Inferring Architectural Designs from Physical Sketches: Application to Daylighting Analysis

Barbara Cutler **Theodore C. Yapo** **Yu Sheng**
Department of Computer Science
Rensselaer Polytechnic Institute
{cutler,yapot,shengyu}@cs.rpi.edu

ABSTRACT

We present the algorithms and implementation of an interactive system to capture and interpret an architectural design from a collection of small-scale physical elements. The user sketches a proposed design by arranging 3D wall modules and simple markers for windows, materials, and other design features on a table. The color-coded elements are captured by a camera mounted above the scene and recognized using computer vision techniques. The architectural design is automatically inferred from this rough physical sketch and a closed, 3D triangle mesh representation is constructed. We apply the system to architectural daylighting analysis using an interactive global illumination simulation that allows designers to explore alternative designs and new technologies for improving the sustainability of their buildings. The participants may interactively redesign the geometry and materials of the space by manipulating the physical design elements and visualize the revised lighting simulation.

Author Keywords

Architectural design, physical sketching, daylighting analysis, and image processing.

ACM Classification Keywords

H.5.1 Information Interfaces and Presentation (HCI): Multimedia Information Systems; Artificial, augmented, and virtual realities.

INTRODUCTION

We introduce a system for physically sketching three-dimensional architectural designs. Our motivating application is the analysis and enhancement of daylighting: the use of windows and reflective surfaces to allow natural light from the sun and sky to provide effective and interesting internal illumination. Appropriate daylighting strategies can reduce energy consumption for electric lighting and create more aesthetically interesting and comfortable architectural spaces. A *heliodon* is a traditional daylighting analysis device in which a small-scale physical model (often 1/4" = 1') is affixed to a platform and rotated relative to a fixed light source that represents the sun (Figure 1a). By studying the distribution of light within the model, the designer gains instantaneous and intuitive qualitative feedback on direct sun penetration and the corresponding indirect illumination [4].

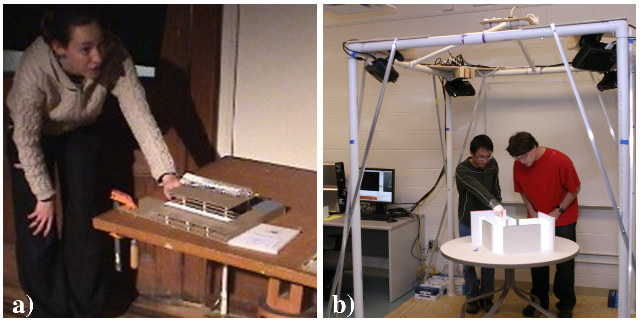

Figure 1. a) Direct light penetration is easily ascertained with the traditional *heliodon*. b) Our *virtual heliodon* allows interactive redesign and enhanced visualization and exploration of the interior spaces.

Our new *virtual heliodon*, shown in Figure 1b), shares many of the features of the traditional heliodon and includes a number of important additional advantages [7]. With a heliodon the user must awkwardly peer through windows (possibly blocking light) or place tiny video cameras within the model; in contrast, the virtual heliodon does not require a physical ceiling on the model, allowing easy inspection of the interior spaces. Furthermore, initial construction and revision of models for the virtual heliodon is faster than for heliodon models because the corners of the model need not precisely align. In the virtual heliodon we "fill in" cracks between walls that in a traditional heliodon would allow light to incorrectly leak into the model. Editing surface materials for a heliodon model is labor intensive and impossible if appropriate scale versions of the material are not available (e.g., Venetian blinds), but with our system surface and window materials can be changed digitally. Our interactive immersive daylighting design system can be used by novice or experienced designers who need not be experts in daylighting technology or advanced graphical simulations.

For many years, virtual reality environments have been dominated by head-mounted displays and CAVE-style environments. Our system builds on research in Spatially Augmented Reality (SAR) [1] for physically-immersive environments that use existing wall and desk surfaces to expand the area of the traditional computer interface [6]. The geometry of the physical objects is known *a priori* and the surfaces are assumed to be a uniform diffuse white material.

IUI 2009 Sketch Recognition, February 8 - 11, 2009, Sanibel Island, FL, USA.

OVERVIEW

Our virtual heliodon system complements modern desktop architectural daylighting design software. The user positions a set of small-scale physical walls within the workspace to "sketch" the 3D geometry of the design. Images captured by a camera mounted above the scene are processed to detect the wall positions. Gaps between the wall are filled to construct a closed 3D mesh. The daylighting solution in the virtual 3D building is computed with respect to both sun and sky illumination. The illumination is then displayed on the physical walls by 4 calibrated projectors.

The efficiency of our interactive global illumination rendering method [2] makes it possible for users to interact with daylighting simulations. The user can explore the high-dimensional configuration space of the design by adjusting the position of wall modules to manipulate the geometry of the building. Through a wireless remote mouse, the user can vary the external conditions such as the sun position (time of day and day of the year) and weather conditions. In addition to visualizing the environment for a single point in time, the system allows time-lapse animations as the sun moves across the sky during the course of a day, throughout the season, or under different weather conditions.

DESIGN ENVIRONMENT

Our system prototype centers around a 42" diameter table with a matte white surface. The table top provides the canvas upon which a designer sketches an architectural design. Our sketching tools consist of a number of *primitives* that the user arranges on the table top to indicate a particular room or building design. The primitives allow for interior and exterior walls, cubicle dividers, different window styles, material properties, and building orientation.

First, the designer manipulates a set of small-scale (1" = 1') planar wall modules with a matte white finish. We provide three different heights: 10' and 8' for walls and 5' for cubicle partitions. The wall tops are covered with strips of colored paper allowing their height to be determined by an overhead camera. Three different styles of window can be placed on any of the planar walls. Window primitives are constructed from simple U-shaped matte white cardstock with their top surface covered in one of three different colors. The window primitives slide over the top of any planar wall to indicate the horizontal placement of the window. Our initial primitives allow the designer a choice of standard, picture (taller than standard), and transom (high, for privacy) window styles. Wide windows may be specified by overlapping two or more window primitives of the same color. As with the wall primitives, the color of the window markers is used to determine their type from images taken by the overhead camera. In addition to windows on walls, the designer may place skylights into the ceiling of the room by placing light blue colored markers on the table surface. The location of these markers is projected upward to the ceiling to indicate the actual skylight position. In addition to the planar walls and dividers, users may indicate curved wall surfaces. We allow curved physical wall primitives that are circular arcs of arbitrary radius and length. More complex curved wall surfaces (an "S"-curve, for example) can be constructed from an appropriate arrangement of circular arcs.

An important consideration for daylighting design is the overall orientation of a room or building with respect to the cardinal geographic directions. Thus, an orange "north arrow" primitive is placed anywhere on the table surface. The direction of the arrow orients the model with respect to the simulated sun and sky conditions. Similarly, the designer exercises control over the surface reflectance properties of the space through a set of material "tokens". A red-outlined token is used to choose the wall material, while a green outline indicates the floor material. A "paint chip" in the center of the token specifies the surface color.

PRIMITIVE DETECTION

To capture the geometry of the design sketch, a calibrated overhead camera acquires images of the table and primitives. Since this single viewpoint precludes accurate photogrammatic determination of object heights, the tops of objects are color-coded to indicate distance above the table surface. The height and window coding colors are carefully selected to allow robust detection in various lighting conditions: red, green, and blue indicate the three wall heights and cyan, magenta, and yellow are used for the different window markers. A nearest-neighbor classifier determines the color class of each pixel in the captured images. The classifier is trained using a target consisting of strips of each color on a museum board background. An image of this target is parsed to find the strips and the background, and a mean RGB value is calculated for each marker color. The background (white) is normalized to compensate for different lighting conditions.

For an image of a new design, we first perform an automatic white balance and exposure compensation. We estimate the illumination on the table surface by averaging the brightest 20% of "white" pixels in the scene, where "white" pixels are determined by thresholding the saturation component of an HSV-transformed image. This method effectively excludes shadows on the table surface from the illumination estimate. We separately calculate averages for the red, green, and blue components of the illumination and divide the image pixels by these values to white balance the image and improve classifier performance under varying illumination. Our initial object detections are performed in image space. We assign each pixel to a color class and perform connected-components analysis to separate the image into discrete objects. Since both wall markers and material tokens share the same red and green colors, material tokens are distinguished by their topological Euler number – material tokens contain a hole, while walls are solid regions.

Next, we estimate the positions of planar wall primitives, which are assumed to be vertical and of the exact, fixed height specified by their color. Since the perspective transformation introduced by the camera preserves straight lines, estimation of wall surfaces can be performed in image space. To estimate a wall's position, we fit a single line to all of the edge pixels of a connected color region. This line, although a poor estimate of either inner or outer wall surface, separates the edge pixels into two sets that can then be individually fit to find both surfaces. In practice, we use an iterative line fitting that rejects points lying more than 2 pixels from the line as outliers; this effectively ignores the effects of pixels lying on the short wall "ends". The length of the wall is then estimated by projecting all colored pixels onto the wall surface lines to find the maximum extents. Finally, using the calibrated camera projection matrix, the 2D wall corner points are back-projected to rays in 3D world coordinates [5]. A simple modification of this scheme improves robustness greatly: each colored wall top stops 1/2" from each end of the wall, allowing for easy separation of physically touching walls during the connected components step; the length of the virtual walls are subsequently extended to match. This scheme allows walls of arbitrary thickness and length and in practice, only a few fixed wall heights are required to model most architectural spaces.

Skylight estimation is similar to that for walls, except that all four edges are required, instead of just the longest two. To estimate skylight edges, we use RANSAC [3] to estimate each edge in turn. Starting with the full set of edge pixels for the object, we use RANSAC combined with a simple robust line fitting algorithm to fit a single line. Pixels which are considered part of this line (their distance falls within a specified threshold) are then removed, and the process is repeated to extract the remaining edges.

While linear objects in the scene can be estimated in image space, pixels corresponding to curved walls must first be back-projected into 3D, since circular arcs are not preserved under perspective transformations. To estimate the curved walls, we first fit a single circle to all edge points of the detected object, and then use the estimated radius to divide the edges into an inner and outer set, similar to our wall estimation algorithm. Once points corresponding to the two arcs are separated, we estimate new radii for each, using a common center point.

To detect the paint chips that specify material properties, we use the camera as a crude spectrophotometer to estimate the RGB values of these colors. To improve color accuracy, the white border around each color swatch is used to perform local white balancing and illumination compensation before an average RGB value for the swatch is computed. Finally, the north indicator arrow is distinguished from other objects in the scene by its orange color. Once classified, the centerline of the arrow is determined by a best-fit line, and one of two possible directions of the arrow chosen by comparing the center of mass of the object to the geometrical center of its extent.

MESH GENERATION

The 3D wall geometry sketched by the user and detected in the previous image processing stage is not a "watertight" model. The walls may have gaps between them and the corners of the physical walls may overlap and extend beyond their natural intersection point. Furthermore, neither ceiling nor windows are explicitly included in the sketch and must be inferred. Thus, the geometry must be processed in order to build a closed 3D triangle mesh model that is appropriate for computing a radiosity global illumination solution. Simple rectangular box rooms may be constructed as the intersection of the halfspaces defined by the four walls. However, this convex hull assumption is not always a suitable assumption for interior architectural design. In particular, it cannot handle the concavities found in L-shaped rooms or rooms with interior wall partitions.

We developed the following algorithm to handle more general room designs. First, all edges of the detected wall and skylight geometry are used to build a 2D arrangement of lines on the plane of the table. Cells of the arrangement that are contained within a wall element are labeled. Next we determine the *percent enclosure* of each of the remaining cells, i.e., what fraction of the 360 degree "view" from the centroid of the cell is obstructed by walls. In a tightly constructed model, with very small gaps between the walls, the percent enclosure for interior cells will be very high. In contrast, when the gaps are larger the percent enclosure value will decrease, especially for cells near the gap. Cells outside of the room will generally have enclosure values much less than 50%. Simply labeling all cells having enclosure values higher than a specified threshold works for many models, but requires careful parameter tuning. Thus, we iteratively select the unlabeled cell with the highest enclosure value and all of the cells that lie in the same combination of halfspaces. Furthermore, rather than halting when an arbitrary percent enclosure threshold is no longer met, the iteration process halts when all walls have been incorporated into the model.

Once all cells have been labeled, we systematically generate a closed, watertight model appropriate for CAD modeling packages and advanced simulations such as physically-accurate rendering. Each interior cell triggers the construction of floor and ceiling triangles. If the cell lies within a skylight, the material of the ceiling triangles is set to glass. The edges of interior cells that border non-interior cells trigger the construction of wall polygons. These wall polygons are subdivided into window polygons as necessary. If the intersection is beyond the edge of a physical wall, the polygons are marked as *fill-in* polygons. Once the closed model is constructed, it is remeshed with a combination of edge split, edge flip, edge collapse, and move vertex operations to arrive at a triangle mesh with approximately 1500 polygons, an appropriate number for interactive radiosity simulations.

RESULTS & FUTURE WORK

From the closed model of the user's design, a physically-accurate daylighting simulation can be calculated (Figure 2) and displayed on the physical model [7]. The gaps and unused portions of the walls are specifically labeled in the model: patches corresponding to the *projection surfaces*, the interpreted *fill-in surfaces* necessary to make a closed model, and the additional unused *physical occlusions* are identified to ensure that the model is correctly rendered. A perspective projection matrix, which represents the projector's intrinsic and extrinsic parameters, is used to render an image of the simulation results for each of the four projectors.

We have presented a physical sketching interface for architectural design. Our small-scale, low-cost physical design environment is practical for academic architectural studios or professional design firms and suitable for installation in a multi-purpose space such as a conference room. This new design tool supports enhanced communication between client and architect and provides a platform for education in sustainable architectural design practice. Preliminary feedback from architecture students about the system has been positive. We plan to conduct formal user studies with our virtual heliodon to compare it with both the traditional heliodon and

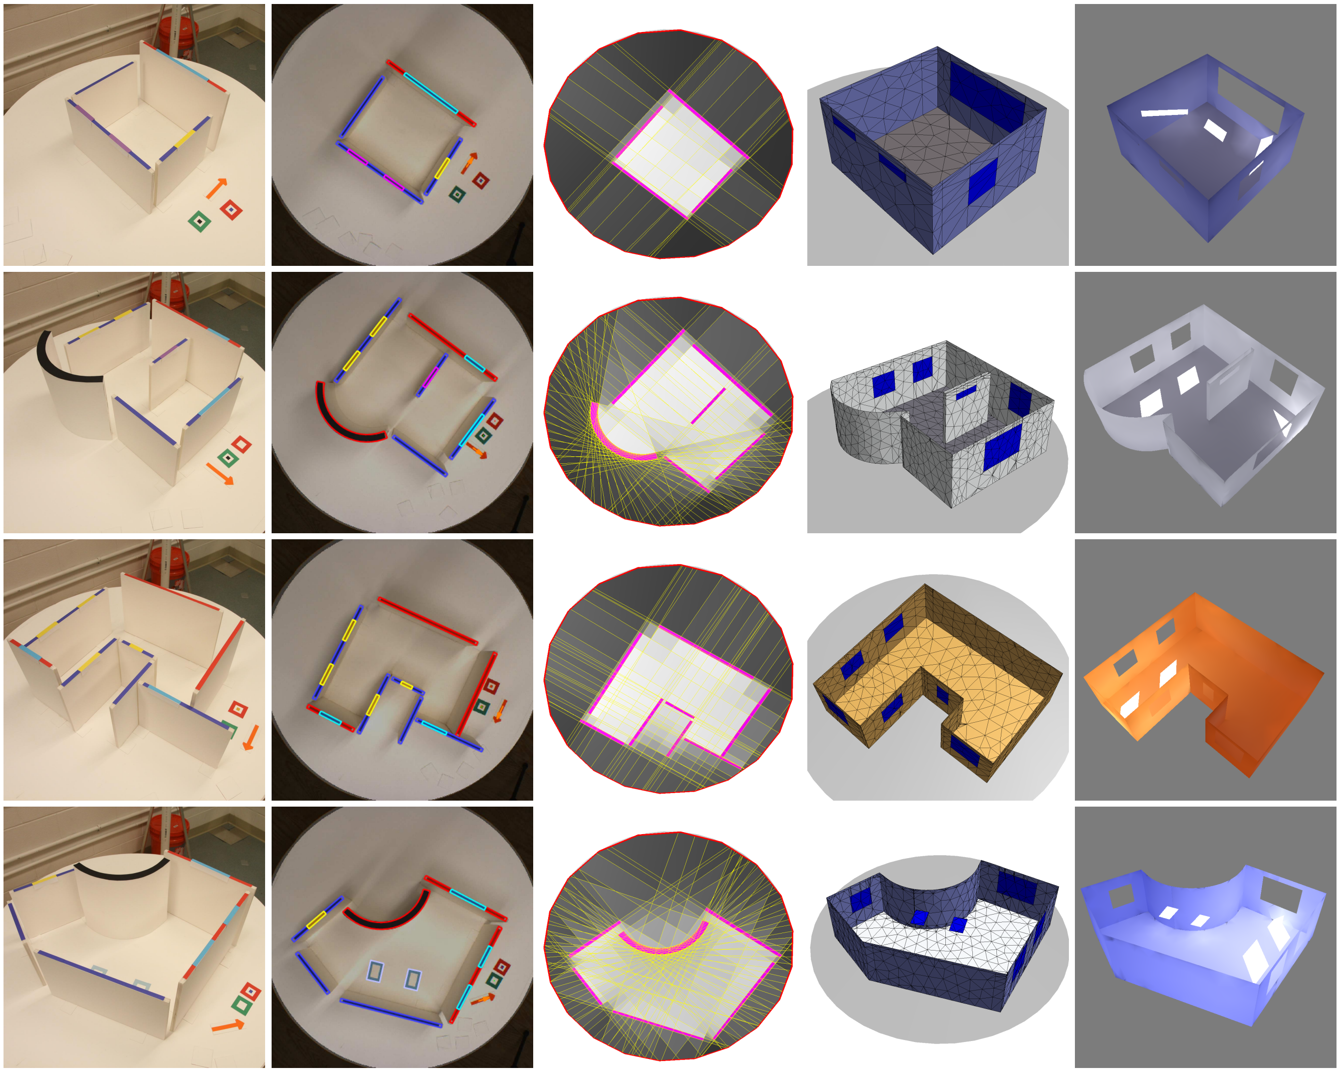

Figure 2. Example designs produced with our system. Left to right: 1) Physical sketch input. 2) Overhead camera image with detected primitives annotated. 3) Arrangement with cells shaded by *percent enclosure*. 4) 3D triangulated inferred design. 5) Rendering of daylighting simulation.

state-of-the-art daylighting analysis software packages. The mechanisms for user interaction in our system can be further enhanced. We envision extending the system with physical control elements or gestures to adjust other aspects of the geometry or simulation. For example, a cloud shape could be used to indicate an overcast sky condition. Or a simple icon placed next to a wall could be used to change the material reflectance properties of the surface or window.

REFERENCES

1. O. Bimber and R. Raskar. *Spatial Augmented Reality: Merging Real and Virtual Worlds*. A K Peters, Ltd., 2005.
2. B. Cutler, Y. Sheng, S. Martin, D. Glaser, and M. Andersen. Interactive selection of optimal fenestration materials for schematic architectural daylighting design. *Automation in Construction*, 17(7):809–823, September 2008.
3. M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981.
4. D. C. Glaser, F. W. Smith, and B. Cutler. Using video for analyzing daylighting simulation tools. In *SimBuild 2006: Building Sustainability and Performance through Simulation*, Aug. 2006.
5. R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521623049, 2000.
6. R. Raskar, G. Welch, K.-L. Low, and D. Bandyopadhyay. Shader lamps: Animating real objects with image-based illumination. In *Rendering Techniques 2001: 12th Eurographics Workshop on Rendering*, pages 89–102, June 2001.
7. Y. Sheng, T. C. Yapo, C. Young, and B. Cutler. Virtual heliodon: Spatially augmented reality for architectural daylighting design. In *Proceedings of IEEE Virtual Reality 2009*, 2009.