IEEE Open Journal of the Communications Society

Received XX Month, XXXX; revised XX Month, XXXX; accepted XX Month, XXXX; Date of publication XX Month, XXXX; date of current version XX Month, XXXX.

Human-Centric 6G RF-EMF Exposure with AI-Assisted Digital Twins in Ten Cities

ROBIN WYDAEGHE [1], GÜNTER VERMEEREN [1], *Member, IEEE*, EMMERIC TANGHE [1], *Member, IEEE*, and WOUT JOSEPH [1], *Senior Member, IEEE*

[1] Department of Information Technology, Ghent University/IMEC, 9052 Ghent, Belgium

CORRESPONDING AUTHOR: Robin Wydaeghe (e-mail: robin.wydaeghe@ugent.be).

ABSTRACT Realistic urban radiofrequency electromagnetic-field (RF-EMF) assessment must connect the field in a city to the power absorbed by the human body. This supports monitoring and epidemiological studies. The goal of this study is to compute direction-aware whole-body specific absorption rate (SAR) along pedestrian routes in ten cities. We combine 360-degree street images, photogrammetric geometry, and AI-derived object and material labels in a human-centric digital twin. Its agentic AI stage uses SAM 3 Agent for material and vegetation labeling. The calculation places possible transmitters along visible rooflines at 15 GHz, traces direct and single-reflection paths, and applies each arrival direction to an anatomical body model. Results are normalized per unit product of source density and effective isotropic radiated power. The ten routes contain 163 observation points, each evaluated with 64 independent runs. The normalized route-median whole-body SAR ranges from 0.00902 to 0.129 $\mathrm{m}^2\ \mathrm{kg}^{-1}$, a factor of 14.31. Direct paths give the largest component at 156 of 157 points with line of sight. The single-reflection specular component is largest at one point, while the single-reflection diffuse component is the only nonzero component at six shadowed points. Increasing the number of runs from 48 to 64 changes the total at any point by at most 0.32%. Validation in a controlled scene gives maximum differences of 1.43% against deterministic quadrature and 0.80% against an independent Sionna RT forward calculation. The reported values apply to the selected routes under the stated source model.

INDEX TERMS AI-assisted digital twins, human-centric wireless systems, ray tracing, RF-EMF exposure, whole-body specific absorption rate.

I. Introduction

REALISTIC radiofrequency electromagnetic-field (RF-EMF) exposure assessment must connect the field in a real environment to the power absorbed by the human body. This link matters for street-level exposure monitoring and for exposure estimates used in epidemiological research. Along a city street, buildings can block the direct field, and facade materials change the reflected power [1], [2]. Human whole-body absorption depends on each arrival direction relative to body orientation [3]. Thus, the city changes the power and directions that reach a pedestrian, while orientation changes how they couple into the body. Total incident power alone therefore does not determine whole-body absorption. A route calculation must retain each arrival direction until the field is applied to the body.

Prior exposure studies provide several parts of this route-to-body link. Ray tracing gives detailed paths in three-dimensional city models [4]. Ray tracing has supported city-scale downlink exposure calculations with known base-station locations [5] and realistic 28 GHz body exposure along an outdoor path [6]. Stochastic geometry represents unknown transmitter locations [7]. SAR conversion factors connect incident fields to absorption [8]. Taken together, these studies provide the main propagation, source, and body components. However, they do not combine those components across cities under one source normalization.

The combined calculation also needs surface materials. Detailed photogrammetric geometry [9], [10] does not identify its surface materials. Street images can supply this information. The Mapillary Vistas data set provides object classes for street scenes [11]. Kamari *et al.* assign material classes from street images to city geometry for millimeter-wave ray tracing [12]. Xia *et al.* combine semantic point-cloud classes with detailed scenes for outdoor ray tracing at 2.8 GHz [13]. These studies show that multimodal scene data can support radio propagation. These scene methods have not been joined with a common source model and direction-aware body calculation across routes in several cities.

To address this gap, we form an AI-assisted, human-centric digital twin from 360-degree street images and photogrammetric city geometry. Mask2Former assigns object labels. SAM 3 Agent performs the agentic AI step by assigning material and vegetation labels on visible surfaces. A common model places possible transmitters along rooflines. The propagation calculation keeps direct, single-reflection specular, and single-reflection diffuse components separate, with their arrival directions, until the body calculation. Together, these stages connect the city scene to body absorption. The propagation model includes one reflection.

The goal of this study is to compute direction-aware human whole-body SAR along pedestrian routes in ten cities at 15 GHz. The routes contain 163 observation points with image-derived surface materials and a body model facing along each walk. Every value is normalized per unit active-source areal density and per unit effective isotropic radiated power (EIRP). The selected-route medians differ by a factor of 14.31. These values describe the 163 fixed observations under the stated source model.

This work makes the following three contributions.

1) Multimodal fusion combines 360-degree street images and city geometry with Mask2Former object labels and SAM 3 Agent material labels in an AI-assisted digital twin. Surfaces without a reliable image label keep their geometry-based material.
2) A normalized roofline source model and an adjoint ray tracer keep direct and once-reflected power separate, with arrival directions, until the field is applied to the body.
3) Results for ten routes quantify whole-body SAR, component shares, numerical convergence, surface-material sensitivity, and controlled single-reflection validation.

II. Methods

A. Study configuration

Fig. 1 shows how the inputs come together at Prague Old Town Square. A 360-degree street image is aligned with a photogrammetric city mesh cropped to a 250 m radius around the route. Image labels are assigned to the visible mesh triangles without changing their geometry. Possible transmitters lie along visible rooflines. Fixed points set the observation positions, and the anatomical body model faces the direction of travel at each point.

With this configuration fixed, Table 1 summarizes the study route in each of the ten cities. Each route follows a connected street corridor covered by aligned 360-degree street images. Fixed observation points include positions between image locations, so the number of images and points can differ. The 163 points are fixed case-study observations that describe the selected routes. A different route would change both position and body orientation.

TABLE 1. Study route in each city, with observation-point counts, route spans, and ray calculation times

City	Points	Span (m)	Time (s)
Brussels	14	87.00	32.16
Ghent	10	49.04	6.59
Krakow	16	116.10	31.76
London	22	121.10	73.77
Madrid	14	73.47	8.06
Mexico City	11	60.79	5.95
Milan	23	128.43	69.02
Prague	22	119.39	39.00
Tokyo Hachiko	16	87.38	10.02
Toulouse	15	82.47	37.87
Total	163	925.17	314.20

B. AI-assisted digital twin

Two image models are applied in sequence. First, Mask2Former assigns a Vistas object class such as building, road, or vegetation to every pixel [11], [14]. Next, SAM 3 Agent performs the agentic AI step by assigning material and vegetation labels inside compatible object regions [15]. Its prompt-guided concepts extend the fixed object classes with material evidence. The camera position and viewing direction locate these labels on visible triangles in the city geometry. When several images cover one triangle, their accepted labels are combined. A triangle changes material only when its object and material labels are compatible and meet the selection criteria. All other triangles keep their geometry-based material. The supplementary material gives the prompts, alignment tests, rejected labels, and assignment rules.

Fig. 2 connects the scene-building and exposure stages. The image and geometry branches first form the human-centric digital twin. Pedestrian locations, body orientations, and possible roofline transmitters complete its route-specific inputs. The final stages compute directional radio arrivals and apply them to the body.

C. Exposure calculation

Following the central path in Fig. 2, every route point uses 15 GHz and a circular scene crop with radius $R_{\mathrm{crop}} = 250$ m and area $A_{\mathrm{crop}} = 196{,}349.54\ \mathrm{m}^2$. The receiver position $\mathbf{x}$ is the ray origin and body reference point. The body faces the local direction of travel. Index i denotes a roofline segment, $r_i(\mathbf{x})$ is its range to the receiver, and $\mathbf{r}$ is a point on the body surface.

The possible roofline transmitters in Fig. 2 represent a future distributed massive MIMO (DMaMIMO) architecture only to define possible access-point locations [6], [16]. Therefore, the model uses visible rooflines as possible transmitter locations. The calculation treats the possible sites as independent power sources and distributes expected transmitters uniformly per unit physical length of the roofline visible

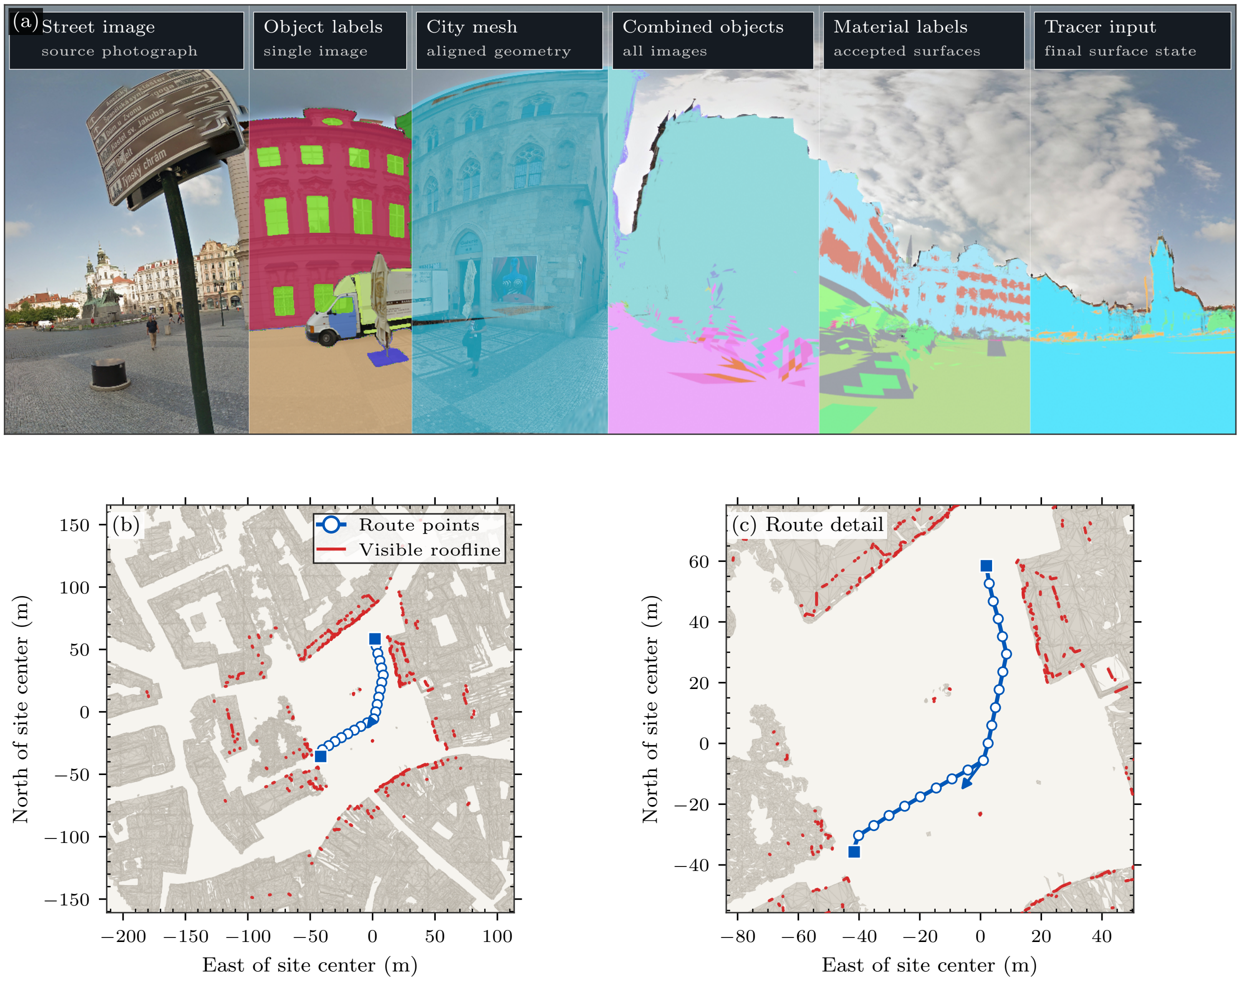

FIGURE 1. **Configuration at Prague Old Town Square. (a) Object and material labels from a 360-degree street image are assigned to the city geometry. (b) Plan view of the 22 observation points and visible roofline. (c) Close-up of the route. The arrow gives the body model's direction of travel.**

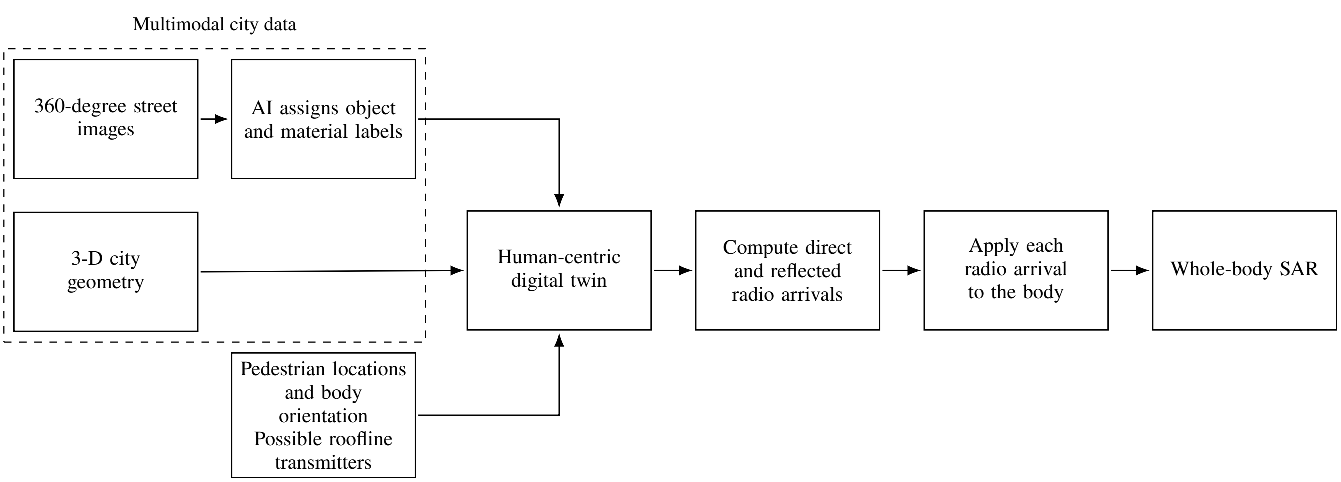

FIGURE 2. **Flowchart from multimodal city data to whole-body specific absorption rate (SAR). Mask2Former assigns object labels, and SAM 3 Agent assigns material labels. These labels, 3-D city geometry, pedestrian locations and body orientations, and possible roofline transmitters form a human-centric digital twin. The propagation and body calculations then give whole-body SAR.**

from the route. The active-site density ρ_A gives an expected count $N_{\text{site}} = \rho_A A_{\text{crop}}$. For a segment with endpoints $\mathbf{p}_i$ and $\mathbf{p}_{i+1}$, its source fraction is

$$p_i = \frac{\ell_i}{\sum_j \ell_j}, \qquad \ell_i = \|\mathbf{p}_{i+1} - \mathbf{p}_i\|_2. \tag{1}$$

The denominator prevents a numerical split of one segment from adding source power. The calculation uses three-dimensional length rather than horizontal length.

An unobstructed reference separates source strength from scene geometry:

$$D_{\text{ref}}(\mathbf{x}) = \sum_i \frac{p_i}{r_i(\mathbf{x})^2}, \qquad S_{\text{ref}}(\mathbf{x}) = \frac{A_{\text{crop}} D_{\text{ref}}(\mathbf{x})}{4\pi}. \tag{2}$$

We apply no visibility test to D_{ref}. All reported values are normalized per unit $\rho_A P_{\text{EIRP}}$. Multiplication by $\rho_A P_{\text{EIRP}}$ gives physical units, but only for a deployment that follows the same roofline transmitter model. The calculation omits transmitters and interactions outside the crop.

The propagation result at each observation point has direct, one-reflection specular, and one-reflection diffuse parts. Let $\mathcal{D}$ contain visible direct paths and $\mathcal{S}_1$ contain accepted one-reflection specular paths. Their normalized powers are $\alpha_a = m_a^{(\text{d})}/D_{\text{ref}}$ and $\beta_b = m_b^{(\text{s})}/D_{\text{ref}}$. The diffuse estimate uses normalized angular-cell powers $\widehat{\gamma}_q = \widehat{m}_q^{(\text{f})}/D_{\text{ref}}$ in $Q = 4096$ fixed angular cells. With $\delta_{\widehat{\mathbf{k}}}$ denoting a unit point mass in arrival direction $\widehat{\mathbf{k}}$, the directional distribution is

$$\mu_{\mathbf{x}} = \sum_{a \in \mathcal{D}} \alpha_a \delta_{\widehat{\mathbf{k}}_a^{(\text{d})}} + \sum_{b \in \mathcal{S}_1} \beta_b \delta_{\widehat{\mathbf{k}}_b^{(\text{s})}} + \sum_{q=1}^{Q} \widehat{\gamma}_q \delta_{\widehat{\mathbf{k}}_q^{(\text{f})}}. \tag{3}$$

The first two sums keep exact path directions and powers. Only diffuse power uses the angular cells. Fig. 3 contrasts the two sampling paths for this diffuse term. *Adjoint ray tracing* launches rays from the observation point and tests visibility from each first surface hit to every roofline source [17], [18]. This source connection is a next-event estimate [19]. The surface law combines unpolarized Fresnel power with a Rayleigh roughness term. The surface law assigns the remaining reflected power to a Lambertian diffuse component, and all components add incoherently. Each sampled path ends after the diffuse reflection. Higher-order specular paths are also omitted. Woody vegetation identified by the images is transparent to rays because the city geometry has no canopy volume. The supplementary material gives the full laws and parameters.

The last two blocks in Fig. 2 make the propagation-to-body handoff explicit. Accordingly, the calculation keeps each arrival direction until body coupling. For outward surface normal $\widehat{\mathbf{n}}(\mathbf{r})$, let $g(\mathbf{r}, \widehat{\mathbf{k}}) = [\widehat{\mathbf{n}}(\mathbf{r}) \cdot (-\widehat{\mathbf{k}})]_+$. The normalized absorbed power density is

$$\widetilde{S}_{\text{ab}}(\mathbf{r}, \mathbf{x}) \equiv \frac{S_{\text{ab}}(\mathbf{r}, \mathbf{x})}{\rho_A P_{\text{EIRP}}} = T_0 S_{\text{ref}}(\mathbf{x}) \int_{\mathbb{S}^2} g(\mathbf{r}, \widehat{\mathbf{k}})\, \mathrm{d}\mu_{\mathbf{x}}(\widehat{\mathbf{k}}). \tag{4}$$

The angular integral sums the direct, specular, and diffuse arrivals while retaining their incoming directions. Here, T_0 is the normal-incidence power-transmission coefficient from the IT'IS tissue database at 15 GHz [20]. It represents incoherent, unpolarized illumination. The function g gives zero absorption where the surface faces away from the incoming direction. We apply this law to the 56,024-element Duke anatomical mesh [21]. With triangle area $A_{\mathbf{r}}$ and body mass m_{body}, the normalized endpoints are

$$\widetilde{P}_{\text{abs}}(\mathbf{x}) = \sum_{\mathbf{r}} A_{\mathbf{r}} \widetilde{S}_{\text{ab}}(\mathbf{r}, \mathbf{x}), \qquad \widetilde{\text{SAR}}_{\text{wb}}(\mathbf{x}) = \frac{\widetilde{P}_{\text{abs}}(\mathbf{x})}{m_{\text{body}}}. \tag{5}$$

$\widetilde{P}_{\text{abs}}$ has units m^2, and $\widetilde{\text{SAR}}_{\text{wb}}$ has units $\text{m}^2\ \text{kg}^{-1}$ per unit $\rho_A P_{\text{EIRP}}$.

D. Numerical settings and route statistics

For each independent run, the calculation launches 200,000 primary rays per observation point and collects single-reflection diffuse power in 4,096 fixed angular cells. Direct and specular paths are not binned on this grid, which does not control the launch directions. The calculations use 64 runs with seeds 7 through 70 and keep cumulative results after 4, 8, 12, 16, 32, 48, and 64 runs. Only the diffuse estimate varies between runs. Whole-body SAR values and body-surface fields are averaged before route statistics are computed. The ten routes give 10,432 directional body fields from 2.086 billion primary rays. Route quantiles use linear interpolation over equally weighted points, and empirical distribution positions are $(\text{rank} - 0.5)/n$. Excess above direct-path power is computed only where direct power is positive. Points with zero direct power remain in the SAR distribution but have no finite excess value. Ray calculation takes 5.95 to 73.77 s per city on one NVIDIA A6000 GPU. These times exclude image acquisition, alignment, depth estimation, and surface-label assignment because their full times were not recorded.

III. Results

A. Validation

Before applying the model to the ten city routes, validation against deterministic quadrature and Sionna RT tests the single-reflection diffuse estimate. Fig. 4 shows the comparison in an open-square scene with 27 sources, 6 receivers, 8 triangles, and one diffuse reflection. The scene has no specular reflection, refraction, or diffraction. Deterministic surface quadrature uses 2,097,152 samples. The adjoint calculation uses 50,000 primary rays for each of four seeds. The independent Sionna RT forward calculation uses 50,000 samples per source for each of three seeds. The maximum adjoint-to-quadrature difference is 1.43% for the reflected component. A separate comparison of total received power gives a maximum adjoint-to-Sionna difference of 0.80%. Fig. 4 also shows the reflected component and its signed

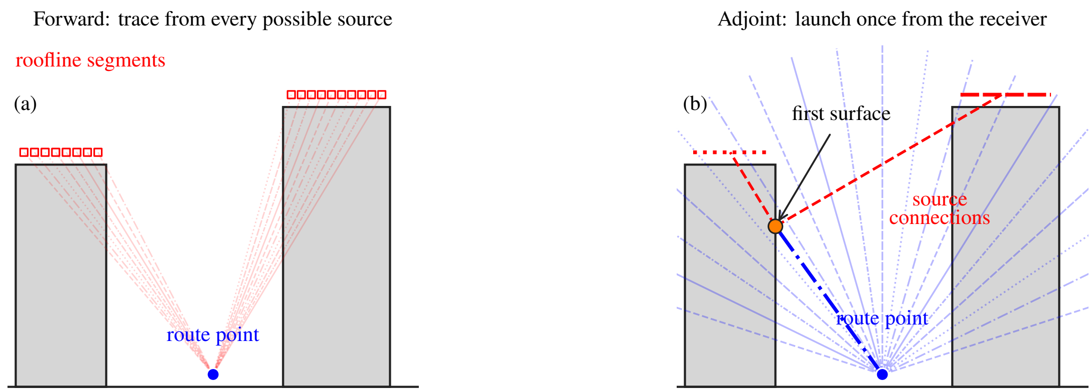

FIGURE 3. **Forward and adjoint sampling for the single-reflection diffuse component. (a) Forward sampling launches rays from every roofline segment. (b) Adjoint sampling launches one set of rays from the observation point and tests visible source connections at the first surface hit.**

percentage difference from quadrature. This test covers diffuse normalization, visibility, inverse-square loss, and cosine terms in a one-reflection scene.

B. Route exposure

After this controlled check, Fig. 5 compares the route distributions and component shares across all 163 observation points. Table 2 complements the figure with exact route quantiles and excess above direct-path power. Normalized whole-body SAR is reported in $\mathrm{m}^2\ \mathrm{kg}^{-1}$ per unit $\rho_A P_{\mathrm{EIRP}}$. A physical value requires multiplication by a deployment's areal source density and EIRP. The figure includes all 163 observation points and marks the six with zero direct and zero single-reflection specular power. The largest selected-route median (Mexico City, 0.1291 $\mathrm{m}^2\ \mathrm{kg}^{-1}$) is 14.31 times the smallest (Milan, 0.009016 $\mathrm{m}^2\ \mathrm{kg}^{-1}$). Mexico City and Tokyo Hachiko have the widest spread along a route and the lowest tails.

Fig. 5(b) shows the additive component shares. Of the 157 points with line of sight, direct power is largest at 156. The single-reflection specular component is largest at Brussels point 2. Mexico City points 0, 1, and 3 and Tokyo Hachiko points 13, 14, and 15 have zero direct and zero specular power. The single-reflection diffuse component is the only nonzero contribution at these six points. Across all 163 points, the separate component medians are 78% direct, 20% specular, and 0.5% diffuse. These medians do not sum to 100% because each is computed separately. The small pooled diffuse median does not represent the six shadowed points.

C. Convergence across runs

Between 48 and 64 runs, the largest pointwise change in total received power is 0.32% at a shadowed point in Mexico City. The next largest change is 0.13% in Tokyo Hachiko. The maximum is below 0.023% on each of the other eight routes, and every route-median whole-body SAR is stable. At one shadowed point in Mexico City, the largest single-run diffuse estimate is 5,738 times the median run estimate. This concentration explains the larger lower-tail variation. This comparison quantifies ray-sampling convergence for the fixed routes. The supplementary material gives the complete 48-to-64-run comparison.

D. Material sensitivity

A paired control for Madrid and Mexico City replaces all image-derived materials with geometry-based defaults. The geometry, route, rooflines, body, seeds, number of rays, and propagation steps remain unchanged. In Madrid, the image-derived case increases normalized whole-body SAR by 5.5%, 5.9%, and 6.4% at q_{10}, q_{50}, and q_{90}. In Mexico City, the image-derived case gives a roughly 305-fold larger q_{10}, a 3.6% smaller median, and a 2.4% larger q_{90}. The large lower-tail ratio occurs at three shadowed points where both values are near zero and only diffuse power is nonzero. Direct power is identical in every pair. In Madrid, image-derived materials increase the median specular component by 54% and reduce the median diffuse component by 94%, but increase the total median by only 5.9%. The control includes both surface parameters and the treatment of woody vegetation as nonblocking. Independent surface measurements are required for a material-accuracy assessment. The supplementary material gives pointwise and component results.

IV. Discussion

Route-median normalized whole-body SAR differs by a factor of 14.31 across the ten selected routes. Fig. 2 links this spread to street width, building height, roofline visibility, surface materials, canopy, and directional body coupling. Each distribution describes its fixed route and observation points. Therefore, the 14.31-fold contrast describes the selected routes under this model.

The single-reflection diffuse component is small at most points, with a pooled median of 0.5%, but it is the only nonzero contribution at the six fully shadowed points. A

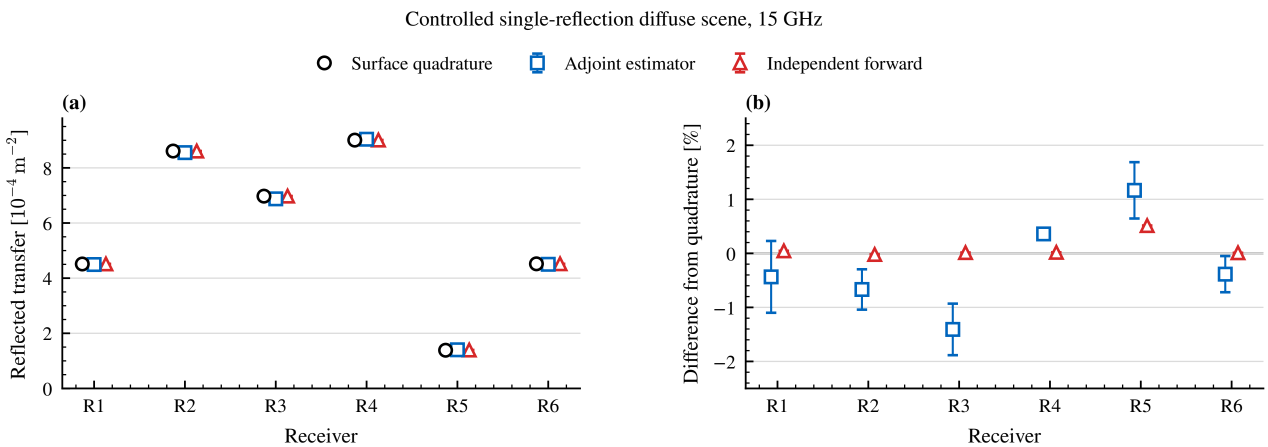

FIGURE 4. Controlled validation of the single-reflection diffuse component. (a) Deterministic surface quadrature, the adjoint estimate, and the independent Sionna RT forward calculation at six receivers. (b) Signed percentage difference from quadrature. Error bars give the standard error across four adjoint seeds and three Sionna seeds.

TABLE 2. Fixed-route exposure summary. Whole-body SAR is normalized per unit $\rho_A P_{\mathrm{EIRP}}$. Excess above direct-path power is computed only at points with nonzero direct power.

	Normalized whole-body SAR [m^2 kg^{-1}]				
City	q_{10}	q_{50}	q_{90}	Median excess above direct [dB]	Shadowed points
Brussels	0.02502	0.03119	0.03371	1.43	0
Ghent	0.05784	0.06205	0.06867	1.16	0
Krakow	0.006717	0.02252	0.04927	0.92	0
London	0.01175	0.01501	0.01744	0.90	0
Madrid	0.02091	0.02238	0.02317	1.55	0
Mexico City	1.01×10^{-6}	0.1291	0.2958	0.69	3
Milan	0.008529	0.009016	0.009861	1.14	0
Prague	0.01216	0.01307	0.01460	1.09	0
Tokyo Hachiko	3.63×10^{-5}	0.009674	0.02520	0.84	3
Toulouse	0.02431	0.03022	0.04292	0.84	0

propagation model without this component would assign zero exposure to those positions. Therefore, the diffuse component is necessary to retain nonzero exposure at these positions.

The paired control shows how image-derived surface information changes the body result. Direct power is identical in both cases. In Madrid, the specular and diffuse changes have opposite signs, while the total route median increases by only 5.9%. The Mexico City median decreases by 3.6%. Its large lower-tail ratio comes from three shadowed points where both estimates are near zero. The image-derived case also treats identified canopy as nonblocking because the city geometry has no canopy volume. The control tests both the surface parameters and vegetation rule. Independent surface measurements are required for a material-accuracy assessment.

This study covers ten selected routes at one frequency with one body model, one roofline source model, and one set of image-derived surface labels. The normalized values can be scaled to a specific deployment when its transmitter distribution matches the source model. The propagation calculation includes direct paths and one specular or diffuse reflection. The controlled comparison covers the diffuse component in a one-reflection scene. Variation across 64 runs quantifies random ray-sampling error. Geometry, image alignment, surface labels, route choice, transmitter placement, body shape, and body orientation form separate sources of uncertainty. The recorded GPU times cover the ray and body calculations.

V. Conclusion

This study contributes a human-centric digital twin that uses SAM 3 Agent for agentic material assignment, a normalized roofline source model, and direction-aware body coupling for RF-EMF exposure assessment. Across 163 observation points on routes in ten cities, normalized route-median whole-body SAR differs by a factor of 14.31. Direct power is largest at 156 of 157 points with line of sight. The specular component is largest at one point, and a diffuse

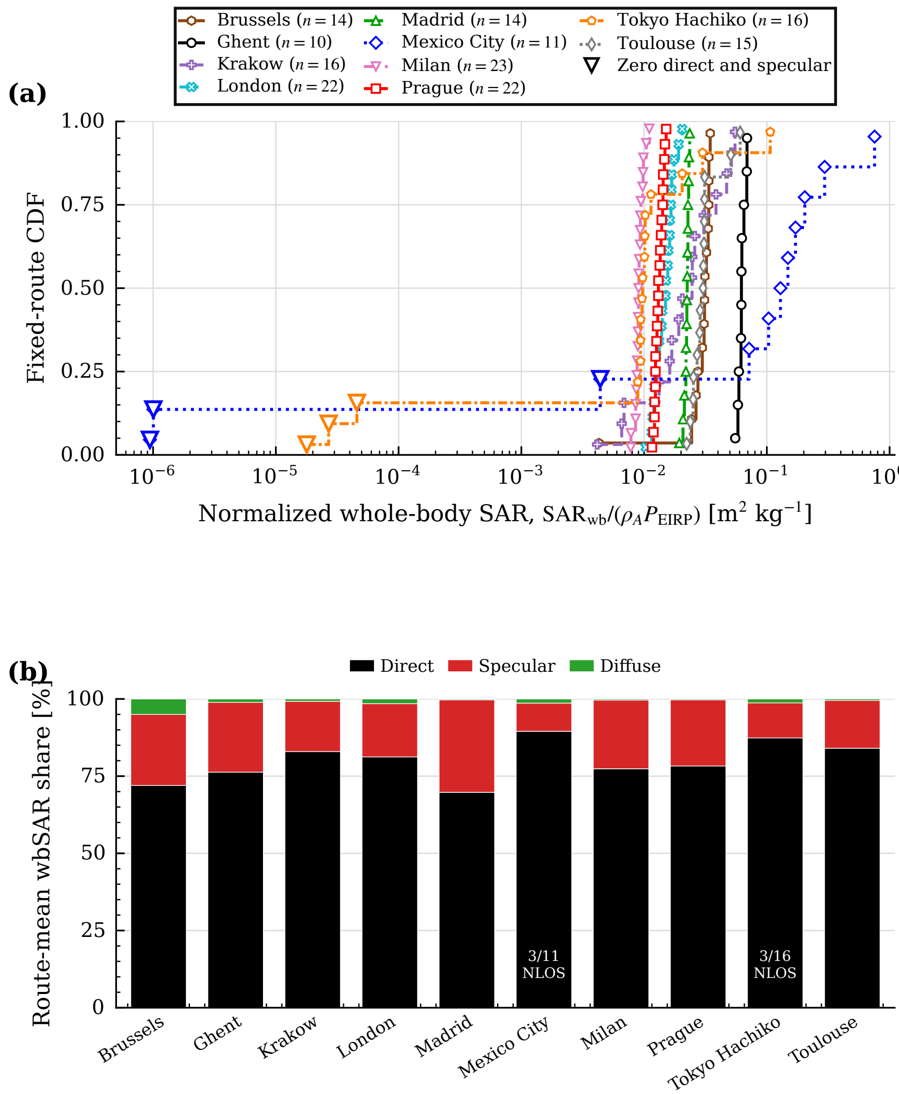

FIGURE 5. Normalized whole-body SAR on the ten fixed routes. (a) Empirical CDFs include all 163 observation points. Hollow triangles mark the six points with zero direct and zero single-reflection specular power. (b) Route-mean component shares. The distributions describe the selected routes under the stated source model.

path is the only nonzero contribution at six shadowed points. Controlled validation gives maximum differences of 1.43% from deterministic quadrature and 0.80% from Sionna RT. Increasing the number of runs from 48 to 64 changes the total at any point by at most 0.32%. These values describe the selected routes under the stated roofline source model.

Image-derived materials increase the Madrid route median by 5.9% and decrease the Mexico City route median by 3.6% relative to geometry-based defaults. Future work will consist of outdoor field validation, measured transmitter distributions, higher reflection orders, and tests of more routes, frequencies, body models, and orientations.

Data and Code Availability

The ten-route results, analysis scripts, and figure scripts are kept in the study repository. A stable public archive with a versioned digital object identifier will be deposited before publication. The 360-degree street images and commercial photogrammetric tiles are governed by their providers' terms and are not redistributed. Their identifiers are retained so the inputs can be obtained again where the licenses allow it.

Acknowledgment

AI tools were used to assist with manuscript preparation. The authors reviewed and verified the scientific claims, numerical values, references, and final text.

REFERENCES

[1] International Telecommunication Union, "Effects of building materials and structures on radio-wave propagation in the range of 1 MHz to 450 GHz," Recommendation ITU-R P.2040-4, 2025.

[2] E. M. Vitucci, V. Degli-Esposti, F. Mani *et al.*, "Tuning ray tracing for mm-wave coverage prediction in outdoor urban scenarios," *Radio Sci.*, vol. 54, no. 11, pp. 1112–1128, 2019, doi: 10.1029/2019RS006869.

[3] International Commission on Non-Ionizing Radiation Protection, "Guidelines for limiting exposure to electromagnetic fields (100 kHz to 300 GHz)," *Health Phys.*, vol. 118, no. 5, pp. 483–524, 2020, doi: 10.1097/HP.0000000000001210.

[4] J. Hoydis, F. Aït Aoudia, S. Cammerer, M. Nimier-David, N. Binder, G. Marcus, and A. Keller, "Sionna RT: Differentiable ray tracing for radio propagation modeling," in *Proc. IEEE Globecom Workshops*, 2023, pp. 317–321, doi: 10.1109/GCWkshps58843.2023.10465179.

[5] M. Leeman, R. Wydaeghe, J. Van der Straeten, S. Goegebeur, G. Vermeeren, and W. Joseph, "City-scale spatio-temporal modeling of 5G downlink exposure of users and non-users by ray-tracing in a real urban environment," *IEEE Access*, vol. 13, pp. 30894–30906, 2025, doi: 10.1109/ACCESS.2025.3541352.

[6] R. Wydaeghe, S. Shikhantsov, G. Vermeeren, L. Martens, E. Tanghe, and W. Joseph, "Hybrid ray-tracing-QuaDRiGa/FDTD method for realistic 28 GHz exposure with 6G CF-MaMIMO in 3D outdoor environments," *npj Wireless Technol.*, vol. 2, no. 1, Art. no. 13, 2026, doi: 10.1038/s44459-026-00031-4.

[7] C. Wiame, S. Demey, L. Vandendorpe, P. De Doncker, and C. Oestges, "Joint data rate and EMF exposure analysis in Manhattan environments: Stochastic geometry and ray tracing approaches," *IEEE Trans. Veh. Technol.*, vol. 73, no. 1, pp. 894–908, 2024, doi: 10.1109/TVT.2023.3307226.

[8] N. Varsier, D. Plets, Y. Corre, G. Vermeeren, W. Joseph, S. Aerts, L. Martens, and J. Wiart, "A novel method to assess human population exposure induced by a wireless cellular network," *Bioelectromagnetics*, vol. 36, no. 6, pp. 451–463, 2015, doi: 10.1002/bem.21928.

[9] Google, "Photorealistic 3D Tiles," 2023. [Online]. Available: https://developers.google.com/maps/documentation/tile/3d-tiles.

[10] vvoovv, "Blosm for Blender: OpenStreetMap, Google 3D cities, terrain," GitHub, 2023. [Online]. Available: https://github.com/vvoovv/blosm.

[11] G. Neuhold, T. Ollmann, S. Rota Bulò, and P. Kontschieder, "The Mapillary Vistas dataset for semantic understanding of street scenes," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 5000–5009, doi: 10.1109/ICCV.2017.534.

[12] A. Kamari, Y. Chae, and P. Pathak, "mmSV: mmWave vehicular networking using Street View imagery in urban environments," in *Proc. 29th Annu. Int. Conf. Mobile Comput. Netw.*, 2023, pp. 1–16, doi: 10.1145/3570361.3613291.

[13] G. Xia, C. Zhou, F. Zhang, Z. Cui, C. Liu, H. Ji, X. Zhang, Z. Zhao, and Y. Xiao, "Path loss prediction in urban environments with Sionna-RT based on accurate propagation scene models at 2.8 GHz," *IEEE Trans. Antennas Propag.*, vol. 72, no. 10, pp. 7986–7997, 2024, doi: 10.1109/TAP.2024.3451214.

[14] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, and R. Girdhar, "Masked-attention mask transformer for universal image segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 1290–1299, doi: 10.1109/CVPR52688.2022.00135.

[15] N. Carion *et al.*, "SAM 3: Segment anything with concepts," arXiv:2511.16719, 2025, doi: 10.48550/arXiv.2511.16719.

[16] H. Q. Ngo, A. Ashikhmin, H. Yang, E. G. Larsson, and T. L. Marzetta, "Cell-free massive MIMO versus small cells," *IEEE Trans. Wireless Commun.*, vol. 16, no. 3, pp. 1834–1850, 2017, doi: 10.1109/TWC.2017.2655515.

[17] A. Behlouli, P. Combeau, L. Aveneau, S. Sahuguede, and A. Julien-Vergonjanne, "Efficient simulation of optical wireless channel application to WBANs with MISO link," *Procedia Comput. Sci.*, vol. 40, pp. 190–197, 2014, doi: 10.1016/j.procs.2014.12.027.

[18] Y. Cocheril and R. Vauzelle, "A new ray-tracing based wave propagation model including rough surfaces scattering," *Prog. Electromagn. Res.*, vol. 75, pp. 357–381, 2007, doi: 10.2528/PIER07061202.

[19] E. Veach, "Robust Monte Carlo methods for light transport simulation," Ph.D. dissertation, Stanford Univ., Stanford, CA, USA, 1997. [Online]. Available: https://graphics.stanford.edu/papers/veach_thesis/

[20] P. A. Hasgall *et al.*, "IT'IS database for thermal and electromagnetic parameters of biological tissues," Version 4.2, 2024, doi: 10.13099/VIP21000-04-2.

[21] A. Christ *et al.*, "The Virtual Family: Development of surface-based anatomical models of two adults and two children for dosimetric simulations," *Phys. Med. Biol.*, vol. 55, no. 2, pp. N23–N38, 2010, doi: 10.1088/0031-9155/55/2/N01.

Robin Wydaeghe received the B.Sc. and M.Sc. degrees in engineering physics from Ghent University, Ghent, Belgium, in 2019 and 2021, respectively. He is currently pursuing the Ph.D. degree in engineering physics with Ghent University. His research interests include computational electromagnetics, numerical assessment of human radiofrequency electromagnetic-field exposure, and propagation modeling for next-generation wireless networks.

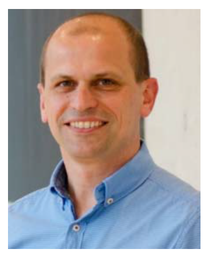

Günter Vermeeren received the M.Sc. degree in industrial engineering from KAHO Sint-Lieven, Ghent, Belgium, in 1998, the M.Sc. degree in electrical engineering from Ghent University, Belgium, in 2001, and the Ph.D. degree in electrotechnical engineering from Ghent University in 2013. Since 2002, he has been a Research Engineer with the Department of Information Technology, Ghent University. His research interests include numerical modeling and measurements of electromagnetic radiation in the domain of radiofrequency dosimetry, electromagnetic exposure, on-body propagation, and medical imaging systems.

Emmeric Tanghe received the M.Sc. and Ph.D. degrees in electrical engineering from Ghent University, Ghent, Belgium, in 2005 and 2011, respectively. In 2015, he became a Part-Time Professor in medical applications of electromagnetic fields in and around the human body. Since 2011, he has been a Postdoctoral Researcher with Ghent University/IMEC, where he focuses on propagation modeling. From 2012 to 2018, he was a Postdoctoral Fellow of FWO-V (Research Foundation-Flanders).

Wout Joseph received the M.Sc. degree in electrical engineering from Ghent University, Ghent, Belgium, in 2000, and the Ph.D. degree in electrical engineering from Ghent University in 2005. Since 2009, he has been a Professor in the domain of experimental characterization of wireless communication systems. His research interests include measuring and modeling electromagnetic fields around base stations for mobile communications, electromagnetic exposure assessment, propagation for wireless communication systems, and antennas and calibration.