Virtual Reality (2007) 11: 15–21
DOI 10.1007/s10055-006-0034-6

ORIGINAL ARTICLE

Jeffrey S. Pierce · Randy Pausch

Generating 3D interaction techniques by identifying and breaking assumptions

Received: 22 November 2004 / Accepted: 15 June 2006 / Published online: 12 December 2006

Abstract Researchers have created 3D interaction techniques for immersive virtual worlds, but existing techniques represent just part of the design space. While exploring other parts of the design space might yield more effective techniques, conducting that exploration is difficult and time-consuming. Analyzing the particular task, user, and hardware characteristics for any given problem is straightforward, but only suggests the shape of a potential technique; generating the technique itself still requires a creative breakthrough. We propose extending existing approaches to generating 3D interaction techniques by focusing more explicitly on identifying and breaking assumptions about the real world to inspire potential technique ideas. We describe our approach, suggest an initial list of assumptions to consider, and present a case study of applying the process to create a technique for navigation with visible landmarks and place representations.

Keywords Virtual reality · 3D interaction · Interaction techniques · Generative methods

1 Introduction

Immersive 3D interfaces offer potential advantages over 2D interfaces, including higher dimensionality input and output and more effective use of proprioception, kinesthesis, and spatial memory. Realizing those advantages, however, requires effective 3D interaction techniques. Researchers have created some techniques, but they represent just part of the possible design space. As such, more effective techniques may exist, but to find them we need to create new techniques.

Creating an effective 3D interaction technique typically entails five steps:

1. Identifying the task that users must perform.
2. Identifying the target user population.
3. Choosing input, output, and tracking hardware.
4. Generating an idea for a new technique that allows users to more effectively complete the task with the available hardware.
5. Evaluating the technique to verify its efficacy.

Completing most of these steps is straightforward. Steps one through three involve either choices or constraints, and designers can draw on well-known methods for step 5. Step 4, generating an idea, is more difficult because it requires inspiration. A designer can know exactly what the task, user population, hardware, and evaluation method will be, but without a good idea he cannot succeed.

Researchers have proposed approaches to help inspire ideas. Considering the task (Tan et al. 2001), user, and hardware characteristics may help shape potential interaction techniques. Alternately, designers can methodically combine existing components of interaction techniques in new ways (Bowman and Hodges 1999). These approaches are primarily evolutionary: they yield new techniques that are similar to existing techniques.

We propose generating ideas by *identifying and breaking assumptions about the real world.* Identifying our assumptions and imagining what is possible if we break them is an established technique for spurring creativity e.g., (Gentner and Nielsen 1996). Concentrating on assumptions about the real world can help designers focus on where virtual reality can productively diverge from, rather than emulate, the real world. This approach, when used as part of the larger process, could help designers create revolutionary new interaction techniques. In the following sections we describe the

J. S. Pierce (✉)
IBM Almaden Research Center, San Jose, CA, USA
E-mail: jspierce@us.ibm.com
Tel.: +1-408-9271282
Fax: +1-408-9273033

R. Pausch
Carnegie Mellon University, Pittsburgh, PA, USA

Table 1 Potential assumptions drawn from previous work

Assumption	Sample technique(s) breaking it
Space is linear and continuous	Locales (Barrus et al. 1996), head-butt zoom (Mine et al. 1997)
The visual properties of a space match the physical properties of a space	Guided navigation (Galyean 1995), force fields (Xiao et al. 1998)
Appearance does not necessarily reflect reality	View-based teleportation (Fisher 1989), flying into a WIM (Pausch et al. 1995)
There is only one world	Worlds within worlds (Feiner and Beshers 1990)
Everyone in the world occupies the same point in time	Local perception filter (Sharkey et al. 1998)
The world is persistent even when not in view	View-dependent culling (Chenney and Forsyth 1997)
Gravity exists	Almost all 3D worlds
Objects do not arbitrarily attract each other	Snap-dragging (Bier 1990)
The intrinsic (e.g., density) and extrinsic (e.g., size, shape) properties of objects are inviolate and distinct	View-dependent geometry (Rademacher 1999), 3D magic lenses (Viega et al. 1995)
Objects are persistent: we cannot create or destroy them on demand	Over-the-shoulder deletion (Mine et al. 1997), voodoo dolls (Pierce and Pausch 2002)
Objects can only exist in one location	World-in-miniature (Stoakley et al. 1995), voodoo dolls (Pierce and Pausch 2002)
"Passive" objects (e.g., a map) cannot affect other objects	Pointing to teleport (Angus and Sowizral 1995), world-in-miniature (Stoakley et al. 1995), voodoo dolls (Pierce and Pausch 2002)
Objects work the same in all contexts	Moding user action by hand position (Mine et al. 1997), voodoo dolls (Pierce and Pausch 2002)
Objects have no built-in knowledge of their properties or functionality	User-centered maps (Darken and Sibert 1993), object associations (Bukowski and Sequin 1995)
Objects (and the world) have no high level semantic memory or history	Almost all 3D worlds
The actions the user can perform and their effects are independent of the current context or task	Rapid controlled movement (Mackinlay et al. 1990), context-sensitive flying (Ware et al. 1997)
We control our viewpoint and actions	Cam droid (Drucker et al. 1995)
We move through the world, the world does not move around us	World-in-hand (Ware and Osborne 1990)
Not all abstractions are represented by physical objects	Field-of-view objects (Hindmarsh et al. 1998)
Effects follow causes	Interactive shadows (Herndon et al. 1992)
Objects work in the familiar, expected way	Vampire mirrors & privacy lamps (Butz et al. 1998)

approach and present a case study illustrating its application.

2 Identifying and breaking assumptions

When creating 3D interaction techniques, many designers implicitly break assumptions. We propose instead explicitly identifying and breaking assumptions, particularly assumptions about the real world. Emphasizing differences from the real world may help designers to more fully explore how to leverage the total control that they possess over virtual worlds.

To apply this approach, designers walk through four steps: choosing a problem to solve, assembling a list of assumptions to break, repeatedly iterating through the list and breaking the assumptions, and finally selecting promising ideas.

Step 1: Choose a problem to solve. The problem can range from very general to very specific. Choosing a general problem facilitates more free-form brainstorming and can theoretically lead to particularly revolutionary techniques. However, we found in practice that it primarily generates impractical or useless ideas. During our initial experimentation with this approach, the primary author chose very general problems (e.g., manipulating a distant object, traversing a large virtual world) and, over 6 months, generated roughly 700 ideas worth recording (and many more not even worth that). However, none of them were sufficiently promising to warrant more than informal evaluation.

We had more success when we chose more specific problems shaped by characteristics of the task, users, or hardware. Breaking up general problems into multiple, specific problems led the primary author to generate roughly 200 smaller ideas worth recording in 4 months and yielded three publishable techniques. While we thus had significantly more success choosing specific problems, we note that unconstrained exploration does have some value. While not directly useful, ideas we generated to solve general problems did influence the ideas we later generated to solve specific problems.

Step 2: Assemble a list of assumptions to break. Because the set of potential assumptions is almost limitless, reducing the candidate assumptions to a manageable quantity is critical. The task, user, and hardware char-

Table 2 Technique ideas for navigation over intermediate distances

Assumption	Idea	Shortcomings
Not all abstractions are represented by physical objects	Reify shortcuts as "lifelines" that attach users to locations; users can return to a location by pulling on its lifeline	Does not scale
Appearance does not necessarily reflect reality	Users look through a telescope, zoom in on a location, and when they look away they are in that location	Does not address occlusion
Objects are persistent: we cannot create or destroy them on command	Users cause a temporary mountain to well up beneath them, allowing them to see (and image plane navigate to) distant objects	In practice view is limited by clipping plane
"Passive" objects cannot affect other objects	Navigation towers that move users (as well as orienting them) when selected	May be occluded
Objects can only exist in one location	Show copies of user in different directions and at multiple distances. Selecting one moves user to its position	Might confuse users and crowd space
"Passive" objects cannot affect other objects	Provide clouds that act like navigation towers	May be difficult to associate clouds with particular locations
The intrinsic and extrinsic properties of objects are inviolate and distinct	Allow users to compress the world toward themselves, navigate, and then expand world	Compressed world may be too cluttered or, if scaled, not provide sufficient accuracy
There is only one world	Provide an alternate version of the world that is more compact and contains only landmarks. To navigate, users switch to the alternate version, navigate to the landmark nearest the desired location, switch back, and then navigate locally	Might be too cumbersome
Space is linear and continuous	Violate occlusion by drawing the navigation towers so that they always appear in front of other objects	Distant towers may be too small
Space is linear and continuous	Landmarks actively aid users: remain visible, maintain a minimum size. Users can image plane navigate to them	May clutter view

acteristics help constrain the assumptions to consider. For example, when trying to create a technique for efficient navigation in very large worlds, breaking assumptions about the properties of space may be effective.

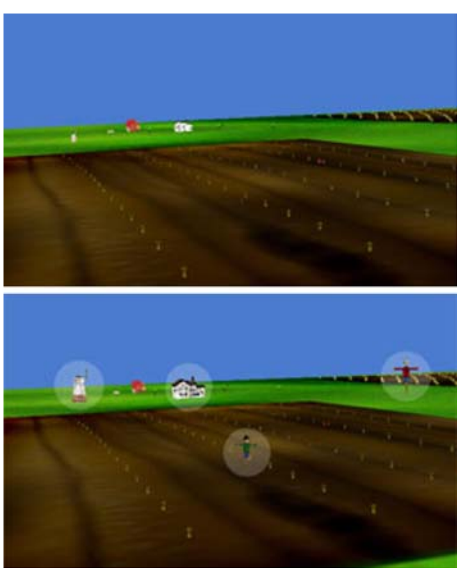

Fig. 1 Views of a farm before (*top*) and after (*bottom*) the addition of visible landmarks

Designers can also reduce the number of assumptions they consider by starting with assumptions that existing techniques break. For our work, we assembled an initial list of assumptions by reviewing the literature and identifying assumptions implicitly broken by existing techniques. We then adjusted our list over time based on our experiences using it.

Some of our initial assumptions carried through to the final list (e.g., *space is linear and continuous* or *the visual and physical properties of a space are identical*). We combined and generalized others that we found too specific to help generate a variety of techniques. For example, we initially had multiple assumptions about object properties: their size remains constant, their shape remains constant, they have mass etc. Rather than focusing too narrowly on a subset of properties, we combined the individual assumptions into the more general assumption that *an object's intrinsic and extrinsic properties are inviolate and distinct*. Finally, we completely eliminated some assumptions. For example, we eliminated the assumption that the world is round because round worlds are already the exception, rather than the rule. By contrast, we kept the assumption that gravity exists as a reminder that simulating physics in virtual worlds, a subject of much research, is not necessarily a good idea.

Table 1 lists the final version of our list. We found three of the listed assumptions particularly useful: space is linear and continuous, objects are persistent, and appearance does not necessarily reflect reality.

Table 3 Technique ideas for navigation over large distances

Assumption	Idea	Shortcomings
Objects work the same in all contexts	Make landmarks visible only within certain locations. Could organize them into a tree or graph where users could only see landmarks a few nodes away	Might be unclear how to reach an arbitrary location
Extrinsic properties of objects are inviolate	Two-handed image plane world scaling: provide dynamic control over resizing and increase maximum scaling factor	Increasing scaling factor decreases accuracy
We move through the world, the world does not move around us	Move a place instead of moving to it. Place a landmark corresponding to a location to move location there	Does not scale
Space is linear and continuous	Divide world into squares and turn it into a sort of chessboard where users can decide to expand or collapse the square they are in or move between squares	Arbitrary divisions might make it hard to find a location
Not all abstractions are represented by physical objects	Present abstract levels (neighborhood, city, state, country) as actual levels. Users can move between nodes at the same level, or step up or down levels	Not clear how to present levels, move between nodes
Objects are persistent: we cannot create or destroy them on demand	User holds a WIM representing the entire world in one hand. With other hand, user reaches in, specifies volume for new WIM, and pulls it out. Repeat to create WIMs for smaller and smaller areas	Might be excessively time-consuming
Objects can only exist in one location	Display multiple WIMs at different scales to user. User chooses desired scale instead of specifying	May need too many WIMs
The visual properties of a space match the physical properties of a space	WIMs representing large areas could show "artist rendition" (symbolic) versions of important areas to aid navigation	Problematic for "flying into the WIM"
Objects are persistent: we cannot create or destroy them on demand	Divide world into sections, create one WIM per section, then embed WIMs for smaller sections into WIM encompassing them, repeat until have WIM for whole world. Allow users to pull out WIM for smaller section from current WIM	May be too cumbersome for users
Not all abstractions are represented by physical objects	Provide landmarks for large places (e.g., US East Coast). When user selects one provide a zooming WIM for that place. User can zoom in and then fly in	Zooming might be too time-consuming
Not all abstractions are represented by physical objects	The hierarchy of places in a world forms a discrete set of focus centers and scales for WIMs. Landmarks can serve as shortcuts for accessing them. Avoids need to zoom	Might not be clear which landmarks correspond to which places
Not all abstractions are represented by physical objects	Present a small set of landmarks and WIMs (some symbolic) around user (determined using place hierarchy), allow users to navigating using them (for WIMs either pull out WIM for lower-level place or fly into the depicted place)	Need to organize world into place hierarchy
Appearance does not necessarily reflect reality	Teleport to place by achieving desired view of WIM representing it	May disorient users

Fig. 2 A view of a city with two symbolic place representations in the foreground

Step 3: Iterate through the list. Take each assumption and break it to generate as many ideas as possible. We found it useful to choose a quantifiable goal (e.g., generate 30 ideas in 30 min) and continue iterating through the list until we reached it. The key for this step is quantity, not quality. Focusing on producing only ''good'' ideas makes it harder to create techniques in new parts of the design space because it reinforces the tendency to rely on useful aspects of known techniques.

Step 4: Select promising ideas. We found it useful to take an initial pass through the generated ideas immediately after each brainstorming session and discard those that were impractical or that clearly violated established interface design principles. We then typically selected a small set (roughly 1 in 20) of the remaining ideas that, based on our experience and/or the research literature, seemed particularly promising. We built quick prototypes of them, and if they still seemed interesting we conducted informal evaluations. We only formally evaluated the small number of ideas that still seemed worthy of publication.

We successfully applied this approach to create three interaction techniques: Voodoo Dolls (Pierce and Pausch 2002), painting Interaction Surfaces (Pierce and Pausch 2003), and navigating with visible landmarks and place representations (Pierce et al. 2004). In the next section we present a case study describing how we applied this approach to create the latter technique.

3 Case study: visible landmarks and place representations

Visible landmarks and place representations facilitate navigation in very large virtual worlds. Visible landmarks keep landmarks visible, allowing image plane navigation (Pierce et al. 1997) relative to them. Actual and symbolic place representations enable travel to distant locations. Users travel using an actual representation (a scaled replica) of a place by holding it to achieve the desired view and teleporting. Users travel using a symbolic representation (an ''artist's conception'') of a large place by reaching into it, pulling out a representation for a smaller place, and then either repeating the step with the new representation (if it is symbolic) or teleporting (if it is actual). The world's semantic place hierarchy determines which visible landmarks and place representations are visible to users in different places.

When we initially set out to create a technique for efficiently navigating large virtual worlds, we chose adults in their 20s and 30s who are proficient with virtual worlds as our target users and a head-mounted display and trackers for the user's head and hands for our hardware. We chose as our initial problem helping users navigate to locations at intermediate distances: far

Fig. 3 To travel to the beach from the city, the user pulls away the symbolic representation of the east coast (*top left*) and then pulls out the symbolic representation of the coast itself (*top right*). The user pulls out the actual representation of the beach (*bottom left*), positions it to get the desired view, and then lets go to teleport (*bottom right*)

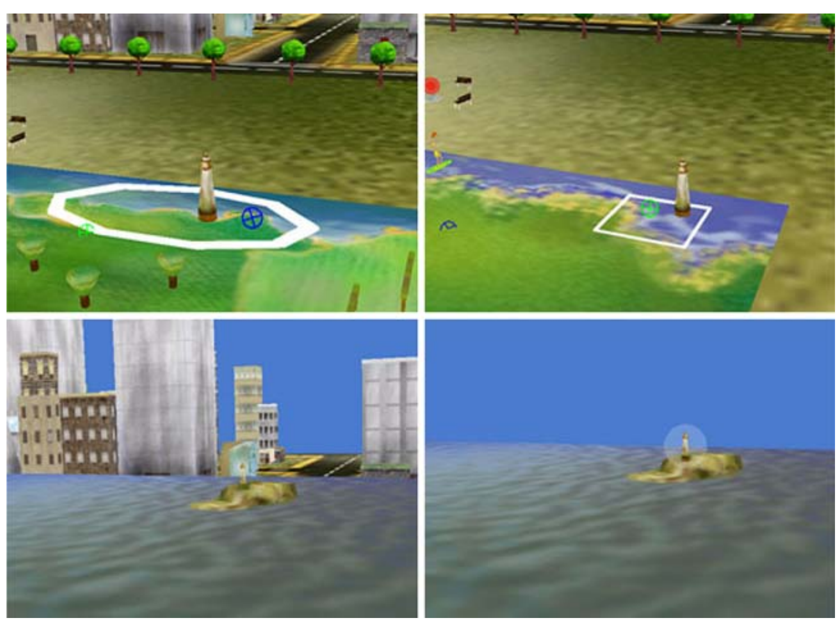

enough that users could not currently see them, but close enough for users to know an approximate distance and direction.

Table 2 lists, in temporal order, some of the ideas that the primary author generated during this stage, their shortcomings, and the employed assumptions. For brevity this list is not exhaustive; we instead present a sample to illustrate the range of ideas we explored and to convey how our ideas converged over time.

While we informally evaluated some of those ideas, the idea that landmarks can serve an expanded role in virtual worlds stood out as both grounded in real world practice and yielding practical benefit. In the real world, landmarks help people locate their objective. When looking for a particular place, people often locate the nearest landmark and navigate relative to it (Lynch 1960). In the real world, the utility of landmarks is limited because at greater distances they get too small to see or other objects occlude them. If we break the assumption that *space is linear and continuous*, we can provide visible landmarks that always remain visible and are always available as targets for image plane navigation (Fig. 1).

While visible landmarks had the potential to aid navigation over intermediate distances, they do not scale to worlds with many landmarks. We thus chose as our next problem how to support navigation over large distances in complex worlds and applied our process again. Table 3 lists the temporally ordered subset of ideas generated by the primary author.

During brainstorming our emphasis slowly shifted to exploring mechanisms to allow users to work with large worlds at a tractable scale, particularly by dividing them up. A key insight was that, as people naturally organize their cognitive map of the world into a semantic hierarchy (e.g., neighborhood, city, state, region, country) (Chase 1983; Stevens and Coupe 1978; Gould and White 1986), we could break the assumption that *not all abstractions are represented by physical objects* and organize the world as a hierarchy of place objects. Subsequently generated ideas contributed other technique elements: showing representations (initially WIMs) rather than landmarks for distant places, providing both symbolic and actual place representations, using the hierarchy to determine which representations to show, embedding representations hierarchically, and teleporting to a place by achieving the desired view of its representation (Figs. 2; 3).

Combining those elements with visible landmarks seemed likely to support efficient navigation over both intermediate and large distances in very large worlds. We built an initial implementation and refined it using informal evaluations, and a summative evaluation confirmed the technique's effectiveness (Pierce et al. 2004).

4 Conclusion

Our case study illustrates that explicitly identifying and breaking assumptions about the real world can successfully generate new 3D interaction techniques. While we recognize that other individuals' degree of success with this approach will vary, our work applying this approach and the large body of existing techniques that implicitly break assumptions suggests that it is promising. Even if the approach offers only a small benefit, exploring the design space for effective 3D interaction techniques is sufficiently difficult that even small benefits are worthwhile. We look forward to other designers and researchers successfully employing this approach, instead of or in addition to their current practices, to create new and unusual techniques.

References

Angus I, Sowizral H (1995) Embedding the 2D interaction metaphor in a real 3D virtual environment. Stereoscopic displays and virtual reality systems II. In: SPIE proceedings, 2409:282–293

Barrus J, Waters R, Anderson D (1996) Locales and beacons: efficient and precise support for large multi-user virtual environments. In: Proceedings of VRAIS 1996, pp 204–213

Bier E (1990) Snap-dragging in three dimensions. In: Proceedings of the 1990 workshop on interactive 3D graphics, pp 193–204

Bowman D, Hodges L (1999) Formalizing the design, evaluation, and application of interaction techniques for immersive virtual environments. J Vis Languages Comput 10(1):37–53

Bukowski R, Sequin C (1995) Object associations: a simple and practical approach to virtual 3D manipulation. In: Proceedings of the 1995 symposium on interactive 3D graphics, pp 131–138

Butz A, Beshers C, Feiner S (1998) Of vampire mirrors and privacy lamps: privacy management in multi-user augmented environments. In: Proceedings of UIST 1998, pp 171–172

Chase W (1983) Spatial representations of taxi drivers. In: Rogers R, Sloboda J (eds) Acquisition of symbolic skills. Plenum, New York, pp 391–405

Chenney S, Forsyth D (1997) View-dependent culling of dynamic systems in virtual environments. In: Proceedings of the 1997 symposium on interactive 3D graphics, pp 55–58

Darken R, Sibert J (1993) A toolset for navigation in virtual environments. In: Proceedings of UIST 1993, pp 157–165

Drucker S, Zeltzer, D (1995) CamDroid: A system for implementing intelligent camera control. In: Proceedings of the 1995 symposium on interactive 3D graphics, pp 139–144

Feiner S, Beshers C (1990) Worlds within worlds: metaphors for exploring *n*-dimensional virtual worlds. In: UIST proceedings, pp 76–83

Fisher S (1989) The AMES virtual environment workstation (VIEW). In: SIGGRAPH 1989 course #29 notes

Galyean T (1995) Guided navigation of virtual environments. In: Proceedings of the 1995 symposium on interactive 3D graphics, pp 103–104

Gentner D, Nielsen J (1996) The anti-mac interface. Commun ACM 39(8):70–82

Gould P, White R (1986) Mental maps. 2nd edn. Allen & Unwin, Boston

Herndon K, Zeleznik R, Robbins D, Conner D, Snibbe S, van Dam A (1992) Interactive shadows. In: Proceedings of UIST 1992, pp 1–6

Hindmarsh J, Fraser M, Heath C, Benford S, Greenhalgh C (1998) Fragmented interaction: establishing mutual orientation in virtual environments. In: CSCW 1998 proceedings, pp 217–226

Lynch K (1960) The image of the city. MIT Press, Cambridge

Mackinlay J, Card S, Robertson G (1990) Rapid controlled movement through a virtual 3D workspace. In: SIGGRAPH 1990 proceedings, pp 171–176

Mine M, Brooks, Jr. F, Sequin C (1997) Moving objects in space: exploiting proprioception in virtual environment interaction. In: SIGGRAPH 1997 proceedings, pp 19–26

Pausch R, Burnette T, Brockway D, Weiblen M (1995) Navigation and locomotion in virtual worlds via flight into hand-held miniatures. In: SIGGRAPH 1995 proceedings, pp 399–400

Pierce J, Pausch R (2002) Comparing voodoo dolls and HOMER: exploring the importance of feedback in virtual environments. In: Proceedings of CHI 2002, pp 105–112

Pierce J, Pausch R (2003) Specifying interaction surfaces using interaction maps. In: Proceedings of the 2003 symposium on interactive 3D graphics, pp 189–192

Pierce J, Pausch, R (2004) Navigation with place representations and visible landmarks. In: Proceedings of IEEE VR 2004, pp 173–180

Pierce J, Forsberg A, Conway M, Hong S, Zeleznik R, Mine M (1997) Image plane interaction techniques in 3D immersive environments. In: Proceedings of the 1997 symposium on interactive 3D graphics, pp 39–43

Rademacher P (1999) View-dependent geometry.In: SIGGRAPH 1999 proceedings, pp 439–446

Sharkey P, Ryan M, Roberts D (1998) A local perception filter for distributed virtual environments. In: Proceedings of VRAIS 1998, pp 242–249

Stevens A, Coupe P (1978) Distortions in judge spatial relationships. Cognit Psychol 10(4):422–437

Stoakley R, Conway M, Pausch R (1995) Virtual reality on a WIM: interactive worlds in miniature. In: CHI 1995 proceedings, pp 265–272

Tan D, Robertson G, Czerwinski M (2001) Exploring 3D navigation: combining speed-coupled flying with orbiting. In: CHI 2001 proceedings, pp 418–425

Viega J, Conway M, Williams G, Pausch R (1995) 3D magic lenses. In: Proceedings of UIST 1995, pp 51–58

Ware C, Fleet, D (1997) Context sensitive flying interface. In: Proceedings of the 1997 symposium on interactive 3D graphics, pp 127–130

Ware C, Osborne S (1990) Exploration and virtual camera control in virtual three dimensional environments. In: Proceedings of the 1990 symposium on interactive 3D graphics, pp 175–183

Xiao D, Hubbold, R (1998) Navigation guided by artificial force fields. In: CHI 1998 proceedings, pp 179–186