

Automated District-Level Energy Demand Modeling Using EnergyPlus Empowered by Digital Twin Technology

Amin Jalilzadeh¹, Azarakhsh Rafiee², Stefan De Graaf³, Thaleia Konstantinou⁴, Peter J.M. van Oosterom⁵

¹ Section GIS Technology, Faculty of Architecture and the Built Environment, TU Delft, Netherlands – A.jalilzadeh@tudelft.nl

² Section GIS Technology, Faculty of Architecture and the Built Environment, TU Delft, Netherlands – A.Rafiee@tudelft.nl

³ Research Department Geodan, Amsterdam, Netherlands – stefan.de.graaf@geodan.nl

⁴ Department of AE+T, Faculty of Architecture and the Built Environment, TU Delft, Netherlands – T.Konstantinou@tudelft.nl

⁵ Section GIS Technology, Faculty of Architecture and the Built Environment, TU Delft, Netherlands – P.J.M.vanOosterom@tudelft.nl

Keywords: Geosimulation, Urban Energy Demand, Digital Twin, EnergyPlus automation, Decision support systems.

Abstract

The global shift toward decentralization and decarbonization in the energy sector demands robust tools for accurately simulating building energy demand at the district scale. Although EnergyPlus is well-regarded for its detailed building-level modeling, it poses challenges when extended to larger districts with diverse building typologies. This paper presents both a conceptual architecture and a working prototype of an automated pipeline that addresses these limitations. Leveraging open-access geo-spatial and non-spatial data from the Netherlands—including over 10 million buildings—the pipeline seamlessly scales EnergyPlus simulations to the district level. The system employs the OGC 3D Tiles standard for efficient streaming and real-time visualization of simulation results across a nationwide 3D Tileset. Implemented in Python and tested on local machines, the pipeline is poised for cloud-based deployment to further enhance scalability and performance. By integrating a digital twin for real-time monitoring and scenario testing, the approach enables efficient bulk operations, interactive decision support, and clearer insights into urban-scale energy consumption patterns. The resulting automation and scalable workflows offer valuable contributions to sustainable urban planning, ensuring that efficiency opportunities can be quickly identified and acted upon at multiple spatial scales.

1. Introduction

The rapid urbanization and the growing urgency for sustainable energy consumption have led to profound transformations in the energy sector. Decentralization and decarbonization are central to these changes, driving the need for advanced tools and methodologies to accurately model energy demand at both building and district scales. Accurate modeling is essential for optimizing energy consumption, resource allocation, and achieving sustainability goals in urban environments (Simpson et al., 2024).

Most simulation tools are designed for individual building simulations, which limits their ability to model district-level energy systems efficiently. EnergyPlus is renowned for its detailed building-level simulations and excels in the modeling of individual buildings (Sola et al., 2018, Muslim, 2021). However, it faces several challenges. First, EnergyPlus requires users to have in-depth knowledge of energy modeling and complex system setups, making it challenging for non-experts or for simpler projects. Additionally, its thermodynamically accurate routines and detailed simulation capabilities demand more effort to configure and manage, which makes it less user-friendly for beginners than simpler simulation tools. EnergyPlus requires users to manually input detailed information and does not make assumptions about Heating, Ventilation, and Air Conditioning (HVAC) systems or other aspects, increasing the complexity for users (Crawley et al., 2001). EnergyPlus also demands detailed building data for accurate simulations, including geometry, HVAC systems, and occupancy data, which is often incomplete or difficult to obtain. The accuracy of results from EnergyPlus depends on the quality of input data, such as building geometries and usage schedules, and a lack of reliable

building-level data can introduce inaccuracies in the simulation results. Finally, simulations in EnergyPlus must be calibrated against real-world data, requiring additional effort to ensure that results align accurately with actual conditions (Ascione et al., 2021).

Consequently, integrating multiple tools like Matlab, Simulink, and City-Sim becomes necessary to manage the complexity and scale of simulating hundreds or thousands of diverse buildings (Robinson et al., 2009). A key challenge at the district level is the diversity in building ages, construction methods, occupancy patterns, and energy systems. For example, older buildings often have different insulation and heating systems compared to newer ones. This diversity complicates modeling, as assumptions for older buildings can vary widely—some may have been upgraded, while others have not. Additionally, manually creating Input Data Files (IDFs) for numerous buildings is impractical and time-consuming without automation (Deng et al., 2023, Wahi et al., 2024).

Data accessibility further complicates district-scale modeling. Detailed building information is often incomplete or inconsistent due to privacy concerns, large-scale data collection, and data-sharing limitations. This lack of reliable data hinders the creation of models that accurately reflect a district's building diversity. Scalability and computational demands also present significant obstacles. EnergyPlus's computational requirements increase substantially when simulating multiple buildings, leading to higher resource use and potential slowdowns. This is especially true for scenarios requiring fine temporal resolution, such as hourly thermal energy profiles, which in large-scale simulations may require high-performance computing resources to manage model complexity (Allegrini et al., 2015).

Recent work demonstrates ways to address these challenges by integrating EnergyPlus with additional tools. For example, (Muslim, 2021) proposes co-simulation using EnergyPlus with Matlab/Simulink, Radiance, or Modelica to handle complex multi-domain interactions. (Glazer, 2016) shows that Eppy (a Python library) can automate modifications to EnergyPlus files, enabling the application of efficiency measures at scale. In a similar vein, (Yue et al., 2021) integrates EnergyPlus with Python and Eppy to reduce setup time and facilitate optimization, slashing simulation periods from ten months to just two days. Other studies highlight Geomeppy's capabilities for 3D geometry handling and thermal zoning, improving realism at an urban scale (Faure et al., 2022). GIS-based methods further enrich these approaches by offering spatial insights into energy consumption (Ali et al., 2018). However, there remains a gap in seamlessly combining these tools—energy simulation libraries, automation scripts, and GIS—under a single framework for district-wide or nationwide analyses.

Digital Twin (DT) technology has emerged as a pivotal solution for bridging these gaps. A DT offers a virtual replica of real-world systems, enabling dynamic monitoring and scenario testing for complex energy infrastructures (Tahmasebinia et al., 2023). By incorporating live data and visualization, DTs can significantly augment EnergyPlus-based pipelines, making district-level simulations more accessible and interactive.

The contribution of this paper is to introduce both the conceptual architecture and a working prototype of a fully automated pipeline for district-level energy demand modeling with EnergyPlus. Unlike most existing approaches, this pipeline spans the entire Netherlands and leverages open-access datasets, applying established standards and documents to ensure data quality across diverse building types. Scripting with Eppy and Geomeppy underpins the automatic generation of IDFs, enabling rapid, large-scale simulations. A further novelty is the integration of DT technology and GIS for real-time visualization, interactive monitoring, and more informed decision-making—all of which reduce the barriers to use for stakeholders without extensive energy-modeling expertise.

The remainder of this paper is organized as follows. Section 2 describes the proposed methodology, including database integration, archetype development, and automated IDF generation. Section 3 presents the results of a pilot study, discussing performance benchmarks and visualization. Finally, Section 4 summarizes the primary conclusions, discusses broader implications, and outlines directions for future research.

2. Methodology

This section presents an automated pipeline for district-level energy demand modeling, highlighting both its conceptual design and practical implementation. By merging open-access datasets with scripted workflows, the pipeline delivers a scalable, efficient, and user-friendly solution for urban energy analysis. Figure 1 provides an overview of the pipeline's core components; in Section 2.4, we further show how these components integrate within a DT architecture.

As shown in Figure 1, the pipeline begins with a comprehensive *Building Database*, which contains detailed information about buildings in the Netherlands, including geometry, age, function, type, and energy labels. *Building Archetypes* are then created by categorizing buildings based on attributes such as

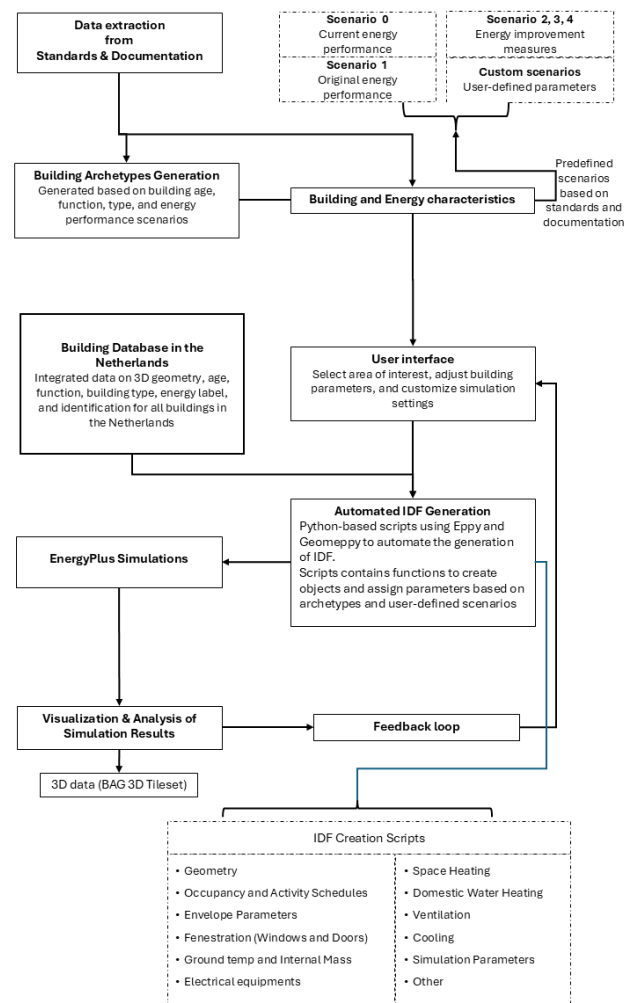


Figure 1. Overview of the Automated Pipeline for District-Level Energy Demand Modeling.

age, function, and energy performance, with key parameters and scenarios derived from national standards. User selections are made through a *User Interface*, where specific areas or buildings can be chosen and custom simulation parameters are defined. Next, *Automated IDF Generation* uses Python scripts (via Eppy and Geomeppy) to produce IDF files that incorporate both archetype-based defaults and user inputs. These files are then processed in the *EnergyPlus Simulations* stage, leveraging scalable computing resources. In the *Results Visualization and Analysis* phase, simulation outputs are visualized through DT dashboards, providing insights into energy consumption trends and performance metrics. Finally, a *Feedback Loop* enables users to refine parameters and rerun simulations, supporting an iterative workflow for policy testing or retrofitting scenarios.

2.1 Building Database in the Netherlands

Accurate building data underpins the pipeline's nationwide scope. Key information comes from several Dutch national organizations, including the Central Bureau of Statistics (CBS) (CBS: Central Bureau of Statistics, n.d.), the Netherlands Enterprise Agency (RVO) (RVO: Netherlands Enterprise Agency, n.d.), the Netherlands Environmental Assessment Agency (PBL) (PBL: Netherlands Environmental Assessment Agency, n.d.), the National Institute for Public Health and

the Environment (RIVM) (RIVM: National Institute for Public Health and the Environment, n.d.), and the Public Services on the Map (PDOK) (PDOK: Public Services on the Map, n.d.). By consolidating publicly accessible records from these organizations into a single PostgreSQL database, we establish a consistent foundation for district-level simulations. In total, approximately ten million building entries are maintained, each linked to its respective district, city, and postal code.

Managing such a large repository involves addressing data completeness and the need for frequent updates. To handle these issues, the database is periodically refreshed via application programming interface (API) connections to the original data providers, ensuring alignment with current conditions. Similar initiatives, such as the "Enriched BAG" project (Enriched BAG to Support Local Energy Issues, n.d.), demonstrate how integrating distributed sources can support local energy planning, although they often cover only a subset of the national building stock.

The Dutch building stock is diverse, spanning residential and non-residential properties. Residential buildings—detached, semi-detached, terraced, and apartments—comprise about 90% of the total. Non-residential categories include offices, health-care facilities, industrial buildings, educational institutions, sports halls, and retail spaces. Each record in the database captures attributes on age, function, type, and, where available, an energy label. Three-dimensional building geometries follow Levels of Detail (LOD) from 1.1 to 2.3, capturing parameters such as roof type, area, slope, height, and volume (3D BAG Dataset, n.d.). To facilitate nationwide 3D visualization, we employ the OGC 3D Tiles standard, which supports smooth streaming and interactive rendering within our DT environment. Where necessary, building orientation was derived through QGIS-based analyses. By centralizing these details and enabling flexible database queries, the system lets users select any region or set of buildings for simulation, forming a robust foundation for district-scale energy modeling.

2.2 Standards and Documents

In the Netherlands, a uniform set of building regulations and national standards ensures consistency in construction practices, thereby influencing energy performance across various building typologies. Such standards provide essential guidelines for modeling energy demand, defining key parameters for simulations, and promoting reliable outcomes in district-scale studies. Since building age, function, and construction practices differ significantly, understanding relevant documents is critical for accurate modeling.

One primary reference is the *Basisregistratie Adressen en Gebouwen (BAG)*, the national register of building footprints, usage types, construction years, and unique identifiers (BAG Viewer, n.d.). The *3D BAG* dataset extends BAG by offering three-dimensional geometries, enabling more detailed simulations of solar gains and shading effects (?). Additionally, the national database of *Energy Performance Certificates (EPC)* tracks energy labels and performance indicators submitted by certified energy advisors (EP-Online: Energy Performance Certificates Database, n.d.). Another key resource is the *Report on Standard and Target Values for Existing Residential Construction*, which, together with the *Voorbeeldwoningen* (example dwellings) datasets, specifies envelope properties, HVAC systems, and other parameters for representative Dutch residences (Report on Standard and Target Values for Existing

Residential Construction, n.d., *Voorbeeldwoningen: Example Dwellings for Dutch Buildings*, n.d.).

Occupancy patterns in Dutch buildings are frequently derived from empirical studies (Guerra-Santin and Silvester, 2017), while the *NTA 8800* standard dictates the official calculation methods for building energy performance (NTA 8800 Standard, n.d.). These calculations guide the assignment of insulation levels, heating system efficiencies, ventilation assumptions, and other variables in EnergyPlus-based simulations. Finally, the *EnergyPlus Documentation* itself remains an indispensable reference, ensuring that modeled objects and parameters align with software requirements and best practices.

2.3 Automated IDF Generation Model

District-scale energy simulation requires generating IDFs for thousands of diverse buildings, which is prohibitively time-consuming if done manually. To address this, our pipeline employs Python libraries *Eppy* and *Geomeppy* for creating and modifying IDFs at scale. *Eppy* provides convenient programmatic access to EnergyPlus objects, while *Geomeppy* extends this functionality to handle 3D geometry. Collectively, these tools enable rapid IDF generation, consistent application of building standards, and automated parameter adjustments.

Within this pipeline, each building is assigned to a particular archetype based on attributes such as age, function, and energy label. Archetypes are then mapped to predefined sets of parameters from sources including the *NTA 8800* and the *Voorbeeldwoningen* datasets. For instance, an archetype might feature a natural ventilation system with no heat recovery, as described in a specific brochure. Relevant EnergyPlus objects—such as `ZoneVentilation:DesignFlowRate`—are then populated with parameters taken directly from official guidelines or inferred from best estimates when data gaps arise.

The flexibility of this automated approach proves valuable for both model calibration and scenario testing. For example, insulation levels, U-values, and other envelope properties can be adapted quickly to reflect different regulatory requirements or retrofitting strategies. Parameters that lack explicit coverage in the documentation, such as unconventional HVAC systems or custom occupancy schedules, can be approximated and refined iteratively. Where official data is incomplete or absent, the pipeline flags these buildings for potential calibration against measured consumption data in later validation steps.

2.4 Integration with Digital Twin Technology

By incorporating DT technology, the pipeline provides an interactive and dynamic environment for visualizing, customizing, and analyzing energy consumption across urban districts. For instance, urban planners can compare the effects of various retrofitting measures, while policymakers can evaluate different energy consumption scenarios to assess proposed regulations and their potential impacts. As illustrated in Figure 2, the overall DT integration comprises three main components: a user interface, backend processing, and results visualization. The user interface is a web-based platform that allows users to select areas of interest by postal code, polygon, or individual building. Users can customize building parameters such as energy systems, insulation levels, and HVAC configurations, then initiate simulations directly from the interface. This setup greatly simplifies the modeling process by hiding the more intricate aspects of simulation configuration, making it accessible even to non-experts.

Backend processing manages the computational workload through a containerized environment for scalability and consistency. Upon receiving user inputs via an API, it generates IDF files according to the specified parameters and local building data, then executes EnergyPlus simulations in parallel. This ensures efficient and reliable performance for large-scale or intensive scenarios. Finally, the completed simulation outputs—including energy demand profiles and key performance indicators like total energy consumption, peak loads, and potential energy savings—are processed and stored in a centralized database. Results flow back to the DT platform via an API, where they can be explored in an immersive, interactive manner through specialized dashboards. This arrangement supports comprehensive scenario analyses, offering decision-makers a clear perspective on outcomes such as carbon reduction potential or overall energy savings.

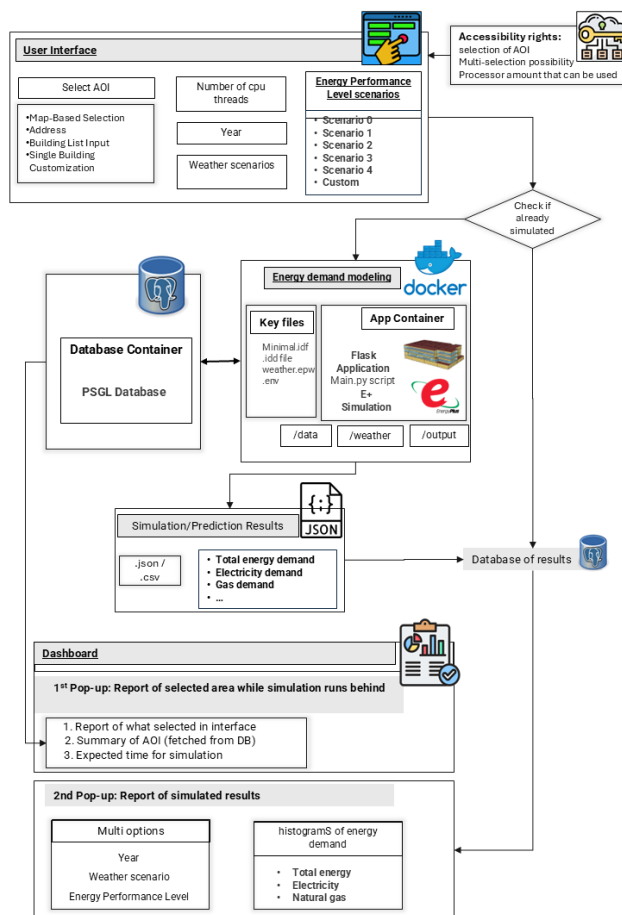


Figure 2. Architecture of the Automated Pipeline Integrating Digital Twin Technology.

2.4.1 Spatial Selection Methods Users can define their area of interest through several approaches. Map-based selection relies on an interactive interface that enables users to outline custom polygons or select administrative boundaries on a digital map. Address or postal code input allows direct specification of geographical regions by entering location details, directing the simulation to one or more discrete areas. Building list input further refines this process by letting users pick specific structures from the DT interface for detailed assessments. Finally, single-building customization supports the simulation of individual buildings under multiple parameter configurations

or scenarios, making it especially suitable for analyzing retrofitting strategies in precise detail. This combination of selection methods ensures that the pipeline can scale from large district planning to targeted, fine-grained analyses.

2.4.2 Building Parameter Customization The pipeline supports a wide range of customization options, offering both static and dynamic settings for simulation scenarios. Static parameters include properties such as insulation levels, window-to-wall ratios, and fundamental material choices, enabling straightforward assessments of building envelopes or compliance with evolving construction standards. Dynamic parameters govern variables like occupancy schedules, equipment usage patterns, and thermostat setpoints, thereby facilitating the modeling of behavioral changes or operational strategies. For calibration or sensitivity analyses, users can specify parameter ranges that define minimum and maximum plausible values, allowing the system to explore uncertainties and their effects on overall energy demand. Additionally, auto-sizing options permit certain parameters to be derived automatically by EnergyPlus based on real-time model conditions, ensuring realistic system sizing. Finally, weather file selection accommodates different climatic conditions, including present-day data or future climate scenarios generated, for example, using KNMI tools in the Netherlands. These extensive customization capabilities make the pipeline adaptable to diverse research objectives, from testing minor improvements to exploring wide-scale urban energy policies.

2.4.3 Scenario Development The pipeline supports the creation of various simulation scenarios, allowing users to evaluate how different interventions and policies affect energy consumption. One commonly used approach is a baseline scenario, which represents the current state of buildings without any improvements or retrofits. This serves as a reference point for comparisons with other scenarios, such as predefined ones based on national standards or specific energy performance levels. For example, Level 0 corresponds to the current energy performance level, while Level 1 reflects the original performance level prior to any upgrades. Levels 2 and 3 encompass common improvement measures, with lower and upper performance thresholds respectively, and Level 4 targets advanced efficiency measures rooted in technically feasible solutions. In addition, users can define energy efficiency measures that incorporate actions like improving insulation, upgrading HVAC systems, or installing double-glazed windows, thus enabling the exploration of potential energy savings. Policy and regulation scenarios model changes in building codes, energy prices, or emissions targets, facilitating policy impact assessments such as the effect of mandating improved insulation for all new constructions. Lastly, user-defined scenarios offer full flexibility to adjust any parameter or assumption for hypothetical what-if analyses, such as switching heating systems to observe shifts in overall energy demand. Because this scenario development process is integrated into an intuitive interface, users of varied expertise levels can configure or select relevant scenarios without modifying the underlying source code.

2.4.4 Simulation Execution and Scalability To handle large-scale simulations across diverse building stocks, the pipeline encloses the necessary software—including EnergyPlus and essential libraries—in containerized environments that maintain consistent performance and simplify cloud deployment. Operated as a web-based API, the system receives configuration details from the user, such as time period and

desired temporal resolution, then generates and executes the corresponding IDF files. This approach supports minute-level precision or annual summaries according to specific project requirements. Parallel processing further improves efficiency by distributing tasks across multiple processor cores, while dynamic resource management adjusts the number of concurrent simulations based on available CPU and memory. As a result, users can simulate hundreds or thousands of buildings concurrently without overloading computing resources. When even larger urban districts or nationwide analyses are required, the pipeline can scale to high-performance computing (HPC) clusters, ensuring that the underlying architecture remains flexible. In addition, the DT platform integrates user-defined access rights for resource control, so permissions can be set for different teams or stakeholders to limit thread counts, define building selection, or govern the number of simulation scenarios.

2.4.5 Visualization and Results Analysis After simulations conclude, the pipeline aggregates key performance indicators such as total energy consumption, peak load, or energy cost savings. These results are then visualized through dashboards and 3D interfaces in the DT environment. Model outputs are plotted directly onto a 3D Tileset, so that users can observe energy consumption patterns for each building in real time as they pan or zoom across the map. This interactive experience allows direct inspection of individual building metrics—such as annual heating demand or peak cooling load—without leaving the digital scene. In addition, users can investigate spatial analysis results to identify clusters of high energy demand or detect specific neighborhoods where certain retrofitting measures may offer the greatest impact. Temporal analysis further enriches decision-making by illustrating how demand fluctuates over hours, days, or seasons, which can inform targeted interventions or grid management strategies. Finally, scenario comparisons enable quick evaluations of multiple interventions or different periods, highlighting the relative effectiveness of each proposed measure in terms of overall consumption and potential energy savings. By presenting these outcomes within an integrated DT framework, the pipeline supports more transparent, data-driven planning and policy development for sustainable urban energy systems.

3. Results and Discussion

This section highlights the outcomes of the research, focusing on the progress made in developing the DT and the associated achievements and challenges. A central innovation of this work is the creation of a DT that encompasses every building in the Netherlands, enabling non-expert users to select specific buildings or districts, configure simulation parameters, and generate energy demand analyses. The model addresses the complexity of district-level simulations by covering a wide range of building types while still delivering timely, detailed results for decision-makers.

3.1 Pilot Project Implementation: Digital Twin Demonstration and Progress

To demonstrate the pipeline's applicability, a pilot project was initiated using a building database that includes all structures in the Netherlands. This database is currently hosted on a dedicated server, while the energy demand modeling component is containerized. At present, simulations are primarily executed on a local laptop, although future plans involve scaling to cloud-based or HPC infrastructures. Figure 3 shows the DT interface used in the pilot, which is still in development. Through

its interactive map, users can select individual buildings or entire districts and then customize parameters such as insulation levels, HVAC configurations, or energy labels. The pipeline returns results in real time, often within minutes, and presents them through dashboards and 3D models for deeper exploration of energy consumption patterns. Despite its ongoing development status, the DT already demonstrates strong potential for handling large datasets and rapidly generating urban-scale simulations.

3.2 Performance Benchmarking

A sequence diagram of user interactions within the pipeline is depicted in Figure 4. Three processes—data retrieval, IDF file generation, and simulation execution—emerged as the most time-consuming components, prompting targeted optimizations. This section details the resulting improvements, focusing on data retrieval strategies and automated IDF generation for EnergyPlus.

3.2.1 Enhanced Data Retrieval and Processing Integrating the DT model with a nationwide building database and predefined scenario tables introduced challenges related to data retrieval times. Whenever a user selects buildings or districts, the system queries a repository containing more than nine million 3D building geometries, then merges the relevant records with archetype data based on the user's chosen configurations. This merged dataset subsequently serves as the foundation for generating IDF files. In early tests, this procedure could take up to a minute, influencing overall responsiveness.

To address this, configuration tables are now embedded directly within the main script as in-memory dataframes rather than being repeatedly pulled from the database. Storing them locally in dataframes cuts down on query overhead and takes advantage of Python's optimized data manipulation libraries, leading to substantial gains in speed. Additionally, updates or parameter changes can be performed within the script itself, avoiding any need to alter the underlying database structure. Although these optimizations have reduced the time to produce IDF files, executing simulations for districts with over 1 000 buildings can still be resource-intensive, and outcomes depend heavily on hardware resources. To balance rapid feedback with the need for comprehensive data, the pipeline offers a two-tiered reporting system. A preliminary report provides immediate, high-level results based on partial computations, followed by detailed results that become available once the full simulation has finished.

3.2.2 IDF File Generation We conducted performance benchmarks on two laptops to evaluate the IDF generation step. The first, a Dell Latitude 5320 featuring an 11th Gen Intel(R) Core(TM) i7-1185G7 processor at 1.80 GHz and 16 GB of RAM, produced an average of 50 IDF files per minute. The second, an Alienware model with a 12th Gen Intel® Core™ i9-12900H CPU at 2.50 GHz and 16 GB of RAM, reached about 90 IDF files per minute. Early tests suggest that shifting to more powerful servers or cloud environments could push these rates even higher. Interestingly, adjusting the CPU thread count did not substantially affect IDF generation, despite low utilization levels in both CPU and I/O. Further investigations are therefore underway to uncover hidden bottlenecks and drive higher throughput. The current implementation uses a `ThreadPoolExecutor` for concurrency, but trials with a `ProcessPoolExecutor` ended in abrupt process terminations, likely due to memory constraints. Future refinements to

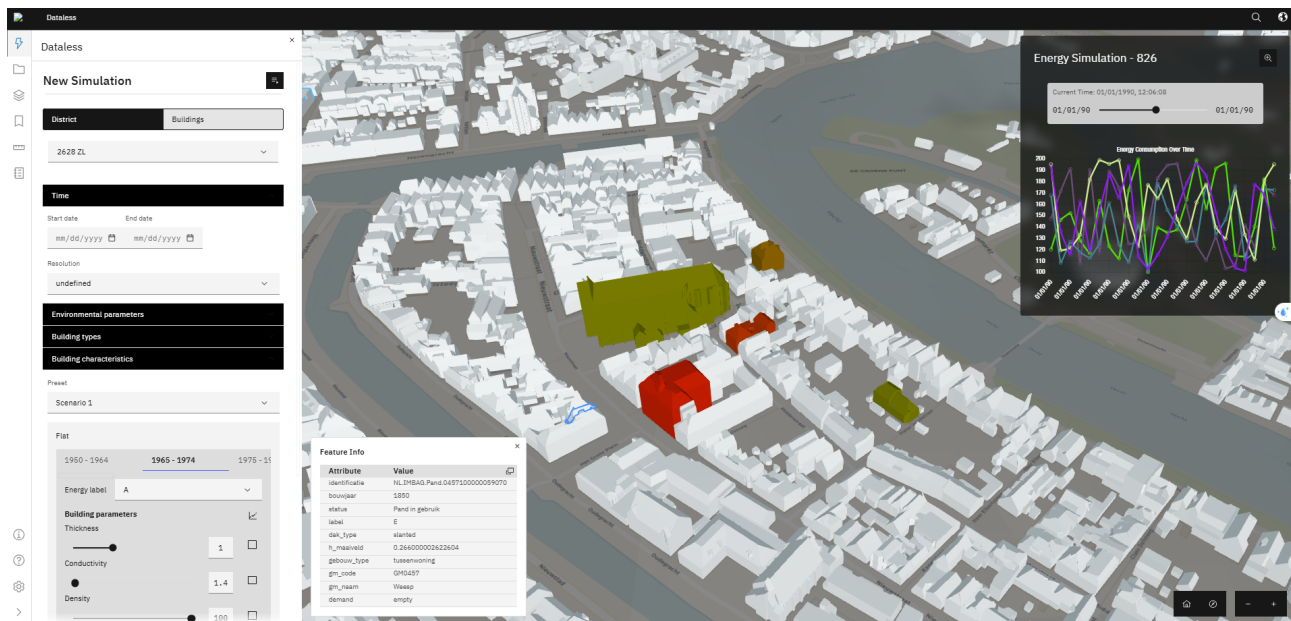


Figure 3. Capture of the Digital Twin project (demo) currently under development.

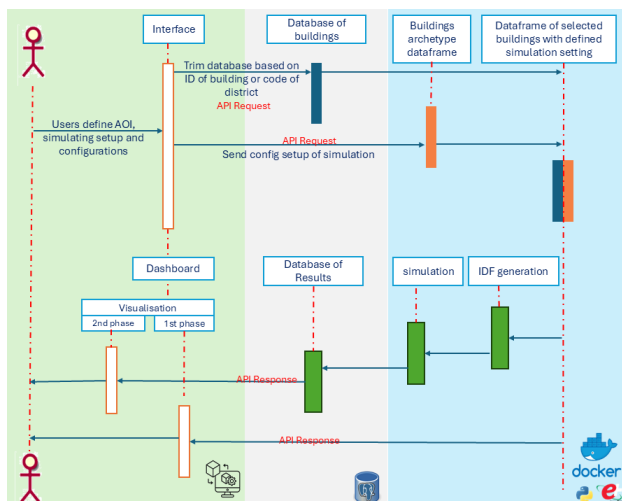


Figure 4. Sequence diagram illustrating an example user interaction within the model.

the codebase will aim to minimize these errors and boost file generation efficiency. Table 1 summarizes representative performance metrics. Abbreviations denote the test device (PC X) and the number of threads (N), with measurements showing the time in seconds to generate a given number of IDF files.

# Bldgs	PC1 (1)	PC1 (4)	# Bldgs	PC2 (1)	PC2 (4)
40	56.3	51.6	47	29.4	29.5
100	128.5	131	96	57.7	67
180	267	237	192	121	117
320	430.4	424.2	288	186	194

Table 1. Performance metrics (in seconds) for IDF generation on two different laptops. Abbreviations: PC X (N) where X is the laptop number and N is the number of worker threads.

3.2.3 Simulation Execution Running EnergyPlus simulations for multiple buildings in parallel presents considerable computational challenges, as large-scale studies can quickly exceed the capacity of single-threaded processing. While in-

creased parallelization is a common route to improved performance, this research also emphasizes containerization for seamless deployment across local or cloud-based servers. Containerization allows the model to exploit available processing resources fully, regardless of platform. Performance tests on a high-end laptop (referred to as “laptop 2”) provide an example of how thread counts influence simulation runtime. In one trial, simulating 55 buildings with eight workers took 55 seconds, whereas using only two workers increased the runtime to 776 seconds. Although laptop 2 supports up to 22 threads, testing was capped at eight to avoid surpassing local hardware limits. In practice, the model can scale to more powerful servers or cloud architectures if faster turnaround times are needed. On the other hand, adding more complexity or detail to the IDF files may increase simulation overhead. Currently, the pilot project separates IDF generation from the simulation phase, producing IDFs before sequentially running simulations. Future plans involve a more concurrent approach that overlaps these steps, aiming to streamline data preparation, simulation runs, and visualization for larger study areas. By processing IDFs, simulation tasks, and visualization outputs simultaneously, the pipeline could reduce total execution time while enhancing overall efficiency.

Threads	Simulations Number	Duration (s)
8	12	75
3	12	105
8	55	403
3	55	776
8	192	1414
2	192	3149
8	286	1766
5	286	2431
2	286	4894
8	432	2487
2	432	6420
5	440	3807

Table 2. Simulation duration by number of threads and total simulations (all times in seconds).

3.3 Advanced Features and Integration with Other Models

Beyond basic parallelism, the Python-based hyperparameter model developed for district-scale energy simulations includes capabilities in integration, optimization, calibration, and user-defined flexibility. A fundamental advantage is the capacity to define parameter ranges, which supports both optimization workflows and calibration against real-world data. By specifying ranges for key attributes—such as insulation thickness or HVAC setpoints—the model can systematically search for configurations that yield high energy efficiency, while calibration benefits from iterative adjustments that bring simulated outputs closer to observed consumption patterns. This functionality also underpins sensitivity analyses, allowing researchers to quantify the impact of individual parameters on overall demand. Another innovative feature is the introduction of controlled randomness in building energy demand. Recognizing that occupant behavior, equipment efficiency, and other factors vary substantially among buildings, the model applies random perturbations within specified bounds. This yields more realistic district-wide simulations, helping planners and grid operators assess supply-demand balances under diverse conditions. Retrofitting measures likewise receive special attention: users can select individual buildings, apply changes to insulation levels or window types, and view results in a concise tabular format that simplifies comparison. To further aid decision-making, the model evaluates each building against performance benchmarks, pinpointing those that require the most urgent improvements.

By uniting parameter flexibility, random demand variations, and retrofit analysis, the hyperparameter model addresses both short-term tasks—such as calibrating a specific site—and long-term strategies aimed at citywide energy efficiency. These features collectively facilitate a more balanced approach to managing supply and demand while offering stakeholders robust tools for planning sustainable urban energy systems.

4. Conclusion

This study demonstrates that automated workflows that integrate EnergyPlus with Python-based tools and DT technology can feasibly scale building energy simulations from single structures to entire districts. By leveraging open-access datasets in the Netherlands, the proposed pipeline reduces the time and effort required for modeling energy demand and enables flexibility in parameter adjustments and scenario design. Key findings suggest that this approach simplifies large-scale simulations, while its real-time visualization and interactive features support informed decision-making—both of which are essential for sustainable urban planning.

Developing the pipeline involved a number of challenges, particularly in obtaining reliable, detailed data. While residential building datasets in the Netherlands are relatively complete and straightforward to incorporate, non-residential data proved more fragmented. Researchers were thus required to piece together schedules, occupancy profiles, and other information from various sources, which remains a time-consuming endeavor. Nonetheless, because residential buildings make up around 90% of the Dutch building stock, the pipeline already delivers reliable simulations at district scales. For non-residential buildings, current data gaps allow for insightful—albeit less granular—analysis. In addition, scripting com-

plex HVAC systems highlighted the need for reverse engineering existing files, illustrating the complexities of customizing advanced EnergyPlus objects. Once developed, however, these scripts can be reused to expedite future modeling efforts. The model currently treats each floor as a single zone, an assumption that may underrepresent energy demand for larger or more intricate buildings. While this approach works well for typical residential structures, improving geometry handling and zoning algorithms is a priority for enhancing simulation accuracy. Another identified challenge stems from the abundance of available datasets. Integrating, cleaning, and verifying multiple standards is labor-intensive and can result in duplicated effort among researchers. This finding underscores a broader call for more unified, well-documented data infrastructures that centralize the necessary energy-related information.

Addressing the wide variety of archetypes likewise proved complex. Dutch buildings span many age groups, functional categories, and energy labels, each with several possible retrofit or performance scenarios. Identifying the correct EnergyPlus objects and parameter values requires extensive consultation of national standards or best-guess assumptions. Although energy labels offer valuable indications of performance, only about five million buildings have such labels on record. In future work, we aim to extend these labels to a larger proportion of the building stock, thereby improving the model's precision and consistency.

Looking ahead, efforts will focus on addressing these limitations to further refine the pipeline's accuracy and utility. Upgrading the data for non-residential buildings is a priority to ensure a more complete representation of the Dutch building sector. Engaging urban planners, engineers, and policymakers in iterative feedback loops will also guide enhancements to the pipeline's interface and features. Additional Python-based libraries, such as the EnergyPlus API or PyEnergyPlus, may streamline aspects of simulation control. Moreover, incorporating real-time weather data, vegetation layers, and infrastructure networks can bolster analyses of renewable energy, solar irradiance, and wind patterns. Plans also involve integrating on-site energy generation, storage systems, and optimization strategies for balancing supply and demand in sustainable districts, thereby extending the pipeline's scope as a comprehensive energy management solution.

Code and Digital Twin Availability

The source code supporting this work is currently under active development and can be accessed on GitHub at https://github.com/amin-jalilzadeh-tu/E_Plus_2030_py. A demo demonstration of the Digital Twin platform is available at <https://dataless.beta.geodan.nl/>.

Acknowledgment

This research received funding from the Dutch Research Council (NWO) and National Natural Science Foundation of China (NSFC) in the framework of the Cooperation China-The Netherlands programme. The author would also like to extend his gratitude to the research group at GEODAN, with special thanks to Henk Scholten, Corentin Kuster, Anne Blankert, Brian de Vogel, and Yneke van Ierse for their invaluable support and assistance in the development of the DT. Additionally, Prateek Wahi's help in the creation of the energy demand model is highly appreciated.

References

- 3D BAG Dataset, n.d. <https://3dbag.nl/en/download>. Accessed: 2024-12-01.
- Ali, U., Shamsi, M. H., Hoare, C., O'Donnell, J., 2018. Gis-based residential building energy modeling at district scale. *Proceedings of the 4th IBPSA-England Conference on Building Simulation and Optimization, Cambridge, UK: 11-12 September 2018*, International Building Performance Simulation Association.
- Allegrini, J., Orehounig, K., Mavromatidis, G., Ruesch, F., Dorer, V., Evins, R., 2015. A review of modelling approaches and tools for the simulation of district-scale energy systems. *Renewable and Sustainable Energy Reviews*, 52, 1391-1404. <https://www.sciencedirect.com/science/article/pii/S1364032115007704>.
- Ascione, F., Bianco, N., Mauro, G. M., Napolitano, D. F., 2021. Knowledge and energy retrofitting of neighborhoods and districts. A comprehensive approach coupling geographical information systems, building simulations and optimization engines. *Energy Conversion and Management*, 230, 113786. <https://www.sciencedirect.com/science/article/pii/S0196890420313091>.
- BAG Viewer, n.d. <https://bagviewer.kadaster.nl/lvbag/bag-viewer>. Accessed: 2024-12-01.
- CBS: Central Bureau of Statistics, n.d. <https://www.cbs.nl/>. Accessed: 2024-12-01.
- Crawley, D. B., Lawrie, L. K., Winkelmann, F. C., Buhl, W., Huang, Y., Pedersen, C. O., Strand, R. K., Liesen, R. J., Fisher, D. E., Witte, M. J., Glazer, J., 2001. EnergyPlus: creating a new-generation building energy simulation program. *Energy and Buildings*, 33(4), 319-331. <https://www.sciencedirect.com/science/article/pii/S0378778800001146>. Special Issue: BUILDING SIMULATION'99.
- Deng, Z., Chen, Y., Yang, J., Causone, F., 2023. AutoBPS: A tool for urban building energy modeling to support energy efficiency improvement at city-scale. *Energy and Buildings*, 282, 112794. <https://www.sciencedirect.com/science/article/pii/S0378778823000245>.
- Enriched BAG to Support Local Energy Issues, n.d. <https://energy.nl/publications/verrijkte-bag-energetische-vraagstukken/>. Accessed: 2024-12-01.
- EP-Online: Energy Performance Certificates Database, n.d. <https://www.rvo.nl/onderwerpen/wetten-en-regels-gebouwen/ep-online>. Accessed: 2024-12-01.
- Faure, X., Johansson, T., Pasichnyi, O., 2022. The Impact of Detail, Shadowing and Thermal Zoning Levels on Urban Building Energy Modelling (UBEM) on a District Scale. *Energies*, 15(4). <https://www.mdpi.com/1996-1073/15/4/1525>.
- Glazer, J., 2016. Using python and eppy for a large national simulation study. *ASHRAE & IBPSA-USA SimBuild Conference 2016*, 7, ASHRAE/IBPSA-USA, 230–237.
- Guerra-Santin, O., Silvester, S., 2017. Development of Dutch occupancy and heating profiles for building simulation. *Building Research & Information*, 45(4), 396–413. <https://doi.org/10.1080/09613218.2016.1160563>.
- Muslim, S. A., 2021. EnergyPlus-Towards the selection of right simulation tool for building energy and power systems research. *Journal of Energy and Power Technology*, 3(3), 1–9.
- NTA 8800 Standard, n.d. <https://www.nen.nl/en/nta-8800-2024-nl-320123>. Accessed: 2024-12-01.
- PBL: Netherlands Environmental Assessment Agency, n.d. <https://www.pbl.nl/>. Accessed: 2024-12-01.
- PDOK: Public Services on the Map, n.d. <https://www.pdok.nl/>. Accessed: 2024-12-01.
- Report on Standard and Target Values for Existing Residential Construction, n.d. <https://www.rijksoverheid.nl/documenten/rapporten/2021/02/28/rapport-standaard-en-streefwaardes-bestaande-woningbouw>. Accessed: 2024-12-01.
- RIVM: National Institute for Public Health and the Environment, n.d. <https://www.rivm.nl/>. Accessed: 2024-12-01.
- Robinson, D., Haldi, F., Leroux, P., Perez, D., Rasheed, A., Wilke, U., 2009. Citysim: Comprehensive micro-simulation of resource flows for sustainable urban planning. *Proceedings of the Eleventh International IBPSA Conference*, 1083–1090.
- RVO: Netherlands Enterprise Agency, n.d. <https://www.rvo.nl/>. Accessed: 2024-12-01.
- Simpson, J. G., Long, N., Zhu, G., 2024. Decarbonized district energy systems: Past review and future projections. *Energy Conversion and Management: X*, 24, 100726. <https://www.sciencedirect.com/science/article/pii/S2590174524002046>.
- Sola, A., Corchero, C., Salom, J., Sanmarti, M., 2018. Simulation Tools to Build Urban-Scale Energy Models: A Review. *Energies*, 11(12). <https://www.mdpi.com/1996-1073/11/12/3269>.
- Tahmasebinia, F., Lin, L., Wu, S., Kang, Y., Sepasgozar, S., 2023. Exploring the Benefits and Limitations of Digital Twin Technology in Building Energy. *Applied Sciences*, 13(15). <https://www.mdpi.com/2076-3417/13/15/8814>.
- Voorbeeldwoningen: Example Dwellings for Dutch Buildings, n.d. <https://www.rvo.nl/onderwerpen/wetten-en-regels-gebouwen/voorbeeldwoningen-bestaande-bouw>. Accessed: 2024-12-01.
- Wahi, P., Konstantinou, T., Visscher, H., Tenpierik, M. J., 2024. Evaluating building-level parameters for lower-temperature heating readiness: A sampling-based approach to addressing the heterogeneity of Dutch housing stock. *Energy and Buildings*, 322, 114703. <https://www.sciencedirect.com/science/article/pii/S0378778824008193>.
- Yue, N., Li, L., Morandi, A., Zhao, Y., 2021. A metamodel-based multi-objective optimization method to balance thermal comfort and energy efficiency in a campus gymnasium. *Energy and Buildings*, 253, 111513. <https://www.sciencedirect.com/science/article/pii/S0378778821007970>.