

**TESTING THE USABILITY OF MICROSOFT  
PHOTOSYNTH, AN APPLICATION WITH A COMPLEX  
AND UNFAMILIAR CONCEPTUAL MODEL**

by

Chirag Xerxes Vesuvala  
Bachelor of Science, University of British Columbia 2005

PROJECT SUBMITTED IN PARTIAL FULFILLMENT OF  
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE

In the  
School of Computing Science

© Chirag Xerxes Vesuvala 2009

SIMON FRASER UNIVERSITY

Spring 2009

All rights reserved. This work may not be  
reproduced in whole or in part, by photocopy  
or other means, without permission of the author.

# APPROVAL

**Name:** Chirag Xerxes Vesuvala  
**Degree:** Master of Science  
**Title of Thesis:** Testing the usability of Microsoft Photosynth, an application with a complex and unfamiliar conceptual model

**Examining Committee:**

**Chair:** **Dr. Stella Atkins**  
Professor of Computing Science

---

**Dr. Arthur E. Kirkpatrick**  
Senior Supervisor  
Associate Professor of Computing Science

---

**Dr. Fred Popowich**  
Supervisor  
Professor of Computing Science

---

**Dr. Lou Hafer**  
Internal Examiner  
Associate Professor of Computing Science

**Date Defended/Approved:** March\_24\_2009



SIMON FRASER UNIVERSITY  
LIBRARY

## Declaration of Partial Copyright Licence

The author, whose copyright is declared on the title page of this work, has granted to Simon Fraser University the right to lend this thesis, project or extended essay to users of the Simon Fraser University Library, and to make partial or single copies only for such users or in response to a request from the library of any other university, or other educational institution, on its own behalf or for one of its users.

The author has further granted permission to Simon Fraser University to keep or make a digital copy for use in its circulating collection (currently available to the public at the "Institutional Repository" link of the SFU Library website <[www.lib.sfu.ca](http://www.lib.sfu.ca)> at: <<http://ir.lib.sfu.ca/handle/1892/112>>) and, without changing the content, to translate the thesis/project or extended essays, if technically possible, to any medium or format for the purpose of preservation of the digital work.

The author has further agreed that permission for multiple copying of this work for scholarly purposes may be granted by either the author or the Dean of Graduate Studies.

It is understood that copying or publication of this work for financial gain shall not be allowed without the author's written permission.

Permission for public performance, or limited permission for private scholarly use, of any multimedia materials forming part of this work, may have been granted by the author. This information may be found on the separately catalogued multimedia material and in the signed Partial Copyright Licence.

While licensing SFU to permit the above uses, the author retains copyright in the thesis, project or extended essays, including the right to change the work for subsequent purposes, including editing and publishing the work in whole or in part, and licensing other parties, as the author may desire.

The original Partial Copyright Licence attesting to these terms, and signed by this author, may be found in the original bound copy of this work, retained in the Simon Fraser University Archive.

Simon Fraser University Library  
Burnaby, BC, Canada

## **ABSTRACT**

How do we test the usability of software whose design features an unprecedented and complex conceptual model? Current usability practices ask users for objective and subjective responses, which do not explicitly test how well a conceptual model is understood. An example of software with a complex conceptual model is Microsoft Photosynth. This is a Web-based tool that stitches together multiple photographs of a particular object into a three dimensional view. I tested its usability in two phases; firstly I asked users to explore a synth and answer questions that would stimulate thought about the conceptual model, and secondly I asked them to perform a complex task that would require a thorough understanding of this model. This two-phase approach allowed me to distinguish usability bugs that were due to misunderstandings of the conceptual model from usability bugs due to the interface's design.

**Keywords: usability test; unprecedented; conceptual model; Photosynth; Microsoft; user interface; design**

**Subject Terms: User interfaces (Computer systems) -- Testing**



To my parents,  
who made all of this possible

## ACKNOWLEDGEMENTS

I'd firstly like to thank my supervisor, Ted Kirkpatrick, for taking a chance on me by accepting my application to graduate school. Not only was I able to fulfil my desire of experiencing grad school and studying Human Computer Interaction, I found my career calling here: usability. I had fun at SFU and will cherish the time that I spent here. Thanks Ted! I'd also like to thank him for his insights into this project, and for editing the initial drafts so thoroughly.

Next, I'd like to thank my sister for picking up my slack at home this semester! Thanks *Roddels*! In return, I promise to put a *little* chilli in everything I cook!

To my friends in Denmark, who encouraged me to apply for grad school while sitting on another continent, and to my friends here in Vancouver, who saw me through this process: Mange Tak, Many Thanks!

Lastly, I'd like to acknowledge the influence my family has had on me to make me the person I am today. Without the support of my parents, sister, grandparents, cousins, aunts, uncles and the rest of my family, I would not be sitting here with the life and opportunities I have.

# TABLE OF CONTENTS

<b>Approval</b> .....	<b>ii</b>
<b>Abstract</b> .....	<b>iii</b>
<b>Dedication</b> .....	<b>iv</b>
<b>Acknowledgements</b> .....	<b>v</b>
<b>Table of Contents</b> .....	<b>vi</b>
<b>List of Figures</b> .....	<b>viii</b>
<b>List of Tables</b> .....	<b>ix</b>
<b>Chapter 1: Introduction</b> .....	<b>1</b>
Goals of this project .....	3
Brief History of Photosynth.....	5
A few features and screenshots of Photosynth .....	7
White Quadrilaterals.....	7
Point Cloud .....	9
Halo .....	11
<b>Chapter 2: The usability testing literature</b> .....	<b>13</b>
Reviewing the current state of Usability Testing.....	13
What is a Conceptual Model? .....	14
Visibility .....	15
Mapping .....	16
Feedback .....	16
Photosynth’s conceptual model and its role in its usability .....	18
Published Work Related to Photosynth.....	19
<b>Chapter 3: Design of a usability test of photosynth</b> .....	<b>22</b>
<i>Photosynth</i> – not <i>Photosharing</i> .....	22
A traditional usability test will give an inaccurate reflection of Photosynth’s usability .....	24
How does one capture a participant’s understanding of the conceptual model? .....	24
Getting the answer without asking the question .....	25
Give the participants a fair chance .....	26
To Task or not to Task.....	27
Define what constitutes success .....	27
Questions that would reveal a participant’s understanding of Photosynth’s conceptual model .....	28
Picking a synth that would allow users to answer my questions .....	31

Applying a conceptual model's understanding to a task.....	32
What if the users did not like the tasks?.....	34
Who would the participants be?.....	35
Test Equipment and Environment .....	37
Procedure .....	38
Task Completion Time.....	38
Loss of Audio Data .....	39
<b>Chapter 4: Results .....</b>	<b>40</b>
Task 1's Results .....	40
Task 2's Results .....	45
Task 3's Results .....	48
<b>Chapter 5: Recommendations.....</b>	<b>51</b>
Cannot Find a Halo .....	51
Users found unexpected zooms disorienting.....	55
Users unaware they could drag a halo.....	58
Users' inefficient navigation techniques .....	59
Problems understanding what the solid white region of halo is for.....	62
Lack of understanding and appreciation for the point cloud.....	63
Users misunderstand the white quadrilaterals .....	67
Users have no problem understanding basics of conceptual model .....	68
Can we trust a user's responses?.....	71
Users should be allowed to pick their own tasks .....	71
Problem with nonrandomized delivery of questions.....	73
Users were confident in their understanding of incorrect concepts.....	73
Summary of Recommendations.....	74
<b>Chapter 6: Conclusion.....</b>	<b>75</b>
<b>Reference List .....</b>	<b>78</b>
<b>Appendices.....</b>	<b>80</b>
Appendix A: TinyURL Referrals.....	81
Appendix B: Documents Presented to Users.....	82

## LIST OF FIGURES

Figure 1: A white quadrilateral a user can click on. They appear when the mouse hovers over the region. 3 corners have been highlighted in red, for clarity. In this example, the 4 <sup>th</sup> corner is not displayed to users. ....	8
Figure 2: The result of clicking on the white quadrilateral shown in Figure 1 .....	8
Figure 3: This is a point cloud of the Basilica di San Marco and Campanile di San Marco, as seen from the Piazza San Marco in Venice. It is composed of many individual dots.....	9
Figure 4: A Halo, as seen at the Palazzo Ducale (Doge’s Palace). A halo indicates that a user can rotate around the object.....	11
Figure 5: The halo, with its subfeatures highlighted. The yellow boxes indicate the limits of the accessible region of this particular halo, where the transition from translucent grey to an almost transparent grey occurs. The green box shows a green region of the halo. The red box indicates the solid white region. ....	29
Figure 6: A screen shot of the first synth that participants explored, the Notre Dame Cathedral in Paris .....	32
Figure 7: A screen shot of the second synth that participants encountered: a house whose floor plan was to be drawn by participants .....	33
Figure 8: A halo appears at the centre of the cathedral, when viewed from a distance ....	51
Figure 9: No halo appears at the centre of the cathedral, when viewed from close up, even though it did when viewed from a distance. ....	52
Figure 10: Halo above the table in the entrance lobby of the house – about to click on green region, expecting to view the table from a 180 degree direction change.....	55
Figure 11: After clicking on the green region highlighted in Figure 11, this is the view presented to users; outside the house and no table in sight. The open door is the front door to the house. ....	56
Figure 12: This halo confused participants because clicking on the green region resulted in seeing a partial view of the fireplace, from the lobby (see Figure 13).....	57
Figure 13: Suddenly, participants were moved into the lobby and shown half the fireplace .....	57
Figure 14: Translucent layers below the image that is in focus, where the mouse is located. The lower borders of the layers have been outlined in red to highlight them.....	60
Figure 15: Example of a view without the point cloud .....	66
Figure 16: Example of a view with the point cloud .....	66

## LIST OF TABLES

Table 1: Summary of participants used in the study .....	36
Table 2: Task Completion Times .....	39
Table 3: Results of Task 1's questions, arranged in descending order of severity .....	41
Table 4: Results of Post-Task1 Questionnaire .....	44
Table 5: Results of Post-Task2 Questionnaire .....	47
Table 6: Results of Post-Task3 Questionnaire .....	49
Table 7: Keywords used to describe what the white quadrilaterals represent .....	67
Table 8: Mapping of TinyURL to their corresponding, real, URL .....	81

## CHAPTER 1: INTRODUCTION

Imagine you are going through the equipment in your research lab and came across your supervisor's old time machine. The machine was built to allow the user a trip to the paranoia that was 1999. You decide to take it for a ride. Next thing you know, Smash Mouth is rocking the billboards.

In the midst of telling your 1999 colleagues of the wonders of the 21<sup>st</sup> century, you feel a buzzing inside your pocket. You realize that your iPhone is reminding you to pick up the groceries tonight (10 years early). Your 1999 colleagues (including your younger supervisor) are very curious about the device. They ask what it is and they all expect you to demo the device for them. However, you suddenly realize that this situation would make for a perfect usability test. Your results might even impress Jakob Nielsen<sup>1</sup>.

You tell the team to quit scrambling; that everyone shall get a chance to use the device in a *usability test*. You tell them that you will speak no further about the device in your hand, lest you reveal too much information and poison the group's naivety. The crowd grows silent, reflecting your solemn tone of voice. Your clearly impressed supervisor makes a mental promise to reward you with an RA grant in your final semester at school.

---

<sup>1</sup> <http://www.useit.com/jakob/> ; 26 March 2009

Everyone lines up outside the usability studio (you are familiar with the equipment because it's the same equipment your lab possesses in 2009). You have the magnetic-tape video camera running and call in the first participant, your supervisor.

"This...is a cell phone" you declare solemnly. "Please place a call to the following number" you say while about to hand the phone over to him. At this very moment you freeze, realizing that you will not get network coverage for an iPhone in 1999. Your usability test of the iPhone will need to focus on other tasks. You come up with a list of 3 possible tasks; play music on the phone, find a particular address on the GPS map you had downloaded and watch the finale episode of the TV show FRIENDS (for the bragging rights that it would entail in 1999).

To avoid errors due to ordering-effect, you randomly pick one of the three tasks, and it happens to be the map finding task. "This...is a cell phone, on which you can view a map of Metro Vancouver ... I mean the GVRD. Please open the map and point me to your home address." You now hand over the device to the participant for the first time.

The participant is experiencing a surge of emotions; awe at holding a device unlike any they have seen, excitement to play with the device, confusion that this cell phone has no dial pad, and all-of-the-above for a cell phone that contains a map of the city. The participant asks you who needs a map on a cell phone, how would he navigate through a giant map on such a small device and whether there is a built-in projector that will display the map on the studio's wall? You tell him that all shall be answered in the debriefing at the end of the session. For now, you only want to test the usability of the device, not other aspects of the user experience.



The screen is currently blank because of your screen-saver settings. The participant doesn't know where to start; where is the "on" switch? He is carefully holding the device by its sides, fingers carefully tucked away from the screen. He raises the device above his head to see its underbelly, thinking there might be an "on" switch there. He checks all the sides but finds nothing that resembles an "on" button. After a minute of investigation he is about to give up when his fingers accidentally touch the screen. The screen lights up and your earlier reminder about the groceries shows up. The participant wonders what it was that caused the screen to come to life. With gradual amazement, he asks aloud if the screen was, *really*, sensitive to touch. He gives his theory a try by touching the "Quit" button on the reminder. Sure enough, the message disappears. The participant has started to form a conceptual model of the workings of the device; not nearly the entire model, but the first inklings have started to occur. He has much to learn about the two-finger usage to zoom in and out of the map, but he has gained the most basic insight about the device...the touch sensitive screen.

### **Goals of this project**

The scenario above hopes to serve as an analogy for the project being reported here. The primary goal of this project is for me to gain experience conducting usability tests on software that is built on *unprecedented* and *non-trivial* conceptual models.

For example, if we were to, today, run a usability test on an email client, most participants would not be encountering the conceptual model for the first time. They anticipate an encounter with an "Inbox", a "Compose Page", perhaps even a "Contacts" list. These concepts might be non-trivial to some, but certainly not unprecedented; most

of today's major email clients are modelled around these concepts. The conceptual model is well engrained.

In fact, the conceptual model of a typical Email client is an extension of a bygone, offline era. For example, in the 1950s, office secretaries were well aware of our modern email client conceptual model. A typical manager in the 1950s had an inbox tray and an outbox tray on their desk. It was the secretary's responsibility to see that any letters or forms in the Outbox tray were mailed accordingly and that any new mail or documents that required the manager's attention were placed on the Inbox tray. Similarly, managers and secretaries had a book of contacts, with people's names, telephone numbers and postal addresses. The similarity between the old conceptual model and today's email conceptual model is quite apparent, and certainly not a coincidence. This is because the designers of early email clients copied a well known conceptual model in order for users to make an easy transition from the old work process to the, familiar, new work process.

Thus, a usability test for such interfaces, whose conceptual model is well engrained, is important but does not challenge the usability practitioner to use methods beyond their basket of past experiences.

Having the opportunity to perform a usability test on interfaces whose conceptual model is complex and unfamiliar is rare. Apple's iPhone<sup>2</sup> and iPod<sup>3</sup> presented us with such an opportunity because they introduced new methods of interaction. Google Earth<sup>4</sup> presented us with such an opportunity because the concept of being able to zoom into any

---

<sup>2</sup> <http://www.apple.com/iphone/> ; 26 March 2009

<sup>3</sup> <http://www.apple.com/ipod/> ; 26 March 2009

<sup>4</sup> <http://earth.google.com/> ; 26 March 2009

point on the planet was novel. Microsoft Surface<sup>5</sup> presented us with such an opportunity because the concept of performing familiar tasks such as paying bills, trip planning, sharing photos and music, all through an electronic table is an unfamiliar one for most people. And now, I believe, Microsoft Photosynth<sup>6</sup> presents itself as an *unprecedented* and *non-trivial* way of organizing and viewing large photo collections. I believe Photosynth is a perfect candidate for me to explore methods of testing the usability of a system which is, relatively, still unknown to the public and which is built around a revolutionary conceptual model. Therefore, this project shall focus exclusively on Photosynth. I shall not be comparing Photosynth to any competitors as that is not the purpose of this project. I think it would be naïve for me to come up with a guideline that other usability practitioners should follow when they test novel conceptual models, because that would require any author to spend years of waiting-around, for unprecedented conceptual models to be released to the public. Only after they have conducted usability tests on multiple *unprecedented* and *non-trivial* conceptual models could they be justified in writing a guideline for others to follow.

### **Brief History of Photosynth**

In 2006, researchers Noah Snavely and Steven Seitz from the University of Washington collaborated with Richard Szeliski, at Microsoft Research, to develop a system called Photo Tourism [Snavely et al. 2006]. This system allowed users to create 3D collections from their regular 2D digital photos. The authors envisioned this product to be used by people who recently returned from a holiday in, say, Rome, where they fell

---

<sup>5</sup> <http://www.microsoft.com/SURFACE/> ; 26 March 2009

<sup>6</sup> <http://www.photosynth.com> ; 26 March 2009

in love with the grandeur and detail of the Pantheon and wished to share their experience with relatives and friends back home. Currently, tourists would take multiple photos of the site and upload all of them, hoping that folks back home would feel the same emotions that they felt when viewing the real object. However, parsing through 2D images doesn't always elicit the emotions hoped for by the photographer. If these 2D images were stitched together and developed into a 3D image, henceforth referred to as a *synth*, viewers back home have a better chance of feeling the rush of emotions experienced by the photographer. To build a synth of value, very many photos of the object in question is needed. The person viewing the synth would need a very high speed internet connection and processor to handle the barrage of data being transferred to their machine. This quickly reaches impractical system requirements.

That is where a Seattle-based company called Seadragon<sup>7</sup> fits in. It developed an algorithm that allows users to view very large images over regular bandwidth and processor speeds. It does this by only transferring pixel data that the user is currently viewing on their screen, and not the data of the entire image. Hence less data is transferred, lowering bandwidth requirements while at the same time the processor does not have to process more data than it needs to.

Microsoft acquired Seadragon in 2006. The merger, along with the Photo Tourism project resulted in the Photosynth project.

In January 2009, Photosynth and CNN conducted a campaign called "The Moment", where they asked attendees of President Barack Obama's inauguration to take photos of the moment when the president was being sworn in, and to email these photos

---

<sup>7</sup> <http://livelabs.com/seadragon/> ; 26 March 2009

to CNN. The collection of images would then be converted into a synth that captures that historic moment as a three dimensional frieze. This campaign was a hugely successful public relations campaign for Photosynth; the event was advertised, for days, on CNN. The resulting synth was demo'ed to the public on air and has since been viewed over 240,000 times online. It also reveals that Microsoft intends Photosynth to be used by a very general audience, and not a specific section of the technology industry (such as computer vision experts).

### **A few features and screenshots of Photosynth**

Here I shall describe three features of the interface that, I feel, require the user to understand the conceptual model before they can appreciate what these features do.


#### **White Quadrilaterals**

The most frequent feature a user of Photosynth will encounter is a white quadrilateral (see Figure 1)<sup>8</sup>. These quadrilaterals appear on the synth as a user's mouse moves over the region that it encompasses. It fades away as soon as the user's mouse leaves the region or if the user brings the mouse to rest within the region. If the user left-clicks anywhere within the region, their vantage point shifts to reflect the orientation of the quadrilateral. In other words, they will shift their view to look 'bang-on' or directly at the quadrilateral. In Figure 1, this means the user's view will strafe slightly to the right and rotate slightly to the left, lining them up parallel to the quadrilateral (see Figure 2). However, there is no guarantee that the new view will maintain the same level of zoom, as can be seen in Figure 2. For Figure 1, this means that clicking within the white


---

<sup>8</sup> <http://tinyurl.com/awt67e> - see appendix for URL expansion

quadrilateral will line up the user parallel to the quadrilateral but might zoom the user out of the scene, rendering the region within the quadrilateral as a small part of a much larger scene.



**Figure 1:** A white quadrilateral a user can click on. They appear when the mouse hovers over the region. 3 corners have been highlighted in red, for clarity. In this example, the 4<sup>th</sup> corner is not displayed to users.



**Figure 2:** The result of clicking on the white quadrilateral shown in Figure 1

## Point Cloud



**Figure 3:** This is a point cloud of the Basilica di San Marco and Campanile di San Marco, as seen from the Piazza San Marco in Venice. It is composed of many individual dots.<sup>9</sup>

Point Cloud (see Figure 3) is a concept that is not unique to Photosynth. It is a general term used in mathematics to represent a set of points described by an  $x, y,$  and  $z$  coordinate. Each photo in the image set is divided into small segments. Images are linked together when the same segment is located on two or more images. In the point

---

<sup>9</sup> <http://tinyurl.com/awt67e> - see appendix for URL expansion

cloud, each point represents a segment where multiple images overlapped. In Figure 3, this means that there were multiple photographs that contained the Basilica di San Marco and the Campanile di San Marco. There would have been photographs that contained people, but none of those people appeared in multiple photos; hence they are not considered part of the model (or Point Cloud) of the Piazza San Marco.

Each point in the point cloud takes on the average colour of the segment that it belongs to. Thus regions which are heavily photographed will have a more accurate colour than points from less frequently photographed segments because less frequently photographed regions will result in sparser dots and less data, making Photosynth unsure of the actual colour of the region. In Figure 3, that results in the lower-central arc of the basilica to be coloured more accurately than, say, the archway on the extreme left, because everyone takes a photograph of the ornate, central archway. There is more data for the central archway and thus the average colour will reflect the real colour more accurately.



## Halo



**Figure 4:** A Halo, as seen at the Palazzo Ducale (Doge's Palace). A halo indicates that a user can rotate around the object<sup>10</sup>

At various points in a synth, the user will be presented with a doughnut-like widget called a Halo (see Figure 4). The halo allows users to rotate around the object that it is attached to. However, the extent to which a user can rotate around an object (and thus the range of the halo) depends on the number of photos provided to the algorithm. If

---

<sup>10</sup> <http://tinyurl.com/awt67e> - see appendix for URL expansion

there are no images, or unrecognizable images, of an object from certain angles, then the user is not able to rotate the halo to those angles. In Figure 4, for example, the user is unable to view the piazza from the vantage point of the palace's balcony because there are no images of the piazza from those angles. Hence, in this case, the halo will not allow the user to rotate beyond a certain range. The inaccessible angles of a halo are *almost* transparent, appearing very faint to the user. The currently viewed angle appears as an opaque white region on the halo. If the user moves their mouse over the accessible sectors of a halo, each possible vantage angle will appear in green while the mouse is hovering over it. All remaining, accessible angles of the Halo appear in a translucent grey colour.

## **CHAPTER 2: THE USABILITY TESTING LITERATURE**

In this chapter, I shall first review the current state of a few aspects of usability testing that I use in my study. Following this, I shall discuss work related to defining what a conceptual model is and how to test complex conceptual models. Finally, I shall review some published work on Photosynth.

### **Reviewing the current state of Usability Testing**

Researchers and practitioners have been wary that a usability test is constrained by how much experience a participant has had using the system being evaluated [Kellogg et al 1987; Spolsky 2001; Whitehead 2006; Nielsen<sup>11</sup>; Tohidi et al. 2006; Bruno et al. 2005 ]. They claim that a usability test might turn into a ‘learnability’ test of the system: how easily can people learn to use the system, rather than how easily could people use the system.

Utility, or usefulness, is a feature that practitioners consider important when they evaluate a system [Wilson 2007; Redish 2007; Howard 2008; Dicks 2002 ]. After all, there is no good having a usable, but useless, system. I was especially intrigued by Wilson’s [2007] suggestion that we ask participants to pick tasks for a usability test. In my study, I asked three of my participants to pick a synth of their own to explore, and to then provide subjective feedback which I compared to feedback received for synths that I had picked for the predetermined tasks.

---

<sup>11</sup> <http://www.useit.com/alertbox/20030825.html> ; 26 March 2009

Work on usability testing techniques is infrequently published in peer reviewed conferences and journals, though some interesting insights can be found in Nørgaard and Hornbæk [2006] where they discuss how usability practitioners conduct think aloud testing, and discuss the implications of these current practices. Tohidi et al. [2006] informs that participants tend to be less critical in an evaluation when presented with only one design for a system. They suggest that in a usability test, practitioners present participants with multiple design alternatives. Tang et al. [2006] provide a useful list of the advantages and disadvantages for recording screen activity and audio during a usability evaluation. Based on their suggestions, I chose to record screen and audio proceedings of the evaluation sessions. Finally, Hard et al. [2008] discuss their finding that traditional usability design guidelines are being grossly violated on the popular, online, social networking site, Facebook. To their surprise, in spite of these violations, most user experiences on Facebook are positive. They suggest that perhaps it is time we update our notions of what constitutes a usable design.

Usability testing is gaining popularity [Chauhan 2006], but an undercurrent of caution has started to flow, warning that not all situations warrant a usability test [Kujala et al. 2000; Molich et al. 2008; Greenberg and Buxton 2008]. In addition, Hertzum and Jacobson [2003] warn readers that performing a think aloud usability test with only one evaluator will almost certainly lead to many usability defects going unnoticed.

## **What is a Conceptual Model?**

Don Norman describes a conceptual model as the knowledge a user has about how a system works. Based on this model, the user can simulate or imagine how a system works, thus predicting the results of their actions [Norman 1990, pp12]. Norman

states that in order for us to gain a good understanding of how the system works – to understand the conceptual model – the interface must make options visible, the controls on the interface must be mapped to its function in a predictable fashion, and the interface should provide feedback to the user’s input. I shall now, briefly, describe these attributes.

### **Visibility**

Quite often, users have a general idea of what needs to be done in order to complete a task on a system. However, translating that idea into an action is not always easy. If users cannot see the controls that they need, it hinders understanding how the system works because they start to doubt the validity of their desired action. They might think that the system does not allow them to perform certain tasks, simply because they could not find the control that would allow them to perform it. This reflects badly on the usability of the system. This sort of situation, where users cannot figure out how to perform an action, often leads to frustration. If this is the first time that the user is interacting with the system, their initial impression of the system will be a negative one, which is often hard to change [Hartmann et al. 2008]. The interface should make as many available options visible as possible. However, we must keep in mind that overwhelming the user with various options needs to also be kept in check. Ultimately, the interface designer needs to decide how many options need to be made available for a particular system, and which options those should be.

## **Mapping**

When a user encounters a control, they need to be aware what function that control performs. If users do not realize a particular control is the one they need, it is a sign of poor usability. An example of good mapping is having a left arrow move the user left. An example of bad mapping is having a square move the user left. A control and its action need to be mapped in a natural fashion, which often is easier said than done.

## **Feedback**

Feedback is also essential when trying to perform a task. A user needs to know whether the action that they just performed, was registered by the system. If this feature is missing, usability is often compromised.

Norman also defines three terms – design model, user model and system image – that have been used by many to describe conceptual models. The design model is a model of how the designer intends the user to interpret and interact with the system. The user model is the model that a user *actually* follows, when interacting with an interface. The system image is the implemented system, which is responsible for conveying the design model to the user. A designer hopes that the system image is able to convey the design model to a user – for them to create a user model that is, ideally, identical to the design model. However, due to deficiencies in visibility, mapping and feedback, among other causes, the design model and the user model seldom match up completely [Norman 1990, pp16].

Aside from Norman, Preece, Rogers and Sharp [2002] present tips for clueing the user into the design model. They talk about conceptual models for activities involving navigation and manipulation of virtual spaces: pertinent to my study of Photosynth. They

advise that we exploit the real world experiences that users are familiar with. Thus we should be able to move to any location in the 3D space, move left, right, forward, backward, zoom into objects, and zoom out of objects, because users perform these functions in the real world and thus expect to be able to perform them in the virtual world. Kellogg and Breen [1987] claim that, with experience, the overlap between the user model and a design model increases. They use the term system model to mean the design model. Their work “involved specifying a system model, deriving user models from groups of users with different amounts of experience with the system, and developing methods for assessing the degree of agreement between users’ and system model.” [Kellogg and Breen 1987]. However, they did not carry out the test on a real system. They presented participants with 51 index cards, each with a verbal label of a text editing tool. Participants were asked to group similar tools together. The number, and content, of groups was noted by the moderator. Through this procedure, the moderator could understand how far or close the participant’s understanding of the tools was from the designer’s model. The problem with such an approach is that participants cannot experiment with a command to understand its function. Their understanding is based entirely on the text printed on a card, and on information that a moderator would provide. I feel that watching participants explore a working system, without providing them with any help, gives a much clearer understanding of real world usage of a system, because in the real world, the user would have to interact with the system without having a moderator by their side. My study compares how much of an overlap exists between a user’s model of Photosynth and the design model of Photosynth, similar to Kellogg and Breen’s goal. However, unlike Kellogg and Breen, I will be conducting the study on the

actual system, not on a paper mock-up or on printouts, for the previously mentioned reason and also because Photosynth does not have commands that users are already familiar with. Thus, even if I did provide users with verbal descriptions or screenshots of Photosynth's controls, it would be an unfair assessment of its usability, as most participants would not have encountered such controls before.

Hsieh et al. [2008] present a system called Photo Navigator, which is similar to Photosynth, in that it allows users to browse their digital photographs in a 3D environment. However, in this system, users do not control or manipulate the environment. It is a completely automated environment, based on a rendering determined by the system. Thus, they do not need to understand the conceptual model to view their photos. They do, on the other hand, need to understand the conceptual model when they are capturing photographs of the scene, as their actions at that time will have implications on the resulting automated slideshow. Photosynth, in fact, has such a feature: the play button on the top right of the screen. What I found interesting from the Photo Navigator paper were the questions that the authors asked participants, during a subjective evaluation of their system. They asked participants to rate Visual Perception, Smoothness, Spatiality, Acceptance and Experience. I found these questions quite useful for a subjective evaluation of a 3D photo viewing environment: a system with a complex conceptual model.

### **Photosynth's conceptual model and its role in its usability**

At its core, understanding Photosynth's conceptual model entails figuring out what actions can be performed in a synth, how does one execute those actions and how does Photosynth work, internally, to accommodate the user. Answers to these questions



depend on a user's prior experience with similar systems. For example, proficient gamers *know* what to expect when they play a game for the first time, because they have preconceived notions of what to expect in a video game, in terms of what actions they might be able to perform (strafe to the left, change weapon, etc.) and how they can perform those actions (press the left arrow key on the remote, press the 'A' button, etc.). However, unlike in video games, Photosynth's users need to have a high-level understanding of how the system works internally. They need to know that they are looking at a collection of photographs that have been stitched together to form a 3D view. More importantly, they also need to be aware that they are limited to view a scene from a *discrete* set of points (i.e. photographs). Unlike in a 3D video game, where players can freely move to any point in the 3D space, Photosynth's users are limited to explore the scene from only those angles from where a photo exists. Preconceived notions, inspired from video games, hamper their conceptualization of how Photosynth internally works.

A test of Photosynth's conceptual model requires the usability professional to test whether participants are a) able to figure out the fact that they are viewing multiple photos and that they understand the limitation of viewing the space from only a discrete set of points, b) able to predict the action of controls within the interface, c) able to find, on the interface, controls that they are looking for, and d) aware that the action they just performed was registered by the system.

## **Published Work Related to Photosynth**

Photo Tourism [Snavely et al. 2006], research from the University of Washington in collaboration with Microsoft Research, is the precursor to Photosynth. This 2006 paper presents an early, yet high fidelity, version of Photosynth. It is referenced on the

Photosynth website as part of Photosynth's history<sup>12</sup>. It “consists of an image-based modelling front end that automatically computes the viewpoint of each photograph as well as a sparse 3D model of the scene and image to model correspondences” [Snavely et al. 2006]. The 3D model is the point cloud, and one of the contributions of this paper is the technique to create the point cloud from various photographs of the scene, taken from various angles and depths. Aside from this technique, the paper describes interface controls that allow users to navigate within the 3D environment. Interestingly, some of the interface features displayed in Photo Tourism, are not present in Photosynth. For example, in Photo Tourism, users are able to drag an outline around an area of the scene that interests them, and the system smoothly transitions to a detailed photo of the highlighted region, if one exists. Also in Photo Tourism, users are able to create and view metadata, or tags, throughout the 3D space. It would be interesting to learn why certain features were dropped in the current version of Photosynth.

In 2008, the same authors, along with Rahul Garg from the University of Washington, published work that is meant to be an advancement of Photo Tourism and Photosynth: “The work presented here is a significant advance over Photo Tourism and Microsoft's related Photosynth” [Snavely et al. 2008]. In this paper, the authors admit some of the usability shortcomings of Photo Tourism and Photosynth:

“key limitations of these previous methods [Photo Tourism and Photosynth] ... [which are] based on the assumption that the scene is well approximated by a planar facades, enabling good results on scenes like the Notre Dame Cathedral, the Trevi Fountain, and Half Dome. These same techniques break down, however, for general objects that contain many sides (e.g., statues, monuments, people, plans, etc.) and for large rotational

---

<sup>12</sup> <http://livelabs.com/photosynth/history> ; 26 March 2009

motions. In addition, navigation in Photosynth and Phototourism is based on the user selecting a photo and moving to it, and does not support the fluid, free-form 6-DOF navigation capabilities common in games and other interactive 3D applications.” [Snavely et al. 2008].

Aside from work by Snavely, Garg, Seitz, and Szeliski, I could not find other work that discusses Photosynth’s interface or usability. Since Photosynth only came out of its beta version, August 20, 2008<sup>13</sup>, I suspect more research on its usability will ensue in the coming year. Through my report, I hope to be one of the first to shed light on Photosynth’s usability.

---

<sup>13</sup> <http://livelabs.com/blog/welcome-to-photosynth/> ; 26 March 2009

## CHAPTER 3: DESIGN OF A USABILITY TEST OF PHOTOSYNTH

In this chapter I shall describe the design of my test of Photosynth's usability.

### ***Photosynth – not Photosharing***

Photosynth differs from other photo sharing websites in fundamental ways.

Firstly, Photosynth's primary purpose is not to share *just any* photograph with your family and friends. Sure, you are sharing photographs with whoever views your synth, but the primary goal here is to share an experience, not a view. On Photosynth you upload photos of a specific environment; your bedroom, the Pantheon, a view of the city. You should not upload photos of your 19<sup>th</sup> birthday party, all the photos from your vacation to India (though a subset of, say, the Taj Mahal would be a good idea) or a single beautiful photo of a sunset. The latter are examples of bad Photosynth usage because there is likely to be very little subject overlap between photographs, rendering it impossible for Photosynth to stitch them together, which in turn misses the purpose of Photosynth: to view your images in a 3D setting. Currently, people are used to uploading *all* their photographs to a single website. Photosynth asks users to rethink this traditional model by uploading multiple photos of a particular subject.

Secondly, Photosynth differs from traditional photo-sharing websites by the fact that users must know that the photos they are about to take will, specifically, be uploaded to Photosynth. They need to know this because they have to take multiple photographs

from multiple angles if they want a smooth synth to be generated. They should also take close-up shots of details to add depth to their synth. Typically, these photos should not be photos of the user or their family and friends. For traditional photo-sharing, no such rules need to be kept in mind when shooting photographs; your photos will be viewed one after another, a well engrained conceptual model for online photo viewing. While this is an important concept to evaluate in a usability test of Photosynth, I will not be covering it in this study because it would require time commitments that I cannot afford.

Thirdly, Photosynth differs from traditional photo-sharing websites because users need to know how to navigate through a synth. The process of viewing the different photos in a synth is more complicated than on, say, Flickr<sup>14</sup>. Users are presented with white quadrilaterals, halos and point clouds; unfamiliar concepts for most users. None of the popular photo-sharing websites use such controls. Hence users need to figure out how these controls work, and I argue that they need to understand *why* these controls work the way they do. The former is tested through a traditional usability test. We ask a user to perform a particular task, and see whether they knew *how* to use the interface. On the other hand, the latter question of *why* a control works the way it does, is often ignored because participants are usually familiar with the conceptual model they are testing. In my study, I test whether participants can understand *why* Photosynth's controls work the way they do; which is equivalent to testing their understanding of the conceptual model.

---

<sup>14</sup> <http://www.flickr.com/> ; 26 March 2009

## **A traditional usability test will give an inaccurate reflection of Photosynth's usability**

The most common approach to a usability test is a task-based approach; moderators ask participants to complete a series of tasks and judge usability on objective and subjective feedback. This method works well for software whose conceptual model users are familiar with. However, for unfamiliar conceptual models, this method would not give a true reflection of the system's usability because participants have to think of a solution to, what might be a familiar task, in a novel way.

Understanding a system's conceptual model is a very important usability hurdle that users must overcome. We, as usability practitioners, should test a user's understanding of the conceptual model as rigorously as any other aspect of the system's usability. Giving away a system's conceptual model during a usability test would be careless of the practitioner. Unless we are sure that users of the software would receive some sort of training, spelling out the conceptual model would skew our test of the system's usability.

## **How does one capture a participant's understanding of the conceptual model?**

This is a tricky question to answer because a practitioner should not reveal the fine details of their study until after the session, lest participants try to perform in a manner that they think helps (or harms) the study. By keeping motives vague, we can limit this behaviour. Also, the term "conceptual model" might not be understood by participants, or understood differently by different participants.

The challenges at this stage were:

- 1) How do I determine the participant's understanding of the conceptual model, without explicitly telling the participant what I was looking for?
- 2) How do I give participants a chance to grasp the conceptual model before I quiz them on it?
- 3) How do I give participants enough *time* to learn how to use Photosynth before I can honestly judge Photosynth's usability [Spolsky 2001]?
- 4) How do I define a user's understanding of the conceptual model as being successful? Different users would probably understand the conceptual model to different degrees

I shall now outline my solutions to each of these challenges.

### **Getting the answer without asking the question**

This was perhaps the key challenge. I would find out that the solution to this challenge led to a solution for the other challenges.

I did not want to directly ask participants how they thought Photosynth works, because 1) that results in answers that can be interpreted in many ways, and 2) I did not want participants to think extra hard about the answer because that would not reflect a real-world situation. I wanted participants to answer the question as naturally as they would if they asked themselves this question.

I decided I would ask users to answer certain questions about three controls that Photosynth uses; White Quadrilaterals, the Halo and the Point Cloud. If users *truly* understood Photosynth's conceptual model, they would be able to answer my questions about these controls.

So this adds a 5<sup>th</sup> challenge: draw up a list of questions regarding the white quadrilaterals, halo and point cloud that would get participants to reveal their understanding of the conceptual model.

### **Give the participants a fair chance**

It is important to give participants a fair opportunity to understand how Photosynth's controls work. I initially considered showing paper printouts of the three controls (white quadrilaterals, halo and point cloud). I wanted to see if participants could understand what these controls would be used for, purely from their appearance within the interface. If participants were able to guess that a white quadrilateral represented a particular viewpoint, by simply looking at one and not interacting with it, this would bode well for Photosynth's overall usability. They seem to 'get' the conceptual model. However I quickly discarded this methodology because I realized that for point clouds or even the halo, just looking at a printout of these controls would not be enough to guess their actions, and more importantly this is not a fair method for letting users understand a system's conceptual model. In the real world, people would interact with the system, not stare at it till they guess a control's function. In addition, presenting paper printouts of the controls draws explicit attention to the controls, which might not be the case in real life. For example, in reality, a user might not encounter a halo until ten minutes into an exploration. In those ten minutes they might have gained an inkling into the conceptual model. Thus showing a halo on paper right at the start of a test session might not be a fair representation of reality.

Hence I decided to allow users to interact with the system while trying to come up with the answers to the questions I ask of them.



### **To Task or not to Task**

A task limits a participant's ability to explore the system *as much as they want to*; something that is essential if they are to understand the conceptual model. For example, even if a participant uses every control in the interface once, it might not be enough exposure for them to understand what is going on. Since different users would require different amounts of exposure to the interface, in order to understand the conceptual model, I decided not to ask users to perform a specific task. Instead I decided that users should be given free reign to explore a synth for as long as they wish to, in order to answer the questions I wanted them to answer. At the same time, I thought it would be fair to set an upper time limit for the exploration because outside of a lab setting, users would give up if they couldn't figure out the answers easily. I decided to give them 30 minutes to explore a synth and come up with responses to my questions.

### **Define what constitutes success**

Once I knew that participants would be exploring a synth and writing down their responses to my predefined questions regarding various controls in the interface, I needed to consider how I could compare these various answers. This was simple. The questions I ask them should have a clear *correct* answer; an answer that does not leave much room for interpretation. This way I could say, with certainty, whether a particular participant answered the question correctly or incorrectly; which in turn would imply whether they understood how and why a control acted the way it did; which in turn implies whether the participant understood the underlying conceptual model; which in turn reveals how usable Photosynth is.

### **Questions that would reveal a participant's understanding of Photosynth's conceptual model**

Here I shall provide a short justification for each question that I asked participants.

#### **What do the white quadrilaterals represent?**


The white quadrilaterals represent the different photographs in a synth. If a user was able to figure this out, they got the biggest hint about the underlying conceptual model.

#### **What is the Halo's function?**

The halo's function in the interface is to allow a user to view an object from multiple angles, easily (see Figure 5). Instead of walking around a scene, users can click on the halo to spin to a new viewpoint. This is one of the main controls a user will encounter. I wanted to see if participants would map its appearance to its function.

#### **When do you expect to see the Halo?**

The answer to this question is that you expect to see a halo when multiple views of an object, from different angles, are available. I asked this question in order to understand whether participants could understand *why* a halo appeared at one point in the synth and not at another.



**Figure 5:** The halo, with its subfeatures highlighted. The yellow boxes indicate the limits of the accessible region of this particular halo, where the transition from translucent grey to an almost transparent grey occurs. The green box shows a green region of the halo. The red box indicates the solid white region.

### **Can you rotate around the church? Why?**

I formed this question after picking the synth that I wanted to present to participants. In the synth I picked, users were not able to move 360 degrees around the church because there were no photos of the church from the back. The halo will always appear as a complete torus, but the inaccessible regions (in this case the region of the halo that represents the back of the church) appear almost transparent (see Figure 5)<sup>15</sup>.

I picked this question because it tested a participant's basic understanding of the conceptual model. If they knew that a synth was just an amalgamation of photographs,

---

<sup>15</sup> <http://tinyurl.com/ceaqlv> - see appendix for URL expansion

then they should have known that not being able to view the church from some angles was because there were no photos of the church from those angles.

**What do the green region and the solid white region on the Halo represent?**

This was asked as two separate questions, but for the reason of similarity, I explain them together. The green region of the halo, which appears under the mouse as you hover over the halo, represents an available angle from which you can view the object. The solid white region of the halo represents the angle via which you are currently viewing the object (see Figure 5).

I picked these two questions because, again, they test whether these participants could map these features to their function.

**What does the size of the green region on the Halo represent?**

You will never have photos of every angle of an object. There will always be discrete angles for which a photo of the object does not exist. The green regions on a halo (i.e. the areas that turn green when the user moves their mouse over the halo) have varying widths, which convey that for all the angles covered by a particular region, only one photo of the object exists.

This question tests whether participants were able to map this feature to its function.

**What do the scattered dots [i.e. the Point Cloud] represent?**

I asked this question because the point cloud is a model that Photosynth derives by analyzing the various photos presented to it. I was curious whether participants would

able to guess that the system looked at each photo and came up with a model of the space. I was also curious whether participants found any use for a point cloud.

### **Why do people who appear at one point in the synth disappear all of a sudden?**


Since a synth is made up of various photographs, some people who appear in one photograph would probably not be in another photograph. If participants could figure out why this was happening, it would be clear that they understood Photosynth's basic conceptual model.

### **Picking a synth that would allow users to answer my questions**

Synths come in all sizes and shapes: you can have a synth made up of five photos, or of over six hundred images. The number of images does not indicate how well connected a synth is; i.e. you could have six hundred photographs with very little in common to each other, forming a very loose synth. On the other hand you could have a synth of five photographs which all have enough in common to be joined to one another. The level of cohesiveness is given a percentage rating, termed 'synthy'. For example a 5% synthy rating is bad because very few images were able to be joined together. On the other hand, a 90% synthy rating is good because most images were able to be stitched together, giving a smooth and elaborate space to explore.

I wanted to pick highly synthy synths. I did not want to present participants with a low quality synth because that would not be fair to participants, who would be cheated of potentially learning the system smoothly and accurately.

I picked a synth that was featured on Photosynth's homepage. Given that it was advertised so prominently, it was fair to participants and to Photosynth. It was a synth of the exterior of the Notre Dame Cathedral, in Paris (see Figure 6)<sup>16</sup>.



**Figure 6:** A screen shot of the first synth that participants explored, the Notre Dame Cathedral in Paris

This synth had a 96% synthy rating and was composed of 512 photographs. It provided participants with ample opportunities to interact with halos, white quadrilaterals, and the point cloud.


### **Applying a conceptual model's understanding to a task**

I decided that a second task would be required to allow a participant to apply their understanding of the conceptual model to a task.

---

<sup>16</sup> <http://tinyurl.com/ceaqlv> - see appendix for URL expansion

This second task was a traditional usability test, in the sense that I presented participants with a task and measured how fast and accurately they perform it.



**Figure 7:** A screen shot of the second synth that participants encountered: a house whose floor plan was to be drawn by participants

I wanted a task that would require participants to use various controls of the interface, multiple times.

I decided on having participants draw a floor plan of the ground floor of someone's house (see Figure 7)<sup>17</sup>. Participants were asked to pay particular attention to the size of rooms (to draw them relatively sized to each other) and to note any furniture they might see in a room. I decided to go with this task because Photosynth has the potential to be used by real estate agents to conduct virtual tours of a house; and it was important that clients be able to navigate through a synth and get a sense for the orientation and size of rooms.

---

<sup>17</sup> <http://tinyurl.com/bcjlro> - see appendix for URL expansion

This task had a clearly defined goal: draw a floor plan. It was also a complex task for participants to perform, given that this was only the second synth they would have ever experienced. Also, it would require participants to use controls like the halo and the white quadrilaterals extensively. If a participant understood the conceptual mode, they would be able to navigate around the house easily, completing the task faster and more accurately than participants who did not understand the conceptual model.

### **What if the users did not like the tasks?**

In my study, I was concerned that participants might not find drawing a floor plan or exploring a Parisian church motivating. I would much rather have a participant work on a task that engages them, than have them work on a task they are doing simply because they are being told to do so. The latter could result in participants giving up prematurely and not taking the task seriously.

To simulate a *real world* scenario more accurately, I had participants search the existing database of available synths, via the Photosynth website, for a synth that they would be curious to explore. I gave them a 15 minute *play* session to find and play with as many synths as they wish. Participants were allowed to quit playing before the 15 minute mark if they so desired.

After this task, I had a questionnaire asking participants which of the various synths they encountered during the entire session (including the two synths I had chosen) was the most *fun* for them. I also asked participants which was the most revealing about how *Photosynth works*.



Unfortunately I only added this task after the 2<sup>nd</sup> participant's session; meaning I had 3 subjects take this task. Comments made by the second participant, regarding not finding the synth of Notre Dame exciting enough (she wanted to go inside the church too) made me realize that I had not compensated for this legitimate concern.

### **Who would the participants be?**

I wanted to have participants represent users that the Photosynth team at Microsoft predicted would use their system. Without access to members of the Photosynth team, the only way I could learn this was through information I could find online.

The biggest advertisement campaign for Photosynth was held on Barack Obama's inauguration day, Jan 20, 2009. Photosynth, in collaboration with CNN, hosted a campaign called The Moment, where they asked *anyone* who was watching the inauguration ceremony, those at the National Mall and those at home, to take a photograph that captures the moment when Obama was sworn into office. Photos taken live at the event were joined to create a synth of the moment, the spot in time, that Obama was sworn in. The photos that were taken at homes across the country were joined to form a synth where the only commonality between the photographs is an image (on television) of Obama's face while taking the oath. This synth was more of a slideshow that showed the scene at various homes at the moment that Obama took his oath.

It is fair to assume that those who took photographs for the campaign would want to view the resulting synth. Hence the synth needed to be "usable" by people of various ages, tech-savvy skills and professions. In the case of The Moment campaign, people

who took the photographs would probably know that the resulting synth is an amalgamation of individual photographs. The folks at Photosynth could assume that in this case, *that* part of the conceptual model would be known to the users. However, I consider this case as an exception; most people who view a synth for the first time would not be told before hand how a synth is created. Hence my test of the role the conceptual model plays in usability is justified.

I wanted to test Photosynth’s usability on people of nontechnical educational and professional backgrounds. I suspected that it would be easier for computer scientists to understand the conceptual model than it would be for non-computer scientists.

**Table 1: Summary of participants used in the study**

<b>Participant Id</b>	<b>Age</b>	<b>Profession</b>	<b>Gender</b>
1	31	Painter	Male
2	44+	Career Consultant	Female
3	44+	Translator	Female
4	20	Arts Student and Copywriter	Male
5	54	Disc Jockey	Male

In order to recruit participants, I placed an advertisement on Craigslist<sup>18</sup>. I decided to pay \$25/hour for a two hour session. I received approximately 130 responses in a 48 hour period. The lesson I learnt here (loud and clear) was that \$25/hour in a time of economic uncertainty will garner a huge response!

I decided to hire only 5 participants for this study. It has been said that 5 participants is often enough to identify most usability bugs in a system<sup>19</sup>. However, I do admit that practitioners who perform tests on novel conceptual models should hire more than 5 participants for their study. In my case, due to time constraints, I thought 5 participants were the most I could accommodate.

See table 1 for a summary of the participants.

## **Test Equipment and Environment**

I decided that an audio transcript and a screen recording should be stored for each session. This would enable me to look back at a session when analysing the results. I also concluded that I would not require video of the user's facial expressions. Camtasia Studio 6.0 was used for the audio transcript and screen recording<sup>20</sup>. The test was carried out on a 3.39Ghz, Pentium 4, Windows XP machine. A 24inch flat screen monitor was used.

The test was carried out in our usability study room. Each session consisted of one participant and the moderator. There were no observers for any of the sessions.

---

<sup>18</sup> <http://vancouver.craigslist.org> ; 26 March 2009

<sup>19</sup> <http://www.useit.com/alertbox/20000319.html> ; 26 March 2009

<sup>20</sup> <http://www.techsmith.com/camtasia.asp> ; 26 March 2009

## **Procedure**

The five sessions were carried out over 3 days. Participants were first informed about the nature of the tasks they would be performing during the session. They were then asked to read and sign the consent form. Following the consent form, I would start the audio and screen capture recording. Participants were next asked to fill out a pre-session questionnaire. I would then read them the description of Task 1, and hand them a copy of the description, which included the 9 questions. I also provided participants with a sheet where they could write down their answers. They were presented with the synth of the Notre Dame Cathedral for this task. After completing this task, participants were asked to fill out a post-task1 questionnaire. I would then read the description for Task 2 and hand them a copy of the description. Participants were also provided with a sheet where they could draw their floor plans. Upon completion of this task, participants were asked to fill out the post-task2 questionnaire. At this point, Participants 1 and 2 had completed the session, and were paid \$50. Participants 3, 4 and 5 were then told an unscripted description of Task 3 and were then allowed to search for a synth of their choice. Following Task 3, participants were asked to fill out post-task3 questionnaire. This completed the evaluation session, and participants were paid \$50. See Appendix B for the documents presented to participants.

## **Task Completion Time**

As can be seen in Table 2, Task 1 took, on average, 30 minutes to complete. I predicted that half an hour would have been enough time for participants to either answer the questions or to give up, but Participants 2 and 5 were willing to try longer and as I

was not short on time, I obliged. They were also the only two participants to give up on Task 2. Participant 5 chose to skip Task 3 altogether due to fatigue.

**Table 2: Task Completion Times**

<b>ID</b>	<b>Task 1 (min)</b>	<b>Task 2 (min)</b>	<b>Task 3 (min)</b>
1	24	25	Not Asked
2	40	30*	Not Asked
3	30	23	14
4	18	20	13
5	43	48*	Participant Chose To Skip Task

\* Participant chose to give up on the task after mentioned duration

## **Loss of Audio Data**

During the analysis phase of my study, I realized that the Camtasia recording of the audio of each session was corrupted. Roughly half an hour of each session had no audio. Unfortunately, for Participant 5's session, the entire recording had corrupted audio. The corruption resulted in white noise being played during the affected regions of the recording. I contacted Camtasia's helpdesk. They acknowledged that this was a known bug in Camtasia Recorder, when recording lossy content. They admitted that there is no recovery for my lost data.

## CHAPTER 4: RESULTS

In this chapter, I shall outline the results of the three tasks I ran during my study of Photosynth's usability.

### Task 1's Results

Task 1 was where participants explored the Notre Dame Cathedral while trying to answer the 9 questions mentioned in Chapter 3. All participants attempted each question, but not all participants could, ultimately, come up with an answer to every question.

Table 3 provides a summary of Task 1's results. In the table, I have arranged columns such that the most urgently required usability fixes are on the left hand side of the table, while the least urgent fix - in fact *no fix* - is listed on the right hand side of the table. Each participant also provided a confidence rating to each of their answers. The average confidence for each question, across all 5 participants, is included in Table 3. For example, Participant 1 answered question 6 incorrectly, Participant 5 could not come up with an answer to Question 6, and Participant 4 answered Question 7 correctly.

Table 3: Results of Task 1's questions, arranged in descending order of severity

	Most Critical									Least Critical			Not a Problem
	Q3	Q6	Q8	Q7	Q5	Q2	Q1	Q4	Q9				
	When do you expect halo?	What is green region's size?	What does point cloud represent?	What is solid white region?	What is green region?	What is halo's function?	What are white quad.?	Can you move 360° around church?	Why do people disappear?				
P1	X	X	X	X	O	O	φ	O	O				
P2	No Idea	X	No Idea	X	O	O	φ	O	O				
P3	X	X	No Idea	No Idea	O	O	φ	X	O				
P4	O	O	O	O	O	O	φ	O	O				
P5	X	No Idea	X	O	O	X	φ	O	O				
Confidence (self reported)	2.6	3.0	2.4	3.0	3.8	4.2	4.2	3.0	4.6				
Struggle (rated by moderator)	5	4	3.5	3.5	1	1.5	0	1	0				

X = Incorrect answer; φ = Incorrect answer, but only on a technicality; O = Correct answer; No Idea = Participant could not come up with an answer. Confidence and Struggle rating out of 5 (0 = none; 5 = extreme)

All five participants figured out that they were looking at a set of photographs that were stitched together. Evidence of this can be seen in Table 3, where all but one participant was able to answer question 4 and all participants were able to answer question 9, two questions that tested a participant's basic understanding of this part of the conceptual model. In fact, the speed and confidence with which they asserted their understanding of the system surprised me.

However, it was evident that this basic understanding of the conceptual model did *not* guarantee successful reasoning for the actions of controls on the interface. It can be said that only Participant 4 had a good understanding of the conceptual model. Participants 1, 2 and 5 got 5 of the 9 questions wrong, or had no idea what an answer should be. Participant 3 did worse, not knowing, or answering incorrectly, 6 of the 9 questions.

The results for question 1 were interesting. All 5 participants mentioned that the white quadrilaterals represent different views of the scene. Even participant 3 mentioned that the white quadrilaterals allowed you to "expand and view what you are looking at". However, the reason I marked this answer wrong for all participants was because no participant mentioned that the white quadrilaterals represented available, individual photos of the scene.

On the other hand, most participants understood the halo's function in the interface (question 2). Participant 5 got this question wrong because he thought a halo would only allow him to travel 180 degrees around an object, as opposed to 360 degrees.



Question 3 was the hardest for all participants. Only Participant 4 got it correct. The others struggled the longest on this question, mostly because it isn't possible to access a halo at will.

Question 6, asking what the size of the green region of the halo represents, was another common problem for participants. While participants did not spend as much effort on this question as they did to answer question 3, again it was the same four participants who got the answer wrong. An interesting point to note here was that, for this question, participants claimed to be more confident about their answers, which might explain why they did not spend as much time trying to answer it.

The point cloud impressed all participants, yet only Participant 4 demonstrated a (very thorough) understanding of what it represented, or even what its use was. The other participants had no idea how these points were generated. Participant 1 and 5 stated that the points represented the cathedral. I thought this was too weak an explanation, and did not accept it as correct. Also, most participants did not see any potential use of the point cloud.

After completing the task, participants were asked to fill out a post-questionnaire (see Table 4). Most participants claimed to have enjoyed the task. They were also fairly confident that they knew how Photosynth works. Unfortunately Participant 2 left the questionnaire incomplete. The other four participants indicated that they would most likely use Photosynth again, and that they would recommend it to friends or family who were tech-savvy. After the first two participants, I decided to amend the questionnaire, forcing participant to think about family or friends who are not comfortable with technology.

**Table 4:** Results of Post-Task1 Questionnaire

<b>ID</b>	<b>Do you know how Photosynth works? 1 = Not at all, 5 = definitely</b>	<b>Was this exploration fun? 1 = not at all, 5 = definitely</b>	<b>Would you ever use Photosynth? 1 = never, 5 = definitely</b>	<b>Would you recommend Photosynth to friends?*</b>	<b>Could you control your movements? 1 = no control, 5 = total control</b>	<b>Is Photosynth user friendly? 1 = not at all, 5 = totally</b>	<b>Will non tech friends find Photosynth fun? 1 = not at all, 5 = yes definitely</b>
1	4	5	5	Yes (4)	3	Not Asked	Not Asked
2	4	5	No Data	No Data	No Data	Not Asked	Not Asked
3	3	5	4	Yes (5)	5	4	4
4	5	4	4	Yes (4)	4	4	2
5	4	3	5	Yes (4)	2	4	2

\* The number in brackets is an indication of friend's confidence with technology (1 = extremely unconfident, 5 = extremely confident)

## **Task 2's Results**

In Task 2, participants were asked to draw a floor plan of a house. For Participants 1 and 2, this task was slightly different: they were asked to list all the furniture they saw in each room and, and only if time permitted, to draw the floorplan. However, I encouraged both these participants to focus more on the floorplan than on the list of furniture, as it was evident that we had enough time.

None of the participants drew a perfect floor plan. Participant 1 drew the best floorplan, but he too missed a door and a fireplace. Also, he admitted having experience drawing floorplans, as he is a painter by profession. Participant 3, who drew the worst floorplan, actually started out drawing a 3D floorplan. I suggested to her that she draw a top-view floor plan: as if she were looking down at the layout from the sky.

In a floorplan, I was not looking for neatness, but rather for accuracy in terms of doors, windows and relative sizes of rooms. Participant 3 drew the most inaccurate floor plan, even though she was one of the three participants who discovered all the rooms of the house. Her floor plan had a Study that wrapped around one part of the house, a kitchen that was oriented incorrectly and that was bigger than it should have been, a piano room that was angular when it should have been rectangular and a missing a corridor outside the (under construction) washroom. She was able to navigate around the house fairly comfortably, but was not able to orient herself correctly in this 3D space.

Participant 2 gave up after half an hour of trying. She had discovered and drawn all parts of the house except the kitchen. She even pointed to the gap on her floor plan

where the kitchen should have been and said that she couldn't find a route into that space of the house.

**Table 5: Results of Post-Task2 Questionnaire**

<b>ID</b>	<b>Did this task change your understanding of how Photosynth works</b>	<b>Was this exploration fun?</b> 1 = not at all, 5 = definitely	<b>Is Photosynth user friendly?</b> 1 = not at all, 5 = totally	<b>Listed every room?</b> 1 = surely not 5 = total confidence	<b>Easy to navigate?</b>	<b>In reality, how many rooms were discovered (out of 6)?</b>	<b>Moderator's rating of participant's floorplan</b> 1 = terrible, 5 = perfect
1	Yes	5	Not Asked	5	Yes	6	4
2	No	5	4	1*	No	5	2
3	Yes	5	4	4	Yes	6	1
4	Yes	4	4	5	No	6	4
5	No	3	3	1*	No	5	2

\* The partici

part chose not to complete the task

Participant 5 drew a better floor plan than Participant 2, but chose to stop the task after 50 minutes of trying. I scored his floor plan the same as Participant 2's floor plan because of the time it took him (Participant 5) to come up with a drawing. It would not have been fair to consider his floor plan "better" than Participant 2's floor plan because he used much more time to come up with his version, and he ultimately also gave up. He was missing the corridor behind the kitchen, the washroom that is under construction, the door that leads into the corridor from the yellow room, the fireplace in the study and the doorway between the kitchen and the living room.

Participant 4, who did the best in Task 1 and the most knowledgeable of the 5 participants coming into Task 2, had problems visualizing the space, but by the end of the task he had realized and fixed his mistakes. For instance, he drew the kitchen before he discovered the yellow room, which shares a wall with the kitchen. Upon discovering the yellow room, he was confused where it fit on the map because the kitchen left no space for this new room. He did pay attention to the widths of corridors and sizes of rooms. He faced problems moving around the house, especially at the passageway between the living room and the kitchen. Overall, he drew a good floorplan, though with slightly less ease than Participant 1 did.

### **Task 3's Results**

Task 3 was presented to Participants 3, 4 and 5 (see Table 6). In this task, participants were asked to explore a synth of their choice, or browse through my list of bookmarked synths, and provide subjective feedback. Prior to this task, I debriefed

participants on Photosynth’s conceptual model and explained how the halo, white quadrilaterals and the point cloud worked. This task’s aim was to see whether participants could navigate more easily around a synth of their choice, powered with by knowledge of Photosynth’s conceptual model. It was also a test of how much participants enjoyed the predefined tasks.

**Table 6: Results of Post-Task3 Questionnaire**

<b>ID</b>	<b>Did this task change your understanding of how Photosynth works?</b>	<b>Which was your favourite synth today?</b>	<b>Which synth taught you the most?</b>	<b>How would you divide a \$100 reward?</b>
3	Yes	Globe (one of my bookmarks)	House Interior	Idea: \$25 Designers: \$25 Photographer: \$50
4	Yes	German town (a synth he searched for)	House Interior	Idea: \$20 Designers + Programmers: \$80 Photographer: \$0
5	N/A *	Notre Dame (from Task 1)	House Interior	Idea: \$30 Designers: \$40 Photographer: \$30

\*Participant 5 did not pick a synth of his own due to fatigue.

Participant 3 did not seem enthusiastic about searching for a synth. Instead she relied on my recommendations. She viewed 4 synths from a list of 21 that I had bookmarked prior to the study. She had wanted to view an additional, 5<sup>th</sup>, synth but due to a technical reason it would not load. She enjoyed exploring the various synths,

especially a synth of the exterior of a house and a synth of a globe. She said that she has a special interest in those subjects. She claimed that this task helped her understand how Photosynth works because “I was allowed to just play”. She claimed the synth of the globe was her favourite, though the synth used in Task 2, the interior of a house, was the best synth for learning Photosynth.

In fact all three participants claimed that the home interior was the synth that taught them the most about how Photosynth works, even though none of them picked that synth as their favourite.

Participant 5 did not browse for a synth because he expressed exasperation after using the system for the first two tasks. He was clearly tired after working on Task 2, and did not want to deal with the frustrations of using the system in a 3<sup>rd</sup> task. He picked the Notre Dame synth as his favourite because of the fine details of the cathedral’s façade that he could zoom into.

Participant 4 picked his favourite synth as one that he found himself. It was of a German town that he had visited.

The last question asked on Task 3’s post questionnaire was who participants would reward for their experience with Photosynth, and to what level each person would get a reward. The results can be seen in Table 7. I shall discuss this in the next chapter: Recommendations.



## CHAPTER 5: RECOMMENDATIONS

In this chapter, I shall provide recommendations that I believe will improve Photosynth's usability. I shall provide the most urgently needed changes first. These recommendations are based on results from Task 1 (see Table 3).




**Figure 8:** A halo appears at the centre of the cathedral, when viewed from a distance

### Cannot Find a Halo

Question 3 in Task 1 resulted in the biggest and most frequent complaint by participants. They were not being able to find a halo when they knew that a halo existed for a particular object. In Task 1, for example, participants knew that a halo that would

allow them to rotate around the church, existed, as can be seen in Figure 8<sup>21</sup>. However, if they viewed the church from a very short distance, no halo would appear (see Figure 9).



**Figure 9:** No halo appears at the centre of the cathedral, when viewed from close up, even though it did when viewed from a distance.

This behaviour was unexpected. Users said that they expected to see a halo in Figure 9. If, at this stage, they wanted to use the halo, they had to zoom out to an acceptable distance, such as in Figure 8, before a halo would appear.

There are a couple of usability issues to discuss here. Firstly, the inconsistency of showing a halo for an object in one view, and then not showing it for another, hinders users' ability to understand the interaction technique because the control is not present at a time when it is clearly expected to be. Secondly, the distance one must be from the

---

<sup>21</sup> <http://tinyurl.com/ceaqlv> - see appendix for URL expansion

object before the halo appears, seems to be random. There is no signal given to the user that indicates how close or far they are from a halo. The user is literally left to move randomly until they encounter a halo. This is terrible treatment of the user's patience. I admit that participants were trying to use the halo a lot in part due to my asking them to do so, but the halo is a common enough control in the interface that even in the real world, participants would be required to interact a lot with the halo. In my study, this was the earliest, biggest and most frequently vented frustration. Question 3 in Task 1 asked participants when it was that they expected to see a halo (see Table 3). Aside from Participant 4, all participants either got this question wrong, or did not have an answer for it, and aside from Participant 4, they all rated it as their least confident response.

Participant 1 stated that "It feels like the halo is the ultimate key, and that should be present throughout the experience." (9:47). Participant 2 says, early in her exploration of Task 1: "I have no idea ... of how do I get it to show the halo." (10:53). She clicks on the *play* button on the top right of the window, to see if that shows her a halo. At the end of Task 2, she claims that she did not learn anything new about the halo because "I still don't know how to make it appear" (1:19:40). This is evidence that even users, not just the usability professional, do not appreciate the lack of visibility in an interface!

Participant 3 struggled a lot with trying to make a halo appear. She also encountered some glitches in the system, where photos would load as an opaque grey rectangle. This caused her to feel even more out of control, stating that technology always goes wrong when it's in her hands. She stated that the halo was shown to users when the system detects that the user is bored. It is a way of prompting you to click on something when it detects that you are not clicking on anything anymore. However she did not write this as

her answer for the question, instead choosing to skip this question. Participant 4 also struggled with this question: "I don't see any halo so I'm going to just go back and try and figure that out. I see a halo [he gets to a photo where the halo appears]; it's sitting there ... [now I ]zip around the side here, there should be a halo there [he sees a halo]. Zip around to the other side, is there a halo? No [he doesn't see a halo]. So from the front perspective there is a halo, and from the right perspective there is halo, which doesn't make sense to me because from the left perspective there's no halo. So that seems kinda incongruous; I don't really understand that." (6:10). Similarly, Participant 5 struggled trying to figure out how a halo is made to appear. He ultimately stated that "the further back [you are] from the object, the more you can use the halo to view different angles" [Participant 5's Answer sheet for Task 1]. This is incorrect because you might have only one picture of the object from a distance, but multiple images of it from close up.

One fix for this annoyance would be to have a static command on the interface that drives the user to the previously viewed halo. Often, participants used the halo to arrive at a particular photo, clicked on a white quadrilateral which led to the halo disappearing, and then simply wanted to get back to the halo to zip to a new angle. A simple and frequent sequence of tasks such as this caused a lot of grief to the participants. Having the system bring them back to the halo they were on would help.

Another fix would be to have a static command on the interface that takes a user to the closest halo. Often users knew that they had seen a halo somewhere in their vicinity, but they struggled to find a photo at the right distance from the object, to bring up the halo. A button that took them to that photo would help greatly.

This usability bug needs to be addressed with the highest priority because it was the most frequently cited complaint.

### **Users found unexpected zooms disorienting**

A confusing encounter for participants using the halo was the fact that different photos in a halo would be at different distances from the object. For example, in Task 2, when participants were trying to rotate around the table at the entrance of the house, some views would take them outside the house (See Figures 10 and 11)<sup>22</sup>. Participants did not notice that the halo, in this case, was actually centred on the front door frame, not on the table.



**Figure 10:** Halo above the table in the entrance lobby of the house – about to click on green region, expecting to view the table from a 180 degree direction change.

<sup>22</sup> <http://tinyurl.com/bcjlro> - see appendix for URL expansion





**Figure 11:** After clicking on the green region highlighted in Figure 11, this is the view presented to users; outside the house and no table in sight. The open door is the front door to the house.

Participants were confused as to why they were moved outside the house. They had not expected such behaviour, when in fact the behaviour should have been expected because you were, correctly, presented with a 180 degree view shift around the door frame.


Another instance of excessive zooming that confused participants was when the halo would show participants a view of the object from a different room (see Figures 12 and 13)<sup>23</sup>. In Figure 12, participants expected to move slightly to the right. They were left confused when they were suddenly taken into the lobby and shown the fireplace from a distance.

---

<sup>23</sup> <http://tinyurl.com/bcjlro> - see appendix for URL expansion



**Figure 12:** This halo confused participants because clicking on the green region resulted in seeing a partial view of the fireplace, from the lobby (see Figure 13)



**Figure 13:** Suddenly, participants were moved into the lobby and shown half the fireplace

## **Users unaware they could drag a halo**

I noticed that only 2 participants learnt to drag the halo in order to navigate around an object: Participants 3 and 4. All other participants would click on green regions of the halo, in order to navigate around an object. However, dragging the halo is much faster and gives the user more control. In order to drag the halo, a user must hold the left mouse button down, on the halo, and move the mouse.

Participant 3 discovered this method of navigation in Task 2, when she wanted to verify that a particular door was behind her (58:40). She expressed strong approval of this feature and noted it twice in the post questionnaire for Task 2 as a very desirable feature.

Participant 4 discovered this feature during Task 1, while trying to figure out the function of the halo (7:15). As he discovered this feature early in his exploration, he used it throughout the session.

The fact that the weakest participant – Participant 3 – expressed great value in this feature (unfortunately towards the end of the session), and that the best performing participant used this feature extensively, and that no other participant discovered this feature, is indicative that this feature needs to be made more apparent to users. Perhaps adding arrow heads that wrap around a halo, would indicate to users that they would be able to drag it.

In the pure point cloud view, all participants enjoyed the freedom they had to view the scene from any angle. Participant 1 asked if he was able to move this freely in the photo mode too, indicating his desire to do so. Participant 2, at the end of the session, also indicated that she liked being able to move around easily in the point cloud mode.



While participants need to be made aware that they will never experience the freedom of movement that they have in the pure point cloud mode, they should be informed of the ability to drag themselves around a synth when photos are being displayed.

### **Users' inefficient navigation techniques**

As mentioned earlier, only 2 participants dragged the halo to navigate:

Participants 3 and 4. The other participants would use the arrow buttons on the borders of the window, and the white quadrilaterals to shift positions. I observed a couple of usability problems encountered by participants during these, frequent, tasks.

The tooltips that appear when you hover over an arrow on the border of the window, appear too slowly. Most times the user has already clicked on an arrow before the tooltip has appeared. This is especially unfortunate for the two downward arrows on the bottom border of the window: one which tilts the view downward, and the other which zooms out. It is not obvious to participants, which of these buttons performs what function, and most simply click the button to see what happens, not even aware that a tooltip would have told them. Similarly, when wanting to rotate to the left or right, participants would click the left or right arrow, not knowing that those buttons take you one step to the left or right, and not rotate you around an object, which is what the halo does. If they had a chance to read the tooltip, they might have realized their misunderstanding, and not have been left confused. This is a simple fix that would reduce participants' anxiety while trying to navigate.

Figure 14<sup>24</sup> is a scene from Task 2. If participants wanted to zoom out this scene, they would move their mouse over the left part of the screen until they found a white quadrilateral that suggested a broader perspective.



**Figure 14:** Translucent layers below the image that is in focus, where the mouse is located. The lower borders of the layers have been outlined in red to highlight them.

However, an easier way of finding the appropriate white quadrilateral would have been by looking at the translucent images around the photo in focus (see Figures 2 and 14). These images have subtle borders that become prominent with a white quadrilateral around it when the mouse is hovered over them. Following the lead of these translucent borders, participants could have found the desired white quadrilateral much faster than randomly searching for one.

---

<sup>24</sup> <http://tinyurl.com/bcjlro> - see appendix for URL expansion

Participant 1 noticed this subtle hint, and appreciated its value. In his post questionnaire for Task 2, he noted that this realization made his navigation around the space much easier. Interestingly, he referred to them as “translucent sheets”, not as surrounding photographs. He also claimed that they kept a history of where he had been in the past. This indicates that he did not realize their conceptual model, but was still able to garner productivity from their use.

Since participants usually used the arrows along the border of the screen and the white quadrilaterals for navigation, when a halo was not present, and they did not use the hints provided by the translucent images, I suggest that hovering the mouse over an arrow, generate a white quadrilateral around the image that would be brought into focus next. This would suggest two clues to the user: what the consequence of clicking on a particular arrow would be, and, perhaps over time, clue the user into the fact that the translucent images are helpful for navigation, because they will see them highlighted whenever they use an arrow.

Participant 5 suggested a very interesting proposal to aid navigation. At the end of the session, after I explained to him how Photosynth works, he suggested that since the point cloud of a space is always known by the system, there should be a small point cloud map inset into the synth, with an indicator that shows your current location on the map. He said that this would have helped him navigate around the house much more easily, since his orientation would have been set by this map. In essence, it would have made my task of drawing the floorplan irrelevant because the map would be a floor plan of the space, but this would, indeed, have helped participants orient themselves better within the space. It would also have told participants their current position if they were suddenly

zoomed out of a scene unexpectedly. They would have been able to recover from the disorientation much faster than they currently did.

### **Problems understanding what the solid white region of halo is for**

Question 7 in Task 1, asking what the solid white region of the halo represented, revealed more confusion about the conceptual model. Participants 1 and 2 got it wrong, while Participant 3 could come up with no answer. Participants 4 and 5 got the right answer. Participant 1 claimed it was an indication of a “zooming limit”, i.e. the photo you are viewing now is the most zoomed-in photo of the church from this angle. His answer was close, but because of his misunderstandings during question 6, his understanding of this feature was incorrect too. He figured out that the solid white region represented the currently viewed angle, but the key information that he thought it conveyed was the distance from the church. Participant 2 claimed that the white region allowed her to zoom out of a church, which is terribly wrong because clicking on the solid white region results in no visible change, let alone zoom you out of a scene. Participant 3 clicked on the solid white region multiple times, and was carefully looking out for any change in the view, that would indicate what the solid white region controls. Currently, you are able to drag the halo by keeping the left mouse button pressed on the solid white region (or any accessible region of the halo, for that matter). When Participant 3 would click on the solid white region, inadvertently she would move the halo a little bit, and would perceive this motion if she missed something: so she would repeat the action to catch what happened. Perhaps a fix for this issue would be to add diagonal red stripes to the solid white region, and to not allow users to be able to drag the

halo by clicking on this region. A red strike through a white background is a cultural symbol for inaccessibility, commonly seen on road signs.

Since participants thought that clicking on this region would result in a change to their view, when it actually does not, perhaps the solid white region should lose its ability to drag the halo; i.e. if it truly acted *dead*, people would interpret that it has no function and that it is inaccessible. This might not indicate its function, but it would prevent users from believing it represented something completely unrelated. If Participant 3 had figured out that this region was inaccessible, she would not have tried clicking on it multiple times. And if this region lost its ability to drag the halo, she would not have seen slight movement whenever she clicked on it. Also, Participants 1 and 2 might not have thought that the solid white region performed a change in view.

Another possible design fix for the solid white region would be to have a tooltip appear when the user hovers their mouse over the region, with text such as “Your current vantage point”. This would give users a quick and clear understanding of the function this feature entails.

### **Lack of understanding and appreciation for the point cloud**

No participant was able to figure out the point cloud’s function, or what it represented (Question 8 in Task 1), without some help on my part. During the session with Participant 1, I realized that if I did not intervene, the idea of the point cloud would be completely lost to him, and I would not gather any valuable usability data for this feature. Thus, I decided to show the scene to him in pure point cloud mode, specifically

for him to get an idea of what the point cloud feature is all about, and to gauge his responses to it. I noticed a similar necessity for the other four participants.

Participants did not pay attention to the point cloud until they were trying to answer the question related to the point cloud. And at this time, the response to the point cloud was a mixture of confusion and apathy. Once I toggled the view to pure point cloud mode, participants expressed positive affect for this feature. They found it visually appealing and they also liked the ability to drag the halo to *any* direction they wished. However, even after viewing and manipulating the scene in pure point cloud mode, participants were not able to tell me what the dots represented or what use they would have in the synth. In essence, the use of the point cloud is that it refines and completes a scene when there are no images for the background of the scene from the current angle (See Figures 15 and 16)<sup>25</sup>. Participant 2 insisted that the point cloud served no purpose. She said that she doesn't "know what I would use this for, other than its kinda fun" (29:35). She admitted that the point cloud view was pretty, but she goes on to say that in the "view with pictures and the dots ... the dots are neither here nor there other than they just serve as a background so that I don't just see black". (30:39). Participant 3 also expressed that she did not see any real value for the point cloud feature: "It is interesting and everything but ... I don't see any use for it." (28:28). She later explained that the only use she could imagine it would have is for photographers, who could use it to "build the image better" (28:50).

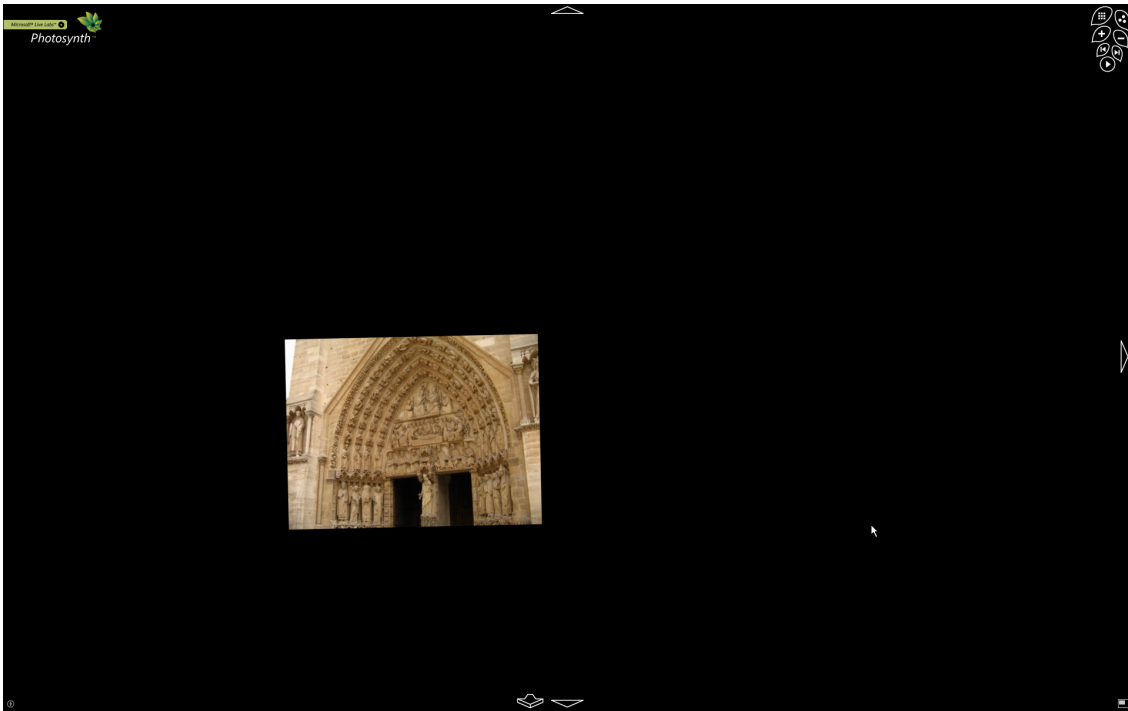
The two youngest participants, 1 and 4, aged 31 and 20 respectively, were the two most excited to view a scene in pure point cloud form. While navigating around the

---


<sup>25</sup> <http://tinyurl.com/ceaqlv> - see appendix for URL expansion

scene in point cloud form, Participant 1 stated: “Oh, that’s neat! That’s much more interesting [than the point cloud was when it was interspersed among photo] ... Isn’t that cool! Neat ... that’s beautiful ... you need to see that applied to a person ... This is a fascinating feature. This is really neat ... Look at that! It is quite extreme. Wow, I love it. Very fun.” (20:10). Participant 4 stated: “Whoa! That’s crazy man. Oh, that’s wicked! ... Oh, that’s crazy. I’ve never seen anything like that except for when I used AutoCAD. That is crazy man! That’s actually pretty wicked, eh? This looks like something game designers would use. I’m just enjoying looking around.” (14:45). None of the other participants expressed such strong positive responses to the point cloud view. In fact, even though Participant 1 found this feature fascinating, he could not figure out what its use was. He noted that it was trying to “represent the feature of the church [and] trying to aid my mind by adding a three dimensionality to the image, but it is not working” (14:20).

Looking at comments posted on Photosynth’s website, it looks like users enjoy viewing a synth in pure point cloud form. Comments posted under synths, inform the reader to view the synth in point cloud form for the thrilling effect. Users also take pride in how complete their point cloud is, because that is an indication of how well they photographed a scene. My study hinted that the younger participants were more excited by, and had more respect for, the point cloud than older participants.



**Figure 15:** Example of a view without the point cloud



**Figure 16:** Example of a view with the point cloud



## Users misunderstand the white quadrilaterals

**Table 7: Keywords used to describe what the white quadrilaterals represent**

ID	My observations of how participants answered Question 1 of Task 1
1	He says the white quads represent points where you can zoom in, which doesn't suggest that he thinks of them as photographs. He restates "zoom points" at the end of the task, even after he knows he is looking at photos. He says it was a tricky question to answer. He didn't get it "right away".
2	She does not mention that each white quadrilateral is a different photo, though she knows that it shows you a different perspective.
3	Does not mention that each quadrilateral is a different photo. She says that the quadrilateral represents an area that you can expand and view, but nothing about it being a photo. She called them spatial areas.
4	Does not mention that each quad was a photo. But he knew he was looking at a bunch of photos. Said the quads represent plains of perspective.
5	He almost never uses the word <i>photo</i> , throughout the experiment. Instead he refers to them as <i>views</i> .

In Task 1, I marked each participant's response to Question 1 as being wrong. As can be seen in Table 7, everyone got the idea right, but no one explicitly stated that the white quadrilaterals represented the available photographs of the scene. On the one hand, this was the first question that all participants answered. Most might not have been aware, at this point, that they were looking at photographs. That might explain why they did not describe the white quadrilaterals as available photographs. On the other hand, before ending the task, I would ask participants if they were sure of their answers, prompting them to go through them again. No one amended their answer to Question 1. Perhaps I am being a stickler, but a zooming area, a spatial area, a plain of perspective, a view ... these need not be photos. Participants might know that they are looking at

photos, but did they know that each white quadrilateral represented a particular photo? I suspect that participants did not easily realize that they were limited to view the scene from fixed locations: the set of photographs. This would explain why they never mentioned that a white quadrilateral is an *available photograph*. They had a preconceived notion that since multiple photographs had been stitched together to form a 3D view, all regions of the 3D space would be accessible. They understood that they were looking at photographs, but they did not realize that the system is not smart enough to fill in the gaps that exist in any photo set.

### **Users have no problem understanding basics of conceptual model**

All participants figured out that they were looking at a collection of photographs that had been stitched together. What's more, they figured this out early during the first task. In Task1, Questions 4 and 9 tested this.

Participant 1 suspected he was looking at photographs, literally 2 seconds after I showed him Photosynth for the first time. He said: "Oh interesting. Its kinda like Cover Flow ... [which] is a Mac feature where it takes photos of album art or photographs [for you to browse through] (3:04)"

Later, when answering Question 9, which asked why people seen at one point of the synth suddenly disappeared, he said aloud, "well that's kinda obvious" (17:45). He explained that this is a "photo tour" (18:08). I asked him to elaborate further, and he replied:

"Well the photographs are taken at different times of day."

"So, are these different photographs?", I asked.

“Yeah, Oh yeah ... See this one, you can tell by the light.”

“So these are all different photographs.”, I restated.

“These are all different photographs.”, he confirmed.

Participants 2 and 4 revealed that they were looking at a collection of photographs, when answering question 4, which asked them whether they were able to view the church from 360 degrees, and why. This question prompted them to reveal, for the first time, that they were looking at photographs. As all participants attempted the questions in the order that they were presented, this shows that participants had already formed a correct conceptual model, early on in the task. I cannot point to the *exact* moment when they came to this conclusion, but question 4 prompted them to voice it. Participant 2 stated that she couldn't go to the back of the church because “there aren't enough pictures; the halo won't let me go that way”. I asked her who was responsible for this lack of pictures. “It would be whoever took the pictures” (14:22). I wanted to confirm that she thought she was looking at a set of pictures. I was careful not to use the word ‘photograph’ because sometimes people think of pictures and photographs differently. I asked her if she was looking at an amalgamation of different pictures. She responded by saying: “Yes, and it was taken by one photographer” (14:33). “And if the photographer was commissioned differently, my guess [is] that there would be a wider variety of pictures to choose from, including some inside and some on the other side of the building on the outside.” (14:46). This showed that by pictures, she did mean photographs.

Participant 4 also thought that he was looking at different photographs. He stated that the 360 degree view was made with a video camera or from a regular camera with

panorama settings. He said that one could load these images into the system to make up the 360 degree view. He says that the video or digital camera “is being moved or angled by somebody” (9:34).

Participant 3 said that she was looking at different photographs, before even reaching question 4. While she was trying to answer question 3, struggling to recover a halo that she lost, I asked her what she thought she was looking at currently. “What am I looking at? I’m looking at a photo, *obviously*.” (9:10), she responded. She then goes on to say that the photo is “almost alive” (9:21). Further, she says that “it is more than a photograph ... it’s almost like an architect would do it, you know ... but obviously this is a *photo*, not an architect [’s drawing].” (9:52). She was vociferous in her amazement with Photosynth’s technology, being able to view the church from different angles by clicking on different parts of the halo.

Participant 5 took the longest to figure out that he was looking at a collection of photographs. Due to missing audio data, I cannot confirm when it was that he commented on his discovery, but I remember thinking that he was taking longer than the others to mention that he was looking at a bunch of photographs. When I quizzed him on it, he said that it might be photographs that he was looking at, but he wasn’t sure. And he also thought that the photos were all clicked at the exact same moment, which must have been very hard to setup, he thought.

Throughout the session, Participant 5 would use the term ‘view’ to refer to the photographs. I suspect that until the end of the session, he was not convinced that he was looking at photographs, but that photos was the best explanation that he could come up with. I suspect this is why he refrained from using the term *photo* very much, throughout

the session; because he might be wrong. Other participants were sure they were looking at different photographs, and used the term *photo* freely, once they were convinced.

### **Can we trust a user's responses?**

Table 3 provides evidence of the confidence each participant had in their answers for Task 1. Looking at Participant 5's confidence ratings, you notice that he ranked each question's confidence at level 3. This, indeed, sounds suspicious. I noticed that Participant 5 rushed through this section of the post-questionnaire, which asked him to rate his confidence in his answers for the 9 questions I asked. He simply ticked level 3 for all 9 questions, without even looking back to see what each question was. I cannot say with absolute certainty that we should discount his confidence ratings, but the potential to do so is high.

### **Users should be allowed to pick their own tasks**

Participants 3, 4 and 5 were presented with Task 3, where they were allowed to pick a synth of their choice, to explore. They were then asked to fill out a post-task questionnaire that covered subjective responses to the test session. Participant 5, due to fatigue, chose not to explore any synths of his choice, and moved straight to the questionnaire.

I found it interesting that Participants 3 and 4 both named their favourite synth of the session as synths that they picked in Task 3. While this might indicate that participants did not enjoy the predefined tasks, and thus did not give it their best, that all three participants picked Task 2's synth of the home interior as the most informative of Photosynth's function, was reassuring.

Additionally, Participants 3 and 4's favourite synths were less synthy than either of Task 1 or 2's synths. They were also composed of fewer photographs. This did not seem to hinder them for being picked as the favourite. This is reassuring for the Photosynth team, because it shows a genuine interest in the product. Users overlook the fact that some synths are less synthy and sparser than others, a common issue with the thousands of synths that have been uploaded. If a participant is able to find a synth of a subject that they are interested in, they seem to enjoy it in spite of its lower quality.

The last question on the post-task questionnaire for Task 3 asked participants who they would reward for their favourite synth, and who would get what percentage of the reward (see Table 6). It did not surprise me that Participant 3 gave half the reward to the user who took the photos, as she had a lot of fun viewing the synths. Her lack of understanding of the point cloud, navigation techniques and the halo must be the reason why the designers did not win the biggest share of the bounty. Participant 4, on the other hand, gave 80% of the reward to the designers and programmers. Among the 5 participants, he understood the technical aspects of Photosynth the best, and chose to show his appreciation to the designers and programmers. He commented that the user who took the photographs deserved no money because they chose to upload a synth: they were not forced to do so. Hence they do not deserve a reward. Participant 5 surprised me the most, because in spite of his apparent fatigue, he rewarded the designers the most. I had expected him to reward the photographer the most, because he often commented on how beautiful the Notre Dame photos were. And, on top of that, he took the longest to complete Task 1 and Task 2, which I thought he would punish the designers for; but apparently not. This might be a good sign for the Photosynth team, because it indicates

that first time users enjoy using the software enough to overlook some of the complaints they have.

### **Problem with nonrandomized delivery of questions**

For Task 1, I presented all participants with the same order of the 9 questions. Perhaps this static order had an effect on responses. For Question 1, which asked participants what the white quadrilaterals represented, randomizing the questions would have given some participants more of an opportunity to interact with the system, before they answered the question. Perhaps then, they would have specified clearly that the white quadrilaterals represented individual photographs.

### **Users were confident in their understanding of incorrect concepts**

Barring participant 5's questionable confidence ratings and each participant's response to Question 1, I noticed that participants were quite confident of answers that were wrong. This can be interpreted as people being misled very convincingly. For example, Participant 1's confidence for answers that he got wrong was 3, 5, 5 and 5 for questions 3, 6, 7 and 8 respectively (see Table 3). Participant 3 was also quite confident of her wrong answers. Let us look at Question 6, which asked participants what the size of the green region of the halo represented. Participants 1, 2 and 3 got the answer wrong, but rated their confidence as 5, 3 and 4 respectively. Participant 5, on the other hand, admitting not having any idea what the answer should be. Participant 4 got the answer correct. Participants 1 and 2 were convinced that the size represented how close to the church you would get if you clicked on that region of the halo. They both said that the larger the green region is, the closer you would end up to the church. They came to this

conclusion because they clicked on a few green regions and, *coincidentally*, experienced this behaviour. Norman, in his book *The Design of Everyday Things*, mentions that people are quick to apply a causal relationship between their actions and the results of those actions [Norman 1990, pp 40]. Participants 1 and 2 happened to pick regions that demonstrated this behaviour, and thus came to their conclusions. If Participants 1 and 2 had been forced to click on more green regions of the halo, they might have noticed that not every large green region zoomed them closer to the church. But, the reality is that users have an upper limit to their patience. Once they reach that limit, if they have not completed their task, they simply give up.

### **Summary of Recommendations**

- Enable users to navigate to a halo of their choice
- Lock the distance from the object, when navigating via a halo
- Indicate to users that they can drag a halo, to rotate around the scene
- Indicate to users what the solid white region of the halo represents
- Indicate navigation via the surrounding photos
- Display tooltips quicker
- Don't count on the point cloud adding value to the user's experience
- Indicate to users that each white quadrilateral is a photo in the image set
- Randomize the questions that quiz the conceptual model understanding
- Look out for confident *wrong* answers when quizzing conceptual model



## CHAPTER 6: CONCLUSION

In this study, I set out to test the usability of Photosynth, a system with an unprecedented and complex conceptual model. I decided that a test of its usability entails a test of how well its conceptual model is understood. To test how well the conceptual model was understood, I wanted to see if participants were a) able to figure out the fact that they were viewing multiple photos and that they understood the limitation of viewing the space from only a discrete set of points, b) able to *predict* the action of controls within the interface, c) able to find, on the interface, controls that they were looking for, and d) aware that the action they performed was registered by the system.

Based on the above 4 points, I conclude that 4 of the 5 participants were not able to grasp the conceptual model – their user models were highly disjoint from the design model; a failure of the system image. The interface was not able to convey the design model to the user.

Although 5 participants is a small number to run the system against, it is *not* an inaccurate sample of a non-technical subset of users; according to Jakob Nielsen, 5 users are often enough<sup>26</sup>.

The subjective responses given by users tell a different story than what my results indicate. All participants said they would recommend Photosynth to friends who are comfortable using technology. In fact, the weakest participant, Participant 3, claimed she would even recommend Photosynth to friends who are not comfortable handling

---

<sup>26</sup> <http://www.useit.com/alertbox/20000319.html> ; 26 March 2009

technology. Participants 3, 4 and 5 all consider Photosynth to be “user friendly” (see Tables 4 and 5) – a term I picked on purpose, for its pop-culture reference. All participants found Task 1 and Task 2 fun. Task 3 was also enjoyed by the two participants who performed it. However, I suspect a lot of this is due to the novelty of using such a unique system. I suspect that, over time, the novelty will wear off, and users’ enthusiasm will die down.

But the team at Photosynth shouldn’t despair. With just a few design changes to the interface, I believe a majority of usability problems will be fixed. Making the halo easier to find, indicating to users that they are able to drag a halo, allowing a user to lock their distance from an object so that they do not get disoriented by unexpected zoom-ins and zoom-outs and representing the currently viewed angle in an updated fashion will fix the common usability problems I noticed. I do not recommend making any changes to the point cloud representation because it did not bother participants: most ignored it as it didn’t get in their way, and the others found it very exciting. It was also encouraging that participants were able to predict accurately and quickly what they expected the halo’s function to be. This means that the halo’s design is accurately mapped to its function.

Snaveley, Garg, Seitz, and Szeliski’s 2008 work upgrading the technology and interface of their *Photosynthesque* project, seems very promising. I suspect that Microsoft will incorporate their changes into Photosynth, just as they did when the product was initially acquired by them. They claim to have addressed the hardships in navigation that participants of this study expressed. For instance, they claim to have stabilized rotation around an object so that users are not zoomed in and out of a scene unexpectedly. They also present users with a list of scenes where they may rotate around

an object; in essence, a list of available halos. A video demonstration of these changes can be viewed on Snavely's homepage, who is now an assistant professor at Cornell University<sup>27</sup>. An interesting follow-up study to my own would be to run my test procedure on the new system, asking the same questions I asked my participants for Task 1, in order to check whether the new design – the new system image – conveys the design model to the user in a better fashion.

According to the Photosynth LiveLab website, the Microsoft Photosynth team has now joined the Microsoft MSN team. They claim that

“Photosynth will begin to become a key part of the experience for MSN's 550 million monthly visitors worldwide. Synths will be prominently featured on MSN.com. To create a more absorbing experience for its visitors, MSN will use synths of popular destinations and notable events in many of the places where static images are used ... today.”<sup>28</sup>

Given Microsoft's plans for Photosynth's future, I believe that more studies of its usability are imperative.

Through this project, I gained valuable experience designing, conducting, analysing and reporting a usability test for a system whose conceptual model is complex and relatively unknown. Occasions to test such systems are infrequent, and I am glad to have capitalized on this, rather rare, opportunity. More such experiences need to be shared by the usability community, to prepare us for the novel and complex interfaces that we are bound to evaluate in our career as usability professionals.

---

<sup>27</sup> <http://www.cs.cornell.edu/~snavely/> ; 26 March 2009

<sup>28</sup> <http://livelabs.com/photosynth/faq/> ; 26 March 2009

## REFERENCE LIST

- Bruno, V., Tam, A., and Thom, J. 2005. Characteristics of web applications that affect usability: a review. In Proceedings of the 17th Australia Conference on Computer-Human interaction: Citizens online: Considerations For Today and the Future (Canberra, Australia, November 21 - 25, 2005). OZCHI, vol. 122. Computer-Human Interaction Special Interest Group (CHISIG) of Australia, Narrabundah, Australia, 1-4.
- Chauhan, V. 2006. Usability professionals: you've come a long way, baby! interactions 13, 2 (Mar. 2006).
- Dicks, R. S. 2002. Mis-usability: on the uses and misuses of usability testing. In Proceedings of the 20th Annual international Conference on Computer Documentation (Toronto, Ontario, Canada, October 20 - 23, 2002). SIGDOC '02. ACM, New York, NY, 26-30.
- Greenberg, S. and Buxton, B. 2008. Usability evaluation considered harmful (some of the time). In Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems (Florence, Italy, April 05 - 10, 2008). CHI '08. ACM, New York, NY, 111-120.
- Hart, J., Ridley, C., Taher, F., Sas, C., and Dix, A. 2008. Exploring the Facebook experience: a new approach to usability. In Proceedings of the 5th Nordic Conference on Human-Computer interaction: Building Bridges (Lund, Sweden, October 20 - 22, 2008). NordiCHI '08, vol. 358. ACM, New York, NY, 471-474.
- Hartmann, J., De Angeli, A., and Sutcliffe, A. 2008. Framing the user experience: information biases on website quality judgement. In *Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy, April 05 - 10, 2008). CHI '08. ACM, New York, NY, 855-864.
- Hertzum, M., Jacobsen, N. E. The evaluator effect: a chilling fact about usability evaluation methods. *International Journal of Human-Computer Interaction*, 15 (1) (2003), 183-204.
- Howard, T. W. 2008. Unexpected complexity in a traditional usability study. *Journal of Usability Studies*, 3, 4 (August 2008), 189-205.
- Hsieh, C., Cheng, W., Chang, C., Chuang, Y., and Wu, J. 2008. Photo navigator. In *Proceeding of the 16th ACM international Conference on Multimedia* (Vancouver, British Columbia, Canada, October 26 - 31, 2008). MM '08. ACM, New York, NY, 419-428.

- Kellogg, W. A. and Breen, T. J. 1987. Evaluating user and system models: applying scaling techniques to problems in human-computer interaction. In Proceedings of the SIGCHI/GI Conference on Human Factors in Computing Systems and Graphics interface (Toronto, Ontario, Canada, April 05 - 09, 1987)
- Molich, R. and Dumas, J. S. 2008. Comparative usability evaluation (CUE-4). *Behav. Inf. Technol.* 27, 3 (May. 2008), 263-281.
- Nørgaard, M. and Hornbæk, K. 2006. What do usability evaluators do in practice?: an explorative study of think-aloud testing. In Proceedings of the 6th Conference on Designing interactive Systems (University Park, PA, USA, June 26 - 28, 2006). DIS '06. ACM, New York, NY, 209-218.
- Norman, D.A. 1990. The design of everyday things. Doubleday, New York, NY.
- Preece, J., Rogers, Y., and Sharp, H. 2002 Interaction design. John Wiley & Sons, Inc, New York, NY.
- Redish, G. 2007. Expanding usability testing to evaluate complex systems, *Journal Of Usability Studies* 2, 3, 102-111.
- Snaveley, N., Garg, R., Seitz, S. M., and Szeliski, R. 2008. Finding paths through the world's photos. In ACM SIGGRAPH 2008 (Los Angeles, California, August 11 - 15, 2008). ACM, New York, NY.
- Snaveley, N., Seitz, S. M., and Szeliski, R. 2006. Photo Tourism: exploring photo collections in 3D. In ACM SIGGRAPH 2006 (Boston, Massachusetts, July 30 - August 03, 2006). ACM, New York, NY.
- Spolsky, J., 2001. User interface design for programmers. Apress: New York, NY.
- Tang, J. C., Liu, S. B., Muller, M., Lin, J., and Drews, C. 2006. Unobtrusive but invasive: Using screen recording to collect field data on computer-mediated interaction. In Proceedings of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work (Banff, Alberta, Canada, November 04 - 08, 2006). CSCW '06. ACM, New York, NY.
- Tohidi, M., Buxton, W., Baecker, R., and Sellen, A. 2006. Getting the right design and the design right. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Montréal, Québec, Canada, April 22 - 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI '06. ACM, New York, NY, 1243-1252.
- Whitehead, C. C. 2006. Evaluating web page and web site usability. In Proceedings of the 44th Annual Southeast Regional Conference (Melbourne, Florida, March 10 - 12, 2006). ACM-SE 44. ACM, New York, NY, 788-789.
- Wilson, C. 2007. Taking usability practitioners to task. *interactions* 14, 1 (Jan. 2007), 48-49.

## **APPENDICES**

## Appendix A: TinyURL Referrals

To increase readability, I used TinyURLs<sup>29</sup> to shorten the long URLs that a synth entails. Here, you may map TinyURL's used in this report with the actual URL that they represent. All URLs were verified to be valid on March 26 2009.

**Table 8: Mapping of TinyURL to their corresponding, real, URL**

Chapter	Description	TinyURL	Actual URL
1	Jonathan's synth, titled San Marco and Beyond	awt67e	<a href="http://photosynth.net/view.aspx?cid=d7d78324-ad5e-41ea-94dd-512580ed4393&amp;i=0:0:6&amp;z=622.1688339456001&amp;g=0&amp;p=-4.69631e-014:-1.34356e-013&amp;m=false&amp;c=18.6503:10.3081:0.183456&amp;d=-0.895512:-1.42306:-1.60214">http://photosynth.net/view.aspx?cid=d7d78324-ad5e-41ea-94dd-512580ed4393&amp;i=0:0:6&amp;z=622.1688339456001&amp;g=0&amp;p=-4.69631e-014:-1.34356e-013&amp;m=false&amp;c=18.6503:10.3081:0.183456&amp;d=-0.895512:-1.42306:-1.60214</a>
3,4	Pierre391's synth, titled Notre Dame De Paris	ceaqlv	<a href="http://photosynth.net/view.aspx?cid=bb015e11-cfb5-4187-ac03-98147a258b34&amp;wa=wsignin1.0">http://photosynth.net/view.aspx?cid=bb015e11-cfb5-4187-ac03-98147a258b34&amp;wa=wsignin1.0</a>
3,4	Reppart's synth, titled Downstairs	bcjlro	<a href="http://photosynth.net/view.aspx?cid=436e315e-af30-4b45-a3c9-1f503ecc8afb&amp;i=0:9:1&amp;z=435.68801060008&amp;g=9&amp;p=0:0&amp;m=false&amp;c=-0.27394:0.000427564:-0.0170394&amp;d=-1.20974:1.20601:1.20751">http://photosynth.net/view.aspx?cid=436e315e-af30-4b45-a3c9-1f503ecc8afb&amp;i=0:9:1&amp;z=435.68801060008&amp;g=9&amp;p=0:0&amp;m=false&amp;c=-0.27394:0.000427564:-0.0170394&amp;d=-1.20974:1.20601:1.20751</a>

Note: All TinyURLS should be preceded by <http://tinyurl.com/>

<sup>29</sup> <http://www.tinyurl.com> ; 26 March 2009

## **Appendix B: Documents Presented to Users**

### **Informed Consent to Participate in a Usability Study**

#### **This study will help make computer software easier to use**

We would like to evaluate how easily and effectively people can use some software. We will use your experiences with the software infer how we might design interfaces better. By participating in this study, you can help make software easier to use. *We are evaluating the software, not your performance;* if you encounter any problems while using the software, it will give us clues as to where the software needs to be improved.

#### **We will ask you to use some office software**

We will ask you to do some basic tasks on Microsoft Photosynth. The tasks will not require you to reveal any personal or confidential information.

Before you start using the software, we will ask you for some basic background information. You can skip any question you do not wish to answer. We will be recording your voice and your screen interactions. After using the software, we will ask you a few more questions, which again you may skip if you do not wish to answer.

#### **We will keep your participation confidential**

This consent form is the only form that will include your actual name. All other paper forms and computer records will only identify you as “Participant N”, where N is a code. There will be no way to connect Participant N with your name.

The audio transcript will include your voice. These tapes will be restricted to members of the usability study team.

We may publish the results of the study in any of several forms: Internal reports to the company that developed it, conference presentations, and formal journal articles. In all of these reports, we will primarily report overall successes and failures with the software. We may also include specific comments you made or incidents that occurred while you used it. In these cases, we will not mention your name, and remove any details that might identify you.

We will maintain confidentiality of your participation and of your data from this study to the full extent allowed by the laws of British Columbia and Canada. Because this study will not ask you about sensitive or confidential information (such as potentially violent acts or child abuse), we consider it highly unlikely that a court will ever compel us to reveal your participation.



We will keep your data for three years. Paper forms will be stored in a locked office or file cabinet. Computer records and the video files will be stored on a secure network that is not publicly accessible. The secure storage will be maintained at the Burnaby campus of Simon Fraser University. Data will be securely destroyed after three years: Paper forms will be shredded and computer files (including the audio transcript) will be completely erased.

**The risks to you are minimal**

The only risks from participating in this study are the minor risks your normally take using a computer for a two hours. If you have serious discomfort using a mouse or keyboard (such as carpal tunnel syndrome), you should not participate.

**You will receive \$50 for participating**

We will present you with \$50 as payment for your participation. You will also have the satisfaction of helping make software easier to use.

**You can withdraw any time without cost**

You can withdraw at any time. If you feel uncomfortable during the study, please tell the experimenter and he or she will end the session. You will still receive the \$50.

At any time after completing the study, you may ask us to withdraw your data. Upon receiving such a request, we will securely destroy all paper forms and computer files (including the video) from your session. If your request comes after we have published something, we will not be able to remove your data from those publications, but your data will not be included in any publications after that.

**You can get the results of the study after it is complete**

We expect to complete the study around April 9 2009. After that date, you may contact Dr. Ted Kirkpatrick (see below for address) to receive a summary of the results.

You can send comments or complaints about the study

If you have any comments or complaints about the study, how it was conducted, or any discomforts from it, you may contact any or all of the following people:

<i>The supervising faculty member:</i> <b>Dr. Ted Kirkpatrick</b> School of Computing Science Simon Fraser University, Burnaby 8888 University Drive Burnaby, BC, V5A 1S6 ted@sfu.ca 778-782-4190	<i>The Director of Research Ethics at SFU:</i> <b>Dr. Hal Weinberg</b> Director, Office of Research Ethics Simon Fraser University 8888 University Drive Burnaby, BC V5A 1S6 hal_weinberg@sfu.ca 778-782-6593
--	--

**Agreement**

**I have read the above and agree to participate in this study:**

Signature \_\_\_\_\_ Date \_\_\_\_\_  
Witness \_\_\_\_\_ Date \_\_\_\_\_

**Pre-session questionnaire**

---

Please answer the following questions as completely as you can. If you feel uncomfortable answering any question, simply leave it blank. Your name will not be recorded on this questionnaire nor during the test. In our records you will be referred to as Participant \_\_\_\_\_.

**1. Please provide us with your age, or indicate your Age Group:** \_\_\_\_\_ years of age.

OR

19-23

24-28

29-33

34-38

39-43

44+

**2. If employed, what is your profession:** \_\_\_\_\_

**3. Online Photo Viewing Habits**

a. Do you own a digital camera? Yes No

b. Do you upload your photos to the internet? Yes No

c. Would you consider your photos artistic? Yes No

d. Please list the websites you upload photos to:

\_\_\_\_\_

e. Do you view family and friends' photos online? Yes No

f. Would you view a stranger's photos online? Yes No

g. Do you like to shoot panoramic photos? Yes No

h. Do you like to view panoramic photos? Yes No

i. How satisfied are you with the experience you get from today's online photo sharing websites? (1=extremely unsatisfied, 5=extremely satisfied)

1

2

3

4

5

N/A

**4. Have you heard of the website Photosynth?** Yes No

a. Have you visited the website Photosynth.com? Yes No

**5. Have you heard of CNN's "The Moment" campaign, (held on Barack Obama's inauguration day)?** Yes No

***Task 1:***

---

Please explore the synth and answer the following questions. You may change your answers as many times as you wish. Let me know when you are done and/or confident with your responses. Feel free to answer the questions in any order. You have upto 30 minutes to complete this task.

- 1) What do the white quadrilaterals represent?
- 2) What is the Halo's function?
- 3) When do you expect to see the Halo?
- 4) Are you able to move 360 degrees around the church? Why do you think this is?
- 5) What does the green region on the Halo represent?
- 6) What does the size of the green region on the Halo represent?
- 7) What does the solid white region of the Halo represent?
- 8) What do the scattered dots represent?
- 9) Why do people who appear at one point in the synth disappear all of a sudden?



8) What do the scattered dots represent?

9) Why do people who appear at one point in the synth disappear all of a sudden?

## ***End-of-Task1 Questionnaire***

---

Please answer the following questions as completely as you can. If you feel uncomfortable answering any of the questions, simply leave it blank. Your name will not be recorded on this questionnaire. In our records, you are referred to as Participant \_\_\_\_.

Do you think you know what the computer is doing in order for you to view this synth?  
(1 = Not at all, 5 = definitely sure)

1      2      3      4      5

If applicable, please explain:

Who or what do you think made this synth? (Microsoft, the programmer, an ordinary user, the computer etc?)

Did you find this task fun?  
(1 = very boring it, 5= really enjoyed it)

1      2      3      4      5

How confident are you with your answers to each question on the handout? Please tick your answer. (1=extremely unsure, 5=extremely confident)

	1	2	3	4	5	<b>I have no idea what the answer is</b>
Question 1						
Question 2						
Question 3						
Question 4						
Question 5						
Question 6						
Question 7						

Question 8						
Question 9						

Would you ever use Photosynth? (1=never, 5=definitely)

1      2      3      4      5

Would you recommend Photosynth to friends?    Yes                  No

If so, what is their confidence with technology?  
(1=extremely timid, 5=extremely confident)

1      2      3      4      5      Not Sure      Not Applicable

Do you think Photosynth is user-friendly?  
(1 = absolutely not, 5 = yes, completely)

1      2      3      4      5

Do you think your family-members or friends who are uncomfortable or intimidated by technology, would find Photosynth *fun* to use?  
(1 = absolutely not, 5 = yes, completely)

1      2      3      4      5

How in control of your movements were you?  
(1=absolutely random, 5=total control)

1      2      3      4      5

***Task 2:***

---

Please explore the following synth and, draw a floor plan of the downstairs. **Pay attention to sizes of rooms and locations of doorways.** Also, indicate the presence of any furniture and fireplaces that you see.



## ***Your Answer For Task 2***

---

Please explore the following synth and, draw a floor plan of the downstairs. **Pay attention to sizes of rooms and locations of doorways.** Also, indicate the presence of any furniture and fireplaces that you see.

## ***End-of-Task2 Questionnaire***

---

Please answer the following questions as completely as you can. If you feel uncomfortable answering any of the questions, simply leave it blank. Your name will not be recorded on this questionnaire. In our records, you are referred to as Participant \_\_\_\_.

Did this task change your understanding of how Photosynth, as a whole, works?

Yes    No

If Yes, please explain what changed and why. If No, please reiterate how you think Photosynth works:

Are you sure you found every room?

(1 = extremely unconfident , 5 = completely confident)

1      2      3      4      5

Did you get a clear sense of the size of rooms?

(1 = absolutely not , 5 = completely clear)

1      2      3      4      5

Would you say that it was easy to navigate around the house? If no, why?

Yes    No    Not Sure(please explain)

It was not easy because:

Did this task change your understanding of what Halo, Point Cloud or White Quadrilaterals do/ what their purpose is? If so, please indicate what you think they now represent, and what made you change your mind:

White Quadrilaterals: Yes                      No

Point Cloud:                      Yes                      No

Halo:                      Yes                      No

Did you find this task fun?  
(1 = very boring it, 5= really enjoyed it)

1      2      3      4      5

Do you think Photosynth is user-friendly?  
(1 = absolutely not, 5 = yes, completely)

1      2      3      4      5

Do you think your family-members or friends who are uncomfortable or intimidated by technology, would find Photosynth *easy to use*?  
(1 = absolutely not, 5 = yes, completely)

1      2      3      4      5

Any other comments (use back of the sheet if you need more space):

### ***End-of-Task3 Questionnaire***

---

Please answer the following questions as completely as you can. If you feel uncomfortable answering any of the questions, simply leave it blank. Your name will not be recorded on this questionnaire.

After exploring synths of your choice, did you get a better understanding of how Photosynth works? If Yes, please explain.

Yes    No  
Yes I did, because:

Amongst the synths that you viewed today, which was your favourite and why?

Amongst the synths that you viewed today, which taught you the most about how Photosynth works, internally?

For your favourite synth, if you were to divide a \$100 reward amongst those responsible for your experience, how much should each of the following teams receive?

The people who came up with the idea:        \$  
The people who came up with the design:     \$  
The user who took the photos:                \$  
Any Others?