

Interactive Furniture Layout Using Interior Design Guidelines

Paul Merrell¹ Eric Schkufza¹ Zeyang Li¹ Maneesh Agrawala² Vladlen Koltun¹

¹Stanford University ²University of California, Berkeley



Figure 1: Interactive furniture layout. For a given layout (left), our system suggests new layouts (middle) that respect the user’s constraints and follow interior design guidelines. The red chair has been fixed in place by the user. One of the suggestions is shown on the right.

Abstract

We present an interactive furniture layout system that assists users by suggesting furniture arrangements that are based on interior design guidelines. Our system incorporates the layout guidelines as terms in a density function and generates layout suggestions by rapidly sampling the density function using a hardware-accelerated Monte Carlo sampler. Our results demonstrate that the suggestion generation functionality measurably increases the quality of furniture arrangements produced by participants with no prior training in interior design.

CR Categories: I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques;

Keywords: furniture arrangement, interior design, layout interfaces, interaction

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

1 Introduction

You are moving into a new home and need to arrange the living room furniture. You have a sofa, armchairs, coffee table, end tables, ottomans, and a media center. What arrangement will create the most comfortable and visually pleasing setting for your home?

Furniture placement is challenging because it requires jointly optimizing a variety of functional and visual criteria. Skilled interior designers follow numerous high-level guidelines in producing furniture layouts [Lyons 2008; Ward 1999]. In a living room for ex-

ample, the furniture should support comfortable conversation, align with prominent features of the space, and collectively form a visually balanced composition. In practice these guidelines are often imprecise and sometimes contradictory. Experienced designers learn to balance the tradeoffs between the guidelines through an iterative trial-and-error process.

Yet most people responsible for furnishing a new home have no training in interior design. They may not be aware of interior design guidelines and they are unlikely to have the tacit knowledge and experience required to optimally balance the tradeoffs. Instead such amateur designers rely on intuitive rules such as pushing large furniture items against the walls. These intuitive rules often lead to functionally ineffective and visually imbalanced arrangements [Lyons 2008]. The resulting furniture layouts “simply don’t look or feel right,” and even worse the amateur designer “can’t pinpoint what the problems are” [Ward 1999].

In this paper, we identify a set of interior design guidelines for furniture layout and develop an interactive system based on these guidelines. In our system, the user begins by specifying the shape of a room and the set of furniture that must be arranged within it. The user then interactively moves furniture pieces. In response, the system suggests a small set of furniture layouts that follow the interior design guidelines. The user can interactively select a suggestion and move any piece of furniture to modify the layout. Thus, the user and computer work together to iteratively evolve the design (Figure 1).

Our approach represents the furniture layout guidelines as terms in a density function and treats manual placement of pieces as subspace constraints. Since the resulting function is highly multimodal, we employ a Markov chain Monte Carlo sampler to suggest optimized layouts. To deal with the substantial computational requirements of stochastic sampling, we use graphics hardware to enable interactive performance.

In summary, our work makes two main contributions. First, we identify and operationalize a set of design guidelines for furniture layout. Second, we develop an interactive system for creating furniture arrangements based on these guidelines. Our results demonstrate that the suggestion generation functionality of our system measurably increases the quality of furniture arrangements produced by users with no prior training in interior design.

1.1 Background

Assisted direct manipulation interfaces have been widely studied in computer graphics, dating back to Ivan Sutherland’s Sketch-Pad [Sutherland 1963]. In the context of architectural design, Harada et al. [1995] describe an interface for creating floorplans. Their interface supports local constraints and invokes discrete local search whenever the user drives the layout into a challenging configuration. Likewise, Michalek and Papalambros [2002] use sequential quadratic programming to optimize an arrangement of rectangles in response to interactive manipulation. These approaches assist the layout of general rectangular arrangements and do not incorporate furniture layout guidelines.

Focusing on furniture layout, Bukowski and Séquin [1995] introduce “object association” constraints that are designed to facilitate direct manipulation of furniture arrangements. For example, the user can constrain a bookshelf to slide along walls without penetration or separation. Xu et al. [2002] present a constraint-based furniture layout system that incorporates pairwise relationships which enforce stability, non-penetration, and alignment. Germer and Schwarz [2009] describe an agent-based procedure for furniture layout. In contrast to these techniques, our approach is based on established layout guidelines employed by practicing interior designers. These guidelines include global criteria such as visual balance, which cannot be expressed as a collection of object associations or pairwise relationships.

Layout problems arise in a number of domains and one common strategy is to use optimization techniques to find a layout that satisfies domain-specific criteria. Researchers have applied this optimization approach to circuit board layout [Sarrafzadeh and Lee 1993], graph layout [Tollis et al. 1998], component layout in product design [Cagan et al. 2002], document layout [Jacobs et al. 2003; Hurst et al. 2009], UI layout [Lok and Feiner 2001; Gajos et al. 2010], label layout [Christensen et al. 1995; Vollick et al. 2007], and architectural floor plan layout [Merrell et al. 2010]. Most of these approaches were developed for off-line layout and do not support direct manipulation or generation of multiple high-quality alternatives.

While many physical and software tools are available for visualizing furniture arrangements [Hendler 1981; Reif 1993; Autodesk 2011], these tools simply alleviate the physical strain of moving furniture pieces to prototype different layouts. The placement of the furniture relies entirely on the user’s expertise, which is often insufficient to produce effective furniture arrangements. Our system assists furniture placement by providing optimized suggestions based on interior design guidelines.

1.2 Overview

Figure 2 gives an overview of interaction with our system. Our interface is inspired by Igarashi and Hughes’ [2001] work on suggestive interfaces. The user begins a layout session by creating a room and populating it with furniture. The furniture items are selected from a library that contains categorized 3D models in canonical orientation. Throughout the session, the user can manipulate the furniture and request suggestions. This is further illustrated in the supplementary video.

The suggestions are generated by sampling a density function, defined over the space of layouts of the current set of furniture in the specified room. The density function is defined using idealized analytical formulations of interior design guidelines, described in Section 2. Sampling is performed with a Markov chain Monte Carlo sampler, parallelized on graphics hardware. The sampler runs in

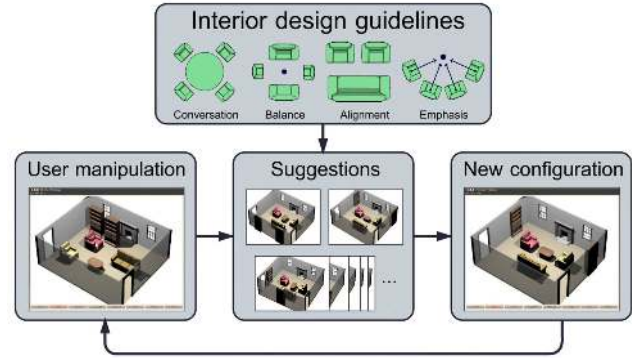


Figure 2: System overview. In response to user manipulation, our system suggests new arrangements that respect user-specified constraints and follow interior design guidelines.

a separate process, so that the user can continue the session while suggestions are being computed.

The suggestions allow the user to quickly experiment with many arrangement options that are already optimized with respect to interior design guidelines. The user can constrain the suggestions by fixing some of the items in place. The constraints simply reduce the dimensionality of the sampled space. This approach allows the user to progressively pin down the desired layout.

2 Furniture Layout Guidelines

Furniture layout design falls under the umbrella of *space planning*, a sub-field of interior design that deals with the allocation of spatial resources [Kilmer and Kilmer 1992; Kubba 2003; Pile 2007; Karlen 2009]. To identify guidelines for furniture layout we have consulted manuals on furniture layout [Talbot 1999; Ward 1999; Sharp 2008] and have interviewed four professional designers who specialize in arranging furniture.

An effective furniture layout must address both functional and visual criteria. The *functional criteria* evaluate how well the layout supports the human activities that take place in the space, such as conversation, rest, or movement. The *visual criteria* concern the perception of the layout as a visual composition.

This section describes some of these criteria and their idealized analytical formulations as terms in a density function. Formally we represent a furniture layout as a tuple $\mathcal{I} = (\mathcal{F}, \mathcal{R}, \mathcal{G})$, where \mathcal{F} is the collection of furniture items placed in the room, \mathcal{R} is a polygon delineating the boundaries of the room, and $\mathcal{G} \subset 2^{\mathcal{F}}$ is a collection of groups of furniture pieces. Such groups can be formed by the user during an interactive layout session with our system.

2.1 Functional Criteria

Functional criteria for furniture layout are based on the constraints imposed by human physiology and the effects of spatial layout on human behavior.

The study of statistical distributions of human physical characteristics, such as body sizes and shapes, is known as anthropometrics. This study establishes guidelines for the necessary clearance around objects and for the proper distances and angles between objects [Panero and Repetto 1975; DeChiara et al. 2001; McGowan and Kruse 2004]. Table 1 lists the anthropometric constraints used in our implementation and Section 2.3 describes our constraint authoring interface.




Figure 3: A real-world furniture layout before (left) and after (right) professional rearrangement (reproduced from [Ward 1999]). On the left, people need to raise their voice to have a conversation when everybody is seated. On the right, conversations are more comfortable and the room has a visual anchor, the fireplace.

Other functional criteria stem from the study of environmental psychology, which examines how the layout of a space affects human activities [Deasy and Lasswell 1985]. Two activities that are affected by furniture layout are conversation and circulation. Conversation is strongly affected by the placement of seats, whose arrangement must support comfortable eye contact and a normal speech volume. Circulation demands that there be sufficient space to comfortably walk to all parts of a room.

Clearance. Many furniture items need open space around them to be accessible and functional. Beds need open space beside them, chairs and bookshelves need space in front of them, and dining room tables need space around their entire perimeter. Table 1(top) lists the clearance constraints used in our implementation. Each constraint specifies the recommended clearance range and direction. To implement these constraints we consider the projection of each furniture item onto the ground plane. We then add desired clearance to the item by taking the Minkowski sum of its projection with a line segment or a disk sized according to Table 1. This defines a set of regions $\mathcal{J}_{\mathcal{F}}$ that delineates furniture items and the clearance around them. We define a clearance violation term $m_{cv}(\mathcal{I})$ that minimizes the overlap between these regions:

$$m_{cv}(\mathcal{I}) = \sum_{\mathbf{f}, \mathbf{g} \in \mathcal{J}_{\mathcal{F}} \cup \{\bar{\mathcal{R}}\}} A(\mathbf{f} \cap \mathbf{g}),$$

where $A(\cdot)$ is an area operator and $\bar{\mathcal{R}}$ is the complement of the room polygon. In our implementation, the regions $\mathcal{J}_{\mathcal{F}}$ are approximated by polygons.

Circulation. An effective furniture layout must support circulation through the room and access to all of the furniture. To evaluate circulation, we adopt the methodology developed for robot motion planning and compute the *free configuration space* of a person on the ground plane of the room [Latombe 1991]. We approximate a person as a disk \mathcal{P} of radius 18" [Panero and Repetto 1975]. We compute the Minkowski sum of \mathcal{P} and the ground plane projection of the collection \mathcal{F} of furniture items as well as the collection \mathcal{W} of wall segments in \mathcal{R} (Figure 5(a)). The free configuration space $\mathcal{C}_{\text{free}}$

Constraint	Distance (in)	Direction
Bedside	36	To the side
Seat	30	In front
Cabinets and shelving	24	In front
Dining table	36	All around
Coffee table to seat	16 - 18	In front of seat
End table to seat	0 - 12	To the back or side of seat
Nightstand to bed	0 - 12	To the side of bed

Table 1: Anthropometric constraints [Panero and Repetto 1975]. Clearance constraints (top) specify recommended amounts of free space around objects. Pairwise constraints (bottom) specify recommended distances and angles between pairs of objects.

is the intersection of the complement of these Minkowski sums:

$$\mathcal{C}_{\text{free}} = \bigcap_{\mathbf{g} \in \mathcal{F} \cup \mathcal{W}} \overline{\mathbf{g} \oplus \mathcal{P}}.$$

The circulation term $m_{ci}(\mathcal{I})$ is defined as the number of connected components in $\mathcal{C}_{\text{free}}$. For efficiency, we approximate this term using graphics hardware, following the approach of Hoff et al. [1999].

Pairwise relationships. Human physiology affects how objects should be positioned with respect to each other. For example, a coffee table should be placed within reach of a seat. Table 1(bottom) lists the pairwise constraints used in our implementation and Section 2.3 describes our authoring interface for specifying additional constraints. We define the pairwise distance for a pair of objects \mathbf{f} and \mathbf{g} as

$$m_{pd}(\mathcal{I}) = - \sum_{\mathbf{f}, \mathbf{g} \in \mathcal{F}} p_{\mathbf{fg}} \cdot t(d(\mathbf{f}, \mathbf{g}), m_{\mathbf{fg}}, M_{\mathbf{fg}}, 2),$$

where $p_{\mathbf{fg}} = 1$ if there is a pairwise constraint between \mathbf{f} and \mathbf{g} and 0 otherwise, $(m_{\mathbf{fg}}, M_{\mathbf{fg}})$ is the range of recommended distances (Table 1(bottom)), and t is a simple objective term defined as

$$t(d, m, M, \alpha) = \begin{cases} \left(\frac{d}{m}\right)^\alpha & d < m \\ 1 & m \leq d \leq M \\ \left(\frac{M}{d}\right)^\alpha & d > M \end{cases}$$

The t function, illustrated in Figure 4, is designed to plateau when d is within the recommended range (m, M) and to gradually decrease as d goes below m or above M . The function has infinite support, to attract the sampler (described in Section 3.1) towards the recommended range regardless of the starting point. The degree of attraction is controlled by the parameter α .




Figure 4: $t(d, 1, 2, 2)$

An analogous pairwise angle term $m_{pa}(\mathcal{I})$ operationalizes the relative direction constraints specified in Table 1(bottom).

Conversation. To support conversation at a normal tone of voice, the seats within a conversation area should be roughly four to eight feet apart [Panero and Repetto 1975]. Our interface allows the user




Figure 5: Notation for interior design guidelines.

to group a collection of furniture items into a conversation area. Given a collection \mathcal{G} of these groups, the conversation distance term is formulated as

$$m_{cd}(\mathcal{I}) = \sum_{S \in \mathcal{G}} \sum_{f, g \in S} q_{fg} \cdot t(d(\mathbf{f}, \mathbf{g}), m_c, M_c, 2),$$

where $m_c = 4$ and $M_c = 8$ feet [Panero and Repetto 1975], and $q_{fg} = 1$ if \mathbf{f} and \mathbf{g} are both seats and 0 otherwise.

The seats should also be angled towards each other to encourage eye contact. The conversation angle term is formulated as

$$m_{ca}(\mathcal{I}) = - \sum_{S \in \mathcal{G}} \sum_{f, g \in S} q_{fg} (\cos \phi_{fg} + 1) (\cos \phi_{gf} + 1),$$

where ϕ_{fg} is the angle between object \mathbf{f} and object \mathbf{g} (Figure 5(b)).

2.2 Visual Criteria

The visual criteria concern the perception of the furniture layout as a visual composition [Arnheim 1974; Poore 1976]. The primary visual rules of thumb used by interior designers are visual balance, alignment, and a dominant point of emphasis.

Balance. The most widely known principle of visual composition is visual balance [Arnheim 1974; Poore 1976; Lok et al. 2004]. The principle is to place the mean of the distribution of *visual weight* at the center of the composition. The visual weight of an element is its perceptual saliency. A common assumption is that larger objects carry more visual weight. Since a furniture arrangement is three-dimensional, visual balance refers to the appearance of the arrangement from multiple viewpoints. A common simplification in practice is to evaluate the distribution of visual weight on the ground plane [Lok et al. 2004]. We define the visual balance cost term as

$$m_{vb}(\mathcal{I}) = \left\| \frac{\sum_{f \in \mathcal{F}} A(\mathbf{f})p(\mathbf{f})}{\sum_{f \in \mathcal{F}} A(\mathbf{f})} - c(\mathcal{R}) \right\|,$$

where $A(\mathbf{f})$ and $p(\mathbf{f})$ are the area and position of \mathbf{f} , respectively, and $c(\mathcal{R})$ is the room's centroid (Figure 5(c)).

Alignment. A basic principle of graphic design that has been extensively employed in automated layout applications is alignment [Sutherland 1963; Lok and Feiner 2001; Jacobs et al. 2003; Vollick et al. 2007]. In furniture arrangement, alignment primarily concerns the orientation of the furniture items relative to each other and to the walls of the room. We define the furniture alignment term as

$$m_{fa}(\mathcal{I}) = - \sum_{S \in \mathcal{G}} \sum_{f, g \in S} \cos(4(\theta(\mathbf{f}) - \theta(\mathbf{g}))),$$

where $\theta(\mathbf{f})$ is the angle of item \mathbf{f} relative to a global coordinate frame (Figure 5(d)). This encourages furniture items within groups to be parallel or at right angles. Furniture items should also be aligned with nearby walls. The wall alignment term is defined as

$$m_{wa}(\mathcal{I}) = - \sum_{S \in \mathcal{G}} \sum_{f \in S} \cos(4(\theta(\mathbf{f}) - \theta_w(\mathbf{f}))),$$

where $\theta_w(\mathbf{f})$ returns the angle of the nearest wall segment to \mathbf{f} .

Emphasis. Another principle of visual presentation that plays a role in interior design is emphasis [Kilmer and Kilmer 1992; Kubba 2003; Pile 2007]. It is generally desirable to have a dominant focal point in the interior, so that the eye can rest without suffering competing demands for visual attention. In residential interiors, the focal point is often a prominent object such as a fireplace, a large window, or an entertainment center.

Our system allows the user to form a group \mathcal{S} of furniture pieces and to associate \mathcal{S} with a focal point \mathbf{p}_S . The focal point is emphasized by orienting the furniture items to face it and by arranging the items symmetrically around it. We define the first emphasis term as

$$m_{ef}(\mathcal{I}) = - \sum_{S \in \tilde{\mathcal{G}}} \sum_{g \in S} \cos \phi_{gp_S},$$

where $\tilde{\mathcal{G}}$ is the set of groups that are associated with focal points and ϕ_{gp_S} is the angle of \mathbf{g} with respect to the focal point \mathbf{p}_S (Figure 5(e)). This encourages objects to face the focal point.

The second emphasis term evaluates the symmetry of groups about their focal points [Kilmer and Kilmer 1992; Kubba 2003; Pile 2007]. Let $s(\mathbf{f}, \mathbf{g}, \mathbf{p})$ be a function that measures the degree of symmetry of items \mathbf{f} and \mathbf{g} about a focal point \mathbf{p} . In our implementation this function is defined in terms of the angles and distances of \mathbf{f} and \mathbf{g} to the focal point:

$$s(\mathbf{f}, \mathbf{g}, \mathbf{p}) = \cos(\theta_{\mathbf{p}}(\mathbf{f}) - \theta_{\mathbf{p}}(r_{\mathbf{p}}(\mathbf{g}))) - \gamma d(\mathbf{f}, r_{\mathbf{p}}(\mathbf{g})).$$

Here γ is a coefficient that determines the relative importance of distance and $r_{\mathbf{p}}(\mathbf{g})$ is the reflection of \mathbf{g} across the symmetry axis defined by \mathbf{p} (Figure 5(e)). The symmetry term is defined as

$$m_{sy}(\mathcal{I}) = - \sum_{S \in \tilde{\mathcal{G}}} \sum_{f \in S} \max_{g \in \mathcal{S}_f} s(\mathbf{f}, \mathbf{g}, \mathbf{p}_S),$$

where $\mathcal{S}_f \subset \mathcal{S}$ is the set of furniture items that match \mathbf{f} . For example, seats match other seats, while they do not match tables.

2.3 Authoring

Our system provides an interface for authoring new constraints, which is useful when adding new types of furniture. All pieces

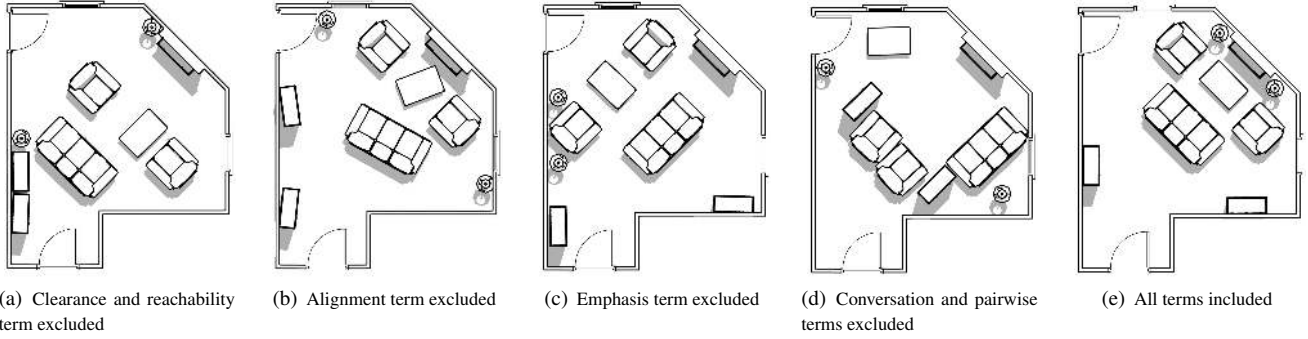


Figure 6: The importance of individual terms in the density function. Lowest-cost samples produced by the sampler when individual cost terms are excluded (a,b,c,d) and when all terms are included (e).

of furniture in the library are categorized by type. New furniture models can be imported and new furniture types can be introduced. Existing constraints can be edited and new constraints can be added. For example, if the set of furniture types is augmented with a billiard table, an appropriate clearance constraint – specifying sufficient clearance for manipulating pool cues – can be added. Similarly, a pairwise constraint for placing filing cabinets near desks can be added. The user can also edit the parameters of each of the constraints, although these parameters were never modified during the evaluation. The authoring interface is demonstrated in the supplementary video.

3 Generating Suggestions

3.1 Density Function and Sampling

We aggregate the objective terms described in Section 2 into a cost function $c_{\mathcal{R},\mathcal{G}}(\mathcal{F})$ that evaluates how much a given furniture arrangement $\mathcal{I} = (\mathcal{F}, \mathcal{R}, \mathcal{G})$ conforms to the interior design guidelines. The cost function is defined as

$$c_{\mathcal{R},\mathcal{G}}(\mathcal{F}) = \sum_i w_i m_i(\mathcal{I}),$$

where $i \in \{cv, ci, pd, pa, cd, ca, vb, fa, wa, sy, ef\}$ and w_i are the mixture weights. In our implementation the weights were set empirically. Our goal is to produce multiple furniture arrangements \mathcal{F} that minimize $c_{\mathcal{R},\mathcal{G}}(\mathcal{F})$, in order to present them as suggestions to the user. (We will omit the subscripts \mathcal{R}, \mathcal{G} in the balance of this section for brevity.) The user can also freeze the placement of specific furniture items. This fixes some of the values of \mathcal{F} and restricts the optimization to the remaining values.

The function c is highly multimodal and is not amenable to exact optimization techniques. We thus employ a Markov chain Monte Carlo sampler to explore the function and produce multiple optimized samples [Gilks et al. 1995]. Specifically, we define a Boltzmann-like density function

$$p(\mathcal{F}) = \frac{1}{Z} \exp(-\beta c(\mathcal{F})),$$

where β is a temperature constant whose importance will be discussed in more detail in Section 3.2 and Z is the partition function that normalizes the distribution. Computing Z is in general intractable. The Metropolis-Hastings algorithm is designed to explore density functions such as p without the need to compute the partition function [Metropolis et al. 1953; Hastings 1970]. The algorithm maintains a current configuration \mathcal{F} and iteratively pro-

poses a modified configuration \mathcal{F}^* that is either accepted or rejected. If the proposal is accepted, \mathcal{F}^* becomes the current configuration. The algorithm iterates until its computational budget is exhausted. All accepted samples are retained for possible use as suggestions.

The acceptance of a proposal move $\mathcal{F} \rightarrow \mathcal{F}^*$ is governed by the Metropolis-Hastings acceptance probability

$$\alpha(\mathcal{F} \rightarrow \mathcal{F}^*) = \min \left(1, \frac{p(\mathcal{F}^*)q(\mathcal{F}|\mathcal{F}^*)}{p(\mathcal{F})q(\mathcal{F}^*|\mathcal{F})} \right), \quad (1)$$

where $q(\mathcal{F}^*|\mathcal{F})$ is the proposal distribution from which a new configuration \mathcal{F}^* is sampled given a current configuration \mathcal{F} . This proposal distribution is key to a successful application of the algorithm. To allow rapid exploration of the density function, the proposal distribution must make both local proposals that make minor modifications to the configuration and global modifications that induce significant rearrangements. Our proposal distribution chooses among three possible proposal moves with equal probability. The first two moves make local adjustments by changing the position or orientation of a single furniture item. The last move swaps two items, which allows for rapid reconfigurations. The proposal moves are as follows:

- Perturb the position of a random furniture item by adding a Gaussian term $\mathcal{N}(0, \sigma^2)$ to each coordinate.
- Perturb the orientation of a random furniture item by $\mathcal{N}(0, \sigma_\theta^2)$.
- Swap the positions and orientations of two randomly selected furniture items.

The standard deviations σ and σ_θ are set empirically. The key advantage of this set of moves is their effectiveness in exploring the space of layouts. Another advantage is that the proposal distribution is symmetric, that is, $q(\mathcal{F}^*|\mathcal{F}) = q(\mathcal{F}|\mathcal{F}^*)$. This allows us to reduce the acceptance probability (1) to the Metropolis ratio, which can be computed directly from the cost function c :

$$\alpha(\mathcal{F} \rightarrow \mathcal{F}^*) = \min \left(1, \frac{p(\mathcal{F}^*)}{p(\mathcal{F})} \right). \quad (2)$$

The algorithm produces a set of sampled layouts. The layouts can be sorted by cost and the lowest-cost ones can be returned as suggestions. However, this can produce very similar suggestions if many of the lowest-cost samples come from a single mode of p . Thus we diversify the list of samples using the Maximal Marginal Relevance criterion, which is commonly used for this purpose in information retrieval [Carbonell and Goldstein 1998].

3.2 Implementation

A key challenge in the application of the Metropolis-Hastings algorithm is its high computational demands. For this reason, applications of global optimization techniques to layout problems have been traditionally restricted to off-line layout [Sarrafzadeh and Lee 1993; Christensen et al. 1995; Cagan et al. 2002]. A recent application of stochastic optimization to architectural layout reported running times of multiple minutes [Merrell et al. 2010]. Such lengthy running times are unacceptable for our system since our goal is to provide responsive suggestions during an interactive layout session. We have thus developed a hardware-accelerated implementation of the sampler by exploiting the massively parallel architecture of graphics hardware [Lee et al. 2010]. Our system generates a variety of optimized suggestions in less than a second, allowing it to update the suggestions immediately in response to direct manipulation operations.

The key to our parallel implementation is the use of a technique called parallel tempering [Geyer 1991]. This technique runs a set of independent Markov chains with different temperature constants β . At a fixed interval two chains are chosen with some probability and propose to swap their current configurations. The proposal is accepted according to the acceptance probability (2). This allows “hot” chains to rapidly explore the configuration space and transfer their state to “cool” chains that can locally optimize within individual modes. Parallel tempering is often used in general applications of MCMC to improve the performance of the sampler without extensive parameter tuning. We use the technique for its additional (and in our case essential) property of parallelizing the sampling process.

Our implementation is illustrated in Figure 7(left). Metropolis-Hastings is performed at the thread block level [Fatahalian and Houston 2008]. The inputs to the algorithm – the room and the given furniture arrangement – are copied to global memory and in turn to each thread block’s shared memory. Within each block, multiple chains with different temperature constants are simulated in parallel by different warps. Exploring chains at the warp level minimizes the need for synchronization, since there is no need for warps to communicate other than to exchange temperature constants. Since warps are able to fit their entire working sets in shared memory, global memory accesses are minimized.

Within each warp, the computation of the density function is performed at the thread level. Each thread is assigned a unique furniture object and is responsible for computing its contribution to the density function. We partition the shared memory between threads to keep memory accesses coherent and avoid memory divergence.

4 Results

We have implemented the presented approach in an interactive furniture layout system. Furniture is populated from a library of 866 furniture models extracted from the Digimation ModelBank Archive. We manually normalized the models to a consistent scale and orientation and organized the library into categories. The library contains chairs, couches, bookshelves, desks, coffee tables, dining tables, lamps, pianos, TVs, and other types of furniture.

In response to a suggestion generation query, the interface chooses the top 36 suggestions returned by the sampler and presents them in groups of 3, as shown in Figure 1. The sampler is run for 10,000 iterations in parallel on each warp. In our experiments, further iterations yielded only very small improvements ($< 0.1\%$) in the density function value, and thus we set 10,000 iterations (per warp) as the computational budget. The hardware-accelerated sampler was implemented using CUDA and tested on an NVIDIA GTX 480 with

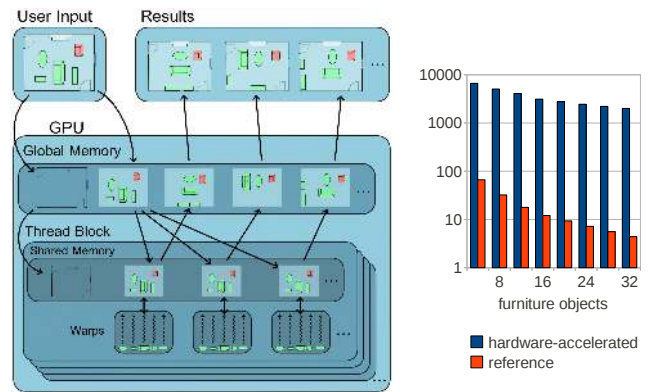


Figure 7: Interactive suggestion generation performance is enabled by exploiting the architecture of graphics hardware (left). Parallel tempered Markov chains are simulated at the warp level, and individual furniture items are assigned to dedicated threads within each warp. The log-scale plot (right) shows the number of density function evaluations performed per second ($\times 1,000$), as a function of the number of furniture items. The hardware-accelerated implementation is over two orders of magnitude faster than the reference sequential implementation.

1.5 GB of global memory. We set the number of thread blocks to 60 and the number of warps to 4. For comparison, we have implemented a sequential version of the sampler, tested on a quad-core 2.66 GHz workstation with 4 GB of memory. The performance comparison is given in Figure 7(right).

Figure 6 demonstrates the importance of individual terms in the density function, as illustrated by the lowest-cost suggestion generated by the sampler for a given configuration when individual terms are ablated. The following weights were used in our implementation: $w_{cv} = 2$, $w_{ci} = 1$, $w_{pd} = 2$, $w_{pa} = 2$, $w_{cd} = 2$, $w_{ca} = 2$, $w_{vd} = 1.5$, $w_{fa} = 2.5$, $w_{wa} = 2.5$, $w_{sy} = 1$, and $w_{ef} = 4$. We experimented with lowering and raising each weight by a factor of two; this mainly affected the order in which the suggested layouts were presented, but not the suggestions themselves. The reported weights were fixed after initial experimentation and were not varied throughout the evaluation.

Informal study. To evaluate the effectiveness of the presented system, we have conducted an informal study with 18 participants with no interior design training. The participants were recruited from the student body of a computer science department in a research university. Each participant was given a five-minute tutorial of the layout interface and was then asked to arrange furniture in five rooms: two living rooms, a game room, a piano room, and an open-wall living and dining room. Some participants did not complete all five arrangements due to time constraints. The initial configurations are shown in Figure 8. Starting from these configurations, the participants were asked to arrange the rooms as they would in their own home. During the study, our system generated suggestions in an average of 0.748 seconds for the room shown in Figure 8(a), 0.734 seconds for 8(b), 0.906 seconds for 8(c), 0.734 seconds for 8(d), and 1.026 seconds for 8(e).

Each arrangement session was performed in one of two randomly chosen conditions: assisted and unassisted. In the assisted condition, the suggestion mode was enabled and the participant could see suggestions as demonstrated in Figure 1 and in the supplementary video. In the unassisted condition, the suggestion mode was disabled. The interface was otherwise identical. In total, 81 ar-

rangements were produced, 39 assisted and 42 unassisted.

In both conditions, participants were free to spend as much time as they desired to complete the task. Layout sessions in the assisted condition lasted 5 minutes and 10 seconds on average. Unassisted sessions were performed in 4 minutes and 50 seconds on average. In the assisted condition, participants were free to use the suggestion generation functionality multiple times. On average, 1.54 suggested layouts were used. Participants could ignore all suggestions, but only did so 12% of the time. After they accepted a suggestion, participants made no further changes to the layout 23% of the time. After accepting the last suggestion, participants made an average of 5 position movements and 4 angle rotations.

The layouts produced by the participants in both conditions were evaluated by two professional interior designers who specialize in furniture arrangement. The evaluation was performed through computerized pairwise comparisons, administered without supervision. The designers were shown pairs of images of layouts produced during the experiment. Each pair contained two layouts of the same room, one created in the assisted condition and one unassisted. The images were presented side by side and their order on the screen was randomized. The conditions were not disclosed to the evaluators, who were told that all layouts were produced with an interface we have developed. The evaluators could click one of three buttons to indicate whether they prefer the layout shown on the left, the one shown on the right, or neither. For each of the five rooms used in the evaluation, all pairs of assisted and unassisted layouts were ranked exactly once. In total, 334 pairwise rankings were performed. The comparison task was partitioned equally between the two designers. The results of the pairwise comparisons are shown in Figure 8. Some of the layouts produced by participants during the study are shown in Figure 9.

5 Discussion

We have presented an interactive furniture layout system that is driven by a set of interior design guidelines. The system can be augmented in a number of ways. First, the layout criteria are currently evaluated over the ground plane and the height of objects in not taken into account; this can be addressed by an easy extension to the analytic formulations of the guidelines. Second, if the room is tightly packed with furniture, with little to no free space, the effectiveness of the sampler is reduced; this can be addressed by a more sophisticated proposal distribution. Third, the presented approach has not been evaluated for non-residential spaces, such as auditoria and conference halls; such spaces may use different layout criteria and different proposal moves. Further, the system can be extended to suggest new furniture that augments the current set, akin to the work of Chaudhuri et al. [2011].

The employed design guidelines were distilled from interviews with practicing interior designers and from published sources. Our operationalization of these guidelines is necessarily simplified. Furthermore, interior designers consider furniture layout alongside the selection of color and material, lighting design, arrangement of art and accessories, carpeting, and detailed traffic patterns through the space. Future work could extend our approach to integrate such criteria, which could be expressed as additional terms in the density function, at the cost of increased dimensionality.

However, we do not advocate an attempt to exhaustively encode the tacit knowledge possessed by professional interior designers. Arrangements produced by different designers exhibit stylistic marks that are recognizable by professionals yet do not seem to be easily amenable to analytical formulations. Such stylistic variability could be learned from data, possibly following the approach of Fisher and

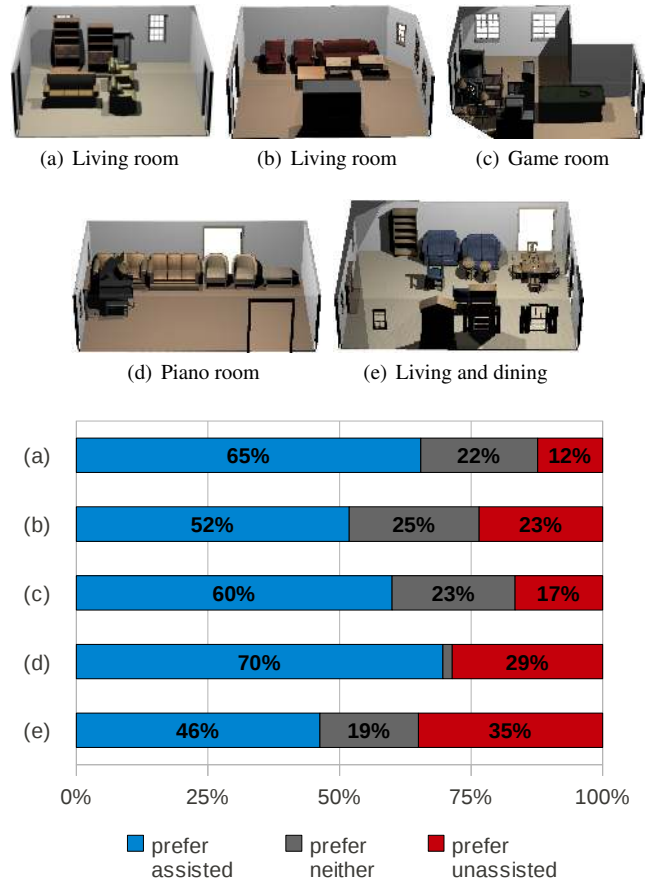


Figure 8: Evaluation outcomes. Eighteen participants arranged furniture in five rooms, starting from initial disorganized configurations (a–e). Layouts produced with and without suggestions were evaluated through blind pairwise comparisons by professional interior designers (bottom). Results for (a–d) indicate a statistically significant preference ($p < 0.01$) for layouts produced with suggestions, according to a two-tailed independent single sample *t*-test.

Hanrahan [2010]. Such data-driven approaches could enable computational representations of a wider scope of functional and aesthetic principles employed by professionals in interior design and other disciplines. Such representations can assist the development of intelligent interfaces for everyday design.

Acknowledgments

We are grateful to Daniel Cohen-Or, Marc Levoy, Sergey Levine, and Philipp Krähenbühl for their comments on this paper. Kristine Moore, Donna Reynolds, Risë Krag, and Jennifer Duchene provided information on furniture layout guidelines. Kristine Moore and Donna Reynolds evaluated layouts produced during the study. This work was supported in part by NSF grants SES-0835601, CCF-0641402, and CCF-0643552.

References

ARNHEIM, R. 1974. *Art and Visual Perception*. University of California Press.




Figure 9: Layouts produced by participants without (left) and with suggestions (right). From top to bottom: Living room a, living room b, game room, piano room, living and dining.

AUTODESK. 2011. *Homestyler*. <http://www.homestyler.com>.

BUKOWSKI, R. W., AND SÉQUIN, C. H. 1995. Object associations: a simple and practical approach to virtual 3d manipulation. In *Proc. Symposium on Interactive 3D graphics*, ACM.

CAGAN, J., SHIMADA, K., AND YIN, S. 2002. A survey of computational approaches to three-dimensional layout problems. *Computer Aided Design* 34, 8, 597–611.

CARBONELL, J., AND GOLDSTEIN, J. 1998. The use of MMR, diversity-based reranking for reordering documents and produc-

ing summaries. In *Proc. SIGIR Conference on Research and Development in Information Retrieval*, ACM.

CHAUDHURI, S., KALOGERAKIS, E., GUIBAS, L., AND KOLTUN, V. 2011. Probabilistic reasoning for assembly-based 3d modeling. In *Proc. SIGGRAPH*, ACM.

CHRISTENSEN, J., MARKS, J., AND SHIEBER, S. 1995. An empirical study of algorithms for point-feature label placement. *ACM Transactions on Graphics* 14, 3, 203–232.

DEASY, C. M., AND LASSWELL, T. E. 1985. *Designing Places for People*. Watson-Guptill.

- DECHIARA, J., PANERO, J., AND ZELNIK, M. 2001. *Time-Saver Standards for Interior Design and Space Planning*, 2nd ed. McGraw-Hill Professional.
- FATAHALIAN, K., AND HOUSTON, M. 2008. A closer look at GPUs. *Communications of the ACM* 51, 10, 50–57.
- FISHER, M., AND HANRAHAN, P. 2010. Context-based search for 3d models. In *Proc. SIGGRAPH Asia*, ACM.
- GAJOS, K. Z., WELD, D. S., AND WOBROCK, J. O. 2010. Automatically generating personalized user interfaces with *Supple. Artificial Intelligence* 174, 12-13, 910–950.
- GERMER, T., AND SCHWARZ, M. 2009. Procedural arrangement of furniture for real-time walkthroughs. *Computer Graphics Forum* 28, 8, 2068–2078.
- GEYER, C. 1991. Markov chain Monte Carlo maximum likelihood. In *Proc. of the 23rd Symposium on the Interface: Computing Science and Statistics*, 156–163.
- GILKS, W., RICHARDSON, S., AND SPIEGELHALTER, D. 1995. *Markov Chain Monte Carlo in Practice*. Chapman & Hall/CRC.
- HARADA, M., WITKIN, A., AND BARAFF, D. 1995. Interactive physically-based manipulation of discrete/continuous models. In *Proc. SIGGRAPH*, ACM.
- HASTINGS, W. K. 1970. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* 57, 1, 97–109.
- HENDLER, M. 1981. *Room and Furniture Layout Kit*. Dover Publications.
- HOFF, III, K. E., KEYSER, J., LIN, M., MANOCHA, D., AND CULVER, T. 1999. Fast computation of generalized Voronoi diagrams using graphics hardware. In *Proc. SIGGRAPH*, ACM.
- HURST, N., LI, W., AND MARRIOTT, K. 2009. Review of automatic document formatting. In *Proc. Symposium on Document Engineering*, ACM.
- IGARASHI, T., AND HUGHES, J. F. 2001. A suggestive interface for 3d drawing. In *Proc. Symposium on User Interface Software and Technology*, ACM.
- JACOBS, C., LI, W., SCHRIER, E., BARGERON, D., AND SALESIN, D. 2003. Adaptive grid-based document layout. *ACM Transactions on Graphics* 22, 3, 838–847.
- KARLEN, M. 2009. *Space Planning Basics*, 3rd ed. John Wiley & Sons.
- KILMER, R., AND KILMER, W. O. 1992. *Designing Interiors*. Thomson Learning.
- KUBBA, S. 2003. *Space Planning for Commercial and Residential Interiors*. McGraw-Hill Professional.
- LATOMBE, J.-C. 1991. *Robot Motion Planning*. Kluwer Academic Publishers.
- LEE, A., YAU, C., GILES, M. B., DOUCET, A., AND HOLMES, C. C. 2010. On the utility of graphics cards to perform massively parallel simulation of advanced Monte Carlo methods. *Journal of Computational and Graphical Statistics* 19, 4, 769–789.
- LOK, S., AND FEINER, S. 2001. A survey of automated layout techniques for information presentation. In *Proc. SmartGraphics*.
- LOK, S., FEINER, S., AND NGAI, G. 2004. Evaluation of visual balance for automated layout. In *Proc. International Conference on Intelligent User Interfaces*, ACM.
- LYONS, G. H. 2008. *Ten Common Home Decorating Mistakes and How To Avoid Them*. Blue Sage Press.
- MCGOWAN, M., AND KRUSE, K. 2004. *Interior Graphic Standards*. John Wiley & Sons.
- MERRELL, P., SCHKUFZA, E., AND KOLTUN, V. 2010. Computer-generated residential building layouts. In *Proc. SIGGRAPH Asia*, ACM.
- 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, 6, 1087–1092.
- MICHALEK, J. J., AND PAPALAMBROS, P. Y. 2002. Interactive design optimization of architectural layouts. *Engineering Optimization* 34, 5, 485–501.
- PANERO, J., AND REPETTO, N. 1975. *Anatomy for Interior Designers*, 3rd ed. Whitney Library of Design.
- PILE, J. F. 2007. *Interior Design*, 4th ed. Prentice Hall.
- POORE, H. R. 1976. *Pictorial Composition*. Dover Publications.
- REIF, D. K. 1993. *Home Quick Planner*. Design Works.
- SARRAFZADEH, M., AND LEE, D. T. 1993. *Algorithmic Aspects of VLSI Layout*. World Scientific.
- SHARP, V. 2008. *The Art of Redesign*. Sharp Redesigns.
- SUTHERLAND, I. E. 1963. *Sketchpad: a man-machine graphical communication system*. PhD thesis, Massachusetts Institute of Technology.
- TALBOTT, C. 1999. *Decorating for Good*. Clarkson Potter.
- TOLLIS, I. G., BATTISTA, G. D., EADES, P., AND TAMASSIA, R. 1998. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice Hall.
- VOLLIK, I., VOGEL, D., AGRAWALA, M., AND HERTZMANN, A. 2007. Specifying label layout style by example. In *Proc. Symposium on User Interface Software and Technology*, ACM.
- WARD, L. 1999. *Use What You Have Decorating*. Penguin Group.
- XU, K., STEWART, A. J., AND FIUME, E. 2002. Constraint-based automatic placement for scene composition. In *Proc. Graphics Interface*.