Research Repository UCD

Title	The social camera : a case-study in contextual image recommendation
Authors(s)	Bourke, Steven, McCarthy, Kevin, Smyth, Barry
Publication date	2011-02-16
Publication information	Bourke, Steven, Kevin McCarthy, and Barry Smyth. "The Social Camera : A Case-Study in Contextual Image Recommendation." ACM, 2011.
Conference details	Paper presented at the 2011 International Conference on Intelligent User Interfaces (IUI2011), February 13-16, 2011, Palo Alto, CA
Publisher	ACM
Item record/more information	http://hdl.handle.net/10197/2951
Publisher's statement	This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in Proceedings of the 16th international conference on Intelligent user interfaces available at http://doi.acm.org/10.1145/1943403.1943408
Publisher's version (DOI)	10.1145/1943403.1943408

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

The Social Camera: A Case-Study in Contextual Image Recommendation

Steven Bourke, Kevin McCarthy, Barry Smyth
CLARITY: Centre for Sensor Web Technologies,
School of Computer Science and Informatics,
University College Dublin (UCD), Belfield, Dublin 4, Ireland.
E-mail: firstname.lastname@ucd.ie

ABSTRACT

The digital camera revolution has changed the world of photography and now most people have access to, and even regularly carry, a digital camera. Often these cameras have been designed with simplicity in mind: they harness a variety of sophisticated technologies in order to automatically take care of all manner of complex settings (aperture, shutter speed, flash etc.) for point-and-shoot ease, these assistive features are usually incorporated directly into the cameras interface. However, there is little or no support for the end-user when it comes to helping them to compose or frame a scene. To this end we describe a novel recommendation process which uses a variety of intelligent and assistive interfaces to guide the user in taking relevant compositions given their current location and scene context. This application has been implemented on the Android platform and we describe its core user interaction, recommendation technologies and demonstrate its effectiveness in a number of real-world scenarios. Specifically we report on the results of a live-user trial of the technology in a real-world tourist setting.

ACM Classification Keywords

H.5.1 Artificial, augmented, and virtual realities:
H.5.2 Graphical user Interfaces (GUI)Prototyping

General Terms

Design, Experimentation

INTRODUCTION

The success of digital cameras means that the world of photography has changed forever. But, the first generation of dedicated point-and-shoot digital cameras represent only the beginning of a much broader revolution. Today many of us carry a digital camera with us everywhere we go; they are a common feature of a modern

mobile phone. This has lead to an explosion of photographic content, which has been created and uploaded to a variety of photo-sharing services. In parallel, considerable research effort has been focused on assisting users when it comes to capturing and managing images. For example, in addition to the auto-exposure setting features of most modern cameras, new advances in face recognition are now being used to help users to improve portrait style photography by auto-focusing on faces in a scene [4].

Recently, cameras have started to become available equipped with location sensing technology and digital compasses and this introduces some interesting new opportunities when it comes to helping novice users to take better quality photos. Indeed, while modern cameras have sophisticated auto-exposure modes (to take care of aperture and shutter speed, flash, etc.) [6] there is little or no support for the end-user when it comes to helping them to compose or frame a scene. Simply put, modern point-and-shoot cameras work well to set the exposure settings that are appropriate to a given scene but they don't help the user when it comes to picking an interesting shot or framing the scene. This is what photographers refer to as the *composition problem* [24] and choosing the right composition is a key ingredient when it comes to taking high-quality photographs.

In this paper, we consider this composition problem as a novel type of recommendation opportunity whereby individual users are prompted through the cameras user interface, as they setup to take a photograph, with nearby examples of relevant, well-composed, previously taken photographs. In other words, we can recommend a short-list of high-quality, well-composed photographs to the user, based on their current location, lighting conditions, etc. in the hope that one of these compositions may usefully guide the photographer with respect to their current photograph. This represents an interesting user experience for a number of reasons, in particular the use of contextual information from the physical world (location, time, lighting, etc.) is related to recent work in *context-aware systems research* [2, 9, 15, 20]. Moreover, this is an opportunity to introduce recommender systems into an existing and ubiquitous consumer technology, namely digital cameras and camera phones, where there is a pre-existing history of so-

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.

IUI'11, February 13–16, 2011, Palo Alto, California, USA.

Copyright 2011 ACM 978-1-4503-0419-1/11/02...\$10.00.

phisticated assistive technologies and thirdly, this work is enabled by recent device advances (e.g. GPS and digital compass technologies as standard in modern camera phones) and online services such as geo-coded image repositories like Panoramio¹, without which this work would not be practically possible.

In this paper then, we describe our initial attempt at developing the *social camera*, which embodies our intelligent interface to assist users with compositional support. The application has been developed on the Android² platform and we describe the user experience, interactions and demonstrate its effectiveness in a number of real-world scenarios. In addition we describe the results of an initial live-user field-trial involving a group of 21 foreign tourists selected at random. In future we would like to perform an evaluation with more users, however due to hardware issues we could not perform this in our initial trial.

BACKGROUND

The work described in this paper brings together ideas from a number of different areas of research including, recommender systems, image retrieval, context-aware systems, mobile computing, computational photography, and the sensor web. These combined techniques provide a powerful user interaction experience which can be used to create a intuitive supportive user interface.

A core element of our social camera system concerns the ability to identify and retrieve relevant images and there has been considerable work to date on image retrieval in general. In particular, content-based image retrieval (CBIR) approaches seek to understand the content of an image by using a variety of image analysis techniques to extract core image features as the basis for matching during retrieval; typically these approaches are necessary when traditional forms of indexing information (such as descriptive text) are unavailable [19]. For example, colour histogram, edge detection, and shape extraction techniques have all been used as the basis for image retrieval; see for example [10, 11, 17]. On their own however these *intrinsic* properties of an image are rarely sufficiently informative to drive an effective image retrieval system and so researchers have begun to look beyond the image towards *extrinsic* forms of information that may be used during retrieval. For example, recently the work of Von Ahn [21] and others have demonstrated how image representations can be greatly enhanced with tags, and how people can be encouraged, and are willing to, provide tags that carry semantically rich information about images. In our work, we are focused on image retrieval in a mobile context and this provides additional sources of informative features that can be used to greatly constrain the retrieval task. For instance, the availability of accurate location information (via GPS sensors) is one obvious source of

context information and this type of information has been harnessed as the basis for tag inferencing during image retrieval [16].

Indeed the availability of rich context information has led to an increasing interest in *context-aware systems*; see [14, 20, 22]. Generally speaking there are different forms of context (as opposed to user preferences) that can be used to guide user interactions within mobile interfaces. Contextual features have long been used by ubiquitous technologies to provide additional support when making computational decisions and supporting user interactions [5]. For example, the early work of Van Setten et al. [20] specifically focused on the role of location information in a mobile recommender systems for tourism applications. In this work they specifically focused on the combination of context-awareness and recommender systems in order to guide tourists around points-of-interest based on their preferences and context. The recommender systems interface would adapt as the user moved through the physical world.

The digital photography field has recently begun to experience the incorporation of contextual features when taking a picture. For instance, context has been used during the picture taking process to add additional effects to pictures as they are taken, the photographer can adapt how these features such as sound and light affect the image in the camera's interface [9]. This helps capture more than the standard visual experience associated with typical digital cameras. A lot of work has been carried out in the post picture taking process to add contextual information to images so that they can be managed and navigated in an easier fashion [13]. Using this technique other users can find more meaningful experiences through using the contextual information during their experience to better associate similar context.

In this paper, we are concerned primarily with a very familiar real-world scenario — helping people to take photographs — and, as we all know, the world of personal photography has changed dramatically over the past ten years, and will continue to do so [7]. We believe that the time is now right for more dynamic and intelligent interfaces to play an important role in a new generation of connected cameras, especially when it comes to providing intelligent assistance to the user. Of course, digital camera manufacturers and users alike are already all too familiar with the important role that sophisticated computational techniques have played when it comes to supporting image capture and processing. A wide range of techniques are now routinely employed by even the most basic of digital cameras when it comes to focus, aperture, speed and other exposure settings. Indeed more recently we have seen a new generation of features, based on on-camera image processing, to further assist the user. For example, sophisticated object and face detection techniques can be implemented in real-time so that prominent objects can be identified

¹www.panoramio.com

²www.android.com

within a scene for improved focus and exposure settings [12]. The dreaded "red-eye" effect can now be removed using a combination of face detection and local colour manipulation [8]. Images that have been blurred by unwanted camera shaking can be repaired to produce sharply focused images [23] and sequences of connected images can be automatically stitched together to form seamless panoramas [3]. The point we wish to make here is that these sophisticated information processing techniques are central to the digital camera revolution and mainstream digital cameras and thus there is already a *platform* that exists which is capable of supporting further innovation in the direction of features to help photographers, from novice to expert, and it is in this context that we propose our social camera system. A system which takes advantage of new location sensing features of modern cameras and the availability of emerging, high-quality, geo-coded image repositories, in order to deliver image suggestions to users as a way to provide framing and composition advice. One area of assistance which has so far been ignored by the camera manufacturers. Previous attempts have been made to computationally calculate photo composition based on attributes such as the relative orientation from which the camera views the subject, subjects may be carefully positioned in the frame to determine visual weight, the distance from the camera and the size of the item in the photo [1, 18]. These approaches use formal verifications to declare what a good composition is, while this approach provides useful insight into creating dynamic, well composed photos. We believe that our social filtering process helps ensure that pictures which are thought of as well composed are offered to the end user in our interface, as opposed to a mathematically well defined image. These approaches present some complications such as good composition being subjective. Also, these approaches were carried out in laboratories and virtual environments, our work is carried out in a real world scenario with random users from the street. This obviously presents differing challenges.

THE SOCIAL CAMERA

To recap, our aim in this work is to develop a interface to assist users when taking pictures that is capable of recommending well-composed photographs to a user, which are relevant to the current location and setting, as a way to help the user take better pictures for themselves. This is achieved by not only using a variety of contextual features but also through the use of social ranking around the pictures quality from panoramio. When designing the user interface for this recommendation technology a number of important considerations must be catered for. Firstly, the entire process should be as intuitive and unobtrusive as possible, taking a picture currently on most camera phones involves a simple point and shoot approach. Secondly, as the work is carried out on a low powered device, we want the entire experience to be as quick as possible. Additionally a number of important ideas should be considered when creating the recommendation aspect of the system: (1)

understanding the user's current context as the basis for a recommendation query; (2) selecting a suitable set of candidate images from an online image repository; (3) ranking these candidates and selecting a short-list for recommendation to the end-user. In this section we describe the form and function of the social camera application, focusing on these three issues in particular.

Architecture

The overall social camera system is divided into 3 main components — the *camera component* or *social camera app*, the *recommendation engine*, and the *image server* — as shown in Figure 1. The camera component is the actual software that runs on the camera. This has been implemented on the Android platform and is responsible for handling the core image capture functions of the camera itself, as well as providing the primary interface between the user, the user's context and the recommendation service. Each time the user points the camera at a scene, the social camera app generates a set of *context features* from the current scene settings. These features include the current time, GPS coordinates, compass direction, lighting conditions, as well as the current camera settings, such as aperture and ISO speed; see Figure 2 for an example of these various context features. In short, these features provide a detailed representation of the current scene context. They represent not just the current location (GPS) but also the direction that the camera is pointing (digital compass) and photos that have similar location and compass features are likely to capture very much the same scene that the user is currently seeing. Features such as camera's ISO, aperture, and exposure time settings are set automatically by the camera device and they capture important information about the lighting conditions that currently exist; images that match in terms of their lighting conditions are therefore likely to be good matches for the current scene, from an exposure viewpoint.

In combination then, these context features provide the basis for image retrieval. In the case of social camera, we rely on a variety of online image repositories, such as Panoramio³ where users have uploaded GPS-tagged photographs, complete with relevant EXIF⁴ (Figure 2) meta-data (ISO, aperture, etc.) and the recommendation engine selects relevant images based on a matching function that compares the current user context to the meta-data stored with the images (Section). This provides a short-list of relevant, high-quality images that can be ranked and presented to the user through the social camera app.

As with any consumer-facing technology, the user experience is a vital success factor and special care and attention has been paid to the development of an simple but powerful user interface for the social camera app. There are three basic parts to the social camera interface: *photo recommendation*, *directional assistance*,

³www.panoramio.com

⁴www.exif.org

Figure 1. System Architecture of the Context-Aware Social Camera

and *framing assistance*. Obviously, the context capture functionality remains invisible to the user and is activated when they point the social camera app at a particular scene. But once a suitable set of recommendations have been located the user is given the option to review these as examples of high-quality images that have been taken nearby; this is the photo recommendation component. If the user chooses an image that they would like then the interface provides on-screen directional assistance to the user to help them to better re-orientate themselves so as to be more closely aligned with the chosen image-scene. Once the user has aligned themselves with the relevant scene then they can receive framing advice, effectively overlaying the chosen image on the current scene as a transparent overlay so that the user can more precisely compose their own photograph. In what follows, we will describe the social camera interface. This includes the location search, photo recommendation, directional assistance, framing assistance and finally a walk through of the recommendation process.

```
<context id="12345" user="abcde">
  <feature name="DateTime"><value>2010:01:17 15:16:21</value></feature>
  <feature name="Lighting"><value>77/10</value></feature>
  <feature name="CompassDirection"><value>15.463</value></feature>
  <feature name="Location">
    <property name="longitude"><value>53-18-26.56N</value></feature>
    <property name="latitude"><value>6-13-21.08W</value></feature>
  </feature>
  <feature name="ISO Speed Ratings"><value>100</value></feature>
  <feature name="ApertureValue"><value>37/10</value></feature>
</context>
```

Figure 2. Example context information.

SOCIAL CAMERA INTERFACE

Figure 3 presents screenshots of the social camera app in action. In this case the user is located near to Tower Bridge in London and in what follows we will summarize a brief walk-through of the assistive technologies in action.

Location Selection

When the social camera app is activated the user is first presented with a location selection screen; Figure 3(a). This screen covers the entire recommendation process during the user session. A user can click and drag to expand the search area which is represented by a coloured circle which grows around the user's current position. Typically the user quickly moves off this screen but it does provide an opportunity to adjust some of the recommendation settings if desired. For example, by default the social camera, as mentioned above, will focus on retrieving images that were taken from positions no more than 50m from their current position. This interface allows the user to easily adjust this default by either extending or contracting the location disc as shown; indeed the user can also use this feature as a way to relocate their current position manually, in order to review photo recommendations from other locations, for example. One beneficial use for this aspect is when the current position the user is in has a low number of high quality photos. Expanding the search can provide better results for them to work from.

Photo Recommendation

Once the user is satisfied with their location, the social camera retrieves a ranked list of images according to the recommendation strategy outlined in the previous section; see Figure 3(b). From an interface standpoint the user can simply cycle through these images until they find one that they like. In this example, let's assume that the user has selected the lower image in Figure 3(b) and wished to take a similar shot. Once the image has been selected the social camera app will adjust the camera's current exposure settings to match those of the current image, allowing for variations in lighting as appropriate; priority is given to aperture in the current system. The selection process involves a colour being associated with each selected image. This is used as a cue for the user to remember positional aspects of the photo in later stages of the picture taking process.

Figure 3. a) Selecting your location. b) Available images from the Photo Recommendation list. c) Viewfinder image with Directional Assistance. d) Increasing the opacity of the recommended image with the Compositional Assistance tool.

Directional Assistance

Having selected the image the user now needs to better position themselves in order to take a similar image. The social camera interface helps the user get into position and as shown in Figure 3(c), a simple form of directional assistance is provided. In the top left corner of the screen, the photographer positioning for the selected images are displayed on a dynamically changing positioning radar. The radar updates as the user moves about and helps the user to get to the correct location. The idea here is for the users to get the colour coded marks which represent each of the selected photo recommendation to the center of the circular radar. Additionally, the user is guided to adjust their position so that the colour-coded recommendation icon presented in the centre of the screen. Is the user's current position is too far off to the left, for example, then the icon will appear on the right-hand side of their screen, and as the adjust their positioning it will move towards the centre of the screen; in Figure 3(c) we can see that the user has adjusted their positioning so as to be aligned just slightly to the left of the retrieved image's position.

Framing Assistance

The directional assistance feature is unlikely to produce a perfect alignment between user and recommended image. There are limits to the accuracy of current GPS location sensing technologies which ultimately mean that the directional assistance feature is useful for some rough positional adjustments. The final feature of the social camera interface is designed to help with the final framing of the user's photograph. Basically, this allows for the current recommended image to be overlaid on the current viewfinder scene as a transparent overlay; the degree of transparency can be adjusted by the user with the on-screen slider.

In this way the user can make some fine-tuned adjustments to their framing and position. In Figure 3(d) for example, we can see that the user has reproduced very closely the recommended image and is now ready to take their own photograph. This is quite advantageous as it also allows external objects such as people or alternative landmarks to be incorporated into a photograph which has already been deemed of high quality composition. Of course at any time the user can just decide to take a photo; they are not compelled to exactly replicate the recommended image.

Recommendation Process

Having described the interface and explained the user interaction with the social camera system we now describe the recommendation algorithm. The summary algorithm is presented in Figure 4. The basic input to the recommender include the context profile (CP) and the number of images to return as recommendations to the end-user (k), typically we return 5 images. The first step of the recommendation engine is to locate a suitable set of images that match in terms of their location and direction properties. This is a key point. There is little benefit to presenting the user with photos, no matter how well composed they are, if the photos bear no resemblance to their current scene and location. Similarly, all things being equal it makes sense to prioritise photos from a given location that have similar directional information. For this reason, during this stage of the recommendation process we retrieve a set of n images (where n is typically 100) such that these images are within 50m of the current location and at a similar time of day; there is little advantage to presenting the user with a night-shot if they are experiencing bright sunshine. Next these images are then scored according to a combination of how close they are to the users current position and the angular difference between their

CB: Context Profile, rs: Result Set, ID: Image Database, EA: ExIfAttributes, SA: Supported ExIf attribute, R: Searchable Range, L: Location

```

1.  Define RecommendImages(CP, k)
2.  rs' ← RequestImagesFromImageProviders (CP, CP.Angle)
3.  rs' ← FindRelevantImagesByTime (RS, CP)
4.  rs' ← ExIfData(RS)
5.  k' ← PopularityRanking(rs,k)
6.  Return k
7.  End

8.  Define RequestImagesFromImageProviders(CP, Angle)
9.  rs' ← ImageDatabase.Query(CP.Location, CP.Radius)
10. Rs.AssignScoreBasedOnImageProviderPopularity
11. For Each (img in RS)
12.   IF(imgWithin Angle)
13.    img.LocationScore=cp.location-
    img.location * 1.4
14.   End
15. End
16. Return RS
17. End

18. Define FindRelevantImagesByTime(RS, CP.Time)
19. For Each (img in RS)
20.   IF(img.Time within (CP.Time + 5hours from
    current time)
21.    timeDifference = CP.Time - img.Time
22.    img.Timescore = (timeDifference / 5) * 1.15
23.   End
24. End
25. Return RS
26. End

27. Define ExIfDataScoring(RS)
28. Count=0
29. For Each (img in RS)
30.   If(img.ExIfAttributes in SA)
31.    Count++
32.   End
33. img.ExIfScore = Count/ SA * 1.25
34. End
35. Return RS
36. End

37. Define PopularityRanking(rs,k)
38. rs' ← rs.retrieveHighestScores(k)
39. Return RS.top(k)
40. End

```

Figure 4. The recommendation algorithm.

direction and the users current direction, as shown in Lines 8-17. Many image repositories allow users to rate images and this information can be used by our recommender to give preference to images that seem to be well liked, on the assumption that such images are likely to be of higher quality; see Line 10. This provides a set of recommendation candidates that are likely to be recognisable within the view of the current user and have been rated as high quality photos by users of the image provider from which the images have been sourced.

Next, we use further scoring functions in order to evaluate the utility of these recommendation candidates. First we score the images based on the time in the meta data (see Lines 18-26). Then we score the images based on how closely their ExIf related settings (Aperture, Light, etc.) match the current user's context features (see Lines 27-36). This will allow us to give a prefer-

ence to photos taken under similar exposure settings; it may be a particularly dull day leading to the need for a longer exposure time or a greater aperture setting, for example.

We now have a set of images that have been taken in the vicinity of the current user, at a similar time of day and these images have been scored according to their precise proximity, exposure settings, and popularity. To produce a final set of k recommendations we take the top rated images which were scored in the previous functions with differing weighting scores.

Discussion

In this section we have briefly described one simple use-case of the social camera app. It might seem that all we are doing here is helping the user to take the same type of photos that have already been taken by others and one could reasonably ask whether this is worthwhile (versus just adding other people's photos to your albums). Our expectations are somewhat different however. We do expect that some people will use social camera in just this way, but rather than just replicate the photos of others who have gone before them, we expect that they will take guidance from the recommender while ultimately taking their own unique picture. For example, some users will wish to reproduce their own image of Tower Bridge but with their own family members in the shot. More generally, however, we hope that this type of interface assistance will serve to gradually improve the ability of users to frame and compose interesting shots and so improve the quality of the photographs they produce in the long-term.

USER STUDY

To evaluate this first version of the social camera we chose to perform a live-user field-trial for foreign tourists in Dublin. The trial took place over a 3 day period (Early 7 to April 9, 2010) and involved randomly selected, individual tourists nearby to four famous Dublin landmarks.

Figure 5. Preferred Camera Type Of Participants

In total 21 tourists participated in the trial: 8 female and 13 male; ranging in age from their late teens to their forties; approximately half of the participants in their

30s. The participants came from 8 different countries including Germany, Sweden, Italy and as far away as Mexico and New Zealand.

Figure 6. Perceived Photography Expertise

Methodology

The trial involved 3 separate stages and typically involved a 15-20 minute interaction with the participant:

1. **Background Survey** — During this stage the users completed a short survey documenting their photographic expertise (beginner, intermediate, expert - expert) and the various tools and equipment they were used to using (whether they used digital or film cameras etc). Figure 5 presents the results of this survey, showing the number of participants listing different types of camera (mobile phone, point-and-shoot compact camera, digital SLR, analog/film camera) as their primary camera, and Figure 6 shows the number of people who rated themselves as beginners, intermediates, and experts.
2. **Social Camera Trial** — Next, the participants were given a brief tutorial of the social camera app interface and an explanation of user interaction was given (remember it has been implemented on a standard HTC Android phone) participants were then asked to try out the interface to take a few pictures in their current location. They were specifically asked to pay attention to the recommended photos, the directional assistance, and the framing assistance interface components. Participants were asked to take photos with and without these assistive interface features. On average the typical user took 2-3 photos.
3. **Outcome Survey** — Finally, having used the social camera, each participant was asked for their views of the key photo recommendation, directional assistance, and framing assistance interface features. The participants were asked to comment on whether they would use such a feature if it was available on their own camera (a simple 'yes' or 'no' response), whether they found the interface to be intuitive and easy to use (on a scale of 1 to 5), and how they judged the quality of the results (again on a scale of 1 to 5).

Results

A summary of the key results from the survey is presented in Figure 7(a-i). In each case we have presented a bar chart of the average ratings provided by the participants, for the key features under investigation, and we have segmented these averages according to photographic expertise (novice vs. intermediate vs. expert) of the participants.

In the photo recommendation feature, we found a very high acceptance rate of this feature, amongst each user group. For example all of the expert users indicated that they would be interested in having access to an interface which provided a list of relevant photos based on their current location to serve as a source of inspiration (Figure 7(a)). In fact we had expected that, overall, this interface feature, and the social camera in general, would appeal to more novice users, but clearly this is not the case with only 65% of novices indicating an interest in this feature. The results were less overwhelming, although still broadly positive, when we look at the ease of use and quality of result metrics Figure 7(b & c). Once again the experts seemed happier with the feature. On average expert users scored its ease-of-use as a 4, compared to just over 3 out of 5 for the novices, and when it came to the quality of the results (that is the quality of the actual photos that were recommended) the experts scored them at about 3.5 out of 5 compared to just over 3 out of 5 for the novices.

Its worth highlighting a few points at this stage. First of all, the interface on the first version of the social camera was used on a low quality screen which meant it did not always work well in bright conditions and not all of the users are familiar with touch-based interfaces. Likewise when it comes to recommendation quality we are limited to the set of geo-coded images that are available for Dublin's tourist hotspots on Panoramio; given the need to filter photos by location and context there are no guarantees that there will be a critical mass of high-quality photos as the basis for recommendation.

Of the three key features under investigation, the directional assistance feature fared least well. A small majority of users indicated they would not be interested in this feature and overall most found it difficult to use and with limited results. Again there are some likely reasons for this. Of all the three interface features the directional assistance feature proved to be very sensitive to bright daylight and many users could not easily see the directional indicator or the radar as they tried to line up the camera. In turn, because of the inherent lack of location accuracy that is built into GPS, it is not possible to guarantee lining up the user with the recommended photo position precisely, even when the user carefully follows the directional advice. Nevertheless, on average, the users rated the quality of the result of the directional assistance at a very respectable 3 out of 5. Directional assistance problems aside, once users did get to roughly line up the social camera with the

Figure 7. Results for the three assistive technologies of Social Camera

target scene they reported strong acceptance (at least among the intermediate and expert users) for the compositional assistance feature. For example, in Figure 7(g) we see that 80% of experts and intermediate users indicated a general acceptance of this feature. Once again there are some interface challenges — on average users rated the features ease of use at about 3 out of 5 — and bright sunlight did tend to interfere, at least on the Android screen, with the transparency of the recommended image. Once again though, users seemed to be broadly happy with the quality of the results — the end photograph that they took — giving a rating of just over 3 out of 5.

Professional Opinion

Finally, we also took the opportunity to have a sample of the images taken by the trial participants, with and without the assistance of the social camera recommendation features, reviewed by a professional photographer. However, due to a technical issue, not all photos saved successfully to the storage device therefore we could only consider users that had all their photos available. Therefore we picked a random sample of 34 photos, 19 photos were produced with the assistance of the

social camera recommendation features, while 15 photos were produced by the social camera app but with its recommendation features switched off. We asked our professional photographer to grade each image according to whether it was properly framed / composed or not; we specifically asked the photographer to consider the composition of the photos only, ignoring other aspects such as focus, exposure, etc. Approximately 67% of the photos taken with the benefit of social camera’s recommendation features were judged to be well framed/composed compared to only 58% of the photos taken without any recommendation assistance (See Figure 8). While these results are based on a small sample size they do suggest that those produced with the assistance of social camera’s recommendation features seem to lead to better compositions, all other things being equal. Of course in reality all other things are not equal: different users may be more or less able to frame their photos effectively which will certainly impact on the quality of the final image.

DISCUSSION

Overall we view the results of this initial user study to be broadly positive, especially given the constraints

Figure 8. Number of well composed photos

under which the study operated. For example, there are a number of factors that mitigate against very positive results:

1. The current social camera app has been developed on an Android phone. Of course this means that the device is a phone first and a camera second and so its ability to take high quality photos is not necessarily a match to that of a dedicated camera.
2. As it turned out there were some visibility problems during the trial that were exacerbated by the limited brightness of the Android display and the novelty of a touch-based interface.
3. As mentioned previously there are fundamental limits when it comes to accurate location sensing and the availability of critical mass of high-quality, relevant images, and these no doubt impacted on overall impressions of the system.
4. And of course we were asking random users to try an entirely different camera and interface to the one they were used to using. Within the constraints of a short trial, users will always face certain challenges in trying to adapt to an unfamiliar interface.

Nevertheless, the results do suggest that there is an interesting opportunity for a recommender system to drive the photography interface on a digital camera, thus playing a novel role in providing users with assistance when it comes to the proper framing and composition of scenes. We asked subjects how often they used other assistive features on their digital cameras. The typical results which came out from this were that many people did use many of the features, but not as expert users and did not always feel as if it was something they required to use when they took photos. Finally, while our work was aimed at helping novices, we found that expert photographers thoroughly enjoyed using the social camera due to the serendipitous photo and location discovery resulting from the image recommendation process. While these photographers did not necessarily feel the need for an assistive feature in taking a picture, they did enjoy being introduced to photographic landmarks that they would have otherwise not

known about. Alternatively we found novice users who did not find the technology useful were typically users who were not comfortable with touch screen technology, or augmented reality applications. Obviously this is not the case for all users who had a less than positive sentiment towards the social camera, but it was certainly a consideration that should be made for future work. Based on our work we identified two core uses for social camera, firstly a tool to help users compose a higher quality photo, these users could be attempting to mimic an image but ideally they could be trying to either improve their own understanding of composition by using high quality images as a guide. Alternatively a family or friend may be trying to include their acquaintances in a well taken picture while on holiday. The second class of user could be someone who uses social camera as a tool to plan photographic outings.

CONCLUSIONS

In this paper we have explored a novel application of a recommender system interface. The basic idea is to provide assistance to users as they prepare to take photographs. Specifically we have argued the need for composition and framing advice so that users can learn how to compose a well-framed photograph. Current digital camera technologies are very much lacking in this regard. As such we have described the design and development of the social camera system. This has been fully developed as an Android app. It provides users with recommendations of well-framed photographs based on their current context (location, direction, lighting conditions) so that the user can easily emulate an image of their choosing and, in the long-run, improve their own photographic competence. We have also described the results of an initial live-user field trial involving random tourists across a number of Dublin's tourist hotspots. The results of the trial suggest that this social camera app has the potential to provide a useful service. A majority of users liked the photo recommendation feature and also the framing assistance. There are some challenges when it comes to improving the user interface and overall user experience but we believe that this is a very interesting and new approach to user interface design and recommender systems.

ACKNOWLEDGMENTS

The authors would like to sincerely thank professional photographer, Paulo Nunes Dos Santos for his expert opinion (www.paulonunesdossantos.com). This material is based on works supported by Science Foundation Ireland under Grant No. 07/CE/11147 CLARITY

REFERENCES

1. William Bares and Byungwoo Kim. Generating virtual camera compositions. In *Proceedings of the 6th international conference on Intelligent user interfaces*, IUI '01, pages 9–12, New York, NY, USA, 2001. ACM.
2. Victoria Bellotti, Bo Begole, and Ed H Chi et al. Activity-based serendipitous recommendations

- with the magitti mobile leisure guide. In *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, CHI '08, pages 1157–1166, New York, NY, USA, 2008. ACM.
3. M. Brown and D.G. Lowe. Recognising panoramas. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pages 1218–1225 vol.2, oct. 2003.
 4. M Davis, M Smith, J Canny, N Good, and S King. Towards context-aware face recognition. *Proceedings of the 13th annual ACM international conference on Multimedia*, pages 483–486, 2005.
 5. Anind K. Dey. Understanding and using context. *Personal Ubiquitous Computer.*, 5:4–7, January 2001.
 6. Fredo Durand and Richard Szeliski. Guest editors' introduction: Computational photography. *IEEE Computer Graphics and Applications*, 27:21–22, 2007.
 7. Alexei (Alyosha) Efros, Ramesh Raskar, and Steve Seitz. Next billion cameras. In *ACM SIGGRAPH 2009 Courses*, SIGGRAPH '09, pages 17:1–17:131, New York, NY, USA, 2009. ACM.
 8. M. Gaubatz and R. Ulichney. Automatic red-eye detection and correction. In *Image Processing. 2002. Proceedings. 2002 International Conference on*, volume 1, pages I–804 – I–807 vol.1, Dec. 2002.
 9. Maria Håkansson, Sara Ljungblad, Lalya Gaye, and Lars Erik Holmquist. Snapshots from a study of context photography. In *CHI '06 extended abstracts on Human factors in computing systems*, CHI '06, pages 333–338, New York, NY, USA, 2006. ACM.
 10. M. L. Kherfi, D. Ziou, and A. Bernardi. Image Retrieval from the World Wide Web: Issues, Techniques, and Systems. *ACM Computing Surveys (CSUR)*, 36(1), 2004.
 11. Ying Liu, Dengsheng Zhang, Guojun Lu, and Wei-Ying Ma. A survey of content-based image retrieval with high-level semantics. *Pattern Recogn.*, 40:262–282, January 2007.
 12. D.G. Lowe. Object recognition from local scale-invariant features. In *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, volume 2, pages 1150–1157 vol.2, August 1999.
 13. Andrés Lucero, Marion Boberg, and Severi Uusitalo. Image space: capturing, sharing and contextualizing personal pictures in a simple and playful way. In *Proceedings of the International Conference on Advances in Computer Entertainment Technology*, ACE '09, pages 215–222, New York, NY, USA, 2009. ACM.
 14. Umberto Pannielo, Alexander Tuzhilin, Michele Gorgoglione, Cosimo Palmisano, and Anto Pedone. Experimental comparison of pre- vs. post-filtering approaches in context-aware recommender systems. *ACM Conference On Recommender Systems*, pages 265–268, 2009.
 15. Moon-Hee Park, Jin-Hyuk Hong, and Sung-Bae Cho. *Ubiquitous Intelligence and Computing*, volume 4611 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.
 16. Benjamin Proß, Johannes Schöning, and Antonio Krüger. iPiccer: automatically retrieving and inferring tagged location information from web repositories. *ACM International Conference Proceeding Series*, pages 1–2, 2009.
 17. Linda; Shapiro and George Stockman. *Query types: Query by example: Using a example image to define what to search for*. Computer Vision. Upper Saddle River, NJ: Prentice Hall., New York, 2001.
 18. C. T. Shen, J. C. Liu, S. W. Shih, and J. S. Hong. Towards intelligent photo composition-automatic detection of unintentional dissection lines in environmental portrait photos. *Expert Systems with Applications: An International Journal*, 36(5):9024–9030, 2009.
 19. Arnold W. M. Smeulders, Marcel Worring, Simone Santini, Amarnath Gupta, and Ramesh Jain. Content-based image retrieval at the end of the early years. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22:1349–1380, December 2000.
 20. Mark van Setten, Stanislav Pokraev, and Johan Koolwaaij. Context-Aware Recommendations in the Mobile Tourist Application COMPASS. In *Adaptive Hypermedia and Adaptive Web-Based Systems*, volume 3137, pages 515–548. Springer Berlin / Heidelberg, 2004.
 21. Luis von Ahn. Human computation. In *Computing Research that Changed the World: Reflections and Perspectives*, CRASS '09, pages 4:1–4:23, Washington, D.C., 2009. Computing Research Association.
 22. Ghim-Eng Yap, Ah-Hwee Tan, and Hwee-Hwa Pang. Dynamically-optimized context in recommender systems. In *Proceedings of the 6th international conference on Mobile data management*, MDM '05, pages 265–272, New York, NY, USA, 2005. ACM.
 23. Lu Yuan, Jian Sun, Long Quan, and Heung-Yeung Shum. Image deblurring with blurred/noisy image pairs. *ACM Trans. Graph.*, 26, July 2007.
 24. Richard S Zakia. *Perception and Imaging*. Focal Press, 1997.