

TicQR: Flexible, Lightweight Linking of Paper and Digital Content using Mobile Phones

Jennifer Pearson¹, Simon Robinson¹, George Buchanan², Matt Jones¹

¹ FIT Lab, Swansea University, ² Centre for HCI Design, City University
{j.pearson, s.n.w.robinson, matt.jones} @swan.ac.uk

Abstract. In this paper we introduce TicQR – a photo-based checkbox-enabled interface which bridges the physical and digital document domains, allowing automatic download or processing of useful data from paper documents. There is a long demonstrated need for people to be able to connect between printed material and digital information and services. By using a combination of image recognition and QR codes we are able to detect user marks on paper documents via a single photograph taken with a standard smart phone. This information can then be used to access the equivalent digital content, save contacts or URLs, or even order goods directly from local retailers.

Keywords: Paper documents, tick boxes, QR Codes

1 Introduction

Digital services now permeate every aspect of our lives. However, printed materials predominate in many everyday situations, and are commonly used to gather or communicate vital information. But paper itself does not connect directly to digital services; for that reason many have sought to create links to the digital domain via a range of technologies. Most existing solutions require specialist equipment to connect paper and digital content, or suffer from relatively coarse information granularity – for example, connecting a whole page or document to a single digital item. While these approaches do indeed connect paper to digital, we argue that it would be preferable to be able to use more commonplace mobile technologies, and at the same time be able to provide finer-grained information, even considering specific marks or annotations by the user on the paper document.

To achieve this goal, we present a novel technologically and interactively lightweight approach that lets readers extract information from a printed document. Our approach – TicQR, allows people to tick boxes to select interesting sections directly on paper, and photograph the page using a standard cameraphone. Given both the document's identity and the specific selections made, TicQR can then undertake a variety of actions depending on how the document's designer has configured the support for their text. We thus provide a rich opportunity for interaction that embraces both print and digital media, but requires no additional hardware or specialist paper.

2 Background

Paper documents have been around for millennia, and despite the ever-growing popularity of digital reading, paper remains ubiquitous. The physicality of paper offers affordances that computer screens or eReaders do not – such as folding and scrib-

bling, for example. Globally, for reasons such as cost, physicality, aesthetics and utility, paper remains a desirable medium in many cases [8]. Yet the digital medium has many advantages, including support for searches across content, links between items, and its physical space benefits. Mixed approaches offer a combination of the strengths of each form.

The TicQR design aims to bridge the physical-digital document gap in a manner that is minimally intrusive to the user. The system allows people to read a paper document, mark which parts they are interested in, and archive them digitally by taking a single photograph using a standard smartphone handset (see Fig. 1). This approach allows multiple ‘clippings’ on a page to be recognised and interpreted simultaneously, then presented to the user for browsing or later use. We envisage many scenarios where TicQR could be useful, some of which are illustrated later in Fig. 2.



Fig. 1. Using TicQR to interact with a newspaper. Left: reading a newspaper, finding and then ticking interesting content. Right: browsing saved items at a later time.

2.1 Related Work

Research into connecting physical and digital media is well established. An early example is the DigitalDesk [10], which combined a desk, and a camera and projectors mounted above, to provide digital augmentation of paper documents. A paper document was placed in a fixed location on the desk, and text was extracted using OCR, allowing cut-and-paste from paper into a digital medium. The DigitalDesk showed the potential benefits of linking print and electronic formats, but required a precisely controlled, calibrated environment.

Digital paper and pens are a more direct combination of printed and electronic media, typified by Anoto’s digital pen technology¹. PaperPoint used Anoto to support real-time annotation of on-screen presentations using a printout of the slideshow [9]. PaperLink [1] used a custom video pen to associate annotations with digital content. In contrast, Papiercraft [4] used Anoto ink markup of the printout of a digital docu-

¹ See: anoto.com

ment to manipulate the digital original. Using Anoto and similar techniques can allow users to trigger specific digital actions from ink on paper documents. However, augmented, patterned paper is required, and a dedicated battery-powered pen must be used to annotate. Printing must also be carefully controlled so that the right layout is associated with each piece of paper.

The *a-book* [5] also uses dedicated equipment to link paper and digital items. The system makes use of a graphics tablet placed under a paper document to allow manipulation (e.g., linking, searching) of handwritten content via a second display. Each of these systems demonstrates the potential for rich interaction and effective usability when using dedicated hardware. However, the associated costs and constraints impact on both developers and users. In response to such approaches, researchers have proposed solutions that use simpler, commonplace technologies, therefore allowing more individuals, organisations and communities to use such techniques.

One common contemporary example is QR codes, which allow direct links to digital material via a standard mobile phone camera. Scanning a QR code typically provides a single piece of information, however – for example, a single contact or a web link. If multiple pieces of information are required, then the user will need to scan multiple codes, or follow web links. Previous solutions to this issue have used one code or other identifier per page of media and then automatically downloaded all information from the page². This approach could easily result in downloading unwanted information, and cause higher user task loads as this content is filtered out.

It is well documented that paper is the preferred medium for annotation (e.g., [6,7]) and other small-scale information work. Our approach takes advantage of this fact, and allows users to mark on the document itself to select items. The mScan project [2] also uses camera phones and users' marks on paper, but to scan a specific, preset form with multiple-choice bubble fields, that is placed in a known position and orientation. While encouragingly reliable, the mScan format only supports very specific marks for data gathering. A dedicated stand is also required to fix the location of a form for analysis, much as with the DigitalDesk (cf. [10]) or a standard flatbed scanner.

To summarise, in contrast to dedicated equipment approaches, our technique uses QR codes to identify documents, and requires only a standard mobile phone camera for image processing. Our approach can also be used on-the-go, allowing the user to browse their clippings in a more 'laid-back' [3] manner than previous QR-powered document scanning designs. Finally, like PapierCraft, Paperpoint and the DigitalDesk, the TicQR technique allows richer physical-digital interaction with paper documents.

3 The TicQR system

The TicQR system offers richer experiences with physical documents by allowing people to tick items they are interested in on the paper itself. Our approach uses a combination of precisely-placed QR codes and image recognition to enable the system to determine both which document has been scanned; and, which checkboxes within the document have been selected. Fig. 2 illustrates several example usage scenarios and corresponding QR code-augmented paper documents that we have created

² For example, see kooaba-shortcut.com or augmentation examples such as aurasma.com

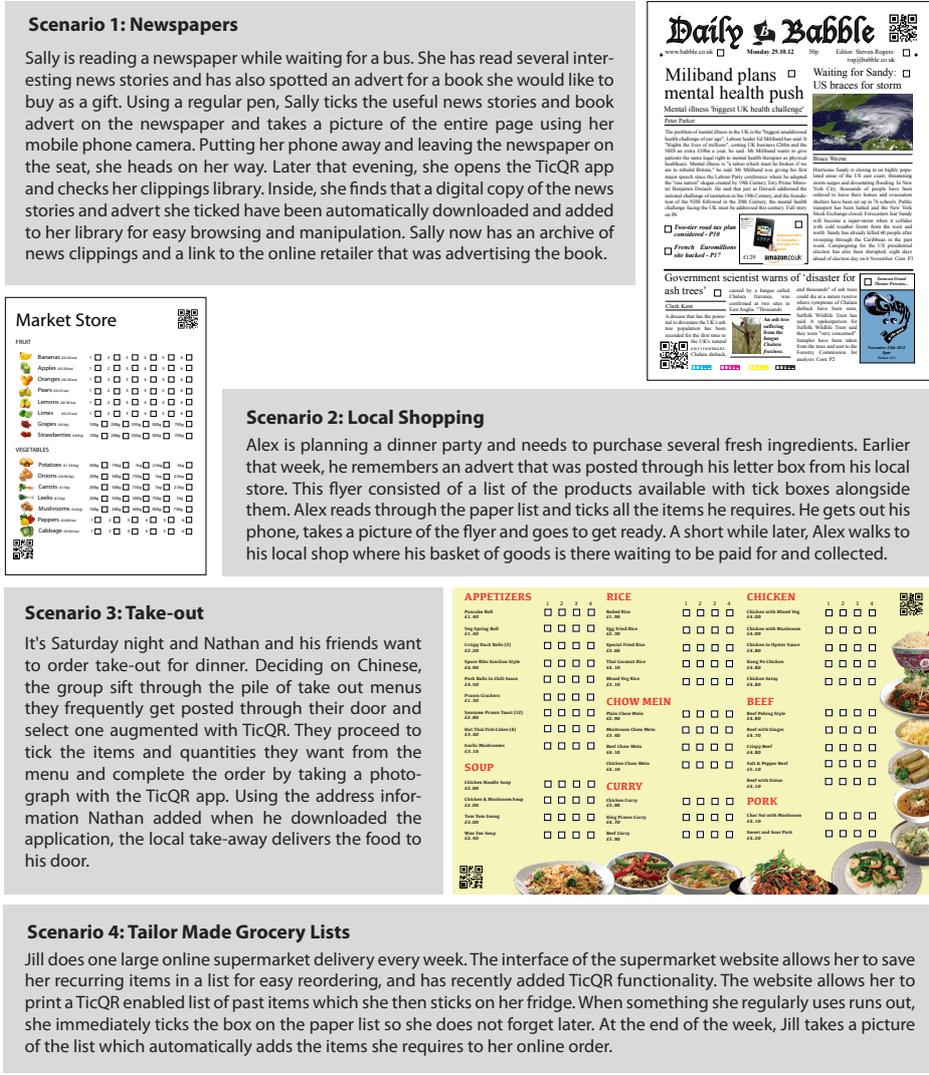


Fig. 2. Potential usage scenarios for the TicQR system. For scenarios 1–3, a TicQR- augmented paper document is also shown. Scenarios 1 and 2 were evaluated in our study.

for the TicQR system. To retrieve digital content from a paper document, the user first positions their cameraphone to take a photo of the item.

The TicQR application automatically takes a photograph when it detects two QR codes which are positioned at opposite corners of the object – one at the bottom left and another at the top right (see Fig. 2). The bottom left code contains the identifier for the item (e.g., the issue and page number of a newspaper), and the top right code is used for image alignment and coordinate calibration. Using the identifier from the QR

code, the application queries an online database to retrieve the item’s metadata, which contains the approximate coordinate of each of its checkboxes.

The current setup requires that the server side (for example, a newspaper publisher) provides a web service that the TicQR application can query to determine checkbox positions, as well as the content that should be attributed to the selection of each checkbox. If at the time the photo was taken there is no internet connectivity, the application simply saves the photo and performs these operations when a connection is available.

When the item’s metadata has been downloaded, the image is then processed to determine if any of its boxes have been marked. When a list of marked boxes has been generated, the application performs actions based on the specific items selected (see Fig. 2 for examples).³ Any downloaded content is saved locally in the application’s clippings repository for later use (see Fig. 1 and Fig. 3). All processing happens in the background of the application – we currently use the OpenCV⁴ library for checkbox processing, which takes less than one second to complete.

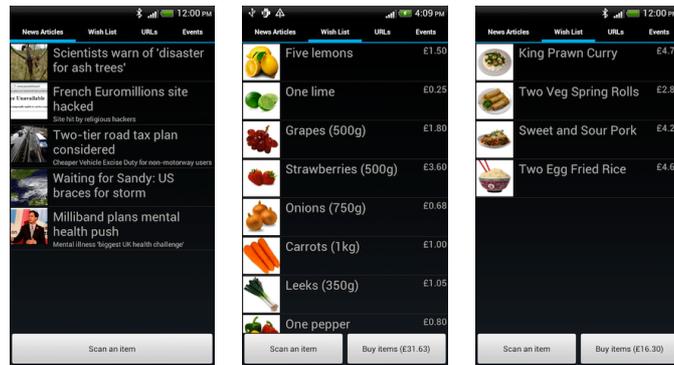


Fig. 3. Example TicQR interfaces. Left: a selection of newspaper clippings; Middle: a grocery order from a local market; Right: a sample order from a take-away menu.

4 Evaluation

We performed a lab study in order to help understand both the TicQR system’s recognition accuracy and its reception by users. 33 participants (16M 17F, aged 18–53) took part in individual 30 min trials. The metrics used were the accuracy rate of tick box detection, and qualitative data in the form of interview responses.

4.1 Prototype Media

We selected two types of prototype media to use in the evaluation – a newspaper (as described in Scenario 1 in Fig. 2), and a shop order form (as described in Scenario 2 in Fig. 2). The newspaper was chosen as its form offers a wide selection of different clipping types (e.g., articles, products, events), and typically includes several types on

³ A video illustrating the technique is also available – see: goo.gl/QBucU

⁴ See: opencv.org

a single page. The mock-up newspaper used in the study included five news stories, one event advert, one product advert, one URL and one contact. The shop order form used was designed as an extensive product list, which allows consumers the freedom to order specific quantities of a particular product, but also has the benefit of “stress-testing” the recognition system with a closely-spaced grid of checkboxes.

4.2 Procedure

At the start of each session the participant was welcomed and led through an ethically-approved consent procedure. Following this, the system was demonstrated once to the participant, and they scanned an example pre-ticked document as a form of training. We explained the two use case scenarios that would be examined in the study to provide participants with a context in which the system could be used. Participants were then given three identical blank copies of each type of prototype document (i.e., six documents in total), and asked to mark a selection of boxes of their choice. After ticking any number of boxes on a form, participants used the TicQR application to take a photograph of the document and either navigate the clippings library (newspaper), or order the list of marked products (order form). Participants were not told which boxes to mark, nor were they instructed on *how* to mark them (e.g., whether to tick, cross, scribble etc.).

We recorded the accuracy rate of the system's checkbox recognition based on the total number of boxes marked, and the number of false positives (i.e., un-ticked boxes detected as ticked) and false negatives (i.e., ticked boxes not detected). To conclude the study we conducted a short semi-structured interview to gather participants' existing document markup behaviours and their opinions of the TicQR system. The participant was then given a gift voucher as token of our appreciation.

4.3 Results and Discussion

Robustness A total of 1324 ticked boxes were recorded over all participants in the study (an average of seven per document). The overall checkbox recognition accuracy was 98.2%. Of the incorrectly recognised checkboxes, 17 were false positives (1.28%) and 7 were false negatives (0.52%). These mis-recognised boxes were rarely due to problems with the algorithm used to detect user marks, however. The majority of false positives were caused by factors which were out of our control, including the way in which participants marked up the paper documents – for example, when the end of a tick from one box inadvertently passed through another (see Fig. 4 (a) and (b) for examples taken from actual study participants).

Another minor issue that caused some of the false positive results during the study involved the way in which participants took the initial photograph. That is, if the photograph was taken when the image was slightly rotated, this could mean that some checkboxes at the edges of the document were only partially included in the photograph. The checkbox recognition algorithm used for the study ignored boxes which were completely outside the image, but detected boxes partially outside the image as ticked (we have since refined the algorithm to mark partial boxes as unticked). Of the 17 false positives observed, 9 were as a result of participants' marks running through multiple boxes, 4 were caused by the corners of the box being occluded due to participants rotating the camera, and a further 4 were true false positives.

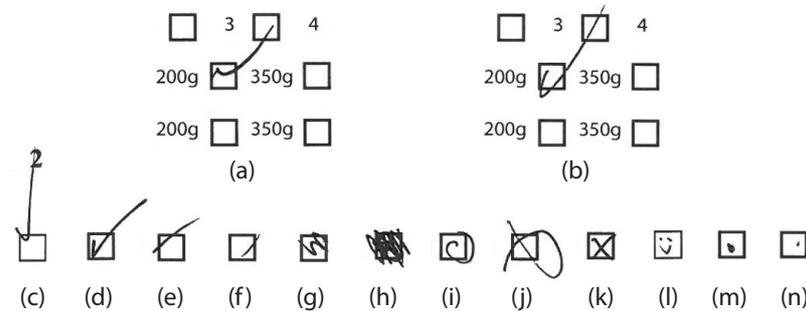


Fig. 4. Sample tick marks made by study participants. Top: participants occasionally ticked multiple boxes with one tick. All other marks except (n) were recognised correctly.

Several of the false negatives recorded were due to participants intentionally trying to fool the system. For example, the mark shown in Fig. 4 (m) was drawn by a participant who stated immediately after drawing it “let’s see how well it gets this one”, while another with a similar marking stated “I want to see if I can break the app”. That said, however, the system was able to pick up on the majority of users’ marks, regardless of how they were drawn. Fig. 4 (c–n) illustrates a representative selection of marks observed during the study, ranging from ticks, to lines, to crosses, scribbles and even a smiley face. With the exception of the small dot shown in Fig. 4 (n), all other marks shown were recognised by the TicQR system.

The majority of inaccuracies within the study were recorded on the order form rather than the newspaper, and in fact all false positives caused by rotation issues or user marks running through multiple boxes were recorded on the order form. We attribute this to the close proximity of the checkboxes to one another and to the edge of the form. This document was designed to test the recognition of the system, and is therefore not necessarily an accurate representation of a real order form. However, to reduce these issues, checkboxes could be placed further apart from each other, and away from the top left and bottom right corners of the document.

Subjective Responses All participants spoke favourably about the TicQR technique, giving an average score of 8.1 out of 10 (lowest 6, highest 10) for the usefulness of the system. Comments made by participants on the topic included: “I do love my shortcuts so anything to make my life easier is great” “it’s very simple and intuitive – I really ‘got’ it”; “it [TicQR] is great – I often take photos of documents myself to look back at later, so an interactive photo to source things you’re interested in is very efficient”; and, “The app does all the legwork for you – you don’t have to trawl the internet to find things”.

A majority of participants (88%; 29 of 33) reported finding themselves in a situation where they had read something on a physical document and later wanted to locate a digital version. The most common reason for this was to allow them to email, search or archive the digital copy. More interestingly, however, when asked if they did actually manage to locate a digital copy only 21% (6 participants) of these 29 had done so without some level of difficulty. The remainder had either not managed to locate a digital copy at all (14%; 4 participants), only managed to locate it on some occasions

(44%; 13 participants), or had found the digital copy only after some time or effort (21%; 6 participants). Some even described alternatives for locating a proper digital version such as taking a photograph of the physical document for later reference (21%; 6 participants).

5 Conclusions and Future Work

The TicQR technique enables automatic downloads of user-selected digital content from a paper document. While previous mobile phone-based methods required either the use of a search engine to locate related digital material, or the download of a complete package of digital content via a single QR code, TicQR affords the freedom to select only the information required via marks on the physical medium itself, and downloads selected content automatically in a single step. In contrast to previous interaction techniques, TicQR uses a standard cameraphone, which can be exploited with a much lower investment by individuals and small organisations.

Our prototype and user study demonstrate that the TicQR technique is practical and effective both technologically and in terms of usability. We have evaluated two concrete use cases for the TicQR technique, and provided several supplementary scenarios where it would also be beneficial. TicQR is thus a powerful and affordable tool for delivering richer interactive services than many existing techniques.

Our next step is to collaborate with a local produce market to provide a TicQR augmented ordering system. We hope that a longitudinal evaluation such as this will provide a fuller picture of the benefits of the system, and show how the combination of physical documents and digital ordering can provide richer user interaction.

6 Acknowledgements

This work was funded by EPSRC grant number EP/J000604/2.

References

1. Arai, T., Aust, D., and Hudson, S. PaperLink: a technique for hyperlinking from real paper to electronic content. In Proc. CHI '97, ACM (1997), 327–334.
2. Dell, N., Breit, N., Chaluco, T., Crawford, J., and Borriello, G. Digitizing paper forms with mobile imaging technologies. In Proc. DEV '12, ACM (2012).
3. Jones, M., Jain, P., Buchanan, G., and Marsden, G. Using a mobile device to vary the pace of search. In Proc. MobileHCI '03, vol. 2795 of LNCS. Springer, 2003, 390–394.
4. Liao, C., Guimbretière, F., and Hinckley, K. PapierCraft: a command system for interactive paper. In Proc. UIST '05, ACM (2005), 241–244.
5. Mackay, W., Pothier, G., Letondal, C., Bøegh, K., and Sørensen, H. The missing link: augmenting biology laboratory notebooks. In Proc. UIST '02, ACM (2002), 41–50.
6. Marshall, C. Reading and Writing the Electronic Book. Morgan and Claypool, 2009.
7. O'Hara, K., and Sellen, A. A comparison of reading paper and on-line documents. In Proc. CHI '97, ACM (1997), 335–342.
8. Sellen, A., and Harper, R. The Myth of the Paperless Office. MIT Press, 2003.
9. Signer, B., and Norrie, M. PaperPoint: a paper-based presentation and interactive paper prototyping tool. In Proc. TEI '07, ACM (2007), 57–64.
10. Wellner, P. Interacting with paper on the DigitalDesk. Commun. ACM 36, 7 (1993), 87–96.