

Illumination-Guided Furniture Layout Optimization

N. Vitsas¹, G. Papaioannou¹ , A. Gkaravelis¹ and A. A. Vasilakis¹ 

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



Figure 1: Lighting significantly affects task performance and is therefore an important factor in interior design. Our method optimizes furniture arrangement by combining illumination goals with functionality constraints, resulting in usable, comfortable layouts.

Abstract

Lighting plays a very important role in interior design. However, in the specific problem of furniture layout recommendation, illumination has been either neglected or addressed with empirical or very simplified solutions. The effectiveness of a particular layout in its expected task performance can be greatly affected by daylighting and artificial illumination in a non-trivial manner. In this paper, we introduce a robust method for furniture layout optimization guided by illumination constraints. The method takes into account all dominant light sources, such as sun light, skylighting and fixtures, while also being able to handle movable light emitters. For this task, the method introduces multiple generic illumination constraints and physically-based light transport estimators, operating alongside typical geometric design guidelines, in a unified manner. We demonstrate how to produce furniture arrangements that comply with important safety, comfort and efficiency illumination criteria, such as glare suppression, under complex light-environment interactions, which are very hard to handle using empirical or simplified models.

CCS Concepts

• Computing methodologies → Graphics systems and interfaces;

1. Introduction

Inverse geometry problems cover a wide area of research that is actively being explored over the past few years. The term encompasses many aspects of geometry optimization via parametric or direct manipulation, driven by specific constraints and goals. Typically a user specifies a set of goals that need to be satisfied by the geometry and a system calculates and proposes valid parameter states that best satisfy them. Automatic and semi-automatic furniture layout is a specific inverse geometry problem relevant to interior design, where a given space must be populated according to functional and aesthetic rules, with either a predetermined or automatically proposed set of furniture pieces, resulting in an optimal arrangement or synthesized environment, respectively.

As will be discussed in Section 2, several furniture arrangement approaches focus on spatial, ergonomic and utilitarian aspects of a layout. However, as indicated by standard practical guidelines [TW13, KSB17], illumination is a core aspect of interior design and measuring the effectiveness of illumination in terms of task-specific target illuminance or intensity levels is an established procedure. Illumination, especially from dominant sources such as sunlight, sky lighting and main fixtures, directly affects comfort levels, task and energy efficiency. For example, a modern living room needs to be designed so that the layout takes advantage of any natural lighting during daytime, while efficiently utilizing artificial lighting in low-light conditions. Likewise, office spaces should be designed and laid out so that lighting conditions favor comfortable

and uninterrupted work for employees, including direct and indirect glare elimination on work surfaces and screens, which can greatly hinder the operator's performance due to eye strain.

Illumination-guided design has often been addressed by very approximate or even empirical approaches in the literature (e.g. [TW13, FB15]). However, in realistic, practical scenarios, light transport in a complex environment with diverse materials can cause the distribution of incident light on object surfaces to vary significantly and is greatly affected by their relative arrangement. Furthermore, important aspects of lighting design, such as glare minimization, depend on the directional characteristics of incident light, something that has often been neglected in the past and can only be addressed by a light transport simulation that encompasses all major light-surface interaction events.

In this paper, we contribute to the current state of the art by introducing detailed and generic illumination constraints to the furniture layout problem. We support multiple types of lighting goals (or *lighting intentions*) and evaluate illumination by physically-based light transport estimators. We combine our illumination-driven approach with established functional constraints from the literature in a unified method that recommends usable furniture layouts (see an example in Fig. 2). We demonstrate how taking into account complex light interactions helps elegantly address highly-directional lighting constraints associated with polished surfaces, gaze or focus direction preference and glare. As a result, the generated furniture arrangement recommendations comply with important safety, comfort and efficiency considerations related to illumination, which are very often hard to address with empirical, simplified or manual approaches. In essence, we transform a tedious, iterative trial and error process into a nearly linear one, providing a starting point for aesthetic adjustments, after addressing comfort and functional aspects of the design. Finally, we adapt a Markov Chain Monte Carlo optimization process to the specific problem at hand performing the following optimizations: We construct a hierarchical mutation strategy that accommodates functional object groups, which naturally fits the problem of furniture layout and helps speed up optimization. We introduce disjoint parameter ranges and show how to effectively incorporate their non-continuous nature in the optimization strategy of our method and only perform the expensive lighting constraint evaluation for object placements that do not violate geometric constraints.

In our system, the user is responsible for providing the geometry of an environment and the objects that must be arranged within its extents. We closely follow the priorities of an interior designer by addressing functionality and comfort constraints and leaving aesthetic interventions last, to be provided by the domain expert. According to standard practical guidelines [TW13, KSB17], measuring the effectiveness of illumination in interior design in terms of task-specific target illuminance or intensity levels is an established procedure. Therefore, the user also provides the desired illumination levels constrained on specific surfaces, directions or volumes attached to the given geometry or placed in free space. Coupled with a plethora of ergonomics and utility guidelines, expressed as geometric constraints, lighting intentions can greatly complement the pipeline of a professional interior designer.

We differentiate this work from light source placement problems, where the primary goal is to establish the number, position and/or orientation of luminaires in a fixed environment. For practical application purposes, interior designers primarily rely on high-power sources, which cannot be altered - unless a thorough and radical renovation of the space is requested - since they constitute either part of the natural lighting or pre-installed building fixtures. In low-light conditions, when the existing illumination sources fail to satisfy our constraints, the user can opt to include the placement of movable light sources in the layout recommendation, which are seamlessly supported by our method.

2. Related Work

In this section, we briefly review prior layout optimization methods that encompass design guidelines and lighting intentions either separately or jointly and stress on the contributions of this work beyond the state of the art.

In the domain of constrained layout optimization, Harada et al. [HWB95] proposed a system for assisted layout design, in which constraints were interactively checked during object manipulation. Violations would trigger a local search for the conflict resolution. Nakajima et al. [NHH*06] presented a specialized interactive method for the specific problem of populating office interiors with furniture and office equipment. More recently, Yu et al. [YYT*11] proposed a fully automatic system for furniture arrangement of indoor scenes. The user provides an initial set of previously furnished interiors, which are used to extract spatial relationship metrics for furniture pieces. The combined metrics are minimized as a cost function via Simulated Annealing in order to find the best arrangement. Fisher et al. [FRS*12] proposed a technique that trains probabilistic models on an existing scene database for both object occurrence and arrangement. These are later used to synthesize new scenes. Fu et al. [FCW*17] used an object relation database along with predefined object categories, in an adaptive method for 3D scene synthesis using activity-associated object relation graphs. In a similar manner, Merrell et al. [MSL*11] proposed a semi-automatic method for interactive layout design. To this effect, a wide range of mathematically-modelled interior design guidelines is supported and combined in a cost function, which

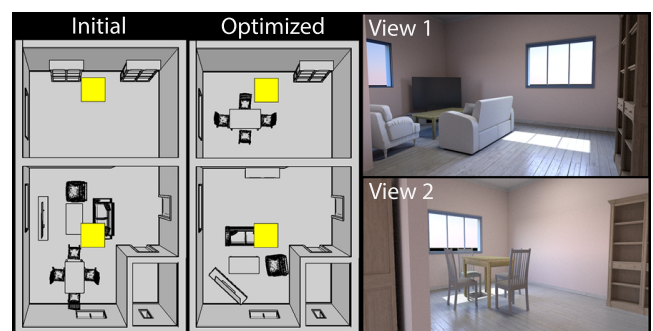


Figure 2: An example scene optimized for reading and dining task illumination and glare minimization (TV screen). The ceiling light fixtures are shown in yellow.

drives a Markov Chain Monte Carlo-style exploration of the transformation parameter space in order to complement the designer's actions. Recently, Ma et al. [MPF*18] proposed a natural language-driven model for 3D scene synthesis.

Several methods target procedural scene population and synthesis. Germer and Schwarz [GS09] described a method to procedurally generate plausible interior layouts of buildings, for use in real-time walkthroughs. The work of Xu et al. [XSF02] attempts to fill an empty interior, one object at a time, utilizing spatial and semantic constraints, which are stored in a semantic database for each object class. More recently, Kán and Kaufmann [KK17, KK18] proposed an automatic method for quick interior environment synthesis using a genetic algorithm and greedy cost minimization, respectively. Their approach combines object selection from a repository as well as material selection for the final configuration to achieve a pleasant color scheme in the generated scene.

Very recently, with the emergence of semantically annotated 3D indoor scene datasets [LSM*18, SYZ*17], machine learning approaches, and particularly deep learning generative techniques, have been applied to both 2D [WSCR18] and 3D [LPX*19, ZYM*18] representations of interior scenes to great effect.

The literature for illumination-guided furniture optimization is very sparse, since none of the above methods includes any lighting constraints. Yamakawa et al. [YDY16] attempted to solve a similar problem to ours, but with an extremely simplified formulation and lighting evaluation approach. The authors proceed to optimize a single scene using object constraints and target illumination levels. Illumination is evaluated using diffuse inter-reflection based on the radiosity method. The authors claim that the contribution of inter-reflections among objects is not significant, despite experimental evidence to the contrary, especially for highly reflective surfaces, densely occupied environments or openings with overhangs. Lastly, they do not account for natural lighting either from the sun or the sky dome.

Illumination as a goal has been central to other forms of inverse design, such as inverse lighting optimization for luminaire placement, car headlight design and opening design. Notable methods in this genre include the works of Kawai et al. [KPC93] for designing the illumination in a static environment, Schwarz and Wonka [SW14] on street light pose optimization, Gkaravelis and Papaioannou [GP16] on the population of interior environments with light sources and Mas et al. [MMP18] on headlight reflector design.

The RADIANCE system from Ward [War94] was developed for the specific demands of lighting design and architecture and was one of the very early attempts to incorporate accurate, physically-based simulation of lighting for tasks other than image synthesis.

3. Method Overview

Given an initial selection of objects for a user-defined interior 3D space, our goal is to find object arrangements that best satisfy interior design guidelines, while respecting lighting intentions set by the designer. To that end, we optimize a cost function, which combines both illumination goals and functional (geometric) con-

straints, over a transformation parameter vector \mathbf{x} . \mathbf{x} is the concatenation of all degrees of freedom \mathbf{x}_i for translation and rotation of entities (Fig. 5 - top right). Generic semantic constraints associating pieces of furniture can be intuitively transformed into geometric ones, as described in Section 5 and also previously discussed in the literature (e.g. [MSL*11]). Please note that certain transformation parameters may be immutable, in order to either respect physical limitations or enforce user-defined invariants. Likewise, certain geometric parts may be completely immovable, solely acting as collision targets and support surfaces for lighting constraints.

We perform arrangement optimization hierarchically, therefore local transformations are defined both at object level and at group level (see illustration in Fig. 3). Groups represent user-defined clusters of objects. For furniture arrangement, these signify functional groups, whose integrity and coherent motion must be respected during optimization. As shown by Xu et al. [XMZ*14], furniture clusters can be inferred from their spatial arrangement in a sample environment or an initial proposed interior design draft. Of course, they may alternatively be explicitly declared through the modelling software, at design time. It is noteworthy that grouping information is not specific to a particular scene. It embodies semantic relations among different types of objects rather than specific geometry. It is therefore inherently transferable and applicable to different objects of similar function.

The importance of hierarchically optimizing the spatial object arrangement is three-fold. First, it allows for parameter space exploration in a meaningful manner, facilitating parameter swapping and mutations, while retaining functional relationships directly and not only through the constraints of the objective function. For example, a dining table set can swap position with a home entertainment cluster of furniture, without breaking intra-group cohesion. Second, the hierarchical parameter space exploration significantly boosts search performance, by both avoiding functionally inadmissible solutions and reducing the dimensionality of the search space (per level). Third, it allows the optimization method to work differently per-level. In our case, we allocate more iterations to intra-group mutations to favor optimizing a particular arrangement of functional groups before moving to another, drastically different layout.

Interior design guidelines are provided in the form of valid local transformation parameter ranges and objective function constraint terms, which are all grouped in a geometric error term $G(\mathbf{x})$. Multiple valid parameter ranges for both the degrees of freedom and the constraints are supported. This formulation makes it especially easy to define constrained motion of objects with respect to multiple other fixed or movable geometric parts. More details are provided in Section 5.

Illumination goals are supplied through a number of user-defined *light samplers*. Each light sampler has an associated range of acceptable illuminance or average luminance values, depending on its type. We define three types of light samplers: *planar*, *volumetric* and *directional*. Planar and volumetric samplers measure the average incoming light at sample positions over a rectangular patch and within a specific bounded volume, respectively. Directional samplers measure the peak average luminance within a specific solid




Figure 3: The various geometric and lighting constraints that can be enforced by our interior layout optimization system. Transformations are hierarchically optimized at group level and then at object level. Geometric constraints can be enabled across object hierarchy levels.

angle around a direction ω and are used for modeling potential gaze directions and defining glare-related constraints.

Luminance or illuminance levels are the standard criteria for the assessment of lighting conditions by interior designers and architects. They are used to determine comfort levels, ease of use, task-level performance and to some extent, aesthetic look for parts of the interior. In fact, acceptable levels for most tasks can be readily found in most architectural lighting guides. Due to the fact that each type of furniture is associated with specific functions, it is easy to reuse previously configured samplers that are semantically annotated based on the desired target surface or space (e.g. reading levels, conversation levels, cooking levels etc.). The reusability of both geometric and lighting constraints across scenes and different pieces of furniture can help even non-trained users configure their own environment using pre-annotated constraints from interior design professionals.

Light samplers can be defined at any hierarchical level, representing movable constraints that follow their associated pieces of furniture or scene-level goals. Details about the illumination measurement are provided in Section 4.

The cost C of a configuration \mathbf{x} is calculated as the weighted combination of the lighting constraints function $L(\mathbf{x})$ and the geometric term $G(\mathbf{x})$. The global weight w can be used to prioritize either functional or illumination goals in the target layout:

$$C(\mathbf{x}) = wL(\mathbf{x}) + (1 - w)G(\mathbf{x}). \quad (1)$$

A key observation regarding our method is that we allow for marginal violations of geometric and lighting constraints in order to favor a wider exploration of the parametric space and reach optimal states in terms of illumination. To this end, for all constraints, we properly parameterise the following penalty function P . P introduces no penalty for constraint values c within the acceptable limits b_{min}, b_{max} , but gradually penalizes values outside this range. Geo-

metric constraint values deviating above a certain threshold b_{rej} result in the problematic parameter vector \mathbf{x} being rejected, i.e. when $P(\cdot) > 1$.

$$P(c, b_{min}, b_{max}, b_{rej}, \alpha) = \begin{cases} k(b_{min} - c)^\alpha & c \leq b_{min} \\ 0 & b_{min} < c < b_{max} \\ k(c - b_{max})^\alpha & c \geq b_{max} \end{cases} \quad (2)$$

In the above penalty function, $k = b_{rej}^{-\alpha}$, where α is the stiffness factor that tunes the rate of error and b_{rej} is the elastic limit of the error, beyond which the particular configuration is rejected as unacceptable. This can be user-provided but is typically a fraction of the acceptable interval length. α values could also differ for each side of the bounds but in this work we only consider equal stiffness factors.

Due to the relaxed penalization of geometric constraints, the recommended layout may require small manual adjustments by the designer to resolve small inconsistencies and also enforce personal aesthetic preferences. Our experiments show that this added relaxation introduces an insignificant amount of error.

It is important to note that the cost of evaluating the geometric part $G(\mathbf{x})$ of the cost function is orders of magnitude lower than the respective cost of evaluating incoming lighting and the resulting illumination error $L(\mathbf{x})$. Therefore, we evaluate $G(\mathbf{x})$ first and proceed to computing $L(\mathbf{x})$ only if the current layout \mathbf{x} does not completely violate any of the geometric constraints. This allows the optimization framework to quickly discard invalid regions of the state space.

4. Illumination Constraints

Sunlight and sky dome illumination are the predominant sources of illumination in architectural design. To account for natural light

due to the sky dome contribution and the sun disk illumination, we used the analytic model and results of [PSS99]. We simulate the sun and sky over a specified time period and bake the average luminance in an environment map. Artificial light sources are physically modeled as area black body emitters and the user provides their geometry and luminous flux in lumens (lm). For dense urban environments or structures that include overhangs, the target space may not be directly lit, but rather achieve strong illuminance levels through indirect light transport. This is why the compliance of the current configuration \mathbf{x} with the lighting intentions is measured by evaluating incoming light at the light samplers via path tracing. We generate paths towards artificial light sources, the sun and the sky dome using multiple importance sampling and portal sampling [UnFK13] on major openings (e.g. windows or skylights).

We support three types of illumination samplers that represent measurements associated with specific desired luminance or illuminance levels:

The **planar** (patch) sampler $S_{patch}(\mathbf{p}, \mathbf{n}, w, h)$ measures the average illuminance in lux (lx) over a number of samples in a quadrilateral patch with normal \mathbf{n} , corner at \mathbf{p} and a width and height w, h . Illuminance is computed over the oriented hemisphere at each point sample. Planar samplers are typically encountered in the bibliography to represent comfort illumination levels on work surfaces and other task areas. Planar samplers can be attached to movable objects or be freely positioned at fixed locations to measure overall illuminance at a specific height.

The simple example in the first row of Figure 4 demonstrates the use of a planar sampler to optimally place the table and bookcase at the most comfortable position in the room for the tasks of dining and reading, respectively, taking advantage of both natural and artificial light.

The **directional** sampler $S_{dir}(\mathbf{p}, \omega, \theta)$ captures incoming light inside a pyramidal frustum of aperture θ , centered around direction ω with the apex at \mathbf{p} . Its primary use is to establish acceptable direct or indirect glare levels (see Fig. 4, second row).

In order to capture the very localized nature of glare we do not average the incoming luminance over the entire frustum but rather split the latter into strata and record the maximum of the average luminance in each one (in nits). Simply measuring the maximum luminance would bias the measurement towards specular noise spikes.

A directional sampler is attached to reflective surfaces, such as TV sets or computer monitor screens, pointing towards their surface or to seats pointing outwards along the front direction in order to measure light as observed by people looking in particular directions. Directional samplers can be also constrained to point to focal points around the scene, such as windows. The use of this sampler is nicely exemplified in Figures 7 and 8.

The **volume** sampler $S_{vol}(\mathbf{p}, \mathbf{u}, \mathbf{v}, \mathbf{w})$ measures the average luminance (in nits) at uniformly distributed point samples within an oriented bounding box centered at \mathbf{p} , aligned with the orthogonal vectors $(\mathbf{u}, \mathbf{v}, \mathbf{w})$ and sides equal in length to the magnitude of the corresponding vectors. Average incoming luminance is estimated over the sphere centered at each volume sample. Volume samplers can be used to establish measurements and reference luminance levels

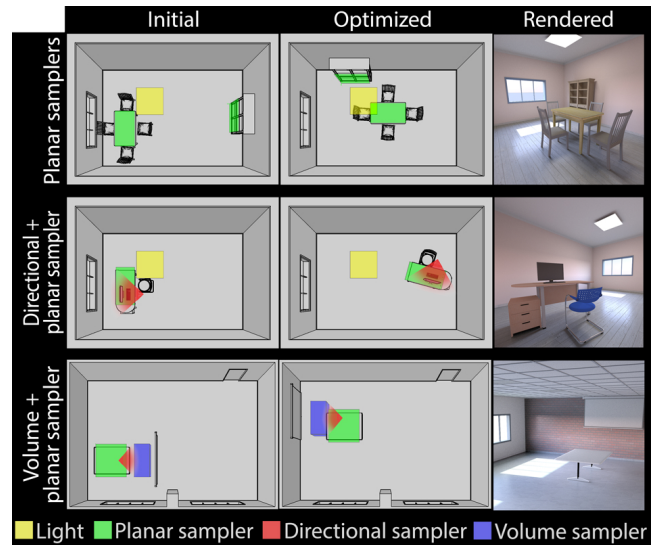


Figure 4: Simple test cases demonstrating the role of each one of the light sampler types.

in large empty volumes in front of or above attention points or areas, such as presentation or theatrical stages, exhibition booths etc.

In the simple example of Figure 4 (third row), we employ a volume light sampler in front of a projection screen to establish a minimum illumination level for the area where the speaker is going to be standing. A complementary directional sampler facing the screen itself enforces low reflected light. The optimizer establishes a projector setup position and orientation so that both contradicting goals are best satisfied.

For patch samplers we measure illuminance at *locations* uniformly distributed over their area. Directional samplers are modeled with an orthogonal frustum subdivided into 8×8 or 16×16 strata with a field of view of $60 - 90$ degrees, depending on the use case. Similarly, for volume samplers we measure average luminance at 200 uniformly distributed positions within their extents. Measurements at each single sample are evaluated using path tracing. We typically generate 1000 paths per sample.

We formulate the deviation of the measured illumination $L(\mathbf{x})$ from the lighting intentions using the penalty function of Eq. 2 on the illumination measurements $\mathcal{L}(S_k, \mathbf{x})$ of each sampler S_k and the current configuration \mathbf{x} :

$$L(\mathbf{x}) = \frac{1}{N_S} \sum_{k=1}^{N_S} w_k P(\mathcal{L}(S_k, \mathbf{x}), \mathcal{L}_{min}(S_k), \mathcal{L}_{max}(S_k), \mathcal{L}_{rej}(S_k), 1). \quad (3)$$

N_S is the number of individual samplers, $\mathcal{L}_{max}(S_k), \mathcal{L}_{min}(S_k)$ specify the acceptable illumination range and $\mathcal{L}_{rej}(S_k)$ the tolerance. We use $\alpha = 1$ for light samplers to allow for a linear penalty function. For lighting intentions, $\mathcal{L}_{rej}(S_k)$ is only indicative of the penalization rate outside $(\mathcal{L}_{max}(S_k), \mathcal{L}_{min}(S_k))$; we do not discard configurations due to illumination deviations, since in realistic scenarios, the desired illumination may not be attainable.

An equal weight w_k is assigned to each sampler, which can also be controlled by the user to bias the importance of lighting intentions. All weights are of course normalized prior to optimization.

5. Design Goals

Design intentions are supported in our framework using two separate mechanisms: a) A set of valid ranges for each degree of freedom in the object hierarchy that is sampled for new states \mathbf{x} and b) a set of rules that functionally constrain the relationship of objects and their clusters, similar in nature to the common practice followed by previous goal-driven furniture arrangement methods.

5.1. Object Placement Limits

Placement constraints are the most straightforward to define and enforce in a scene. An object or group, i.e. a geometric *entity*, moves on the 2D support plane defined by its initial placement, pivoting around its axis of gravity, when allowed.

Placement constraints are strict. Consequently, they are not part of the cost function. Instead, plausible configurations are sampled from the union of valid parameter space partitions \mathbf{X}_m representing a set of potentially disjoint continuous intervals that map to tuples of compatible object parameters (see Fig. 5). For example, allowed 3D translation ranges are defined per entity as the union of multiple local bounding boxes. This formulation makes it especially easy to define constrained motion of objects with respect to multiple other fixed or movable geometric parts. For instance, a bookcase can be placed against two different and possibly disjoint wall sections. These constraints are directly described in our system for each hierarchical object level via scripting.

5.2. Relational and Functional Constraints

To evaluate the geometric term $G(\mathbf{x})$ of the objective function in Eq. 1, we define relationships between objects in the same spirit as [MSL*11]. We uniformly express our constraints using the penalty function of Eq. 2 and present individual details and improvements to prior work in the resulting cost functions that follow. All constraints are either manually specified or automatically extracted during scene modeling and stored in a declarative format. We also allow constraints to be established across object hierarchy levels.

Alignment. For every object, we define one or more front-facing directions that represent either access directions to the furniture piece or a practical side for alignment with other objects. For a given combination of objects, every available front direction is considered during optimization and the best candidate is used. These discrete directions are entered as disjoint parameter sub-spaces X_m for this variable in the configuration vector \mathbf{x} . We use a user-defined angular tolerance $\theta_{m,n}$ for the alignment between two front vectors \mathbf{u}_m and \mathbf{u}_n and use Eq. 2 to define the alignment cost function G_A for a pair of front directions $(\mathbf{u}_m, \mathbf{u}_n)$:

$$G_A(\mathbf{x}, \mathbf{u}_m, \mathbf{u}_n) = P(\mathbf{u}_m \cdot \mathbf{u}_n, 0, \cos \theta_{m,n}, \cos \frac{\theta_{m,n}}{8}, 2). \quad (4)$$

The penalty in $G_A(\cdot)$ is quadratic and the configuration rejection limit becomes a fraction of $\theta_{m,n}$ to reflect the strictness of the alignment rule.

Distance. Pairwise distance constraints are defined via a minimum and maximum allowed distance ℓ_{min} and ℓ_{max} between the centers $\mathbf{c}_s, \mathbf{c}_t$ of two objects s and t . For distances within these limits, the constraint is considered fully satisfied. A quadratic penalty is applied otherwise:

$$G_D(\mathbf{x}, s, t) = P(\|\mathbf{c}_s - \mathbf{c}_t\|, \ell_{min}, \ell_{max}, \ell_{rej}, 2). \quad (5)$$

We empirically set the rejection bound ℓ_{rej} to 5% of the initial valid range.

Overlap. Instead of rejecting configurations with any overlap between objects, we accept a small partial pairwise overlap, though penalized, in order to retain the continuity of the cost function. We measure the overlap between two objects s and t by the intersection of their discretely sampled volumes V_s and V_t and associate them with the penalty:

$$G_P(\mathbf{x}, s, t) = P(V_s \cap V_t, 0, 0, V_{ref}, 2), \quad (6)$$

where V_{ref} is a fixed penetration tolerance, typically set to 5% of the smallest of two volumes. The object's volumetric samples are generated using the method by [KPT99] on the GPU at load time. The method's intrinsic inability to properly voxelize internal cavities, actually works to our advantage here, since removing internal voxels should not reduce the penetration value, which is solely affected by the shell of the objects. $V_s \cap V_t$ is simply calculated by counting the overlapping voxels after any transformations are applied.

Focus and Conversation. As indicated by previous research, one can establish rules for communication and mutual visibility between two occupied objects as well as preferential focus on specific parts of the environment or task. We show here that these constraints can be mapped to a combination of distance and alignment ones, dispensing with the requirement of a separate model.

Focus between objects s, t can be described as an alignment constraint between the front vector of s and the direction towards a target object t :

$$G_F(\mathbf{x}, s, t) = G_A\left(\mathbf{x}, \mathbf{u}_s, \frac{\mathbf{c}_t - \mathbf{c}_s}{\|\mathbf{c}_t - \mathbf{c}_s\|}\right). \quad (7)$$

For pre-determined object groups, we declare which piece is the center of attention and define distance and focus constraints indirectly. Picking the central piece of a furniture cluster as the focal point, is not always the right choice, since many arrangements are asymmetrical by design.

Communication requires that end-points are roughly facing each other with a tolerance ℓ_{conv} . In essence, a communication constraint translates to a mutual focus rule combined with a distance limiter:

$$G_C(\mathbf{x}, s, t) = \frac{1}{3} (G_F(\mathbf{x}, s, t) + G_F(\mathbf{x}, t, s) + G_D(\mathbf{x}, s, t)). \quad (8)$$

The geometric term of the objective function $G(\mathbf{x})$ is the weighted and normalized sum of all the individual penalty terms $G_A(\cdot)$, $G_D(\cdot)$, $G_P(\cdot)$, $G_F(\cdot)$ and $G_C(\cdot)$ over the respective sets of constraints.

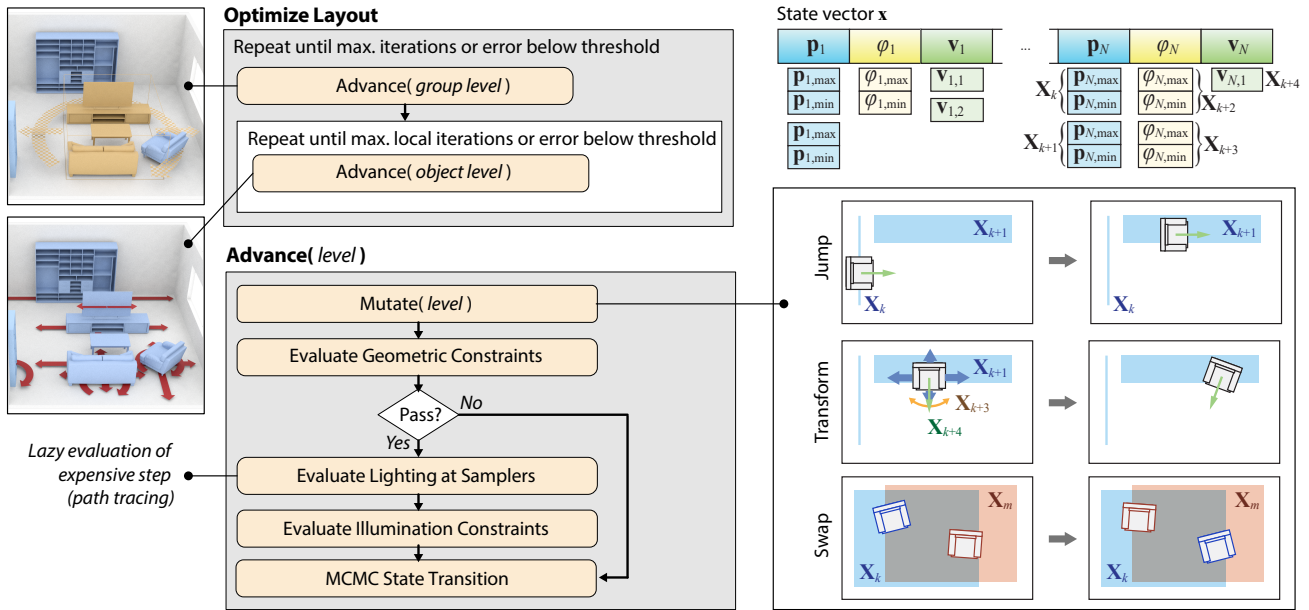


Figure 5: Overview of the two-level layout optimization process.

6. Optimization

The furniture layout optimization problem is characterized by a large parameter space. Additionally, Eq. 1 is a fairly expensive, highly discontinuous objective function, which should be evaluated as sparsely as possible. Such cases are known to benefit from stochastic processes as they allow the escape from local minima and balance between exploration and exploitation. We base our optimization process on the Metropolis algorithm, but adopt a mutation strategy tailored to the particular nature of our segmented parameter domain. We also use a nested optimization scheme, reflecting the inherent hierarchical object organization and coherently mutate states of dependent objects. Finally, lighting is evaluated only for states that satisfy the geometric constraints. Geometric error functions are evaluated in the CPU due to their low overhead, in contrast to light samplers, which are evaluated using path tracing entirely on the GPU for optimal performance. The entire optimization procedure is summarized in Figure 5 - left.

Within each level, optimization proceeds using one of the following events chosen with equal probability: translation, rotation, position swap and parameter interval jump (see examples in Fig. 5 - right). The process is initialized with group- and object-level transformation parameters being drawn uniformly from their respective intervals. The state transition from the current parameter vector \mathbf{x} is performed according to the following steps:

- Choose an event type.
- If the event is a swap operation, it also affects another entity in the same hierarchy branch, which is also marked as having performed a swap operation, simultaneously. The swap is performed only if the switched positions still fall within a valid interval for both entities. Swapping is performed only for positions, since mutually exchanging the orientation of two entities is not a meaningful operation for furniture layout.

- In case of a translation or rotation, first, uniformly select an entity and simultaneously mutate the corresponding transformation parameters x_i, \dots, x_j , e.g. its x, y, z translation offsets. For continuous parameter domains, generate a new sample $X_i^{(k)}$ for each parameter x_i according to the Normal distribution $\mathcal{N}(X_i^{(k)} | X_i^{(k-1)}, \sigma_i)$, rejecting parameters outside the interval. The standard deviation σ_i effectively dictates the exploration rate in each parameter interval. Discretized parameter domains, e.g. discrete possible orientations, are sampled uniformly.
- If the parameter space consists of multiple intervals and a jump operation has been selected, move to a new, uniformly selected interval and initialize the translation and rotation parameters with uniformly selected values within that interval.

σ_i is empirically set to a percentage of the interval's size (25% in our implementation). This percentage is fixed for all parameter intervals, therefore the normal distribution spread adapts to the interval size.

The above mutation strategy ensures that state transition probabilities are symmetrical, since all mutation events are performed in an exclusive manner, interval jumps use a global uniform distribution, rotations and translations use the (symmetrical) Normal distribution, and a swap operation is reversible by nature.

The algorithm repeats for a fixed number of iterations or until a minimum error threshold is achieved. In our experiments, we typically use from 10^3 to 10^6 iterations depending on the scene's geometric complexity. At each iteration, the cost function is evaluated and the result is accepted or rejected according to the Metropolis step. If an arrangement is found that perfectly satisfies all constraints, that is, has a cost function of 0, the optimization is immediately terminated and the current arrangement is returned.




Figure 6: Effectiveness of hierarchical optimization. Grouping significantly improves convergence, when compared with a single-level arrangement.

7. Method Evaluation

For our experiments we first tested the validity and effectiveness of each one of the lighting samplers in relative isolation, using simple, targeted experiments, as shown in Figure 4. Next, we evaluated the method with increasingly complex scenarios, where multiple degrees of freedom and constraints at both group and object level were present. Throughout the evaluation, we requested feedback from a professional interior designer, whose comments are provided in the relevant experiments, in order to assess the importance, compliance and usefulness of our method. As shown by previous research in lighting design [WSL*19], keeping the professional in the loop is very important for validation.

7.1. Simple Experiments

In the first row of Figure 4, two patch samplers are defined, one on the dining table and one on the bookshelf, targeting ample light for reading (400-600lx). From the interior design perspective, we only require that the table be aligned with the walls and the bookshelf to rest on one of the four walls. Sunlight hits the floor at a characteristic bright patch, which clearly needs to be avoided. Using the planar lighting constraints, both the table and the bookshelf are optimally positioned away from the bright spot. The table set went under the artificial light source (chairs have no own degrees of freedom here) and the bookshelf moved close to the window taking advantage of the natural lighting, thus fulfilling the desired lighting conditions.

In the second row of Figure 4 we solve the glare problem. The position and orientation of an office desk and its computer screen are optimized given two constraints: a) adequate illumination on the work surface (desktop) in the range 300-500lx and b) glare avoidance in the form of a directional sampler facing the computer screen ($< 100\text{ nits}$). No alignment constraints were specified for this particular example. In the resulting configuration, the desk is positioned so that it avoids direct and indirect glare from both the natural and artificial light source.

The third example introduces the volume sampler, complemented by the other types of samplers. A conference table with a ceiling-mounted projector and projection screen needs to be positioned in such a way that it satisfies three lighting conditions: a) low reflected light towards the audience, b) adequate illumination in the space in front of the projection screen for the presenter to be clearly visible and c) comfortable illumination level on the table top. Low

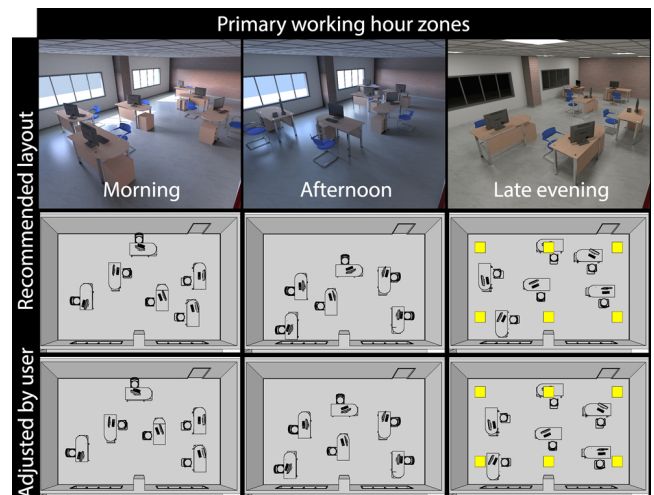


Figure 7: Office space furniture optimization for adequate task lighting and glare minimization in different primary office hours. Morning: strong sunlight - 27sec, 10^4 iterations. Afternoon: indirect and skylight illumination - 33sec, 10^4 iterations. Late evening: artificial lighting - 15sec, 3753 iterations.

reflected light is handled by a directional sampler pointing at the screen. The method positions the screen very close to a wall where only skylight illumination hits the surface and only at oblique angles. The directional sampler also avoids direct glare from placing the screen in front of the windows. The minimum desired illumination level for the presenter is attained by indirect sunlight and direct sky lighting traversing the volume sampler bounds.

The significant impact of hierarchical grouping to the optimization convergence versus a flat object organization is demonstrated in Figure 6. Incorporating all semantically relevant pieces of furniture in groups and allowing permutations at both hierarchical levels, is more intuitive and orders of magnitude faster to converge compared to solely relying on constraints for the same effect.

7.2. Extended Experiments

Computational cost aside, layout optimization, where lighting is sampled over drastically different lighting conditions, spanning the entire day, can only result in a conservative, sub-optimal configuration, which does not respect the typical use of the space. Instead, we concluded that for most practical scenarios, it is more beneficial to focus on a specific time interval corresponding to the space utility (e.g. office hours, event hours).

Figure 7 demonstrates the importance of lighting conditions in an office space layout, where six workstations need to be arranged following both interior and lighting design principles. We optimized the office furniture for three different time intervals, corresponding to three potential primary activity zones for the space. The three cases represent distinct experiments (as it is impractical to change furniture layout during the day). The orientation and position on the desktop of each monitor are both adjustable. The arrangement must provide adequate circulation distances between

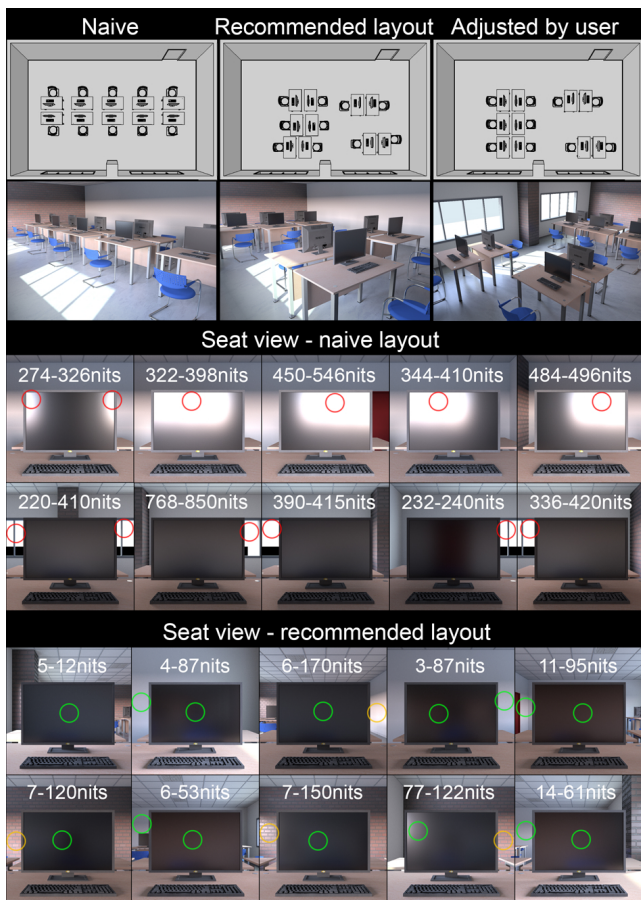


Figure 8: Comparison of a common symmetrical arrangement of workstations in an office space against an illumination-driven layout optimized by our method for glare suppression. The layout after minor aesthetic corrections by an interior designer is included, along with peak glare ranges.

workstations and each monitor should focus on the operator's chair (focus constraints). Desks must be also aligned with the walls. Access to the door is guaranteed using an invisible blocker. Six planar samplers on the desk surfaces set the minimum task illuminance for reading (300-500lx). Additionally, one directional sampler is placed at eye level facing each computer screen, to minimize direct and indirect glare.

The above opposing goals cannot practically be satisfied simultaneously for all scenarios shown, especially in the morning and afternoon intervals, where desktop-level illuminance varies significantly due to artificial lighting being switched off. The morning lighting scenario includes direct sunlight causing very bright illuminance levels near the windows and high glare. In the afternoon scenario, the room is only illuminated by sky lighting and indirect sunlight bouncing off nearby buildings and external structures. Here, the moderate illuminance allows the desks to move closer to the windows in order to take advantage of the natural light, while still avoiding glare. In the nighttime scenario, only the six overhead




Figure 9: Layout optimization with both fixed and movable light sources. Top: the ceiling fixtures alone do not satisfy the reading illuminance levels, which are complemented by stand lights to achieve a more uniform and brighter illumination on table tops - avg. 35sec, 7000 iterations.

panel lights (3200lm each) illuminate the scene. In this example, a more uniform distribution of the furniture is achieved, due to the balanced indoor lighting. Despite the uneven illumination present in the above scenarios, our method manages to adequately comply with the constraints, proposing configurations that, although divergent from typical aesthetics-driven options, are very effective at complying to the desired illumination comfort levels. The interior designer examined the recommended layouts and despite their unconventional setup, after aesthetically adjusting the solutions, positively commented on their effectiveness.

A more regular and tightly-spaced layout is pursued in the computer laboratory example of Figure 8, where workstations are clustered in islands of two desks each. We evaluate the layout recommended by our method against a typical symmetrical arrangement encountered in such a space, confirmed by the interior designer as a generally acceptable one. We optimize the layout with similar constraints to Figure 7 and compare the resulting glare level, measured as average direct and reflected luminance in the directional samplers. The method drives the desks away from areas accessible by direct sunlight and properly orients the desks and screens to avoid glare (the limit was set to 100nits), while still taking advantage of the indirect illumination close to the window to satisfy the desktop illuminance levels.

The experiments in Figures 2 and 9 show a two-room apartment, under different lighting conditions, where we demonstrate a) the simultaneous use of functional and illumination constraints in a tight space, b) the use of sub-space partition jumps needed for furniture pieces to switch room and c) the seamless integration of movable, complementary light sources into the framework. All three types of samplers are present, serving different functional and aesthetic

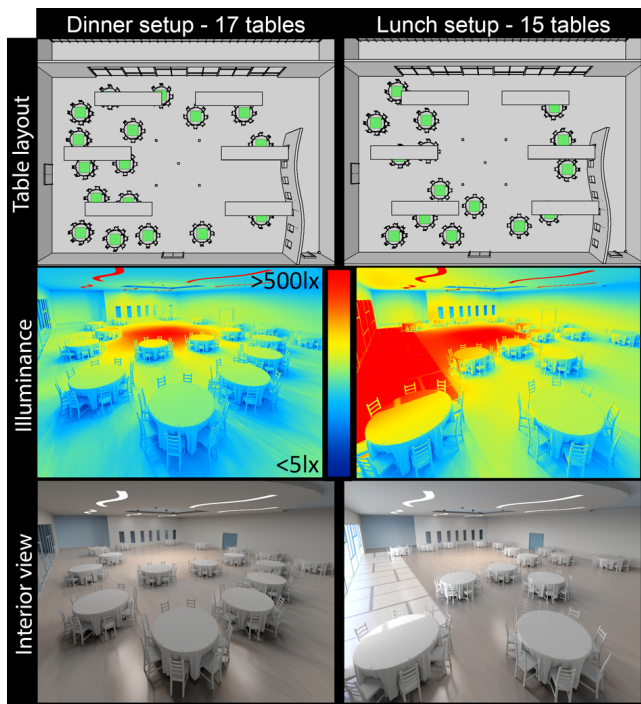


Figure 10: Optimization of a table layout guided by illuminance levels for dining (150-350lx) in different time of day ranges. Illuminance is measured on each table top (highlighted in green). The middle row shows the effective illuminance in the scene, for reference. Execution time varies from 2 mins to 5 mins and iterations are between 10^4 and 10^5 .

goals. Hierarchical optimization has an obvious advantage here, since it allows for functional groups to perform a jump simultaneously and not blindly search for a configuration that encompasses focus and distance constraints. In Figure 9, our method was able to recommend a significantly different layout variation, when additional light sources were introduced, rather than merely perform a local adjustment, achieving illuminance levels closer to the target ones.

The final experiment in Figure 10 presents an event planning application scenario, where the layout is, in practice, manually changed, according to a client's needs, including different number of tables and different time of day for each event. The space is a large social event room with 6 irregularly-shaped composite panel lights for ambient illumination and 5 bright spotlights in the (central) dance floor section. One side of the room has a large glazed section allowing strong natural lighting in at noon. In terms of geometric constraints, the tables need proper spacing for clearance and access to doorways. Clearance is handled by extending the bounds of the table sets and doorways and main walk paths are avoided using invisible blocking geometry. A table and its chairs are treated as a single entity, since for the particular scenario they represent a semantically indivisible object. Each table also has an associated patch sampler with task illuminance set for dining and conversation (150-350lx).

The proposed nighttime scenario in Figure 10-left shows a plausible and realistic setup for the available tables, mainly guided by the lighting goals. The available space around the dance floor is effectively utilized, while tables clear the dance area, without the need to explicitly place any blocking geometry there, due to the high resulting illuminance from the spotlights. The daytime scenario also demonstrates how our framework can still achieve nearly optimal results, despite the significant reduction in usable floor space, due to the presence of large, overbright areas.

The execution time, reported for each example in the respective figure, is obviously affected by the scene complexity, mainly due to the light evaluation. In addition, heavily geometrically-constrained scenarios lead to a high rejection rate, wasting many optimization cycles prior to evaluating the illumination for a valid state. However, this does not necessarily translate to proportional increase in run time, since constraint checks are significantly faster than lighting measurements.

8. Discussion and Future Work

In this work we introduced illumination constraints coupled with physically-based light transport to the furniture layout problem, capable of handling from simple task-related illuminance levels to glare and volumetric ambience. We showed how such illumination constraints can be elegantly coupled with established functional goals and demonstrated the importance of hierarchical optimization in speeding up convergence and easily tackling scenarios with disjoint placement intervals. Our test cases indicate that lighting goals can drastically affect the layout of interior spaces in ways that cannot be defined through geometric constraints alone, nor are they easy to achieve with manual experimentation. At the very least, the proposed layouts can be used by an interior designer as a starting point to further refine the furniture arrangement aesthetically.

One expected limitation of the method is that, whereas scenarios with attainable illumination goals converge quite fast, unrealistic or contradicting lighting intentions may have a significant impact on both the convergence speed and the quality of the solutions. In such cases, the optimizer reaches the maximum number of iterations, oscillating between practically sub-optimal solutions. Another missing feature of the current version is that we do not address any statistical illumination constraints over the planar and volume samplers, such as uniformity or minimum/peak acceptable levels.

From the feedback we received from the interior designer, a general observation was that an automatic method that respects both functional and lighting constraints would nicely complement a professional's pipeline, which, in her case, involves 3D modeling and lighting setup in Autodesk 3DS Max and Chaos Group V-Ray. In scenarios like the apartment in Figures 9 and 2, the designer characterized the configurations as plausible and commended the method for the varied and non-obvious results that respected all constraints. In the ballroom experiment, the designer singled-out the importance of automation in such a large, complex environment. All results were compliant with the target constraints and designer intervention was minimal, at most.

Future directions and improvements include the investigation of machine learning for light field encoding and preference-based layout proposition, e.g. [LLL*19], since such approaches have produced promising results for scene synthesis tasks. To accommodate a wider range of user requirements and different design workflows, a broader use case study would greatly benefit our work. Finally, we are looking forward to extending our method for urban planning tasks, where different sets of design guidelines apply.

Acknowledgements

This work was supported by the Hellenic Foundation for Research and Innovation (HFRI) and the General Secretariat for Research and Technology, <https://doi.org/10.13039/501100003448>, under the HFRI PhD Fellowship grant with GA No. 1545. The monitor and sofa models were downloaded from [BlendSwap](#).

References

- [FB15] FERNÁNDEZ E., BESUIEVSKY G.: Inverse Opening Design with Anisotropic Lighting Incidence. *Computers & Graphics* 47 (2015), 113–122.
- [FCW*17] FU Q., CHEN X., WANG X., WEN S., ZHOU B., FU H.: Adaptive Synthesis of Indoor Scenes via Activity-associated Object Relation Graphs. *ACM Trans. Graph.* 36, 6 (Nov. 2017), 201:1–201:13.
- [FRS*12] FISHER M., RITCHIE D., SAVVA M., FUNKHOUSER T., HANRAHAN P.: Example-based Synthesis of 3D Object Arrangements. *ACM Trans. Graph.* 31, 6 (Nov. 2012), 135:1–135:11.
- [GP16] GKARAVELIS A., PAPAIOANNOU G.: Inverse Lighting Design using a Coverage Optimization Strategy. *The Visual Computer* 32, 6 (Jun 2016), 771–780.
- [GS09] GERMER T., SCHWARZ M.: Procedural Arrangement of Furniture for Real-Time Walkthroughs. *Computer Graphics Forum* 28, 8 (2009), 2068–2078.
- [HWB95] HARADA M., WITKIN A., BARAFF D.: Interactive Physically-based Manipulation of Discrete/Continuous Models. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1995), SIGGRAPH '95, ACM, pp. 199–208.
- [KK17] KÁN P., KAUFMANN H.: Automated Interior Design Using a Genetic Algorithm. In *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology* (New York, NY, USA, 2017), VRST '17, ACM, pp. 25:1–25:10.
- [KK18] KÁN P., KAUFMANN H.: Automatic Furniture Arrangement Using Greedy Cost Minimization. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (March 2018), pp. 491–498.
- [KPC93] KAWAI J. K., PAINTER J. S., COHEN M. F.: Radiotimization: Goal Based Rendering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1993), SIGGRAPH '93, ACM, pp. 147–154.
- [KPT99] KARABASSI E.-A., PAPAIOANNOU G., THEOHARIS T.: A Fast Depth-Buffer-Based Voxelization Algorithm. *Journal of Graphics Tools* 4, 4 (1999), 5–10.
- [KSB17] KARLEN M., SPANGLER C., BENYA J. R.: *Lighting Design Basics, 3rd Edition*. John Wiley & Sons, 2017.
- [LLL*19] LIANG W., LIU J., LANG Y., NING B., YU L.-F.: Functional workspace optimization via learning personal preferences from virtual experiences. *IEEE Transactions on Visualization and Computer Graphics* 25 (2019), 1836–1845.
- [LPX*19] LI M., PATIL A. G., XU K., CHAUDHURI S., KHAN O., SHAMIR A., TU C., CHEN B., COHEN-OR D., ZHANG H.: GRAINS: Generative Recursive Autoencoders for Indoor Scenes. *ACM Trans. Graph.* 38, 2 (Feb. 2019), 12:1–12:16.
- [LSM*18] LI W., SAEEDI S., MCCORMAC J., CLARK R., TZOUMANIKAS D., YE Q., HUANG Y., TANG R., LEUTENEGGER S.: InteriorNet: Mega-scale Multi-sensor Photo-realistic Indoor Scenes Dataset. In *British Machine Vision Conference* (2018).
- [MMP18] MAS A., MARTÍN I., PATOW G.: Heuristic Driven Inverse Reflector Design. *Computers & Graphics* 77 (2018), 1–15.
- [MPF*18] MA R., PATIL A. G., FISHER M., LI M., PIRK S., HUA B.-S., YEUNG S.-K., TONG X., GUIBAS L., ZHANG H.: Language-driven Synthesis of 3D Scenes from Scene Databases. *ACM Trans. Graph.* 37, 6 (2018), 212:1–212:16.
- [MSL*11] MERRELL P., SCHKUFZA E., LI Z., AGRAWALA M., KOLTUN V.: Interactive Furniture Layout Using Interior Design Guidelines. *ACM Trans. Graph.* 30, 4 (July 2011), 87:1–87:10.
- [NHH*06] NAKAJIMA T., HASHIMOTO S., HARUYAMA K., NAKAMURA T., OSANA Y.: Office Layout Support System using Interactive Genetic Algorithm. In *2006 IEEE International Conference on Evolutionary Computation* (July 2006), pp. 56–63.
- [PSS99] PREETHAM A. J., SHIRLEY P., SMITS B.: A Practical Analytic Model for Daylight. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1999), SIGGRAPH '99, ACM Press/Addison-Wesley Publishing Co., pp. 91–100.
- [SW14] SCHWARZ M., WONKA P.: Procedural Design of Exterior Lighting for Buildings with Complex Constraints. *ACM Trans. Graph.* 33, 5 (Sept. 2014), 166:1–166:16.
- [SYZ*17] SONG S., YU F., ZENG A., CHANG A. X., SAVVA M., FUNKHOUSER T.: Semantic Scene Completion from a Single Depth Image. In *2017 IEEE Conference on Computer Vision and Pattern Recognition* (July 2017), pp. 190–198.
- [TW13] TREGENZA P., WILSON M.: *Daylighting: Architecture and Lighting Design*. Routledge, 2013.
- [UnFK13] UREÑA C., FAJARDO M., KING A.: An Area-Preserving Parametrization for Spherical Rectangles. *Computer Graphics Forum* 32, 4 (2013), 59–66.
- [War94] WARD G. J.: The RADIANCE Lighting Simulation and Rendering System. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques* (NY, USA, 1994), SIGGRAPH '94, ACM, pp. 459–472.
- [WSCR18] WANG K., SAVVA M., CHANG A. X., RITCHIE D.: Deep Convolutional Priors for Indoor Scene Synthesis. *ACM Trans. Graph.* 37, 4 (2018), 70:1–70:14.
- [WSL*19] WALCH A., SCHWÄRZLER M., LUKSCH C., EISEMANN E., GSCHWANDTNER T.: Lightguider: Guiding interactive lighting design using suggestions, provenance, and quality visualization. *IEEE Transactions on Visualization and Computer Graphics* (10 2019).
- [XMZ*14] XU K., MA R., ZHANG H., ZHU C., SHAMIR A., COHEN-OR D., HUANG H.: Organizing Heterogeneous Scene Collections Through Contextual Focal Points. *ACM Trans. Graph.* 33, 4 (July 2014), 35:1–35:12.
- [XSF02] XU K., STEWART J., FIUME E.: Constraint-Based Automatic Placement for Scene Composition. In *Proceedings of the Graphics Interface 2002 Conference* (May 2002), pp. 25–34.
- [YDY16] YAMAKAWA T., DOBASHI Y., YAMAMOTO T.: Efficient Simulation of Furniture Layout Taking into Account Lighting Environment. In *Proceedings of the 29th International Conference on Computer Animation and Social Agents* (New York, NY, USA, 2016), CASA '16, ACM, pp. 173–179.
- [YYT*11] YU L.-F., YEUNG S.-K., TANG C.-K., TERZOPOULOS D., CHAN T. F., OSHER S. J.: Make It Home: Automatic Optimization of Furniture Arrangement. *ACM Trans. Graph.* 30, 4 (July 2011), 86:1–86:12.
- [ZYM*18] ZHANG Z., YANG Z., MA C., LUO L., HUTH A., VOUGA E., HUANG Q.: Deep Generative Modeling for Scene Synthesis via Hybrid Representations. *ACM Trans. Graph.* (to be published) (2018).