

Make it Home: Automatic Optimization of Furniture Arrangement

Lap-Fai Yu¹ Sai-Kit Yeung¹ Chi-Keung Tang² Demetri Terzopoulos¹ Tony F. Chan² Stanley J. Osher¹
¹University of California, Los Angeles ²Hong Kong University of Science and Technology



Figure 1: Left: Initial layout where furniture pieces are placed arbitrarily. Middle and right: Two synthesized furniture arrangements optimized to satisfy ergonomic criteria, such as unobstructed accessibility and visibility, required of a realistic furniture configuration.

Abstract

We present a system that automatically synthesizes indoor scenes realistically populated by a variety of furniture objects. Given examples of sensibly furnished indoor scenes, our system extracts, in advance, hierarchical and spatial relationships for various furniture objects, encoding them into priors associated with ergonomic factors, such as visibility and accessibility, which are assembled into a cost function whose optimization yields realistic furniture arrangements. To deal with the prohibitively large search space, the cost function is optimized by simulated annealing using a Metropolis-Hastings state search step. We demonstrate that our system can synthesize multiple realistic furniture arrangements and, through a perceptual study, investigate whether there is a significant difference in the perceived functionality of the automatically synthesized results relative to furniture arrangements produced by human designers.

CR Categories: I.3.7 [Computing Methodologies]: Computer Graphics—Three-Dimensional Graphics and Realism;

Keywords: Procedural modeling, interior design, interior generation, interior modeling, virtual reality, stochastic optimization

Links: [DL](#) [PDF](#)

1 Introduction

Whereas in recent years numerous publications have appeared demonstrating the automatic modeling of building exteriors and facades, the automatic generation of realistic indoor configurations has not yet received the attention that it deserves. With the growing popularity of social virtual worlds and massively-multiplayer

online games that feature large quantities of realistic environmental content, automated procedural methods for synthesizing indoor environments are needed, as it would be too tedious and impractical to model every indoor scene manually. Currently, such indoor modeling is usually simplified or even ignored, which severely limits the realism of many virtual environments.

A realistic indoor scene is typically populated by several different kinds of furniture objects, but only a few of the many possible spatial arrangements of these objects are functional and livable. For example, the front of a television or computer screen should not be blocked, since it is supposed to be visible. Furthermore, most of the objects in the scene should be accessible to human inhabitants. On the other hand, one object is often placed on top of another object, such as a vase on a table, so there exists a hierarchical relationship among the two objects if we regard the carrier object as the parent and the supported object as its child.

While the aesthetic and creative process of interior design would best be done by professional interior designers, our goal is to create software capable of automatically generating furniture arrangements for complex indoor scenes that are optimized to respect important ergonomic factors. This technique would be useful in multiplayer online games and other graphics applications requiring fully automatic interior design with a high degree of realism. The system that we present in this paper achieves this goal in two stages:

First, our system extracts spatial relationships on the placement of furniture pieces from user-supplied exemplars of furnished indoor scenes. This step is done only once, in advance. The acquisition of examples and subsequent extraction of spatial relationship should not be costly, given that many virtual worlds feature user-created content and collaborative design. A scene is then initialized with furniture pieces randomly placed at arbitrary positions and orientations. Here, the furniture placement is almost always unlivable, with objects that are wrongly-located (e.g., a bookshelf is placed at the center of the room rather than against a wall) or wrongly-oriented (e.g., a television screen is facing the wall), and furniture is usually blocking pathways between doors.

Given an arbitrary initial arrangement, such as the one shown in Figure 1(left), optimizing a furniture arrangement subject to human ergonomics is not an easy task, since the search space can be prohibitively large. To address this issue, in the second stage, the initial layout will be adjusted iteratively by minimizing a cost function that accounts for factors, such as human-accessibility, visibility, pairwise object relationships, and so forth, wherein the spatial relation-




Figure 2: Examples of furniture arrangement. Left: An unsatisfactory spatial arrangement resulting from the neglect of human ergonomics considerations; note that the furniture objects are packed together near the upper-left corner and are blocking the door. Right: A satisfactory arrangement with realistically positioned furniture objects that are accessible, do not obstruct the door, and include a television that is readily viewable from a well-positioned chair.

ships extracted from the exemplars are encoded as prior cost terms. We demonstrate that the overall cost function can be optimized by simulated annealing with a Metropolis-Hastings state-search step. From the random initial arrangement in Figure 1, the optimization produces the two synthesized example scenes shown in the figure.

We furthermore perform a perceptual study that adopts a subjective forced-choice approach to investigate whether people have a preference based on perceived functionality between our synthesized results and arrangements produced by human designers.

2 Related Work

We will first discuss the scarce existing research on generating functional furniture arrangements and then review other relevant work ranging from floor-plan generation and interior design metrics to applications requiring (rapid) generation of livable indoor scenes, such as virtual worlds and artificial life.

2.1 Furniture Arrangement

Previous systems that generate furniture arrangements [Kjølaas 2000; Akazawa et al. 2005; Germer and Schwarz 2009; Larive et al. 2004; Sanchez et al. 2003] require manual control or intervention, or do not adequately consider ergonomic factors, which makes them susceptible to generating uninhabitable arrangements. Figure 2 shows an unsatisfactory arrangement where ergonomic factors are neglected in the interior design.

To generate a furniture layout, Kjølaas [2000] represents a given room as a nested hierarchy of rectangular templates, which are swapped by eight predetermined mutation functions. Empty boxes are placed in front of doors and windows to represent free space. However, the approach is limited to rectangular rooms, and each template plus the set of corresponding parameters must be carefully designed.

Akazawa et al. [2005] use a semantic database to explicitly store furniture spatial relationships in order to synthesize new arrangements. Our furniture representation is similar in its use of furniture object interrelationships with parent-child hierarchies where each object is represented as a bounding box. Unlike our approach, however, their inter-object contact constraints must be manually specified in the database.

Germer and Schwarz [2009] take a similar approach, regarding each furniture object as an agent seeking to attach itself to a parent object. Since the parent-child relationships of each object must be

manually defined, however, this task will become prohibitive as the number of objects grows. Furthermore, as ergonomic factors such as good accessibility and visibility are disregarded, unrealistic and uninhabitable configurations suffering undesired physical or visual blocking are unavoidable.

2.2 Floor Plan Generation

Relevant floor plans can be generated before planning the layout of furniture pieces. Recently, Merrell et al. [2010] proposed a data-driven method to generate residential building layouts. In their work, 120 examples of architectural programs are used to train a Bayesian network that captures the relationships among different rooms. Given certain user requirements as priors, the Bayesian network can then be used to generate a floor plan, which is iteratively modified to incorporate desirable human factors. The result is a synthesized floor plan for a residential building.

We also note the work of Chun and Lai [1997], which encapsulates architectural design knowledge into an expert system. The system modules can be used to evaluate floor plan and furniture arrangement according to government regulations and interior design guidelines, providing suggestions for changes.

2.3 Interior Design Metrics

Although interior design involves creative solutions that can be fairly subjective, a set of quantifiable design criteria has long been accepted in the industry. Specifically, such criteria determine whether the design is functional and suitable for human inhabitants. Panero and Zelnick [1979] conducted a detailed study on human dimensions and ergonomics, by carefully defining metrics such as height, width, reachable-range, and visibility, which are believed to be conducive to functional and comfortable designs. For example, a television should maintain a certain distance from the normal viewing area (e.g., a sofa) depending on the dimensions of its screen. Viewing from an oblique angle should also be avoided for the sake of the viewer’s comfort. Note that [Ching and Binggeli 2005] describes these human factors as the “prime determinants” of interior design, emphasizing that while average measurements should be used, flexibility should be exercised to satisfy specific user needs. In optimization terms, such guidelines can be interpreted as soft constraints. In [Mitton and Nystuen 2007; Ching and Binggeli 2005], the importance of accessibility in furniture placement is noted, which is a common consideration in decorating rooms with different purposes. Ching and Binggeli [2005] illustrate how pathways connecting doors may affect human movement and interior furnishing. In general, a pathway connecting doors should be a short path that facilitates movement while leaving considerable floor areas for furnishing. The width of a pathway should depend on the habitant’s body width, with possible amendments when designing homes for the physically-challenged to allow wheelchair movement.


2.4 Related Applications

The realistic synthesis of spatial arrangements of objects can tremendously facilitate virtual world modeling. For example, [Shao and Terzopoulos 2007] demonstrated a large-scale virtual model of a train station populated by numerous autonomous pedestrians. The mobile human agents can perceive the environmental objects they encounter and respond to them appropriately. An automatic means for properly placing various different kinds of environmental content in the scene would be useful in this context.

Collaborative design spaces have been used to assist object modeling [Talton et al. 2009], and they are commonly used in con-


Extraction Step

Positive Examples



Relationship Extraction:
1) Spatial
2) Hierarchical
3) Pairwise

Initialization



Optimization Step

Cost-term Definitions

Update Scene by a Step

No

Terminate?

Yes

Synthesis Finished




Figure 3: Overview of our approach.




Figure 4: Left: A television, its bounding box, and six surfaces; Right: A candelabrum on a table; the table is a first-tier object and the candelabrum is a second-tier object.

structuring virtual worlds, such as Second-Life and many massively-multiplayer online games. The first stage of our approach entails the extraction of spatial and hierarchical relationships from positive furnished examples, which is a more practical approach compared to the manual specification of such relationships, especially for scenes where there are hundreds of different kinds of objects.

There are numerous efforts in the modeling and synthesis of cities and building exteriors. In [Chen et al. 2008; Müller et al. 2006; Parish and Müller 2001] a procedural modeling approach was used to realistically synthesize streets, buildings, and cities, by which parameters such as height and age can be specified to guide the synthesis. The result is a realistic city model populated with buildings. Other efforts [Müller et al. 2007; Xiao et al. 2008] also employ image-based approaches to model facades. Such techniques can be adopted in Google Earth, Bing Maps, and other applications that enable users to zoom into street views and navigate the exteriors of texture-mapped buildings in 3D.

3 Furniture Relationship Extraction

Figure 3 provides an overview of our approach, which is divided into two stages: (1) the extraction of spatial, hierarchical, and pairwise relationships from positive examples and (2) the synthesis of novel furniture arrangements through optimization. This section describes the first stage.

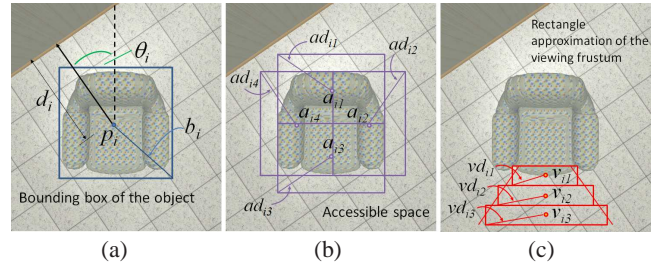


Figure 5: An example object i . (a) Length d_i measures the distance of the object center p_i to its nearest wall. Angle θ_i is the orientation of the object relative to the nearest wall (or the tangent plane if the wall is nonplanar). Length b_i gives the diagonal of the bounding box. (b) The object has 4 accessible spaces centered at a_{i1} , a_{i2} , a_{i3} , and a_{i4} respectively. (c) A viewing frustum associated with the object is represented by 3 rectangles centered at v_{i1} , v_{i2} , v_{i3} . Quantities ad_{ik} and vd_{ik} denote the corresponding diagonal lengths of the rectangles.

3.1 Object Representation

Optimizing furniture arrangement into a realistic and functional indoor configuration involves considerable complexity, taking into account various interacting factors, such as pairwise furniture relationships, spatial relationships with respect to the room, and other human factors. An effective representation that captures the necessary spatial relationships is needed.

Bounding surfaces: Similarly to [Germer and Schwarz 2009; Kjølås 2000], each object in the scene is represented by a set of bounding surfaces (it can be a simple rectangular bounding box or a convex hull to deal with more complex spatial arrangements). Figure 4 shows an example object (television) represented by a bounding box whose six surfaces are labeled 1 to 6. Apart from the top and bottom surfaces, we search for the “back” surface of every object, which is the surface closest to any wall. Other surfaces are labeled as “non-back” surfaces. The back surface is used to define a reference plane for assigning other attributes.

Center and orientation: Figure 5(a) shows the key attributes of an object—center and orientation, denoted by (p_i, θ_i) , where p_i denotes the (x, y) coordinates and θ_i is the angle relative to the nearest wall (defined as the angle between the nearest wall and the back surface). An optimized furniture arrangement $\{(p_i, \theta_i)\}$ involving all objects i is one that minimizes our cost function defined in the next section.

Accessible space: For each surface of the object, we assign a corresponding accessible space (see Figure 5(b)). We define a_{ik} to be the center coordinates of accessible space k of object i . The diagonal of the region is measured by ad_{ik} , which is used to measure how deep other objects can penetrate into the space during optimization. The size of the accessible space is set from available examples or given as input related to the size of a human body. If the space is very close to the wall in all the examples, the corresponding surface need not be accessible; otherwise, we set it to be the dimension of an average-sized adult if such a measurement is not given.

Viewing frustum: For some objects, such as the television and painting, the frontal surface must be visible. We assign a viewing frustum to this particular surface. Given an object i , its viewing frustum is approximated by a series of rectangles with center coordinates v_{ik} , where k is the rectangle index. vd_{ik} is the diagonal of the rectangle, which is useful in defining the penetration cost akin to that for the accessible space. Figure 5(c) provides an example.

Other attributes: Other attributes are involved in the optimization process. Referring again to Figure 5(a), the distance from p_i to its nearest wall is defined as d_i ; the diagonal from p_i to the corner of the bounding box is defined as b_i (the current implementation is a rectangle). We also record the z -position z_i of the object.

Note that to simplify the optimization process, the translation step considers the (x, y) -space only. In other words, an object’s z -position is fixed as the z -position of the surface of its first-tier parent. Nevertheless, the z -position can still change in the swapping step, when a second-tier object changes its first-tier parent and is placed on a different surface. Possible collisions in the z -dimension will still be considered when evaluating accessibility and visibility costs. For example, an overlap between a chair and a bed in the (x, y) space is penalized, while that between a wall clock and a bed is not, as the former involves collision in the z -dimension but the latter does not. Thus, the chair tends to move away from the bed in the (x, y) space, whereas the wall clock does not.

3.2 Learning Prior Relationships

Given the above object representation, the following furniture relationships are extracted automatically from positive input examples.

Spatial relationships: The key prior relationships are the distance of an object to its nearest wall d_i and its relative orientation to the wall $\bar{\theta}_i$. They are respectively estimated as the clustered means of input examples, where we can assign one of the clustered means as \bar{d}_i and $\bar{\theta}_i$ respectively for object i during optimization. The number of clusters can be preset or estimated by [Grunwald 2007].

Hierarchical relationships: Given two objects A and B , object A is defined as the parent of B (and B as the child of A) if A is supporting B by a certain surface. Figure 4 shows a candelabrum on top of a table. The table is hence the parent of the candelabrum, and the candelabrum is the child of the table.

Suppose an example room populated by furniture objects is given. With the room itself regarded as the root, all objects directly supported by the floor or the wall are defined as “first-tier objects” (e.g. bed, table, clock on the wall). All objects supported by a surface of a first-tier object (e.g., a vase on top of a cupboard) are defined as “second-tier objects”. A room configuration is thus represented by a hierarchy of relationships. For simplicity, our optimization considers only first-tiers and second-tiers, which should cover most objects of interest.

Pairwise relationships: Certain objects, such as a television and a sofa or a dining table and chairs, interact with each other in pairs subject to pairwise orientation and distance constraints. Each pairwise relationship can be set by clicking the corresponding objects

in the UI, after which the mean relative distance and angle are extracted from the examples for use as pairwise constraints. Table 1 summarizes the pairwise relationships used for each scene in our experiments in Section 5.

4 Furniture Arrangement Optimization

Given the spatial relationships extracted as described above, our goal is to integrate this information into an optimization framework with a properly defined cost function quantifying the quality of the furniture arrangement. Given an arbitrary room layout populated by furniture objects, the synthesized arrangement should be useful for virtual environment modeling in games and movies, interior design software, and other applications.

The search space of our problem is highly complex as objects are interdependent in the optimization process. The furniture positions and orientations depend on numerous factors, such as whether the object should be visible or accessible. It is very difficult to have a global optimization scheme or a closed-form solution that yields a unique optimum.

To tackle this problem, we resort to stochastic optimization methods, specifically, simulated annealing [Kirkpatrick 1984] with a Metropolis-Hastings state-search step [Metropolis et al. 1953; Hastings 1970] to search for a good approximation to the global optimum. Note, however, that given a room, a set of furniture objects, and the prior spatial and hierarchical relationships, numerous acceptably-good configurations will be possible. This is the rationale for finding a good approximation in a reasonably short time, rather than searching exhaustively over the complex search space in order to find the global optimum of the cost function. The evaluation of interior decoration results can be subjective; hence, we will perform a perceptual study to validate the realism of our synthesized results.

4.1 Simulated Annealing

Simulated annealing is a computational imitation of the (physical) annealing process, which gradually lowers the temperature of a heat bath that controls the thermal dynamics of a solid in order to bring it into a low-energy equilibrium state. Theoretically, the algorithm is guaranteed to reach the global minimum at a logarithmic rate given a sufficiently slow cooling schedule [Geman and Geman 1984]. Using such a slow cooling schedule is impractical, however. Nevertheless, it has been widely used to find quasioptimal configurations in circuit design, operations, and many scientific problems. As in the work on floor-plan generation [Merrell et al. 2010], we found that simulated annealing with the simple Metropolis criterion [Chib and Greenberg 1995] is effective in our problem of optimizing configurations in the space of possible furniture arrangements. For additional details about the simulated annealing method, refer to [Schneider and Kirkpatrick 2006; Liu 2008; Aarts and Korst 1989].

By analogy, the furniture objects in our application are regarded as the atoms of a metal being annealed—they are initially “heated up” to allow flexible rearrangement, and refine their configuration as the temperature gradually decreases to zero. At each temperature, the Metropolis criterion is used to determine the transition probability. It employs a Boltzmann-like objective function

$$f(\phi) = e^{-\beta C(\phi)}, \quad (1)$$

where the state of the system $\phi = \{(p_i, \theta_i) | i = 1, \dots, n\}$ represents a furniture configuration comprising the positions p_i and orientations θ_i of each of the n furniture objects, C is the cost (energy) function, which will be defined in Section 4.3, and β is inversely



Figure 6: Furniture arrangement optimization from a random initial configuration (left). As the optimization process proceeds, the furniture configuration is iteratively updated until it achieves an optimized final arrangement ϕ^* in 25,000 iterations (right).



Figure 7: Left: A pathway connecting doors. Right: A pairwise constraint between the television and the sofa.

proportional to the temperature, increasing over the iterations as the system anneals from a high temperature to a low temperature. At each iteration, a new furniture configuration ϕ' , or “move”, is proposed, and it is accepted with probability

$$\alpha(\phi'|\phi) = \min \left[\frac{f(\phi')}{f(\phi)}, 1 \right] \quad (2)$$

$$= \min \left[\exp(\beta(C(\phi) - C(\phi'))), 1 \right]. \quad (3)$$

Note that the Metropolis criterion can accept moves that increase the cost, which enables the method to avoid becoming stuck at local minima.

Figure 6 depicts an example furniture optimization process. We typically initialize the furniture objects in random positions and orientations, a configuration that typically has very high energy. The supplemental videos include animations of the optimization process.

4.2 Proposed Moves

To explore the space of possible arrangements effectively, the proposed move $\phi \rightarrow \phi'$ involves both local adjustment, which modifies the current arrangement, and a global reconfiguration step that swaps objects, thereby altering the arrangement significantly.

Translation and Rotation: The basic move of the optimization modifies the position of an object and its orientation. For the purposes of the furniture arrangement problem, 2D translation and rotation transformations suffice to configure objects into practicable arrangements, since in most cases furniture objects stand upright on the floor due to gravity. In addition, we found in practice that performing translation and rotation separately gives a more stable optimization. In mathematical terms, an object i or a subset of objects is selected and updated with the move $(p_i, \theta_i) \rightarrow (p_i + \delta p, \theta_i)$ or $(p_i, \theta_i) \rightarrow (p_i, \theta_i + \delta \theta)$, where $\delta p \sim [\mathcal{N}(0, \sigma_p^2) \mathcal{N}(0, \sigma_p^2)]^T$ and $\delta \theta \sim \mathcal{N}(0, \sigma_\theta^2)$, with $\mathcal{N}(\mu, \sigma^2) = (2\pi\sigma^2)^{-1/2} e^{-(x-\mu)^2/2\sigma^2}$ a normal (Gaussian) distribution of mean μ and variance σ^2 . The variances σ_p^2 and σ_θ^2 , which determine the average magnitude of the moves, are proportional to the temperature.

Swapping Objects: To enable a more rapid exploration of the arrangement space and avoid becoming stuck in local minima, a move involving the swapping objects in the existing arrangement may be proposed. Two objects of the same tier are selected at random and their positions and orientations are interchanged: $(p_i, \theta_i) \leftrightarrow (p_j, \theta_j)$ for objects i and j . Object swapping usually changes the cost significantly, thereby leading to considerable rearrangement of the configuration.

Moving Pathway Control Points: Given two doors, multiple pathways are possible. By moving the control points of the pathway, which is represented as a cubic Bezier curve, the pathway can change its course to avoid colliding with furniture objects. As shown in Figure 7, the free space of a pathway is represented by a series of rectangles along the curve. Thus, pathways can also be regarded as “furniture objects” whose control points may be modified, and a move can be defined as the translation of a pathway control point in a certain direction.

With the aforementioned moves, given a floor-plan and a fixed number of furniture objects that define the solution space, the configuration of a furniture object (p_i, θ_i) has a positive probability to move to any other configuration (p'_i, θ'_i) . Given the annealing schedule, the solution space is explored more extensively with larger moves early in the optimization, and the furniture configuration is more finely tuned with smaller moves towards the end.

4.3 Cost Function

The goal of the optimization process is to minimize a cost function that characterizes realistic, functional furniture arrangements. Although it is often difficult to quantify the “realism” or “functionality” of a furniture arrangement, the following basic criteria should not be violated.

Accessibility: A furniture object must be accessible in order to be functional [Mittou and Nystuen 2007; Ching and Binggeli 2005]. In Section 3.1, we defined for every face of an object an accessible space determined from prior examples and the dimensions of the human body (see Figure 5). To favor accessibility, the cost increases whenever any object moves into the accessible space of another object. Suppose object i overlaps with the accessible space k of object j , the accessibility cost is defined as

$$C_a(\phi) = \sum_i \sum_j \sum_k \max \left[0, 1 - \frac{\|p_i - a_{jk}\|}{b_i + ad_{jk}} \right]. \quad (4)$$

Note that we simplify the move by dropping the optimization of orientation θ_i , only measuring the relative distance. Our experiments revealed that this simplification suffices to ensure accessibility and more easily prompts the overlapping object to move away.

Visibility: Some objects, such as a television or a painting, impose strict requirements on the visibility of their frontal surfaces, since their fundamental functionality is compromised if their fronts are blocked by another object. For every such object that must be




Figure 8: The effect on the automatic arrangement (b) of the furniture in (a) resulting from the omission of individual cost terms: Disregarding human ergonomics results in unrealistic synthesized arrangements that are not livable in several ways; e.g., (c) the furniture objects are colliding, (d) a potted plant is blocking the television and the armchair, (e) the work-chair is too far from the desk, (f) the armchair is facing away from the television, (g) the desk and work-chair are blocking the door, (h) furniture objects are too far from the wall, (i) objects are randomly oriented.

visible, we associate with it a viewing frustum (see Figure 5). Similar to the accessibility constraint, whenever another object moves into some object’s viewing frustum, the cost increases in order to discourage the move. As discussed in Section 3.1, for an object j with a viewing frustum we approximate the frustum by a series of rectangles whose center coordinates are defined as v_{jk} . If object i overlaps with the visibility approximation rectangle k of object j , the visibility cost is defined as

$$C_v(\phi) = \sum_i \sum_j \sum_k \max \left[0, 1 - \frac{\|p_i - v_{jk}\|}{b_i + v d_{jk}} \right]. \quad (5)$$

Note that it is similar to the accessibility cost C_a , where the accessible space of object j is replaced by the viewing frustum.

Pathway Connecting Doors: Another important criterion involves pathways between doors [Ching and Binggeli 2005]. The placement of furniture objects such that they block doors should obviously be inhibited. However, a room configuration with circuitous and narrow pathways should also be avoided. To strike a balance, we assume that a pathway in a typical living environment should be smooth, and we define its locus by a cubic Bezier curve, where the free space of the pathway is approximated by a series of rectangular objects, as shown in Figure 7. Thus, the movement of furniture objects into the rectangles is penalized. Apart from moving furni-

ture objects, the pathway itself can be adjusted by translating the control points of the Bezier curve. Because a pathway should be free of obstacles and thus visible, the pathway cost C_{path} can be defined similarly as C , defined in Eq. (5), and applied to the series of rectangles along the pathway.

Prior: The prior cost controls the similarity between the new configuration and configurations seen in the examples. According to Section 3.2, we extract for each furniture object its prior distance and orientation to the nearest wall ($\bar{d}_i, \bar{\theta}_i$). Alternatively, for any new furniture object that is absent from the positive examples, the user can manually assign the prior. Given a new room layout, the current furniture arrangement will be compared with the prior by

$$C_{\text{pr}}^d(\phi) = \sum_i \|d_i - \bar{d}_i\| \quad (6)$$

$$C_{\text{pr}}^\theta(\phi) = \sum_i \|\theta_i - \bar{\theta}_i\|, \quad (7)$$

where d_i and θ_i can be computed from the current p_i ; i.e., finding the distance and relative angle to the nearest wall.

Pairwise Constraint: The pairwise constraint is applied between two furniture objects with a specific pairwise relationship; e.g., the television should be facing the sofa as shown in Figure 7, and a bed-

	Number of Objects	Pairwise Relationships	Number of Iterations	Total Time (sec)
<i>Living Room</i>	20	television & sofa	20000	22
<i>Bedroom</i>	24	television & armchair, desk & work chair	20000	48
<i>Restaurant</i>	54	chair & dish set, chair & table	25000	219
<i>Resort</i>	30	easel & stool, drum & chair, guitar & chair, couch & tea table	42000	126
<i>Factory</i>	51	work desk & chair, supervisor’s desk & chair	42000	262
<i>Flower Shop</i>	64	none	22000	376
<i>Gallery</i>	35	chair & chair	18000	88

Table 1: Computation times are measured on a 3.33GHz Intel Xeon PC. Spatial and hierarchical relationships are extracted automatically from positive examples. Each pairwise relationship can be set by clicking the corresponding objects in the UI, whereupon the mean relative distance and angle are extracted from the examples for use as pairwise constraints.

side table should be close to a bed. It thus encodes the natural affinity of certain furniture objects in the optimized result. We define the pairwise constraint $C_{\text{pair}}^d(\phi)$ and $C_{\text{pair}}^\theta(\phi)$ by simply replacing the distance and orientation to the wall in the prior cost as defined by Eqs. (6) and (7), with the desired distance and orientation between the pair of objects.

Given the above costs, we define the overall cost function as

$$\begin{aligned}
C(\phi) = & w_a C_a(\phi) + w_v C_v(\phi) + w_{\text{path}} C_{\text{path}}(\phi) \\
& + w_{\text{pr}}^d C_{\text{pr}}^d(\phi) + w_{\text{pr}}^\theta C_{\text{pr}}^\theta(\phi) \\
& + w_{\text{pair}}^d C_{\text{pair}}^d(\phi) + w_{\text{pair}}^\theta C_{\text{pair}}^\theta(\phi). \quad (8)
\end{aligned}$$

The w coefficients determine the relative weighting between the cost terms; in practice, we set $w_a = 0.1$, $w_v = 0.01$, $w_{\text{path}} = 0.1$, $w_{\text{pr}}^d = w_{\text{pair}}^d = [1.0, 5.0]$, and $w_{\text{pr}}^\theta = w_{\text{pair}}^\theta = 10.0$. The effect of omitting individual terms is depicted in Figure 8.

The optimization formulation can be readily extended to second-tier objects—optimization is performed to move second-tier objects on the supporting surfaces provided by their first-tier counterparts in the same way that furniture objects move over the floor space of a room, which is regarded as the root in the hierarchy. However, second-tier objects will attach to their first-tier parents if they are not already attached when the optimization begins.

5 Results

Figure 9 shows typical input exemplars that serve in extracting furniture relationships. For each scene, we build five exemplars which cover the most common types of furniture objects. The furniture objects used in the input exemplars for relationship extraction may differ in appearance from those used in the synthesis, although they are of the same type.

To demonstrate the efficacy of our optimization approach, we tested it on seven different scenes, the *Living Room* and *Bedroom* examples shown in earlier figures, plus the five additional scenes, *Factory*, *Flower Shop*, *Gallery*, *Resort*, and *Restaurant*, shown in Figure 10. Table 1 tabulates the computational complexity, running time, number of iterations, and the pairwise relationships used in each scene. Note that the respective positions and orientations of the windows, doors, and ceiling fans are fixed and not updated during the optimization unless otherwise stated.

For each scene, we synthesized three different furniture arrangements; the same view of each synthesized arrangement is shown in Figure 10 for comparison. Two additional views of each synthesized arrangement are included in the perceptual study, which will be detailed in the next section. All the images for the input exemplars and synthesized arrangements can be found in the supplementary materials supporting this paper.

The *Factory* scene in the figure shows the efficacy of the pairwise constraint. By modifying the weights of the pairwise distance and orientation terms, different groupings of work desks and chairs are obtained. The accessibility and visibility constraints acting together prevent the door and poster from being blocked. The *Flower Shop* provides a striking example of the effect of the pathways constraint, which maintains a clear path between the doors despite the dense coverage of the remainder of the room by flowers. For this scene, we change the position of the main door in each synthesis to illustrate different path generation solutions. The accessibility constraint also prevents the cashier from being blocked. We modeled the *Gallery* scene based on an image of the Yale University Art Gallery. The scene consists of a non-rectangular room supported by numerous pillars. Our synthesis result suggests a new interior arrangement for the gallery, where optimizing visibility and accessibility helps avoid obstruction of the pictures and information counter. The *Resort* provides another example of a non-rectangular room. Using pairwise constraints between the easel and the stool, our system automatically generated a area dedicated to painting. We choose a different mean position for the sofa in Synthesis 2 so that the optimized location is farther from the wall. The *Restaurant* example illustrates the significance of the pairwise relationship on both first-tier and second-tier objects. With the use of a concentric spatial relationship between the chairs and table extracted from the exemplars, different numbers of chairs are correctly oriented and evenly distributed around their respective tables and each dish-set is near and properly oriented to its corresponding chair.

6 Perceptual Study

We performed a perceptual study to evaluate the realism and functionality of the furniture arrangements synthesized by our interior design system. Our null hypothesis H_0 was that users perceive no significant differences in the functionality of the synthesized arrangements relative to those produced by a human designer given the same rooms and sets of furniture objects. The alternative hypothesis H_1 was that users did perceive significant differences. Our experiment was conducted using a subjective, two-alternative, forced-choice preference approach patterned after the one reported in [Jimenez et al. 2009].

6.1 Participants

25 volunteer participants were recruited who were unaware of the purpose of the perceptual study. This number of participants was comparable with similar studies in which 16 users were recruited [Jagnow et al. 2008; Jimenez et al. 2009]. The participants included 18 males and 7 females whose ages ranged from 20 to 60. All the subjects reported normal or corrected-to-normal vision with no color-blindness and reported that they are familiar with the indoor scenes to be tested in the study. 14 subjects reported that they did not have any expertise in interior design.



Figure 9: Typical input examples for different scenes.



Figure 10: Selected views of our synthesized results. Top to bottom: Factory, Flower Shop, Gallery, Resort, Restaurant.




Figure 11: Synthesis results obtained without enforcing a selected constraint. Left: No distance constraint; e.g., in the Living Room, the couch is not placed against the wall and, in the Factory, some work-chairs are placed far from their respective work desks. Right: No orientation constraint; e.g., in the Living Room the television is oriented at an awkward angle against the wall and, in the Factory, some work-chairs are oriented arbitrarily.

6.2 Data

The synthesis results shown in Figure 10 were compared against furniture arrangements designed by humans. To assess the significance of priors and pairwise constraints, we produced two additional synthesis results by respectively setting $w_{pr}^d = w_{pair}^d = 0$ and $w_{pr}^\theta = w_{pair}^\theta = 0$. Figure 11 shows selected views of the two additional synthesized examples of the five scenes. Note that the positions of objects mounted on the walls, such as paintings and posters, are fixed in the examples.

6.3 Procedure

The study was conducted in a manner similar to the traditional practice adopted in industry, where interior decorators present their design alternatives to customers and request their preference. It involved static 2D image viewing rather than 3D scene navigation so as to eliminate differences due to varying degrees of skill among the participants in using navigation software. The viewing of video was avoided because, as our preliminary experiments showed, repeated



Figure 12: A screenshot used in our perceptual study. Each participant was shown 70 pairs comprising a synthesized arrangement and an arrangement created by a human designer given the same room and set of furniture objects. Left: Overhead and 2 different views of a synthesized furniture arrangement. Right: Corresponding views of a human-designed furniture arrangement.

video viewing easily causes fatigue.

Figure 12 shows a screenshot used in our perceptual study for pairwise comparison. The left and right color plates respectively show three views of a furniture arrangement, one synthesized by our system and the other created by a human designer. Each participant viewed a total of 70 trials (5 paired comparisons \times 7 scenarios \times 2 trials).

Participants were encouraged to ask any questions prior to the study. After completing a consent form and questionnaire, they were given a sheet indicating the task description:

“This test is about selecting a color plate from a pair of color plates, and there are 70 pairs in total. Each plate shows three views of a furniture arrangement. You will be shown the plates side-by-side with a grey image displayed between each evaluation.

Your task in each evaluation is to select the arrangement in which you would prefer to live, stay, work, visit, etc., depending on the primary function of the room, by clicking on the color plate. You can view the test pair for an unlimited amount of time, but we suggest that you spend around 15 seconds on each set before making your selection.”

The color plates were presented to each participant in a different random order. Counterbalancing was used to avoid any order bias—each paired comparison was assessed twice by each participant, where in half of the trials the synthesized arrangement is displayed as the left plate and as the right plate in the other half.

6.4 Outcome and Analysis

The primary goal of the experiment was to validate the quality of the furniture arrangements synthesized by our system relative to that of arrangements designed by humans. If human-designed arrangements are not clear winners over the synthetic ones, then our system may be considered successful.

The collected preference outcomes were analyzed to determine if any statistically significant trend exists. To this end, we first adopted the Chi-square nonparametric analysis technique. A one-sample Chi-square includes only one dimension, such as is the case in our perceptual study. The obtained $(A_E/A_1, A_E/A_2, A_E/A_3, A_E/A_4, A_E/A_5)$ frequencies were compared to an expected 25/25 (50 for each comparison) result to ascertain whether this difference is significant. The Chi-square values were computed and then tested for significance. Table 2 tabulates the survey results. Overall, they indicate that the furniture arrange-

Scene	A_E/A_1		A_E/A_2		A_E/A_3		A_E/A_4		A_E/A_5	
	χ^2 -value	p -value	χ^2 -value	p -value	χ^2 -value	p -value	χ^2 -value	p -value	χ^2 -value	p -value
<i>Living Room</i>	1.210	0.271	0.010	0.920	0.010	0.920	5.290	0.021	10.89	0.001
<i>Bedroom</i>	0.810	0.368	7.290	0.007	0.010	0.920	4.410	0.036	20.09	0.000
<i>Factory</i>	0.490	0.484	1.690	0.194	0.810	0.368	13.69	0.000	20.25	0.000
<i>Flower Shop</i>	0.090	0.764	9.610	0.002	6.250	0.012	0.090	0.764	10.89	0.001
<i>Gallery</i>	0.250	0.617	3.610	0.057	0.090	0.764	1.690	0.194	3.610	0.057
<i>Resort</i>	0.010	0.920	2.890	0.089	0.090	0.764	9.610	0.002	12.25	0.000
<i>Restaurant</i>	3.610	0.057	0.250	0.617	1.690	0.194	8.410	0.004	2.890	0.089

Table 2: Chi-square analysis (degrees of freedom = 1, level of significance = 0.05). A_E , A_1 , A_2 , A_3 are, respectively, the example arrangement and synthesis results 1, 2, and 3 in Figure 9 and Figure 10. A_4 and A_5 are the respective synthesis results without distance and orientation considerations. Values shown in boldface indicate significant differences.

Scene	A_E/A_1 odds	A_E/A_2 odds	A_E/A_3 odds	A_E/A_4 odds	A_E/A_5 odds
<i>Living Room</i>	(1.377)	5.506	5.506	0.016	0.000
<i>Bedroom</i>	(2.135)	0.002	5.506	0.042	0.000
<i>Factory</i>	3.050	(0.818)	(2.135)	0.000	0.000
<i>Flower Shop</i>	4.894	0.000	0.005	4.894	0.000
<i>Gallery</i>	4.020	0.102	4.894	(0.818)	0.102
<i>Resort</i>	5.506	0.223	4.894	0.000	0.000
<i>Restaurant</i>	0.102	4.020	(0.818)	0.000	0.223

Table 3: Odds on the null hypothesis H_0 over the alternative hypothesis H_1 . Values shown in boldface favor H_0 , indicating no significant difference; values shown in parentheses are inconclusive; other values favor H_1 .

ments created by humans are *not* clearly preferred over the arrangements A_1 , A_2 , and A_3 , when all the cost terms participate in the optimization that synthesizes the furniture arrangements. For the A_E/A_1 , A_E/A_2 and A_E/A_3 pairs, among the 21 synthesized arrangements, only 3 showed a significant difference ($p < 0.05$) inasmuch as most of the participants were able to identify the human-designed arrangement in these cases.

Second, we adopted a Bayesian analysis [Gallistel 2009; Rouder et al. 2009] to determine whether the number of participants who selected the synthesized layout was what would be expected by chance, or if there was a preference pattern. For each scene, we assumed that the participant had a probability P of picking the human-designed arrangement, and that the results of different trials of the same scene were independent of each other. Based on these assumptions, we used a binomial distribution to model the results, where the only parameter was P . Then H_0 has $P = 0.5$ and H_1 has $P = [0, 1]$. We computed the odds O on H_0 over H_1 . According to [Rouder et al. 2009], $O > 3$ shows evidence favoring H_0 whereas $O < 1/3$ shows evidence favoring H_1 , while other odds values are inconclusive.

Table 3 tabulates the odds computed. For the A_E/A_1 , A_E/A_2 , and A_E/A_3 pairs among the 21 synthesis results, 10 favor H_0 indicating the lack of a significant perceived difference between the furniture arrangements synthesized by our system and the human-designed arrangements, 6 favor H_1 indicating a significant difference, and 5 are inconclusive.

Orientation vs distance: Most users chose the human-designed arrangement when the distance or orientation constraint was inhibited, and it was easier for users to detect the difference when we inhibited the orientation term than when we inhibited the distance term. Omitting orientation constraints yields bad results in practice, which suggests that a greater weight can be applied in penalizing orientation deviation during optimization.

7 Discussion and Future Work

Although our framework espouses optimization as a means of synthesizing realistic furniture arrangements, it provides users the flex-

ibility to control furniture placement that respects furniture functionality and interior design aesthetics. For instance, the pairwise constraint promotes the even distribution of chairs around a circular table in the *Restaurant* example, which is a typical case of radial balance or symmetry [Ching and Binggeli 2005; Malnar and Vodvarka 1992]. Our framework also demonstrates its effectiveness in a “tight fit” scenario, where many functional groupings of furniture (e.g., work desks and chairs) are possible as in the *Factory* example, as well as in a “loose fit” scenario, where the placement is more flexible and furniture types are more diverse, as in the *Resort* example. The framework is also flexible enough to cater to specific needs related to human factors, which may be readily encoded into the accessibility and pathway terms in order to generate livable furniture arrangements. Note that we make the implicit assumption that the perimeter of a room is long enough to accommodate all the furniture objects that ought to be placed against walls. Violating this assumption may lead to local suboptima or failure cases, where in the resulting layout some furniture objects that should be positioned against walls may be placed at a distance from the nearest wall that is already occupied by another object.

Given our automatic tool for synthesizing furniture arrangements plus existing methods for synthesizing floor-plans, buildings, and cities, we can potentially create and model virtual worlds much faster and with much less human effort. The resulting interior enrichment would enhance the level-of-detail and, therefore, the realism of large-scale buildings in virtual worlds, which are becoming increasingly ubiquitous in motion pictures and interactive games.

The future extension of our framework will focus on functional issues, including consideration of interior lighting design and the acoustic qualities of a synthesized furniture configuration, as well as on subjective, aesthetic issues, among them the selection of furniture styles and colors consistent with design concepts such as balance, harmony, and emphasis [Ching and Binggeli 2005; Malnar and Vodvarka 1992]. Furthermore, pairwise relationships between objects in exemplars should be analyzed in a more sophisticated manner from the perspective of unsupervised machine learning.

8 Conclusion

We have introduced a framework for the automatic synthesis of furniture layouts, avoiding manual or semi-automated interior design approaches that are impractical in graphics applications requiring full automation. We believe that our work is the first to consider in a comprehensive manner human factors, among them accessibility, visibility, pathway constraints, and so forth. We have demonstrated the effectiveness of our automated interior design approach in generating arrangements for various scenarios, and our results have been deemed by human observers to be perceptually valid in functionality compared to arrangements generated by human designers.

Acknowledgements

We are grateful to the anonymous reviewers for their constructive comments. The Google 3D Warehouse provided us with a rich source of 3D models, which helped us tremendously in speeding up our modeling process. Our special thanks go to Shawn Singh for his expert advice and efforts in video editing, scene-modeling, and rendering, to Yu-Wing Tai for his helpful suggestions in running experiments and on a draft of this paper, to Michael S. Brown for narrating our demonstration video, to Howard Alexander Greene for his advice on video editing, to Jan Adamec for providing a modified version of his Room Arranger software, which enabled our early testing, to Wenzhe Hu for technical advice on stochastic optimization, and to Lap-Fai Lee for his advice on data analysis in our perceptual study. The research reported herein was supported in part by the Hong Kong Research Grant Council under grant no. 620309, RICE/MURI Award 443948-SN-80050, NSF Award Nos. 443948-TH-22487 and 0905671, and ONR grant N00014-09-1-0105. Lap-Fai Yu is supported by a Sir Edward Youde Memorial Fellowship from the Government of Hong Kong.

References

- AARTS, E., AND KORST, J. 1989. *Simulated Annealing and Boltzmann Machines: A Stochastic Approach to Combinatorial Optimization and Neural Computing*. Wiley, New York, NY.
- AKAZAWA, Y., OKADA, Y., AND NIJIMA, K. 2005. Automatic 3D scene generation based on contact constraints. In *Proc. Conf. on Computer Graphics and Artificial Intelligence*, 593–598.
- CHEN, G., ESCH, G., WONKA, P., MÜLLER, P., AND ZHANG, E. 2008. Interactive procedural street modeling. *ACM Trans. Graph.* 27, 3.
- CHIB, S., AND GREENBERG, E. 1995. Understanding the Metropolis-Hastings algorithm. *The American Statistician* 49, 4, 327–335.
- CHING, F. D., AND BINGGELI, C. 2005. *Interior Design Illustrated, 2nd Ed.* Wiley, New York, NY.
- CHUN, H. W., AND LAI, E. M.-K. 1997. Intelligent critic system for architectural design. *IEEE Trans. Knowl. Data Eng.* 9, 625–639.
- GALLISTEL, C. R. 2009. The importance of proving the null. *Psychological Review* 116, 2, 439–453.
- GEMAN, S., AND GEMAN, D. 1984. Stochastic Relaxation, Gibbs Distributions and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6, 6, 721–741.
- GERMER, T., AND SCHWARZ, M. 2009. Procedural arrangement of furniture for real-time walkthroughs. *Computer Graphics Forum* 28, 8, 2068–2078.
- GRUNWALD, P. 2007. *The Minimum Description Length Principle*. MIT Press.
- HASTINGS, W. 1970. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 97–109.
- JAGNOW, R., DORSEY, J., AND RUSHMEIER, H. 2008. Evaluation of methods for approximating shapes used to synthesize 3D solid textures. *ACM Trans. Appl. Percept.* 4, 4, 1–27.
- JIMENEZ, J., SUNDSTEDT, V., AND GUTIERREZ, D. 2009. Screen-space perceptual rendering of human skin. *ACM Trans. Appl. Percept.* 6, 4, 1–15.
- KIRKPATRICK, S. 1984. Optimization by Simulated Annealing: Quantitative Studies. *Journal of Statistical Physics* 34, 5, 975–986.
- KJØLAAS, K. A. H. 2000. *Automatic Furniture Population of Large Architectural Models*. Master's thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Cambridge, MA.
- LARIVE, M., ROUX, O. L., AND GAILDRAT, V. 2004. Using meta-heuristics for constraint-based 3D objects layout. In *Proc. Conf. on Computer Graphics and Artificial Intelligence*, 11–23.
- LIU, J. S. 2008. *Monte Carlo Strategies in Scientific Computing*. Springer, New York NY.
- MALNAR, J. M., AND VODVARKA, F. 1992. *The Interior Dimension: A Theoretical Approach to Enclosed Space*. Wiley, New York, NY.
- MERRELL, P., SCHKUFZA, E., AND KOLTUN, V. 2010. Computer-generated residential building layouts. *ACM Trans. Graph.* 29, 181:1–181:12.
- METROPOLIS, N., ROSENBLUTH, A. W., ROSENBLUTH, M. N., TELLER, A. H., AND TELLER, E. 1953. Equation of state calculations by fast computing machines. *Journal of Chemical Physics* 21, 1087–1092.
- MITTON, M., AND NYSTUEN, C. 2007. *Residential Interior Design: A Guide to Planning Spaces*. Wiley, New York, NY.
- MÜLLER, P., WONKA, P., HAEGLER, S., ULMER, A., AND VAN GOOL, L. 2006. Procedural modeling of buildings. *ACM Trans. Graph.* 25, 3, 614–623.
- MÜLLER, P., ZENG, G., WONKA, P., AND VAN GOOL, L. 2007. Image-based procedural modeling of facades. *ACM Trans. Graph.* 26, 3, 85.
- PANERO, J., AND ZELNICK, M. 1979. *Human Dimension and Interior Space: A Source Book of Design Reference Standards*. Watson-Guptill, New York, NY.
- PARISH, Y. I. H., AND MÜLLER, P. 2001. Procedural modeling of cities. In *Proc. Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*, ACM, New York, NY, 301–308.
- ROUDER, J. N., SPECKMAN, P. L., SUN, D., MOREY, R. D., AND IVERSON, G. 2009. Bayesian t-tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin and Review* 16, 225–237.
- SANCHEZ, S., ROUX, O., LUGA, H., AND GAILDRAT, V. 2003. Constraint-based 3D-object layout using a genetic algorithm. In *Proc. Conf. on Computer Graphics and Artificial Intelligence*.
- SCHNEIDER, J. J., AND KIRKPATRICK, S. 2006. *Stochastic Optimization (Scientific Computation)*. Springer-Verlag, Berlin.
- SHAO, W., AND TERZOPoulos, D. 2007. Autonomous pedestrians. *Graphical Models* 69, 5-6, 246–274.
- TALTON, J. O., GIBSON, D., YANG, L., HANRAHAN, P., AND KOLTUN, V. 2009. Exploratory modeling with collaborative design spaces. *ACM Transactions on Graphics* 28, 5, 167:1–167:10.
- XIAO, J., FANG, T., TAN, P., ZHAO, P., OFEK, E., AND QUAN, L. 2008. Image-based façade modeling. *ACM Trans. Graph.* 27, 5, 1–10.