

Mutually Reinforcing Systems

John Ferguson
Department of Computing
Science
University of Glasgow, UK
john@dcs.gla.ac.uk

Marek Bell
Department of Computing
Science
University of Glasgow, UK
marek@dcs.gla.ac.uk

Matthew Chalmers
Department of Computing
Science
University of Glasgow, UK
matthew@dcs.gla.ac.uk

ABSTRACT

This paper demonstrates strategies for designing mobile games with by-products in order to allow the acquisition of specific data. A mobile game with by-products called *EyeSpy* and a photo website called *Realise* will be used as examples to demonstrate these strategies. The *Realise* website allows users to browse geographically tagged photos and make specific requests for new ones. In the *EyeSpy* game, players use mobile phones to tag geographic positions with photos and text. *EyeSpy* players can earn points from validating each others' tags by visiting tag locations and attempting to 'confirm' them. If players go to the correct location, both the player confirming the tag and the player who created it will gain points. This creates game content for *EyeSpy* and provides more refined results for the *Realise* website. In this way, both the systems mutually reinforce each other.

Categories and Subject Descriptors

H.5.m [Information Interfaces and Presentation]: Miscellaneous

Keywords

Human computation, mutually reinforcing systems, games with by-products, mobile multiplayer games, mobile photography.

1. INTRODUCTION

It has been demonstrated that 'human computation' games (such as the ESP Game [3]) can be effective in producing by-products as a side effect of play. These games often produce generalized by-products and are not designed for the benefit of a specific system. While this approach may maximize the potential usefulness of the by-products, there will be occasions when specific data is required that may not have been produced. This approach of creating general by-products often leads users to come up with the same solutions for similar problems. Players may see generalities in

the game play and feel they have solved the problem before. This can lead to a feeling of monotony. The similarity in the designs of many games with by-products means that a successful method of making specific requests for data could be adapted to work on many games.

General improvement could be achieved by designing systems which work more closely with each other. A system that is good at producing one kind of data can share that strength with other systems, or two systems that produce the same type of by-product can be used to form a larger dataset. If a system is bad at gathering one kind of data, it makes sense that it should be able to request that data from another system which is better suited to the task.

Mutually reinforcing systems have the ability to receive and act on requests and the ability to make requests of each other. By working together to improve the efficiency and quality of by-product creation, they would each do a better job of achieving their goals. The loop that they would form, where each system can make requests of the other, as well as responding to requests and sharing data, is a strong element of mutually reinforcing systems. This loop is the reinforcing element that allows the improvement of function.

In order to explore requesting specific data and making the games more diverse and enjoyable for the players, we developed two systems that can be described as mutually reinforcing. This means that each system provides a benefit to the other such that each of them better achieves its goals.

We have augmented a game called *EyeSpy* [1] so that it can deal with specific requests for data from a second system, called *Realise*. In *EyeSpy*, users created text and photographs, described as 'targets,' which were associated with geographic locations. The geo-tagged photographs would then have text associated with them (the text attached to the same location). Other users would then have to visit these locations to confirm the targets. However, they had no means of knowing where the targets were, other than the photographs or texts themselves. This design encourages people to produce targets with enough detail to be located easily. *Realise* is a website that allows users to browse the geo-tagged photographs produced by *EyeSpy*, see their location on a map and search them based on nearby text targets. It also allows users to make requests for new photographs by describing the photograph they want to the *EyeSpy* players. These appear in the game as a special target type that requires the players to create a photograph that matches the description they have been given.

By allowing users to request specific photographs from the *EyeSpy* players, *Realise*'s data set is broadened and its

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.

KDD-HCOMP'10 July 25, 2010, Washington, DC, USA.
Copyright 2010 ACM 978-1-4503-0222-7 ...\$10.00.

usefulness is increased. It does not have to rely solely on the players' decisions about what should be in the data set and can allow for any holes in the data set to be filled. Also, by introducing an element into the game that is not defined by the players themselves, the game will have more diversity and become less monotonous for the players.

2. DESIGN

The rules in EyeSpy encourage photographs that are created specifically for use by other people because the tags must be confirmed. In a sense, the photos are created specifically so that they will be easy for others to find.

The EyeSpy game was created such that players would create a set of photographs that were geo-tagged and had text associated with them. However, the text might not necessarily describe the photographs themselves, but the area in which they were taken. This would mean that searching the set of images on the Realise website for a specific term would return not only images that the term described, but also images from the surrounding area. It also means that a new photograph taken in an area which had previously been tagged with text would immediately be labeled with this information.

The Realise website allows users to make requests for photographs when the existing dataset does not satisfy their needs.

This most likely happens when users search for an image and do not get the results they want. They can then request that better results are produced. These requests can be very specific, but they can also rely on local knowledge or human decision making, such as, 'a tree,' or, 'the nearest subway to Glasgow University.' In this sense, the website is no longer just an image search engine with the option to request that 'holes' in the dataset be filled. An entirely new function has been added where labeling of images is descriptive in a subjective way. This ability would not have been possible without being part of a mutually reinforcing system loop. This is generally avoided in games with by-products because objective data is considered more valuable (leading to agreement rules between randomly paired players in order to create such objectivity).

While the Realise user may be happy with a photo because it fulfilled their particular needs, that does not mean it is labeled in a useful way for other users. The existing game mechanic can be helpful in this instance. Since the photo that was taken for the priority target had location information, it can still be used with the existing text targets for that location. However, we still need to know if the image is generic enough to be part of a general search for that area. This can be achieved by putting the image back into the game as a regular photo target. This also means that players who make priority targets will have the opportunity to get points from the photo a second time when it becomes a photo target. If enough players confirm the photo as a photo target, it means it is recognizable enough to be used as a search result on the Realise website.

Varied game experience

EyeSpy was designed to be a self contained system. An entirely new game mode needed to be added for it to function as part of a mutually reinforcing system loop. A new type of target was added to the game. This 'priority' target had different rules from the existing target types and was supplied

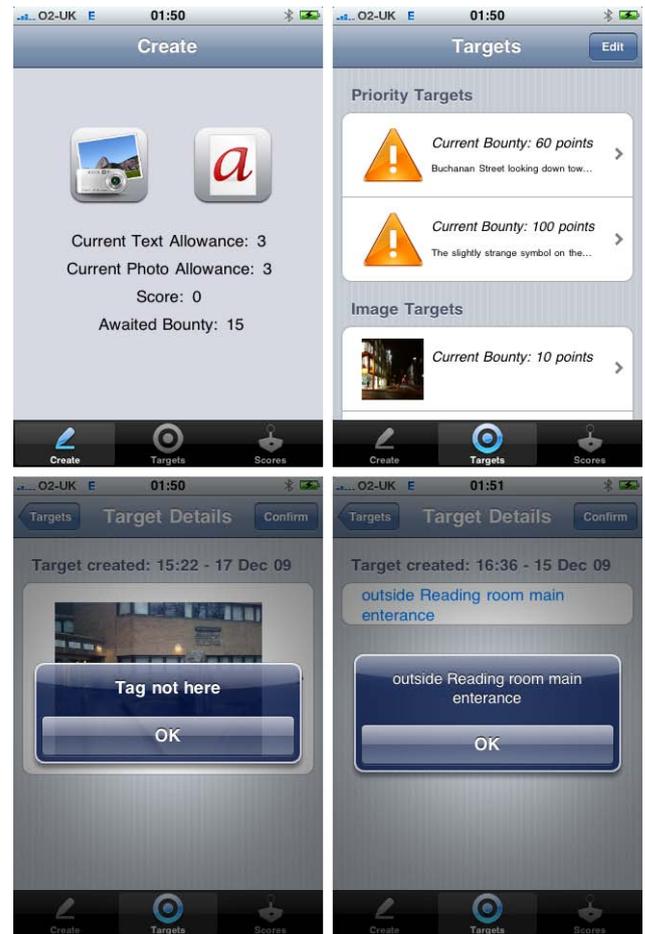


Figure 1: Screen shots from EyeSpy showing, clockwise from top left: the create screen (where players make image and text targets), the targets screen (where players browse the targets to be confirmed), a text target screen (where users can confirm a text target and see its full text), and an image target screen (where users can confirm an image target) showing an unsuccessful confirmation attempt.

by the website users rather than the game players. When the website users make a request, it comes into the game as a priority target. This is a textual description of the kind of photo for which the website user wants a picture. Unlike the confirm procedure for the other target types (where the player is checking to see if they are in the same place the target was created), the priority targets require the players to create a photo that matches the priority target's description. Since the Realise website supplies no location information with a request, priority targets can be confirmed in any location. The photo is then sent back to the Realise user who requested it to see if it is what they wanted. The EyeSpy player will not receive any points until the Realise user has accepted the photo as being what they requested.

Allowing the EyeSpy players to fulfill the requests of the website users creates a different game dynamic than when they are creating and confirming targets amongst themselves. The website users make requests for photographs that are not already in the system. This means there is a

greater likelihood that the requests will be for things that the players would not pick on their own, which makes the game more varied. Additionally, the website users may make requests which call upon the EyeSpy players to give subjective opinions which are not supported by the existing game dynamic.

Equilibrium between systems

The website users create bounties for their requests. Each day that a Realise user logs into the website, they are given points which are added to a total. These points can be assigned to requests for photos. EyeSpy players will receive these points if they successfully create the required photo.

The number of points that a Realise website user is assigned is based on the overall system activity for the previous day. This means that if a smaller percentage of website requests were dealt with by the game players, the number of points that the website users are assigned will be less. This means that users will want the same number of photos, but will have fewer points to spend. This measure is used to encourage people to request fewer photos until the EyeSpy players can catch up with the requests already in the system. In assigning requests to players, the system will also favor requests that have more points assigned to them. This adds more value to the points as it indicates that the more that Realise users assign to a request, the more quickly they are likely to receive a response. This should encourage more points being used on a smaller number of requests.

The system also gives each Realise user an allowance of requests that can be made each day. This allowance does not accumulate and is reset every day. This allowance is also adjusted automatically based on the overall activity in the system. If there are greater numbers of requests awaiting photos, there will be fewer requests that are allowed to be made. In some instances where there are too many requests awaiting photos, this may mean that a Realise website user is not allowed to make any requests on a given day.

These measures mean that each Realise user has a fair chance of getting photos for their requests. This method of valuing the work between systems in a mutually reinforcing system loop is important. The systems should form a partnership without one becoming too dominant, or else the equilibrium will collapse and the benefits of mutually reinforcing systems might be lost.

Using multiple systems for validation

If a requested photograph is satisfactory, the Realise user will pay the bounty (giving the EyeSpy player the points). This allows the Realise user to access the full sized photograph (they must make their determination of whether a photograph is satisfactory based on a smaller thumbnail version). This has the double effect of validating that the photograph matches the request text (which can now be used as a label for it) and giving the player and website user what they wanted (respectively, points and a full sized photo).

The photograph itself will return to the game as a regular photo target so that it can be confirmed by other players as well before it is added to the usable data set.

The website users and game players can also monitor each other for inappropriate behavior. The website users can report inappropriate photographs that appear on the website, and the game players can report inappropriate requests made of them. Any person who has enough targets or re-

quests reported can then be banned from using the systems. This banning is not automatic. Once the threshold has been reached, an administrator is alerted to review the activity and determine if a ban is warranted. This prevents game players from colluding to try and get another player banned by flooding the system with reports of inappropriate behavior when there is none. It is, however, unlikely that this would happen because only a target or request can be flagged, not a person. Since the person who creates a target is kept anonymous, it would be hard for other players to target anyone specifically. Admittedly, banning people will not prevent them from creating a new account and repeating their abusive behavior, but they will have to start with no points to spend on the website, or at the bottom of the leader board in the game. More extreme methods of identification could be included to make people accountable for their actions (such as sending an authentication code in the post to someone's home address, making it difficult to register multiple times). However, unless a serious problem of abusive behavior is demonstrated in the system, this would merely detract from the experience of the users.

In addition to getting targets which are more interesting to confirm, using the second system to validate requests adds a new scoring system to EyeSpy that would not have been possible without mutually reinforcing systems. The new priority targets act like a bonus in the game. The number of priority targets received and how many points they are worth will appear to be random to the EyeSpy players, making the game play less monotonous. Additionally, the requests themselves will not always be the same objective type as in the rest of the game. They may now include non-location based requests for specific objects, requests which have potentially more than one correct answer and requests which may require the player to give an opinion.

3. IMPLEMENTATION

The systems each have a client and a database. The EyeSpy game is currently implemented as a phone-based game using the phone's GPS hardware and camera to create geo-tagged photographs and text. The EyeSpy database contains all these geo-tagged photographs and text so that other systems can access the dataset through a web service.

The Realise website accesses the EyeSpy web service so that it can display the photographs to the website users. It also contains its own database storing user and request information. In addition to this, the EyeSpy web service allows the website to add requests to the EyeSpy game.

The Realise website was implemented as a search engine. When a user searches using a keyword, the website sends a request to the EyeSpy web service telling it the keyword and how the results should be ordered. The EyeSpy web service then checks to see if that keyword exists in any of the text labels in the dataset. For every matching label that is found, the EyeSpy web service returns images that were taken in the same area as the matching label.

In order to allow the Realise website users to see where a photo was taken, the Google Maps API was used to create a pin showing the location. The main search results page for Realise can be seen in figure 2.

4. TESTING THE SYSTEMS

Three trials were carried out to test the augmented version

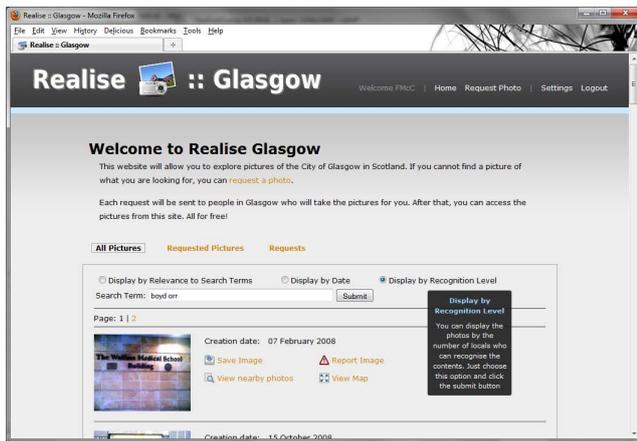


Figure 2: The Realise website showing results for a user entered search. For each result, users can see the full sized image, save the image, report the image, see the image’s location on a map or see photos that are near it. The results can be ordered by relevance, date or recognizability (the number of times the image was confirmed by EyeSpy players).

of EyeSpy and the Realise website, however, the addition of a bounty system was not included in the first trial. Initially, all priority targets were worth the same number of points, but the Realise website users were still required to confirm them. However, the Realise users quickly flooded the system with requests and confirmed almost all the images that were submitted whether they were accurate or not. This led to the inclusion of the bounty system in the second and third trials. In these latter trials, the problem of flooding was reduced and photos were only confirmed when they represented what the Realise users actually wanted.

The three trials consisted of a total of 46 people (13 female and 33 male): the first trial consisted of 10 EyeSpy players and 5 Realise users, the second trial consisted of 18 users split into two groups of 9 (one group of users who had lived in the game area for more than three years and one group of users for less than three years), and the third trial consisted of 13 players split into two groups of 7 people and 6 people (7 users who had lived in the game area for more than three years and 6 for less than three months). This selection process gives a broad range of users with varying levels of local knowledge. The first trial lasted two weeks and the users maintained their roles for the whole trial. In the second and third trials, the two groups each took on the role of EyeSpy player or Realise user for half the trial (11 days), then the two groups swapped for the second half. The second and third trials lasted three weeks each. The EyeSpy game was not seeded with targets in any of the trials, but the Realise website was seeded with all the data that had been gathered from the previous trials, as well as the data gathered from earlier EyeSpy trials [1]. In each trial, the participants got paid £10 a week with winners of the game being paid double. In the first trial, the participants came from a mixture of professions, but most were University students studying a mixture of subjects. In the second and third trials, the participants were all University students with most coming from Computing Science with only four studying non-computing

subjects. In total, there were five groups of people in the trials who knew each other well: three groups of two users, and two groups of three users. This means that 34 of the trial participants (74%) did not know the other participants.

The trials produced 427 images, 290 text labels and 212 requests. 123 of the images (29%) were produced as the result of requests made to EyeSpy from the Realise website. Out of a total of 717 images and text labels, 60% were images and 40% were text. On average, users produced 16 targets during the game: 9 images and 7 text labels. However, on average, about a third of those images were produced to fulfill the direct needs of a Realise website user. This strong level of directed effort is a major benefit of mutually reinforcing systems as it ensures that the by-products produced are those which are needed.

5. DISCUSSION

While the Realise website is currently limited in ways for the user to interact, it shows the potential in creating a photograph website where users can request new photographs. The website could be used to gather local information about a place before you visit it (such as asking for a picture of a hotel so that you know it when you see it), or to get a photograph of a place you visited before when you forgot to bring your camera. This idea of pre-visiting and post-visiting locations is an existing concept in e-tourism [2].

Adding EyeSpy to a mutually reinforcing system loop has resolved some of the problems that the game suffered from: the by-products produced are now specifically requested by the system that uses them (so they are closer to what is required and less broad), the game is less monotonous because of the randomness and challenge of the new priority targets (these are like bonus targets because the players cannot guarantee that they will always have them during play), and players feel more competitive because the priority targets are worth a variable number of points depending on what the Realise website user assigned to the request. This means that the game encourages finer grained results where required and is now more enjoyable to play because it is less monotonous. By working together in a mutually reinforcing system loop, Realise can now request specific data from EyeSpy, and EyeSpy becomes more diverse and enjoyable to play.

6. ACKNOWLEDGMENTS

This research was funded by UK EPSRC (GR/N15986/01, EP/E04848X/1, EP/F035586/1).

7. REFERENCES

- [1] M. Bell, S. Reeves, B. Brown, S. Sherwood, D. MacMillan, J. Ferguson, and M. Chalmers. Eyespy: Supporting navigation through play. In *CHI 2009*, April 2009.
- [2] B. Brown and M. Chalmers. Tourism and mobile technology. In *ECSCW 2003: Proceedings of the eighth european conference on computer supported cooperative work*, pages 335–355, Helsinki, Finland, 2003. Dordrecht: Kluwer Academic Press.
- [3] L. von Ahn and L. Dabbish. Labeling images with a computer game. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 319–326, New York, NY, USA, 2004. ACM Press.