• **Article** •

Opening Design using Bayesian Optimization

Nick Vitsas[1*] , Iordanis Evangelou[1] , Georgios Papaioannou[1] , Anastasios Gkaravelis[1]

1. *Department of Informatics, Athens University of Economics and Business, Greece*

*** Corresponding author:** vitsas@aueb.gr

Received: - **Accepted:** -

Supported by the Hellenic Foundation for Research and Innovation (Project No. 7310).

Abstract Background Opening design is a major consideration in architectural buildings during early structural layout specification. Decisions regarding the geometric characteristics of windows, skylights, hatches, etc., significantly affect the overall energy efficiency, airflow, and appearance of a building both internally and externally. **Methods** In this work, we employ a goal-based, illumination-driven approach to opening design using a Bayesian optimization approach based on Gaussian processes. A method that enables designers to conveniently set lighting intentions in conjunction with the qualitative and quantitative characteristics of the desired openings is proposed. **Results** The parameters are optimized within a cost-minimization framework to calculate geometrically feasible, architecturally admissible, and aesthetically pleasing openings of any desired shape while taking into account the designer's lighting constraints.

Keywords Inverse Geometry; Lighting Simulation; Scene Synthesis

1 Introduction

Accurate light-transport simulation has been aided by significant improvements in the convergence speed in recent years. This is largely attributed to the more advanced statistical light-transport solvers, denoising algorithms, and inclusion of hardware-accelerated ray tracing in commodity GPUs. This has unlocked a substantial potential for its utilization in demanding applications of light simulation (such as architectural lighting) in terms of accuracy and consistency. In this work, we leverage photorealistic simulation via interactive path tracing to propose an automatic solution to the opening design problem that maps effectively to the luminance distribution characteristics of natural light scattering.

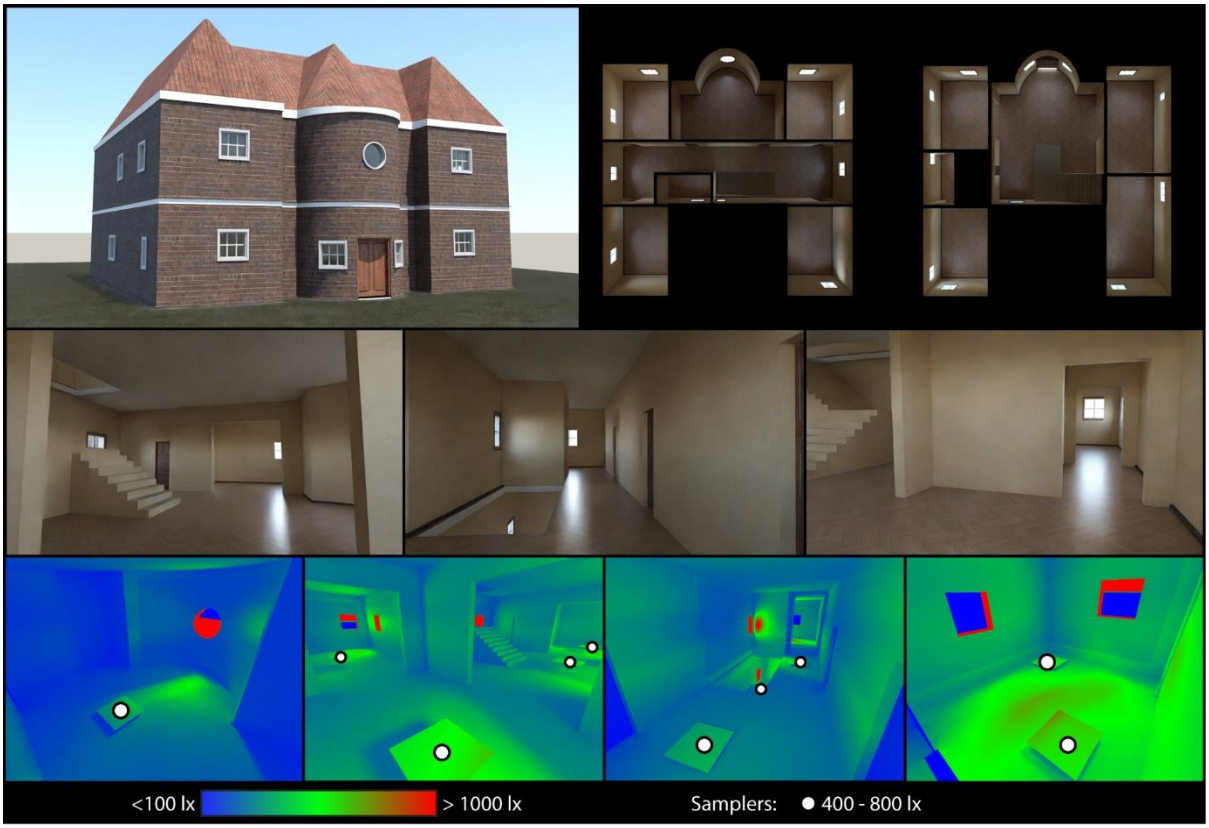

Figure 1 Openings computed on a two-story house using our method. In this example, 27 openings were optimized using 17 planar samplers for a total of 44 optimization parameters. Sky lighting is set for office hours for the summer–fall seasons in Capetown.

Opening design is an important subspace in architectural building design. An opening is a geometric formation that connects the interior of a building to its exterior. In our context, the problem of designing such openings is approached as an *inverse geometry problem*, i.e., a goal-driven definition of the design parameters (commonly referred to as the *inverse opening problem*).

Modern opening designs place human factors at the center of the design process. This is because energy and efficiency considerations do not ensure visual comfort or pleasing aesthetic results. For example, in workspaces where employee comfort, reliability, and safety are pivotal in day-to-day activities, an inconsiderate window design can cause discomfort, eye strain, inadequate privacy, and psychological stress. Therefore, occupant requirements are important aspects of the opening design process. For years, designers have used empirical and simulation methods to assess the performance of proposed configurations. Previous studies [1, 2, 3] have highlighted the importance of the window position and size to the overall efficiency of a building's design both in terms of energy consumption and residential comfort.

With regard to daylight availability, one of the most commonly used methods is the *daylight coefficient* (DC) approach. It was originally proposed by Tregenza and Waters [1]. Other metrics include the *daylight factor*, *daylight autonomy* measure, and *useful daylight illuminance* (UDI) [4].

Glare is also an adverse effect that should be minimized. Quantifying the risk and levels of glare has been studied extensively [5, 6]. A standard metric for glare under artificial lighting was published by the

International Commission on Illumination CIE under the name "unified glare rating" (UGR). Fisekis et al. [5] approached discomfort glare specifically for windows and acknowledged the limitation of excluding natural light sources from the UGR. Wienold and Christoffersen [6] proposed the "daylight glare probability" (DGP) for the specific case of daylighting.

In this work, we primarily focus on the design of openings as natural light sources and their potential to maximize natural light usage and minimize the requirements for artificial light sources during the daytime (**Figure *1***). We describe and develop a method as well as the respective working system that combines robust photorealistic rendering, a comprehensive set of tools for describing the designer's illumination requirements and aesthetic preferences, and an optimization framework adapted to the opening design problem.

The contributions of this work are as follows: a) a problem formulation based on Bayesian optimization, which acknowledges the designer's workflow and allows for a convenient expression of each part of their design process; b) accurate physically-based lighting evaluation during optimization by utilizing maximum luminance environment maps for glare prevention; and c) support for arbitrary opening designs using daylighting systems of any shape, size, and orientation.

The remainder of this paper is organized as follows. In Section 2, we cover related work on opening design and the background of sky modeling and Bayesian optimization. Section 3 describes the necessary elements of the goal-driven opening design process. Section 4 presents an in-depth analysis of the proposed method followed by an experimental evaluation in Section 5. Finally, we conclude with a method recap and future directions in Section 6.

2 Background and Related Work

2.1 Inverse Opening Design

An inverse problem can be described as the process of determining the cause of a phenomenon by starting from its effects. In this context, several publications have contributed to the field of inverse design with direct implications for many aspects of 3D scene synthesis. This section covers major contributions to the subject of light-driven design, with an emphasis on aspects pertaining to the *inverse opening problem*.

Mahdavi and Berberiodu-Kallivoka [7] investigated the potential of integrating computational light simulations with computer-aided daylighting modeling during building design. They designed and proposed a tool in which the variations in a performance indicator (e.g., daylight factor, average illuminance levels, and uniformity factors) can be translated directly to modifications in a set of design-related variables (e.g., geometry configuration and material properties). The authors adopted an inverse approach for design exploration. The variations in the design space are tested and adopted only if the subsequent effect on the performance indicator is desirable.

Tourre et al. [8] were among the first to propose an approach that used an accurate light simulation for opening design, albeit with a highly simplistic setup. The potential opening areas are meshed in so-called "opening elements." Each opening element is considered as an intermediate and anisotropic light source. Its lighting contribution to the interior of the building is computed as an image through a pinhole camera simulation.

Fernández and Besuievsky [9] defined a constrained problem for roof skylights and artificial Lambertian light emitters. They considered two categories of light constraints: a) geometric restrictions and b) lighting intentions. The first category considers the imposed constraints that the light source needs to achieve (such as the light size, aspect ratio, source spacing, or symmetries). For example, the building construction constraints regarding the installation of skylight sources on a roof may impose a regular distance between skylights and their alignment along a given axis. For lighting intentions, they considered the valid intervals of light intensity that each user-specified surface is allowed to reflect. The artificial and natural light sources are modeled as Lambertian emitters. Similarly, all the materials are required to be Lambertian to facilitate a GI solution using a modified Radiosity method.

Subsequently, Fernández et al. [10, 11] focused on the opening design problem and extended it to support anisotropic light incidence. Each potential opening was divided into several smaller patches, and each contribution was approximated using a pinhole camera. They again used the Radiosity method to calculate the contribution of the patch from a static sky or environment map and used it to determine the best opening shape. The resulting contribution of the opening was the sum of all the elements. The support for arbitrary lighting incidence rendered this approach a good candidate in an urban context [12]. Here, the Radiosity method was used to compute daylighting that includes inter-reflections with the exterior components of a building. The inputs to the process are a sky model, hourly daylight data, an urban model with a selected objective building, and an interior room model with predefined locations for openings as discretized surface patches. The Radiosity system was solved to selectively open the optimal patches according to various daylighting metrics.

One of the first methods using a robust light transport solver for the inverse opening design problem was presented by Kalampokis et al. [13]. They described a generic, physically-based method that determines the number, location, and shape of openings given the designer's lighting intentions, the candidate surfaces for hosting openings, and a description of the scene. Similar to previous work, they used an arbitrarily fine quantization of the opening domain into candidate opening *elements*. The opening domain is virtually "diced" using constructive solid geometry (CSG) subtraction based on ray tracing during the estimation of the contribution of each element. Contrary to previous work, these openings have actual physical characteristics that are respected by the global illumination solution used. The illumination was measured at user-defined sampling locations. Here, the illumination goals were provided in the form of irradiance levels. The contribution of each opening to each sampling point was stored in a matrix. The cumulative lighting contributions of multiple light paths crossing the opening domain were considered to be additive. This enabled the formulation of the problem as a least mean squares optimization problem using non-negative

values. It was then rounded to a binary solution vector containing "on" and "off" elements. The resulting binary configuration of the elements was used as the initial state of the a genetic algorithm. This phase stochastically mutated the solution, thereby favoring error minimization and shape coherence. This effectively incorporated structural constraints in the formed opening patterns, such as the compactness, number, and boundary shape of the openings.

2.2 Daylight Simulation

In many environments, natural light from the sky dome can be the dominant factor in a scene's illumination. Specifically, in the context of building openings, effective utilization of daylight is the most important factor for a building's energy profile and the comfort of its inhabitants. The importance of realistic and controllable sky illumination in daylighting simulations and rendering has resulted in several analytical models of the sky luminance distribution. Considering the high cost and difficulty of simulating atmospheric phenomena, these are parametric sky models represented by simple analytical expressions that can efficiently generate plausible natural radiance maps of the sky. Analytical sky models have been actively researched for several years.

A few have been formally adopted by the International Commission on Illumination (CIE) [14]. Many well-known and widely used models in computer graphics are based on variations of the Perez formula [15]. These include the widely used models by Preetham et al. [16] and the Hosek and Wilkie [17] clear-sky models. These provide analytic formulas for the upper hemisphere.

2.3 Bayesian Optimization

Significant challenges are encountered during the optimization stage in the realm of non-convex function minimization. In particular, when an evaluation function is computationally expensive, the acquisition of informative partial derivatives can be highly intricate or non-existent.

The Bayesian optimization theory [18] provides an elegant, computationally efficient, and data-efficient solution to these issues. It has achieved significant success in the context of hyperparameter optimization [19]. This technique leverages the probabilistic modeling of uncertainty directly in the function space to efficiently explore (both globally and locally) the search space and identify feasible solutions to unknown objectives.

More formally, given a multidimensional variable space X and an arbitrary function $f: X \rightarrow \mathbb{R}$ (which cannot be expressed analytically, is non-convex, and is computationally intensive to evaluate), we wish to determine a configuration $x^{\star} \in X$ such that

$$x^{*} = arg_{x \in X} \min f(x), \quad (1)$$

To solve this, a Bayesian optimizer employs an iterative scheme that can be summarized in two steps. First, it receives a finite precomputed subset $D = \{(x_i, f(x_i))\}_{i=1}^{n}$ of sampled configurations in conjunction with their associated values computed from the actual function and constructs a probabilistic model that expresses a posterior distribution over the function space of the given dataset $P(f \mid D)$. In practice, to maintain a computationally tractable expression, this process is typically modeled by an explicit joint distribution with typical examples being the *Gaussian process* [20] and *student-t process* [21] using learnable kernel parameters for the covariance matrix.

In the second phase, the optimizer iteratively refines the internal surrogate function encoded by the posterior using a risk minimizer known as the *acquisition* function. Specifically, we propose the next sampling point based on

$$x_{n+1} = arg_x min\, a\,(x; D_n), \quad (2)$$

where the function $a : X \longrightarrow \mathbb{R}$ leverages the uncertainty imbued in the posterior distribution to guide the exploration of the search space towards promising regions, while also exploiting already known ones.

Given a new proposal x_{n+1}, the function $y_{n+1} = f(x_{n+1})$ is evaluated, and the result is backpropagated to augment the dataset $D_{n+1} = \{D_n, (x_{n+1}, y_{n+1})\}$. Thereby, the statistics of the model are updated. It can be proved that by transforming the problem into a process of identifying a sequence of proposals with an appropriate definition of an acquisition function, the optimizer converges to a function minimum. Nevertheless, there are various types of acquisition functions [19]. These have different ways of balancing exploration and exploitation, depending on the task at hand.

The notable contributions using this optimization scheme span diverse fields. These range from a purely machine learning context such as deep learning [22], to direct applications in robotics [23] and reinforcement learning [24], as well as computer graphics-related domains such as visual design appearance and interpolation [25] and procedural animation [24]. We refer interested readers to a recent survey [19] that contains a comprehensive analysis of this framework and its immediate application.

In this work, we employed Bayesian optimization as a derivative-free optimizer to minimize a computationally expensive and non-analytic function related to the design of architectural openings. Modeling the uncertainty over the entire function domain is a flexible tool for deriving guiding decisions via acquisition functions that balance both exploitation of local space (similar to the case of Markov chain Monte Carlo algorithms) and attempts to undertake global exploration steps based on samples extracted from a data-learned posterior distribution that would result in more informative regions of the hidden function.

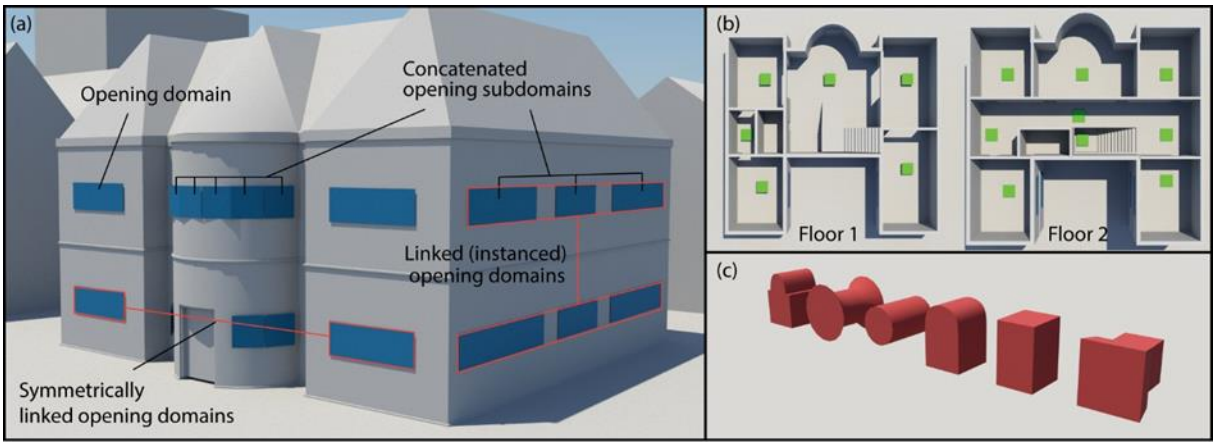

Figure 2 Opening design elements of our method. a) Static geometry and opening domains. b) Illumination goals (planar samplers shown in green). c) Opening template examples.

3 Opening Design Specification

3.1 Design Elements

To develop a process that maps intuitively to a designer's pipeline and requires simple tool interactions, we modeled the goal-driven design process with four classes of elements (***Figure 2***). These are used to describe the geometry, design constraints, and illumination goals.

Static geometry. This is the geometry of the main architectural model in which the openings should be placed. This also includes the surrounding environmental elements that contribute to the proper illumination evaluation (***Figure 2a***). The geometry can be modeled in any application and imported into the opening design system. Any part of the static geometry can host openings (see Opening Domains). The only practical limitation is that it behaves correctly under constructive solid geometry (CSG) operations.

Opening domains. These establish user-provided admissible regions for the opening positions. They are represented as arbitrary rectangles in space that can be disjointed or concatenated to form larger and more complex regions. The static geometry is referenced by this component to apply the opening configuration. ***Figure 2a*** depicts the definition and use of various standalone or grouped regions to establish potential opening domains on the walls of a two-story building.

Opening templates. These meshes are used to cut out the openings on the static geometry. The opening templates are oriented according to the opening domains and are situated within their combined boundaries. In sharp contrast to prior art [13], these can be of an arbitrary shape as long as these are watertight, in order to perform valid CSG operations on the target geometry. Such opening templates are shown in ***Figure 2c***. This enables the design of intricate openings and daylighting systems.

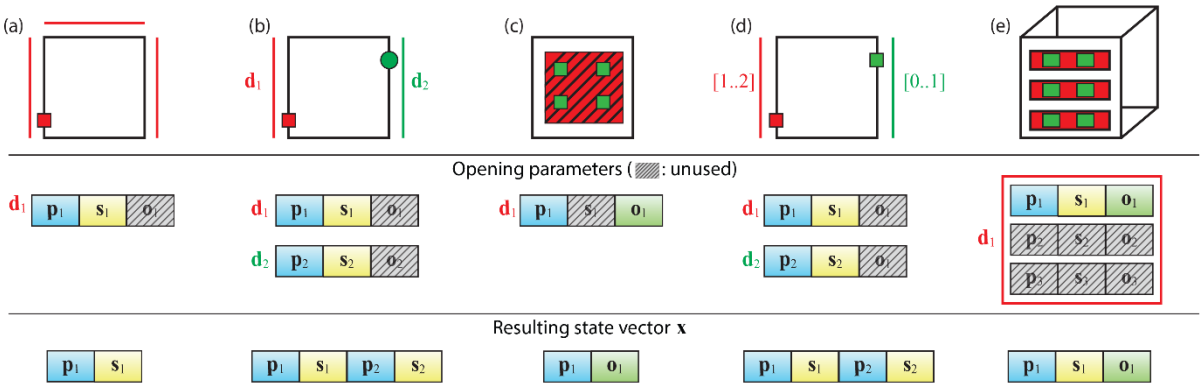

Figure 3 Design space formulation examples. Opening parameters are added selectively to the state vector for optimization. a) A template applied on a group of connected domains. b) Two templates of different types on separate domains. c) A template with adjustable position and spacing. d) Two templates applied on constrained domains with different number of potential openings. e) Three instanced domains sharing the (p_1, s_1, o_1) parameters.

The position and orientation of opening templates are constrained to follow the shape and extents of the opening domains. Meanwhile, their dimensions are optimized independently. If the designer requires, additional constraints with regard to spacing, uniformity, or cross-domain similarity or symmetry can be set by the user for artistic control. To support a wide range of potential opening formations, the final shape of the template is governed by three sets of parameters: a 2D translation vector p_t on the parametric space of the opening domain, 2D scale vector s_t (width, height), and 2D spacing offset vector o_t. The third parameter is used to generate additional openings of an identical scale s_t at positions $\{p_t + io_t\}_{i=1}^{n}$. Here, n is calculated such that the number of the offset openings is within the bounds of the domain. o_t can accommodate negative values. The sign effectively determines the offset direction. Not all parameters and dimensions need to be under optimization depending on the desired degrees of freedom (***Figure 3***).

Illumination goals. The desired illumination levels are set using two types of *samplers*: *planar* and *view*. Planar samplers (***Figure 2b***) are used to set illuminance goals on rectangular patches in space (typically on surfaces). Meanwhile, view samplers are used primarily to set glare prevention goals. Setting lighting intentions in the form of target illumination ranges is a standard practice for lighting-design professionals. Planar samplers require users to define a range of reasonable irradiance values (E_{min}, E_{max}).

Similar to Vitsas et al. [26], a view sampler is a component that measures the maximum average luminance (in nits) over a discretized frustum. It is implemented as a regular perspective camera. View samplers typically require a maximum permitted radiance level L_{max} (minimum of zero). However, for specialized applications (e.g., those that have specific uses of reflectors, etc.), a nonzero lower bound L_{min} can be specified.

Aesthetic and functional constraints. In many designs, the building symmetry, floor-to-floor opening alignment, and spacing constraints should be respected out of purely aesthetic considerations.

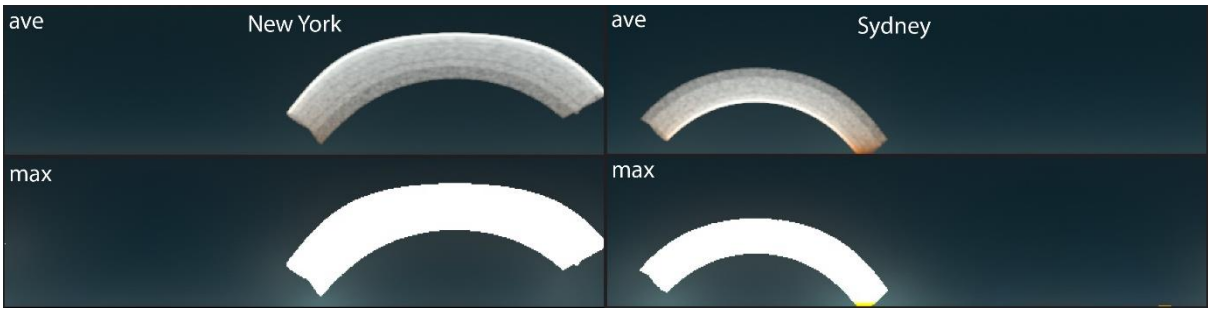

Figure 4 Examples of average and maximum radiance maps of a working day (9:00 a.m–18:00 p.m.) from January to June in the southern hemisphere (Sydney) and northern hemisphere (New York). The luminance is suppressed for clarity.

To accommodate similar architectural decisions while simultaneously simplifying the optimization process by reducing the target parameters, we allow for a wide range of constraint specifications.

We begin with *the grouped domains* in ***Figure 3***a. These comprise a single concatenated domain that is sampled as a continuous parameter space. This practically accommodates "jumps" between disjoint regions. ***Figure 3***b shows the capability to use multiple disjoint domains and groups with separate parameters for optimization. ***Figure 3***c illustrates an example configuration optimized for both cutter position (upper left green square) and spacing offset vector. ***Figure 3***d illustrates how the parameter ranges of different groups (here, the number of openings) can be optimized jointly. Finally, in ***Figure 3***e, we employ parameter space *instancing* of the locally parameterized opening templates. The mechanisms allow for versatile opening configuration constraints and a rapid exploration of the design space. Here, the symmetry, alignment, and uniformity are the most important aspects. These also help substantially minimize the parameters under optimization, thereby facilitating faster convergence.

3.2 Sky modeling

Because our main objective is to optimize a design based on natural light flow, it is important to consider the geospatial location as well as the temporal domain. Lighting evaluation may become meaningless unless it is spread over daytime intervals relevant to the space occupancy and function. These should also be examined over the course of one year and for the specific geospatial location and orientation of a building. The most convenient method to model the contribution of the sky is to use captured sky images. Such captures are intrinsically realistic, may include cloud coverage and building surroundings, and are typically encoded on an HDR environment map. However, this procedure is impractical and potentially unreliable in advanced simulation scenarios. We used the clear-sky dome [27] and solar radiance [17] model proposed by Wilkie and Hosek. The user selects the time interval in months, days, or hours. The module generates a map (parameterized in hemispherical coordinates) of the average and maximum radiance across the sky (*Error! Reference source not found.*). Customized environment maps from external artificial light sources can also be imported. These environment maps can be computed once and reused.

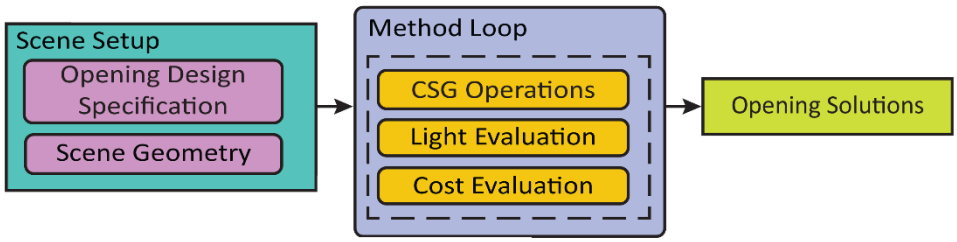

Figure 5 Overview of our opening design method.

4 Method Description

4.1 Method Overview

The inputs for the proposed method are the elements of an architectural design described in Section 3. The method generates a number of valid opening recommendations. The resulting opening design system consists of several stages. These are briefly highlighted in (***Figure** 5*).

As part of the design cycle, the desired illumination goals are specified inside the three-dimensional scene. The opening domains are provided by the user in the form of parametric surfaces at the architecturally admissible parts of the model. Next, several applicable opening templates are selected for each opening domain in conjunction with the constraints on the opening properties that can impose certain aesthetic or structural restrictions (Section 3.1). The independent opening template parameters are concatenated into a single parameter vector x and optimized using a custom approach based on Bayesian optimization (Section 4.4).

For each parameter vector instance, opening templates are applied to the scene geometry as cutters using constructive solid geometry subtraction. Moreover, the illumination is evaluated based on the predetermined lighting conditions (Section 4.2). Our method is orthogonal to the selection of the underlying illumination computation algorithm. However, path tracing was selected to better estimate the complex light transport phenomena.

4.2 Illumination Evaluation

To evaluate the appropriateness of a given opening configuration x_i, after all the CSG operations were applied to the geometry, photometric statistics are collected from the samplers distributed manually across the scene by the designer (which function as virtual sensors). Both types of samplers utilize photorealistic simulation via path tracing to capture light responses from the environment.

As mentioned in Section 3.2, the sky dome and sun disk illumination contributions are captured in environment maps over a sampled timeframe of daylighting intervals across the year for a particular location. One environment map expresses the average value, and the other registers the maximum luminance at each azimuth and elevation (*Error! Reference source not found.*). The maps are prepared as a preprocessing step after adjusting all the daylighting parameters. We use the average sky luminance while computing the surface illuminance for planar samplers. This is because similar to the bibliography, we are

interested in a statistical measure for lighting-goal satisfaction. However, for glare (view samplers), we need to monitor the peak luminance (maximum luminance environment map). This was because average values would be suppressed significantly by the averaging, thereby eliminating problematic daylighting events.

4.3 Cost Function

This step entails the transformation of the evaluated lighting at the samplers into a meaningful normalized error given a candidate opening configuration x. Because an analytical form of the objective function is not readily available, we should quantify the recommended opening configuration for optimization by aggregating the response of each sampler and measuring how effectively the proposed configuration aligns with the desired design goal. We define the cost function as

$$C_{illum}(S,\, x) \;=\; \sum_{k=1}^{|S|} \phi\big(L(S_k, x), s_k^{min},\, s_k^{max}\big),\ (3)$$

where x is the evaluated domain input sample (opening configuration) and L is the measurement of the sampler S_k (illuminance or peak luminance). s_k^{min} and s_k^{max} represent the illumination goal limits (expressed in appropriate units) depending on the type of the sampler S_k. Finally, $\phi(\cdot)$ maps the result to a normalized value with respect to the limits using the *sigmoid* function:

$$\phi(x, y_0, y_1) = \begin{cases} 2\sigma\big(-k_{lb}(x - y_0)\big) - 1 & x < y_0 \\ 2\sigma\big(k_{ub}(x - y_1)\big) - 1 & x > y_1 \\ 0 & otherwise \end{cases}, \qquad (4)$$

where the hyperparameters k_{lb} and k_{ub} control the elasticity of the function below and above the range limit, respectively. This prevents a flat cost function response for out-of-range parameters, thus improving the convergence. We used $k_{lb} = 10^{-2}$ and $k_{ub} = 10^{-3}$ for all the experiments.

Candidate openings can overlap depending on the current parameter state x since the position and scale for each opening are independent coordinates unless the openings are explicitly instanced. To mitigate this, we incorporate an additional term into the cost function. Specifically, we utilize the *Intersection-over-Union* function to compute the percentile volume overlap between each template pair (T_i, T_j), as follows:

$$C_T \;=\; \sum_{i<j} \frac{V\big(T_i \,\cap\, T_j\big)}{V\big(T_i \cup T_j\big)}, \qquad i, j =, \ldots, |T|,\ (5)$$

where T is the set of all active opening templates (Section 3.1) and the function $V(\cdot)$ measures the volume in world space coordinates. The resulting penalty increases the overall cost of the surrogate objective,

thereby deterring the optimizer from selecting the configurations. Using this term, the global cost function becomes

$$C(S,T,x) = C_{illum}(S,x) + wC_T(x), \quad (6)$$

where the hyperparameter w controls the impact of the penalty function on the overall cost. It was set to 10 for all the experiments.

4.4 Cost Minimization

In our method, mapping to the Bayesian framework outlined in Section 2.3 is relatively straightforward. The data samples composing the set D are defined by the number of opening configuration states x_i in conjunction with their associated costs evaluated using Equation 6. Our goal is to identify the optimum opening configuration $x^\star$ by minimizing Equation 1. Therefore, we adopted an online training scheme in which only a few initial samples, $|D| = 30$ in our case, are used to fit the initial model. The training samples are generated by Latin-hypercube stratified sampling of the opening configuration parametric domain.

For the posterior distribution $P(f \mid D)$, we choose a zero-mean Gaussian process with an isotropic Matérn kernel [28], with the hyperparameter $\nu = 5$. In the subsequent lighting and cost evaluation iterations, we query the Bayesian optimizer for the next exploration step to identify the most promising configurations using Equation 2. We guide the exploration phase of our framework using the Expected Improvement [29] acquisition function without hyperparameters. Then, at a fixed step-size of 10 iterations, we refit the entire model using all prior knowledge and iteratively apply this process until a satisfactory solution is found or the maximum number of 100 steps is reached.

4.5 Design Solution Recommendations

The valid generated solutions that satisfy the constraints and achieve the illumination goals are stored, and the K best is presented to the designer. K is a user-defined option. The designer can then select the most appropriate goal according to the aesthetic or functional goals that could not have been defined stringently otherwise. Because the optimizer can produce many similar valid solutions, we clustered these using the K-means clustering algorithm. Here, the distance corresponds to the $\ell 2$-norm on the state vector x. After clustering, we select the solution closest to the mean of the cluster as the representative solution. We chose the K-means method. This is because it exhibits high performance and generates separable clusters, which is a key characteristic for visually distinctive design variations. It is important to note here that the designer can freely adjust the parameters of the resulting openings, "freeze" a subset of the parameters, and iteratively refine the design by re-optimizing the remaining ones.

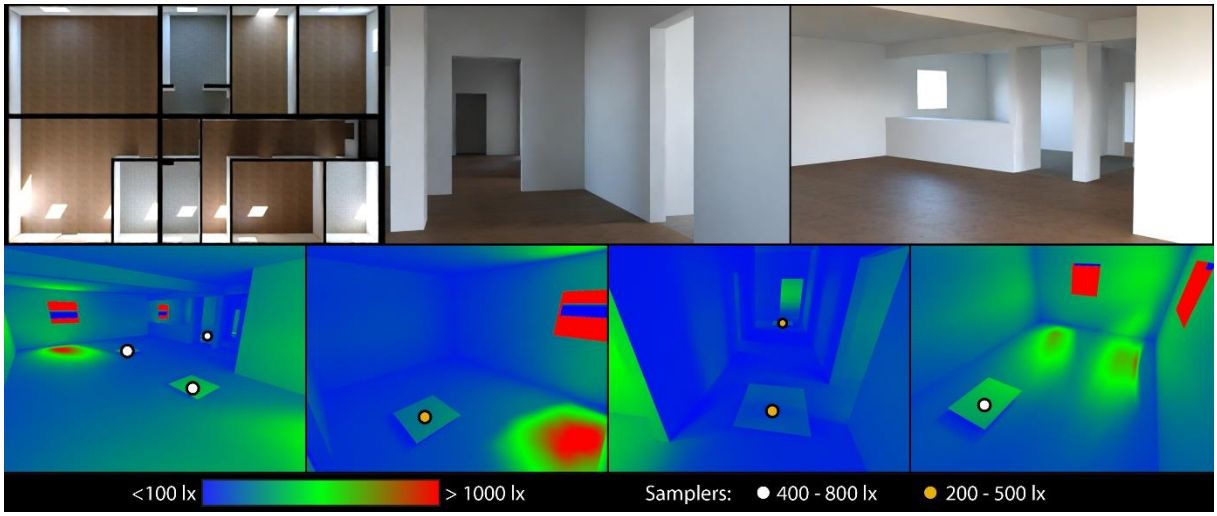

Figure 6 Opening design configuration proposed by our method. The top row illustrates the optimized opening locations. In the bottom rows, the heatmaps quantify the light contributions at various sampling locations. The scene is optimized for 14 openings using 2 view and 11 planar samplers for a total of 41 optimization parameters. Sky lighting is set for office hours and the summer–fall seasons in New York.

5 Implementation and Evaluation

5.1 Implementation Details

We constructed our opening design system as a plugin on top of the existing world modeling and interactive rendering pipeline offered by Unreal Engine 5 [30] using the native physically-based path tracing simulator for global illumination. We used the native Unreal Engine plug-in for the CSG operations applied to the scene geometry. This enabled us to fully utilize the existing integrated real-time modeling and rendering lifecycle of the engine; utilize the built-in hardware-accelerated path tracer, world building, and asset management facilities; and focus on the tasks at hand.

The planar samplers discussed in Section 3.1 are sampled regularly over their surface. For each sample, a batch of random, cosine-weighted rays is traced within the scene to capture the illuminance values. Each batch contains 32 rays. For the view samplers discussed in that section, we typically use a frustum resolution of 16 × 16 pixels with 32 samples per pixel.

To implement the Bayesian optimization framework, we adapted a publicly available implementation [31].

5.2 Evaluation

We tested our method on a collection of architectural instances of varying complexity in terms of illumination, geometric and aesthetic goals. These scenes are shown in ***Figures 6, 7, 8,*** and ***9***. The performance measurements were recorded on an NVIDIA RTX 3080Ti with 12 GB of VRAM and an Intel i7 12700 K CPU with 32 GB of RAM. We provide the full source code of our proposed method and the datasets presented in this work at https://github.com/cgaueb/ods.

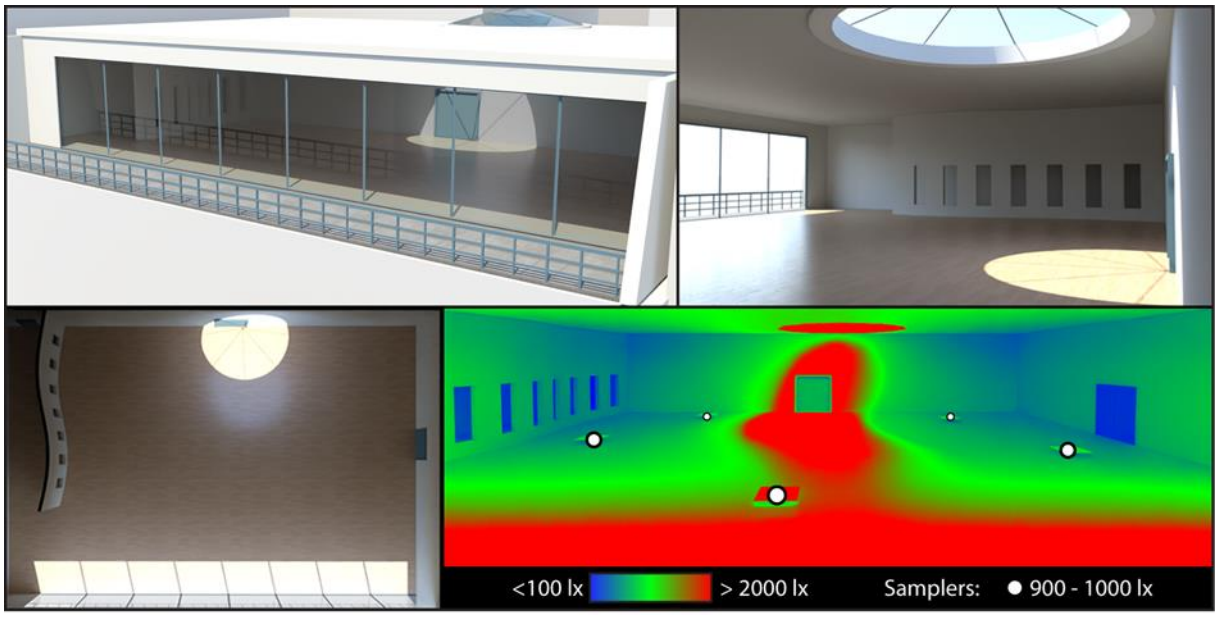

Figure 7 The scene is optimized for two openings using two view and five planar samplers for a total of seven optimization parameters. Sky lighting is set for office hours for the spring–summer seasons in Rome.

Qualitative Evaluation. To stress the opening domain configuration and record the quality of the optimized solutions, we conducted experiments on distinct cases. These included scenes with high degrees of freedom to test the space exploration performance and those with more stringent and potentially more realistic design constraints.

A representative experiment of a typical design workload is presented in the **Apartment** scene (***Figure***). It is a standard apartment in which the designer pre-determines the functional layout of the space. The front door leads to a living room that forms a unified space with the kitchen. This part of the house requires an additional door leading to the outside as well as adequate lighting for various important human tasks (cooking, reading, item location, etc.). A series of patch samplers were distributed to achieve uniformly high illuminance (400–800 lx). Additionally, a view sampler was incorporated to suppress the adverse glare caused by the nearby openings. This was performed to reserve space for potential sitting or entertainment spots. The other parts of the house were intended for bedrooms, bathrooms, and study spaces. Planar samplers were used to set compatible illumination goals. A view sampler was positioned such that it allowed for the generation of a home office space with low glare. For the bedrooms, a marginally lower illumination goal was set (200–500 lx). Opening domains were set on the exterior walls of each room. The opening template is a box parameterized over its scale and position. The resulting configurations and illumination levels are shown in ***Figure*** . The opening design framework respects all the geometric constraints while satisfying the lighting goals. The goal attainment percentage is 92.8%.

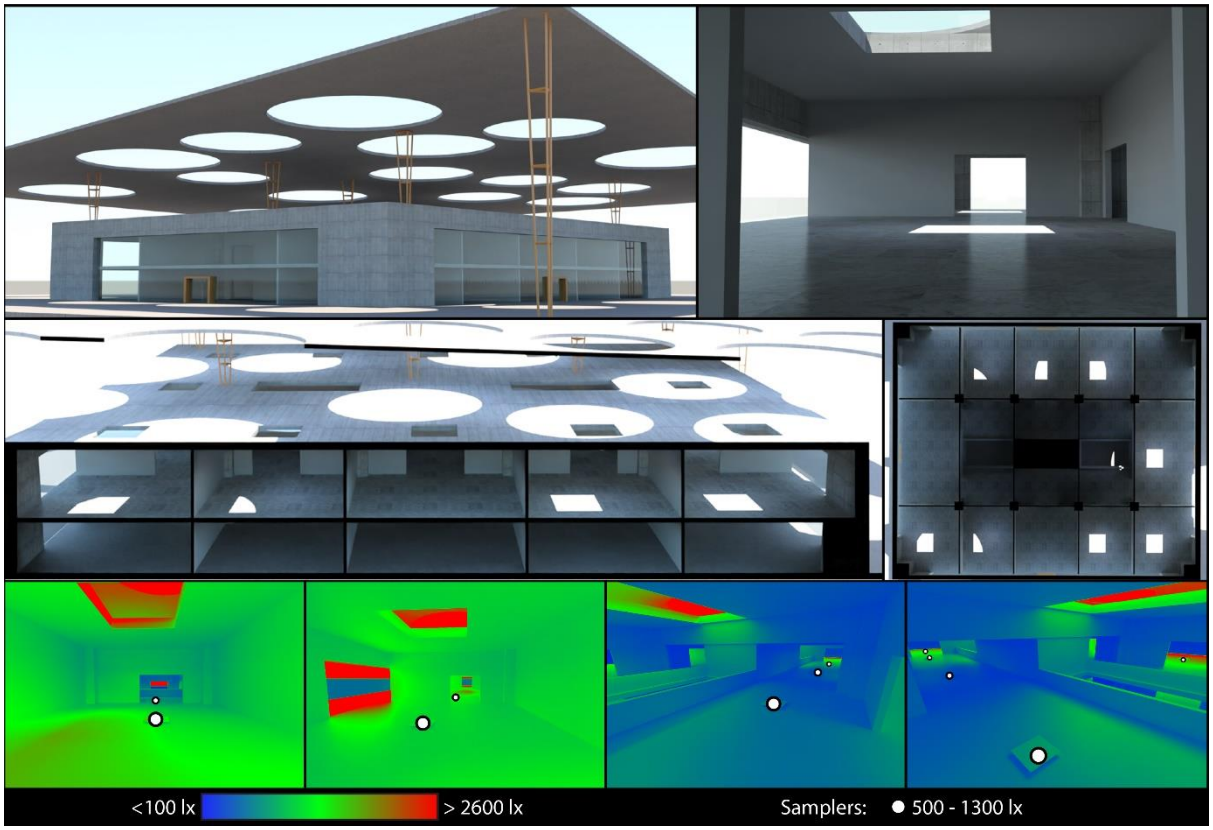

Figure 8 The scene is optimized for 15 openings using 20 planar samplers for a total of 45 optimization parameters. Skylighting is set for office hours for the summer–fall seasons in Capetown.

The **Hall** scene (***Figure 7)*** represents a social event room. This is a scenario in which the geometric constraints disagree with the lighting goals. In particular, a large opening in front of the building should provide convenient access to the balcony. This requirement generates a bright zone. However, it leaves the back of the room underlit. To counter this effect, a cylindrical opening was positioned so that adequate lighting penetrated the deeper parts of the room. Planar samplers (5) with a target illuminance of 900–1000 lx were used to establish appropriate lighting conditions for the entire event room. In addition, two view samplers with a luminance range of 0–1000 nt were employed to suppress the adverse glare caused by the large balcony door. The entire front wall and ceiling were indicated as opening domains, and the sizes and positions of the two openings were optimized. The resulting configuration naturally created an overhang to suppress excess glare, while light penetrated to the back through the skylight. A goal-attainment percentage of 87.6% was achieved.

The **Mall** scene (***Figure 8***) simulates a renovation action in which the initially inferior skylight design is retrofitted using a newly added canopy. The lighting goal was to decrease the excess light from the glass roof panels. The only geometric constraint was the cylindrical shape of the opening. In this example, the overlap cost is more relevant than in the other examples. Even with 15 opening templates, overlap was prevented, whereas the additional construction fixed the inconsistencies of the earlier design. A collection of 20 planar samplers with a target range of 500–1300 lx was spread over the building to maintain adequate illumination primarily for the shopping spaces on the first floor. The resulting configuration achieved a goal attainment percentage of 82.0%.

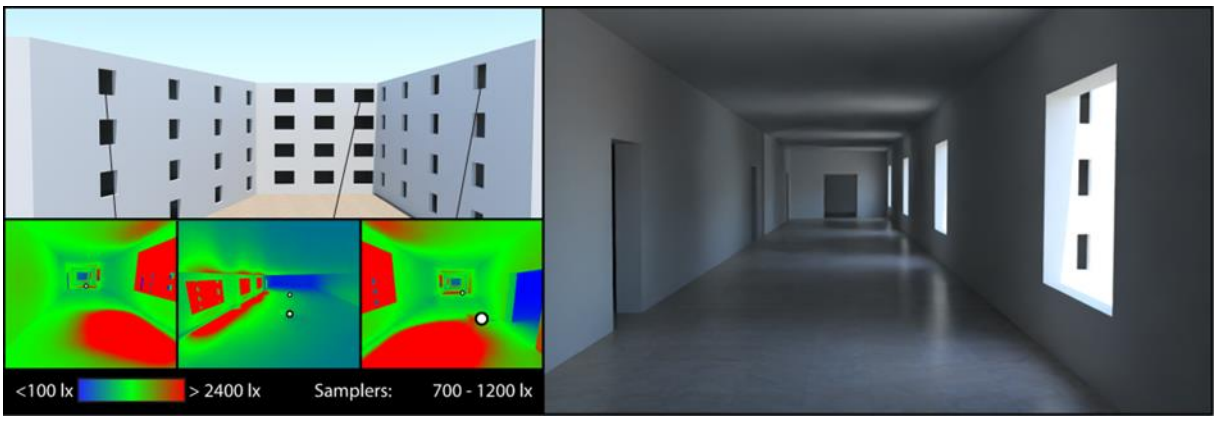

Figure 9 The scene is optimized for variable number of openings using 32 planar samplers for a total of 10 optimization parameters. Sky lighting is set for office hours for the spring–summer seasons in Rome.

In the **University Building** (***Figure 9***), the illumination goal was to provide sufficient lighting to the hallways of the four-story compound. We used 32 planar samplers (eight per floor) with a target range of 700–1200 lx. This range ensured a well-lit interior for all the building corridors. The opening domains were associated with the exterior walls of each wing and the main building. Achieving a configuration that satisfies all the floors presents a challenge. This is because the occlusion factor from the surrounding environment and the building varies substantially among the floors.

The opening templates were optimized over all the parameters, namely, the position, scale, and spacing offset. Enabling the optimization of the spacing offset practically solves for the number of required openings. To ensure symmetry and alignment between the floors, we used instancing (discussed in Section 3.1) on the vertical plane. We employed instancing for the domains of the two opposite wings as well to obtain more regular results. This is an example of a large number of samplers placing an excessive demand on our method both in terms of lighting evaluation and cost minimization, because the lighting goals on higher floors are significantly more convenient to satisfy than those on lower floors. The resulting solution attained this goal in 85.5% of the cases.

Finally, the **Mansion** scene (***Figure 1***) is a complete example that takes advantage of all the available constraints (instancing, symmetry, and spacing) and uses numerous illumination targets. The opening domains on the façade and backside use instancing in both vertical and horizontal planes to achieve symmetry. The windows use a box-shaped opening template except for a circular central window over the main door. The sidewalls employ instancing and domain-grouping methods to accommodate several design options. Planar samplers (15) are placed throughout the interior. The rooms have a target illuminance range of 700–1000 lx, whereas the hallways have an ancillary illumination target of 300–600 lx. The process achieved a goal-attainment rate of 95.7%.

Performance Measurements. The execution time of the proposed method is spent on three distinct modules. First, CSG operations are required to map the state vector to the actual openings. This is trivial to

apply. The medium-complexity models require less than 2 ms for each template on average. A path-tracing simulation is executed for each sampler. For the sampler configuration discussed in Section 5.1, the average time per scene to obtain the simulation result and apply the denoising post-processing step is 12 ms for each sampler. Finally, the Bayesian optimizer requires less than 1 ms (on average) for the initial fitting of the model and subsequent refitting stages. Querying the optimizer for new samples requires 20 ms on average. Each opening design estimation session took less than 2 min (on average) to complete. Please consider that most of the time was consumed for updating Unreal Engine's internal state and rather than by the stages described above.

6 Conclusions

In this work, we presented a goal-driven method for designing openings in architectural scenes using the Bayesian optimization theory. Using the proposed opening specification, the user is presented with a flexible and expressive framework that accounts for both lighting goals and functional or aesthetic geometric constraints. A complete implementation was integrated into a high-performance rendering engine to support accurate photometric evaluation. This enabled us to define the appropriate types of samplers that exploit photorealistic light simulation. Finally, we incorporated physical daylight measurements from existing sky models that account for the complete temporal trajectory of the sun disk.

Limitations. We evaluated several strategies to minimize the number of parameters required for optimization while simultaneously satisfying the geometric constraints. However, in scenarios with an excessive number of parameters, the results of the proposed framework deteriorated in terms of quality ("the curse of dimensionality"). Needless to say, contradicting or unrealistic goals have a significant impact on the quality of the proposed solutions.

Future Work. In this work we did not include any direct energy consumption metrics in our cost function. Carlos et al. [32] demonstrated results while optimizing the window size for both visual comfort and energy consumption. They concluded that optimizing for only one goal hinders performance with regard to the other goal. Windows optimized exclusively for visual comfort produce large energy-consumption patterns. However, optimizing window size for low energy consumption does not satisfy any predetermined visual acceptance criteria. In the future, we will include metrics regarding energy consumption within our cost-minimization framework by combining natural and artificial lighting solutions in a unified framework. Finally, we would like to investigate data-driven methods based on deep neural networks that have shown very promising results in scene synthesis [33]. Similar approaches could help encode the designer's process within the network and be used as generative models. However, this would require a professionally curated and extensive dataset.

Acknowledgements. This research was funded by the Hellenic Foundation for Research and Innovation (H.F.R.I.) under the "3rd Call for H.F.R.I. Research Projects to support Post-Doctoral Researchers" (Project No. 7310).

References

1 R. M. J. Bokel, "The effect of window position and window size on the energy demand for heating, cooling and electric lighting," in *Building Simulation*, 2007.

2 P. Tregenza and M. Wilson, Daylighting: architecture and lighting design, Routledge, 2013.

3 E. Ochoa Morales, M. B. C. Aries and E. J. Loenen, "Considerations on the design optimization criteria for windows providing low energy consumption and high visual comfort," *Applied Energy,* vol. 95, p. 238–45, 2012.

4 A. Nabil and J. Mardaljevic, "Useful daylight illuminance: a new paradigm for assessing daylight in buildings," *Lighting Research & Technology,* vol. 37, pp. 41-57, 2005.

5 K. Fisekis, M. Davies, M. Kolokotroni and P. Langford, "Prediction of discomfort glare from windows," *Lighting Research & Technology,* vol. 35, pp. 360-369, 2003.

6 J. Wienold and J. Christoffersen, "Evaluation methods and development of a new glare prediction model for daylight environments with the use of CCD cameras," *Energy and Buildings,* vol. 38, pp. 743-757, July 2006.

7 A. Mahdavi and L. Berberidou-Kallivoka, "A generative simulation tool for architectural lighting," *Evaluation,* vol. 2, 1995.

8 V. Tourre, J.-Y. Martin and G. Hégron, "An Inverse Daylighting Model for CAAD," in *Proceedings of the 24th Spring Conference on Computer Graphics*, New York, NY, USA, 2008.

9 E. Fernández and G. Besuievsky, "Technical Section: Inverse Lighting Design for Interior Buildings Integrating Natural and Artificial Sources," *Comput. Graph.*, vol. 36, December 2012.

10 E. Fernández and G. Besuievsky, "Inverse opening design with anisotropic lighting incidence," *Computers & Graphics,* vol. 47, pp. 113-122, 2015.

11 E. Fernández, J. P. Aguerre, B. Beckers and G. Besuievsky, "Optimizing Window Shape for Daylighting: An Urban Context Approach," in *Eurographics Workshop on Urban Data Modelling and Visualisation*, 2016.

12 G. Besuievsky, E. Fernández, J. P. Aguerre and B. Beckers, "A radiosity-based methodology considering urban environments for assessing daylighting," *Journal of Physics: Conference Series,* vol. 1343, p. 012156, November 2019.

13 K. Kalampokis, G. Papaioannou and A. Gkaravelis, "A Generic Physically-Based Approach to the Opening Design Problem," in *Proceedings of the 37th Annual Conference of the European Association for Computer Graphics: Short Papers*, Goslar, 2016.

14 C. T. Committee and others, "Spatial distribution of daylight-luminance distributions of various reference skies," 1994.

15 R. Perez, R. Seals and J. Michalsky, "All-weather model for sky luminance distribution—Preliminary configuration and validation," *Solar Energy,* vol. 50, pp. 235-245, 1993.

16 A. J. Preetham, P. Shirley and B. Smits, "A Practical Analytic Model for Daylight," in *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, USA, 1999.

17 L. Hošek and A. Wilkie, "Adding a Solar-Radiance Function to the Hošek-Wilkie Skylight Model," *IEEE Computer Graphics and Applications,* vol. 33, pp. 44-52, 2013.

18 D. R. Jones, M. Schonlau and W. J. Welch, "Efficient Global Optimization of Expensive Black-Box Functions," *J. of Global Optimization,* vol. 13, p. 455–492, December 1998.

19 B. Shahriari, K. Swersky, Z. Wang, R. P. Adams and N. de Freitas, "Taking the Human Out of the Loop: A Review of Bayesian Optimization," *Proceedings of the IEEE,* vol. 104, pp. 148-175, 2016.

20 C. E. Rasmussen and C. K. I. Williams, Gaussian processes for machine learning., MIT Press, 2006, pp. I-XVIII, 1-248.

21 A. Shah, A. Wilson and Z. Ghahramani, "Student-t Processes as Alternatives to Gaussian Processes," in *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, Reykjavik, 2014.

22 J. Snoek, H. Larochelle and R. P. Adams, "Practical Bayesian Optimization of Machine Learning Algorithms," in *Advances in Neural Information Processing Systems*, 2012.

23 R. Martinez-Cantin, N. de Freitas, E. Brochu, J. A. Castellanos and A. Doucet, "A Bayesian exploration-exploitation approach for optimal online sensing and planning with a visually guided mobile robot," *Autonomous Robots,* vol. 27, pp. 93-103, 2009.

24 E. Brochu, V. M. Cora and N. de Freitas, "A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning," *ArXiv,* vol. abs/1012.2599, 2010.

25 Y. Koyama, I. Sato, D. Sakamoto and T. Igarashi, "Sequential Line Search for Efficient Visual Design Optimization by Crowds," *ACM Trans. Graph.,* vol. 36, July 2017.

26 N. Vitsas, G. Papaioannou, A. Gkaravelis and A. A. Vasilakis, "Illumination-Guided Furniture Layout Optimization," *Computer Graphics Forum,* vol. 39, pp. 291-301, 2020.

27 L. Hosek and A. Wilkie, "An Analytic Model for Full Spectral Sky-Dome Radiance," *ACM Trans. Graph.*, vol. 31, July 2012.

28 M. G. Genton, "Classes of Kernels for Machine Learning: A Statistics Perspective," *J. Mach. Learn. Res.*, vol. 2, March 2002.

29 E. Vazquez and J. Bect, "Convergence properties of the expected improvement algorithm with fixed mean and covariance functions," *Journal of Statistical Planning and Inference,* vol. 140, pp. 3088-3095, 2010.

30 Epic Games, *Unreal Engine,* 2023.

31 R. Martinez-Cantin, "BayesOpt: A Bayesian Optimization Library for Nonlinear Optimization, Experimental Design and Bandits," *Journal of Machine Learning Research,* vol. 15, p. 3915–3919, 2014.

32 C. E. Ochoa, M. B. C. Aries, E. J. van Loenen and J. L. M. Hensen, "Considerations on design optimization criteria for windows providing low energy consumption and high visual comfort," *Applied Energy,* vol. 95, pp. 238-245, 2012.

33 K. Wang, M. Savva, A. X. Chang and D. Ritchie, "Deep Convolutional Priors for Indoor Scene Synthesis," *ACM Trans. Graph.*, vol. 37, July 2018.