Deep Umbra: A Generative Approach for Sunlight Access Computation in Urban Spaces

Kazi Shahrukh Omar, Gustavo Moreira, Daniel Hodczak,
Maryam Hosseini, Nicola Colaninno, Marcos Lage, and Fabio Miranda

Abstract—Sunlight and shadow play critical roles in how urban spaces are utilized, thrive, and grow. While access to sunlight is essential to the success of urban environments, shadows can provide shaded places to stay during the hot seasons, mitigate heat island effect, and increase pedestrian comfort levels. Properly quantifying sunlight access and shadows in large urban environments is key in tackling some of the important challenges facing cities today. In this paper, we propose Deep Umbra, a novel computational framework that enables the quantification of sunlight access and shadows at a global scale. Our framework is based on a conditional generative adversarial network that considers the physical form of cities to compute high-resolution spatial information of accumulated sunlight access for the different seasons of the year. We use data from seven different cities to train our model, and show, through an extensive set of experiments, its low overall RMSE (below 0.1) as well as its extensibility to cities that were not part of the training set. Additionally, we contribute a set of case studies and a comprehensive dataset with sunlight access information for more than 100 cities across six continents of the world. Deep Umbra is available at urbantk.org/shadows.

Index Terms—Urban computing, Urban analytics, Sunlight access, Shadow, Generative adversarial networks.

✦

1 INTRODUCTION

The urban population surpassed the rural one in 2007, and by 2030 it is expected to represent 60% of the global population and 68% by 2050 [1]. This leads to rapid urban development around the world, requiring sustainable solutions, social equity, and urban regeneration [2]. Urban growth implies two types of development: horizontal (through the use of land for new infrastructure, e.g., buildings, parks) and vertical (through densification and construction of taller buildings). Denser cities, however, can impact urban environmental conditions and quality of life, influencing solar access and shadows in outdoor spaces and recreational areas [3]. The appropriate management of shadows then plays a vital role in maintaining the citizens' quality of life. As previously highlighted [4], it is crucial to have efficient methods that allow stakeholders to (1) comprehensively analyze the shadow and sunlight access in a region, and (2) democratize the planning process, allowing communities and the general audience to be aware of the potential disruptions of their *right to light*.

In previous years, this problem has been tackled from different angles and by different domains [5], [6], [7], [8]. However, such studies were oftentimes not scalable beyond a handful of cities (or even just neighborhoods), only quantifying shadows over a limited set of timeframes. At the

core of our research is the understanding that increasing awareness about this pressing issue and supporting sustainable urban development efforts require new methods and datasets that enable the *comprehensive* and *fast* analysis of shadows in cities around the world over long periods of time. By comprehensive, we argue that it is not enough for analyses to be constrained to a handful of timestamps, but there is a need to *accumulate* shadows and generate data inventories with the percentage of sunlight access at a given point over a long period of time – e.g., a given point is under shadow for 70% of time from sunrise to sunset. By fast, we argue that the aforementioned accumulation computations should be performed as fast as possible, as to allow for large-scale interactive and exploratory analyses. Previous approaches tackled these challenges by ray tracing to accumulate shadows over time [4], but still with performance results that limit the scale and scope of analyses. Popular GIS tools, for example, take hours to compute the shadow / sunlight access data for a single large city, such as New York City (NYC). Such performance limits the scale and scope of analyses by urban experts.

This research is supported by the growing availability of urban data describing the built environment (in our case, building heights) as well as recent developments in machine and deep learning that enable image-to-image translation tasks. With sufficient training data, we frame our problem around tile images with height information and propose a *height tile→accumulated shadow tile* translation task. We also take into account a city's latitude and time of the year, as these two factors greatly impact the shape of shadows. With this in mind, we propose Deep Umbra, a novel computational framework that uses a conditional generative adversarial network (GAN) for the quantification of sunlight access and shadows over long periods of time and at a global scale. Our framework considers the physical form

- K. Omar, G. Moreira, D. Hodczak, and F. Miranda are with the University of Illinois Chicago.
  E-mail: {komar3, gmorei3, dhodcz2, fabiom}@uic.edu
- M. Hosseini is with the Massachusetts Institute of Technology.
  E-mail: maryamh@mit.edu
- N. Colaninno is with the Polytechnic University of Milan.
  E-mail: nicola.colaninno@polimi.it
- M. Lage is with the Universidade Federal Fluminense.
  E-mail: mlage@ic.uff.br

of cities to compute high-resolution spatial information of *accumulated* sunlight access for the different seasons of the year. As shadows also depend on the latitude of a city, a key aspect of our research is to ensure transferability to varying locations around the world. We enforce this by training our model in urban areas from different continents and with disparate physical forms. Our experiments show how this approach leads to good results in cities outside the training set. Moreover, in order to enable research in domains interested in the impact of shadows and sunlight access in public spaces, we make our outcomes publicly available at urbantk.org/shadows.

Contributions. In this paper, we take a step towards enabling the fast accumulation of shadows over time using a conditional generative adversarial network. We first describe how shadow accumulation can be quantified over geographical regions. We then introduce a conditional GAN framework to efficiently accumulate shadows over time. This approach is then used to compute accumulated shadow information for over 100 cities in the world. We also report two case studies performed by urban experts (co-authors of the paper) using the data. These studies highlight the opportunities created by such data, not only for fine-grained analyses across neighborhoods of a city, but also large-scale exploratory analyses across different cities. Our contributions can then be summarized as follows:

- We propose Deep Umbra, a computational framework for the quantification of sunlight access and shadows over long periods of time with a speed up of $6\times$ against the state of the art [4].
- We introduce the Global Shadow Dataset, a comprehensive dataset with the accumulated shadow information for over 100 cities in 6 continents.
- We show experimental evaluation demonstrating the accuracy of our framework, including its transferability to cities outside of the training set.
- In collaboration with urban experts, we demonstrate the usefulness of Deep Umbra and the Global Shadow Dataset through a set of case studies evaluating sunlight access over street networks and parks.

This paper is organized as follows: In Section 2 we review important related work; Section 3 briefly reviews information regarding the accumulation of shadows; Section 4 and Section 5 present Deep Umbra and the accumulation shadow data, respectively; In Section 6 we present a detailed evaluation of Deep Umbra; Section 7 presents the Global Shadow Dataset; Section 8 highlights two case studies. Finally, Section 9 concludes our paper and presents future research directions.

2 RELATED WORK

In this section, we review previous works highlighting the significance of shadows to different domains, shadow computation techniques and GAN-assisted urban analytics. **Study of sunlight access and shadows**. Due to population growth, urban constructions have increasingly started to propagate vertically, creating the problem of uneven solar access in cities. The concept named "Solar Envelope" was first introduced by Knowles [9]. His idea was to combine urban building design with the impact of shadow and solar accessibility. Studies have extended upon this idea of solar envelope by proposing vertical limits and distance regulations for sites under construction to ensure sunlight access for dwellings in cities [10], [11], [12]. In addition, design layouts for buildings have been proposed in several works aiming at utilizing solar potential [13], [14], [15].

Sunlight access and impact of shadow are key factors for a comfortable microclimate at street level for pedestrians. Studies have extensively discussed the contribution of the street network design and solar orientation to outdoor shadow impact and thermal comfort levels [16], [17]. In major cities, the surge of high-rise buildings is also sinking the public spaces into deepening shadow, calling into question the protection of sunlight in open spaces. Different studies have been conducted to properly measure sunlight access in parks and other public recreational areas [18], [19], [20].

By making available both our model as well as a comprehensive global dataset with shadow information, our contributions can greatly increase the spatial coverage of previous works and offer the first steps towards a more comprehensive understanding of this pressing problem.

Shadow computation techniques. At the core of our paper is determining whether a point is under shadow or not for one or more timestamps. This problem has been comprehensively explored in computer graphics [21]. In general, real-time shadow computation can be performed through *shadow mapping*, an image-based approach in which the 3D scene is projected onto a 2D plane situated between the scene and the light source. In contrast, ray tracing is a robust and popular approach in which computation is carried out by tracing rays from the light source to the objects [22]. More recently, deep learning approaches have been proposed for the computation of shadows [23], [24]. These techniques, however, are constrained to single timestep shadows, i.e., calculating scene shadows at a particular time. While single-step computation might be useful for real-time shadow assessment at particular timesteps, analyses using the accumulation of shadows over multiple timesteps pave the way for environmental impact studies [25] and city-scale urban planning [26]. Deep Umbra's focus is therefore on performant *accumulation* of shadows over time, which requires innovative approaches to handle not only the spatial aspect (i.e., where casters of shadows are located), but also the temporal aspect (i.e., how shadows are moving over time) of the task.

In recent years, several open and proprietary solutions have emerged proposing analyses of the interplay between built environment and shadows, including online services [27], [28] and GIS tools [4], [29], [30]. JveuxDuSoleil [28] is a web-based tool for single timestep shadows. ShadeMap is a closed source web-based tool that allows for the accumulation of shadows, though the on-the-fly computations are constrained to areas currently being rendered on screen, limiting comparisons across neighborhoods or cities. Both QGIS and ArcGIS are fully-fledged GIS tools, but suffer from poor performance when accumulating shadows. Google has recently released a paid API [31] that provides roof's solar potential for a number of cities. Our previous tool, Shadow Profiler [4], while providing significant gains when compared with previous efforts, still suffer from high latency. Deep Umbra, on the other hand, offers performance
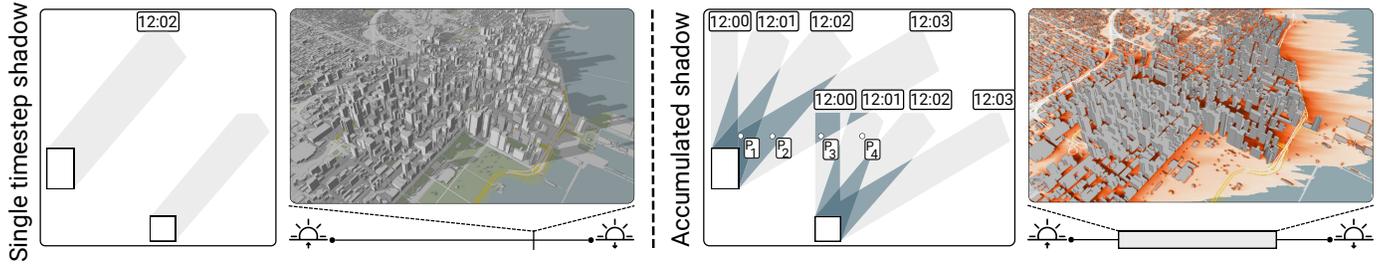
Fig. 1. Left: Example of a single timestep shadow. Right: accumulated shadows. By accumulating shadows between a time range, we can comprehensively analyze the impact of buildings on the sunlight access of public spaces. In the illustration, P_1 is 100% of the time under shadow, P_2 and P_3 50%, and P_4 25%. The rightmost image shows the resulting accumulation when considering a time range between 10 AM and 4 PM.

gains and generalizability that allows for the accumulation of shadows over multiple cities, going beyond the scale tackled by previous approaches.

Deep Umbra also adds to the growing number of open-source and low-cost solutions for urban planning. These solutions include gamification approaches [32], [33] and visual analytics toolkits [34], [35]. Recently, Yap et al. [36] presented an extensive survey of open-source tools that support different stages of the urban planning process.
GAN-assisted urban analytics. The increase in urban data, combined with recent advancements in generative adversarial networks [37], [38], has led to a number of works in the context of urban analytics [39]. Wang et al. [40] proposed a generative adversarial model for urban planning configuration, framing the problem as a task of learning land-use configurations given surrounding spatial contexts. Mokhtar et al. [41] developed a conditional GAN model that is capable of generating pedestrian wind flow approximations for a variety of metropolitan topologies. Zhang et al. [42], Zhang et al. [43], and Yuan et al. [44] use GANs for traffic tasks. Recently, Wu et al. [45] reviewed applications of GANs in various urban tasks. Our work is motivated by the growth of works using GANs to tackle urban-specific problems. Deep Umbra is the first work that uses a generative adversarial approach to model shadow accumulation leveraging building height information and transferring knowledge across cities with varying urban morphologies.

3 BACKGROUND

The study of shadow and sunlight access has been a core consideration in several domains, from architecture, urban planning, and civil engineering, to occupational therapy and environmental sciences. A standard approach in these domains is to perform *single-step shadow analysis*. In other words, only compute the shadow cast on public spaces for a very limited number of timesteps – oftentimes only for a single timestep. As we have highlighted in our previous work [4], there is a need by urban experts for a more comprehensive approach, moving beyond single-step to multi-step (i.e., accumulated) shadow analysis. For example, when comparing the sunlight access in two regions, rather than comparing across individual timestamps (e.g., shadows at every hour), an urban expert can simply compare the accumulated shadow in these two regions. Data from the accumulation of shadows can inform different urban analyses, such as (1) urban heat island effect [25], (2) city

planning (e.g., where to locate new parks) [26], and (3) public policy (e.g., restrictions on building heights) [47].

The aforementioned concept is illustrated in Figure 1. Shadows considering a single timestep (left) offer a limited and superficial glimpse at the condition of sunlight access and shadow in public spaces. However, by accumulating shadows at multiple timesteps (right) we have a much more comprehensive understanding of the state of shadows and sunlight access in public spaces. The example shows the accumulation of shadows for 360 minutes, between 10 AM and 4 PM. Each pixel value is given by the sum of the shadows in each minute inside a time range:

$$A_{i,j}^{s,e} = \sum_{h=s}^{e} Shadow_{i,j}(h)$$

where (i, j) is a point in the 2D plane, s and e are the start and end accumulation times, and $Shadow$ is a function that returns whether the point is under shadow or not at a particular timestep h.

In our example, the final accumulated value of shadows are then mapped to a colorscale, with darker shades of red indicating a larger amount of shadow coverage in the considered period. Even though our proposal can be applied in the computation of arbitrary time ranges, in this paper we focus on three specific days: (1) summer solstice on June 21, (2) March equinox on March 20, (3) winter solstice on December 21. Two important considerations: it is enough to only consider a single equinox for both fall and spring and a summer solstice in the Northern Hemisphere is equivalent to a winter solstice in the Southern Hemisphere. For each one of these days, we will consider accumulations one hour after sunset and one hour before sundown, totaling 720 minutes for (1), 540 for (2) and 360 for (3).

This comprehensive approach was also featured in a *New York Times* article [48], stressing the need for the generation of this type of data not only for domain-specific purposes, but also for the engagement of a broader audience in this pressing issue. In that light, three key aspects of our research are the ability to (1) leverage crowdsourced data that is globally available, (2) adhere to open and reproducible standards, and (3) ensure the transferability of our outcomes to cities in different countries and with varying built environment characteristics.

4 METHODOLOGY

In this section, we describe the main components of our Deep Umbra framework. At its core, Deep Umbra consists of
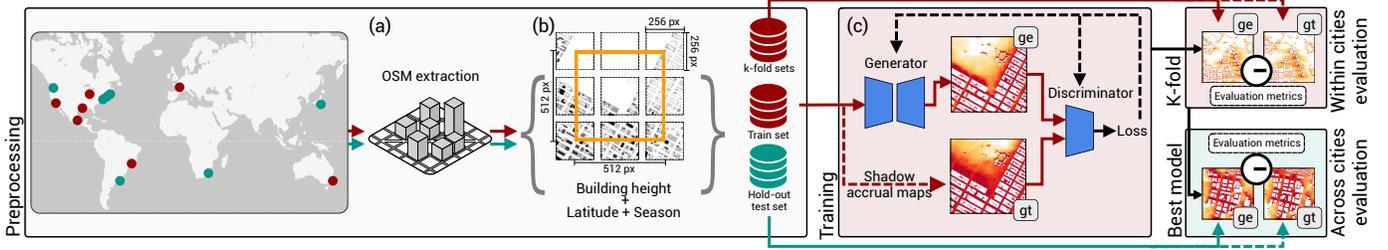
Fig. 2. Overview the different components in the Deep Umbra framework. Step 1, left: The preprocessing component is responsible for extracting building height information from OpenStreetMap [46] (a) and generate 512×512 tiles with this information (b). Step 2, right: The preprocessed data is then used to train a conditional generative adversarial network to map height $\rightarrow$ accumulated shadow, also taking into account season of the year and latitude of the tile (c). Cities used in the training set are shown as ● on the map, and cities used for testing are shown as ●; **gt** refers to ground truth data and **ge** GAN generated data.

two components: (1) a pre-processing component that leverages crowdsourced information (e.g., OpenStreetMap [46]) to generate tiles with building height information; (2) a conditional GAN model to infer the accumulated shadow values over different geographical regions and seasons.

4.1 Framework

The overall framework of Deep Umbra is shown in Figure 2. In Step 1, we extract building height information from datasets that describe the urban morphology in a city. In this paper, we leveraged OpenStreetMap (OSM) data, but our approach is agnostic to the source, also being compatible with point cloud data and already-existing digital elevation models. The geometry description is used to generate 256×256 image tiles, with each pixel containing the average building height at that particular location. To account for the fact that buildings from a tile can also cast shadows on neighboring tiles, we adopt a padding strategy. That is, for each 256×256 tile with height information (i.e., b), our framework pads it with neighboring height tiles, generating a 512×512 height tile that covers the central tile plus parts of its nine adjacent tiles. Figure 3 details this strategy.

In Step 2, we train a model that is able to infer accumulated shadow information given the tiles from the previous step. In this step, we leverage the 512×512 padded tiles. In essence, our objective is to learn a function $\mathcal{G} : B \times L \times T \rightarrow A$, where B is the space containing building information (i.e., height), L represents the spatial information characteristic of a tile (i.e., latitude of a tile), and T is the set of seasons. A will then be the space of completed accumulated shadows. Deep Umbra requires paired input and ground truth images for training. As such, underlying our efforts is the understanding that *data generation*, in our case, is straightforward yet computationally expensive. Given a pre-processed tile (from Step 1), we can leverage off-the-shelf ray-tracing architectures to generate training data. In our case, we make use of *Shadow Accrual Maps* [4], our previously introduced shadow accumulation algorithm. Next, we describe both the model and architecture used to implement Deep Umbra.

4.2 Conditional GAN model

In order to predict the shadow tiles, we leverage the conditional GAN architecture [49] to synthesize tiles with accumulated shadow. A generator in a conditional GAN learns a mapping from input data x and random noise z to an output image a, i.e., $G : x, z \rightarrow y$ [38]. A discriminator D will be trained to distinguish between synthesized and real tiles. More specifically, in our case, the generator G will consider: (1) height tiles $b \in B$ padded with neighboring tiles (as previously mentioned and highlighted in Figure 3); each 512×512 padded tile b represents the building morphology of the location covered by the tile. (2) spatial information $l \in L$ of the tiles, with the associated latitudes of the tile. (3) temporal information $t \in T$, with one of the three considered accumulation periods (summer solstice, March equinox, winter solstice). With (1), Deep Umbra takes into account building morphology, while with (2) and (3) it takes into account the apparent movement of the sun across the sky (i.e, the sun path) at different latitudes and seasons. G will then generate shadow tiles $\hat{a}$ seeking to mimic the ground truth from that region ($a \in A$). In short, we have:

$$x = (b, l, t)$$
$$\hat{a} = G(x)$$

The discriminator D is trained to classify real and synthetic tiles when compared to the ground truth, whereas the generator G tries to fool the discriminator by generating images as close to the ground truth as possible. To increase diversity in the generation task, G uses dropouts for implicit noise [50]. This contest between D and G will be formulated through an objective that contains a reconstruction loss $\mathcal{L}_r$ and an adversarial loss $\mathcal{L}_{adv}$, detailed next.

Reconstruction loss. A reconstruction loss will try to reduce the pixel-wise error between ground truth and the generated tile. Traditional approaches usually rely on L_1 or L_2 distances as the reconstruction loss [38], [51]. In our work, we experimented with both and found out that L_1 performs better than L_2. However, both failed to capture the sharpness and detail of long shadows and instead produced blurry results, especially for the longer winter shadows at higher latitudes. To minimize these issues, we augmented the usual L_1 loss with a structural similarity loss function (SSIM loss) to better capture details in spatially adjacent pixels [52]. In addition, we also examined an edge detection-based Sobel loss function to preserve the sharpness in longer shadows [53]. Following previous works that combine multiple loss functions [24], [53], [54], [55], we performed a set of experiments and found out that a combination of L_1, SSIM, and Sobel losses attains the best results at generating accumulated shadow information. We discuss these in more
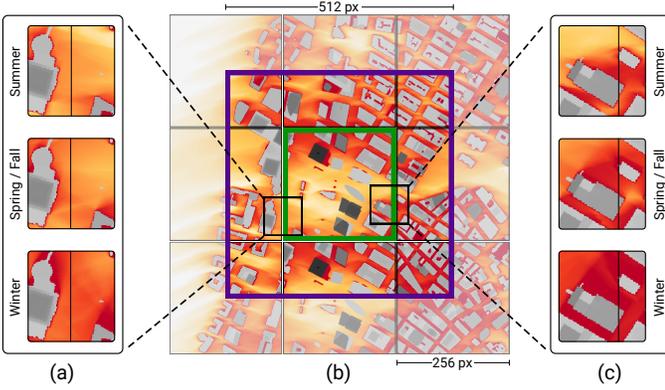
Fig. 3. Example where shadows are cast across tiles. In (b), the two small squares highlight areas where shadows from neighboring buildings are cast onto the center tile. This is further highlighted for different seasons in (a) and (c), with vertical black lines indicating the boundaries between tiles. To account for accumulated shadows from neighboring tiles, during training, a 256×256 height tile is padded with parts of adjacent tiles, resulting in a 512×512 tile (larger purple square in (b)). Model results are cropped back to the original size of 256×256 (smaller green square in (b)).

Fig. 4. Distribution of average tile heights for the train and test sets.

detail as part of the ablation study in Section 6. In short, $\mathcal{L}_r$ is given by:

$$\mathcal{L}_r(G) = \mathcal{L}_1(G) + \mathcal{L}_{SSIM}(G) + \mathcal{L}_{Sobel}(G)$$
$$\mathcal{L}_1(G) = \|a - \hat{a}\|_1$$
$$\mathcal{L}_{SSIM}(G) = 1 - SSIM(a, \hat{a})$$
$$\mathcal{L}_{Sobel}(G) = (Sobel(a) - Sobel(\hat{a}))^2$$

Adversarial loss. The competition between discriminator D and generator G can be formulated by the adversarial loss function as follows:

$$\mathcal{L}_{adv}(G, D) = \mathbb{E}_{x,a}[log D(x, a)] + \mathbb{E}_x[log(1 - D(x, \hat{a}))]$$

During training, the generator G tries to minimize L_{adv}, whereas the goal of discriminator D is to maximize it. This means that the generated shadow tile $\hat{a}$ should be as close to the ground truth a as possible. The discriminator D takes as input either the generated shadow tiles $\hat{a}$ or the ground truth tiles a to determine whether $\hat{a}$ is real or not.

Full objective. The overall objective function of our conditional GAN is composed of the adversarial loss $\mathcal{L}_{adv}$ and reconstruction loss $\mathcal{L}_r$ and can be expressed as following:

$$G^* = \arg \min_G \max_D \mathcal{L}_{adv}(G, D) + \lambda \mathcal{L}_r(G)$$

where λ controls the relative importance of reconstruction loss compared to the adversarial loss. We use $\lambda = 100$.

4.3 Architecture & implementation

To satisfy our requirements and similar to previous conditional GAN approaches, we leverage PatchGAN [38] as the discriminator D. For the generator, we tested two architectures: U-Net [56] and ResNet [57], ultimately choosing ResNet as our generator G. It is important to stress that conditional GANs are difficult to train and sensitive to hyperparameters. Given that Deep Umbra is the first attempt to tackle the problem of generating accumulated shadow maps with spatial and temporal information, we chose to leverage well-known architectures (PatchGAN + ResNet), which in turn allowed us to rely on already established
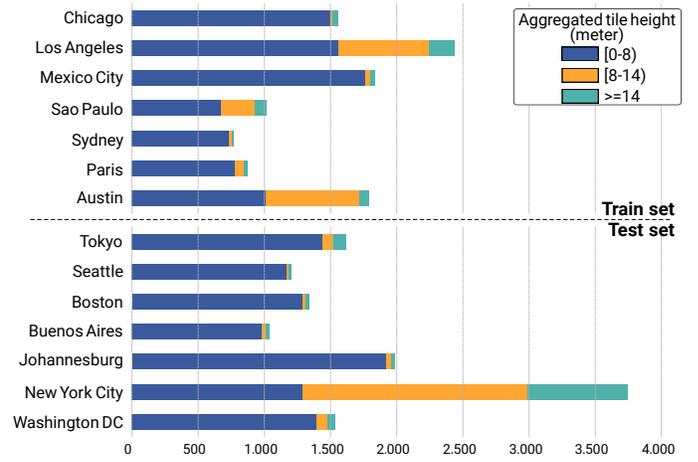
hyperparameters and training procedures. In addition, Section 6 presents a detailed ablation study. In it, we discuss the selection of the architecture and the loss function used in our final model.

Deep Umbra's data $(b, l, t, a, \hat{a})$ follows the tile map format to facilitate implementation in the web environment (e.g., web maps). As previously mentioned, the input and ground truth data are padded tiles with size 512×512 to account for shadows from neighboring tiles. The entire training process (Figure 2 (right)) uses these 512×512 tiles. To adhere to the tile map size of 256×256, we make use of a cropping operator C that removes the outer tiles from the model result ($\hat{a}$), leaving only the central 256×256 tile ($C(\hat{a})$), as depicted in Figure 3.

We use ReLU activation in the generator and Leaky ReLU in the discriminator layers, also incorporating stride in the discriminator for downsampling and fractional stride in the generator for upsampling. Moreover, we apply batch normalization in both the generator and discriminator. Adam optimizer is used with a learning rate of 0.0002 and a momentum of 0.5.

5 DATA DESCRIPTION

We now describe the pipeline to extract height information from OSM data, and the process to compute the ground truth data used to train Deep Umbra.

Building height data. For each city of interest, we first determine the bounding box that covers the entire spatial extension of the urban area and use this information to query OSM for building geometry information. After that, we generate tiles (with size 256×256 pixels) covering the entire region and storing the mean building height in each pixel – we normalize this data considering a maximum building height of 500 meters. In this paper, we focus on tile maps using zoom level 16. This zoom level gives us a good tile resolution of ~ 2.3 meters per pixel. Although we rely on OSM for height information, we realize that crowdsourced initiatives have inherent limitations in terms of spatial coverage. However, we note that our approach can leverage any data source that can generate reliable height information – such as point clouds and high-resolution digital elevation models.

TABLE 1
Results of ablation experiments with different generator architectures. For SSIM, higher values are better.

Generator	RMSE	MAE	MSE	SSIM
U-Net	0.0706	0.0308	0.0073	0.8619
ResNet	**0.0693**	**0.0303**	**0.0071**	**0.8629**

TABLE 2
Results of ablation experiments with different loss functions and ResNet generator. For SSIM, higher values are better.

Loss function	RMSE	MAE	MSE	SSIM
L_1	0.0693	0.0303	0.0071	0.8629
L_2	0.0703	0.0328	0.0072	0.8446
SSIM + L_1	0.0667	0.0292	0.0067	0.8718
SSIM + Sobel	0.0653	0.0296	0.0065	0.8657
SSIM + Sobel + L_1	**0.0643**	**0.0289**	**0.0063**	**0.8724**

Ground truth shadow data. To create the ground truth data to train our model, we use our previous work [4], a GPU-based ray-tracing technique that is optimized for accumulating shadow values. Due to its interpolation approach, it has considerable performance gains even when compared with industry standards, such as Nvidia's OptiX.

6 EXPERIMENTS & RESULTS

To evaluate our framework, we performed experiments considering a large set of cities, taking into account different continents and building morphologies. Our evaluation can be divided into four parts. First, we performed an ablation study to properly evaluate and choose the best configurations for Deep Umbra. Second, we quantitatively evaluated the best configuration from the ablation study on a number of cities across different continents, including cities outside the initial training set. Third, we qualitatively evaluated Deep Umbra, highlighting both general cases as well as large error cases. Finally, we compared Deep Umbra against existing tools and techniques, including commercial and open-source ones. We report root mean square error (RMSE), mean absolute error (MAE), mean squared error (MSE), and structural similarity (SSIM), considering ground truth tiles as a reference to the generated shadow tiles.

To ensure generalizability, our quantitative experiments follow two approaches: first, tiles from the same set of cities are used for the training and test sets (i.e., *within cities evaluation*), and second, tiles from two disjoint sets of cities are used for the test and training sets (i.e., *across cities evaluation*). The cities were carefully picked to ensure that they had buildings with different heights and morphologies. Figure 4 shows the distribution of mean building height per tile for the train and test cities used in the experiments. To account for the scant number of tiles with large mean building height ($>= 14$ meters), we sample the training data with 50% tiles with mean building height below or equal to the 50th percentile and the rest above the 50th percentile.

We implemented our model using TensorFlow. The experiments were executed on a desktop computer with an Intel i7-13700KF CPU, 32 GB of RAM, and an NVIDIA GeForce 3080 with 10 GB of RAM.

Fig. 5. Results of ablation study using U-Net and ResNet generator architecture and L_1 loss function. RMSE and SSIM scores with respect to ground truth are highlighted below each image (for SSIM, higher values are better). Note that ResNet architecture produces significantly better results, especially for longer shadows that are distant from the building. Top: Highlights with areas with longer shadows. Bottom: Zoomed areas.

6.1 Ablation study

As a baseline, we considered a simple model with a U-Net generator and L_1 loss function that solely utilizes building height (b) as input and does not consider latitude (l) and season (t). The RMSE, MAE, MSE, and SSIM scores for the baseline model are 0.1171, 0.0556, 0.0208, and 0.7729, respectively. The next configurations incorporate all three parameters (b, l, t) as the input for the model. Our ablation study uses the *within cities evaluation* setting (see Section 6.2).

Although U-Net is the most commonly used architecture for image-to-image translation [38], Hu et al. [58] have shown that residual-based generators outperform U-Net in certain tasks. Thus, to verify Deep Umbra's effectiveness, we tested two generator architectures: U-Net [56] and ResNet [57]. In both cases, we added L_1 distance along with the adversarial loss to generate outputs closer to the ground truth. The results, summarized in Table 1, reveal that both configurations outperformed the baseline model and that the ResNet architecture had better results than U-Net. We further visually inspected the results and found that, while both architectures perform fairly well for shadows closer to the buildings, ResNet performs significantly better in inferring longer shadows, especially for tiles from the winter season (see Figure 5).

Next, we tested alternative loss functions with the ResNet architecture to further evaluate the model. L_1 and L_2 distances are standard loss functions used in image-to-image translation tasks [38], [51]. In addition, we tested combinations of SSIM and Sobel losses to (1) preserve the structural integrity and (2) improve the sharpness of shadows, especially in long shadows. L_1 and L_2 distances only focus on error sensitivity, whereas SSIM loss, as an alternate measure, evaluates the structural, luminance, and contrast difference between ground truth and generated shadow images [52]. Furthermore, Sobel loss tends to reduce the inherent blurs in shadows [53]. The results, summarized in Table 2, reveal that a combination of SSIM, Sobel, and L_1 loss produces the best outputs. We also visually inspected

| Ground truth | L1 | L2 | SSIM + L1 | SSIM + Sobel | SSIM + Sobel + L1 |

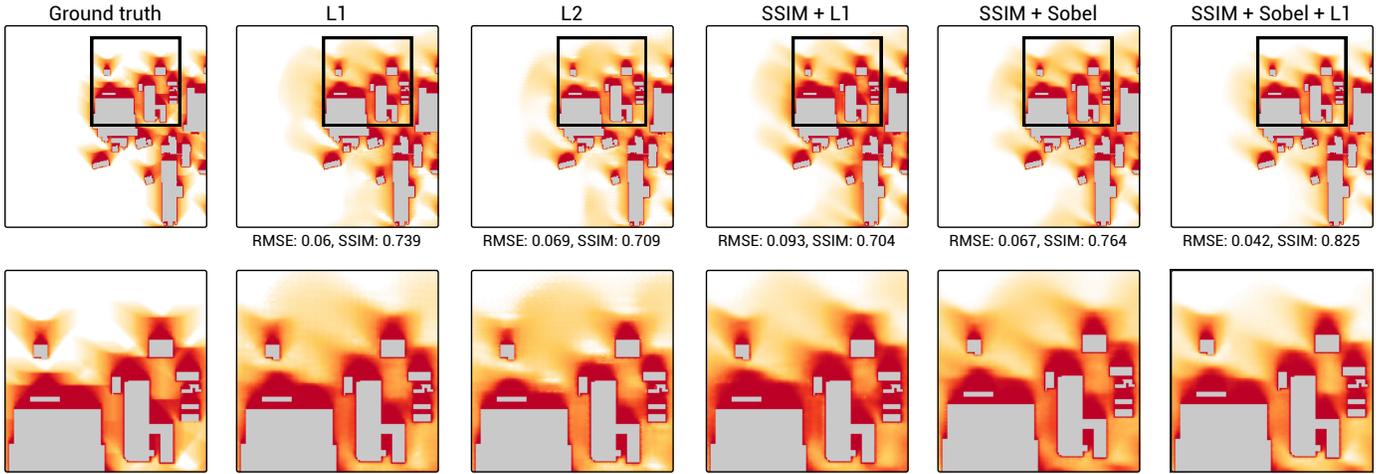| | RMSE: 0.06, SSIM: 0.739 | RMSE: 0.069, SSIM: 0.709 | RMSE: 0.093, SSIM: 0.704 | RMSE: 0.067, SSIM: 0.764 | RMSE: 0.042, SSIM: 0.825 |

Fig. 6. Results of ablation study testing different loss functions with the ResNet architecture. RMSE and SSIM scores with respect to the ground truth are highlighted below each image (for SSIM, higher values are better). SSIM and Sobel losses combined with L_1 distance generate more precise accumulated shadows, whereas combinations of other loss functions are blurrier. We note that complex shapes (such as the Y-like pattern due to the sun's winter path) are more accurately represented with the aforementioned combination. Top: Highlights with areas with longer shadows. Bottom: Zoomed areas.

the results and found out that by combining SSIM and Sobel with L_1 loss, the generated images most closely depicted the ground truth compared to L_1 loss being used alone (see Figure 6). For the next experiments, we used the ResNet model coupled with SSIM, Sobel, and L_1 losses.

6.2 Quantitative evaluation

To quantitatively assess Deep Umbra's error, we performed two evaluations. In the first evaluation, we selected tiles from seven cities to create training and test sets. Then, in the second evaluation, we used the trained model in seven other cities (outside the training set).

Within cities evaluation. In this evaluation, we selected seven cities from four different continents and with varying building morphologies: Los Angeles, Chicago, Austin, Mexico City, São Paulo, Paris, and Sydney (● circles in Figure 2 (left)). In total, we generated ground truth (i.e., ray-traced) data for 20,260 height tiles of these cities. Then, we randomly selected 750 tiles from *each* city in *each* season (i.e., 750×3 tiles per city). We then used k-fold cross validation with $k = 5$ for training and testing. Table 3 shows the average RMSE, MAE, MSE, and SSIM across all 5 runs. RMSE is fairly low at 0.0631, which in practical terms means that, when accumulating 360 minutes for the winter season, the error would be equal to approximately 23 minutes.

Across cities evaluation. In this second evaluation, we selected seven different cities to assess the generalizability of the model trained in the *within cities evaluation*. The seven selected cities were Washington DC, NYC, Boston, Seattle, Johannesburg, Buenos Aires, and Tokyo. In this evaluation, we also selected 750 tiles from each city in each season (i.e., 2,250 tiles per city). Table 4 presents the RMSE, MAE, MSE, and SSIM in each one of the seven target cities – again the model was not trained with data from these cities. High-density cities such as Washington DC, NYC, and Tokyo present the highest RMSE (∼0.08), and Tokyo, in particular, presents the lowest SSIM value. Given the importance of street-level analyses for urban accessibility [59], [60], [61], we further tested Deep Umbra only considering tile pixels

TABLE 3
Within cities analysis. For SSIM, higher values are better.

	RMSE	MAE	MSE	SSIM
K-fold ($k = 5$)	0.0631	0.0276	0.0063	0.9000

TABLE 4
Across cities analysis. For SSIM, higher values are better.

Target city	RMSE	MAE	MSE	SSIM
Washington DC	0.0786	0.0318	0.0091	0.8819
NYC	0.0818	0.0374	0.0084	0.8792
Boston	0.0743	0.0309	0.0081	0.8985
Seattle	0.0597	0.0339	0.0043	0.8578
Johannesburg	0.0281	0.0102	0.0014	0.8749
Buenos Aires	0.0441	0.0146	0.0033	0.9160
Tokyo	0.0837	0.0432	0.0093	0.7981
Average	0.0643	0.0289	0.0063	0.8724

TABLE 5
Across cities performance analysis, considering street networks.

Target city	RMSE	MAE	MSE	SSIM
Washington DC	0.0262	0.0053	0.0013	0.9579
NYC	0.0289	0.0067	0.0012	0.9593
Boston	0.0263	0.0054	0.0013	0.9618
Seattle	0.0249	0.0067	0.0009	0.9452
Johannesburg	0.0095	0.0017	0.0003	0.9697
Buenos Aires	0.0162	0.0027	0.0005	0.9739
Tokyo	0.0447	0.0135	0.0029	0.9172
Average	0.0252	0.0060	0.0012	0.9550

that fall within streets (and not parks or other facilities). The results (Table 5) show that, when only considering streets, our results show even lower errors (mean RMSE of ∼0.02).

6.3 Qualitative evaluation

We also highlight a set of qualitative results from our trained model using tiles from different cities. Figure 7 shows relevant images for NYC, Johannesburg, and Seattle, including
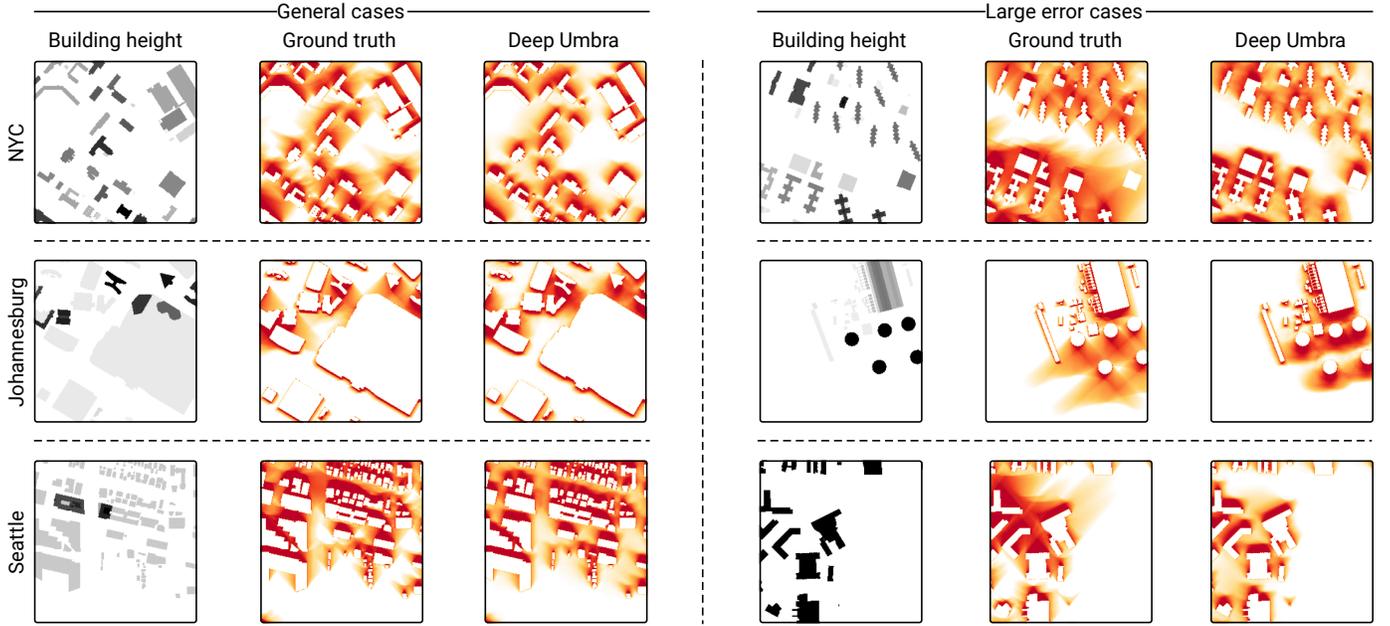
Fig. 7. Results with tiles from NYC, Johannesburg, and Seattle. The examples come from different seasons of the year, but use the same colorscale where darker shades of red correspond to greater shadow coverage during the accumulation period. The left column shows general cases where the RMSE is within the interquartile range. The right column shows large error cases where the RMSE is larger than $1.5\times$ the interquartile range.

building height input, ground truth and Deep Umbra results. In these cases, it is possible to notice that the generated shadow values more closely match the ground truth when pixels are located closer to buildings (Johannesburg, left; Seattle, left). Yet, when moving away from them, many of the nuances and sharpness of the shades are lost (NYC, right; Johannesburg, right). Figure 7 also highlights large error cases where the RMSE is larger than $1.5\times$ the interquartile range. We note, however, that these cases represent less than 5% of the total cases in our test set, underscoring how well our framework is able to model shadow accumulation.

6.4 Comparison with existing tools & techniques

We also compared Deep Umbra with existing tools and techniques across five dimensions: accuracy, performance, support for what-if analyses, support for shadow accumulation, and whether the solution was open or not. In collaboration with urban experts (co-authors of the paper), we selected five already-existing tools and techniques for comparison: (1) Shadow Accrual Maps [4], a GPU-based technique for the interactive accumulation of shadows at city scale and achieves a speedup of $10\times$ when compared to Nvidia's OptiX. (2) JveuxDuSoil [28], a web-based service for the assessment of single timestep shadows. (3) ShadeMap [27], a web-based service that supports the accumulation of shadows. (4) QGIS [30] and ArcGIS Pro [29], two popular GIS tools in urban planning and architecture that also support shadow accumulation. (5) Google Solar [31], an API that allows for the querying of hourly shade information. Table 6 summarizes the comparison.

Regarding time performance, Deep Umbra and Jveux-DuSolei are the only ones that achieve real-time rates, i.e., the computations are performed in less than ∼0.05 second. Lower times are particularly important for two reasons: First, to drive visual analytics systems, as increased latency

reduces users' rate of observations, generalizations and generation of hypotheses [62]. Second, to enable the fast computation of accumulated shadows for larger geographical regions. For example, to compute the shadow ground truth data for the 7,000 tiles in NYC, Shadow Accrual Maps took more than 12 minutes for a single season. Such a situation is even more impractical when considering GIS tools, such as QGIS. Using the Urban Multi-scale Environmental Predictor (UMEP) extension, it took 3 hours to accumulate shadows for a single season. Deep Umbra, on the other hand, was able to compute the accumulated shadows for the entire city within 105 seconds, a speed up of $6\times$ compared to Shadow Accrual Maps, and $102\times$ compared to QGIS.

Support for what-if analyses was also considered in our comparison, i.e., the ability to replace buildings and recompute analyses to assess shadow impact of proposed developments. JveuxDuSoil and Google Solar do not support such analyses. JveuxDuSoil also does not support shadow accumulation. Of all techniques and solutions, Deep Umbra is the only one to support real-time accumulation of shadows for what-if analyses. Considering accuracy, previous tools support the exact computation of shadows (considering user-specified temporal and spatial resolutions). While not being exact, Deep Umbra provides an alternative with fairly low error and significant time performance gains. Deep Umbra, together with QGIS, is the only solution that is open source and freely available for code inspection and modification.

7 GLOBAL SHADOW DATASET

As part of our work, we have also made available a large collection of accumulated shadow tiles for over 100 cities (in 6 continents), computed using Deep Umbra. The complete dataset contains 999,807 map tiles, totaling over 16 GB. The

TABLE 6
Comparison between different shadow assessment tools and techniques. We classify as *interactive* the solutions that compute results in less than ∼0.15 second, and *real-time* the ones with results below ∼0.05 second. *Open source* refers to software that is freely available for code inspection.

Techniques & tools	Performance	What-if support	Accumulation	Accuracy	Open source
Shadow Accrual Maps [4]	Interactive	**Yes**	**Yes**	**Exact**	No
JveuxDuSoleil [28]	**Real-time**	No	No	**Exact**	No
ShadeMap [27]	Interactive	**Yes**	**Yes**	**Exact**	No
QGIS [30]	Non-interactive	**Yes**	**Yes**	**Exact**	**Yes**
ArcGIS Pro [29]	Non-interactive	**Yes**	**Yes**	**Exact**	No
Google Solar [31]	Offline (pre-computed)	No	**Yes**	**Exact**	No
Deep Umbra (this work)	**Real-time**	**Yes**	**Yes**	∼0.06 RMSE	**Yes**

selected cities cover six continents and have different building morphologies and urban street networks. We hope that the data will provide new opportunities for urban experts to perform not only fine-grained analyses (i.e., *what is the best location for a certain facility?*) but also large-scale comparisons and exploratory analyses across neighborhoods or cities (e.g., *what is the park with least shadow amount?* or *what is the city with least shadow amount?*). Our objectives are twofold: (1) facilitate research on the topic of sunlight and shadow quantification by providing a dataset that can be compared to and augmented; and (2) enable research across different domains, particularly urban planning, and public health. We have also made available a web-based interface that can be used to effortlessly browse and visualize the Global Shadow Dataset. The dataset currently contains accumulated shadow information, but the accumulated sunlight can be obtained by simply computing the complement of the accumulated shadow value. The dataset and web viewer can be accessed at urbantk.org/shadows.

8 CASE STUDIES

In this section, we demonstrate the application of our proposed framework through a set of case studies performed in collaboration with urban experts, both co-authors of this paper. First, to highlight the scale and opportunities brought forward by Global Shadow Dataset, Figure 8 visualizes the mean shadow accumulation for parks in 110 cities. As expected, cities near the equator present similar values across seasons. Cities in high latitudes, however, have vastly different values between summer and winter. Specifically, Paris exhibited relatively high shadow values during winter. The first case study then analyzes the impact of building density on winter sunlight access in several parks in the Vaugirard neighborhood in Paris.

We further illustrate the utility of our work by measuring shadow accumulation on street networks, commonly utilized in urban accessibility studies [59], [60], [61]. Figure 9 visualizes the distributions of overshadowed streets (i.e., streets with mean shadow accumulation value above 75%) over three seasons for four cities. We can see that for these cities, the shadow distribution in the winter is more skewed towards higher values than summer/spring, partially due to the sun's lower elevation during the season. Our second case study is then based on an in-depth analysis of the shadow accumulation on streets aggregated by the nearby land use type (e.g., residential, commercial, industrial) in two of the largest cities in the US: Chicago and NYC.

TABLE 7
Percentage of parks that have high sunlight access (mean shadow accumulation value below 25%), moderate sunlight access (between 25% and 50%), partially shadowed (between 50% and 75%), and overshadowed (above 75%).

Season	Mean shadow accumulation value			
	< 25%	[25%,50%]	[50%,75%]	> 75%
Summer	66.6%	30.7%	2.4%	0.3%
Spring	61.7%	25.5%	9.2%	3.4%
Winter	32.1%	23.9%	19.0%	25.0%

8.1 Sunlight access in Parisian parks

Health researchers have revealed that access to sunlight in winter is as essential as shade in summer. A significant percentage of the population suffers vitamin D deficiency at the end of winter and consequently becomes vulnerable to diseases, including diabetes, heart disease, and osteoporosis [63]. With the rapid growth in urbanization, sunlight access during winter within high-density municipal areas has been deemed critical more than ever [26], [64], [65], [66]. However, people living or working in high-density environments, more often than not, have limited access to private green open spaces. And as development intensifies, the overshadowing of existing parks increases, creating tension between urban growth and the preservation of sunlight access [26].

In this case study, we investigated sunlight access in parks of Paris during winter. Table 7 shows the distribution of shadows in parks during the three seasons. Throughout winter, 25% of the parks in Paris are overshadowed, i.e., they have a mean shadow accumulation value above 75%. Also, there is a discrepancy in sunlight access comparing the distributions over the three seasons. The overshadowing of open urban spaces in Paris can largely be attributed to its densely built environment and narrow alleyways [67], [68].

We further looked into six public parks in the district of Vaugirard and divided them into four categories based on their mean shadow accumulation percentile score. The categories are: high sunlight access (below 25%), moderate sunlight access (between 25% and 50%), partially shadowed (between 50% and 75%), and overshadowed (above 75%). Figure 10 visualizes shadow accumulation during winter for these parks using our data. Among them, two parks are partially shadowed, namely, Square du Clos Feuquières (59%), and Square Duranton (70%). Only Jardin Élisabeth Boselli has high sunlight access (22%), whereas
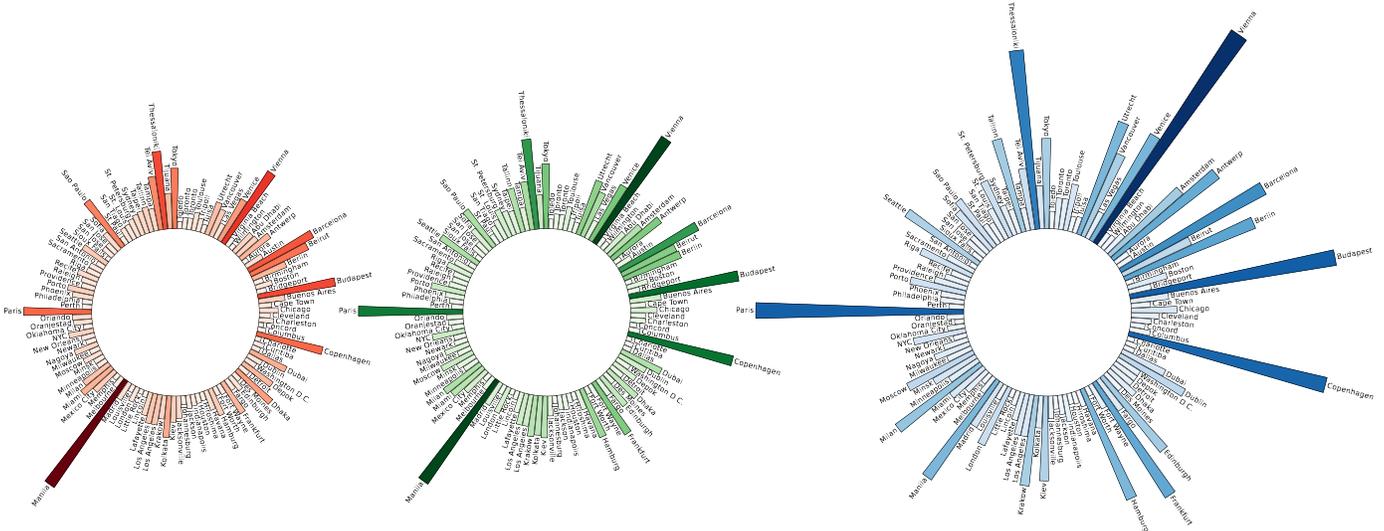
Fig. 8. Comparison of mean shadow accumulation scores across parks in 110 cities during the summer (left), spring (middle), and winter (right).
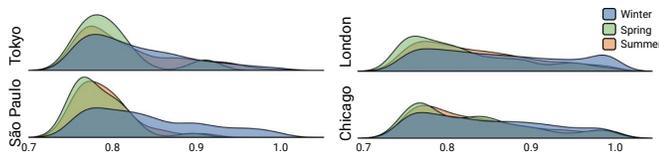

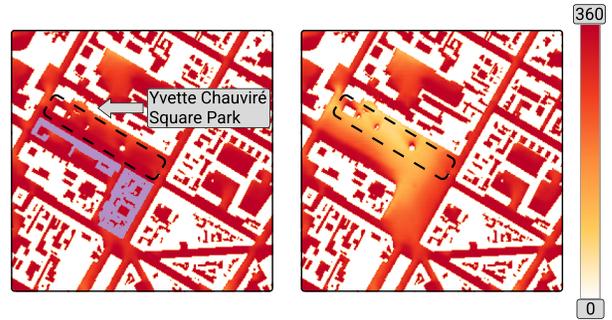Fig. 9. Distribution of shadow accumulation in overshadowed streets.

Fig. 11. Comparison of the impact of nearby buildings on the overshadowing of Yvette Chauviré Square Park (left) with an alternate scenario of without the buildings (right). The buildings considered are highlighted in purple. Removing these buildings too close to the park improves sunlight access by a significant amount.

8.2 Analyzing land use-level street shadows

The right to access sunlight and the impact of tall buildings on restricting this access in dense urban areas have long been the subject of popular debates [18], [48]. Two main contributing factors to shadows on streets are street width and the dimension (i.e., height and bulk) of the surrounding buildings [16], [69]. To investigate how the sunlight access pattern changes with land use in Chicago and NYC, in this case study we first looked at the average amount of accumulated shadow during the winter. Since we are interested in the impact of building morphology on shadow accumulation, we created street-level shadow data by aggregating shadow values at a street-segment level, using data from OSM [46]. Each street segment considers a buffer of 5 meters, and we computed the mean accumulated shadow per segment by averaging pixel values within the buffered region. For each segment, we then assigned a land use type by spatially joining buffered streets with official land use datasets for Chicago [70] and NYC [71]. The results show no significant difference between the average shadow score in different land uses. NYC has, on average, a slightly higher mean accumulated shadow, the maximum being 94% in residential areas compared to 89% for the same land use in Chicago. The distributions were also close, which does not reveal any specific pattern.
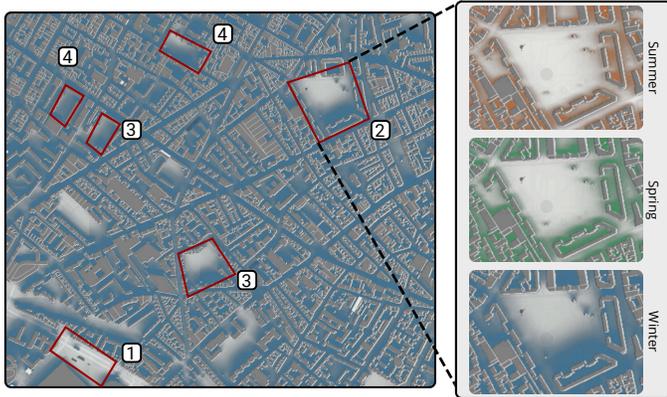

Fig. 10. Visualizing accumulated shadow cast over parks in Paris. Left image highlights six parks divided into four categories of sunlight access during winter: (1) high sunlight access, (2) moderate sunlight access, (3) partial overshadowing and (4) significant overshadowing. Right images highlight shadow accumulation for one park (Square Saint Lambert).

Square Saint-Lambert has moderate sunlight access (50%). Conversely, Jardin Marguerite Boucicaut (78%) and Square Yvette Chauviré (89%) lie in the overshadowed category.

To inspect the impact of nearby buildings on the overshadowing of Square Yvette Chauviré, we performed a what-if scenario analysis by removing buildings too close to the park. Figure 11 shows both the removed buildings (left) and the resulting shadows (right), highlighting Deep Umbra's ability to assist in urban planning activities. Deep Umbra can help policy makers and government agencies in several ways. They can better identify locations for new parks within the municipality, and introduce sunlight protection policies considering new development scenarios.
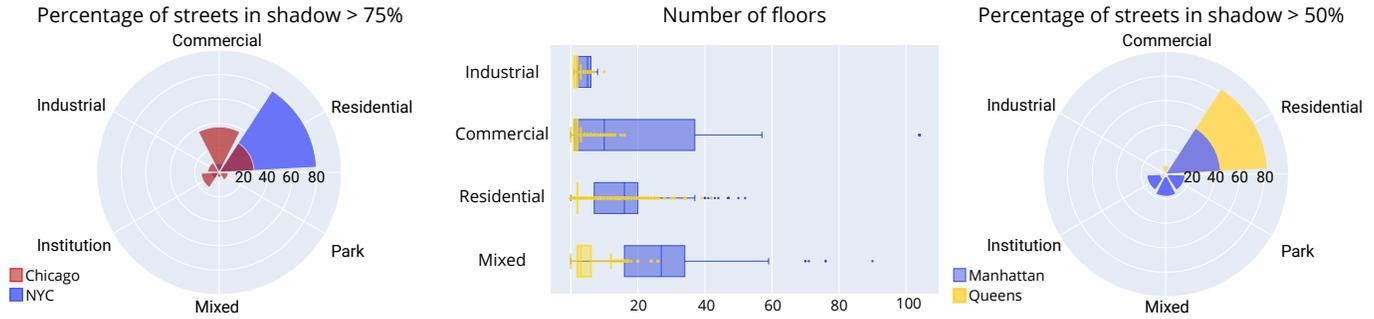
Fig. 12. The percentage of streets in overshadowed areas per land use for Chicago and NYC during the winter (left). The distribution of building heights (middle) and the percentage of streets that are partially or overshadowed in Manhattan and Queens (right).

Next, we computed the percentage of streets falling in the overshadowed range (above 75%) for each land use type in both cities. Here, we see a more drastic difference. Although in NYC, mixed-use areas generally have a higher average of shadow score compared with residential ones, as Figure 12(left) shows, residential streets comprise 78% of the total number of streets that are overshadowed during the winter. However, in Chicago, distribution is more balanced: overshadowed streets are primarily concentrated in commercial areas (39% of the total number of streets), while residential areas comprise 25.5% of total overshadowed streets. Streets around parks and open spaces in Chicago receive less sunlight than those in NYC (7% vs. 2%).

Concerns around equitable sunlight access age back to 1916, when NYC's first zoning resolution was adopted. To provide equitable access to sunlight, the regulations enforce restrictions on the height and bulk of the buildings. In NYC, the Queens borough's zoning regulations restrict it to mainly up to eight-story residential buildings, while in Manhattan tall skyscrapers dominate dense residential and mixed-use areas. But do height restrictions in Queens provide residents with more sunlight access?

In Queens, the mean accumulated shadow value in residential streets is 82%, with a maximum of 96%, while in Manhattan, we observed a lower average of 79% and a lower maximum of 91%. We looked at the percentages of streets falling within our ranges (Figure 12 (right)). The publicly available PLUTO dataset provides information about the height and bulk of buildings across NYC [71]. Using that data, we computed the distribution of building height per land use for each borough to have a more vivid picture and a fair comparison. Figure 12 (middle) shows that buildings in Manhattan have a much wider height range across all relevant land uses (we excluded parks and industrial areas from this analysis since they have very different zoning regulations) and generally are much taller than their counterparts in Queens. However, the radial plot (right) reveals an interesting pattern: although Manhattan buildings are much taller than Queens, 83% of the shadowed streets of Queens are placed in residential areas, while this number drops to 43% in Manhattan during the winter. This pattern, in huge part, is due to the bulky residential buildings and row houses built densely in Queens, which as the analysis shows, cast more shadows than the taller but skinnier buildings – a pattern backed up by previous studies [4].

Using our computed shadow data and the publicly available land use and building datasets, we could uncover an interesting and unexpected pattern in the Queens borough of NYC, which raises questions regarding the equitable access to sunlight and bulk and height regulations already in place. This analysis can be further replicated and extended to inform policy makers and residents in other cities about the built environment's impact on their daily lives.

9 CONCLUSION AND FUTURE WORK

In this paper, we presented Deep Umbra, a novel framework that enables the quantification of shadow and sunlight access in urban environments on a global scale. We leveraged a conditional generative adversarial network to enable the real-time computation of OSM-like map tiles with shadow accumulation. For a city such as NYC, it only takes 105 seconds to generate shadow tiles for the entire city, making it a viable and valuable solution for large-scale urban shadow and sunlight analysis.

We also reported a series of experiments demonstrating the effectiveness of our approach, particularly taking into account cities from different continents. Due to the data training variety, the model can be used for different seasons in any city of the globe (given available building height data). The model was evaluated using root mean square error, mean squared error, and mean absolute error, also considering the transferability to cities outside the training domain. Moreover, we have made available the Global Shadow Dataset, a comprehensive dataset covering more than 100 cities over six continents. By doing this, we hope to facilitate research in multiple domains interested in shadow and sunlight access computation and alleviate data acquisition friction to benefit city planning activities as well as facilitate the development of concrete applications and studies that have shadows in cities as its object of analysis.

From an urban domain perspective, integrating advanced shadow-casting tools, such as Deep Umbra, is fundamental in climate-resilient urban planning. The dynamic interplay of shadows, contingent upon the sun's seasonal path, significantly influences the energy dynamics and thermal comfort within urban environments. During the winter months, these tools empower policymakers to optimize access to sunlight, thereby mitigating energy consumption and enhancing urban comfort. Conversely, precise shadow calculations foster thermal comfort during summer by strategi-

cally shading pedestrian pathways and outdoor spaces. This holistic approach to shadow management is a cornerstone for resilient urban planning and design.

However, there are still some unaddressed challenges and limitations that we plan to tackle in future work. First, even though the amount of building height data available through OSM is extensive, several cities in the developing world lack good coverage. Therefore, we plan to explore alternatives such as utilizing widely available data (such as remote sensing) to enhance currently available height data. Second, Deep Umbra does not consider urban greenery, which impacts urban climate. Deciduous trees, for instance, offer a dual advantage, permitting sunlight penetration in winter while shielding against direct solar radiation in summer. In future work, we plan to extend our framework to also take into account trees. Third, while the findings of this study offer valuable insights, they should be interpreted with prudence when applied to policy formation or decision-making at the city level. Given the complexities and dynamic nature of urban environments, results may not fully encapsulate every variable and scenario. Therefore, we advise that our results and model be used as a guiding framework rather than definitive solutions, encouraging further investigation and adaptation to specific urban contexts. Lastly, we plan to incorporate Deep Umbra into GIS tools, such as QGIS. This will open new opportunities for studies in urban accessibility, targeting the identification of black ice, and urban heat island effect.

REFERENCES

[1] UN Department of Economic and Social Affairs. (2019) World urbanization prospects 2018: Highlights. [Online]. Available: https://population.un.org/wup/Publications/Files/WUP2018-Highlights.pdf

[2] X. Ding, W. Zhong, R. G. Shearmur, X. Zhang, and D. Huisingh, "An inclusive model for assessing the sustainability of cities in developing countries–trinity of cities' sustainability from spatial, logical and time dimensions (tcs-sltd)," *Journal of Cleaner Production*, vol. 109, pp. 62–75, 2015.

[3] J. M. Barr. (2019) Skyscrapers and shadows: The value of sunshine in the city. [Online]. Available: https://buildingtheskyline.org/city-shadows

[4] F. Miranda, H. Doraiswamy, M. Lage, L. Wilson, M. Hsieh, and C. T. Silva, "Shadow accrual maps: Efficient accumulation of city-scale shadows over time," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 3, pp. 1559–1574, 2018.

[5] R. Compagnon, "Solar and daylight availability in the urban fabric," *Energy and buildings*, vol. 36, no. 4, pp. 321–328, 2004.

[6] E. Krüger, F. Minella, and F. Rasia, "Impact of urban geometry on outdoor thermal comfort and air quality from field measurements in Curitiba, Brazil," *Building and Environment*, vol. 46, no. 3, pp. 621–634, 2011.

[7] L. Jiang, Y. Hua, C. Ma, and X. Liu, "Sunchase: Energy-efficient route planning for solar-powered evs," in *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2017, pp. 383–393.

[8] A. Vulkan, I. Kloog, M. Dorman, and E. Erell, "Modeling the potential for pv installation in residential buildings in dense urban areas," *Energy and Buildings*, vol. 169, pp. 97–109, 2018.

[9] R. L. Knowles, "The solar envelope," *Solar L. Rep.*, vol. 2, p. 263, 1980.

[10] F. De Luca and T. Dogan, "A novel solar envelope method based on solar ordinances for urban planning," in *Building Simulation*, vol. 12, no. 5. Springer, 2019, pp. 817–834.

[11] H. H. Alzoubi and A. A. Alshboul, "Low energy architecture and solar rights: Restructuring urban regulations, view from jordan," *Renewable Energy*, vol. 35, no. 2, pp. 333–342, 2010.

[12] F. O. R. Pereira, C. A. N. Silva, and B. Turkienikz, "A methodology for sunlight urban planning: a computer-based solar and sky vault obstruction analysis," *Solar energy*, vol. 70, no. 3, pp. 217–226, 2001.

[13] C. Hachem, P. Fazio, and A. Athienitis, "Solar optimized residential neighborhoods: Evaluation and design methodology," *Solar Energy*, vol. 95, pp. 42–64, 2013.

[14] A. Vartholomaios, "The residential solar block envelope: A method for enabling the development of compact urban blocks with high passive solar potential," *Energy and Buildings*, vol. 99, pp. 303–312, 2015.

[15] J. Kanters and M. Wall, "A planning process map for solar buildings in urban environments," *Renewable and Sustainable Energy Reviews*, vol. 57, pp. 173–185, 2016.

[16] E. Andreou, "The effect of urban layout, street geometry and orientation on shading conditions in urban canyons in the mediterranean," *Renewable Energy*, vol. 63, pp. 587–596, 2014.

[17] J.-Y. Deng and N. H. Wong, "Impact of urban canyon geometries on outdoor thermal comfort in central business districts," *Sustainable Cities and Society*, vol. 53, p. 101966, 2020.

[18] K. Q. Seelye, "Building boom in Boston casts shadows on history and public space," *The New York Times*, vol. 11, 2017.

[19] F. Costamagna, R. Lind, and O. Stjernström, "Livability of urban public spaces in northern Swedish cities: The case of Umeå," *Planning Practice & Research*, vol. 34, no. 2, pp. 131–148, 2019.

[20] R. Zhu, L. You, P. Santi, M. S. Wong, and C. Ratti, "Solar accessibility in developing cities: A case study in kowloon east, hong kong," *Sustainable Cities and Society*, vol. 51, p. 101738, 2019.

[21] E. Eisemann, M. Schwarz, U. Assarsson, and M. Wimmer, *Real-time shadows*. CRC Press, 2011.

[22] J.-H. Nah and D. Manocha, "Sato: Surface area traversal order for shadow ray tracing," in *Computer Graphics Forum*, vol. 33, no. 6. Wiley Online Library, 2014, pp. 167–177.

[23] S. Zhang, R. Liang, and M. Wang, "Shadowgan: Shadow synthesis for virtual objects with conditional adversarial networks," *Computational Visual Media*, vol. 5, no. 1, pp. 105–115, 2019.

[24] D. Liu, C. Long, H. Zhang, H. Yu, X. Dong, and C. Xiao, "Arshadowgan: Shadow generative adversarial network for augmented reality in single light scenes," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 8139–8148.

[25] J. K. Tan, R. N. Belcher, H. T. Tan, S. Menz, and T. Schroepfer, "The urban heat island mitigation potential of vegetation depends on local surface type and shade," *Urban Forestry & Urban Greening*, vol. 62, p. 127128, 2021.

[26] Hodyl and Co. (2018) Central Melbourne sunlight report. [Online]. Available: https://www.hodyl.co/projects/sunlight-to-public-open-space

[27] Shademap: the definitive source of shade data for the planet. [Online]. Available: https://shademap.app/

[28] JveuxDuSoleil: a web-app to simulate the shadow of cities and detect sunny terraces. [Online]. Available: https://jveuxdusoleil.fr

[29] ArcGIS Pro. [Online]. Available: https://www.esri.com/en-us/arcgis/products/arcgis-pro/overview

[30] QGIS: a free and open source geographic information system. [Online]. Available: https://qgis.org/en/site/

[31] Google Solar API. [Online]. Available: https://developers.google.com/maps/documentation/solar

[32] M. Angelidou and A. Psaltoglou, "Social innovation, games and urban planning: An analysis of current approaches," *International Journal of Electronic Governance*, vol. 11, no. 1, pp. 5–22, 2019.

[33] I. Kavouras, E. Sardis, E. Protopapadakis, I. Rallis, A. Doulamis, and N. Doulamis, "A low-cost gamified urban planning methodology enhanced with co-creation and participatory approaches," *Sustainability*, vol. 15, no. 3, p. 2297, 2023.

[34] G. Moreira, M. Hosseini, M. N. A. Nipu, M. Lage, N. Ferreira, and F. Miranda, "The Urban Toolkit: A grammar-based framework for urban visual analytics," *IEEE Transactions on Visualization and Computer Graphics*, 2024.

[35] W. Yap, R. Stouffs, and F. Biljecki, "Urbanity: Automated modeling and analysis of multidimensional networks in cities," *npj Urban Sustainability*, vol. 3, no. 1, p. 45, 2023.

[36] W. Yap, P. Janssen, and F. Biljecki, "Free and open source urbanism: Software for urban planning practice," *Computers, Environment and Urban Systems*, vol. 96, p. 101825, 2022.

[37] Y. Zhang, Y. Li, X. Zhou, X. Kong, and J. Luo, "Curb-gan: Conditional urban traffic estimation through spatio-temporal generative adversarial networks," in *Proceedings of the 26th ACM SIGKDD*

International Conference on Knowledge Discovery & Data Mining, 2020, p. 842–852.

[38] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2017, pp. 1125–1134.

[39] Y. Zheng, W. Wu, Y. Chen, H. Qu, and L. M. Ni, "Visual analytics in urban computing: An overview," *IEEE Transactions on Big Data*, vol. 2, no. 3, pp. 276–296, 2016.

[40] D. Wang, Y. Fu, P. Wang, B. Huang, and C.-T. Lu, "Reimagining city configuration: Automated urban planning via adversarial learning," in *Proceedings of the 28th International Conference on Advances in Geographic Information Systems*, 2020, p. 497–506.

[41] S. Mokhtar, A. Sojka, and C. C. Davila, "Conditional generative adversarial networks for pedestrian wind flow approximation," in *Proceedings of the 11th Annual Symposium on Simulation for Architecture and Urban Design*, 2020, pp. 1–8.

[42] Y. Zhang, Y. Li, X. Zhou, X. Kong, and J. Luo, "TrafficGAN: Off-deployment traffic estimation with traffic generative adversarial networks," in *2019 IEEE International Conference on Data Mining (ICDM)*, 2019, pp. 1474–1479.

[43] X. Zhang, Y. Li, X. Zhou, and J. Luo, "cGAIL: Conditional generative adversarial imitation learning—an application in taxi drivers' strategy learning," *IEEE Transactions on Big Data*, vol. 8, no. 5, pp. 1288–1300, 2020.

[44] Y. Yuan, Y. Zhang, B. Wang, Y. Peng, Y. Hu, and B. Yin, "STGAN: Spatio-temporal generative adversarial network for traffic data imputation," *IEEE Transactions on Big Data*, vol. 9, no. 1, pp. 200–211, 2022.

[45] A. N. Wu, R. Stouffs, and F. Biljecki, "Generative adversarial networks in the built environment: A comprehensive review of the application of gans across data types and scales," *Building and Environment*, p. 109477, 2022.

[46] OpenStreetMap. [Online]. Available: https://www.openstreetmap.org

[47] Central Park Sunshine Task Force Committee. (2015) Central Park sunshine taskforce report. [Online]. Available: https://www.cb5.org/cb5m/resolutions/2015-may/may-2015_10

[48] Q. Bui and J. White. (2016) Mapping the shadows of New York City: Every building every block. [Online]. Available: https://www.nytimes.com/interactive/2016/12/21/upshot/Mapping-the-Shadows-of-New-York-City.html

[49] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014.

[50] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, vol. 25, 2012.

[51] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros, "Context encoders: Feature learning by inpainting," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2536–2544.

[52] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.

[53] S. Paul, B. Jhamb, D. Mishra, and M. S. Kumar, "Edge loss functions for deep-learning depth-map," *Machine Learning with Applications*, vol. 7, p. 100218, 2022.

[54] L. Wang, L. Wang, and S. Chen, "ESA-CycleGAN: Edge feature and self-attention based cycle-consistent generative adversarial network for style transfer," *IET Image Processing*, vol. 16, no. 1, pp. 176–190, 2022.

[55] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro, "High-resolution image synthesis and semantic manipulation with conditional GANs," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 8798–8807.

[56] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical Image Computing and Computer-assisted Intervention*. Springer, 2015, pp. 234–241.

[57] S. Jian, H. Kaiming, R. Shaoqing, and Z. Xiangyu, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision & Pattern Recognition*, 2016, pp. 770–778.

[58] J. Hu, W. Yu, and Z. Yu, "Image-to-image translation with conditional-GAN," *CS230: Deep Learning, Spring*, 2018.

[59] M. Hosseini, M. Saugstad, F. Miranda, A. Sevtsuk, C. Silva, and J. Froehlich, "Towards global-scale crowd+AI techniques to map and assess sidewalks for people with disabilities," in *Proceedings of the CVPR 2022 AVA (Accessibility, Vision, and Autonomy Meet) Workshop*, 2022.

[60] J. E. Froehlich, Y. Eisenberg, M. Hosseini, F. Miranda *et al.*, "The future of urban accessibility for people with disabilities: Data collection, analytics, policy, and tools," in *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility*, 2022.

[61] M. Hosseini, A. Sevtsuk, F. Miranda, R. M. Cesar Jr, and C. T. Silva, "Mapping the walk: A scalable computer vision approach for generating sidewalk network datasets from aerial imagery," *Computers, Environment and Urban Systems*, vol. 101, p. 101950, 2023.

[62] Z. Liu and J. Heer, "The effects of interactive latency on exploratory visual analysis," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 12, pp. 2122–2131, 2014.

[63] M. F. Holick, "Sunlight and vitamin D for bone health and prevention of autoimmune diseases, cancers, and cardiovascular disease," *The American journal of clinical nutrition*, vol. 80, no. 6, pp. 1678S–1688S, 2004.

[64] The Parks Council, NYC. (1991) Preserving sunlight in New York City's parks: A zoning proposal. [Online]. Available: https://bethletain.squarespace.com/s/Preserving-Sunshine-Zoning.pdf

[65] M. Mondon. (2015) NYC councilman wants to preserve Central Park sunshine. [Online]. Available: https://nextcity.org/urbanist-news/central-park-sunshine-high-rise-shadows-new-committee

[66] N. Filatoff. (2021) Let there be sunlight ... in the public parks of high-density communities! [Online]. Available: https://www.pv-magazine-australia.com/2021/09/20/let-there-be-sunlight-in-the-public-parks-of-high-density-communities

[67] E. Sciolino. (2003) Paris journal; call it the city of darkness, and give it vitamin D. [Online]. Available: https://www.nytimes.com/2003/01/06/world/paris-journal-call-it-the-city-of-darkness-and-give-it-vitamin-d.html

[68] A. Small. (2016) Why are European cities so dense? [Online]. Available: https://www.bloomberg.com/news/articles/2016-10-27/why-european-cities-still-have-more-dense-development

[69] R. Zhu, M. S. Wong, L. You, P. Santi, J. Nichol, H. C. Ho, L. Lu, and C. Ratti, "The effect of urban morphology on the solar capacity of three-dimensional cities," *Renewable Energy*, vol. 153, pp. 1111–1126, 2020.

[70] The Chicago Metropolitan Agency for Planning (CMAP). Land Use Data. [Online]. Available: https://www.cmap.illinois.gov/data/land-use

[71] New York City Department of City Planning (NYC DCP). PLUTO and MapPLUTO. [Online]. Available: https://www1.nyc.gov/site/planning/data-maps/open-data/dwn-pluto-mappluto.page

Kazi Shahrukh Omar is a Computer Science PhD student at UIC, interested in urban computing and visualization.

Gustavo Moreira is a Computer Science PhD student at UIC, interested in visualization and visual analytics.

Daniel Hodczak is an undergraduate student at UIC, interested in designing frameworks for efficient urban research.

Maryam Hosseini is a Postdoctoral Associate at the City Form Lab at MIT. She received the Ph.D. in Urban Systems from Rutgers University. She works in the intersection between visualization, computer vision and urban science.

Nicola Colaninno is an Assistant Professor at the Polytechnic University of Milan. He holds a Ph.D. in Urban and Architectural Mgmt. and Valuations from UPC. His research focuses on urban planning and urban climate analysis.

Marcos Lage is an Associate Professor in the Dept. of Computer Science at UFF. He received the Ph.D. in Applied Mathematics from PUC-Rio. His research interests include visual computing and topological data structures.

Fabio Miranda is an Assistant Professor in the Dept. of Computer Science at UIC. He received the Ph.D. in Computer Science from NYU. His research focuses on large-scale data analysis, data structures, and urban data visualization.