









HeatMat: Simulation of City Material Impact on Urban Heat Island Effect

M.Reinbigler^{1,2} , R.Rouffet² , P.Naylor³ , M.Czerkawski^{3,4} , N.Dionelis³ , E.Brunet¹ , C.Fetita⁵  and R.Martin² 

¹SAMOVAR, Inria Saclay, Télécom SudParis, Institut Polytechnique de Paris, Palaiseau, France

²Adobe Research, Paris, France

³ESA/ESRIN, φ-lab, Frascati, Italy

⁴Asterisk Labs, London, UK

⁵SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, Palaiseau, France




Figure 1: We propose a pipeline for estimating the facade composition of buildings in a city, linking this information with their thermal properties to feed our 2.5D heat transfer simulator. First, we combine OpenStreetMap data with geolocated street-view 360° images to identify the best image for each building, which we input into a VLM to query the facade composition. We represent the city fully, including its materials and geometry, in randomly accessible 2D maps, which serve as inputs for our simulator that models coupled heat transfers.

Abstract

The Urban Heat Island (UHI) effect, defined as a significant increase in temperature in urban environments compared to surrounding areas, is difficult to study in real cities using sensor data (satellites or in-situ stations) due to their coarse spatial and temporal resolution. Among the factors contributing to this effect are the properties of urban materials, which differ from those in rural areas. To analyze their individual impact and to test new material configurations, a high-resolution simulation at the city scale is required. Estimating the current materials used in a city, including those on building facades, is also challenging. We propose HeatMat, an approach to analyze at high resolution the individual impact of urban materials on the UHI effect in a real city, relying only on open data. We estimate building materials using street-view images and a pre-trained vision-language model (VLM) to supplement existing OpenStreetMap data, which describes the 2D geometry and features of buildings. We further encode this information into a set of 2D maps that represent the city's vertical structure and material characteristics. These maps serve as inputs for our 2.5D simulator, which models coupled heat transfers and enables random-access surface temperature estimation at multiple resolutions, reaching a 20x speedup compared to an equivalent simulation in 3D.

CCS Concepts

• **Computing methodologies** → Computer vision; **Physical simulation**; • **Applied computing** → Computer-aided design;

1. Introduction

The Urban Heat Island (UHI) effect designates the temperature gap between an urban area and its surrounding rural area. As a conse-

quence of global warming, the UHI is amplified, both temporarily and in intensity. For example, by considering the 65 largest cities in the USA in 2024, 33.8 million Americans were exposed to UHIs with a gap in temperature of 8°F or more [Cen04]. Overheating

impacts both the environment by disturbing and degrading it, and human health by increasing mortality and the occurrence of respiratory diseases. Studying this phenomenon and its causes is therefore essential. Among the causes, urban materials' thermal properties are known to play a significant role in heat retention, which are adjustable parameters. Thus, we focus this study on providing city planners with insights into their material selection independently from the other causes.

To study the UHI effect, surface temperature can be monitored thanks to remote sensing measurements covering the entire Earth, but at a low spatial resolution (100m/pixel) and very low temporal resolution (one data point every few days). The air temperature can also be an indicator, but is captured by very scarce meteorological stations. If it is sufficient for capturing and predicting the global heat flux, it is not accurate enough to study how urban materials locally affect the surface temperature and reinforce the UHI effect.

High-resolution surface temperature prediction can be performed using heat transfer simulators that evaluate the impact of different variables of the system in the resulting heat maps. Software based on physical equations solvers, such as ENVI-met [Em24], has proven its efficiency in simulating heat transfers in the field of city digital twins, as well as Monte Carlo based approaches, such as Stardis [PBC*24], but they both present one main drawback: the need for a 3D digital-twin model with material descriptions, which is a tedious task to obtain at the scale of an entire city. On the other hand, OpenStreetMap (OSM) [Ope04] database contains the vector graphics description of the 2D maps of worldwide cities with a large coverage and easy to query via an API. Street-view images are omnipresent in dense cities, and can be used to estimate the facade materials per building, for a precise and automated way to describe the urban materials.

To remove the need of a 3D digital twin, we propose to encode the city in a set of top-view 2D maps including the description of its verticality and material information to form a 2.5D city model. It is coupled with a novel, Monte Carlo based, heat transfer simulator that runs as a pixel shader on a GPU, leveraging our 2.5D city representation. Our 2.5D simulator follows the line of work of [CdB*24] and [BBC*23] for coupled heat transfers, with direct and indirect solar irradiative, radiative, convective, and conductive transfers, with a 20x speedup and comparable results.

Our contributions can be summarized as follows:

- **HeatMat:** A novel approach allowing the study of the impact of materials on the UHI effect in a real city using only 2D data (OSM and street view images).
- A 2.5D heat transfer simulator running on GPU implementing coupled heat transfers (direct and indirect solar irradiative, radiative, convective, and conductive transfers) based on a Monte Carlo method that produces top-view heat maps of the surface temperatures, given the 2.5D city representation and its location on Earth, the date, time and ambient temperature. We achieve a 20x speedup with comparable results to SOTA.
- A 2.5D representation of a city, accounting for its geometry (height, signed distance field), and materials (top view, and facade materials procedural representation), via a set of textures, randomly-accessible from a GPU shader.

- A VLM-based pipeline for estimating the facade materials statistics of buildings from street views, enabling a procedural representation of the facade as a complement of OSM building data.
- Validations of the heat transfer implementations involving conductive transfer, individually or coupled, with two different implementations, namely the Finite Difference Method and the Monte Carlo based Stardis prototype, as well as with satellite measurements.

2. Background and Related works

2.1. Heat transfer simulation

The heat transfers can be simulated using two families of methods.

Deterministic methods. Numerical methods based on physical equation solving have been extensively used to simulate heat transfer in urban areas. These are based on the discretization of the 3D environment. Tools are available at different scales: Urban canopy models [SLZ*16, Mas00, KKKK01, MCR02, Em24] [TC08, SKD*19, LRM23], and Building energy models [TNZ14, BMMK*14, CHKG08, ESR03, KMW19, LP19, MR20] [PE14, RM01, RHL*09]. However, these require meshing, and a predefined level of detail to build a model: e.g., a whole coarse city model, *i.e.*, urban canopy model, or a detailed building model. The combination of model scales leads to approximations.

Monte Carlo methods

Using Monte Carlo methods for heat transfer simulation has been studied to model radiative transfers, leveraging photon tracing techniques [FHR*23, FLZ*25]. It can be combined with a differentiable modelization to optimize certain parameters, such as the new building geometry choices presented in [FLZ*25]. On the other hand, individual conductive transfer equations solving has been implemented and recently improved by the introduction of the Walk on Stars [SMGC23, MSCG24] that ensures only one boundary intersection per walk step, thus improving results quality. It is compatible with Dirichlet and Neumann boundary conditions [SMGC23] and further extended to Robin boundary conditions [MSCG24]. [TAB*23] has demonstrated theoretically the possibility of combining individual transfer models in one random path to couple radiative-convective-conductive transfers, which is also compatible with computer graphics techniques for efficient transfer solving. In addition, [IBC*22] proved the independence of these methods from the geometry size. These properties enable complex physical phenomena modeling, such as combining an atmospheric model with an urban model, while preserving complex multiscale scenes [VHd*22]. [PBC*24] implemented the theoretical work of [TAB*23]. [BBC*23] added a post-mortem replay of the simulation, and [CdB*24] added the solar irradiation contribution.

Our work proposes to leverage the Monte Carlo based model from [BBC*23] and [TAB*23] to build a coupled heat transfers simulator, but instead of relying on a 3D city representations, it inputs 2D random-accessible maps describing the city, implying adaptations of conductive transfers from 3D to 2D.

2.2. Material retrieval

Fine-grained material estimation that recovers the Spatially Varying Bidirectional Reflectance Distribution Function (SVBRDF) of a quasi-planar surface has been explored, relying on neural architecture to solve this ill-posed problem. In the literature, researchers have imposed lighting constraints with flash-light [DAD*18, DDB20, LSBE24, CWZ*24] and used in the wild image capture [MRR*22, VMR*24, VSPS24], or both [ZK21, SP23, ZGL*25].

For large-scale material estimation, satellite, aerial, and street view images are often used. [BKI18] generates a 3D city model that is automatically textured, also performing automatic building reconstruction. However, it requires the acquisition of airborne images at a 10 cm resolution or more and a flight altitude of 800 to 2500 m. Recent papers detecting materials use several sources, including LIDAR [HMB22], satellite for roof material detection [KBKL21], and manual surveys. Automated methods at large scale exist and are based on ground views: (i) Google photorealistic tiles, providing only RGB textures [Pla24]; (ii) LuxCarta geospatial training and simulation products [Car24]; (iii) Geometry from Street-View [ZWA*16]; and (iv) [RBD23] using street-view images and Computer Vision. However, [RBD23] does not statistically represent facade material distribution, and thus requires training to extend to new cities. In this paper, we focus on the physical properties of materials and on a statistical facade representation that allows procedural element arrangement and connection to appearance and/or a physical material database.

2.3. Procedural facade

Due to the inherent regularity and repetition of building facades, procedural modeling has been an effective paradigm for representing architectural facades. Split grammars and L-systems have been extensively explored to represent city elements [PM01, WWSR03, MWH*06, Pog19], further extended with inverse procedural modeling ability [MZWVG07, SBM*10, WYD*14, CKP*24]. More recently, the emphasis has shifted towards hybrid neuro-symbolic models that combine the structural interpretability of grammars with the expressivity of data-driven learning [PSPSM24, PSM25].

Our work hierarchically combines simple shape elements to model the facade procedurally, and could be compatible with more advanced approaches.

2.4. City twin for UHI detection

Remote sensing data are extensively used to model the UHI effect. [Sma06] uses Landsat data to model the relationship between surface reflectance (how surfaces interact with light) and thermal conditions (surface temperature) across urban environments. Recently, remote sensing data are more commonly used, such as using LiDAR data to approximate 3D urban morphology [AAF20, BSK*23], Sentinel-2 data for Normalized Difference Vegetation Index [AAF20, BSK*23], or Landsat for approximating land surface temperature [AAF20, BSK*23, LW23, KSS*21]. Microclimate Vision [FKY*24], integrates different sources of data,

including satellite observations, to predict high-resolution microclimate data for cities [FKY*24].

UHI detection using city twin is exploited by [Em24], which models many more physical phenomena than heat transfers. It is based on a voxel representation. It can also import OSM geometries to be extruded and voxelized but with aliasing effects. The ENVI-met software has been adopted by studies to analyze urban design and geometry influence on urban heat [RGL18, ARF21, TF23], but its radiosity model involves slow execution and in a very restricted area in the free version. Computational fluid dynamics based models are also proposed to analyze albedo influence on urban heat [LCAGDM*22]. Some models simulate heat transfers coupled with atmosphere effects, e.g., (i) [RMBG23] includes a mass transfer model in addition to coupled heat transfers to include a climate model, and (ii) [Mas00] models energy transfers in a simplified 3D city geometry, including an atmospheric model.

In this work, we propose to isolate the effect of materials on the UHI phenomenon by performing a high-resolution simulation of the heat transfers from the city materials and geometry description extracted automatically from open-source data.

3. City acquisition

As highlighted by previous works, Monte Carlo methods are commonly adopted to solve heat transfers in an urban geometry, leveraging rendering techniques thanks to their ability to scale to the geometry size. However, current implementations require a 3D representation of a city, which is a tedious task at the scale of a real city. We propose to rely on 2D information only, OpenStreetMap data for the city map, coupled with street-view images (Sec. 3.1), from which we estimate the facade materials using a VLM-based pipeline (Sec. 3.2). We encode the city into 18 2D textures (Sec. 3.3), including the facade description (Sec. 3.4), to allow for running the simulation as a simple pixel shader and gain performance. The whole pipeline is summarized in Fig. 2.

3.1. Coupling street-views with building footprints

Several open data sources are available, including Mapillary, which provides both user-captured street-view images as well as a set of curated datasets. Among these, we select the Metropolis dataset [Map24], captured in Detroit, Michigan, with a 360° camera on a geolocated car. As we focus on buildings, we use the geolocation metadata to deduce the field of view (FOV) cone of each subsection corresponding to the left and right views from the car position, likely to contain buildings. We query the OSM database with this cone, with a far distance large enough to cross a standard street, to retrieve all building footprints intersecting it. We discard from the query results the buildings fully occluded by the closest buildings to the camera position. We then augment the Metropolis metadata with, for each captured view, the visible buildings, their rotation, as well as their relative position, and coverage in the image. We further use this information to identify the best view of each building, by scoring the building's proximity to the camera, frontal positioning and coverage. This allows for a reduction by a factor of 10 in the number of images to process, for a sober use of the following material estimation pipeline.




Figure 2: Pipeline for material retrieval. From left to right: Starting with street-view 360° images and geolocated camera positions, we couple both information and elect, for each building, the 360° image that best represents it. We rectify this image and extract the building facade of interest, segment it, and query a VLM to estimate the material and proportions of each element in the facade.

3.2. VLM-based facade composition estimation

We propose a pipeline for estimating the materials in a facade using image processing and state-of-the-art pre-trained models. It starts with the recovery of the left and right rectified pictures from the 360° equirectangular views from Mapillary, using the Perspective Fields method [JZH*23], computed for the four azimuth angles (0, 90, 180, 270) of the panorama and averaged to uniformly rectify it, accounting for an 85° vertical field of view. Only the left and right images are kept from the rectification, and each image is then sliced to extract only the building view corresponding to our metadata. Each sliced image is then segmented using GroundedSAM [RLZ*24] that we query with the prompt "facade" to keep only the buildings. We crop the segmented image to the bounding box of the mask, to focus the VLM attention.

The segmented image is provided as input to a pretrained VLM model, Llava 1.6 [LLLL24] in our experiments, which can be easily replaced. We design our prompt so that the VLM returns, in JSON format, the material and proportion of each important facade element, separated for the ground level (main material, doors, windows, frames) and the upper levels (main material, windows, frames, shutters). We augment our building metadata with this data.

For validating the facade materials and proportions estimated, we manually built a ground truth dataset made of 50 street views that represent one building facade, for which we label the materials and proportions manually, by carefully analyzing the image. An ablation of the use of GroundedSAM, and measures of the material estimation with IoU and accuracy metrics, using three different VLMs, are presented in the Suppl. Tab. 1.

3.3. Top view data encoding into texture

The first set of maps describes the features of the building visible from a top view, to be stored in a 2D map of the city. Among those features, there are the identifier (id) of each building, its height, and the material composition of the roofs and the ground. We create the maps by applying a set of operations on the footprint polygons using the building metadata. We rasterize the area of interest with a cell size of the resolution of the simulation chosen in meters. When a cell is at the boundary of multiple buildings, we store the infor-




Figure 3: Illustration of the rasterization of footprint buildings. In each pixel rasterized (left) we keep only one segment that represents a facade, to allow for the discrete SDF representation (right). At the corners (middle) we simplify the geometry to connect the two segments coming from the adjacent facades.

mation of the building with the largest overlap. We pack our scalar data into RGBA 16-bit textures to minimize the texture sampling later in the simulator.

3.4. Facade representation and encoding

To represent the verticality of the city, not visible from a top view, we define a methodology to first encode the facade geometry while removing the aliasing effect due to the rasterization of the footprint polygons for a more accurate intersection detection.

Thus, we represent the facade orientation using the couple (distance, normal). The distance corresponds to the space between the pixel centroid and the closest facade, which forms a discrete SDF to be used for Sphere Tracing. The normal corresponds to the facade normal, expressed as the azimuth angle of the facade. To ensure one normal and one distance per pixel for the external facade, we modify the geometry by approximating corners by one segment connecting the two points of intersection to the grid cell. Figure 3 illustrates this approximation to preserve intersection continuities.

To ensure coherence in the sampled facade material, we build a procedural function that associates with each position (u,h) in the facade plane the building component to which it belongs (door, window, etc.), further matched with the corresponding material in




Figure 4: The u map and perimeter maps allow for retrieving the relative position of a 3D point in the plane of a facade, and for guaranteeing a consistent facade representation in the simulator.

that facade (metal, glass, etc.). To allow a random accessible way to get the relative position u along a facade plane, we parameterize it as an interpolation between 0 and 1, during the polygons rasterization. At simulation time, when querying the facade material at a point $P(x,y,z)$, h is given by $P.z/\text{height}$, and u is obtained by adding to the queried u value of the u map the facade length, divided by the whole perimeter length. Details are available in Fig. 4. Using a procedural sampling function, based on the 2.5D encoded material maps, we identify the facade component to which our point belongs and the material id.

To retrieve the facade material properties, we use the first 8 bits for the material id and the following 8 bits for the percentage. Table 1 in Supplemental Material presents the material thermal properties database, aggregated from several sources [fBERCoNM, Sco10, CYP91, Joh21, Too, KOIK21, Ome].

2.5D city maps and a visual of the procedural facade sampling are illustrated in Fig. 5. This facade sampling, coupled with the procedural representation of facade and the 2D random-accessible geometry encoding (Umap, perimeter map, height map, SDF, and material information) allows to simulate different facade elements along the vertical axis, which is novel to the best of our knowledge.




Figure 5: 2.5 representation of a city using a set of maps, with a top view (roofs, ground), building facade materials (left), city geometry height map, signed distance field, and facade normals (middle). It enables procedural modeling of a facade (right) when intersecting points on a facade in the heat transfer simulator.

4. 2.5D heat transfer simulation

The set of 2D maps generated making this information random-accessible, offers the opportunity to build a simulation running on a GPU as a pixel shader. Thus, it takes advantage of computer graphics techniques, such as sphere tracing and walk-on-sphere, enabling an efficient intersection computation. The simulation, based on Monte Carlo methods, computes the heat transfers in an urban area, including direct and indirect solar irradiation, radiative, convective, and conductive heat transfers. Each one is solved independently as described in Sec.4.1. The main mathematical and physical equations presented hereafter are extracted from Caliot *et al.*'s paper[CdB⁺24], including how Monte Carlo methods and computer graphics techniques can solve the problem in a 3D geometry. Our goal is to estimate the surface temperature T_s using Monte Carlo estimator $\tilde{T}_s(\vec{y}_0, t_0)$ where

$$T_s(\vec{y}_0, t_0) \approx \tilde{T}_s(\vec{y}_0, t_0) = \frac{1}{N} \sum_{k=1}^N W_k \quad (1)$$

W_k is the Monte Carlo method weight of one realization defined in Eq. 2

This weight can be decomposed as the sum of the possible path end temperatures, to which the weight $W_{o,k}$ is added (Eq. 2). The weight contains T_I when the path reaches the initial time τ_I via the conductive transfer. Otherwise, it is composed of T_D when it reaches a Dirichlet boundary, a known solid temperature, T_F at a solid-fluid interface corresponding to the fluid temperature involved in the convective transfer or T_{sky} when a radiative ray ends in the sky. $\delta_{condition}$ designates the Dirac function : it is worth 1 if the condition is met, 0 otherwise.

$$W_k(\vec{y}_0, t_0) = \delta_{t_n \leq \tau_I} T_I + \delta_{t_n > \tau_I} \left\{ \begin{aligned} &\delta_{\vec{y}_n \in \partial D_S} T_D(\vec{y}_n, t_n) \\ &+ \delta_{\vec{y}_n \in \partial D_{S,F}} T_F(\vec{y}_n, t_n) \\ &+ \delta_{\vec{y}_n \in \partial D_{sky}} T_{sky}(\vec{y}_n, \vec{\omega}_n, t_n) \end{aligned} \right\} + W_{o,k} \quad (2)$$

$W_{o,k}$ refers to the solar contribution to the temperature estimation. It is composed of two parts: the direct solar irradiation $W_{o,d}$ and a second one corresponding to the indirect contribution along the indirect path (Eq. 3). n_o designates the number of Robin boundary conditions met, $n_{r,j}$ to the number of reflections on the indirect path and h_T the sum of each heat transfer mode coefficient.

$$W_{o,k} = \frac{1}{h_T} \sum_{j=1}^{n_o} \left\{ W_{o,d}(\vec{y}_j) + \left(\sum_{m=1}^{n_{r,j}} \delta_{\vec{y}_m \in \partial D_{S,F}}^L W_{o,r,d}^L(\vec{y}_m) \right) + W_{o,sky}(\vec{y}_{n_{r,j}+1}) + W_{o,r,d}^F(\vec{y}_{n_{r,j}+1}) \right\}. \quad (3)$$

The direct solar radiation is defined in Eq. 4 as a function

of the solar direction and normal vectors $\vec{\omega}_d$ and \vec{n}_0 , the direct solar irradiation D_o , and the material emissivity ϵ , *i.e.* material's ability to radiate absorbed heat. Based on Kirchhoff's law, *emissivity = absorptivity* for all wavelengths, *i.e.* for opaque material *emissivity = 1 - albedo*.

The indirect solar irradiation contribution can be decomposed into three components: the diffusion of solar irradiation at each bounce on a Lambertian surface (Eq. 5), and two possible path end conditions. One end condition corresponds to an indirect path ray reaching the sky, thus adding a weight defined by Eq. 6, modeling an atmospheric sky model where it is proportional to the direct radiative solar intensity $I_{o,d}$. The second end condition corresponds to the reflection of the ray on a specular surface with a reflected ray contained in the sun half-angle. It depends on the material emissivity and the direct radiative solar intensity $I_{o,d}$ (Eq. 7).

$$W_{o,d}(\vec{y}_0) = \delta_{\vec{y}_d \in \partial D_d} \epsilon(\vec{y}_0) |\vec{\omega}_d \cdot \vec{n}_0| D_o(\vec{y}_0) \quad (4)$$

$$W_{o,r,d}^L(\vec{y}_i) = \epsilon(\vec{y}_0) \delta_{\vec{y}_d \in \partial D_d} |\vec{\omega}_d \cdot \vec{n}_i| D_o(\vec{y}_i) \quad (5)$$

$$W_{o,sky}(\vec{y}_i) = \epsilon(\vec{y}_0) \pi \delta_{\vec{y}_{i+1} \in \partial D_{sky}} f_{sky}(-\vec{\omega}_i) I_{o,d}(\vec{y}_i) \quad (6)$$

$$W_{o,r,d}^F(\vec{y}_i) = \epsilon(\vec{y}_0) \pi \delta_{\vec{y}_i \in \partial D_{S,F}}^F \delta_{\vec{y}_{i+1} \in \partial D_d} I_{o,d}(\vec{y}_i) \quad (7)$$

Sections 4.1 and 4.2 describe how we implemented the solving of these equations in practice in 2.5D, inspired by previous works [CdB*24, BBC*23] and including adaptations to the 2D conductive model and data representation specific to this paper.

4.1. Individual heat transfers design

4.1.1. Conductive heat transfer

In our 2.5D representation, only top-view information are accessible which implies we do not model the building interiors. Thus, our conductive heat transfer implementation leverages the Walk-On- δ -sphere, similarly to [BBC*23], but in 2D dimension on surfaces, meaning across the facades, roofs, and ground surfaces. Like us, Caliot *et al.* and Bati *et al.* compute the conductive path based on the Walk-on- δ -Sphere techniques, *i.e.* the Walk-on-Sphere method, but instead of estimating the radius of the sphere when tangent to the surface, we move with a sphere with a fixed radius δ . Both followed Tregan *et al.* [TAB*23] statement saying that the estimation matches accurately the exact solution when $\delta = \frac{L}{20}$, where L is the characteristic length of the problem. In the 3D context, the characteristic length of the conductive transfer is the wall thickness. To adapt it to our 2D case, we propose to define the characteristic length as $\min(\text{width}_{pattern}, \text{height}_{pattern})$ size where the pattern corresponds to the unit facade arrangement being used in the facade sampling. In addition, as we remove a dimension along the facade, the coupled conductive path algorithm requires adjusting the exponential sampling mean of the rewind in time as follows:

$$\frac{\rho c \delta^2}{6k} \rightarrow \frac{\rho c \delta^2}{4k}$$

with ρ the material density, c the heat capacity, and k the thermal conductivity.

4.1.2. Radiative heat transfer and solar irradiation

The radiative heat transfer is divided into two categories: the terrestrial radiative transfer generated by objects' temperature and the solar radiative transfer resulting from the sun irradiation. In both cases, we need to explore the space, which is implemented using the sphere tracing technique based on a discrete Signed Distance Field (SDF) stored in the SDF map used as input. The intersection computations are performed in 3D with a building geometry described by its heightfield and the facade description, as discussed in Sec. 3.4. For the solar radiative transfer, to better represent reality, we simulate the real sun path depending on time and date.

4.1.3. Convective transfer

The volume of air temperature in an urban area is substantial, the inertia required for the air temperature to be increased globally under the influence of the solids present in the scene is large. Thus, we can consider that the convective transfer contribution to the Monte Carlo estimator of the surface temperature equals the air temperature at the local time, similarly to Caliot *et al.* work [CdB*24].

4.2. Heat transfers coupling




Figure 6: Coupling of radiative, convective, conductive, as well as solar direct and indirect irradiation.

The coupling, inspired by the Bati *et al.* pseudo-code [BBC*23], is completed by the solar contribution described by Caliot *et al.* [CdB*24], with adaptations due to our 2.5D setting. We start our computation at the location corresponding to each pixel centroid, at an altitude matching the height map value. Then, at each solid-fluid interface reached, we compute the solar irradiation contribution to the final temperature. When at a solid-fluid interface, we randomly select a transfer mode using the material properties dependent probability to select a radiative, a conductive, or a convective heat transfer as described in Caliot *et al.* [CdB*24]. When a convective transfer is sampled, the path ends and returns the air temperature at the time reached by the simulation, plus the solar contribution accumulated along the path. If the radiative heat is sampled, then a ray is sent in a randomly sampled direction using a cosine distribution. If it hits a geometry, a new sampling of the heat transfer type is computed. Otherwise, the path is ended and returns the corresponding to the sky temperature, added to the accumulated solar Monte Carlo weight. Finally, if the conductive

transfer is selected, we need to adapt it to our surface transfer instead of volumetric transfer. Consequently, as the conductive path evolves on the facade surface, it remains at the solid-fluid interface at each step. Thus, contrary to the literature method, we do not sample a new heat transfer type once we reach an interface, but after a number of steps to be parametrized. At each step of the conductive path, a rewind in time is performed accounting for the material property. If the initial time is reached, the conductive path ends the simulation path by returning the initial temperature at the current location together with the solar contribution. An illustration of all transfers at stake is available in Fig. 6.

5. Evaluation and performances

We evaluate individually our 2D implementation of the conductive transfer compared to the Finite Difference Method (Sec. 5.1), and we validate its coupling with radiative and convective transfer (Sec. 5.2) and against satellite measurements (Sec. 5.3).

5.1. Individual conductive transfer evaluation

We compare our 2D implementation of the conductive transfer against the surface Finite Difference Method. The simulation consists of a Carbon 1% Steel heated continuously at its two vertical extremities to reach 1°C and maintained at 0°C at the horizontal extremities. We compare the result at the stationary state. The initial temperature is set to 0°C in the whole rectangle except the top and bottom row, heated at 1°C. Due to the method's stochasticity, the results only approximate the solution, thus introducing small deviations. When computing the scaled difference as defined in [CdB*24], we observe an error of 1% compared to the reference, which is in the same order of magnitude as reported in [CdB*24].



(a) Finite Difference Method (b) Our Walk-on- δ -Sphere based method

Figure 7: Conductive transfer comparison at the stationary state presents a mean scaled difference of 1%.

5.2. Coupled heat transfers evaluation

We compare our coupling implementation to the available 3D Monte Carlo based simulator Stardis [PBC*24], which models the same transfers as our method except for the solar contribution. We set all initial conditions according to Suppl. Tab. 4. To be able to run the simulation on the same geometry, we created a 2.5D representation of 210x320 pixels of the Stardis starter pack with a 0.5m per pixel resolution. As for Stardis, the 3D model weighs 11.4MB while our corresponding 2.5D representation weighs 1MB. Figure 8 presents the top-view heatmap outputs in the same color range. Due to the rendered perspective and to balcony-related effect, as we do not model those elements, the resulting map cannot be superimposed on our heatmap. Thus, we estimate the average

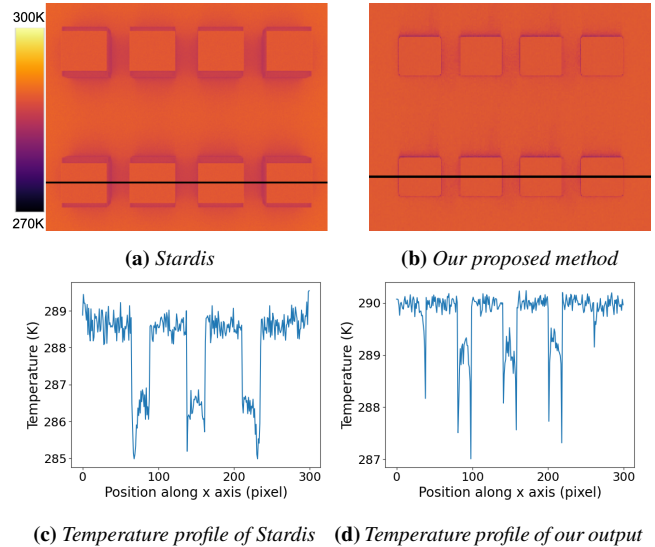


Figure 8: Coupled heat transfer comparison, without solar irradiations, from a top view heatmap and via temperature profile along the displayed line. Our 2.5D method provides comparable results to Stardis, which runs on a 3D geometry.

temperature per area of interest: roofs, ground far from buildings, and in-between buildings. The values deviate respectively from 0.1%, 0.4%, and 0.2% from Stardis values, which correspond to less than a Kelvin, which should not affect the ability to detect heat islands perceptible by human beings. The temperature profiles along the horizontal line shows the three down peaks, corresponding to the drop in temperatures between the building which are in the shadow. The two additional peaks in our result correspond to the left and right façades, hidden by the perspective on the Stardis result. The left drop is larger than the right one as it is less exposed to the sun. We can also note that our simulator tends to heat a bit more than Stardis but follows the same overall evolution.

In Tab. 1, we compare the performance of our 2.5D simulator against Stardis on this geometry, with varying samples per pixel (fewer samples, more noise) on a laptop equipped with a Processor 11th Gen Intel(R) Core(TM) i7-11800H (2.30 GHz) and an NVidia GeForce RTX 3080 Laptop GPU. The 20x speedup is obtained by writing the simulator as a GPU shader, using random accessible maps for representing the city. Our simulator scales linearly with the maps resolution and the number of samples per pixel.

Simulator / SPP	125	250	500	1000
Stardis	11min	21min41	44min11	90min10
Ours w/o irradi.	42sec	1min21	2min41	5min15

Table 1: Simulator performances. Ours is 20x faster.

5.3. Comparison to Satellite Land Surface Temperature

Collecting ground truth data of surface temperatures at high temporal and spatial resolutions is a complex task, which is why being able to simulate it is so crucial for providing urban planners with more insights. For instance, the Landsat satellite captures surface




Figure 9: Detroit district simulation, July 27, 2024, at 12:15 pm with the same initial conditions as in Sec. 6. experiment. The gradients of our downscaled simulation (b) are very similar to the Landsat capture (d) at the same day and time.

temperatures at 100m/pixel every 16 days. We found one record for Detroit around 12:16 pm on July 27, 2024, retrieved the ambient temperature at that time, and ran our simulator with these settings. We applied several concrete versions with different emissivities directly derived from the albedo estimation performed on satellite images. We generate the set of 18 198x167 pixels maps with a spatial resolution of 3m/pixel for a broad overview. Fig. 9 presents our simulation, and the Landsat record, as well as Fig. 10. Darkest roofs correspond to lightest concrete with highest albedo, whereas hottest one corresponds to darkest concrete version. We performed bilinear downsampling and upsampling on our heatmap to enable a direct comparison at the same resolution. Considering that Landsat measurements are accurate at about 3K close and is lower in urban areas than in rural areas, our results in absolute value are in general slightly hotter but the overall gradient is very similar, confirming our intuition.

Similar experiments have been performed on several other dates with low cloud coverage. Fig. 10 illustrates that our simulator output follows similar heat distribution as the Landsat heatmap, The shift in absolute values can be explained by the Landsat images themselves that do not include any radiometric calibration, atmospheric correction nor emissivity correction, which can introduce some gaps. In addition, weather conditions, vegetation or anthropogenic activities are also not studied as we aim at isolating material impact independently of the other parameters.

6. Application to urban planning

To demonstrate how urban planners can benefit from this simulator, we present two studies that might be of interest for them. The first study examines the temporal evolution of surface temperature over time to identify areas of heat retention (Sec. 6.1). The second study explores the impact of material changes on heat transfers (Sec. 6.2). In order to highlight the impact of facade materials in those experiments, in this section we applied the same concrete to all building roofs.

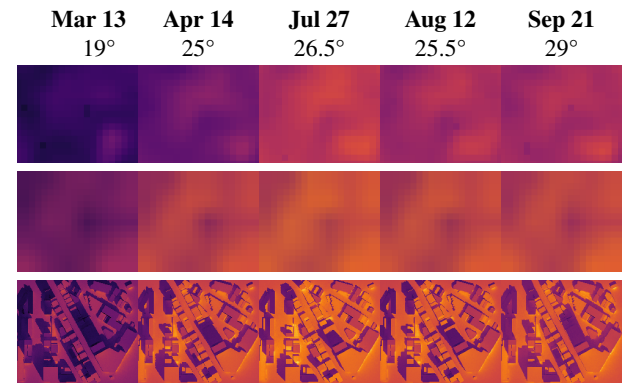


Figure 10: Comparison of Landsat Surface Temperature at different dates in 2024 in Detroit (top) with our simulation (bottom), run at the corresponding date, time and air temperature. We downscale our result to the same resolution as Landsat (100m/pixel) to compare (middle). The satellite records Detroit around 12:16 pm, local time. Results are all presented with the inferno range 290K-338K, as in Fig. 9. The overall gradients are very similar to Landsat ones, with a consistent shift in absolute value.

6.1. Temporal evolution study

We simulate the surface temperature evolution on June 17th, 2024, in a Detroit neighborhood described by 18 181x245 pixels maps at 0.5m/pixel resolution. The air temperature goes from 22°C (295K) up to 35°C (308K). The experimental conditions are summarized in Fig. 12, and Fig. 11 shows the heatmap evolution over the day.

Each output of Fig. 11 corresponds to the steady state of one simulation ahead of time, with as initial surface temperature the result of the simulation at the previous time. We notice a heat peak at 2pm, with an overall surface temperature of 330K, hotter than the air temperature, which is explained by the absorption of solar irradiation. The shadows are in accordance with the hourly sun position, and we can observe hotter areas on the ground next to the facades that receive direct solar irradiation. At 10pm, the overall temperature is the same as the air temperature, around 292K, which is also expected as the sun does not contribute anymore to surface heating. It can be noted that the surfaces at 9pm are hotter than at 6am even if the air temperature and the solar irradiation are similar, due to heat absorption by the building material.




Figure 12: Experimental conditions of the time-varying air temperature and the solar irradiation.




Figure 11: Temporal evolution of surface temperature over one day of one building block of Detroit.

6.2. Material change study

To highlight the capability of our simulator, we investigate the influence of materials on the UHI effect in a neighbourhood of the city of Detroit with a spatial resolution of 0.5m per pixel for fine-grained examination, resulting in 18 181x245 pixels maps. The simulation runs with a starting point on June 6, 2024 at noon with an assumed air temperature of 308K. The material disposition considered is the one extracted with our pipeline, before applying a change in the facade material to observe material properties' impact on the resulting surface temperatures. In the simulation shown in Fig. 13, we init the simulation a) with all building facades made of concrete as the main material, keeping the same percentage estimated, and b) with limestone instead. We don't modify the roof and ground materials, that we set respectively to concrete and asphalt. We can observe in c) that limestone induced an increase in the surface temperature, up to 2.4K in some areas. Limestone has a higher emissivity than concrete (See Suppl. Tab. 3), which implies a higher thermal absorptivity and a lower reflectivity in the visible spectrum. Thus, the limestone facades tend to irradiate more via the indirect sun path, leading to higher local temperatures near the exposed facades. Moreover, knowing that the map is north-facing, and the simulation is running for 2 pm in Detroit, the shadow positions are consistent with the sun's position and their sizes with the district height map.

Our method scales well to bigger area as highlighted in Fig. 9 which represents the Detroit district heatmap and its corresponding heightmap around our current study area with a spatial resolution of 3m/pixel, thus covering 592x499m, but with diverse roof albedos.

Thus, a user can, for any city where street-view images and OpenStreetMap information are available, model a neighborhood




Figure 13: Simulation of a facade material change in Detroit. a) The main material of all facades is set to Concrete. b) The main material of all facades is set to Limestone. c) Difference in temperature between a) and b). With the higher emissivity of Limestone, more heat is reflected to the roofs next to a high building, with a difference of 2.4K in some areas.

and test different material compositions to select the one with the lowest thermal impact. This could benefit urban planners to study material influence independently of other factors, such as the anthropic heat, the vegetation or the weather.

7. Discussion and Future Work

In this paper, we introduced a pipeline that estimates the composition of a city's building facades, coupled with 2D city features, which we encode into a set of textures to feed into our 2.5D coupled heat transfer simulator. This approach enables the simulation of the surface temperature of an urban area, and to test the impact of new material configurations on the UHI effect. The evaluation showed consistent results compared to satellite measurements, as well as to the 3D thermal simulator Stardis. Using 2.5D data instead of a 3D model enables 20× faster simulations due to the random accessibility of the maps and the reduced complexity introduced by our representative procedural facade generation. Moreover, 2D representations are an order of magnitude simpler for neural models to process than 3D ones. This paves the way for future dataset generation and learning-based approaches for UHI effect prediction.

Though our approach is conceptually generalizable to any city, our solution relies heavily on the availability of open data. Consequently, the set of materials considered is limited, as literature describing both optical and thermal properties of the same material is rare. Similarly, the coverage of the Mapillary database is spatially variable as it depends on community contributions. It can be mitigated by coupling this data with other sources, like cities' own open data. Our solution depends on the VLM performance in the detection of building materials. In this study, we opted for LLaVA 1.6, but it can be substituted by any other VLM or model specifically trained for this task. Our implementation of coupled conductive transfer currently allows for heat transfer within a bounded surface, namely the facade of a building, the roof, or the ground. The transition between surfaces and the change of path direction at corners are not yet included and could be beneficial to the simulator.

In the future, we could extend our facade material estimation pipeline, by considering aerial or satellite images to automatically retrieve both roofs and ground materials. Finally, to handle potential scalability issues induced by hardware specifications, our approach could benefit from a multi-resolution approach: start with a low resolution satellite land surface temperature image of a hot sunny day and determine the hottest areas of the city. Then, in those areas, run our simulation at low-resolution (3m/pixel), to identify problematic building blocks before applying our current high-resolution pipeline (0.5m/pixel) and explore different material alternatives to mitigate the local temperature increase.

8. Acknowledgments

We thank Yannick Hold-Geoffroy, Théo Thonat, Elie Michel and Tamy Boubekeur for their insightful feedbacks and their technical support. The authors also acknowledge the financial support of Hi! PARIS Center (<https://www.hi-paris.fr>).

References

- [AAF20] ASADI A., AREFI H., FATHIPOOR H.: Simulation of green roofs and their potential mitigating effects on the urban heat island using an artificial neural network: A case study in austin, texas. *Advances in Space Research* 66, 8 (2020), 1846–1862. URL: <https://www.sciencedirect.com/science/article/pii/S0273117720304737>, doi:<https://doi.org/10.1016/j.asr.2020.06.039>
- [ARF21] ACHARYA T., RIEHL B., FUCHS A.: Effects of albedo and thermal inertia on pavement surface temperatures with convective boundary conditions - a CFD study. *Processes* 9, 11 (2021). URL: <https://www.mdpi.com/2227-9717/9/11/2078>, doi:10.3390/pr9112078. 3
- [BBC*23] BATI M., BLANCO S., COUSTET C., EYMET V., FOREST V., FOURNIER R., GAUTRAIS J., MELLADO N., PAULIN M., PIAUD B.: Coupling conduction, convection and radiative transfer in a single path-space: Application to infrared rendering. *ACM Trans. Graph.* 42, 4 (July 2023). URL: <https://doi.org/10.1145/3592121>, doi:10.1145/3592121. 2, 6
- [BKII18] BUYUKDEMIRCIOGLU M., KOCAMAN S., ISIKDAG U.: Semi-automatic 3d city model generation from large-format aerial images. *ISPRS International Journal of Geo-Information* 7, 9 (2018). URL: <https://www.mdpi.com/2220-9964/7/9/339>, doi:10.3390/ijgi7090339. 3
- [BMMK*14] BEAUSOLEIL-MORRISON I., MACDONALD F., KUMMERT M., MCDOWELL T., JOST R.: Co-simulation between esp-r and trnsys. *Journal of Building Performance Simulation* 7, 2 (2014), 133–151. ARRAY(0x56102554e968). URL: <https://publications.polymtl.ca/12751/>, doi:10.1080/19401493.2013.794864. 2
- [BSK*23] BAGYARAJ M., SENAPATHI V., KARTHIKEYAN S., CHUNG S. Y., KHATIBI R., NADIRI A. A., ASGARI LAJAYER B.: A study of urban heat island effects using remote sensing and gis techniques in kancheepuram, tamil nadu, india. *Urban Climate* 51 (2023), 101597. URL: <https://www.sciencedirect.com/science/article/pii/S2212095523001918>, doi:<https://doi.org/10.1016/j.uclim.2023.101597>. 3
- [Car24] CARTA L.: Mapping your world in 3d, 2024. <https://www.luxcarta.com/projects/3d-city-model-sao-paulo-brazil/> [Accessed: (05/09/2024)]. 3
- [CdB*24] CALIOT C., D'ALENÇON L., BLANCO S., FOREST V., FOURNIER R., HOURDIN F., RETAILLEAU F., SCHOETTER R., VILLEFRANQUE N.: Coupled heat transfers resolution by monte carlo in urban geometry including direct and diffuse solar irradiations. *International Journal of Heat and Mass Transfer* 222 (2024), 125139. URL: <https://www.sciencedirect.com/science/article/pii/S001793102301284X>, doi:<https://doi.org/10.1016/j.ijheatmasstransfer.2023.125139>. 2, 5, 6, 7
- [Cen04] CENTRAL C.: Urban heat hot spots in 65 cities, 2004. <https://www.climatecentral.org/climate-matters/urban-heat-islands-2024> [Accessed: (12/11/2024)]. 1
- [CHKG08] CRAWLEY D. B., HAND J. W., KUMMERT M., GRIF-FITH B. T.: Contrasting the capabilities of building energy performance simulation programs. *Building and Environment* 43, 4 (2008), 661–673. Part Special: Building Performance Simulation. URL: <https://www.sciencedirect.com/science/article/pii/S0360132306003234>, doi:<https://doi.org/10.1016/j.buildenv.2006.10.027>. 2
- [CKP*24] COGO E., KRUPALIJA E., PRAZINA I., BEĆIROVIC S., OKANOVIĆ V., RIZVIĆ S., MULAHASANOVIĆ R. T.: A Survey of Procedural Modelling Methods for Layout Generation of Virtual Scenes. *Computer Graphics Forum* (2024). doi:10.1111/cgf.14989. 3
- [CWZ*24] CHENG J., WANG L., ZHANG L., GAO F., ZHANG J.: Single-image SVBRDF estimation with auto-adaptive high-frequency feature extraction. vol. 124, p. 104103. URL: <https://www.sciencedirect.com/science/article/pii/S0097849324002383>, doi:<https://doi.org/10.1016/j.cag.2024.104103>. 3
- [CYP91] CLARKE J., YANESKE P., PINNEY A. A.: The harmonisation of thermal properties of building materials, 1991. https://www.aivc.org/sites/default/files/airbase_6124.pdf [Accessed: (12/11/2024)]. 5
- [DAD*18] DESCHARENTRE V., AITTALE M., DURAND F., DRETTAKIS G., BOUSSEAU A.: Single-image SVBRDF capture with a rendering-aware deep network. *ACM Trans. Graph.* 37, 4 (July 2018). URL: <https://doi.org/10.1145/3197517.3201378>, doi:10.1145/3197517.3201378. 3
- [DDB20] DESCHARENTRE V., DRETTAKIS G., BOUSSEAU A.: Guided fine-tuning for large-scale material transfer. *CoRR abs/2007.03059* (2020). URL: <https://arxiv.org/abs/2007.03059>, arXiv:2007.03059. 3
- [Em24] ENVI-MET: Leading 3d modelling software for urban cooling and climate adaptive planning, 2024. <https://envi-met.com/> [Accessed: (05/09/2024)]. 2, 3
- [ESR03] ESRU O.: The ESP-r System for Building Energy Simulation: User Guide Version 10 Series. *University of Strathclyde: Glasgow, Scotland* (2003). URL: https://www.esru.strath.ac.uk/Documents/ESP-r_userguide.pdf. 2
- [fBERCoNM] FOR BUILDING ENERGY RESEARCH U.-I. J. C., (CBERD) D., OF NEW M., (MNRE) R. E.: Thermo-physical-optical property database of construction materials. https://carbse.org/administrator/upload/PDF/_1716812084.pdf [Accessed: (12/11/2024)]. 5
- [FHR*23] FREUDE C., HAHN D., RIST F., LIPP L., WIMMER M.: Precomputed radiative heat transport for efficient thermal simulation. *Computer Graphics Forum* 42, 7 (2023), e14957. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14957>, doi:<https://doi.org/10.1111/cgf.14957>. 2
- [FKY*24] FUJIWARA K., KHOMIAKOV M., YAP W., IGNATIUS M., BILJECKI F.: Microclimate vision: Multimodal prediction of climatic parameters using street-level and satellite imagery. *Sustainable Cities and Society* 114 (2024), 105733. URL: <https://www.sciencedirect.com/science/article/pii/S2210670724005584>, doi:<https://doi.org/10.1016/j.scs.2024.105733>. 3
- [FLZ*25] FREUDE C., LIPP L., ZEZULKA M., RIST F., WIMMER M., HAHN D.: Inverse simulation of radiative thermal transport. *Computer Graphics Forum* 44, 2 (2025), e70048. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.70048>, doi:<https://doi.org/10.1111/cgf.70048>. 2
- [HMB22] HEISEL F., MCGRANAHAN J., BOGHOSSIAN A.: ScanR: A composite building scanning and survey method for the evaluation of materials and reuse potentials prior to demolition and deconstruction. *IOP Conference Series: Earth and Environmental Science* 1078 (09 2022), 012012. doi:10.1088/1755-1315/1078/1/012012. 3
- [IBC*22] IBARRART L., BLANCO S., CALIOT C., DAUCHET J., EIBNER S., EL-HAFI M., FARGES O., FOREST V., FOURNIER R., GAUTRAIS J., KONDURU R., PENAZZI L., TREGAN J.-M., VOURC'H T., YAACOUB D.: Advection, diffusion and linear transport in a single

- path-sampling Monte-Carlo algorithm : getting insensitive to geometrical refinement. working paper or preprint, Oct. 2022. URL: <https://hal.science/hal-03818899.2>
- [Joh21] JOHRA H.: *Thermal properties of building materials - Review and database*. No. 289 in DCE Technical Reports. Department of the Built Environment, Aalborg University, 2021. doi:10.54337/aau456230861.5
- [JZH*23] JIN L., ZHANG J., HOLD-GEOFFROY Y., WANG O., BLACKBURN-MATZEN K., STICHA M., FOUHEY D. F.: Perspective Fields for Single Image Camera Calibration. 17307–17316. URL: <https://doi.org/10.1109/CVPR52729.2023.01660>, doi:10.1109/CVPR52729.2023.01660.4
- [KBKL21] KIM J., BAE H., KANG H., LEE S. G.: Cnn algorithm for roof detection and material classification in satellite images. *Electronics* 10, 13 (2021). URL: <https://www.mdpi.com/2079-9292/10/13/1592>, doi:10.3390/electronics10131592.3
- [KMKK01] KUSAKA H., KONDO H., KIKEGAWA Y., KIMURA F.: A Simple Single-Layer Urban Canopy Model For Atmospheric Models: Comparison With Multi-Layer And Slab Models. *Boundary-Layer Meteorology* 101, 3 (Dec. 2001), 329–358. doi:10.1023/A:1019207923078.2
- [KMW19] KATAL A., MORTEZAZADEH M., WANG L. L.: Modeling building resilience against extreme weather by integrated CityFFD and CityBEM simulations. *Applied Energy* 250 (2019), 1402–1417. URL: <https://www.sciencedirect.com/science/article/pii/S0306261919308499>, doi:https://doi.org/10.1016/j.apenergy.2019.04.192.2
- [KOIK21] KOUMBEM W. N. D., OUE'DRAOGO I., ILBOUDO W., KIENO P.: Numerical study of the thermal performance of three roof models in hot and dry climates. *Modeling and Numerical Simulation of Material Science* 11 (01 2021), 35–46. doi:10.4236/mnsms.2021.112003.5
- [KSS*21] KAMARAJ N. P., SHEKHAR S., SIVASHANKARI V., BALASUBRAMANI K., PRASAD K. A.: Detecting heat-inducing urban built-up surface material with multi remote sensing datasets using reflectance and emission spectroscopy. *Remote Sensing of Environment* 264 (2021), 112591. URL: <https://www.sciencedirect.com/science/article/pii/S0034425721003114>, doi:https://doi.org/10.1016/j.rse.2021.112591.3
- [LCAGDM*22] LOPEZ-CABEZA V. P., ALZATE-GAVIRIA S., DIZ-MELLADO E., RIVERA-GOMEZ C., GALAN-MARIN C.: Albedo influence on the microclimate and thermal comfort of courtyards under mediterranean hot summer climate conditions. *Sustainable Cities and Society* 81 (2022), 103872. URL: <https://www.sciencedirect.com/science/article/pii/S2210670722001998>, doi:https://doi.org/10.1016/j.scs.2022.103872.3
- [LLLL24] LIU H., LI C., LI Y., LEE Y. J.: Improved baselines with visual instruction tuning, 2024. URL: <https://arxiv.org/abs/2310.03744>, arXiv:2310.03744.4
- [LP19] LUZIANI S., PARAMITA B.: Autodesk green building studio an energy simulation analysis in the design process. *KnE Social Sciences* (2019), 735–749. doi:10.18502/kss.v3i21.5007.2
- [LRM23] LIU D., RIST F., MICHELS D. L.: Urbanflow: Designing Comfortable Outdoor Areas. In *2023 Annual Modeling and Simulation Conference (ANNSIM)* (2023), pp. 628–644. URL: <https://arxiv.org/abs/2204.01117>.2
- [LSBE24] LUO X., SCANDOLO L., BOUSSEAU A., EISEMANN E.: Single-image SVBRDF estimation with learned gradient descent. In *Computer graphics forum* (2024), vol. 43, Wiley Online Library, p. e15018. doi:10.1111/cgf.15018.3
- [LW23] LIU Y., WANG Z.: Research progress and hotspot analysis of urban heat island effects based on cite space analysis. *Land* 12, 6 (2023). URL: <https://www.mdpi.com/2073-445X/12/6/1154>, doi:10.3390/land12061154.3
- [Map24] MAPILLARY: Mapillary metropolis dataset, 2024. <https://www.mapillary.com/dataset/metropolis/> [Accessed: (12/11/2024)].3
- [Mas00] MASSON V.: A physically-based scheme for the urban energy budget in atmospheric models. *Boundary-layer meteorology* 94 (2000), 357–397. doi:10.1023/A:1002463829265.2,3
- [MCR02] MARTILLI A., CLAPPIER A., ROTACH M. W.: An urban surface exchange parameterisation for mesoscale models. *Boundary-layer meteorology* 104 (2002), 261–304. doi:10.1023/A:1016099921195.2
- [MR20] MAHMOUD H., RAGAB A.: Impact of urban geometry on cooling loads in egypt : Case study: A social residential compound in new aswan city. In *2020 9th International Conference on Power Science and Engineering (ICPSE)* (2020), pp. 71–75. doi:10.1109/ICPSE51196.2020.9354366.2
- [MRR*22] MARTIN R., ROULLIER A., ROUFFET R., KAISER A., BOUBEKEUR T.: MaterIA: Single Image High-Resolution Material Capture in the Wild. vol. 41, pp. 163–177. doi:10.1111/cgf.14466.3
- [MSCG24] MILLER B., SAWHNEY R., CRANE K., GKIOULEKAS I.: Walkin' robin: Walk on stars with robin boundary conditions. *ACM Trans. Graph.* 43, 4 (July 2024). URL: <https://doi.org/10.1145/3658153>, doi:10.1145/3658153.2
- [MWH*06] MÜLLER P., WONKA P., HAEGLER S., ULMER A., VAN GOOL L.: Procedural modeling of buildings. *ACM Trans. Graph.* 25, 3 (jul 2006), 614–623. URL: <https://doi.org/10.1145/1141911.1141931>, doi:10.1145/1141911.1141931.3
- [MZWVG07] MÜLLER P., ZENG G., WONKA P., VAN GOOL L.: Image-based procedural modeling of facades. *ACM Trans. Graph.* 26, 3 (July 2007), 85–es. URL: <https://doi.org/10.1145/1276377.1276484>, doi:10.1145/1276377.1276484.3
- [Ome] OMEGA: Emissivity of common materials. <https://www.omega.co.uk/literature/transactions/volumel/emissivityb.html> [Accessed: (12/11/2024)].5
- [Ope04] OPENSTREETMAP: Openstreetmap, 2004. <https://www.openstreetmap.org/> [Accessed: (12/11/2024)].2
- [PBC*24] PENAZZI L., BLANCO S., CALIOT C., COUSTET C., EL HAFI M., FOURNIER R., GAUTRAIS J., GOLJANEK-JEDRZEJCZYK A., SANS M.: Path integrals formulations leading to propagator evaluation for coupled linear physics in large geometric models. *Computer Physics Communications* 294 (2024), 108911. URL: <https://www.sciencedirect.com/science/article/pii/S0010465523002564>, doi:https://doi.org/10.1016/j.cpc.2023.108911.2,7
- [PE14] PENG C., ELWAN A.: An outdoor-indoor coupled simulation framework for Climate Change-conscious Urban Neighborhood Design. *Simulation* 90, 8 (2014), 874–891. doi:10.1177/0037549714526293.2
- [Pla24] PLATFORM G. M.: Photorealistic 3d tiles overview, 2024. <https://developers.google.com/maps/documentation/tile/3d-tiles-overview/> [Accessed: (05/09/2024)].3
- [PM01] PARISH Y. I. H., MÜLLER P.: Procedural modeling of cities. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2001), SIGGRAPH '01, Association for Computing Machinery, p. 301–308. URL: <https://doi.org/10.1145/383259.383292>, doi:10.1145/383259.383292.3
- [Pog19] POGRZEBACZ V.: A graph grammar for modelling of 2d shapes, July 2019. URL: <https://www.cg.tuwien.ac.at/research/publications/2019/pogrzebacz-2014-gggm/>.3
- [PSM25] PLOCHARSKI A., SWIDZINSKI J., MUSIALSKI P.: Pro-dg:

- Procedural diffusion guidance for architectural facade generation, 2025. URL: <https://arxiv.org/abs/2504.01571>, arXiv:2504.01571. 3
- [PSPSM24] PLOCHARSKI A., SWIDZINSKI J., PORTER-SOBIERAJ J., MUSIALSKI P.: Façaid: A transformer model for neuro-symbolic facade reconstruction. In *SIGGRAPH Asia 2024 Conference Papers* (New York, NY, USA, 2024), SA '24, Association for Computing Machinery. URL: <https://doi.org/10.1145/3680528.3687657>, doi:10.1145/3680528.3687657. 3
- [RBD23] RAGHU D., BUCHER M. J. J., DE WOLF C.: Towards a 'resource cadastre' for a circular economy – urban-scale building material detection using street view imagery and computer vision. *Resources, Conservation and Recycling* 198 (2023), 107140. URL: <https://www.sciencedirect.com/science/article/pii/S0921344923002768>, doi:<https://doi.org/10.1016/j.resconrec.2023.107140>. 3
- [RGL18] RIDHA S., GINESTET S., LORENTE S.: Effect of the shadings pattern and greenery strategies on the outdoor thermal comfort. *International Journal of Engineering and Technology* 10 (04 2018), 108–114. doi:10.7763/IJET.2018.V10.1043. 3
- [RHL*09] ROBINSON D., HALDI F., LEROUX P., PEREZ D., RASHEED A., WILKE U.: CITYSIM: Comprehensive Micro-Simulation of Resource Flows for Sustainable Urban Planning. In *Proceedings of the Eleventh International IBPSA Conference* (2009), pp. 1083–1090. doi:10.26868/25222708.2009.1083-1090. 2
- [RLZ*24] REN T., LIU S., ZENG A., LIN J., LI K., CAO H., CHEN J., HUANG X., CHEN Y., YAN F., ZENG Z., ZHANG H., LI F., YANG J., LI H., JIANG Q., ZHANG L.: Grounded SAM: Assembling open-world models for diverse visual tasks. URL: <https://arxiv.org/abs/2401.14159>. 4
- [RM01] ROBERTS A., MARSH A.: ECOTECT: environmental prediction in architectural education. pp. 342–347. doi:10.52842/conf.ecaade.2001.342. 2
- [RMBG23] RUIZ M., MASSON V., BONHOMME M., GINESTET S.: Numerical method for solving coupled heat and mass transfer through walls for future integration into an urban climate model. *Building and Environment* 231 (2023), 110028. URL: <https://www.sciencedirect.com/science/article/pii/S0360132323000550>, doi:<https://doi.org/10.1016/j.buildenv.2023.110028>. 3
- [SBM*10] ST'AVA O., BENEŠ B., MĚCH R., ALIAGA D. G., KRIŠTOF P.: Inverse procedural modeling by automatic generation of l-systems. *Computer Graphics Forum* 29, 2 (2010), 665–674. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2009.01636.x>, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8659.2009.01636.x>, doi:<https://doi.org/10.1111/j.1467-8659.2009.01636.x>. 3
- [Scol0] SCOTT K. T.: Property tables and charts, 2010. <https://cecs.wright.edu/people/faculty/sthomas/htappendix01.pdf> [Accessed: (12/11/2024)]. 5
- [SKD*19] SKAMAROCK W., KLEMP J., DUDHIA J., GILL D., LIU Z., BERNER J., WANG W., POWERS J., DUDA M., BARKER D., HUANG X.-Y.: A description of the advanced research wrf model version 4. doi:10.5065/1dfh-6p97. 2
- [SLZ*16] SHUI T., LIU J., ZHANG P., LIU S., ZHAO Z.: Development of an urban canopy model for the evaluation of urban thermal climate with snow cover in severe cold regions. *Building and Environment* 95 (2016), 160–170. URL: <https://www.sciencedirect.com/science/article/pii/S0360132315301232>, doi:<https://doi.org/10.1016/j.buildenv.2015.09.016>. 2
- [Sma06] SMALL C.: Comparative analysis of urban reflectance and surface temperature. *Remote Sensing of Environment* 104, 2 (2006), 168–189. Thermal Remote Sensing of Urban Areas. URL: <https://www.sciencedirect.com/science/article/pii/S0034425706001817>, doi:<https://doi.org/10.1016/j.j.rse.2005.10.029>. 3
- [SMGC23] SAWHNEY R., MILLER B., GKIOULEKAS I., CRANE K.: Walk on stars: A grid-free monte carlo method for pdes with neumann boundary conditions. *ACM Trans. Graph.* 42, 4 (July 2023). URL: <https://doi.org/10.1145/3592398>, doi:10.1145/3592398. 2
- [SP23] SARTOR S., PEERS P.: Matfusion: A generative diffusion model for svbrdf capture. In *SIGGRAPH Asia 2023 Conference Papers* (Dec. 2023), SA '23, ACM, p. 1–10. URL: <http://dx.doi.org/10.1145/3610548.3618194>, doi:10.1145/3610548.3618194. 3
- [TAB*23] TREGAN J., AMESTOY J., BATI M., BÉZIAN J. J., BLANCO S., BRUNEL L., CALIOT C., CHARON J., CORNET J.-F., COUSTET C., D'ALENÇON L., DAUCHET J., DUTOUR S., EIBNER S., MOUNA E. H., EYMET V., FARGES O., FOREST V., FOURNIER R., YAACOUB D.: Coupling radiative, conductive and convective heat-transfers in a single monte carlo algorithm: A general theoretical framework for linear situations. *PLoS one* 18 (04 2023), e0283681. doi:10.1371/journal.pone.0283681. 2, 6
- [TC08] TRUSILOVA K., CHURKINA G.: The response of the terrestrial biosphere to urbanization: land cover conversion, climate, and urban pollution. *Biogeosciences* 5, 6 (2008), 1505–1515. URL: <https://bg.copernicus.org/articles/5/1505/2008/>, doi:10.5194/bg-5-1505-2008. 2
- [TF23] TABATABAEI S. S., FAYAZ R.: The effect of facade materials and coatings on urban heat island mitigation and outdoor thermal comfort in hot semi-arid climate. *Building and Environment* 243 (2023), 110701. URL: <https://www.sciencedirect.com/science/article/pii/S036013232300728X>, doi:<https://doi.org/10.1016/j.buildenv.2023.110701>. 3
- [TNZ14] THIERRY NOUIDUI M. W., ZUO W.: Functional mock-up unit for co-simulation import in energyplus. *Journal of Building Performance Simulation* 7, 3 (2014), 192–202. URL: <https://doi.org/10.1080/19401493.2013.808265>, arXiv:<https://doi.org/10.1080/19401493.2013.808265>, doi:10.1080/19401493.2013.808265. 2
- [Too] TOOLBOX E.: www.engineeringtoolbox.com [Accessed: (12/11/2024)]. 5
- [VHd*22] VILLEFRANQUE N., HOURDIN F., D'ALENÇON L., BLANCO S., BOUCHER O., CALIOT C., COUSTET C., DAUCHET J., HAFI M. E., EYMET V., FARGES O., FOREST V., FOURNIER R., GAUTRAIS J., MASSON V., PIAUD B., SCHOETTER R.: The “teapot in a city”: A paradigm shift in urban climate modeling. *Science Advances* 8, 27 (2022), eabp8934. URL: <https://www.science.org/doi/abs/10.1126/sciadv.abp8934>, arXiv:<https://www.science.org/doi/pdf/10.1126/sciadv.abp8934>, doi:10.1126/sciadv.abp8934. 2
- [VMR*24] VECCHIO G., MARTIN R., ROULLIER A., KAISER A., ROUFFET R., DESCHAINTE V., BOUBEKEUR T.: ControlMat: A Controlled Generative Approach to Material Capture. vol. 43, Association for Computing Machinery. URL: <https://doi.org/10.1145/3688830>, doi:10.1145/3688830. 3
- [VSPS24] VECCHIO G., SORTINO R., PALAZZO S., SPAMPINATO C.: MatFuse: Controllable Material Generation with Diffusion Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), IEEE, pp. 4429–4438. URL: <https://doi.org/10.1109/CVPR52733.2024.00424>, doi:10.1109/CVPR52733.2024.00424. 3
- [WWSR03] WONKA P., WIMMER M., SILLION F., RIBARSKY W.: Instant architecture. *ACM Trans. Graph.* 22, 3 (July 2003), 669–677. URL: <https://doi.org/10.1145/882262.882324>, doi:10.1145/882262.882324. 3
- [WYD*14] WU F., YAN D.-M., DONG W., ZHANG X., WONKA P.: Inverse procedural modeling of facade layouts. *ACM Trans. Graph.*





33, 4 (July 2014). URL: <https://doi.org/10.1145/2601097.2601162>, doi:10.1145/2601097.2601162. 3

[ZGL*25] ZHU P., GUO J., LIU Y., SUN Q., WANG Y., XU K., LIU L., GUO Y.: Appearance-aware Multi-view SVBRDF Reconstruction via Deep Reinforcement Learning. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers* (2025), Association for Computing Machinery. URL: <https://doi.org/10.1145/3721238.3730718>, doi:10.1145/3721238.3730718. 3

[ZK21] ZHOU X., KALANTARI N. K.: Adversarial single-image SVBRDF estimation with hybrid training. In *Computer Graphics Forum* (2021), vol. 40, Wiley Online Library, pp. 315–325. doi:10.1111/cgf.142635. 3

[ZWA*16] ZAMIR A., WEKEL T., AGRAWAL P., WEI C., MALIK J., SAVARESE S.: Generic 3d representation via pose estimation and matching. vol. 9907, pp. 535–553. doi:10.1007/978-3-319-46487-9_33. 3

HeatMat: Simulation of City Material Impact on Urban Heat Island Effect *Supplemental Material*

M.Reinbigler^{1,2}, R.Rouffet², P.Naylor³, M.Czerkawski^{3,4}, N.Dionelis³, E.Brunet¹, C.Fetita⁵ and R.Martin²

¹SAMOVAR, Inria Saclay, Télécom SudParis, Institut Polytechnique de Paris, Palaiseau, France

²Adobe Research, Paris, France

³ESA/ESRIN, ϕ -lab, Frascati, Italy

⁴Asterisk Labs, London, UK

⁵SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, Palaiseau, France

1. Material estimation pipeline validation

To measure the performance of our pipeline and validate our approach, we use multiclass IoU and accuracy metrics. As our facade element segmentation is purely metadata and not localized pixels as in a common segmentation task, we consider as true positive the common percentage of an element with the right material compared to ground truth. Each group of facade elements (ground or upper levels) must add up to 100% in the VLM output, and we enforce it afterward if they don't, by scaling the results. Table 1 presents the ablation regarding the benefits of GroundedSAM and shows performances obtained with Llava1.6 [LLLL24], InternVL2 [CWT*], and ClaudeAI [The]. Even if ClaudeAI has the best results, we pursue with Llava 1.6 to use an open model in our experiments. We use the following prompt:

Could you please describe the main building present in the image, by giving the material and proportion of each facade element ?
Output text as a JSON.
Separate the ground level and the upper levels.

For the ground level, provide:

- 1) the main facade material name and percentage of presence,
- 2) the window material name and percentage of presence,
- 3) the window frame material name and percentage of presence,
- 4) the door material name and percentage of presence.

For the upper levels, provide:

- 1) the main facade material name and percentage of presence,
- 2) the window material name and percentage of presence,
- 3) the window frame material name and percentage of presence,
- 4) the shutter material name and percentage of presence.

Here is the list of different materials:

Brick
Aluminium
Concrete
Metal
Glass
Wood
Terracotta
Limestone
Stone
Cement
Asphalt
Slate

Return only a json like this:

```
'''json
{
  "ground_level": {
    "main_facade_material": "",
    "main_facade_material_percentage": "%",
    "windows_material": "",
    "windows_material_percentage": "%",
    "frames_material": "",
    "frames_material_percentage": "%",
    "doors_material": "",
    "doors_material_percentage": "%",
  },
  "upper_levels": {
    "main_facade_material": "",
    "main_facade_material_percentage": "%",
    "windows_material": "",
    "windows_material_percentage": "%",
    "frames_material": "",
    "frames_material_percentage": "%",
    "shutters_material": "",
    "shutters_material_percentage": "%",
  }
}
'''
```

We illustrate in Fig.1 the process from the equirectangular geolocated image of Mapillary dataset (a.), that we rectify using the Perspective Fields method to get a straight building (b.), before calling

Model	Llava	Llava	InternVL2	Claude
Input	input	SAM	SAM	SAM
Multi-class IoU	0.375	0.416	0.392	0.47
Accuracy	0.521	0.563	0.521	0.59

Table 1: Material estimation performances. The first two columns demonstrate the gain in performance of GroundedSAM. The two others compare Llava with two other VLM: InternVL2 and ClaudeAI. The latter performs best.

GroundedSAM to get the segmentation mask containing facades, that we use to crop the image (c.). With this input, when we query Llava1.6 with our prompt, we get the following material proportion estimation:

```
'''json
{
  "ground_level": {
    "main_facade_material": "Brick",
    "main_facade_material_percentage": "45%",
    "windows_material": "Glass",
    "windows_material_percentage": "25%",
    "frames_material": "Aluminium",
    "frames_material_percentage": "20%",
    "doors_material": "Wood",
    "doors_material_percentage": "10%"
  },
  "upper_levels": {
    "main_facade_material": "Brick",
    "main_facade_material_percentage": "45%",
    "windows_material": "Glass",
    "windows_material_percentage": "30%",
    "frames_material": "Aluminium",
    "frames_material_percentage": "20%",
    "shutters_material": "Wood",
    "shutters_material_percentage": "5%"
  }
}
'''
```




Figure 1: From 360° Mapillary image (a.) to sliced rectified facade (b.) to cropped building facade (c.).

2. VLM pipeline accuracy impact on simulation results

To illustrate the impact of an inaccurate facade composition, using the same simulation parameters as in Fig. 9, we compared the ground truth material composition to four configurations: (1) both material types and percentages are incorrect, (2) wrong materials only, (3) wrong percentages only, (4) the material composition provided by our pipeline. We compared each configuration results to the ground truth using the RMSE (See Table 2). Our pipeline result

Config.	Both wrong	Mat. only	Perc. only	Ours
RMSE (K)	2.79	2.72	2.13	1.92

Table 2: Study the impact of wrong material predictions on the temperature, over 4 configurations, compared to GT.

is accurate enough to detect material impact and capture variations, performing better than arbitrary composition. In addition, the most impactful mistake is the misdetection of material nature, rather than inaccurate percentages.

3. Comparison to Local Climate Zone materials

As an alternative to our VLM pipeline for material estimation, we also consider relying on coarser-grained information from existing databases. In particular, the Local Climate Zone (LCZ) classification [DKM*22] characterizes urban areas as a combination of one of ten standardized build types (density, size, and types of buildings) and one of seven land cover types (soil kind and vegetation if any). LCZ data are provided by area as an aggregate of the characteristics of buildings in that area. In comparison, our VLM pipeline provides material estimation for each building individually, therefore offering more fine-grained modeling.

To evaluate the benefits of our VLM pipeline over LCZ data, we compare the absolute errors compared to our hand-made ground truth with both our VLM pipeline material description and the uniform one obtained with LCZ data. Figure 2 presents the results of the comparison, with an RMSE of 0.59 with the VLM pipeline, and 0.67 (13% higher) when using LCZ materials. As expected, when using LCZ data the errors close to the facades are bigger than when using data from our VLM pipeline. Indeed, LCZ provides only aggregated data on facade materials, while our method more accurately models the specifics of each building.




Figure 2: Heatmap obtained with ground truth materials (a), and difference of absolute errors between heatmaps relative to the Ground Truth with LCZ and VLM materials (b). Positive difference means bigger error with LCZ materials, negative difference means bigger error with VLM materials.

4. Material database

The material properties considered in the experiments on the Detroit city are listed in Tab. 3. The information has been concatenated from various data sources [fBERCoNM, Sco10, CYP91, Joh21, Too, KOIK21, Ome].

5. Common experimental settings for all experiments

In all the experiments mentioned in the context of the paper, the results presented correspond to about 5000 paths initiated per pixel.

Material	Th. capacity (J/K/kg)	Conductivity (W/m/K)	Density (kg/m ³)	Emissivity
Brick	790	0.9	1920	0.93
Aluminium	903	237	2702	0.03
Concrete	880	1.4	2300	0.88
Steel	456	15.6	7913	0.85
Glass	840	1	500	0.93
Wood	1880	0.12	450	0.9
Terracotta	1800	0.8	780	0.6
Limestone	1000	1.7	2200	0.95
Stone	840	2.68	2550	0.87
Cement	920	0.43	1283	0.54
Asphalt	1000	0.5	1700	0.94
Slate	1000	2.2	2400	0.97

Table 3: Material thermal and optical properties.

When involved, the radiative and indirect irradiation paths can bounce at most 30 times. For the conductive path, 700 steps are performed before switching to another heat transfer mode. The number of steps is set so that enough of the facade space is explored to reach the initial time to account for conductive transfer in the estimated temperature. It is set based on empirical observations against the reference provided by satellite imagery for a short time period of a few hours. We perform at most 100 transitions between heat transfer modes if the path doesn't reach an end condition. The path is thus considered invalid and is not involved in the average computation.

6. Stardis experimental settings




Figure 3: 3D geometry input of Stardis, converted into a set of 2D maps for comparison.

For the comparison with the Stardis simulator, we consider the simple city geometry provided by the Stardis Starter Pack [MS] without the balconies and the lake elements, as shown in Fig. 3. It is composed of eight buildings 20m long, 15m large, and 13.7m high. Buildings of a row are separated by 10m, and the two rows by 35m. The initial conditions of each material in the model are described by Tab. 4. In the context of our simulator, we set the facade sampling parameters, extracted from the 3D model, gathered in Tab. 5. In addition, the air temperature is set to 300K and the sky temperature to 280K.

7. Study of the convergence of our simulator

To study the convergence of our simulator, we compute the same simulation with an increasing number of pixels and compute the

Material	Th. capacity (J/K/kg)	Conductivity (W/m/K)	Density (kg/m ³)	Emissivity	Initial Temp. (K)
Wall	2000	1.5	2500	0.9	273
Glazing	2000	0.9	2500	0.9	273
Ground	2000	1	2500	0.9	273

Table 4: Material thermal and optical properties and their initial temperature as input of our simulator for the comparison to Stardis.

Pattern Width (m)	6.6
Pattern Height (m)	2.7
Max door width (m)	4
Facade perimeter (m)	68.6
Facade height (m)	13.7
Upper Floors Window Coverage (%)	6
Ground Floor Window Coverage (%)	6
Upper Floors Facade Coverage (%)	94
Ground Floor Facade Coverage (%)	94
Ground Floor Door Coverage (%)	0

Table 5: Procedural facade sampling parameters used for the comparison to Stardis.

FLIP [ANAM*20] error against the image obtained after the equivalent of 15000 samples per pixel in Stardis. The FLIP indicator enables to quantify the differences between rendered images and corresponding ground truths, which corresponds to the image assumed converged in our case. Figure 4 demonstrates that after 5000 samples per pixels, the difference between two rendered images decreases slowly, which can constitute a sweet spot for convergence.

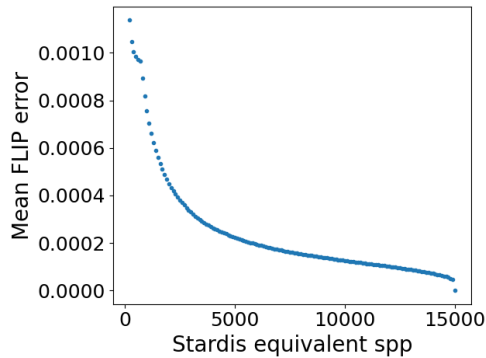


Figure 4: Evolution of the FLIP value according to the number of sample per pixels equivalent in Stardis simulation.

8. The colorized procedural facade generation

Figure 5 details the process of generating the procedural facade and introduces how it can produce a more realistic rendering using our pipeline. As already described in the paper, the procedural facade is generated based on the facade composition extracted from the street-view image of the facade using a Vision Language Model (VLM). In addition, it also relies on the facade perimeter and the

building height in order to create a facade pattern that could be superimposed on extruded 3D geometry like wallpaper. The facade unit element pattern to be repeated depends on two extra parameters. The pattern width, meaning the portion of the wall in which we expect to have one window with its surrounding elements like a frame or shutters horizontally. The pattern height corresponds to the floor height. These values are constants and specified for each city according to its main features. In addition, for a more realistic rendering, the segmented main facade extracted by the Segment Anything Model is given as input of a color palette extractor, thus providing the set of the 4th majoritarian colors in the facades as well as their occurrence.

By combining both the procedural facade and the color palette, we can procedurally generate a simplified version of the real facade statistically coherent, to be further rendered with the right colors. It could be even a basis to create a realistic view by applying each material its own procedural textures, parametrized by the extracted colors, which could also embed architectural features specific to the studied city.

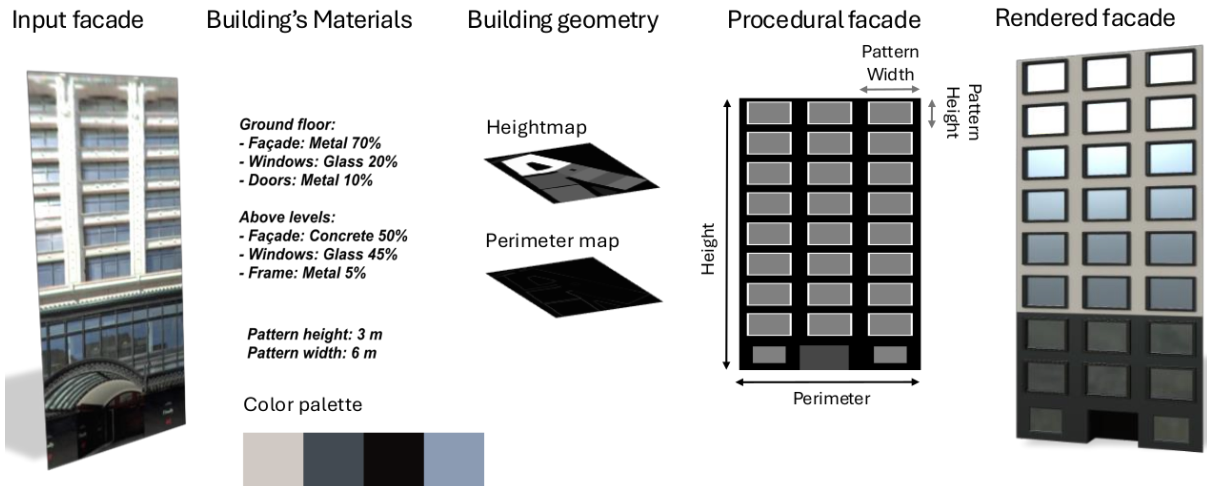


Figure 5: A full pipeline providing a realistic rendering of the facade by also extracting the color palette from the street-view image. Realism can be further improved by combining it with procedural material.

References

- [ANAM*20] ANDERSSON P., NILSSON J., AKENINE-MÖLLER T., OSKARSSON M., ÅSTRÖM K., FAIRCHILD M. D.: Flip: A difference evaluator for alternating images. *Proc. ACM Comput. Graph. Interact. Tech.* 3, 2 (Aug. 2020). URL: <https://doi.org/10.1145/3406183>, doi:10.1145/3406183. 4
- [CWT*] CHEN Z., WANG W., TIAN H., YE S., GAO Z., CUI E., TONG W., HU K., LUO J., MA Z., ET AL.: InternVL2. <https://internvl.github.io/>. 2
- [CYP91] CLARKE J., YANESKE P., PINNEY A. A.: The harmonisation of thermal properties of building materials, 1991. https://www.aivc.org/sites/default/files/airbase_6124.pdf [Accessed: (12/11/2024)]. 3
- [DKM*22] DEMUZERE M., KITNER J., MARTILLI A., MILLS G., MOEDE C., STEWART I. D., VAN VLIET J., BECHTEL B.: A global map of local climate zones to support earth system modelling and urban-scale environmental science. *Earth System Science Data* 14, 8 (2022), 3835–3873. URL: <https://essd.copernicus.org/articles/14/3835/2022/>, doi:10.5194/essd-14-3835-2022. 3
- [fBERCoNM] FOR BUILDING ENERGY RESEARCH U.-I. J. C., (CBERD) D., OF NEW M., (MNRE) R. E.: Thermo-physical-optical property database of construction materials. https://carbse.org/administrator/upload/PDF/_1716812084.pdf [Accessed: (12/11/2024)]. 3
- [Joh21] JOHRA H.: *Thermal properties of building materials - Review and database*. No. 289 in DCE Technical Reports. Department of the Built Environment, Aalborg University, 2021. doi:10.54337/aau456230861. 3
- [KOIK21] KOUMBEM W. N. D., OUEÐRAOGO I., ILBOUDO W., KIENO P.: Numerical study of the thermal performance of three roof models in hot and dry climates. *Modeling and Numerical Simulation of Material Science* 11 (01 2021), 35–46. doi:10.4236/mnsms.2021.112003. 3
- [LLLL24] LIU H., LI C., LI Y., LEE Y. J.: Improved baselines with visual instruction tuning, 2024. URL: <https://arxiv.org/abs/2310.03744>, arXiv:2310.03744. 2
- [MS] MESO-STAR: Stardis starter pack v0.2. <https://www.meso-star.com/projects/stardis/downloads/Stardis-Starter-Pack-0.2.0.tar.gz> [Accessed: (12/11/2024)]. 3
- [Ome] OMEGA: Emissivity of common materials. <https://www.omega.co.uk/literature/transactions/volumel/emissivityb.html> [Accessed: (12/11/2024)]. 3
- [Scol10] SCOTT K. T.: Property tables and charts, 2010. <https://cecs.wright.edu/people/faculty/sthomas/htappendix01.pdf> [Accessed: (12/11/2024)]. 3
- [The] The Claude 3 Model Family: Opus, Sonnet, Haiku. DOI. 2
- [Too] TOOLBOX E.: www.engineeringtoolbox.com [Accessed: (12/11/2024)]. 3