
Landspotting: A Serious iPad Game for Improving Global Land Cover

Tobias STURN¹, Dietmar PANGERL¹, Linda SEE²,
Steffen FRITZ² and Michael WIMMER¹

¹Vienna University of Technology/Austria · tobias.sturn@vol.at

²International Institute for Applied Systems Analysis, Vienna/Austria

This contribution was double-blind reviewed as full paper.

Abstract

Current satellite-derived land cover products, which are very important for answering many crucial research and policy-related questions, show huge disagreements. In this paper we present a serious game for the iPad with the purpose of improving global land cover data. We describe the game, discuss the design decisions made and outline the challenges faced while developing the game. We evaluate how well the players are able to annotate land cover by comparing the game against expert validations collected using the Geo-Wiki tool and provide evidence that games can be a useful way to increase the quality of global land cover.

1 Introduction

At present, there is no single satellite-derived global land cover product available that is accurate enough to answer important research and policy relevant questions for humanity such as “Does the climate change?” or “How much additional land is available to grow biofuels?”, or to tackle problems of food security (FRITZ & SEE 2006).

In the last decade, three global satellite-derived land cover products have been created: MODIS, GlobCover and GLC-2000. A pixel-by-pixel comparison of these three different land cover datasets reveals huge areas of the world where they do not agree. For example, the three products have an overall forest and cropland disagreement of 893 million hectare, and estimates of cropland differ by up to 20%, which makes these maps very unreliable (FRITZ et al. 2011).

To improve these land cover products, the Geo-Wiki tool (FRITZ et al. 2012) has been developed to crowd-source user validations and to create a new and better hybrid map. Unfortunately, it is quite difficult to gather enough validations without giving the users any incentives to validate land. Therefore, the goal of the Landspotting project is to use games to crowd-source land cover validations. The incentives for the players to contribute should be the same as for people to play games, namely that they are fun and engaging. Using games to improve global land cover datasets has not been attempted before and it represents an exciting new way to demonstrate how games can be used for a serious purpose.

In this paper, we present an iPad game for improving global land cover. We describe the decisions that were made in the design of the game and we describe the challenges that were faced while developing this game. We also explain how we try to make the game fun and engaging and how we give the player incentives to categorize the Earth correctly. An evaluation is given regarding how well the players are able to annotate land cover by comparing the game against expert validations collected using Geo-Wiki. These results provide evidence that games can be a useful way to help increase the quality of global land cover.

The main contributions of this work are:

- To analyse how well humans can annotate land cover.
- To analyse how well humans can distinguish between different land cover types.
- To validate the differences between experts (via Geo-Wiki) and non-experts (the players in the game)
- To compare the performance of a wiki with a serious game.

2 Related Work

Today it is very common to utilize platforms in which volunteers help science by making tasks which only humans can perform. For example, in Geo-Wiki, people are asked to annotate land cover. A project which is very similar to Geo-Wiki is Galaxy Zoo (LINTOTT 2011), which invites the general public to visually inspect and classify galaxies via the Internet. Over 150,000 participants have already given over 60 million classifications, which makes Galaxy Zoo the largest astronomical collaboration in history. The results show that the general public can reliably classify large sets of galaxies with a similar accuracy to professional astronomers.

Platforms like Galaxy Zoo and Geo-Wiki do not exploit gamification methods such as high-score lists or other achievements. This may be one reason why the number of Geo-Wiki users is still small, i.e. 1,500 users. In order to engage more participants, so called serious games (GWAP, games with a purpose) have been developed that try to make the task fun. One of the most famous examples is FoldIt (COOPER 2010), a puzzle-game in which the players fold proteins. For example, within ten days, the FoldIt players accurately approximated the crystal structure of M-PMV, an AIDS-causing simian virus, which had been an unsolved problem for scientists for the last 15 years. Just recently, with the help of the 230,000 FoldIt players, a novel algorithm for protein folding has been found which outperforms previously published methods (KHATIB 2011).

As evidenced by FoldIt, games can indeed help to create impressive scientific results. Moreover, several types of mechanics for collecting correct meta-information through games have been developed. However, there are still no serious games which have achieved the tremendous success of games like Farm Ville or Angry Birds, with hundreds of millions of players per month. As has already been described by Markus Krause (KRAUSE 2010, SMEDDINCK 2011), the currently tightly coupled GWAPs do not use the whole scope of game-design elements. Instead, he proposes a disjoint game design where the task to be solved should be part of the mechanics, but no longer the dominant element. Therefore, our goal in the Landspotting project is to take successful games as reference and to copy these

already proven game mechanics to create games with a scientific purpose that appeal to a large audience. In the game presented in this paper, we have therefore copied the successful game mechanics of very popular tower defence games. The game and the design decisions made are described in more detail in the following section.

3 Game Design

This chapter describes the design process and the design decisions made for how to get correct land cover data and how to make the Landspotting game fun and appealing to reach high player numbers. We have decided to create a number of games and select from these games the most promising ones on which to concentrate further. Therefore, we have developed four game prototypes, from which one will be discussed in this paper. Our team consists of 2 professors, 2 project assistants and 5 students, which imposes restrictions on graphics, 3D models, sound effects, storyboard and character design, and thus the explored game ideas have to fit within these resource restrictions.

3.1 Disjoint Game Design

The first decision was to use a disjoint game design (KRAUSE 2010), i.e., one where the actual task of ‘Landspotting’ the earth is only part of the game mechanics and is no longer the dominant element. Through this we are not limited to the task itself and are able to use the full spectrum of game design. In order to make fun and engaging games, we emulate already successful games and try to add the additional ‘Landspotting’ task as seamlessly as possible to these already proven games. In the game described in this paper, we try to emulate the proven tower defence mechanics such as in *Plants vs. Zombies* (POPCAP GAMES) in which enemies approach a city and the player has to build defences like walls, traps and warriors to defend the city against the enemies.

3.2 For Which Platform?

We analysed several platforms and determined how feasible they are for a Landspotting game prototype. Game consoles were discarded due to lack of the required development kits. We have taken the following platforms into consideration: PC standalone application, standalone browser games, Facebook browser games, iPad, iPhone, and Android phones. We found that for this game the AppStore with the iPad would be the most promising option for several reasons. In the last couple of years with the arrival of the iPad, the use of tablets has increased tremendously. A recent study has revealed that 31% of all US internet users use tablets to browse the web. By the end of 2013, this number is forecast to increase to 47% (ONLINE PUBLISHER ORGANISATION 2012). Also the player numbers on the AppStore are very impressive. For example, the game *Draw Something* (ZYNGA) peaked at over 15 million users per day. Finally, the iPad is the perfect device to draw land cover maps.

3.3 2D or 3D

In a series of early prototypes, a number of 3D games were developed. One example is an archaeology game in which the player has to find hidden treasures like dinosaur bones or

diamonds in a 3D scene with different tools like a pick or a shovel. Depending on the tools used, the land cover data could be estimated. However, it soon became apparent that 3D was not suitable for determining land cover. Most terrain, especially those where most disagreements in the datasets occur (e.g. Africa), is rather flat, and therefore the elevation information is not important. Also, close to the ground, very good texture quality is needed. Otherwise the texture cannot be recognized anymore because of the high magnification. Finally, much more time is needed for creating assets like 3D models. Therefore, we decided to concentrate on 2D games instead. This also makes sense for an iPad game, where the majority of games are in 2D.

3.4 Ensuring Correct Data

One of the big challenges for a GWAP is getting correct data from the players, since gamers will be looking for limitations in the game mechanics. One important aspect is the main incentive for playing the game, i.e. to have fun, so providing the correct data should be the most rewarding strategy for a player in this respect.

We have analysed different approaches to guarantee data correctness (AHN 2004, AHN 2007, LAW 2007, CHIEN-JU 2010). We decided to use output agreement. The more the players agree on the land cover they have given as output, the higher scores they achieve and the faster they progress in the game. We try to use presents and achievements as positive reinforcements to educate the players to choose the correct land cover on the earth. Besides ensuring data correctness, we try to collect as detailed land cover information from the players as possible. We do this by letting the players annotate the land with simple drawing tools like a brush instead of using a simple classification. Thus, we gain very detailed land cover maps, containing more information than simple annotations.

3.5 Game Mechanics

The game presented in this paper is an iPad tower defence game which emulates the very successful game *Plants vs. Zombies* (POPCAP GAMES), which sold more than 300,000 copies in the first nine days. A screenshot of the game can be seen in figure 1. A player has to defend his or her city from quickly approaching enemies by building warriors, towers, walls, and traps to stop these threats from reaching the city.

The way we integrated Landspotting was via the idea of an “exploration task”: for building something to defend against enemies reaching the city, the player has to explore land to gather resources which are needed for building. The player draws on a land cover image using standard image editing tools (brush, fill bucket). For drawing these land cover maps, the iPad is thereby the perfect tool because it feels much more natural to draw with the fingers than with the mouse. The more the player agrees with the ratings of a Geo-Wiki expert about the land cover, the more resources they collect and the quicker they are able to finish building something and therefore the better they can defend their city. Through this mechanic (*output agreement*), we make sure that the players have an incentive to correctly designate the land cover because they have to agree with the experts from Geo-Wiki to succeed.



Fig. 1: Screenshot of the game Plants vs. Zombies

To motivate the players to play the game and to compete with their friends, Game Center is supported in the game, which provides a leaderboard and achievements to unlock. The game also integrates Facebook and Twitter support. Furthermore, sound effects and music have been added to make the game more immersive.

4 Results

In this section we analyse the data the players generated with the iPad Game. We show how well humans can annotate and distinguish land cover and validate how well the players perform compared to Geo-Wiki.

4.1 Overview

The game was released on January 7, 2013 on the App Store. In the first 2 weeks the game was downloaded 351 times and collected 1371 land cover validations from the players from 196 games that were begun. The area of one land cover pixel is 2 km^2 , which means that land cover data of 2742 km^2 has been created by the players.

4.2 Game and Geo Wiki Comparison

The results collected through the game were compared against the expert ratings gathered by Geo-Wiki. The comparison between the game and the experts was undertaken by 2 independent validators who categorized each comparison (table 1) into one of the following categories:

- Game and Geo-Wiki wrong
- Game wrong, Geo-Wiki correct

- Geo-Wiki wrong, Game correct
- Both equally correct
- Both correct, Geo-Wiki better
- Both correct, Game better
- Unknown: It is not possible to say something about the land cover.

Table 1: Comparing the game (players) against Geo-Wiki (experts)

Category	Validator 1	Validator 2	Only Ratings where both agree
Total Validations	1227	1255	695
Game and Geo-Wiki wrong	46 (3.8%)	27 (2.2%)	8 (1.2%)
Game wrong, Geo-Wiki correct	289 (23.6%)	230 (18.3%)	170 (24.5%)
Geo-Wiki wrong, Game correct	71 (5.8%)	85 (6.8%)	37 (5.3%)
Both equally correct	604 (49.2%)	472 (37.6%)	400 (57.6%)
Both correct, Geo-Wiki better	60 (4.9%)	194 (15.5%)	32 (4.6%)
Both correct, Game better	89 (7.3%)	214 (17.1%)	42 (6%)
Unknown	68 (5.5%)	33 (2.6%)	6 (0.9%)

As can be seen in table 1, the two independent validators came to nearly the same conclusions and therefore the data presented here are quite certain. In the majority of cases (57.6%) both the game and Geo-Wiki are both equally correct. In 24.5% of cases the game is wrong and Geo-Wiki is correct whereas only in 5.3% of cases Geo-Wiki is wrong and the game is correct. Only in 1.2% of cases both the game and Geo-Wiki are wrong.

Table 2: Results of the game (players) and Geo-Wiki (experts)

Category	Validator 1	Validator 2	Only Ratings where both agree
Geo-Wiki correct	89.9%	90.8%	93.5%
Geo-Wiki wrong	10.1%	9.2%	6.5%
Game correct	71.1%	79%	74.2%
Game wrong	28.9%	21%	25.8%

Table 2 depicts the results of the game and Geo-Wiki. The experts using Geo-Wiki were correct in 93% of cases while 74% of the data collected by the game was correct. This means that Geo-Wiki performs better than the game which is unsurprising since the game does not take the task of “Landspotting” seriously. However, it has to be noted that the users of Geo-Wiki were also incorrect at times (6.5%).

4.3 Land-Cover Agreement

In order to determine how land cover agreement and the rating correspond to each other if the game and Geo-Wiki are correct or wrong, we have calculated the percentages the game and Geo-Wiki land cover agree with each other for each category. In figures 2, 3 and 4 the percentages of how much the land covers of the game and Geo-Wiki agree (0%-100%) are depicted for those cases where both are rated by the validators as equally correct, Geo-Wiki correct and the game wrong or the game correct and Geo-Wiki wrong.

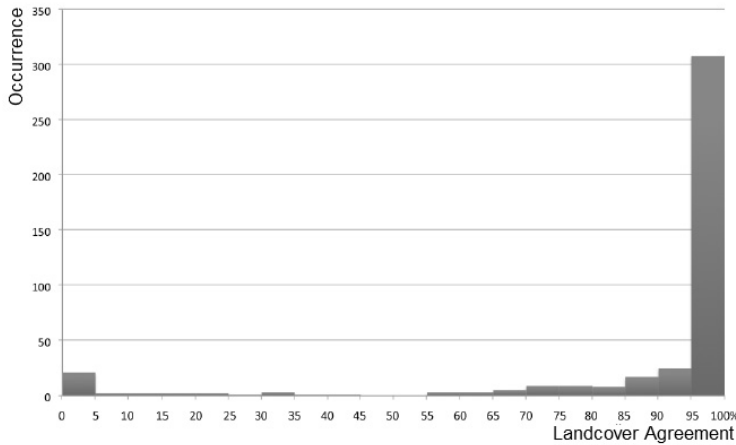


Fig. 2: Land-cover agreement between the game and Geo-Wiki for both equally correct

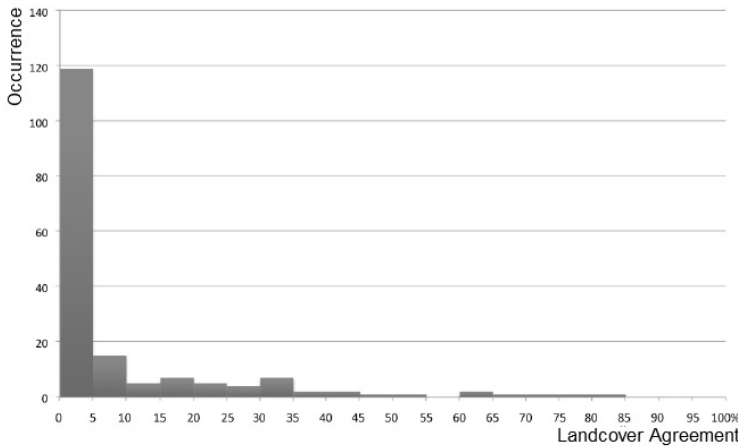


Fig. 3: Land-cover agreement between the game and Geo-Wiki for Geo-Wiki correct and game wrong

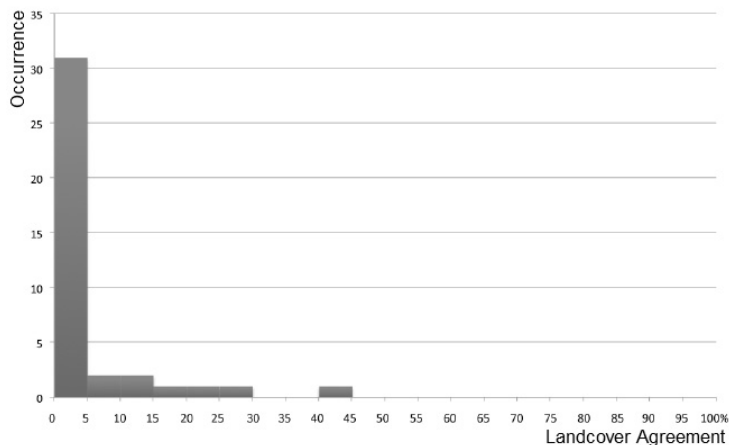


Fig. 4: Land-cover agreement between the game and Geo-Wiki when Geo-Wiki is wrong and game correct

Table 3: Average land cover agreement between the game and Geo-Wiki

Category	Land Cover Agreement
All	47%
Both Wrong	26%
Game wrong, Geo-Wiki correct	8%
Game correct, Geo-Wiki wrong	4%
Both equally correct	89%
Both correct, Geo-Wiki better	61%
Both correct, Game better	60%
Don't know	16%

Table 3 shows the average land cover agreement for the categories. As can be seen, the land cover agreement is the best (89%) when both the game and Geo-Wiki are rated equally correct. It is the lowest when only the game or Geo-Wiki is correct (4-8%). This also makes sense because the two obviously disagree but one of them is correct. The agreement is also low (16%) when it is impossible to judge which land cover is better, which is also a reasonable finding. It can thus be concluded that agreement can serve as a very important clue as to whether the data are correct or not and that therefore humans can be a useful tool for improving information on global land cover.

4.4 Overlapping Land Cover

In order to find out how difficult it is for humans to distinguish between different land cover types, we have calculated a land cover overlap matrix which shows the percentage of overlap between land covers identified in the game and in Geo-Wiki (table 4). This matrix has been calculated from a subset of the data in which both the game and Geo-Wiki are rated by the two validators as equally correct.

Table 4: Overlapping of land cover between the game and Geo-Wiki

	Grass	Trees	Shrubs	Culti.	Barren	Ice	Stone	Water	Urban	Unk.
Grass	70%	14%	4%	3%	6%	0%	0%	2%	1%	0%
Trees	2%	90%	5%	1%	0%	0%	0%	1%	1%	0%
Shrubs	4%	19%	66%	2%	7%	0%	0%	2%	0%	0%
Cultivated	1%	1%	0%	96%	0%	0%	0%	0%	2%	0%
Barren	2%	0%	3%	0%	95%	0%	0%	0%	0%	0%
Ice	0%	0%	0%	0%	0%	95%	5%	0%	0%	0%
Stone	0%	0%	1%	0%	71%	29%	0%	0%	0%	0%
Water	1%	2%	3%	1%	1%	0%	0%	92%	0%	0%
Urban	2%	3%	0%	7%	0%	2%	0%	1%	85%	0%
Unknown	35%	2%	63%	0%	0%	0%	0%	0%	0%	0%

The land cover “Stone” and “Unknown” are only available in the game and not in Geo-Wiki. It can be seen that the categories “Cultivated”, “Barren”, “Ice” and “Water” have huge overlaps (96%, 95%, 95%, 92%) with themselves which means that they are very easily distinguishable. It was expected that “Water”, “Barren” and “Ice” were easily recognisable because they have clearly visible properties. However, it is quite surprising that players are also that good at distinguishing cultivated land. The smallest overlap with themselves have the categories “Grass”, “Shrubs” and “Stone” which means that these types are hard to distinguish. “Stone” for example has the most overlap with “Barren” and “Ice”. It is very difficult for the players to distinguish “Barren” from “Stone” because they are ambiguous, “Stone” can be “Barren” and vice versa. It seems to be most challenging for humans to distinguish between the categories “Shrubs”, “Grass” and “Trees”. It is very difficult to judge based on satellite images if the land cover is “Shrubs” or “Trees” because of the missing height information. Some satellite images are of a coarser resolution so it is quite difficult to decide if it is “Grass” or “Trees”.

5 Conclusions

An iPad game has been presented for improving global land cover. In the first 2 weeks the game was downloaded 351 times and 1371 land cover drawings (2742 km²) were collected from players across 196 games. The data collected by the game were directly compared against results collected using Geo-Wiki, which were also rated by two independent validators. The results showed that 74% of the data collected by the game were correct whereas in Geo-Wiki, correct answers were provided in 93% of the cases. This means that

although Geo-Wiki performs better than the game, the collected data can be combined to create much more reliable information because, as has been shown, agreement can serve as a good hint as to whether the data are correct or not. Furthermore it has been shown that some land covers are more easily differentiable than others for example “Water”, “Barren” and “Ice” are easily recognisable, whereas “Shrubs”, “Grass” and “Tree” can be quite hard to distinguish.

References

- AHN, L. & BABBISH, L. (2004), Labeling images with a computer game. Proceedings of ACM SIGCHI, 319-326.
- AHN, L., KEDIA, M. & BLUM, M. (2007), Verbosity: A game for collecting common-sense knowledge. Proceedings of ACM SIGCHI, 75-78.
- CHIEN-JU, H., TAO-HSUAN, C. & JONG-CHUAN, L. (June 2010), Kisskissban: A competitive human computation game for image annotation. ACM SIGKDD Newsletter, 12, 21-24.
- COOPER, S., TREUILLE, A., BARBERO, J. & POPOVIC, Z. (2010), The challenges of designing scientific discovery games. FDG, 40-47.
- FRITZ, S. & SEE, L. (2006), A method to compare and improve land cover datasets: Application to the glc-2000 and modis land cover products. IEEE Transactions on Geoscience and Remote Sensing, (July 2006), 1740-1746.
- FRITZ, S., SEE, L., MCCALLUM, I., SCHILL, C., OBERSTEINER, M. & VELDE, M. (2011), Highlighting continued uncertainty in global land cover maps. Environmental Research Letters, 6:4:044005.
- FRITZ, S., MCCALLUM, I. & PERGER, C. (2012), Geo-Wiki: An online platform for improving global land cover. Environmental Modelling & Software, 31, 110-123.
- KHATIB, F., COOPER, S. & BAKER, D. (2011), Algorithm discovery by protein folding game players, Proceedings of National Academy of Sciences USA.
- KRAUSE, M., TAKHTAMYSHEVA, A., WITTSTOCK, M. & MALAKA, R. (2010), Frontiers of a paradigm: exploring human computation with digital games. Proceedings of the ACM SIGKDD Workshop on Human Computation, 22-25.
- LAW, E., AHN, L., DANNENBERG, R. & CRAWFORD, M. (2007), TagATune: A game for music and sound annotation. Proceedings of the Eighth International Conference on Music Information Retrieval, 361-364.
- LINTOTT, C., SCHAWINSKI, K. & VANDENBERG, J. (September 2008), Galaxy Zoo: morphologies derived from visual inspection of galaxies from the sloan digital sky survey. Monthly notices of the Royal Astronomical Society, 389, 1179-1189.
- ONLINE PUBLISHERS ASSOCIATION (2012), A portrait of today’s tablet user wave.
- POPCAP GAMES, Plants vs. Zombies. <http://www.popcap.com>.
- SMEDDINCK, J. & KRAUSE, M. (September 2011), Human computation games: A survey. Proceedings of the 19th European Signal Processing Conference.
- ZYNGA, Draw Something. <http://www.omgpop.com>.