High resolution solar potential computation in large scale urban areas by means of semantic 3D city models

Longxiang Xu[1], Camilo León-Sánchez[2], Giorgio Agugiaro[2], Jantien Stoter[2]

[1] Delft University of Technology, L.Xu-18@student.tudelft.nl
[2] 3D geoinformation group, Delft University of Technology, Faculty of Architecture and the Built Environment, Department of Urbanism, Julianalaan 134, 2628BL Delft, The Netherlands, c.a.leonsanchez,g.agugiaro,j.e.stoter@tudelft.nl

Keywords: 3DCM, Solar Potential, Shadows analysis, Raytracing

Abstract

Solar energy is becoming increasingly important with the transition towards green and sustainable energy. Predicting solar irradiance is one of the core steps to optimise solar energy utilisation when planning and scheduling power grids. Accurate solar irradiance prediction can also help forecast microclimate conditions, enabling the analysis of citizens and planning of optimal intervention strategies for heating or cooling behaviour. This paper discusses a novel approach to computing the solar potential of buildings at the city level with promising scalability using semantic 3D city models. Experiments are conducted at different locations in the Netherlands. We evaluate our results by comparing them to the statistical Dutch data, and CitySim shows huge discrepancies in summer.

1. Introduction

The World Bank reports that in 2023, 56% of the global population resides in urban areas, with projections rising to 70% by 2050 (The World Bank, n.d.). This surge in urban population density poses significant challenges for energy and climate systems in cities. Effective urban planning and market adjustments, such as in the carbon market, are vital for accurately quantifying energy demand and supply. Misestimates can result in better decisions and better planning in energy systems. Accurately determining solar irradiance is essential, as it influences distributed solar power grids, environmental conditions, and occupant behaviour. For instance, high solar irradiance can increase indoor temperatures, raising cooling demands. Thus, precise estimation of solar radiation for specific areas is crucial.

However, computing the solar radiation hitting surfaces of urban objects (e.g. a building, a shed, or the solar panels placed on top of them) is rather challenging due to the peculiar geometrical characteristics of the study object (surface tilt and inclination), the overall topography of the site that may lead to shadowing, and also the solar position. As a result, the solar irradiance on two similar surfaces in a city can be very different, even if they are very close.

It is crucial to account for geographic location and surroundings to analyse solar potential in complex urban environments effectively. Semantic 3D city models [3DCM] provide a structured dataset that offers both geometric and semantic representations of urban features, crucial for accurate solar irradiance calculations on various surfaces, including factors like orientation, tilt, and shadow effects (Agugiaro et al., 2020).

This paper introduces a workflow for computing solar irradiance at the city level with high temporal resolution using 3DCM. We utilise the 3DBAG dataset (Peters et al., 2022), which contains 3D models of all buildings in the Netherlands, available in multiple formats and levels of detail [LoD] (Biljecki et al., 2016). Specifically, our workflow processes data in the CityJSON format (Ledoux et al., 2019), although it is compatible with any CityJSON-encoded dataset. The result is a highly detailed output of solar irradiance values, achieving sub-meter spatial and hourly temporal resolutions. This approach notably avoids assumptions about urban morphology while remaining computationally efficient for consumer-level devices.

2. Method

The computation of the solar potential involves two main sections: shadowing and solar irradiance. The workflow diagrams depicting the processes are presented in fig. 1.

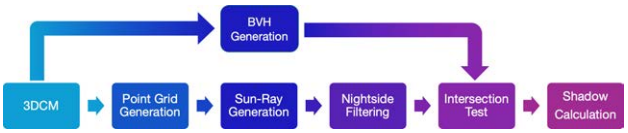

(a) Shadowing calculation

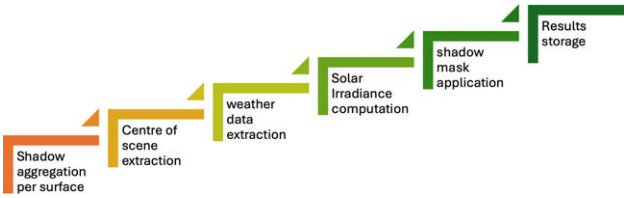

(b) Solar Irradiance calculation

Figure 1. General overview of the workflows

Our method starts by extracting the centre coordinates of the input 3DCM. These coordinates are necessary to calculate the sun's position throughout a typical year, recording those hours when the sun is above the horizon. Additionally, We compute the orientation and inclination of all boundary surfaces of the buildings. Then, surfaces are progressively divided into smaller triangles to ensure uniform samples and a more detailed analysis. To deal with sliver triangles, we divide them further by

the corner of the acute angle; more details can be found at (Xu et al., 2024).

2.1 Shadow Calculation

The five-step shadow computation process (fig. 1(a)) starts with generating a grid on building surfaces and extracting barycentric coordinates. Sun-ray vectors are then created, pointing in the sunlight's direction. Concurrently, nightside filtering ((Wang et al., 2023)) removes self-occluded surfaces, and a Bounding Volume Hierarchy (BVH) ((Meister et al., n.d.)) is constructed. The final step is an intersection test between sun-rays and the BVH to identify shadowed surfaces.

2.2 Solar Irradiance Calculation

In this research, we use the weather data from the locations of the weather stations of the Royal Netherlands Meteorological Institute (KNMI, 2024). Figure 2 presents the main components of the solar irradiance values. We use the solar beam and the horizontal diffuse radiation components in our computations. We merge each point in the grid with the irradiance vector, resulting in a time series of values.

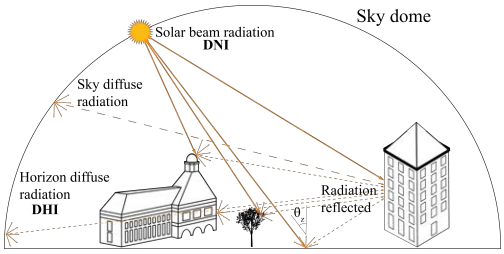

Figure 2. Solar irradiance component

It is essential to utilise a transition model based on meteorological data to analyse the solar irradiance impacting tilted surfaces. This data provides measurements of solar irradiance that reach the ground, and these measurements are expressed by eq. (1) (Loutzenhiser et al., 2007a).

$$GHI = DHI + DNI \cdot \cos(\theta_Z) \tag{1}$$

Global Horizontal Irradiance [GHI], Direct Normal Irradiance [DHI], and Diffuse Horizontal Irradiance [DNI], θ_Z represents the solar zenith angle. DNI is divided into sky and horizon diffuse components fig. 2. The general form of the transition model is formulated as eq. (2).

$$I_{S,\beta} = I_{S,dir,\beta} + I_{S,diff,\beta} + I_{S,refl,\beta} \tag{2}$$

Where $I_{S,dir,\beta}$, $I_{S,diff,\beta}$, $I_{S,refl,\beta}$ represent the direct beam irradiance, sky diffuse irradiance, and ground reflected irradiance. $I_{S,\beta}$ represent the total solar irradiance and β stands for surface inclination. The determination of the direct beam solar irradiance is articulated as eq. (3).

$$I_{S,dir,\beta} = M_{shadow} \cdot DNI \cdot \cos\delta \tag{3}$$

M_{shadow} represents the binary shadow mask. In this study, we aggregate the shadow mask for all sample points on each surface and calculate the average value of the mask. The shadow mask value ranges from 0 to 1, where 0 indicates that the surface is completely shadowed.

The methods to estimate the sky diffuse component of total solar irradiance have been extensively studied over the past decades. In general, different models require varying amounts of input data. In our study, we opt for the Isotropic (Kamphuis et al., 2020) and the Perez (Perez et al., 1987, Perez et al., 1988, Perez et al., 1990) models. *Isotropic model* considers the sky a uniform diffuse radiation source. The motivation for choosing this simple model is that our study focuses more on accurately estimating direct beam solar irradiance, which is significantly affected by shadows. The diffuse components of solar irradiance in the isotropic model can be formulated by eq. (4).

$$I_{S,diff,\beta} = DHI\frac{1+\cos\beta}{2} \tag{4}$$

The ground-reflected solar irradiance in a similar simplified setting can be determined as follows:

$$I_{S,refl,\beta} = I_{S,tot} \cdot \gamma_{refl} \cdot \left(\frac{1-\cos\beta}{2}\right) \tag{5}$$

Perez model is an empirical model adopted widely and tested to show promising alignment with ground truth data (Loutzenhiser et al., 2007b). Compared with the Isotropic, Perez requires more parameters such as airmass. Eq. 6 details the basic form of this model (Loutzenhiser et al., 2007a):

$$E_d = DHI \times \left[(1-F_1)\left(\frac{1+\cos(\theta_T)}{2}\right) + F_1\left(\frac{a}{b}\right) + F_2\sin(\theta_T)\right] \tag{6}$$

Where F_1 eq. (7) and F_2 eq. (8) are complex empirically fitted functions that describe circumsolar and horizon brightness, respectively. Additionally:

- $a = \max(0, \cos(AOI))$
- $b = \max(\cos(85°), \cos(\theta_Z))$
- AOI is the angle of incidence between the sun and the plane of the array
- θ_T is the array tilt angle from horizontal

$$F_1 = \max\left[0, \left(f_{11} + f_{12}\Delta + \frac{\pi\theta_Z}{180°}f_{13}\right)\right] \tag{7}$$

$$F_1 = \max\left[0, \left(f_{11} + f_{12}\Delta + \frac{\pi\theta_Z}{180°}f_{13}\right)\right] \tag{8}$$

The f coefficients are defined for specific bins of clearness (ε), which is defined as eq. (9):

$$\varepsilon = \frac{(DHI + DNI)/DHI + \kappa\theta_Z^3}{1 + \kappa\theta_Z^3} \tag{9}$$

where: κ is a constant equal to 1.041 for angles in radians, or 5.535×10^{-6} for angles in degrees.

3. Experiments and Results

For each location, we calculate the hourly shadowing and solar irradiance.

3.1 Shadow Calculation

Implemented in C++, our shadow calculation leverages BVH construction from the book "Ray Tracing: The Next Week" ((Ray Tracing in One Weekend Book Series, 2020)). The model operates on Ubuntu 20.04 through WSL on Windows 11, utilizing a computer with an Intel(R) Core(TM) i7-12700H CPU and 32GB RAM. We enhance performance by parallelizing the process across 20 CPU threads.

All simulations are applied on three different test locations in the Netherlands: Groningen, Maastricht, and Utrecht see fig. 3. These cities, distributed across the Netherlands, exhibit a variety of urban morphologies. The corresponding 3DBAG tiles and their details are listed in table 1. Our experiments are computed for 2023.

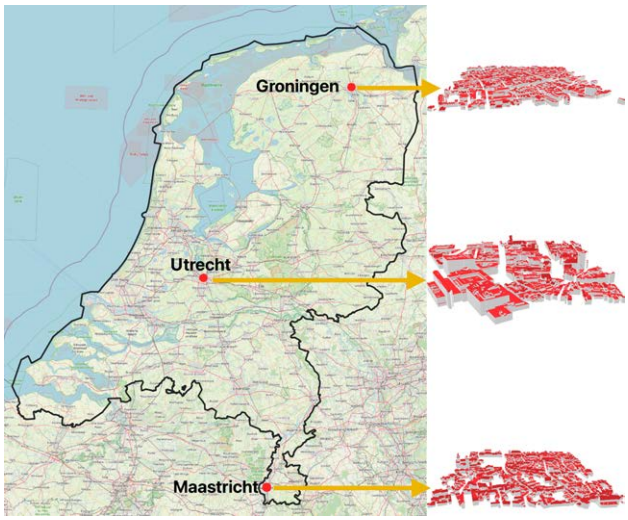

Figure 3. Study Areas

Our initial filtering is based on the sun's location above the horizon, which leads to ~ 4.400 hours (epochs) in which the sun is visible, instead of the $8.760h$ of a typical year (365×24).

The experiment involves shadowing calculation using the LoD2.2 data per tile. This calculation is performed on the roofs and building facades, leveraging the semantic information provided by the input 3DCM. Table 2 displays the results of the shadow calculations for all the selected tiles.

An additional experiment was conducted on tile 10-880-109, which is the primary area for shadow calculation. In contrast, all surrounding tiles were used solely for intersection tests (see Groningen(+) in Tables 1 and 2), using 3D building models of the same level of detail (LoD 2.2). We perform this experiment to include the buildings that may lie close to the borders of the tile that is being analysed. For that reason, the tile of interest (*10-880-109*) and its surrounding tiles are included in the BVH (see fig. 1(a)). By doing this, we perform two analyses: the influence of the surrounding buildings and the scalability of our method.

Table 1. Study Area Information. Columns indicate the number of buildings, surfaces, and computation hours respectively.(+) indicates the inclusion of the 8 neighbouring tiles.

Tile	City	Building Num	Surface Num	Hours Num	Extent (meter)
10-880-1098	Groningen	925	84581	4462	500×500
9-656-44	Maastricht	3048	82882	4443	1000× 1000
10-490-596	Utrecht	664	73944	4459	500× 500
10-880-1098(+)	Groningen(+)	5521	469867	4462	1500× 1500

Table 2. Experiment results. Where "PG Num" number of grid points; "Duration" total computation time; "D/Pt" computation time per grid point for a year. Other columns follow the definitions in Table tab:tileinfo.

Tile	Building Num	Surface Num	PG Num	Duration (s)	D/Pt (ms)
Groningen	925	84581	525270	5234	9.96
Maastricht	3048	82882	518757	3137	6.05
Utrecht	664	73944	484247	4802	9.92
Groningen(+)	5521	469867	525270	10398	19.80

3.2 Solar Irradiance Calculation

Hourly solar irradiance values for the study areas are calculated using the Python package pvlib (Holmgren et al., 2018). We use two weather input datasets for our experiments. The first dataset is the API PVGIS (European Commission, 2024), which downloads the weather data for typical meteorological years as a typical meteorological year (TMY3) files from a given set of coordinates (obtained from the centre of the tile). These data contain the required parameters for the *isotropic model* computation. We use the built-in functions of pvlib to calculate the extra required parameters of the *Perez model*. The second dataset corresponds to the data from the *climate.OneBuilding.Org* project (Lawrie and Crawley, n.d.), which offers climate data designed to support building simulations such as EnergyPlus files (EPW) available from a variety of organisations and countries (in our case, KNMI).

We calculate the solar irradiance for each point in the computed grid point of the shadowing calculation. further details are available at (Xu et al., 2024). The results are aggregated per surface using a weighted average. Moreover, we compute a shadowing mask per surface for each epoch. The resulting value is applied as a correction factor of the direct beam solar irradiance as outlined in eq. (3). Finally, we aggregate the solar irradiance values for each month, as shown in Figure 10, so we can compare them with the statistical values set by the Dutch standard NTA8800:2023 (NEN, 2024); this standard specifies the method for assessing the energy performance of buildings in the Netherlands. The statistical values provided in the standard are categorised by the orientation and inclination of the boundary surfaces that define the thermal hull. Since our method computes these values, the results of the experiments are comparable to the official ones.

To further evaluate our method. We perform solar potential analysis using CitySim (Emmanuel and Jérôme, 2015). It is a building energy simulation tool (BES) that models building performance to compute, among other applications, the solar potential or the energy demand for heating/cooling buildings at a city level. Although it uses its own data format (XML CitySim), a graphical user interface (GUI) named CitySim Pro supports importing from other data formats such as dxf or CityGML. We

use the latter as the input dataset to perform the solar potential analysis by transforming the CityJSON files from the 3DBAG using citygml-tools (Nagel and Deininger, 2024).

4. Results

Figure 5 to fig. 9 show scatter plot as a comparison of the results obtained by the *Perez model* and the NTA8800 values at each of the test locations. Each point corresponds to a surface in the study area.

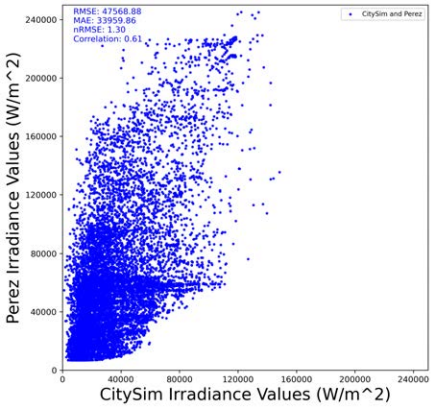

Figure 4. Scatter plot of Monthly Solar Irradiance Values in Groningen between Perez model results and CitySim results using the EPW as input weather data

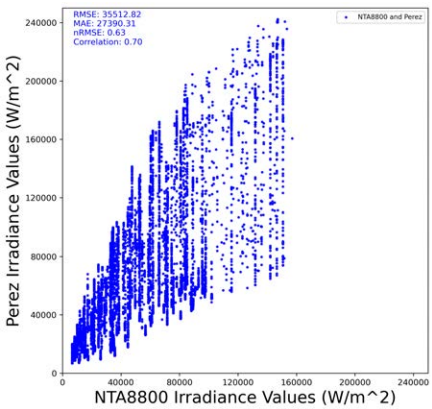

Figure 5. Scatter plot of Monthly Solar Irradiance Values in Groningen between Perez model results and NTA8800 using the EPW as input weather data

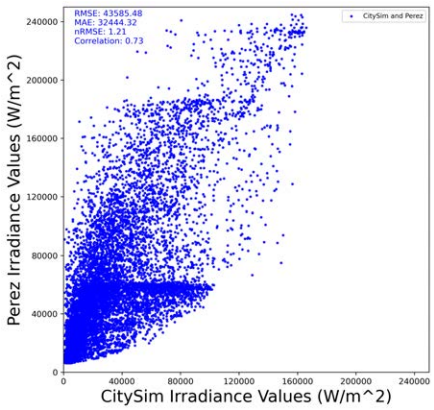

Figure 6. Scatter Plot of Monthly Solar Irradiance Values in Utrecht between Perez model results and NTA8800 using the EPW as input weather data

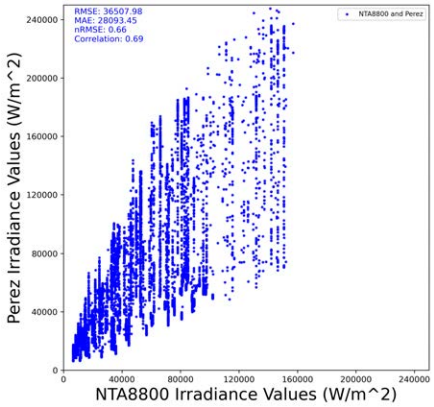

Figure 7. Scatter Plot of Monthly Solar Irradiance Values in Utrecht between Perez model results and NTA8800 using the EPW as input weather data

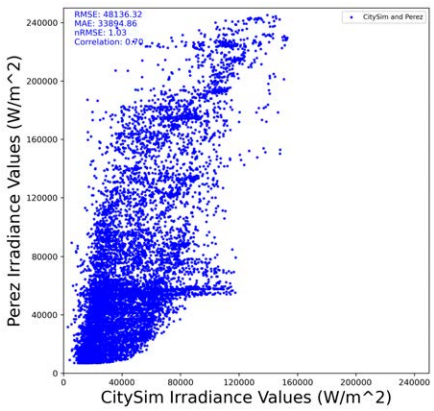

Figure 8. Scatter Plot of Monthly Solar Irradiance Values in Maastricht between Perez model results and CitySim results using the EPW as input weather data

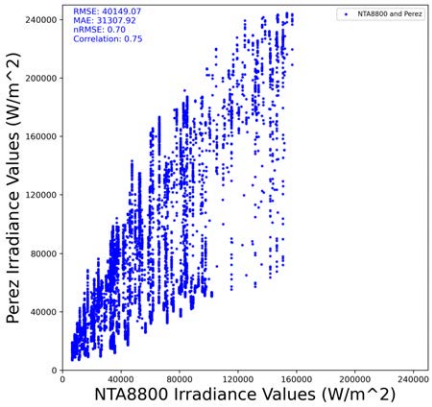

Figure 9. Scatter Plot of Monthly Solar Irradiance Values in Maastricht between Perez model results and NTA8800 using the EPW as input weather data

The correlation coefficients in the scatter plots (fig. 4 -fig. 9) vary from 0.66 to 0.81, indicating a moderate to strong linear relationship between our method and the NTA8800 norm. Despite potential differences arising from variations in shadow modelling, irradiance calculation approaches, and meteorological data sources, there is an alignment in the patterns of irradiance captured by both methods.

From fig. 10 to fig. 15 show bar diagrams of the average solar irradiance values per test location. Values are presented per month so that the user can better compare our method against the norm values.

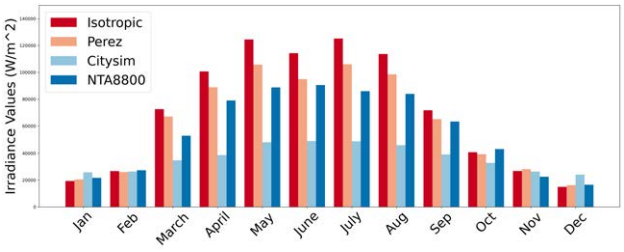

Figure 10. Groningen, bar plot of Monthly Solar Irradiance Values comparison using TMY3 as input weather data

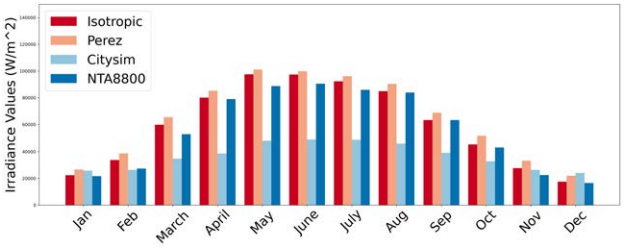

Figure 11. Groningen, bar plot of Monthly Solar Irradiance Values comparison using EPW as input weather data

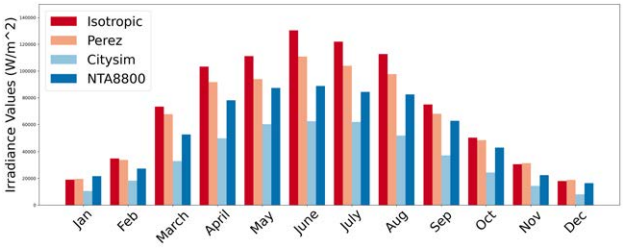

Figure 12. Utrecht, bar plot of Monthly Solar Irradiance Values comparison using TMY3 as input weather data

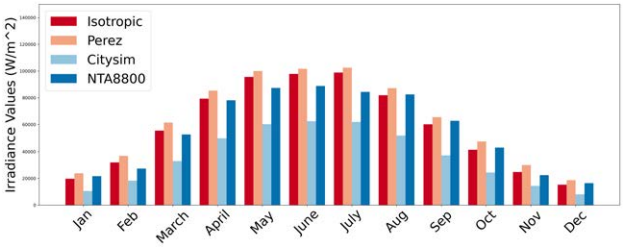

Figure 13. Utrecht, bar plot of Monthly Solar Irradiance Values comparison using EPW as input weather data

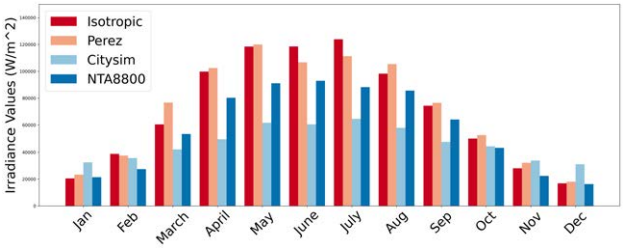

Figure 14. Maastricht, bar plot of Monthly Solar Irradiance Values comparison using TMY3 as input weather data

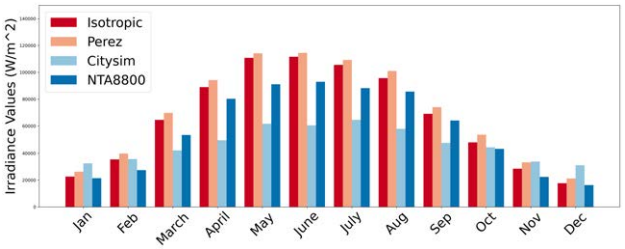

Figure 15. Maastricht, bar plot of Monthly Solar Irradiance Values comparison using EPW as input weather data

CitySim average values are usually the lowest in the three study area locations. In the case of Groningen with the EPW3 data, the values can be 40% lower compared to the other computing methods, but those values are the highest in December. For all locations, our computation values using the TMY3 dataset are significantly higher, showing us the impact the input weather dataset can have in the final output. The difference between our average results is significant in the case of the NTA8800 norm statistical values. However, it is essential to note that the norm values are statistical approaches that do not consider urban morphology as our method.

Error metrics and the bar plots reveal that warmer seasons exhibit more significant discrepancies between NTA8800 values

and our computed values. More experiments and analyses need to be conducted to reveal the reason for the phenomena. However, the solar irradiance estimation between historical standard and high-resolution simulation can significantly differ.

Finally, fig. 16 to fig. 19 show the solar irradiance per surface in May and December in the test location Utrecht. All images use the same colour palette from blue (lowest value) to yellow (highest value).

(a) May

(b) December

Figure 16. Utrecht. Isotropic model output visualisation

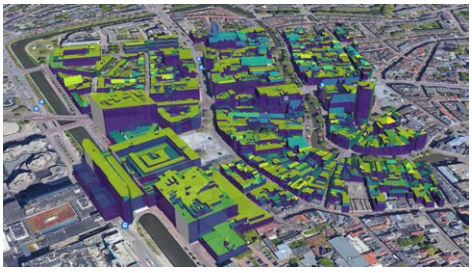

(a) May

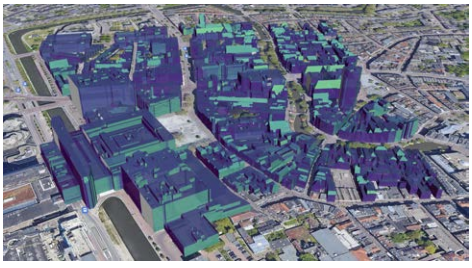

(b) December

Figure 17. Utrecht. Perez model output visualisation

(a) May

(b) December

Figure 18. Utrecht. NTA8800 statistical values visualisation

(a) May

(b) December

Figure 19. Utrecht. CitySim output visualisation

5. Discussion

Through simple CPU parallel processing, it takes less than 20 milliseconds to compute shadows annually per urban point. Our method efficiently handles additional shadow-casting surfaces without significantly affecting ray computation time, demonstrating scalability across diverse urban environments.

Table 2 shows that the computation time per ray—ranging from 6 to 9 milliseconds— does not positively correlate with the

number of shadow-casting surfaces. This may be due to overlapping BVH nodes, where rays intersect multiple child nodes, reflecting scene-specific conditions. Yet, these results align with the $O(\log n)$ theoretical time complexity, affirming the efficiency and scalability of using BVH for shadow computation.

Employing BVH means that our vector tracing approach avoids assumptions typical in other shadow projection techniques, such as fixed ratios relative to the height of surrounding buildings. BVH accurately handles complex urban and mountainous terrains by integrating digital terrain models and ensures precise shadow calculations on minor features such as corners, proving versatile for indoor and outdoor environments.

The current shadow computation method uses basic CPU parallel processing and is highly parallelisable, indicating significant potential for faster computations. Given the independence of ray intersection tests and the capabilities of modern GPU cores, this approach could see a several hundred-fold increase in speed.

The current experiments use limited semantic data on buildings, but this is expandable. Recording intersections will enable analysis of occlusion patterns, and incorporating more semantics will help focus computations on selected urban elements such as buildings, trees, and city furniture. The comparison of irradiance values between the NTA8800 and those estimated using our method highlights significant differences between purely statistical approaches and analytical methods that account for urban morphology.

The correlation coefficients $(0.61 - 0.75)$ show a moderate to strong linear relationship between our method and outputs from NTA8800 or CitySim, indicating a basic alignment despite differences in shadow modelling, irradiance calculations, and meteorological data. Error metrics and bar plots (fig. 10 - fig. 15) indicate more significant discrepancies in summer compared to winter. During winter, shadowing lessens energy loss due to prevalent overcast conditions blocking direct solar irradiance (Tuononen et al., 2019). Consequently, NTA8800 values, which ignore shadow effects, align more closely with our output computations in winter. In contrast, clear skies in summer accentuate shadowing effects, widening the gap between NTA8800 values and our estimates.

One limitation of our method is its basic handling of light ray diffusion and reflection, as shown in fig. 2. Our approach simplifies the treatment of these factors. However, it still has the potential to accurately calculate the Sky View Factor, which is crucial for estimating diffuse and ground-reflected irradiance. Enhancements in the ray-triangle intersection test pipeline could achieve this accuracy. However, due to limited parallelization capabilities in the current development, using the vector-tracing pipeline to calculate the Sky View Factor would demand excessive computational resources. Instead, employing existing rendering techniques that simulate the fish-eye camera view could be a viable alternative for computing the Sky View Factor, as suggested by (Hämmerle et al., 2011, Miao et al., 2020).

6. Conclusions and future work

Our proposed method enhances shadow calculation using semantic 3D city models and vector tracing, avoiding reliance on culling assumptions. It incorporates nightside filtering to exclude computations for surfaces not facing the sun efficiently.

Our method's rudimentary handling of light ray diffusion and reflection is a limitation, as depicted in fig. 2. While this approach simplifies scenarios, it still holds promise for precisely calculating the Sky View Factor, which is crucial for accurate diffuse and ground-reflected irradiance estimates. Enhancing the ray-triangle intersection test pipeline could achieve this precision. However, due to current development constraints on parallelization, using the vector-tracing pipeline for this calculation would be computationally demanding. Alternatively, fish-eye camera simulation rendering techniques could be used to calculate the Sky View Factor, as explored in previous studies (Hämmerle et al., 2011, Miao et al., 2020).

We anticipate high accuracy, based on the speed of the intersection tests of BVH without assumptions. We aim to analyse further discrepancies caused by input weather data, results aggregation, solar irradiance models, or surface characteristics (orientation, inclination). Additionally, we plan to validate our irradiance calculations against actual ground truth data to ensure reliability.

References

Agugiaro, G., González, F. G. G., Cavallo, R., 2020. The city of tomorrow from . . . The data of today. *ISPRS International Journal of Geo-Information*, 9(9).

Biljecki, F., Ledoux, H., Stoter, J., 2016. An improved LOD specification for 3D building models. *Computers, Environment and Urban Systems*, 59, 25–37.

Emmanuel, W., Jérôme, K., 2015. A verification of CitySim results using the BESTEST and monitored consumption values. *Building Simulation Applications*, 215–222. ISBN: 9788860460745.

European Commission, 2024. Pvgis typical meteorological year (tmy) generator.

Hämmerle, M., Gál, T., Unger, J., Matzarakis, A., 2011. Comparison of models calculating the sky view factor used for urban climate investigations. *Theoretical and applied climatology*, 105, 521–527.

Holmgren, W. F., Hansen, C. W., Mikofski, M. A., 2018. pvlib python: A python package for modeling solar energy systems. *Journal of Open Source Software*, 3(29), 884.

Kamphuis, N., Gueymard, C., Holtzapple, M., Duggleby, A., Annamalai, K., 2020. Perspectives on the origin, derivation, meaning, and significance of the isotropic sky model. *Solar Energy*, 201, 8–12.

KNMI, 2024. Automatic Weather Stations.

Lawrie, L. K., Crawley, D. B., n.d. Development of global typical meteorological years (TMYx). Publication Title: http://climate.onebuilding.org.

Ledoux, H., Arroyo Ohori, K., Kumar, K., Dukai, B., Labetski, A., Vitalis, S., 2019. CityJSON: a compact and easy-to-use encoding of the CityGML data model. *Open Geospatial Data, Software and Standards*, 4(1). https://opengeospatialdata.springeropen.com/articles/10.1186/s40965-019-0064-0. Publisher: Open Geospatial Data, Software and Standards _eprint: 1902.09155.

Loutzenhiser, P. G., Manz, H., Felsmann, C., Strachan, P. A., Frank, T., Maxwell, G. M., 2007a. Empirical validation of models to compute solar irradiance on inclined surfaces for building energy simulation. *Solar Energy*, 81(2), 254–267. https://www.sciencedirect.com/science/article/pii/S0038092X06000879.

Loutzenhiser, P. G., Manz, H., Felsmann, C., Strachan, P., Frank, T., Maxwell, G., 2007b. Empirical validation of models to compute solar irradiance on inclined surfaces for building energy simulation. *Solar Energy*, 81(2), 254–267.

Meister, D., Ogaki, S., Benthin, C., Doyle, M. J., Guthe, M., Bittner, J., n.d. A survey on bounding volume hierarchies for ray tracing. *Computer Graphics Forum*, 40, Wiley Online Library, 683–712. Issue: 2.

Miao, C., Yu, S., Hu, Y., Zhang, H., He, X., Chen, W., 2020. Review of methods used to estimate the sky view factor in urban street canyons. *Building and environment*, 168, 106497.

Nagel, C., Deininger, K., 2024. citygml4j/citygml-tools: Collection of tools for processing CityGML files.

NEN, 2024. NTA 8800:2024 nl. Technical report, NEN.

Perez, R., Ineichen, P., Seals, R., Michalsky, J., Stewart, R., 1990. Modeling daylight availability and irradiance components from direct and global irradiance. *Solar energy*, 44(5), 271–289.

Perez, R., Seals, R., Ineichen, P., Stewart, R., Menicucci, D., 1987. A new simplified version of the Perez diffuse irradiance model for tilted surfaces. *Solar energy*, 39(3), 221–231.

Perez, R., Stewart, R., Seals, R., Guertin, T., 1988. The development and verification of the perez diffuse radiation model. Technical report, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States); State Univ ....

Peters, R., Dukai, B., Vitalis, S., van Liempt, J., Stoter, J., 2022. Automated 3D Reconstruction of LoD2 and LoD1 Models for All 10 Million Buildings of the Netherlands. *Photogrammetric Engineering and Remote Sensing*, 88(3), 165–170. _eprint: 2201.01191.

Ray Tracing in One Weekend Book Series, 2020.

The World Bank, n.d. World development indicators | DataBank.

Tuononen, M., O'Connor, E. J., Sinclair, V. A., 2019. Evaluating solar radiation forecast uncertainty. *Atmospheric Chemistry and Physics*, 19(3), 1985–2000.

Wang, X., Zhang, X., Zhu, S., Ren, J., Causone, F., Ye, Y., Jin, X., Zhou, X., Shi, X., 2023. A novel and efficient method for calculating beam shadows on exterior surfaces of buildings in dense urban contexts. *Building and Environment*, 229, 109937. https://linkinghub.elsevier.com/retrieve/pii/S0360132322011672.

Xu, L., León-Sánchez, C., Agugiaro, G., Stoter, J., 2024. Shadowing Calculation on Urban Areas from Semantic 3D City Models. T. H. Kolbe, A. Donaubauer, C. Beil (eds), *Recent Advances in 3D Geoinformation Science*, Springer Nature Switzerland, Cham, 31–47.