EUROGRAPHICS 2026 / B. Masia and J. Thies
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 45 (2026), Number 2

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

M.Reinbigler[1,2], R.Rouffet[2], P.Naylor[3], M.Czerkawski[3,4], N.Dionelis[3], E.Brunet[1], C.Fetita[5] and R.Martin[2]

[1]SAMOVAR, Inria Saclay, Télécom SudParis, Institut Polytechnique de Paris, Palaiseau, France
[2]Adobe Research, Paris, France
[3] ESA/ESRIN, ϕ-lab, Frascati, Italy
[4] Asterisk Labs, London, UK
[5] SAMOVAR, 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 Llav1.6 [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 Perpective Fields method to get a straight building (b.), before calling 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%"
}}
```
```

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.*

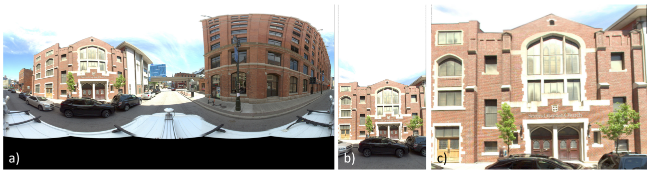

Figure 1: *From 360° Mapillary image (a.) to sliced rectified facade (b.) to cropped building facade (c.).*

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.

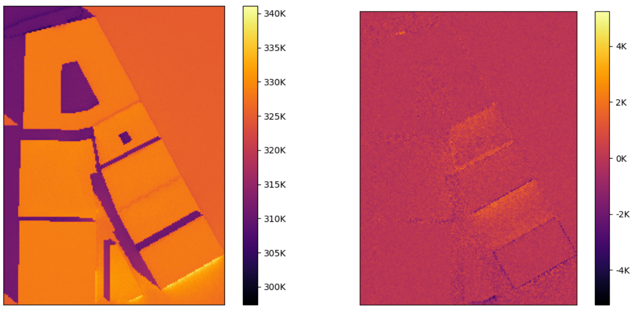

(a) *Ground truth materials* **(b)** *Difference of absolute error*

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

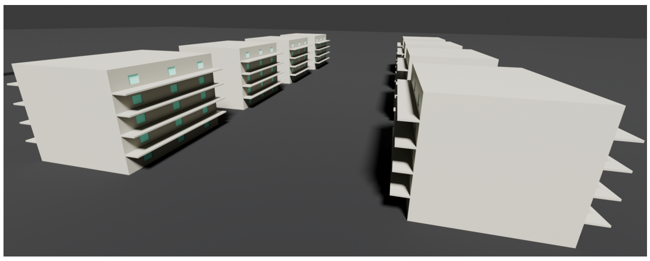

Figure 3: *3D geometry input of Stardis, converted into a set of 2D maps for comparison.*

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.

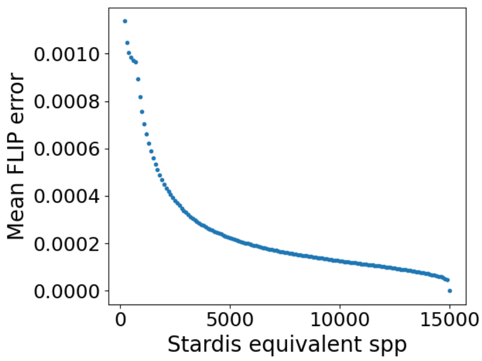

Figure 4: *Evolution of the FLIP value according to the number of sample per pixels equivalent in Stardis simulation.*

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.

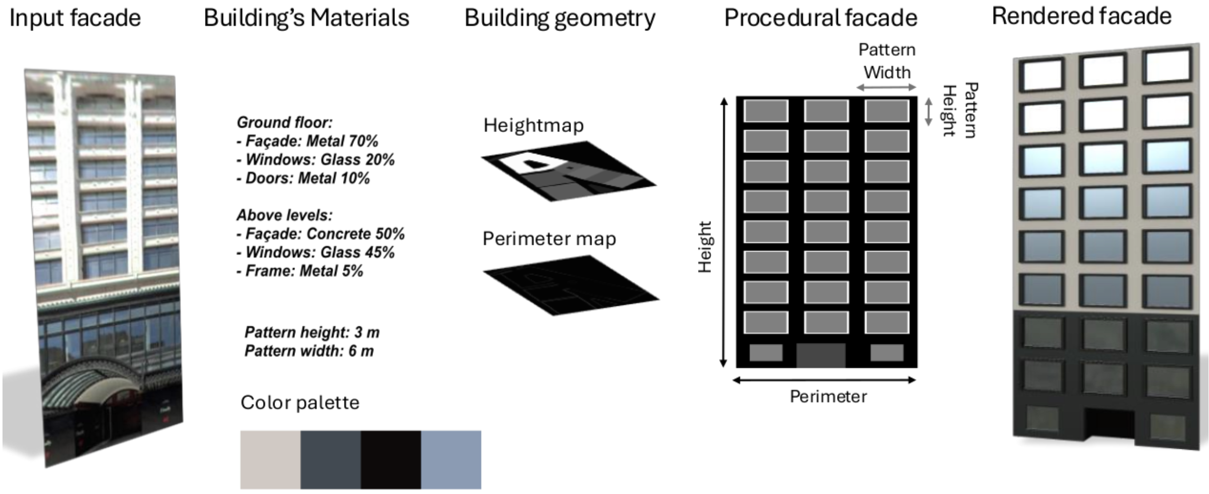

Figure 5: *A full pipeline providing a realistic rendering of the facade by also extracting the color palette from the street-view image. Realism can be further improved by combining it with procedural material.*

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., KITTNER 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., OUÉDRAOGO 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/volume1/emissivityb.html [Accessed: (12/11/2024)]. 3

[Sco10] 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