

Developing a Building Identification Tool to Support Mass Deep Energy Retrofits

HETONG SHEN
McGill University

FRÉDÉRIC VERRIER-PAQUETTE
McGill University

FRANK SUERICH-GULICK
McGill University

LOÏC HO-VON
McGill University

RUOQI WANG
McGill University

DANIEL CHUNG
University of Toronto

THOMAS KING
McGill University

GABRIELLE GOLDMAN
McGill University

MICHAEL JEMTRUD
McGill University

Keywords: Building Stock Model, Building Data, Typologies, Building Envelope, Machine Learning, Data Augmentation Process, Digital Tools, Mass Deep Energy Retrofits (DERs), Industrialized Construction, Building Information Modeling

Deep energy retrofits (DERs) are recognized as a critical strategy in reducing the building sector's greenhouse gas (GHG) emissions^{1,2}. Compared to shallow retrofits, DERs reduce energy demand of existing buildings by up to 50 to 80%. In a fabric first over-cladding approach to DERs, the insulation and airtightness of the building envelope are significantly improved, and the mechanical systems are replaced with smaller, ultra-high-efficient ones. To contribute in a meaningful way to reaching Canada's GHG reduction targets, a substantial portion of the existing building stock must undergo deep retrofitting⁴. To achieve this, the DER process—from identification of candidate buildings to the industrialized manufacture and installation of panelized envelope solutions—must be streamlined, affordable, and scalable. A key barrier to development is lack of access to accurate building data qualifying and quantifying the condition of the existing building stock, which is needed to assess the pool of buildings and building aggregations with a high potential for deep retrofit³. Specifically, data characterizing building envelopes—from construction, form, material, and openings to renovation history and overall fitness—are unavailable for the full building stock and currently can only be obtained through labor-intensive on-site inspections. This paper presents the framework and methodology of a new web application, the Building Identification Tool (BIT), which allows one to remotely source and survey candidate buildings with a high potential for prefabricated panelized DER (PPDER). BIT allows users to augment a preprocessed geo-referenced building dataset with information relevant to PPDER, facilitating the identification of a scalable pipeline of projects that are amenable to mass customizable over-cladding solutions. The paper also explores how modern machine learning methods could eventually be used to automate certain stages of the data augmentation process.

BACKGROUND AND INTRODUCTION

Existing financing and incentive programs in the Canadian building energy efficiency market have largely resulted in ad hoc, singular retrofit actions. These typically include the installation of high efficiency building technologies, such as heat pumps, LED lighting, and weatherization measures, which can result in 30-50% energy savings. By comparison, a Deep Energy Retrofit (DER) approach can achieve energy reductions of 50-80%, where the driving factor is a major upgrade in the building envelope that dramatically reduces heating and cooling demand⁴. Beyond energy savings, DERs generate numerous benefits ranging from enhanced occupant comfort to better indoor air quality, and long-term climate resilience⁵⁻⁷. Despite these advantages, very few DERs are currently realized annually although being a critical component in national efforts to meet decarbonization targets⁸. This is largely due to capacity deficiencies in terms of financing and implementation that could be at least partially alleviated with a streamlined process for identifying and determining aggregation scenarios suited to mass customization, bulk purchasing, and the efficiencies of scale of production.

Prefabricated Panelized DERs (PPDERs) aim to expedite and systematize retrofits through an integrative design and manufacturing process generally known as *over-cladding*. In this approach, panels are manufactured off-site and packaged to include structural support and attachments, thermal insulation, hygrothermal control membranes, flashings, fenestration, and cladding. These exterior retrofits are coupled with high-efficient retrofit-ready active systems. By focusing resources on off-site industrialized processes, PPDERs streamline implementation by minimizing waste, on-site construction time and tenant disruption, and by integrating most steps of the design and construction process into a single, high-quality controlled workflow and team.

The current challenge for mass PPDER is that buy-in from diverse stakeholders cannot be achieved without data on the existing building stock to identify market-viable aggregation opportunities¹³⁻¹⁵. Specifically, data are needed to characterize building envelopes and super-structures at the individual building-level within typologically consistent sets of buildings. These are difficult, if not impossible to obtain at scale with current labor-intensive inspection processes: if, and when, such information is

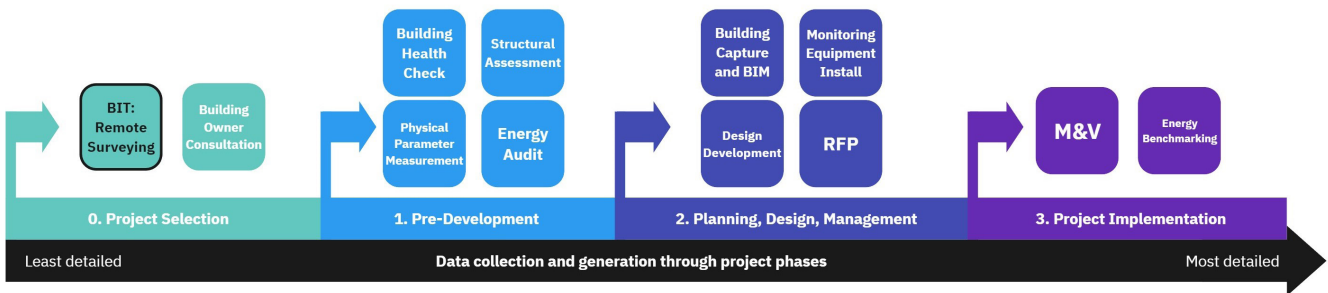


Figure 1. Workflow of PPDER development. The evaluation and decision-making process to design, manufacture, and install a PPDER solution has multiple steps with increasing levels of detail. BIT serves as the first step in this workflow, providing a quick and simple tool capable of enhancing low resolution data with the information needed to identify subsets of PPDER candidates for further analysis..

available through large portfolio owners, datasets are often decentralized, proprietary, and contain information asymmetries. The absence of a consistent and sufficient datasphere remains a persistent hurdle to scaling PPDER effectively.

Some existing approaches attempt to understand the demand and potential impact of DERs at a building stock scale with limited data, including building stock models (BSMs), which often assess the characteristics, energy usage, and performance of the existing building stock. Diverse model types exist, ranging from top-down and statistical approaches to bottom-up, physics-based models¹⁶⁻¹⁸. According to Langevin et al. (2020), bottom-up BSMs are best suited to evaluate the impact of new technologies when high spatial granularity is needed, as they are formulated to measure outputs at the individual building level according to physical simulations^{19,20}. However, while the level of detail of BSM techniques is increasing^{21,22}, existing bottom-up models continue to rely on aggregate data or representative building archetypes that may be helpful for high-level energy policy decisions but are deficient for on-the-ground retrofit aggregation, owner investment, and implementation decision-making²³.

Street-view and satellite imagery from Google Maps are a promising means for collecting data about individual buildings at a large scale^{24,25}. They have been used to generate building function maps at regional and urban scales in North America²⁶, predict building function and architectural styles in Mexico²⁷, and analyze building façade materials and reusable components in Barcelona and Zurich²⁸. Satellite images have also been used independently to extract and detect the presence of buildings^{29,30}, or to detect roof and material type³¹. They have yet to be used to assess wholistic conditions that determine suitability for retrofit.

The Building Identification Tool (BIT) is a web application that collects data relevant to PPDER suitability for existing buildings at an individual-building level, to assist in identifying promising

candidates and developing pipelines of PPDER projects (Figure 1). BIT guides users through a survey of questions regarding each building's characteristics based on street-view and satellite images complemented with data extracted from an existing dataset such as the property assessment roll. The collected data are then compiled and stored in a database, allowing all buildings of a selected subset or type of building to be roughly and quickly characterized for PPDER suitability. These data can then be used to identify both individual candidate buildings and potential building aggregation clusters, and from there, to quantify the potential for PPDER of a given building stock of the same typology and/or construction method.

Building types are deemed to have PPDER potential based on five primary criteria: 1) structural and tectonic consistency that would engender a similar, and thus repeatable, over-cladding technical solution; 2) the sheer number of such buildings across the province capable of establishing a pipeline; 3) a vintage of building stock that would present a high number of "anyways" renovation scenarios; 4) societal and public sector relevance; 5) stable public sector investment opportunity at the municipal, provincial, and/or federal level.

Based on these criteria, BIT was developed to identify 1) pre-fabricated steel buildings typical of community centers, curling/hockey rinks, and light manufacturing; 2) three- to five-story light-framed or concrete superstructure multi-unit residential buildings (MURBs) typical of 1950-80's social housing programs that produced a large number of comparable buildings. A distinct database subset is constructed for each building type, to be processed separately. These building types were strategically chosen, as they lend themselves well to simplified over-cladding retrofit solutions due to their simple form factor. Many buildings of both types are owned by not-for-profit entities, which are more likely to have explicit GHG emission goals and dedicated renovation budgets intended to protect and enhance the

longevity and resilience of their building stock and increase the well-being of its occupants.

Social housing MURBs have additional advantages for early phase PPDER, as clusters of similar MURBs of the same age are often owned and operated by the same entity, making them optimal candidates for early phase pilot projects. In addition, federal agencies are currently allocating substantial funds to realize pilot project DERs of social housing buildings¹². The research team is working closely with Quebec's housing agency, the *Société d'habitation du Québec* (SHQ), which oversees all social housing in the province. The goal is to develop a pipeline and workflow adapted to their MURB building stock and collaborate on a series of PPDER pilot projects.

This paper describes the following aspects of BIT:

- An overview of BIT's structure and methodology, including the initial database construction, preprocessing steps, the user survey and the user interface;
- a description of preprocessing steps specific to 1) prefabricated steel buildings, and 2) social housing MURBs;
- an initial presentation and discussion of how machine learning methods can be applied to automate parts of the retrofit candidate selection process.

BIT STRUCTURE AND WORKFLOW

Figure 2 shows the data processing workflow leading to BIT's initial database, and how data collected through surveying augments the database. To accelerate data augmentation and make optimal use of users' time, machine learning is used to preselect and order the candidate buildings shown to users. In addition, a collection of machine learning methods is explored to automate surveying, with preliminary results presented in the last sections of the paper.

INITIAL DATABASE CONSTRUCTION

The initial database is populated using data from the 2022 Quebec property assessment roll (hereafter "the Roll")³³, which covers buildings for the entire province of Quebec and is updated every three years. The buildings are organized into evaluation units ("units", for short) containing one or more adjacent buildings. Each unit has a unique ID, street name, primary land-use code (CUBF) and tax-related property information. Additional fields of interest are available only for some units: the full address, construction year, physical adjacency type (e.g., row house, single-family detached, etc.), the number of floors and the number of dwellings in the building.

Roll data fields are provided in XML and GIS Shapefile format. Custom parsers are used to query and combine data provided in both formats and store them in an SQL database. The availability

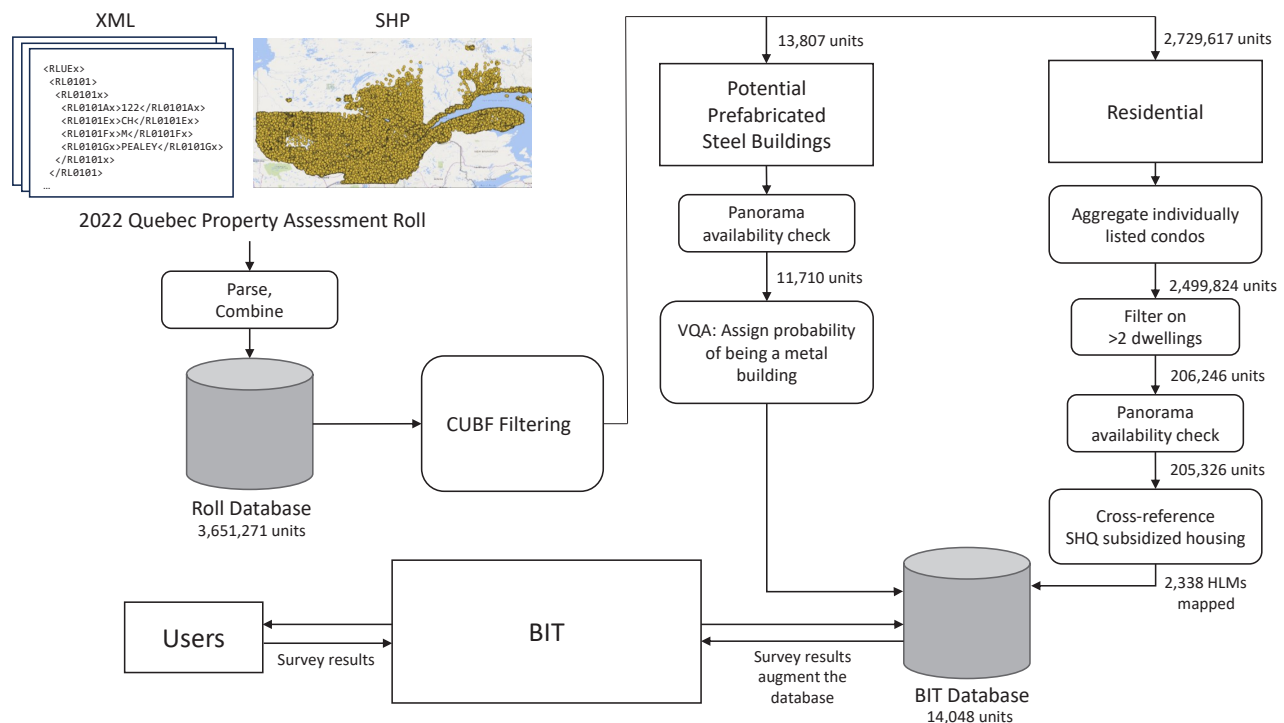


Figure 2. Diagram of BIT's data preprocessing and augmentation flows. Starting in the top left, data are extracted from the Roll and split into two streams, potential steel buildings and residential buildings, each of which undergoes various preprocessing steps before storage in the initial BIT database (bottom right). BIT then shows the data to users for evaluation and augments the BIT database with their survey responses.

of Streetview panoramas for each unit is verified; only units yielding a panorama are included. The CUBF code is used to select all units that might fall under one of the two desired building types: prefabricated steel community buildings, and social housing MURBs. The subsets for the two building types undergo further preprocessing separately.

BIT USER INTERFACE

The BIT user interface shows data from the initial database to users for each evaluation unit and prompts them to answer survey questions for that unit. Django ModelForms³⁶ is used on the backend to easily construct a fillable survey from its database model, along with some custom fields and JavaScript code to improve user experience. The system is designed with extensibility, effectiveness, and suitability in mind so that adding or removing questions is easy for future versions of the survey.

When accessing BIT online for the first time, a user is directed to a landing page that explains the tool’s context and objectives, and where they are asked to create an account. Once registered, the user is taken to the survey page where data and images for the first evaluation unit are displayed in two side-by-side interactive Streetview and Satellite Map windows, with a pin indicating the unit’s coordinates from the Roll, as shown in Figure 3.

The survey (Table 1) appears in a separate tab as a scrollable list of questions. A user will commonly toggle back and forth between the survey and building views, and may virtually move around the building if obstructions occur from one viewing

angle in Streetview, as they answer questions. A time travel functionality is implemented, enabling users to view the building conditions’ evolution through time, which may indicate façade renovations or reveal better views if vegetation obstructs the view during certain seasons. Once a user submits the survey, the responses are captured and associated to the evaluation unit in the database. The user is then directed to complete the survey for the next evaluation unit.

As the user completes the survey, they are encouraged to record screenshots of the Streetview each time they observe new façades, change viewing angles, or notice something interesting. The Satellite view is automatically captured each time the user switches to the survey tab, and in cases where the user does not take screenshots, the last state of the Streetview is captured at survey submission. All images are uploaded asynchronously to servers to ensure the application remains responsive. This results in a dataset of images labelled with survey responses which can be used to finetune or train computer vision models to automate all or parts of the survey response process and reduce the need for human intervention in this phase of data collection.

BIT SURVEY & DEVELOPMENT

The current version of the user survey (Table 1) focuses on traits that can be gleaned from Streetview and satellite imagery that influence suitability for over-cladding or indicate an imminent need for retrofit. These traits complement or are redundant with information typically collected by professionals using Building Condition Assessments and energy audits to produce

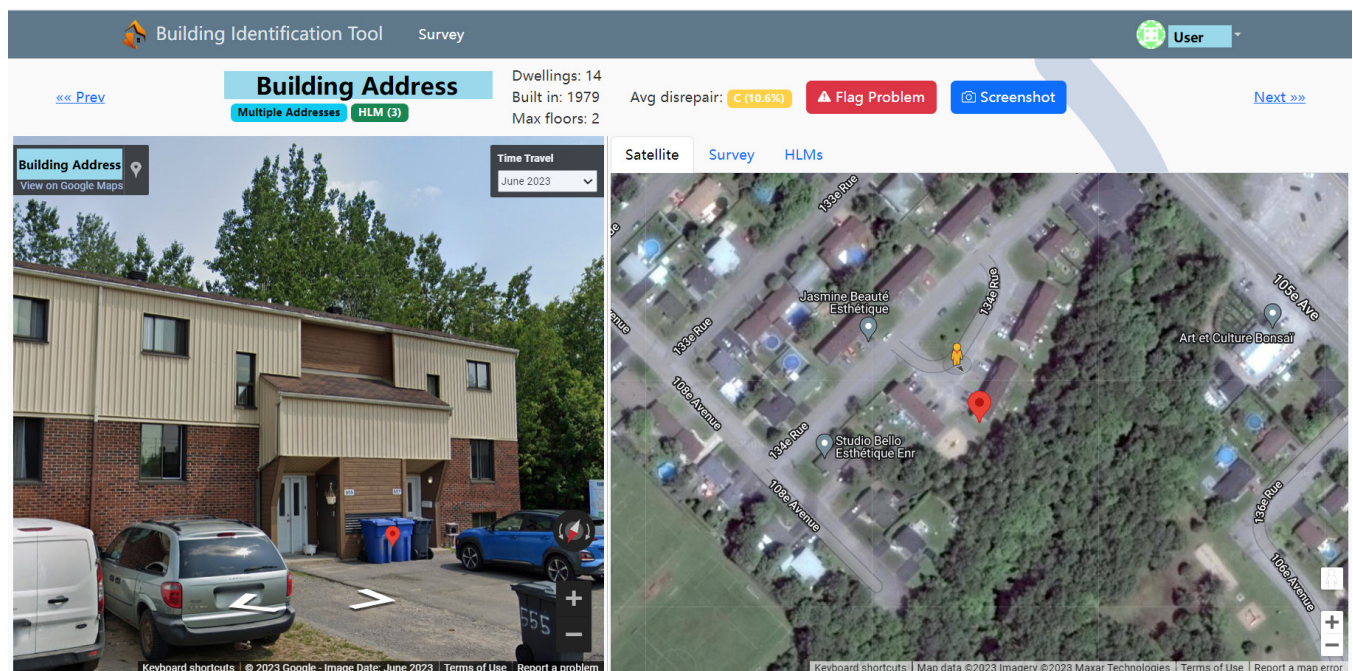


Figure 3. The BIT user interface displaying a sample unit. The Streetview image appears on the left, and satellite image on the right. The tabs above the satellite image allow the user to toggle between the satellite image and the survey. For HLM building types, the “HLMs” tab lists the HLMs present in the evaluation unit. Building information extracted from the property assessment roll is displayed at the top. The ‘Screenshot’ button is used to take additional screenshots from the Streetview, and the ‘Flag Problem’ button is used to flag units that have no building.

and organize systematic roadmaps for deep retrofits (Building Renovation Passports)³².

Though BIT has been tested exclusively with architecture students as users, the goal is to eventually recruit users from the public. The survey questions are therefore designed to be simple, with a small number of answer options. For each question, visual examples are available through a pop-up information bubble, as well as an explanation of its relevance to the PPDER process. Answer options such as “unsure” and “other (specify)” allow a user to skip over tricky questions or note atypical configurations.

To construct an initial coarse dataset of potential prefabricated steel community buildings, all Roll units with CUBF codes for educational, recreational and public gathering spaces are extracted, with the idea that these are most likely to contain this type of building. This approach casts a wide net: many units coded as recreational areas are parks and sports fields that do not contain a building and only a small proportion of these units are prefabricated steel buildings.

A preprocessing step is developed to rank the units and determine which ones are most likely to contain a metal building. An image-text multimodal deep learning model⁹ is employed to perform a visual question answering (VQA) task on the unit’s Streetview panoramas. Two sets of questions are directed to the VQA to estimate 1) the probability $P(B)$ that a building is present in the unit, and 2) the probability $P(M)$ that a metal-clad building exists in the unit, as shown in Figure 4. $P(M)$ and $P(B)$ for each unit

are computed by averaging all the values produced by the VQA model for every Streetview panorama obtained for that evaluation unit, including those available from multiple angles and time-periods. Each question is considered as a binary test and Bayes’ theorem is applied using the empirically measured sensitivity, specificity and prevalence to calculate the probabilities¹⁰.

The estimated $P(B)$ and $P(M)$ values for each unit are theWn used to determine the order in which they are shown to users. Units most likely to have a metal building (high $P(B)$ and $P(M)$) are shown first, followed by those likely to have a non-metal building, and those least likely to have a building of any kind (low $P(B)$ and $P(M)$). Visual inspection of a random sample of the highest and lowest scored evaluation units shows promising results, with around 75% of the highest rated units being true positives, and none of the lowest rated ones being false negatives.

PREPROCESSING FOR MULTI-UNIT RESIDENTIAL BUILDINGS

The initial database of MURBs is constructed through multiple preprocessing steps. The initial dataset is obtained by selecting all units in the Roll with a residential CUBF code. The resulting set of 2,729,617 units contains MURBs, single-family homes and individually listed condominiums, many of which appear in a same building. As a first step, the individual condominiums are aggregated into a single evaluation unit representing the full building using a composite key of the latitude, longitude, address, municipality fields. The address and municipality fields are likely redundant for this operation. The result is the aggregation

Constructs	Item	Question	Choices
Compartments	1	Does the building appear part of a self-similar cluster. If so, how many buildings are in the cluster?	Number of buildings in cluster; No
	2	Does the building have a simple footprint?	Yes; No
Building Configuration and Characteristics	3	Does the building have a simple volumetric form?	Yes; No
	4	How many storeys above-ground does the building have?	Number of storeys; Unsure
	5	What best describes the roof geometry?	Flat; Low pitched; High pitched; Curved; Complex; Unsure
	6	Does the building have a basement?	Yes; No; Unsure
Obstructions and Accessibility	7	Select any and all obstructions to machine access around the building.	Important trees or landscaping; Overhead wires, blocking access; Buildings; Other (specify);
	8	Are there significant appendages to the building faces? Select all that apply.	Roof overhangs/eaves; Balconies; Porches/stoops; Exterior vestibules; Other (specify);
Building Exterior and Window	9	What type of exterior cladding does the building appear to have? Select all that apply.	Concrete; Plaster; Wood; Vinyl; Curtain Wall; Metal; Brick masonry; Stone masonry; Unsure; Other (specify)
	10	Are the building façades in poor condition and in need of replacement?	Yes; No; Unsure
	11	Does glazing make up more than 40 % of the total area of all visible façades?	Yes; No; Unsure
	12	Are there very large and/or irregularly shaped windows?	Very large; Irregularly shaped
Year of Construction and Condition	13	Does this building look new and/or recently renovated?	Newly built; Recently renovated

Table 1: Building Identification Survey

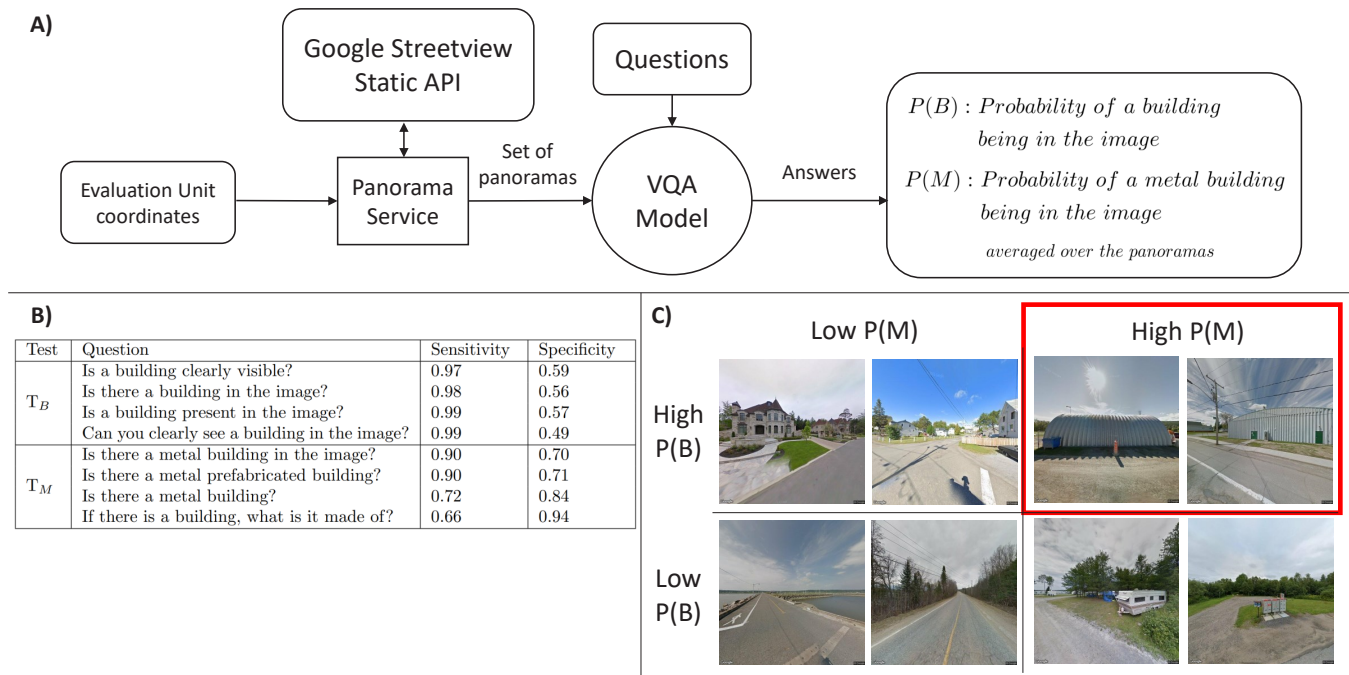


Figure 4. (A) Diagram showing the VQA workflow. A set of panoramas is obtained from the Streetview API and submitted to the VQA model along with a set of questions. $P(B)$ and $P(M)$ are estimated and averaged over the panoramas of the unit. (B) The list of questions asked for each test. Sensitivity indicates how often the model accurately predicts the presence of a building, and specificity, its absence. (C) An example of model results for images from eight different units. Images in the upper row are rated by the model as more likely to have a building according to Test T_B and those in the right-hand quadrants are rated as more likely to have a metal building (Test T_M). Units with high $P(B)$ and $P(M)$ values (top right quadrant) are most likely to include a metal building.

of 248,713 individually listed condos into 20,054 MURBs. By then removing all buildings that include fewer than 3 dwellings, the total number of residential MURBs of interest is thus reduced to 206,246. While this resulting dataset will be valuable for future pipelines of privately owned MURBs, the building subset of immediate interest includes only the social housing MURBs that are owned and operated by the SHQ, which are locally referred to as *HLMs*, from the French term *habitations à loyer modique*. The steps required to extract these from the larger MURB database are explained next.

PREPROCESSING FOR SOCIAL HOUSING MURBS

The second building type under study using BIT are social housing MURBs owned by the SHQ. The SHQ maintains a publicly available dataset of its HLMs, including an index that indicates their state of disrepair. This index is of strategic importance for retrofit planning, as the added costs of realizing PPDER for buildings requiring major renovations should be marginal relative to conventional retrofits.

The initial database of social housing MURBs is produced by cross-referencing the HLM dataset with the MURBs dataset described in the previous section. While this operation is conceptually simple, it is complicated in practice by issues such as inconsistent abbreviations, incomplete addresses, and multiple buildings located within a single Roll unit. For example, units

located in northern communities are underrepresented in the BIT database due to the absence of Streetview imagery in the Google API and the omission of the municipality field in the Roll for some units. When direct lookups of SHQ addresses in the Roll, approximated “fuzzy” matching is used on the municipality and street names, and the street numbers are checked for inclusion within the Roll unit range. As a result, 4,637 out of 7,373 entries in the SHQ dataset are successfully cross-referenced to the Roll data. Filtering for buildings with three or more dwellings yields 3,140 HLMs mapped to 2,338 unique evaluation units.

SURVEY AUTOMATION - MACHINE LEARNING EXPERIMENTS

Given the large size of the datasets to be evaluated, it is desirable to have as many parts of the BIT workflow performed by machine learning as possible instead of by human users. As a first step towards evaluating this potential, answers from a machine learning model are compared to answers collected from users about the roof type and number of floors in the current MURBs BIT dataset. The open-source BRAILS framework is used, which provides pre-trained models for these purposes^{34, 35}, working from Streetview screenshots saved by BIT users.

The BRAILS NFloorDetector model detects the number of floors using the EfficientDet-D4 architecture, by drawing a bounding box around detected floors and counting the number of bounding



Figure 5. Examples of number of floors image prediction.

Left to right: (1) Correct prediction of 3 storeys. (2) Incorrect prediction of 3 storeys (correct value is 2). The building erroneously included floors detected on the neighbouring building (boxes that are not vertically aligned) in its count. (3) Incorrect prediction of 5 storeys (correct value is 3). Images with a steep lateral viewing angle decrease model performance.

boxes, as shown in Figure 5. The BRAILS results were only accurate 52% of the time compared to the user responses for 343 MURBs in the BIT dataset. The model often overestimates the number of floors when an adjacent building is visible, as shown in the middle image. This issue partly arises because the model, developed using data from New Jersey's single-family (50.0%), multi-family (32.2%), and commercial buildings (17.8%), may not be suitable for Quebec's MURBs or multi-family buildings, given that the model was not trained for this specific building type. Additionally, the BIT screenshots taken by BIT users had more obstructions such as trees and vehicles, adjoining buildings and acute lateral viewing angles. The model performance might be improved by more suitable and higher quality images or if it can be adjusted to consider only those bounding boxes that are vertically aligned in its count.

To test roof classification prediction, images from the BRAILS dataset sourced from OpenStreetMaps were used³⁴ to train the model. The dataset consists of 6,000 labeled satellite images, with 2,000 examples each for "flat", "gabled", and "hipped" roof types. "Gabled" and "hipped" categories were merged into a single "not flat" category to align with the roof type categorization in the BIT survey, forming a dataset with 2000 examples of "flat" roofs and 4000 examples of "not flat" roofs. The dataset was randomly split into 5000 images for training and 1000 images for testing. The resulting model trained on ResNet-18³⁵ had a training accuracy of 89.3 %, and a testing accuracy of 86.5 % on the BRAILS dataset. It had an accuracy of 80% when tested on 60 BIT satellite images.

DISCUSSION

The workflow and tool described here were developed in an iterative way by a team of researchers from computer science, architecture and engineering with a diverse range of experiences and expertise. While many of the processes and

components may appear simple, they are the result of multiple cycles of discussion, development, testing, and refinement. The survey and interface were adapted based on feedback from many team members, collaborators and students to improve clarity and user experience. Few programming details and challenges related to querying, preprocessing and analyzing data are described here for the sake of brevity and to engage with the intended audience of architects; such details will be covered in a subsequent publication.

Among the limitations of BIT affecting its accuracy is that Streetview imagery for some units may be outdated. In particular, imagery for remote locations may date from many years or even a decade ago. There is also potential for inaccuracies in the collected data since BIT gathers data through user input. However, the resulting database will be further augmented and verified through additional rounds of information gathering in the path towards identifying a PPDER pipeline, so false positives from BIT will be identified and eliminated along the way.

To scale up the data collection process, the research team intends to release BIT to the public and explore ways to incentivize citizens to participate in the research effort, in a "citizen science" manner, including gamifying the process and developing a communications campaign with a positive mission narrative. Different surveys may be developed and directed to different user categories, based on their reported or assessed level of expertise. Future work includes the progressive development of machine learning models, and an automated survey workflow integrated with a PPDER suitability scoring function to rank building candidates. This will entail detecting and analyzing building features from images and roll data and automatically answering survey questions based on this information.

CONCLUSION

PPDERs are emerging as a vital component in reducing GHG emissions from the building sector in Canada. By streamlining the retrofit process—beginning with the identification of aggregate candidate buildings—and focusing on high-quality, controlled manufacturing and installation workflows, PPDERs hold the potential not only for significant energy reductions but also for enhancing occupant comfort, indoor air quality, and resilience.

The BIT application addresses the lack of suitable large-scale datasets for PPDER assessment by generating bottom-up building data through novel data collection techniques. It augments an initial database with Google Streetview, satellite images, and user-generated data. BIT is currently employed to analyze the existing MURBs stock and prefabricated metal community buildings in Quebec to identify suitable candidates for PPDER interventions. Machine learning is used to automate parts of the workflow to quickly characterize and rank a large volume of buildings.

BIT is part of a larger suite of custom tools and approaches that are in development with collaborators to support the rapid expansion of PPDERs. Development of BIT will continue within this larger suite, as these tools are expected to play a critical role in facilitating strategic investment and developing the necessary capacity to realize PPDERs at scale. The successful scaling of PPDERs will be a critical step towards achieving national and global environmental goals, again underscoring the importance of this research in the broader context of energy efficiency and climate action.

ACKNOWLEDGEMENTS

This work was supported by the National Sciences and Engineering Research Council of Canada (NSERC Alliance grant), Hydro-Québec, and Quebec's Ministry of the Environment, the Fight Against Climate Change, Wildlife and Parks. Several people provided valuable suggestions and input in producing the tool, including Antoine Mathys, Olivia Raymond, Madeleine Lachance, Ella Fortney, Nick Rudnicki, Karine Lavigne, Mike Coillot, and students from the McGill School of Architecture.

ENDNOTES

1. Hoicka, C. E., & Das, R. "Ambitious Deep Energy Retrofits of Buildings to Accelerate the 1.5°C Energy Transition in Canada." *Canadian Geographies / Géographies Canadiennes* 65, no. 1 (2020): 116-127. <https://doi.org/10.1111/cag.12637>.
2. Göswein, V., Silvestre, J. D., Sousa Monteiro, C., Habert, G., Freire, F., & Pittau, F. "Influence of Material Choice, Renovation Rate, and Electricity Grid to Achieve a Paris Agreement-Compatible Building Stock: A Portuguese Case Study." *Building and Environment* 195 (2021): 107773. <https://doi.org/10.1016/j.buildenv.2021.107773>.
3. Sanhudo, L., Ramos, N. M. M., Poças Martins, J., Almeida, R. M. S. F., Barreira, E., Simões, M. L., & Cardoso, V. "Building Information Modeling for Energy Retrofitting – A Review." *Renewable and Sustainable Energy Reviews* 89 (2018): 249–260. <https://doi.org/10.1016/j.rser.2018.03.064>.
4. Torrie, R., Bak, C., Heaps, T. "Building Back Better with a Bold Green Recovery Synthesis Report." Available at: <https://www.corporateknights.com/leadership/building-back-better-bold-green-recovery-synthesis-report/>
5. Francisco, P. W., Jacobs, D. E., Targos, L., Dixon, S. L., Breyse, J., Rose, W., & Cali, S. "Ventilation, Indoor Air Quality, and Health in Homes Undergoing Weatherization." *Indoor Air* 27, no. 2 (2016): 463–477. <https://doi.org/10.1111/ina.12325>.
6. Wolkoff, P. "Indoor Air Humidity, Air Quality, and Health – An Overview." *International Journal of Hygiene and Environmental Health* 221, no. 3 (2018): 376–390. <https://doi.org/10.1016/j.ijheh.2018.01.015>.
7. Laski, J. "The Business Case for Green Buildings." In *Marketing Green Building Services*, 82–98. (2018). <https://doi.org/10.4324/9780080942315-10>.
8. Frappé-Sénéclauze, T.-P. "A Deep Retrofit of Homes and Buildings Is the Megaproject Canada Needs: Post-COVID Is an Opportune Time to Invest Nationally in Climate-Oriented Upgrades That Will Improve Efficiency and Move Us Toward a Low-Carbon Future." *Policy Options* (2020). <https://policyoptions.irpp.org/magazines/may-2020/a-deep-retrofit-of-homes-and-buildings-is-the-megaproject-canada-needs/>.
9. Summaira, J., Li, X., Shoib, A. M., Li, S., & Abdul, J. "Recent Advances and Trends in Multimodal Deep Learning: A Review." *arXiv.org* (2021): <https://arxiv.org/abs/2105.11087>.
10. Mathai A, Parikh S, Chandra Sekhar G, Thomas R. Understanding and using sensitivity, specificity and predictive values. *Indian J Ophthalmol.* 2008 Jan-Feb;56(1):45-50. doi: 10.4103/0301-4738.37595. PMID: 18158403; PMCID: PMC2636062.WW
11. Green Retrofit Economy Study – CAGBC. Available at: www.cagbc.org/wp-content/uploads/2022/06/Green-Retrofit-Economy-Study.pdf.
12. Canada Greener Affordable Housing. CGAH. (n.d.). Available at: <https://www.cmhc-schl.gc.ca/professionals/project-funding-and-mortgage-financing/funding-programs/all-funding-programs/canada-greener-affordable-housing-program>
13. Muller, M. F., Esmaniotto, F., Huber, N., Loures, E. R., & Cangiuglieri, O. "A Systematic Literature Review of Interoperability in the Green Building Information Modeling Lifecycle." *Journal of Cleaner Production* 223 (2019): 397–412. <https://doi.org/10.1016/j.jclepro.2019.03.114>.
14. Nägeli, C., Camarasa, C., Jakob, M., Catenazzi, G., & Ostermeyer, Y. "Synthetic Building Stocks as a Way to Assess the Energy Demand and Greenhouse Gas Emissions of National Building Stocks." *Energy and Buildings* 173 (2018): 443–460. <https://doi.org/10.1016/j.enbuild.2018.05.055>.
15. Kavgić, M., Mavrogianni, A., Mumovic, D., Summerfield, A., Stevanovic, Z., & Djurovic-Petrovic, M. "A Review of Bottom-Up Building Stock Models for Energy Consumption in the Residential Sector." *Building and Environment* 45, no. 7 (2010): 1683–1697. <https://doi.org/10.1016/j.buildenv.2010.01.021>.
16. Wiesmann, D., Lima Azevedo, I., Ferrão, P., & Fernández, J. E. "Residential Electricity Consumption in Portugal: Findings from Top-Down and Bottom-Up Models." *Energy Policy* 39, no. 5 (2011): 2772–2779. <https://doi.org/10.1016/j.enpol.2011.02.047>.
17. Swan, L. G., & Ugursal, V. I. "Modeling of End-Use Energy Consumption in the Residential Sector: A Review of Modeling Techniques." *Renewable and Sustainable Energy Reviews* 13, no. 8 (2009): 1819–1835. <https://doi.org/10.1016/j.rser.2008.09.033>.
18. Oladokun, M. G., & Odesola, I. A. "Household Energy Consumption and Carbon Emissions for Sustainable Cities – A Critical Review of Modeling Approaches." *International Journal of Sustainable Built Environment* 4, no. 2 (2015): 231–247. <https://doi.org/10.1016/j.ijse.2015.07.005>.
19. Lawrence Berkeley National Laboratory. Available at: https://etapublications.lbl.gov/sites/default/files/developing_a_common_approach_for_classifying_building_stock_energy_models.pdf.
20. A Methodology for Spatial Modelling of Energy and Resource Use Of Buildings In Urbanized Areas. Available at: publications.lib.chalmers.se/records/full-text/205110/local_205110.pdf.
21. Krarti, M., & Dubey, K. "Review Analysis of Economic and Environmental Benefits of Improving Energy Efficiency for UAE Building Stock." *Renewable and Sustainable Energy Reviews* 82 (2018): 14–24. <https://doi.org/10.1016/j.rser.2017.09.013>.
22. Nägeli, C., Jakob, M., Catenazzi, G., & Ostermeyer, Y. "Towards Agent-Based Building Stock Modeling: Bottom-up Modeling of Long-term Stock Dynamics Affecting the Energy and Climate Impact of Building Stocks." *Energy and Buildings* 211 (2020): 109763. <https://doi.org/10.1016/j.enbuild.2020.109763>.
23. Krarti, M., Aldubyan, M., & Williams, E. "Residential Building Stock Model for Evaluating Energy Retrofit Programs in Saudi Arabia." *Energy* 195 (2020): 116980. <https://doi.org/10.1016/j.energy.2020.116980>.
24. Explore Street View and Add Your Own 360 Images to Google Maps. Google. Available at: www.google.com/streetview/.
25. Google Map Overview. Google Earth, Google, Available at: www.google.com/earth/about/.
26. Kang, J., Körner, M., Wang, Y., Taubenböck, H., & Zhu, X. X. "Building Instance Classification Using Street View Images." *ISPRS Journal of Photogrammetry and Remote Sensing* 145 (2018): 44–59. <https://doi.org/10.1016/j.isprsjs.2018.02.006>.
27. Taoufiq, S., Nagy, B., & Benedek, C. "HierarchyNet: Hierarchical CNN-Based Urban Building Classification." *Remote Sensing* 12, no. 22 (2020): 3794. <https://doi.org/10.3390/rs12223794>.
28. Raghu, D., Markopoulou, A., Marengo, M., Neri, I., Chronis, A., & De Wolf, C. "Enabling Component Reuse from Existing Buildings through Machine Learning, Using Google Street View to Enhance Building Databases." *CAADRIA proceedings, 2022*. <https://doi.org/10.52842/conf.caadria.2022.2.577>.
29. Ghanea, M., Moallem, P., & Momeni, M. "Building Extraction from High-Resolution Satellite Images in Urban Areas: Recent Methods and Strategies against Significant Challenges." *International Journal of Remote Sensing* 37, no. 21 (2016): 5234–48. <https://doi.org/10.1080/01431161.2016.1230287>.
30. Sirko, W., Kashubin, S., Ritter, M., Annkah, A., Bouchareb, Y. S. E., Dauphin, Y., ... Quinn, J. "Continental-Scale Building Detection from High Resolution Satellite Imagery." *arXiv.org* (2021): <https://arxiv.org/abs/2107.12283>.
31. Kim, J., Bae, H., Kang, H., & Lee, S. G. "CNN Algorithm for Roof Detection and Material Classification in Satellite Images." *Electronics* 10, no. 13 (2021): 1592. <https://doi.org/10.3390/electronics10131592>.
32. Sesana, M. M., & Salvalai, G. "A Review on Building Renovation Passport: Potentialities and Barriers on Current Initiatives." *Energy and Buildings* 173 (2018): 195–205. <https://doi.org/10.1016/j.enbuild.2018.05.027>.
33. The province of Quebec, Canada, is used as a case study for BIT's operations. We use property assessment roll data from: *Données Québec (2022) Tous les fichiers des rôles d'évaluation foncière du Québec 2022. Gouvernement du Québec. Available at: https://www.donneesquebec.ca/recherche/dataset/roles-d-évaluation-foncière-du-québec/resource/93548503-59b0-4721-b1a8-df98f447aa5c*
34. Cetiner, B., Wang, C., McKenna, F., Hornauer, S., Gao, Y., Yu, S. X., Taciroglu, E., & Law, K. H. "NHERI-SimCenter/Brails: Release v3.0.0." *Zenodo* (2022): <https://doi.org/10.5281/zenodo.7132010>.
35. He, K., Zhang, X., Ren, S., & Sun, J. "Deep Residual Learning for Image Recognition." *arXiv.org* (2015): <https://arxiv.org/abs/1512.03385>.
36. "Django." Django Project. <https://docs.djangoproject.com/en/4.2/topics/forms/>.
37. "15. Spatial Indexing - Introduction to Postgis." PostGIS. <https://postgis.net/workshops/postgis-intro/indexing.html>.