Reconstructing the World in 3D: Bringing Games with a Purpose Outdoors

Kathleen Tuite[1], Noah Snavely[2], Dun-Yu Hsiao[1], Adam M. Smith[3], and Zoran Popović[1]

[1]Center for Game Science
Department of Computer Science & Engineering
University of Washington
{ktuite,dyhsiao,zoran}@cs.washington.edu

[2]Department of Computer Science
Cornell University
{snavely}@cs.cornell.edu

[3]Department of Computer Science
UC Santa Cruz
{amsmith}@cs.ucsc.edu

ABSTRACT

We are interested in reconstructing real world locations as detailed 3D models, but to achieve this goal, we require a large quantity of photographic data. We designed a game to employ the efforts and digital cameras of everyday people to not only collect this data, but to do so in a fun and effective way. The result is PhotoCity, a game played outdoors with a camera, in which players take photos to capture flags and take over virtual models of real buildings. The game falls into the genres of both games with a purpose (GWAPs) and alternate reality games (ARGs). Each type of game comes with its own inherent challenges, but as a hybrid of both, PhotoCity presented us with a unique combination of obstacles. This paper describes the design decisions made to address these obstacles, and seeks to answer the question: Can games be used to achieve massive data-acquisition tasks when played in the real world, away from standard game consoles? We conclude with a report on player experiences and showcase some 3D reconstructions built by players during gameplay.

Categories and Subject Descriptors

K.8.0 [**Personal Computing**]: General – Games; I.4.8 [**Image Processing and Computer Vision**]: Scene Analysis – Shape

Keywords

3D reconstruction, alternate reality, ARG, computer vision, games with a purpose, GWAP, game, photography, virtual world

1. INTRODUCTION

In the past few years, there has been increasing interest in building 3D models of entire cities for navigation, visualization, and planning applications. Recently, computer vision algorithms have been developed for completely automatic 3D reconstruction of buildings and landmarks from large photo collections [13], including images scraped from photo sharing sites such as Flickr. However, the resulting models are incomplete, as only "popular" viewpoints of popular landmarks are well-represented on Flickr. Reconstructing complete, high-resolution models of contiguous neighborhoods and whole cities requires many more photos, distributed more uniformly, than currently exist on such sites.

In PhotoCity, a hybrid alternate reality game with a purpose, we exploit the mobility of everyday players to collect carefully composed photographs of locations in an area of interest. Player photos include not just the standard, front-facing views of a location, but photos from a wide variety of viewpoints, allowing us to build a complete reconstruction. Bridging active collection of imagery with the human desire to playfully compete has required the design of a new gameplay form and forced the development of new computer vision technologies that can interact with players and usefully direct their effort.

We favor models reconstructed from photos because of the accuracy and realism such techniques provide. These techniques, borrowed from computer vision, are also cheaper and more automatic than hiring a 3D modeling artist.

Digital reconstructions are sensitive to the camera placement of the original images and resulting models often contain gaps and holes where additional views are needed. The data from a single pass with a vehicle-mounted camera array is often not dense enough to admit the quality of reconstructions we seek and is limited to where the vehicle can drive.

In PhotoCity, we have approached the problem of collecting useful data by actively involving players in an iterative growing and refining process. To play, players go to locations in the real world that correspond to locations under construction in the virtual world and take photos. We direct players to focus their attention on the gaps and fringes in a partial reconstruction in exchange for in-game rewards. This work primarily focuses on the technical challenges involved in bringing a complex vision system into contact with players. While we intend players to have fun playing PhotoCity, we have by no means exhausted the depth of gameplay that could be realized in this kind of system.

Unique to PhotoCity are the combined challenges of designing an effective game with a purpose (GWAP) and an alter-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

FDG 2010, June 19-21, Monterey, CA, USA Copyright 2010 ACM 978-1-60558-937-4/10/06... $10.00

nate reality game (ARG), using the real world as its gaming platform. This pairing requires both getting skilled players out in the world and communicating game rules and feedback to them. In this paper we explain how, by enticing players to pursue their own goals, we achieve our purpose, and describe the solutions to major challenges that arose.

2. RELATED WORK

2.1 GWAPs and ARGs

The phrase "game with a purpose" is most often associated with von Ahn's program of creating online computer games that harness human intellect to solve large-scale computational problems [15]. Established GWAPs such as von Ahn's ESP game [14] or Cooper's Foldit [5], which invite players to textually describe the contents of photos or to fold models of proteins into low-energy configurations respectively, use human intelligence as a stand-in for poorly-performing automated methods for the same problems. PhotoCity, while partially reliant on automatic algorithms, also relies on humans to capture the right set of photographs. In contrast to earlier GWAPs, our game actively engages the human player's body as well as their mind.

While GWAPs marry a real-world purpose to the mechanics of a fictional, on-screen world, ARGs aim to bring elements of the fictional world back out into meatspace where people live and work. In PhotoCity, the fictional world is literally modeled after the real world. Games such as *I Love Bees* [2] and *The Beast* [1] use familiar elements of the real world such as payphones and even in-person meetings with actors. These ARGs tend to be predominantly story-driven with a few human "puppet masters" [9] directing the game, while in PhotoCity, the game server maintains rigid control of the game.

On the frontier between ARGs and GWAPs there are games like Jane McGonigal's *Cruel 2 B Kind* [11]and *CryptoZoo* [10] which aim to engage the player's body for societal purposes such improving public kindness and exercise habits. These games bring purpose out into the real world where fictional elements are scattered to entice players, but they are missing the computational aspect of traditional GWAPs that results in large, useful datasets.

2.2 3D Reconstruction from Photos

The core technology behind the game is the computer vision system for automatically turning 2D photos into a 3D model, then updating that model as more photos are contributed. This system relies on three major components. The first is a batch reconstruction tool, for taking a set of images and producing an initial 3D model. The second is an algorithm for taking a new image and updating the model by adding newly visible points. These two stages produce a relatively sparse "point cloud" derived from structure from motion techniques [6]. The final component is a multi-view stereo system for taking this sparse point cloud and generating a much denser model that is more suitable for rendering in a graphics pipeline. We now briefly describe each of these components.

Batch reconstruction. Each model in our game starts as a small "seed" reconstruction created from about fifty photos. To build an initial 3D model from these initial seed photos, we use a batch reconstruction method derived from the Bundler structure from motion (SfM) system of Snavely *et al.*, [13]. This tool takes an unordered collection of photos and automatically computes a camera position and orientation for each image, as well as a sparse point cloud reconstruction of the scene itself (similar techniques are used in Microsoft's Photosynth tool [12]). An example sparse reconstruction is shown in Figure 1(a). A key technology behind Bundler is the ability to detect and match distinctive feature points (e.g., the corner of a window, the texture on an inscription) between different images.

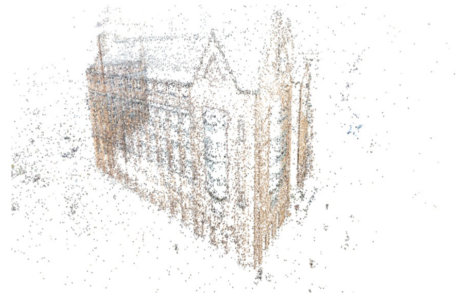

(a) Sparse point cloud

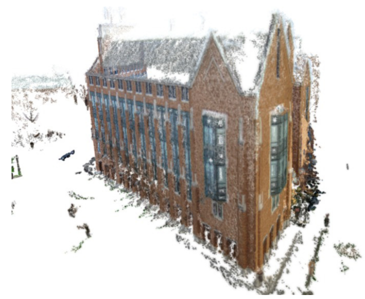

(b) Dense point cloud

Figure 1: Two views of a 100-photo model

Incremental updates. When a user uploads a photo to add it to a particular model, we need a fast way to determine where that photo was taken (and which direction the camera was pointed), while adding newly visible points to the model (which will add to the user's score). To do this, we match the 2D feature points of the new photo to the 3D points in the current model and then use pose estimation techniques [8] to find where the new photo fits into the scene (this optical localization process is 10–100 times more accurate than GPS). We then find nearby photos and use the combined information of new and old photos to synthesize new 3D geometry, thereby "growing" the model.

Dense reconstruction. The batch and incremental SfM techniques described above maintain a very sparse point cloud representation of the scene, which is not intended to be dense enough to produce attractive renderings. In order to densify the model, we apply the patch-based multi-view stereo (PMVS) system of Furukawa and Ponce to each sparse 3D model [7]. The PMVS system finds many additional pixel matches between the images, generating a much denser point cloud, such as the one shown in Figure 1(b). Because multi-view stereo is very time- and memory-intensive, we

Figure 3: Gameplay cycle

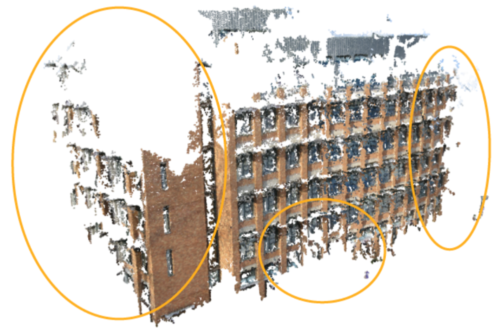

Figure 4: 30-photo seed with holes

run PMVS on each 3D model as a batch process only once a day, generating updated results that can be viewed by the users.

Interactive data collection. What sets PhotoCity apart from techniques that build reconstructions from photos scraped from the Internet [13] is the involvement of players in the incremental process, and the guiding of players to take photos that have never been taken before. Even the largest batch technique by Agarwal [3], which can process hundreds of thousands of photos and reconstruct very large models, produces disjoint reconstructions corresponding to separate popularly photographed locations. To get a truly complete model, these disjoint models need to be expanded and merged together. Due to the extreme computational and data collecting efforts required, this can only happen as an incremental process with humans in the loop.

3. GAME MECHANICS.

With the purpose of our game being to collect the kinds of images that would yield a high-quality reconstruction, we were faced with the challenge of translating this task into an inviting gameplay experience. In our players' eyes, PhotoCity is about capturing virtual flags anchored in real-world locations and vying for nominal ownership of familiar landmarks.

The core mechanic of our game involves players inspecting the state of the game world on a map, taking photos at locations of promising in-game value, and uploading the photos, and then observing the results of their play. Through repeated cycles of this process, summarized in Figure 3, players introduce new geometric points in the reconstructions that correspond to points of detail in the real world.

Accumulating points lets players control flags and own models in the game's fictional world. We will briefly go through each element of the game: where the models come from, where the flags come from, how exactly a player receives points, and the requirements for capturing flags and taking ownership of models.

A model is a partial building that starts off as a seed generated from a batch of photos of the real building. Recall that the number of photos used to make a seed is about fifty. As a result, models in their starting state only span one face or one corner of a building and have rough edges and large holes where data for the building has yet to be captured. Figure 4 shows the size of a seed made from thirty photos. The seeds are added either by the game designers or the players themselves, but go through a manual approval process that rejects offensive material and low quality reconstructions.

To anchor the model to the world, we manually align an overhead view of the model with a satellite map (Figure 7(c) shows one such anchoring). Flags are automatically placed on the model so that each flag has a map location corresponding to a position on a wall, or sometimes a tree or cluster of non-building points. To take a photo of a flag, one must face the building the flag is on and stand far enough away get distinctive features of the building in the shot. As the models grow, new flags automatically appear at the edges of the models for players to capture. If the flags at a model's edge are already captured, a player can work to grow the model enough to spawn a new flag.

The typical player looks at the map of the game, identifies the flags she wishes to capture, and then takes photos of that portion of the building. For a photo to earn points, it must (1) overlap with enough existing points in the model and (2) overlap the empty space next to a model. The first requirement allows the photo to connect with the model and have its position within the model automatically calculated. The second requirement, that the photo look beyond the existing model, allows new points to be added to that void as soon as there are enough other photos to triangulate the 3D positions of those features. A photo can add up to several thousand new points (shown in Figure 6(c)).

To earn points at a flag, a player simply generates new 3D points in a column around a flag with a radius of about thirty feet. 3D points translate directly into game points. The team with the majority of points at a flag captures that flag. Flags often accumulate tens of thousands of points, so flags that one team owns but that have fewer than 2,000 points total are considered "disputed" because they are relatively easy for another team to steal.

While teams control flags, individual players control individual models. If one player has over 10,000 points at a model (roughly equivalent to making a major contribution at five different flags) and has more points than any other player, that player owns that model.

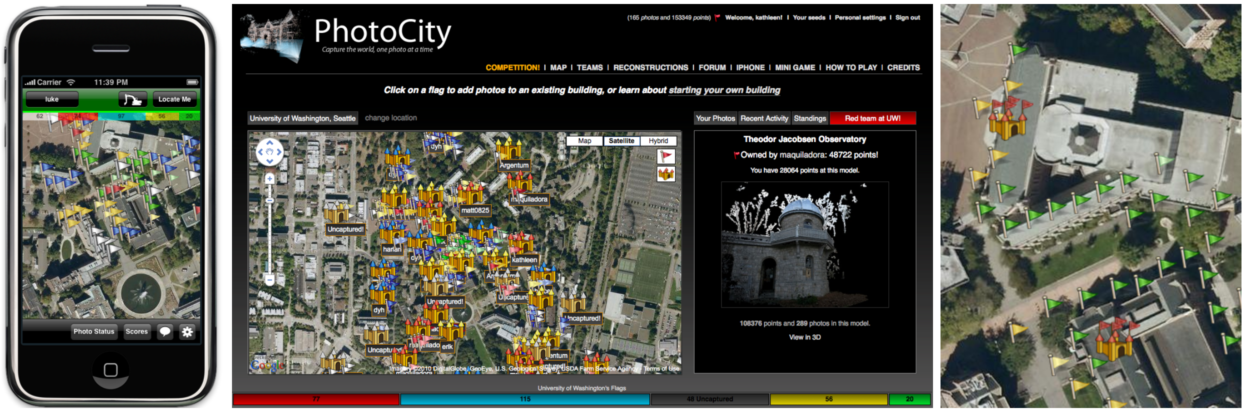

(a) iPhone (b) PhotoCity in the web browser (c) Flags on map

Figure 2: PhotoCity screenshots

The photography must happen outdoors, in the vicinity of certain buildings, but the game can be played in two different ways. Players can bring their cameras inside and use a high-bandwidth link to upload photos to the website, or players can play using a custom iPhone application that lets them submit photos and get feedback in one seamless process. Figures 2(a) and 2(b) show the two game interfaces.

Putting all of these elements together, let us follow the activity of a hypothetical player, Sam, through one cycle of play. Sam checks the game's map display, finding a model on the map that he could gain control over with only a few thousand more points. Sam takes his digital camera outside and takes a series of photos. Sweeping a wide arc around the target building, taking a new picture every few steps, he collects perhaps 200 photos in ten minutes. He knows this string of very-similar photos will match well, producing a sufficient number of new 3D points for him to capture the model. Further, his tour included shots of nearby buildings which might cause the target model to expand. Returning to a web browser, Sam uploads his photos in a large batch (having selected the model he intends to improve from the map). Even as the photos are uploading, the game server begins accumulating new geometry into the model, visibly updating Sam's score. Minutes later, enough photos have been integrated to name Sam the new owner of the target model. Sam now proceeds to show his friend-and-rival, Tim, the recent conquest on the map display, starting a new cycle of play for both of them. Tim is already plotting how to capture the new flags that have spawned at the edges of Sam's model.

These mechanics result in a virtual world, populated by 3D reconstructions, which grows organically through continued player interaction.

4. DESIGN CHALLENGES

Designers of alternate reality games face the challenges of getting players out to play, and of integrating the fantasy game world into the real world. Separately, games with a purpose must be enjoyable experiences that also produce valuable results. PhotoCity must do both at the same time. Our hybrid AR-GWAP must sufficiently inform players of the alternate reality and equip them with the tools and feedback to make them effective, while also being fun to play.

4.1 GWAP Challenges

The purpose of PhotoCity is to collect photos that densely cover a scene. To achieve this purpose, we must design the game to attract and motivate players and make players perform the correct tasks, design challenges that are common to all GWAPs.

Attracting and motivating players. The more people PhotoCity appeals to, the more people will participate and contribute photos. Simply growing 3D models without any game-related incentives may have worked, but it would have only attracted a certain type of person already interested in photography or 3D modeling. Bartle [4] describes four types of players, the *explorer*, the *achiever*, the *killer*, and the *socializer*, and we believe the game we have built PhotoCity into will appeal to all of them.

At its core, PhotoCity is well suited to *explorers* who not only get to chart out the virtual 3D world themselves by starting new seeds, as shown in Figure 5, but also experiment with camera placement and learn new strategies for taking effective photos. We wanted PhotoCity to appeal to a wider audience, however, and have more traditional game elements, so we added flags for teams to capture and allowed individual players to own models. The flags and the models give *achievers* something to collect and give the *killers* resources to fight over. Lastly, the we cater to the *socializers* by providing a chat room, integrating with Facebook to share accomplishments.

The lesson here is that while our original concept for PhotoCity held appeal for a certain type of player, we had to augment the design to widen our potential audience. The different players can all play the same game without interfering with each other (explorers can still explore while killers roam) though we realize the highly competitive players may intimidate new or casual players. One solution would be to let players self-organize into groups with similar play styles.

Making player input count. A big challenge in designing

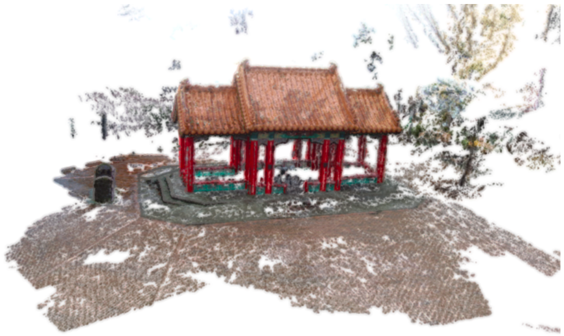

Figure 5: Seed model added by a player

any GWAP is to get the players to actually perform the right tasks. In our case, we want players to take photos all over campus such that the photos can be used to build a 3D reconstruction.

We would like to collect photos of all sides of a building, so instead of giving the instructions, "take photos of this building", which would likely result in a lot of similar photos of the most popular facade, we use flags to direct the users to specific parts of the buildings. The game automatically places flags along the walls of a reconstruction and colors the flag by the the team that currently owns it and by the number of points. This way, players can identify the parts of the building that are already thoroughly photographed and the edges or holes in the reconstruction that need work.

A new flag cannot be captured with just a single photo. Similarly, a battle over a flag cannot last forever: the 3D model becomes saturated with points and additional photos do not add new information. This makes the flags a limited resource; players can battle over flags (contributing different views of that part of the building) but eventually they must work on new areas of the model if they want to continue gaining points.

We wanted the players to accomplish the task of photographing specific areas, so we placed a flag resource in the game that the players could only get by performing the task. The value of the resource was also useful for generating the right amount of activity in a certain area. The only issue was that other resources in the game (overall points, model ownership) encouraged players to keep submitting photos of saturated areas.

4.2 ARG Challenges

We classify PhotoCity as an alternate reality game because it is a fictional, virtual world that players change by taking actions in the real world. The challenges specific to all ARGs include getting players out into the real world and conveying the fictional game world to them.

Getting Players Outside. Many of the challenges of designing an alternate reality game arise from the fact that the gaming platform is the real world. In a traditional video game, many players who start the game may not make it past the first level due to confusion, boredom, or other reasons. In our game, a player can sign up fully intending to go out and take photos at the next opportunity, but uncontrollable external conditions such as weather and daylight can discourage them from play.

One solution to getting the player outside, with camera, and playing the game, is to put the game on a mobile phone. That way, if a player is walking down the street or on her way to class, if she is near a building that is active in the game and happens to think about the game, she can pull out her phone and start playing. We chose the iPhone as the first platform because of its popularity and camera (2+ megapixels).

For the players who do not have iPhones, the game website actually provides a richer portal to the game. Since having players go outside with their cameras is already asking a lot, we strive to make the rest of the experience as easy as possible. The website allows players to select any flag on a particular model on the map and upload all their photos of that model through that flag in a large batch. Players seem to enjoy seeing the game absorb hundreds of their photos at a time.

The lesson here is to remove as many obstacles as possible. Designers should strive to turn daunting or mundane tasks (such as uploading large numbers of photos) into enjoyable activities, or at least hide or minimize them.

Conveying the Fictional World. We must convey enough information about the virtual world to the players to show or teach them what to do. In addition to the flag and model information shown on the map, we have three additional display types.

The first type of display is an image that highlights which points on a model correspond to a particular flag. The points that are in a fixed radius around the map location of a flag contribute to the score of a flag, and these are the points we color in the image. Figure 6(a) shows what the player would see for a newly created flag on the edge of a model, directing them to take photos that include the points in the orange rectangle to capture that flag.

The second is the player contribution image, which is generated every time a photo successfully registers with a model. In this image, the model is rendered from the same viewpoint as the player's camera so she can verify that the photo registered correctly. This image also highlights the new points the photo added. In Figure 6(c), the original photo is placed next to the contribution image with the 800 new points highlighted in green. Even if players ignore all other instructions, we hope that if they take enough photos and see these contribution images, they can come to understand how to take high-scoring photos.

The third type of display is an interactive 3D viewer that lets the user view the sparse point cloud of a given model. A static image of a point cloud is often hard to interpret, but player-controlled movement provides the necessary depth cues to understand the point cloud as a 3D shape rather than a blob of 2D points. The points in the interactive viewer also appear in the order in which they were added to the model and show the user exactly how the model expanded.

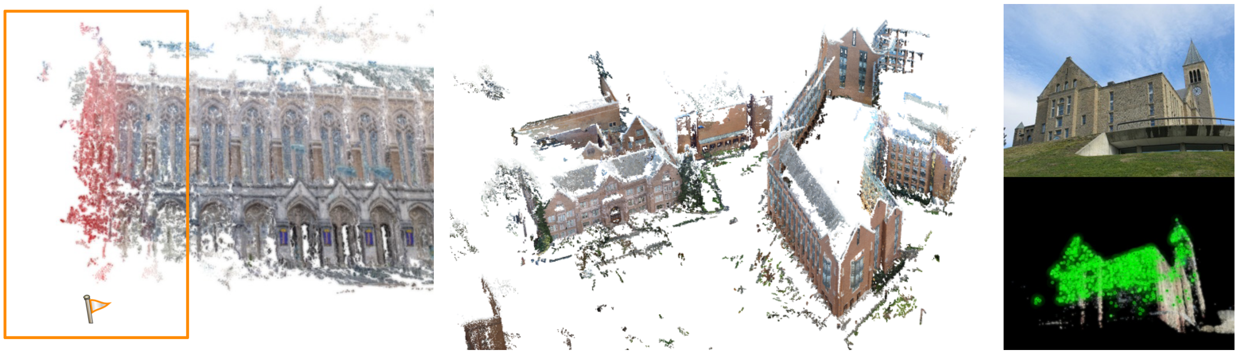

(a) Flag-specific visualization

(b) Dense, multi-building model

(c) Player contribution (800 points)

Figure 6: Visualizations to help the player take effective photographs

	May Trial	Dec. Trial
Signed Up	20	104
Submitted Photos	15	25
Successfully Added Photos	9	21
Median Number of Photos	80	40
Maximum Number of Photos	649	5867
Active >1 Day	no data	17
Active >1 Week	n/a	14
Number of Pro Players*	4	5
Photos Submitted by Pros	722	12330
Successful Pro Photos	410	8820
Number of Novice Players	11	20
Photos Submitted by Novices	942	2050
Successful Novice Photos	312	1500
Photo Success Rate	42%	72%
Success Rate: Novice Only	33%	73%
Success Rate: Pro Only	57%	72%

*Game designers and two top players

Table 1: Numerical Evaluation

We found it useful to expose a variety of visualizations to players. Different players latched on to different types of displays; some players primarily used the map to identify where to take photos, while other players relied more on the flag visualization images.

5. RESULTS

To evaluate how well we overcame these design challenges, we ran two game trials, one in May 2009 that lasted for one week, and one beginning in December 2009 that lasted for two months. Both trials took place on the University of Washington campus using the same set of building seeds, which were reset to their initial states before each trial. We report player statistics for both trials and present the reconstruction results from the second trial.

5.1 Game Trials

The first trial was publicized through a department mailing list and attracted 20 players, 15 of which took photos, and 9 of which successfully added points. The only game interface at the time was the website, a predecessor to the one pictured in Figure 2(b). This trial lasted one week and the reconstruction progress was duplicated and surpassed by players in the second trial.

To promote the second trial, we ran an advertisement on Facebook targeted at local students who listed photography as an interest. Thousands of people saw the advertisement, a fraction of those signed up for a game account, and a fraction of those actually uploaded photos and played. The biggest problem is that this game is unfamiliar and people do not know what to expect. Without coaching, new players are reluctant to take the hundreds of photos per session that expert players regularly do. We realized that we need to make it more apparent that it is highly beneficial to the player and the game to take many photos.

Foldit, a similarly complex game with a purpose, has two additional resources for new players that we hope to build for PhotoCity: videos of the game in action, and a set of introductory levels that teach the necessary skills. A tutorial would be especially useful to new PhotoCity players, and to our goal of retaining as many players as possible, as it would engage them right away without requiring them to wait for the right photography conditions.

By introducing the iPhone in the second trial as an alternate way to play, we expected the cycle of play to switch from about a day (photographing during the day, uploading at night) to something much tighter. However, some players still migrated toward the website, even when they used their iPhones as cameras, because it was easier for them to take many photos and upload them from a computer than upload them over the mobile network. We need to take this into consideration in the future and find a way to balance mobile phone bandwidth with the mechanics of our game.

Between the two trials, players appear to have gotten better at taking successful photos. This is perhaps due to a better user interface on the website (see Figure 2(b)) and to better visualizations of the models, including rendering the dense models (Figure 1(b)). Late in the second trial, we introduced the new visualization (pictured in Figure 6(a)) that highlights the shape of a specific flag zone. In talking to players, those that began the game after this addition started off less confused and more confident about what initial actions to take than previous players.

One surprise was that we expected the experienced "pro"

(a) Detail

(b) Circumnavigated

(c) Model growth from overhead, 1200 photo difference

Figure 7: Reconstruction results

(expert) players to have a better ratio of successful photos to all photos, but in fact, that ratio (about 72%) is the same for both pro and novice players. This can be explained in two ways. First, given the right feedback, players can learn to improve their effectiveness on their own. Second, the more experienced players try to make bigger, riskier moves, such as expanding buildings around corners or onto nearby buildings, which can be challenging and require a large number of photos.

5.2 Reconstructions

We started with 12 seed models at the beginning of the second trial, listed in Table 5.3. The 13th model was added during the game by one of the players. The first two columns show the number of photos and points in the initial seed models and the last two columns show the growth of each model. Figure 7(c) shows the growth from overhead of the model *Fountain-facing Corner of EE*.

Two very exciting things happened with the reconstructions that demonstrated the success of the game. First, as shown in Figure 7(b), players circumnavigated several buildings, despite obstacles like trees and hedges at the corners of buildings that were hard to reconstruct. Second, players learned how to expand models onto other physically disconnected buildings by taking photos of one building with the second building in the background. Players actually expanded seeds to touch and overlap nearby seeds. Figure 6(b) shows a seed that started as just the corner of the angular building expanding to include several adjacent buildings.

Towards our goal of harnessing player effort to build 3D reconstructions by taking photos, the 25 players in our second trial submitted over 11,000 new photos generating over 3,000,000 new sparse-model points.

5.3 Reactions to Game Elements

Did the PhotoCity players have fun? We interviewed several of them and also observed their interactions with other players.

New players who managed to score points in their first set of photos felt good about themselves, even if not all of their photos were successfully added to a model. The positive effect of a matching photo was greater than the negative effect of a failed photo.

Players who started the game when there were a large number of white (uncaptured) flags were very excited to see the flags change to their team color. However, players who came to the game when many photos had already been taken and most of the flags captured were overwhelmed and felt less able to make a difference (a common occurrence in multiplayer online games).

Despite the game being driven by the mere acquisition of points, flags, and buildings, players discovered "expert strategies" that both made them more interested in the game, and allowed them to acquire vast quantities of points. For example, when one player had built a sufficiently dense model (so dense that new photos would match but not add new points), a second player came up with a strategy to reclaim the flags by expanding nearby models into the flag-space already inhabited by the first model. This interesting emergence was a direct result of attaching abstract mechanics to locations in the physical world.

6. CONCLUSION

We have described PhotoCity, a game with a purpose that draws its players outdoors to participate in the construction of a virtual world. In creating a game at the intersection of two genres, alternate reality games and games with a purpose, we encountered challenges on a new scale. PhotoCity is both vastly more data intensive than other GWAPs and more process intensive than traditional ARGs, even requiring the development of new algorithms and deployment of them in a distributed system to maintain interactivity. Initial trials of our game have already resulted in a collection of highly-detailed 3D models, demonstrating that hybrid AR-GWAPs, played out in the real world, can be effective at collecting vast quantities of data, provided the games are interesting, fun, and equip players with the right tools and training to accomplish the task at hand.

By sharing the design of our game, we hope to stimulate further discussion on games with a purpose that extend into the real world. A great many computational problems can be solved in the presence of masses of data, so that cheaply and effectively collecting this data quickly becomes the primary issue. The reader is invited to visit the game's reconstruction gallery[1] and join our ongoing effort to reconstruct the world in 3D by submitting his or her own images.

[1] `http://photocitygame.com/reconstructions.php`

Model	Starting Photos	Starting Points	Ending Photos	Ending Points
CSE Front Entrance	58	18,131	515	148,189
Electrical Engineering	30	13,133	1,230	380,102
James J. Hill Statue	48	51,461	406	224,783
Southeast Corner of CSE	27	9,597	935	304,030
Fountain-facing Corner of EE	68	38,439	979	321,805
Commodore Apartments	41	9,973	581	66,777
Allen Library	92	28,945	840	131,558
Engineering Library	74	8,598	584	62,095
Guggenheim	101	85,239	1,394	622,355
Mechanical Engineering	46	27,022	741	273,889
Mary Gates (West Side)	32	36,891	1,488	761,777
Suzzallo Library	41	36,566	1,526	492,871
Hing Hay Park*	73	46,965	116	54,333

*Seed added by player during trial

Table 2: Model Growth

7. ACKNOWLEDGMENTS

Many thanks to the efforts of Sylvia Tashev, Nadine Tabing and Juliet Bernstein, and to our star players, Erik, Daniel, and Yuki. We also thank Intel, Google, and Microsoft for their support.

8. REFERENCES

[1] 42 Entertainment. The beast, 2001. `http://www.42entertainment.com/beast.html`.

[2] 42 Entertainment. I love bees, 2004. `http://www.42entertainment.com/bees.html`.

[3] S. Agarwal, N. Snavely, I. Simon, S. M. Seitz, and R. Szeliski. Building Rome in a day. *International Conference on Computer Vision*, 2009.

[4] R. Bartle. *Designing Virtual Worlds*. New Riders Games, 2003.

[5] S. Cooper, A. Treuille, J. Barbero, D. Baker, Z. Popovic, and D. Salesin. Foldit, 2008. `http://fold.it`.

[6] F. Dellaert, S. Seitz, C. Thorpe, and S. Thrun. Structure from motion without correspondence. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'00)*, June 2000.

[7] Y. Furukawa and J. Ponce. Accurate camera calibration from multi-view stereo and bundle adjustment. *Int. J. Comput. Vision*, 84(3):257–268, 2009.

[8] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521623049, 2000.

[9] J. McGonigal. The puppet master problem: Design for real-world, mission based gaming. In P. Harrigan and N. Wardrip-Fruin, editors, *Second Person: Role-Playing and Story in Games and Playable Media*. The MIT Press, 2007.

[10] J. McGonigal. Cryptozoo, 2009. `http://cryptozoo.ning.com/`.

[11] J. McGonigal and I. Bogost. Cruel 2 b kind, 2006. `http://www.cruelgame.com/`.

[12] Microsoft. Photosynth, 2008. `http://photosynth.net`.

[13] N. Snavely, S. M. Seitz, and R. Szeliski. Photo tourism: Exploring photo collections in 3d. In *SIGGRAPH Conference Proceedings*, pages 835–846, New York, NY, USA, 2006. ACM Press.

[14] L. von Ahn. Esp game, 2003. `http://www.gwap.com/gwap/gamesPreview/espgame/`.

[15] L. von Ahn. Games with a purpose. *Computer*, 39:92–94, 2006.