Inference of Building Archetypes and Schedules from City of Victoria Open Data

Garrett E. S. Therrien[1], Theo Christiaanse [1], Ralph Evins [1]
[1]University of Victoria, Victoria, Canada

Abstract

Building-by-building modelling of districts using existing archetypes has become common practice for creating detailed district models. But, this technique scales poorly: detailed data on each building in a large area is usually non-existent or not publicly available. However, there are open data sources, such as zoning and building footprint area, commonly available from cities and municipalities that allow the assignment of archetypes. This paper presents a method using municipality open data sets and the Building Technology Assessment Program (BTAP) archetypes to create an urban district building stock model, enabling detailed analysis of a city's building stock. The method developed will be demonstrated on City of Victoria open data, and the model will be verified by comparing daily electricity use.

Introduction

Urban building energy modelling (UBEM), or modelling of the building stock of city or other urbanized region, is a generally underdeveloped field of study inside building modelling. Yet, with the growing challenges of climate change and the matching growth cities' climate change policies, it is a critical field, able to provide guidance on the best climate change mitigation policies.

As late as 2016 UBEM was referred to as a 'nascent field' by Reinhart and Cerezo Davila (2016) who, in their review of UBEM, laid out what remains a typical, data-intensive process for building out a city building stock model. In essence, this typical process relies on having detailed building data (including hourly energy use) for a large segment of the buildings in a city, clustering these by some process (in the beginning it was often done by manual judgement; now machine learning tools are more common, as shown in Deru et al. (2011); Reinhart and Cerezo Davila (2016)) to create data-driven archetypes, and filling in any gaps with constructed archetypes (archetypes built from the building codes, such as the U.S. DOE archetypes). This process also shows why so little study has been given to UBEM: the data sets simply do not exist in the vast majority of cases, especially in North America. Europe has several data projects which have provided material for case-studies of UBEM; unfortunately these are not generalisable as they are tied to the datasets used. (Ali et al., 2019; Cerezo et al., 2017; Reinhart and Davila, 2019; Swan et al., 2009; Pasichnyi et al., 2019; Deru et al., 2011; Mata et al., 2014)

However, in the last few years there have been advances in parallel fields in computer science, providing a suite of machine learning concepts and tools for inferring and tuning models from very little data. These tools are just beginning to be applied to UBEM for the construction of stock models, usually by calibrating an existing archetype set. Notably, Nagpal et al. (2019) and Nagpal and Reinhart (2018) developed a calibration technique that works on city blocks or large buildings. Tardioli et al. (2020) recently took this work further, doing work with building location in GIS as well as using Bayesian inference techniques to improve data-driven archetypes. Unfortunately, these techniques use a large amount of data as their initial base — in essence what has been developed is a better way of filling in the gaps and building out a data-driven model and not a method for working in a data-sparse environment and filling in the gaps in constructed archetypes. (Ali et al., 2019; Cerezo et al., 2017; Reinhart and Davila, 2019; Nagpal and Reinhart, 2018)

Some of the most recent work in this field, though, carried out by Roth et al. (2020) takes the calibration ideas in Nagpal et al. (2019) to the next logical step, using them to refine a mix of constructed and data driven archetypes. This work shows it is possible to take a mix of highly specific building data for only part of city, general archetypes, and a relatively data-sparse city with open data and calibrate a UBEM stock model from this data. While their paper relied on having a rich data source for validation and cross-checks, now that the concept and technique has been proved, it is of interest how much of such data is necessary — the less data needed, the more generally applicable the technique. Furthermore, because they were uncertain if the concept would work, they used a multi-faceted method to test several different ways of performing calibration; this process, while effective, could prove a bar to wide spread deployment. (Nagpal and Reinhart, 2018; Nagpal et al., 2019; Tardioli et al., 2020)

And wide spread deployment is highly desirable — detailed may provide municipalities with the tools to develop targeted evidence based policies to lower carbon emissions. But, because such data sets are the result of years of work, often contain confidential information, and, if developed by a company, are usually proprietary, these detailed building stock models are rare, forcing cities to use rough estimates (such as multiplying the square footage of building types by the relevant archetypes, refered to later in this paper as the 'scaled' method) to predict the impact of their building policies. (Nagpal and Reinhart, 2018; Reinhart and Cerezo Davila, 2016; Swan et al., 2009; Cerezo et al., 2017; Mata et al., 2014)

Fortunately, most cities can obtain basic building information (building heights, footprint areas, and general uses) that could form the base of a more generic archetype-based model. Most of the methods used in Roth et al. (2020) can be separated from the data-rich baseline; in place of data-driven archetypes a more generic archetype can be used. City-scale aggregate electricity use data (for tuning) is relatively easy to request, as it offers no privacy issues in the way building-by-building data would. (Swan et al., 2009; Nagpal and Reinhart, 2018; Nagpal et al., 2019; Kristensen et al., 2018; Deru et al., 2011; CanmetENERGY, 2020)

Therefore, on the basis of this past work, this paper seeks to prototype a simpler, more portable method of archetype calibration for an UBEM stock model. This method we propose uses a minimum of specialised data: data used will be limited to open and common data, such as municipal GIS data and easily requested aggregate electricity use. Because of this, it will focus on 'bottom-up' methods, starting with assumptions and refining on the basis of data, rather than the 'top down' method more generally employed, starting with data and using assumptions to fill in the gaps. Rather than attempting to build something perfect, this paper will focus on 'good enough', striving for simple methods that are reasonably accurate. Most municipalities are facing difficult decisions with no tools, and this paper seeks a method to fill in this gap, to find the minimum amount of data and model parameters that can be used to develop a UBEM stock model.

Methods

The methods deployed in this paper comprise three distinct stages:

1. Build: Create archetypes: select base archetypes and input data in preparation for the next two stages.
2. Fit: Generate building stock options, map them to the archetypes created in phase 1, and select of parameters for tuning and fitting a surrogate model to each archetype to speed up computation of energy use in the genetic algorithm.
3. Tune: Use a genetic algorithm to tune the parameter values selected in 2 to fit the real-world time sequence.

The goal of this process is a set of parameters for the input archetypes which create a daily energy time series which matches the actual time series from BC Hydro.

Input Data Selection

For this study, we use various different data made available by the City of Victoria and BC Hydro. The data sets can be put into two categories. The first is building size, use, and classification from the City of Victoria. The second is the total electricity use of the downtown core in the form of hourly time series data.

The first set of data was taken from City of Victoria open data portal (City of Victoria, 2020) , which has a tremendous amount of information. The data set of interest to us is the Building Rooflines 2019 data set (the Rooflines dataset), with information on the building use (allowing sorting into archetypes), and building size and shape (allowing scaling of energy use).

The second set of data — energy use data — was requested from BC Hydro. This data is sufficiently aggregated that it was easy to request — there are no privacy issues with electricity use data from a large area, as it is impossible to de-aggregate to individual buildings. It contains aggregated hourly data of every commercial building in the City of Victoria.

The last input needed is the archetypes. These were selected from the Building Technology Assessment Program (BTAP) (CanmetENERGY, 2020), a Canadian version of the U. S. Department of Energy (DOE) standard reference buildings produced by Natural Resources Canada (NRCAN). To these archetypes a 'No Energy' archetype was added, having no energy use at any time, for vacant buildings and parking lots. (Deru et al., 2011)

For this study, the BTAP 2011 National Energy Building Code (NEBC 2011) implementation was used as base buildings, as they represented a compromise position: while there are many older buildings in Victoria's downtown core, 2011 was the earliest

code implemented in BTAP. These archetypes were assigned to the buildings in the downtown core using the 'Actual Use' given in the Rooflines data set, and, in the case of buildings with multiple sizes, storeys. (CanmetENERGY, 2020)
As the BTAP archetypes are commercial only, this study is limited to commercial data only.

Generation of the Building Stock Model

To generate the building stock we will translate the building stock data extracted from the Rooflines data set to BTAP archetypes. First of all, each actual use type — the City's classification of the building by its use — is translated into a building type. Then these building types are assigned BTAP archetypes. While this may seem redundant, as several building types may have the same BTAP archetype, in later tuning steps these building types will have unique deviations from the base BTAP archetype.
This division into building types and assignment into archetypes is done by engineering judgement. Where there was a need to determine if a building was small, medium, or large, the number of storeys was used, such as for office buildings.
There are two obvious ways to assign the buildings into building types, and both were tried: straight assignment into archetypes (where BTAP archetypes plus the "no energy" archetype become building types, so there are only 13 building types matching the 13 base archetypes), and straight assignment of building types by actual use type (where the City's actual use type categories become the building types.)
After refinement of parameters and the addition of light and plug load, a third, simplified, assignment was tried, between these two extremes, both to reduce the number of variables in the GA and to better reflect reality. For example there were several categories of small office in the actual use type set, ranging from a generic small office to a government building, that were the same type of building type — a small office building used for desk work. Such building types could be profitably consolidated into one and were in this simplified actual use type run.
As the BTAP models are electric heat, the percent of buildings using gas heat was located in the NRCan statistics for building stock in Canada, and overall energy use was reduced by this amount, 76.7%. (Government of Canada, 2020)

Parameterisation and Energy Use

Parameters were decided upon based on previous research into archetypes and tuning (Swan et al., 2009; Reinhart and Cerezo Davila, 2016; Ali et al., 2019; Cerezo et al., 2017; Reinhart and Davila, 2019; Swan et al., 2009; Pasichnyi et al., 2019; Deru et al., 2011; Mata et al., 2014), with the parameters that had the consistently highest impact on energy use used:

- Heating Set Point – Overall heating set point. In later runs this was changed for weekday only, during the day.
- Cooling Set Point – Overall cooling set point. In later runs this was changed for weekday only, during the day.
- Insulation Value – value of the wall insulating materials in RSI.
- Light and Plug Load – average sum of all lighting and plug loads in the building, on a schedule determined by the archetype.

It is noted that these parameters oversimplify; often archetypes have, for instance, different set points for weekend and weekday which are not taken into account in the initial or final heating and cooling set point setup. However, tuning all set points often leads to an excessive number of variables being tuned; in the case of the actual use type tuning it is in the thousands of variables (69 actual use types, some with as many as 40 variables) so simplification must be done.
Light and plug load was added after runs the first set of runs, in an attempt to further improved the archetypes.
Once these parameters were decided, it was relatively easy to parameterise the archetypes and determine energy use. However, as EnergyPlus (EP) runs are quite slow, and each run of the genetic algorithm will need many simulation runs, it was decided to train surrogates in place of EP for use in the genetic algorithm. Each surrogate model is based on a regression neural network method. All surrogate models were constructed using TensorFlow with input parameters proportional to the building type division and 365 output parameters with a learning rate of 0.001 using 4000 epochs. The model additionally has four hidden layers with the first having twice the size of the input parameter nodes, the second 183 nodes, and the third and fourth 365. The inputs and outputs are normalized using the training data, which consists of a subset 80 percent of 50 EP runs. The input parameters for these EP runs are based on random set points produced by Latin hypercube sampling, while the outputs are a year's electricity use time series data. The surrogate model showed high R2 values on the test data in the range of 0.99, due to the simplicity of the model.

Genetic Algorithm Tuning

To tune the buildings, a genetic algorithm was used. Time series energy use for the city, as predicted by the input parameters, a list of the buildings in the city, and the building type models, was compared to the actual time series energy use of the city using cumulative RMSE. This error was minimised, as it

Run Name	Building Type Division	Weekday Set Points Only ((Y)es or (N)o)	Lights and Plugs ((Y)es or (N)o
Archetype	Archetype	N	N
Actual Use Type	actual use type	N	N
Four Parameter Archetype	Archetype	Y	Y
Four Parameter Actual Use Type	actual use type	Y	Y
SAU	Simplified Actual Use Type	Y	Y

Table 1: Setup for Each Genetic Algorithm Run

was desired to have as little difference between the predicted and actual energy time series.
See table 1 for a breakdown of which parameters were used by which runs. Note, this table does not show the scaled-only run as it was not put through the GA. Common values to all runs are:

- Heating Set Point Range: 5-15°C
- Cooling Set Point Range: 25-35°C
- Insulation Value Range: 0-10°C
- Light and Plug Load (for runs used, see table 1): 5-30 W/m^2
- Genetic Algorithm Population: 20
- Genetic Algorithm Evaluations: 600
- Number of GA runs: 20

Results Interpretation and Error Calculation

To gain a larger sample of parameters and prove convergence, the GA will be run 20 times and the mean value for each parameter calculated. This mean value will be used to predict a final daily energy time series. Errors will be quantified using the root mean square error (RMSE), showing how far the predicted series deviates from the actual (i.e. the spread of the residuals) and the mean bias of the predictions (showing over or under prediction.) Additionally, scatter plots will be plotted of predicted vs actual results, allowing easy visualisation and comparison of the over and under prediction of each run.

Results

Results of all runs are in fig. 1, which shows the predicted energy uses and the actual energy uses. Errors fig. 3, which shows the mean root mean square error (RMSE) (showing how far the predicted series deviates from the actual, i.e. the spread of the residuals) and the mean bias of the predictions (showing over or under prediction) (ideally both of these are zero), and in fig. 2 which shows the SAU time series, which was the best time series calibrated. Convergence for the SAU run is shown in fig. 6.
To further quantify over- and under-prediction, fig. 4 was plotted, showing the predicted vs actual values. The red line in each plot is the line of equality; ideally the scatter plot would be concentrated along this line. The rows are method (straight scaling, three parameter, and four parameter, from top to bottom), and the rows are the building type division employed (archetype, actual use type, and simplified actual use type, from left to right.)
The parameter values for the four parameter actual use and the SAU runs are compared in fig. 5; these are the two best runs and have very similar output time series and identical errors, making a detailed comparison of their output desirable.

Discussion

While precise results were not arrived at, several useful conclusions can be drawn. The process was started with with the baseline case of straight scaling of archetypes, then progressed through tuning archetypes and building types, with steady improvements. A final improvement was obtained by refining parameters, and the interpretation of results was improved by creating the simplified actual use building type. The base case of scaling can be seen in fig. 1 (the orange line; the blue line is the actual energy use.) It can be seen this technique is sub-par; it vastly under predicts, though it does maintain a reasonable weekday/weekend relationship. Therefore, you cannot just scale the energy use by floor area and buildings using electric heat, as was done to generate the orange line, but rather need to adjust the actual underlying values.
The next stage was to tune one building type per archetypes for all the buildings. i.e. splitting the city up into 13 building types directly corresponding to the archetypes (12 of these are BTAP archetypes, 1 is the 'No Energy' Archetype.) This is the green line fig. 1. As shown in fig. 3, this is extremely inaccurate, having both high bias and high RMSE values, due to this method's over prediction of energy use, as shown in fig. 4.
To improve the results, the City of Victoria Roofline data set actual use types were used as building types, creating 41 types; note six of these have a base 'No Energy' archetype, and will have an energy use of zero. Each building type was assigned a base archetype, but each building type was tuned individually, giving the GA 105 parameters to tune (compared to the 36 used in the archetype division.) Both these parameter sets exclude the 'No Energy' archetype building types (whose energy use is zero).
The results of this actual use building type division is much better than the straight archetypes; as shown fig. 3 and in fig. 1 the time series (the red line) is much closer than one for straight archetype scaling. This is due to essentially eliminating over prediction, as shown in fig. 4 and fig. 3; unfortunately, as also shown in these figures there is a pronounced tendency to under predict, mostly on the weekends.
It was hypothesized this was due to issues with the base load, and two changes were made. One, instead of using a universal heating and cooling set point, only the weekday heating and cooling set points were changed; and two, light and plug load was introduced. For the sake of completeness, this was computed using archetype divisions, but (as expected), this

Figure 1: All BTAP Archetype Based Energy Prediction Series

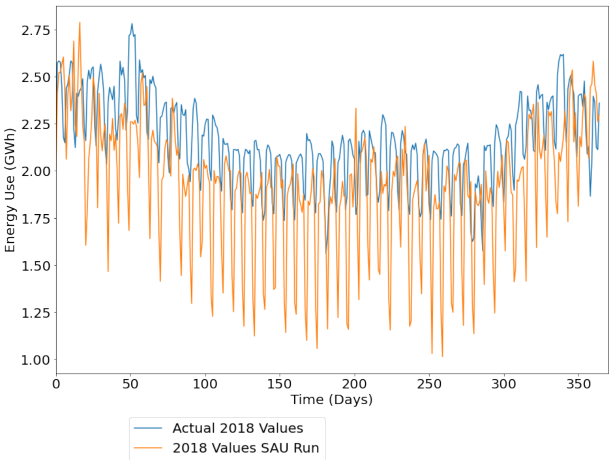

Figure 2: SAU Energy Series

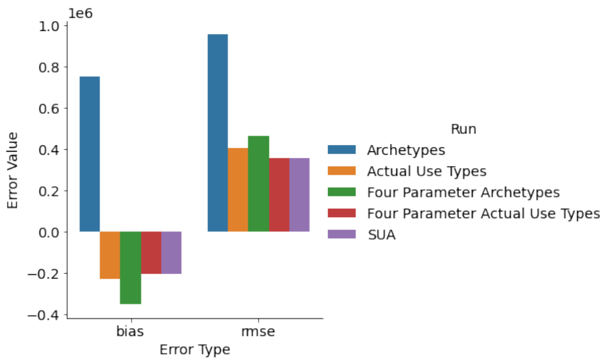

Figure 3: Bias and Root Mean Square Error for BTAP Archetype Based Predictions

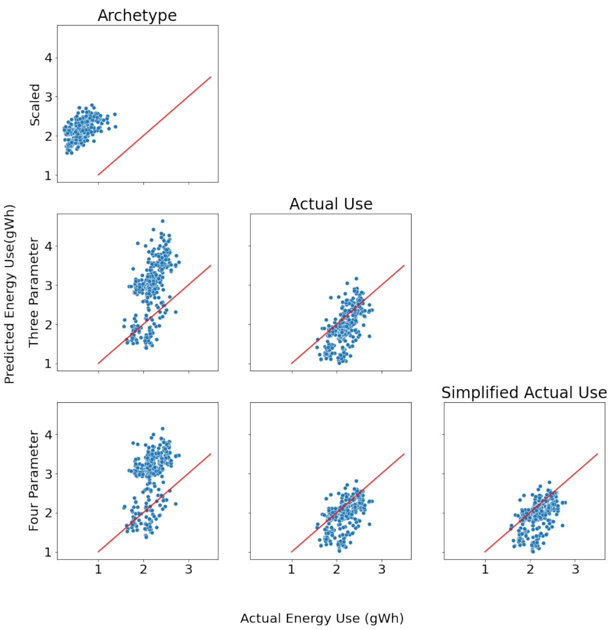

Figure 4: Predicted vs Actual Energy Use

was a poor prediction. However, it is notable that these changes did improve the significantly improve the prediction from the previous, three parameter, archetype prediction — as shown in fig. 4 the over prediction was consolidated, and the errors shown in fig. 3 were much reduced. Unfortunately, the improvements in the actual use type building division prediction using four parameters were more modest; while the predicted vs actual plot in fig. 4 shows a lessening of the under prediction, the error plots in fig. 3 show that the overall improvement was only slight.

Furthermore, this actual-use tuning actually divides the buildings too much, creating a situation where some categories only had a few buildings in them, too few for good tuning by the GA or statistical analysis. Hence, at this stage, the simplified actual use building types were created. This set has fewer building types than actual use type (by, for example, merging various kinds of small stores together), but more categories than pure archetype binning (which had poor results due to having too few building types for the GA to have a sufficient number of variables to tune.) It was possible to reduce the city to 23 building types, of which only one was No Energy, leaving 22 building types for a total of 66 parameters.

The simplified actual use type did not improve the errors as shown in fig. 4 and fig. 3, but as shown in fig. 5 it did improve the overall quality of the results — there are fewer outliers and generally smaller bars with shorter whiskers with simplified actual use type compared to actual use building types, making results from the simplified actual use more precise.

While it was briefly attempted to increase the heating and cooling set point divisions, to allow for a weekend daytime heating set point, the preliminary results were quite bad, showing errors in excess of the simplified actual use runs, and this work was finished there, as it was felt that the limits of the accuracy of the method with the current data had been reached.

Potential improvements in this method are to be found in either more accurate base archetypes (for instance archetypes with gas heating and tuning with a gas time series as well as an electric one), or in using a detailed data set of a subset of buildings (such as Roth et al. (2020) did. However, Roth et al. (2020) is a top-down approach, using assumptions to fill in the gaps, as opposed to our bottom-up approach, which uses data to refine assumptions, and hence is not a method that can be used in a truly data-poor area.) This work does, however, show the potential viability of a bottom up approach, as the opposed to the mostly top-down approach used in previous papers. It also shows the need for at least some detailed building data to flesh out the assumptions and basic, publicly available, data in

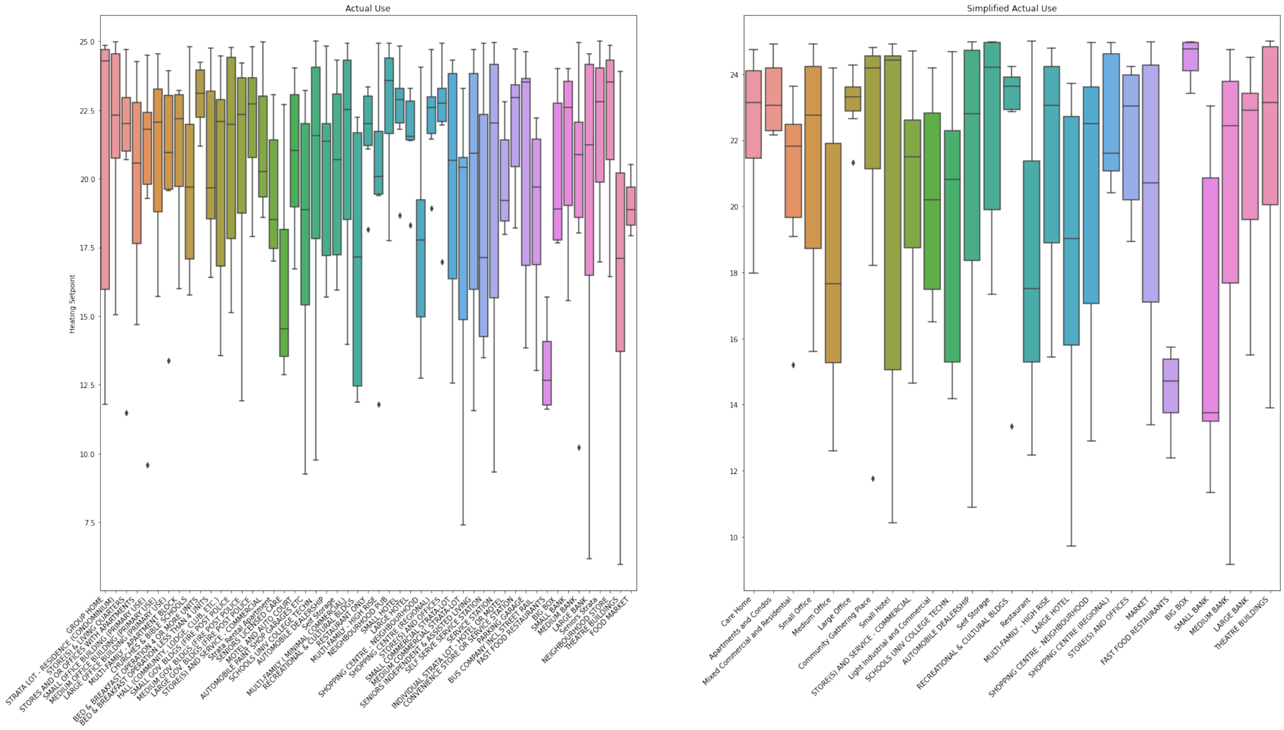

Figure 5: Actual Use Type and SAU Run Variable Results by Parameter for Heating Setpoint

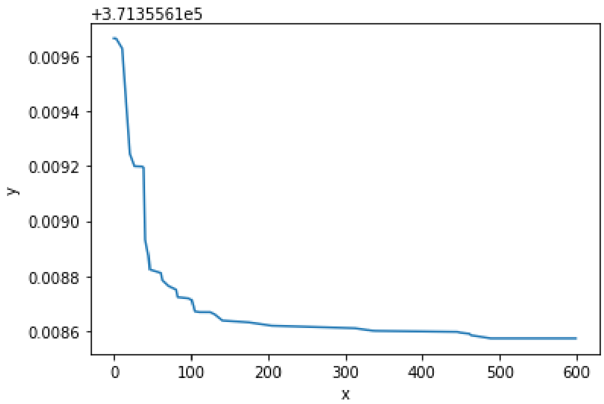

Figure 6: SAU Run Convergence

a bottom-up approach. Determining the amount of detailed data needed to produce an acceptable model in a bottom-up approach is left for future work.

Conclusion

While a straight scaling of generic archetypes is inaccurate, it is possible to create a time series which is close to accurate with very little data and a set of generic archetypes, proving the viability of a bottom-up, data-poor, approach. However, a truly acceptable bottom-up approach will require at least some detailed building information to allow final refinements of the archetypes.

References

Ali, U., M. H. Shamsi, C. Hoare, E. Mangina, and J. O'Donnell (2019, November). A data-driven approach for multi-scale building archetypes development. *Energy and Buildings 202*, 109364.

CanmetENERGY (2020, June). Building Assessment Technology Platform Github. https://github.com/canmet-energy/btap.

Cerezo, C., J. Sokol, S. AlKhaled, C. Reinhart, A. Al-Mumin, and A. Hajiah (2017, November). Comparison of four building archetype characterization methods in urban building energy modeling (UBEM): A residential case study in Kuwait City. *Energy and Buildings 154*, 321–334.

City of Victoria (2020). City of victoria open data portal. http://opendata.victoria.ca/.

(2011, February). *U.S. Department of Energy Commercial Reference Building Models of the National Building Stock*.

Government of Canada, N. R. C. (2020, January). Commercial/Institutional Sector British Columbia and Territories[1] Table 24: Space Heating Secondary Energy Use and GHG Emissions by Energy Source.

Kristensen, M. H., R. E. Hedegaard, and S. Petersen (2018, September). Hierarchical calibration of archetypes for urban building energy modeling. *Energy and Buildings 175*, 219–234.

Mata, E., A. Sasic Kalagasidis, and F. Johnsson (2014, November). Building-stock aggregation through archetype buildings: France, Germany, Spain and the UK. *Building and Environment 81*, 270–282.

Nagpal, S., C. Mueller, A. Aijazi, and C. F. Reinhart (2019, January). A methodology for auto-calibrating urban building energy models using surrogate modeling techniques. *Journal of Building Performance Simulation 12*(1), 1–16.

Nagpal, S. and C. F. Reinhart (2018, August). A comparison of two modeling approaches for establishing and implementing energy use reduction targets for a university campus. *Energy and Buildings 173*, 103–116.

Pasichnyi, O., J. Wallin, and O. Kordas (2019, August). Data-driven building archetypes for urban building energy modelling. *Energy 181*, 360–377.

Reinhart, C. and C. C. Davila (2019, April). Urban building energy modeling. In J. L. Hensen and R. Lamberts (Eds), *Building Performance Simulation for Design and Operation* (2 ed.)., pp. 696–722. Second edition. | Abingdon, Oxon ; New York, NY : Routledge, 2019.: Routledge.

Reinhart, C. F. and C. Cerezo Davila (2016, February). Urban building energy modeling – A review of a nascent field. *Building and Environment 97*, 196–202.

Roth, J., A. Martin, C. Miller, and R. K. Jain (2020, December). SynCity: Using open data to create a synthetic city of hourly building energy estimates by integrating data-driven and physics-based methods. *Applied Energy 280*, 115981.

Swan, L. G., V. I. Ugursal, and I. Beausoleil-Morrison (2009, June). A database of house descriptions representative of the Canadian housing stock for coupling to building energy performance simulation. *Journal of Building Performance Simulation 2*(2), 75–84.

Tardioli, G., A. Narayan, R. Kerrigan, M. Oates, J. O'Donnell, and D. P. Finn (2020, July). A methodology for calibration of building energy models at district scale using clustering and surrogate techniques. *Energy and Buildings*, 110309.