Napkin Sketch - Handheld Mixed Reality 3D Sketching

Min Xin*
University of Calgary

Ehud Sharlin†
University of Calgary

Mario Costa Sousa‡
University of Calgary

Abstract

This paper describes, Napkin Sketch, a 3D sketching interface which attempts to support sketch-based artistic expression in 3D, mimicking some of the qualities of conventional sketching medium and tools both in terms of physical properties and interaction experience. A portable tablet PC is used as the sketching platform, and handheld mixed reality techniques are employed to allow 3D sketches to be created on top of a physical napkin. Intuitive manipulation and navigation within the 3D design space is achieved by visually tracking the tablet PC with a camera and mixed reality markers. For artistic expression using sketch input, we improve upon the projective 3D sketching approach with a one stroke sketch plane definition technique. This coupled with the hardware setup produces a fluid sketching experience.

CR Categories: I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques; H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—Artificial, augmented, and virtual realities

Keywords: sketch-based design, 3D design, mixed reality

1 Introduction

Sketching has long been recognized as an essential process of cognition and design [Do and Gross 1997]. It is a staple in various creative fields such as architecture and engineering especially for the preliminary stages of design and ideation. However, despite its importance, it has been difficult to integrate the expressive and freeform "spirit" of sketching in computer design systems. Most designers prefer using conventional sketching tools and mediums such as pencil and paper over computer interfaces for jotting down ideas or exploring new concepts. This is mainly due to the rigidity and heavy weight nature of typical computer-supported design systems both in terms of input methods and interaction styles. Such systems hinder rather than support the creative process because the critical thinking of users can be easily disrupted by having to attend to premature decisions or complex interfaces. However, with advantages such as ease of editing and dynamic 3D visualizations to counteract existing flaws, the challenge is to create a computer-supported design system which offers the same affordances, simplicity, portability, flexibility, and fluidity of conventional sketching tools and mediums. In this paper, we present our attempt at a simple 3D sketching interface, Napkin Sketch (Figure 1), which allows users to create 3D sketches using a tablet PC as the design tool and a piece of ordinary paper as the design medium. Our contributions are twofold with both a light weight hardware setup and its complimentary easy to use software interface.

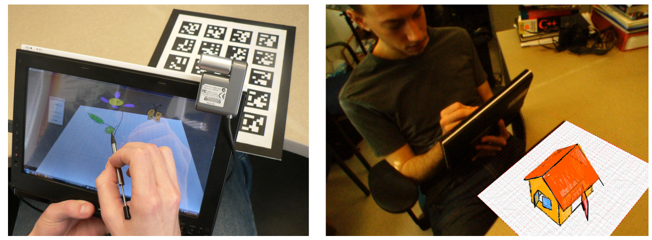

Figure 1: *Creating 3D sketches on top of a physical napkin*

2 Related Work

Computer-supported sketching and 3D design is a broad and well researched area. Topics range from direct 3D design and sketching in space [Sachs et al. 1991; Kiyokawa et al. 1997] to quasi-3D explorations of projected sketches [Tolba et al. 1999]. The challenge of creating intuitive design experiences similar to pencil and paper can be approached from many different directions. Our system makes improvements and advancements in two sub areas: 3D input devices and projective 3D sketching.

One approach to improving the 3D design experience and avoiding the complexity of conventional 2D interfaces for 3D design is to use 3D interfaces and interaction techniques. Various systems allow users to directly create and manipulate sketches and designs in 3D with the help of 3D tracking. Sachs et al. [1991] introduced the concept of "design directly in 3D" with the 3-Draw system which makes use of a pair of handheld 6 degrees of freedom trackers. One tracker is used as a reference for the 3D design space, while the other is used as a stylus. These direct manipulation systems are simple in concept and allow unconstrained interaction and exploration within the 3D design space, but they often require heavy weight and expensive setup and can be difficult to use due to the lack of a clear design medium and haptic feedback. Our system draws from the advantages of direct manipulation achieved through 3D tracking for camera movements but employs a much more light weight and portable solution, using visual tracking and a handheld display. The semi-immersive technique of mixed reality used in our system is also adopted by other works such as Cheok et al.'s [2002] curve and surface modeler. It uses a visually tracked glove and design setting to create and manipulate geometry with hand gestures. Unlike some virtual reality systems, one advantage of mixed reality is its ability to aid the user in spatial understanding while avoiding disorientation since the virtual world is anchored on top of the physical world. Although our system uses mixed reality and 3D tracking for camera manipulation, we have not explored the technique of direct 3D sketching due to the haptic feedback problems. Sketching is traditionally a 2D interaction technique, and moving a stylus in 3D may be unnatural for users. Therefore, we wanted users to always sketch on a 2D surface very much like they would on paper. We are aware of one 3D input device that shares this philosophy. The

*e-mail:mxin@ucalgary.ca
†e-mail:ehud@cpsc.ucalgary.ca
‡e-mail:smcosta@ucalgary.ca

3D Tractus [Lapides et al. 2006] uses a tablet PC placed on top of a mechanical table which can be moved up and down to record the third dimension. By moving the table and drawing with the pen on the tablet, complex non-planar curves can be generated. However, since the 3D Tractus design space can only be explored through axis-aligned vertical volume slices, it is not suitable for arbitrary 3D designs.

The software interface and the dominant sketching interaction experience of our system are closely related to works involving the use of raw sketches as the visual representation of 3D designs. Often, the method for creating these sketches is through the projection of 2D strokes onto 3D surfaces (typically planes). Users can sketch in perspective as if in 2D and still produce the desired 3D representation provided the 2D strokes are projected to the appropriate positions in the 3D design space. Sketchpad+ [Piccolotto 1998] allows users to draw strokes with a pen on a large tilted digital design table and generate 3D sketches by projecting them onto user defined grids in 3D. The positioning and orientation of these grids are specified using standard 3D manipulation operations such as rotation and translation. Another relevant work to our Napkin Sketch system is Mental Canvas [Dorsey et al. 2007]. This system is designed to allow architects to organize concept drawings in 3D. Architects first make several regular 2D sketches of their design from different viewpoints. Later, these are fused together to generate a 3D sketch representation. 3D planes are first defined in the appropriate locations using standard 3D manipulation operations. Selected 2D strokes are then pushed or projected onto these planes. This system does not technically support 3D sketching but rather makes the process of transferring from 2D to 3D more efficient. Although projective 3D sketching provides the closest sketching interaction experience to pencil and paper, a major bottleneck of this approach is the placement of the surfaces or planes which the sketched strokes are projected on. The systems described above make use of standard 3D manipulation operations to position and orient the surfaces and planes in the design space. These operations are often slow and non-trivial, potentially disrupting the users' creative process. Napkin Sketch improves upon this problem with a one stroke sketch plane definition technique to allow users to quickly continue sketching as they switch between projection surfaces.

3 Sketch Pad and Napkin

The systems described in the previous section are primarily implemented for stationary desktop settings. Although some can be ported to run on mobile computing devices like tablet PCs, their interfaces often do not function fluidly on such platforms. Many make heavy use of keyboard controls which would be difficult to access on tablet PCs, and most use standard keyboard and mouse camera manipulation which often requires the use of special modes and extra interface controls to perform rotation and translation in a serial fashion. With portability, fluidity, and simplicity as major goals of Napkin Sketch, we took several unconventional approaches to the design of the hardware setup. We feel a good sketching tool should always be at hand, and we believe a variety of mobile platforms such as tablet PCs, PDAs, camera phones, UMPCs, and portable gaming devices can eventually rival the ubiquity of pencil and paper. Therefore, one obvious design choice is the use of an ultra portable tablet PC which can be easily held in one hand as the sketch pad (Figure 1). What is less orthodox in our design is the use of a physical napkin, printed with mixed reality markers, to represent and anchor the 3D design space. The approach is similar to many handheld mixed reality systems. The tablet PC is equipped with a front facing camera delivering live video of the physical space. It is then used as a mixed reality view port to reveal the virtual design space superimposed on top of the napkin in

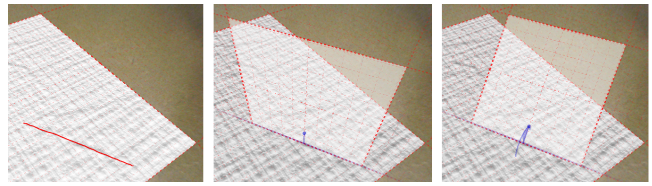

Figure 2: *One stroke technique (left to right: frame creation stroke, created perpendicular frame, and frame rotation)*

the physical environment (Figure 1). Changing the relative positions of sketch pad and napkin changes the viewpoint of the design space, facilitating natural camera manipulation through the direct mapping. For example, to view sketched content in detail, users can zoom in by physically moving closer to the napkin with the sketch pad. Research has shown that users spend a significant amount of time inspecting drawings in a computer-supported pen-based system [Lim 2003], and this would be especially true for mobile interfaces due to their small displays. We are hoping users can take advantage of the tangible camera controls to achieve a more fluid design experience. To interact with our system, users place the napkin on a surface such as a table, move around it to select a comfortable view of the design space, and then sketch on the surface of the sketch pad to record their designs on the napkin. Users can also change the color and size of their strokes through standard GUI controls. Overall, the combination of the sketch pad and the napkin provides a light weight and relatively inexpensive mobile sketching interface which is easy to set up and use.

4 Projective 3D Sketching

The main design approach for our software interface is the notion of freeform interaction. This concept is introduced in Moran et al.'s work [1995] on implicit structures for pen-based systems. Keeping the structures of representations implicit or temporary is the essence of freeform interaction. We feel this approach fits our design goal because it allows the system to be flexible and easy to use. One major advantage of freeform interaction is its ability to support ambiguity which is a highly valued feature of pencil and paper. By incorporating freeform interaction in 3D sketching, users can suggest different designs with unconstrained strokes to layout their rough ideas and not worry about having to deal with explicit geometric structures before they are ready. Following the concepts of freeform interaction, we are using projective 3D sketching as the main method of creating 3D representations in our system. Unlike 2D sketching, this technique requires two steps. Users first define a 3D surface in the design space where their sketches will be projected, and then they sketch on the tablet PC as they view the 3D design space on top of the napkin. The recorded 2D strokes are projected onto the 3D surface in a way such that the projected 3D stroke looks identical to the original 2D stroke. This visual correspondence along with the directness of sketching on the tablet PC provides a natural sketching experience for users. Because no explicit structures need to be interpreted for the sketches, users are free to use solid, overlapping, and stippled strokes to indicate contours or hatching and scribbling to suggest surfaces.

One important element in our interface is the concept of frames. These are temporary 3D surfaces which are placed in the scene for the sketches to be projected on (Figure 2). In terms of freeform interaction, these are the implicit structures we can define to guide the sketching process. The simplest frame is a plane. Cognitively, users can think of the surfaces of the 3D design they wish to sketch as frames or parts of frames, or the frames themselves can be interpreted as a sort of flexible canvas that can be spatially positioned

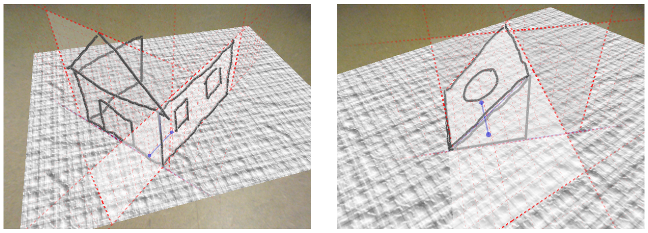

Figure 3: *Intuitive frame switching with the one stroke technique and the frame extension display approach*

anywhere in the scene. For example, a vertical plane can be temporarily perceived as the wall of a house. Users can then sketch the appropriate boundaries of the wall on the plane to make the idea more concrete or scribble and fill in the area to further solidify the concept. In Napkin Sketch, only one frame is active or can be sketched on at any time during the sketching process. This avoids confusion as frames can possibly occlude each other. To maintain visual coherence with the physical space, frames can only be placed on top of the napkin. Frames with infinite extent are truncated at their intersection with the surface of the napkin. This ensures that no strokes can be sketched underneath the napkin. As noted before, the definition of 3D surfaces for sketch projection is a major bottleneck of the projective sketching technique as users must often stop sketching, correctly position a surface, and then resume sketching again. In Napkin Sketch, we attempt to minimize the interruption this process can cause to users. Our one stroke technique allows users to quickly create and position plane frames by simply sketching. The concept of our one stroke technique (Figure 2) is to take advantage of the implicit geometric relationships of the strokes in a sketch. Strokes rarely float in space, visually disconnected from others. For example, in the sketch of a cube, the lines sketched for the front face of the cube are perpendicular to the receding lines of the side face of the cube. Therefore, when users finish sketching the front face of the cube on one plane frame, it is natural for them to wish to switch to a plane frame that is perpendicular to the previous plane and also intersects the previous plane at one of the edges sketched for the face (Figure 3). This would allow them to quickly continue sketching the side face of the cube. The one stroke technique supports this process. Users can sketch frame creation strokes, activated via a mode switch. This stroke is projected on the current active frame just like regularly sketched strokes. A relatively straight stroke is interpreted as a straight line, and a new plane frame is created which includes the line and is perpendicular to the plane frame that the line is sketched on. For example, if a frame creation stroke is sketched on the napkin ground plane, a vertical plane frame would be created at the location where the stroke is made and oriented to be parallel with the stroke. However, the one stroke technique only allows new plane frames that are perpendicular to the previous plane frame to be created in one step. Therefore, we also allow users to rotate the created plane frame along the axis specified by the sketched frame creation stroke to cover all possible plane frame orientations. This is achieved by manipulating graphical handles on the plane frames (Figure 2). We feel that generating perpendicular plane frames with one stroke is a logical choice because many common everyday objects have contours with perpendicular relationships. Our technique for defining plane frames allows users to transition fluidly from one plane to another as they follow the contours and surfaces of their design.

Users can cycle through previously visited frames using two arrow buttons for going backward and forward. They can go to previous frames to modify parts of the sketch made on those frames and use them as starting points for creating new frames. This allows the system to better deal with the unpredictability associated with the work flow of different users. Since it is easy to become lost in virtual 3D space without proper visual cues, we have tried to carefully design the appearance of the plane frame in our system. Because plane frames are infinite, they are difficult to visualize. In practice, quads must be rendered to indicate their position in 3D. We try to infer the proper size of the newly created plane frames from the frame creation stroke. New plane frames are created with the same width as the creation stroke and twice the width of the stroke is set as the height. The display boundaries of the plane frames are rendered as stippled red lines, and they gradually fade into space to indicate the infinite nature of the plane. The creation stroke is replaced with a straight blue stippled line which indicates the plane frame's axis of rotation. The quad itself is highlighted in translucent white if the plane frame is the current active frame. We also display the previous plane frame in translucent gray as a 3D visual cue. Since the napkin ground plane is a critical visual cue for anchoring the 3D design space, we reorient the visual boundaries of the plane frames so that the top and bottom edges of the frame are always parallel to the napkin. We also try to extend the boundaries of the plane frame so the bottom edge touches the napkin, provided the plane is not parallel to the napkin, and the extension does not make the boundaries too expansive (Figure 3). This allows users to easily find the surface of the napkin if required and helps them judge the 3D positions of their sketches. Faint grid lines are also displayed on the plane frames to provide a perspective cue and help users sketch straight lines.

5 Preliminary Assessment

We have performed an informal preliminary usability evaluation of Napkin Sketch with seven participants (4 male and 3 female) of varying artistic and drawing skills recruited from the lab. They were provided with a brief tutorial on how to use the interface, but no explicit instructions were given on what they should create with the system. We observed the participants sketching their desired creations and conducted open-ended interviews afterward. The results are positive (see Figure 4 for sample sketches), but we have also discovered some major areas for improvement. We are excited to see that users interacted with our interface in ways which are similar to conventional pencil and paper. Without explicit instructions, users took advantage of the flexibility of the system to perform sketch operations which are analogous to those in conventional 2D sketching. For example, one user started his sketches, drawing in very thin gray lines and later went back to emphasize the lines with thicker and darker strokes. This display of over sketching demonstrates the natural top-down design thinking. Other users appropriated the ability to sketch large strokes as a way to provide occlusion for suggested surfaces. They filled in surfaces with colored scribbles to block out strokes seen in the background. This shows the concept of freeform interaction. The users did not explicitly define surfaces but rather used visual representations to suggest their presence. We are also happy to see that many of the interface features we designed are useful to the users. For example, most of them found the grid lines displayed on the plane frames to be useful for drawing in perspective and measuring, all of them made heavy use of the frame cycling capability, and many of them found the 3D visual cues beneficial, commenting on features such as displaying the previous frame and extending the plane frame visual boundaries to the napkin as being helpful. Many users also commented on the intuitiveness of the tangible camera controls. Despite the success of the perspective and 3D visual cues, users still became lost on occasion as to how to connect lines in space or make proper sketches in perspective. Perspective sketching is not as intuitive as we anticipated. In fact, most of the users mentioned that they do not normally sketch in perspective even in 2D. Because of this,

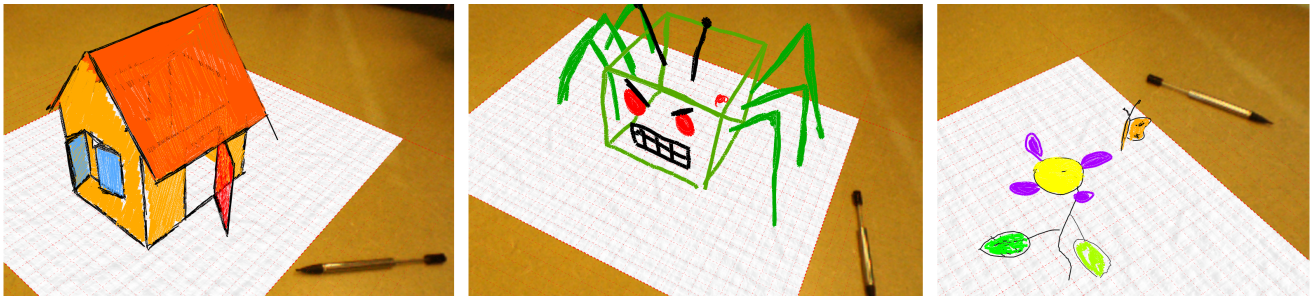

Figure 4: *Interaction examples from our preliminary evaluation*

many of them frequently changed the view of the scene using the napkin to find a visual representation which minimized perspective distortion. One user even sketched explicit visual cues to help him connect lines in 3D. The system can definitely benefit from a richer set of visual guides which can possibly be added by the users to assist them in perspective sketching. Although the direct camera manipulation with the napkin worked well with the exception of minor tracking jitters, we found it to be awkward to use in two situations. One was when users wished to view the scene in profile, and the other was when they wished to view the scene from overhead. Viewing the scene in profile means the camera must be located at an extreme angle to the napkin, where the markers are hardly visible, and tracking deteriorates. Viewing the scene from overhead often causes users to stand up which was at times awkward or uncomfortable. This issue can be solved by providing better visual guides to make users more confident about sketching in perspective or allowing them to temporarily sketch in the profile and overhead views and resume napkin tracking of the scene when they are done. One last interesting finding is that users demonstrated different types of work flows when defining planes and sketching. Some followed connected frames in sequence, and others used one frame as the reference frame and always went back to that frame for creating new frames. One user wanted to create a set of frames all at once and cycle through them to sketch without having to stop and create new frames, but this type of work flow is not supported by our system. A solution to this may be to allow users to explicitly save frames. This finding demonstrates the diversity of the ways users sketch and hints at the need for more flexibility in future prototypes of Napkin Sketch.

6 Conclusion and Future Work

The approaches we have taken for the design of our Napkin Sketch system appear promising. We are planning to conduct more formal evaluations with skilled artists and designers, but judging from the limited yet insightful findings from our preliminary evaluation, it seems users are adopting the interface in ways similar to pencil and paper rather than along the lines of using conventional 3D modeling software. One telling sign is that users are taking a variety of approaches to interact with the system. This indicates that the users' creative thinking is not restricted by the structure of the interface. Therefore, although 3D sketching may take longer than conventional 3D modeling software to construct certain designs, the future direction of our system is not to find an optimal method for 3D sketching in terms of efficiency but rather to provide users with more resources for creating design solutions. For instance, allowing them to quickly select previous frames they have worked on by selecting existing sketches would provide more flexibility for their work flow. We also need to explore better 3D visualization techniques and try to understand where bottlenecks occur in the users' 3D perception. Finally, we would like to look at collocated collaborative sketching with several users working together on a single 3D design task through the use of napkins and multiple individual sketch pads, similar to the way a collocated group of people may jot and examine ideas by sketching on napkins.

References

CHEOK, A. D., EDMUND, N. W. C., AND ENG, A. W. 2002. Inexpensive non-sensor based augmented reality modeling of curves and surfaces in physical space. In *ISMAR '02: Proceedings of the 1st International Symposium on Mixed and Augmented Reality*.

DO, E. Y.-L., AND GROSS, M. D. 1997. Inferring design intentions from sketches-an investigation of freehand drawing conventions in design. In *CAADRIA '97: Proceedings of Computer Aided Architectural Design Research in Asia*.

DORSEY, J., XU, S., SMEDRESMAN, G., RUSHMEIER, H., AND MCMILLAN, L. 2007. The mental canvas: A tool for conceptual architectural design and analysis. In *PG '07: Proceedings of the 15th Pacific Conference on Computer Graphics and Applications*.

KIYOKAWA, K., TAKEMURA, H., AND YOKOYA, N. 1997. Manipulation aid for two-handed 3-d designing within a shared virtual environment. In *HCI '97: Proceedings of Human Computer Interaction International*.

LAPIDES, P., SHARLIN, E., SOUSA, M. C., AND STREIT, L. 2006. The 3d tractus: A three-dimensional drawing board. In *TABLETOP '06: Proceedings of the First IEEE International Workshop on Horizontal Interactive Human-Computer Systems*.

LIM, C.-K. 2003. An insight into the freedom of using a pen: Pen-based system and pen-and-paper. In *Proceedings of the 6th Asian Design International Conference*.

MORAN, T. P., CHIU, P., VAN MELLE, W., AND KURTENBACH, G. 1995. Implicit structure for pen-based systems within a freeform interaction paradigm. In *CHI '95: Proceedings of the SIGCHI conference on Human factors in computing systems*.

PICCOLOTTO, M. A. 1998. *Sketchpad+ Architectural Modeling through Perspective Sketching on a Pen-based Display*. Master's thesis, Cornell University.

SACHS, E., ROBERTS, A., AND STOOPS, D. 1991. 3-draw: A tool for designing 3d shapes. *IEEE Comput. Graph. Appl. 11*, 6.

TOLBA, O., DORSEY, J., AND MCMILLAN, L. 1999. Sketching with projective 2d strokes. In *UIST '99: Proceedings of the 12th annual ACM symposium on User interface software and technology*.