

Cropland Capture: An Effective Game for Collecting VGI

¹Tobias Sturn, ¹Christoph Perger, ¹Linda See, ¹Ian McCallum, ¹Carl Salk, ²Michael Wimmer and ¹Steffen Fritz
¹International Institute for Applied Systems Analysis ²TU Vienna, Computer Graphics Institute

1 Introduction

There are a number of global satellite-derived land cover products available, e.g. MODIS [9], GlobCover [7], and the GLC-2000 [8]. These products are important for monitoring, assessment and modeling purposes, yet when compared with one another, they show huge spatial disagreements. To improve these land-cover products, the Geo-Wiki application [1] has been developed for crowd-sourcing land cover using high resolution satellite imagery. Through a number of different crowdsourcing campaigns, more than 250,000 land cover validations have been collected, which have been used for both the development of hybrid land cover products and for validating existing maps. However, it has been quite challenging to gather data through these campaigns, where incentives have been small prizes and co-authorship on scientific publications.

In order to attract more people to take part in crowdsourcing tasks, serious games are often used. These so-called “GWAP” (games with a purpose) have already helped science in impressive ways. One of the most famous examples is FoldIt [2, 3], a puzzle-like game in which the players fold proteins. Within ten days, the FoldIt players have, for example, accurately determined the crystal structure of M-PMV, an AIDS-like virus infecting apes, which had been an unsolved problem by scientists in spite of the 15 years effort. Just recently, with the help of the 230,000 FoldIt players, a new algorithm for protein folding has been developed which outperforms previously published methods [4].

Therefore, as a way of increasing crowd participation in Geo-Wiki, we have developed a number of serious games [5, 6] to improve global land cover maps. In this paper we present our latest game called Cropland Capture, a simple-cross platform application for collecting volunteered geographic information (VGI) related to cropland. We first give an overview of the game and some game design choices. We then present some results from the game and provide a short conclusion.

2 Cropland Capture

Cropland Capture is a cross platform game played in a browser or on mobile devices like the iPhone, iPad and/or Android devices. It can be downloaded from the AppStore or GooglePlay Store or can be played in a browser from this location:
<http://www.geo-wiki.org/games/croplandcapture/>

In the game, players are presented with an image (either satellite images or ground-based pictures) and are asked to whether they see any cropland. They can answer “yes”, “no” or “maybe” if they are unsure. On mobile devices the players can swipe the image very easily to the correct category, located on different sides of the screen. In the browser version the players can click on the categories or use the cursor keys to play more efficiently.

The player scores one point for each correct answer and loses one point for each wrong answer. Answers of ‘maybe’ result in neither gain nor loss of points. To determine the correct answer, we used output agreement. This means that the more tplayers agree on an answer, the more we assume their answer to be correct. We started with a small pool of images validated by

experts. In order to automatically grow the pool of images, 10% of the images players got have not been classified at all. For those images we assumed that the answer from the first player was correct.

We implemented several quality assurance features to ensure that players who answer randomly cannot influence the results. Only players who correctly classify the images will receive unclassified images. Moreover, only the answers from these players will then be uploaded to the database and be used in the output agreement calculation. Another quality assurance feature implemented is that the ratio of cropland to non-cropland images that each player receives is roughly 50%. This ensures that players who always choose the same answer consistently will not progress and their score will stay the same. To further improve data quality, players can challenge an answer if they think they are correct but have been penalized for an incorrect answer. These images are then sent to an expert for analysis. If the players are actually deemed to be correct, then they will gain 5 points but if they are wrong, they will lose 3 points.

A central design goal of Cropland Capture was efficiency. The browser version of Cropland Capture, for example, runs 4 background threads with each loading the images so there is no waiting time for the players. In the browser version players can use the arrow keys to validate very quickly. In the mobile versions the players swipe the images to the correct category, which is also very efficient. Once enough information has been collected about an image we then take the images out of the game in order to avoid redundant classifications that do not provide new information about the presence of cropland.

We have provided intrinsic and extrinsic incentives for the players to play Cropland Capture. As an intrinsic motivator we continually highlight how much of the Earth they have helped us to classify and that they are helping us with scientific research. This incentive is intended to make players feel good while playing the game. As extrinsic motivators we have added a leaderboard and we have offered two forms of prizes. Each week during the last five weeks of the 25 week competition, one answer was randomly chosen and the player who submitted the answer was awarded a prize such as a fitness monitor or a compass. Therefore, each additional classification increased a player’s chance of winning the weekly prize. The second set of prizes was awarded at the end by a draw. In order to qualify for the draw, the top 3 players in each week were automatically entered. Scores were reset at the beginning of each week so players would have a new chance to be entered into this final draw. After the Cropland Capture competition was finished we randomly picked three winners from these top weekly winners who then became our overall winners and were awarded prizes such as smart phones and tablets.

Cropland Capture was launched on November 15th 2013 via a media campaign, with press releases, blogs and a twitter account set up for the game. From there, the game was picked up on a blog by National Public Radio, an article in the Guardian and reported with interviews with Geo-Wiki staff on German radio. The Geo-Wiki network was also contacted by email with a new monthly newsletter to provide regular updates on the game. The game ran for 25 weeks and finished on May 9th 2014. Some results of the game are summarized in the next section.

3 Results

Cropland Capture was played by 3,014 players who together provided 4,567,211 classifications of 187,673 unique images. Of these, 98,411 were satellite images with scopes ranging from 250 m to 1 km² and 89,232 were landscape pictures.

We collected information on the type of device used to play the game. The players made 578,331 ratings from an iPhone5, 616,537 from other iPhones, 698,762 from iPads and 1,636,627 from the browser version and 1,036,853 from Android devices. This was quite surprising as we thought that the browser version would be by far the most successful platform. As it turned out, the iOS version collected 1,893,630 classifications, which is more than the browser version. This is surprising as we did not get featured on the Apple App Store which would explain this result. It is also surprising that the iOS version collected nearly twice as many ratings as the Android version although the Android phone market is much bigger with nearly 80% market penetration.

The average amount of time per image classified was about 1500 milliseconds. Interestingly, we found a strong negative correlation between the time needed for making a classification and the output agreement. Ratings faster than 1500 ms had a higher certainty than ratings slower than 1500 ms. This certainty drops further in the 3500-5000 ms range where the average certainty is lowest. Interestingly, for ratings slower than 5000 ms the certainty rises slightly. This means that speed is an important indicator of players' performance and should be taken into account in future games in order to more efficiently remove images with high certainty.

The speed analysis also shows that the browser-based ratings are faster than other platforms. Also quite interesting is that all classifications faster than 500 ms were done in the browser version and all of those were only images classified as non-cropland. This shows that it is much faster to say that there is no cropland in an image than the opposite. Also interesting is that a big portion of these images were ground-based pictures, which is contrary to the expectation that landscape pictures are harder to judge than satellite images. No significant time differences could be found for the different zoom levels at which the data were collected, which means that the players needed the same amount of time to validate, for example, a 250 m² or a 1 km² pixel on a satellite image.

Figure 1 shows the number of images in different categories of player agreement. The majority of images have a > 90% agreement, indicating that it was easy to identify the presence or absence of cropland. Interestingly there is also a small peak where no agreement can be reached, which means that there are also some images which were very difficult to classify.

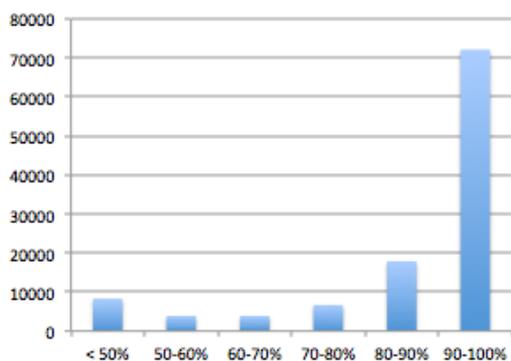


Figure 1: Agreement among players. Only images with more than 10 answers were included.

In order to determine whether landscape or satellite images are easier to classify, we have calculated the player agreement on positions where we have both satellite images and landscape pictures and compared them to each other. The results show that the landscape pictures have an average agreement of 93% while the satellite images had 95% agreement. This means that satellite images are slightly easier to classify, although not significantly so. Analyzing the agreement for different satellite image zoom levels (250 m = avg. agreement 95.148%, 500 m = 95.248%, 1 km² = 95.57%), it can also be concluded that there are no meaningful differences between them.

4 Conclusion

We have presented Cropland Capture, a simple cross platform game for efficiently collecting VGI. The game has been able to collect a huge amount of data from a relatively small number of players. The agreement among the players makes it clear that people are capable of classifying images in terms of land cover. We have presented some first results and have shown some potential ways in which the effectiveness of the game can be further improved, for example, by also taking into account the time needed for making a classification.

We will now use the data collected for improving our current global hybrid cropland map which integrates many existing cropland products. The data will also be used to improve cropland extent estimates in countries where maps have known problems and in further validation of this and other products.

Given the success of this game, we plan to use this type of approach for gathering information on other land cover types in the future.

References

- [1] Fritz, S., I. McCallum, C. Schill et al.: Geo-Wiki.Org: The use of crowdsourcing to improve global land cover, *Remote Sensing*. 1(3):345-354, 2009.
- [2] S. Cooper, F. Khatib, A. Treuille, Z. Popovic, et al. Predicting protein structures with a multiplayer online game. *Nature*, 466:756-760, 2010.
- [3] S. Cooper, A. Treuille, J. Barbero, Z. Popovic, et al. The challenge of designing scientific discovery games. *FDG*, pages 40-47, 2010.
- [4] F. Khatib, S. Cooper, D. Baker, and others. Algorithm discovery by protein folding game players. *Proceedings of National Academy of Sciences USA*, 2011.
- [5] T. Sturn, M. Wimmer, P. Purgathofer, S. Fritz. Landspotting - Games for Improving Global Land Cover. *FDG 2013*, Chania, Greece, May 2013.
- [6] Sturn, T., D. Pangerl, L. See et al. Landspotting: A serious iPad game for improving global land cover. *GI-Forum*, Salzburg, Austria, July 2 2013.
- [7] Defourny, P., C. Vancustem, P. Bicheron et al. GLOBCOVER: A 300m global land cover product for 2005 using ENVISAT MERIS time series. *Proceedings of the ISPRS Commission VII Mid-Term Symposium: Remote Sensing: from Pixels to Processes*. Enschede NL, 2006.
- [8] Fritz, S., E. Bartholomé, A. Belward et al. *Harmonisation, mosaicing and production of the Global Land Cover 2000 database (Beta Version)*. Office for Official Publications of the European Communities, Luxembourg, 2003.
- [9] Friedl, M.A., D. Sulla-Menashe, B. Tan et al. MODIS Collection 5 global land cover: Algorithm refinements and characterization of new datasets, *Remote Sensing of Environment*, 114(1): 168-182, 2010.