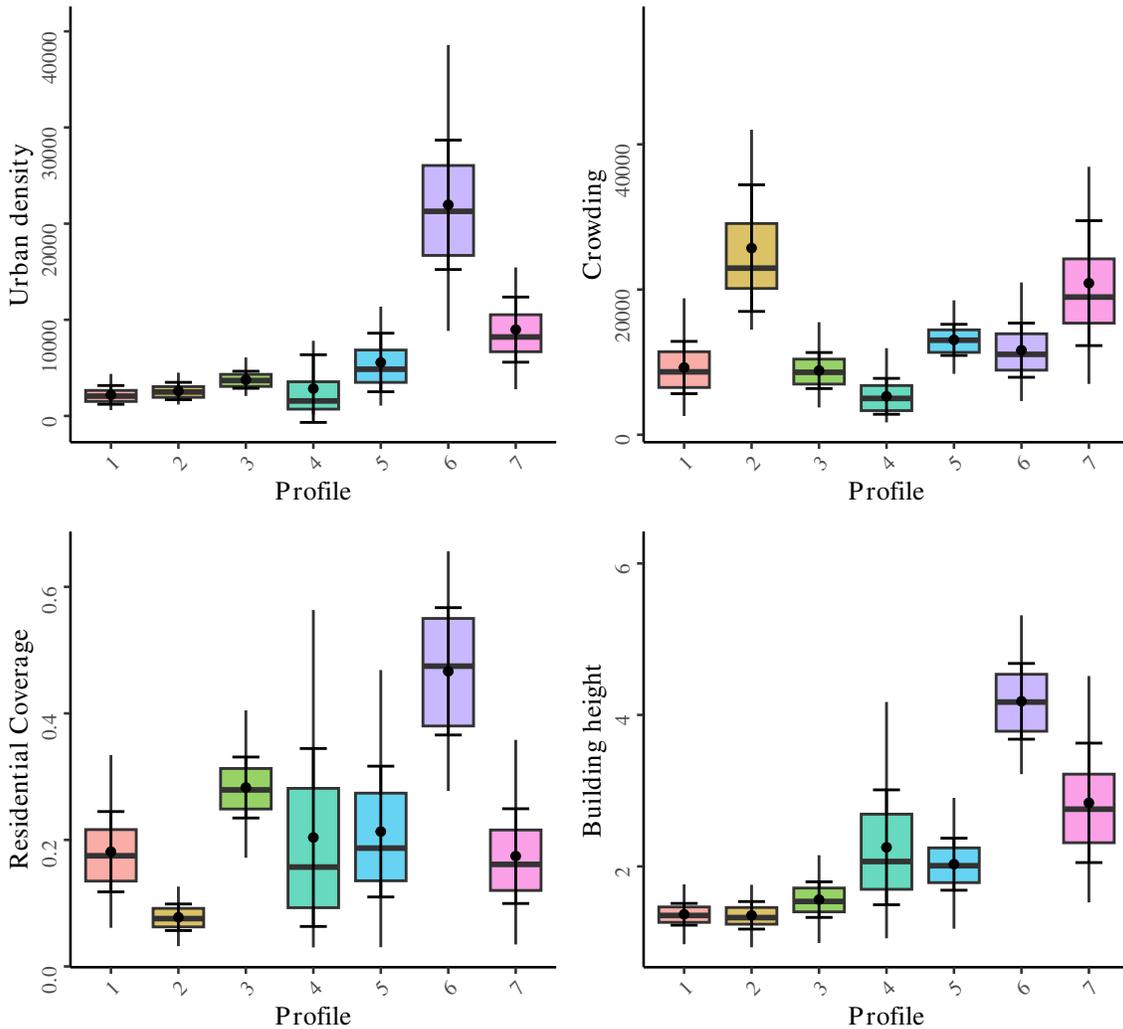
Note: The figure displays the bin scatter plot of residential coverage and crowding as well as building height (indicated by darker shades of red).

Table A-3: Number of Neighborhoods by Profile and Urban Area

Profile	1	2	3	4	5	6	7
Bergen	134	4	40	66	73	6	18
Bodø	25	0	6	16	3	0	0
Hamar	36	6	3	23	6	0	0
Haugesund	16	35	2	15	17	0	7
Kristiansand	89	10	16	24	10	0	1
Kristiansund	11	10	1	5	5	0	2
Lillehammer	9	1	2	5	0	0	0
Molde	20	0	3	3	0	0	0
Oslo	842	52	363	208	148	123	155
Stavanger	112	1	51	41	38	0	6
Tromsø	15	6	0	7	4	0	2
Trondheim	102	1	71	43	60	4	29
Ålesund	20	14	1	11	8	0	0

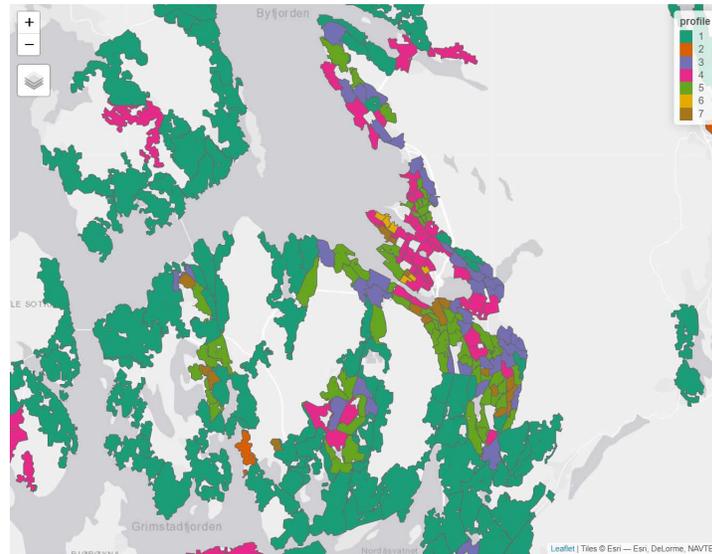
Note: The table shows the number of neighborhoods per urban area that fall in each of the 7 profiles. More information on the density characteristics of each profile is provided by [Table 3](#).

Figure A-12: Box Plots of the Within-Profile Variation of Neighborhood Density



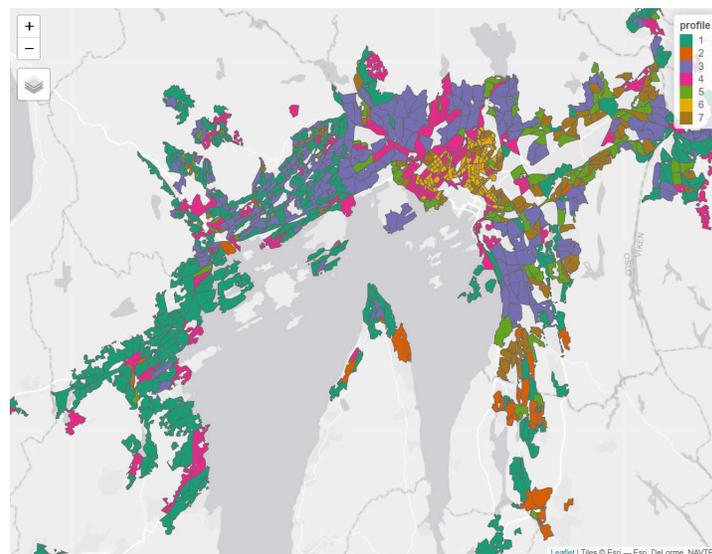
Note: This figure presents the distribution of urban density or its components across different profiles, visualized through box plots. The boxplot displays data spread by profiles, emphasizing the middle 50% of data between the first (25th percentile) and third quartiles (75th percentile), with a line for the median, indicating the midpoint. Whiskers show the range, reaching up to the highest and lowest values. Black points mark the mean and black bars the standard deviation for each profile.

Figure A-13: Latent Profiles in Bergen



Note: The figure shows a map of Bergen with the neighborhoods as classified by the latent profiles, see [Table 3](#) for more details on the profiles.

Figure A-14: Latent Profiles Oslo



Note: The figure shows a map of Oslo with the neighborhoods as classified by the latent profiles, see [Table 3](#) for more details on the profiles.

Table A-4: Neighborhoods with the Lowest Latent Profiles Uncertainty

Profiles	Uncertainty	Name	Latitude	Longitude
1	0.0091	Voiebyen-Nordvest (Kristiansand)	58.1085	7.9532
1	0.0094	Hosle Sør 04 (Oslo)	59.9303	10.5797
1	0.0096	Rekkevik (Oslo)	59.0254	10.0720
1	0.0096	Nordeide (Bergen)	60.3171	5.2799
2	0.0117	Tangen-Åskollen 13 (Oslo)	59.7113	10.2548
2	0.0118	Fuglevik 3 (Oslo)	59.1950	10.9336
2	0.0125	Mosbron (Haugesund)	59.2529	5.2057
2	0.0143	Raglatua (Haugesund)	59.3946	5.3176
3	0.0820	Nygård (Oslo)	59.9491	10.7553
3	0.0833	Solslett (Bodø)	67.2877	14.4094
3	0.0858	Nyhavn (Bergen)	60.4194	5.3083
3	0.0891	Søndre Hellerud (Oslo)	59.8705	10.8062
4	0.0000	Bragernes Sentrum 6 (Oslo)	59.7482	10.1997
4	0.0000	Mart'nsplassen (Hamar)	60.7953	11.0844
4	0.0000	Osnes Nedre (Ålesund)	62.3438	5.8299
4	0.0000	Gaustad (Oslo)	59.9522	10.7185
5	0.0315	Filipstad (Oslo)	59.9090	10.7153
5	0.0500	Foldal (Trondheim)	63.3868	10.4161
5	0.0653	Mulen (Bergen)	60.4049	5.3288
5	0.0758	Steinkjelleren (Bergen)	60.3980	5.3278
6	0.0070	Tøyen rode 2 (Oslo)	59.9189	10.7642
6	0.0092	Gamle Aker rode 4 (Oslo)	59.9267	10.7489
6	0.0093	Tøyen rode 2 (Oslo)	59.9146	10.7708
6	0.0098	Grønland rode 4 (Oslo)	59.9154	10.7646
7	0.0000	Industriområdet 1 (Stavanger)	58.8503	5.7396
7	0.0000	Danvik-Fjell 14 (Oslo)	59.7184	10.2281
7	0.0000	Hagaløkka (Oslo)	59.8308	10.4268
7	0.0000	Rosenborg 18 (Trondheim)	63.4353	10.4145

Note: For each of the seven latent density profiles, the table shows the five neighborhoods from the country-wide sample that have the lowest uncertainty of belonging to the profile. It indicates the uncertainty, the identification number of the corresponding *grunnkrets*, the name of the neighborhood with its urban area in brackets, as well as its latitude and longitude.

Table A-5: Elasticities of Within-City Density and Its Components with Socio-Economic Outcomes and Interactions

	(1)	(2)	(3)	(4)	(5)
Depend.Var:	mean income	income ineq.	age mean	kid share	sick notes mean
<i>Panel C: The components of urban density with interactions</i>					
ln(residential cover.)	-0.70*** (0.08)	-1.28*** (0.26)	-10.16*** (2.51)	0.11*** (0.03)	0.96*** (0.15)
ln(building height)	0.68*** (0.23)	1.04* (0.57)	6.32 (5.83)	-0.25*** (0.07)	-1.82*** (0.31)
ln(crowding)	-0.43*** (0.04)	-0.78*** (0.12)	-5.08*** (1.09)	0.05*** (0.01)	0.57*** (0.06)
ln(residential cover.) × ln(building height)	0.33** (0.13)	0.31 (0.32)	4.82 (3.58)	-0.10*** (0.04)	-0.85*** (0.19)
ln(residential cover.) × ln(crowding)	-0.15*** (0.02)	-0.26*** (0.06)	-2.59*** (0.54)	0.02*** (0.01)	0.22*** (0.03)
ln(building height) × ln(crowding)	0.17*** (0.05)	0.23* (0.12)	2.21* (1.26)	-0.02* (0.01)	-0.36*** (0.06)
triple interaction	0.05* (0.03)	0.03 (0.07)	2.11** (0.82)	-0.02** (0.01)	-0.16*** (0.04)
R ²	0.50	0.13	0.17	0.43	0.30
Fixed-effects	Yes	Yes	Yes	Yes	Yes
Observations	3,322	3,322	3,322	3,322	3,322

Note: The table complements Table 4 by reporting regression results of socio-economic outcome variables on its individual components including interactions. For readability, log density measures are given per square meter. *mean income* is the log of income per capita, *income ineq.* is income inequality captured by the coefficient of variation of income, *age mean* is average age within a neighborhood, *kids share* denotes the share of children and teenagers under 18, and *sick notes* the average number of sick notes. Standard errors are clustered at the *kommune* (municipality) level, see Section 3.1 for detail. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.