

Mobile Cosmetics Advisor: An Imaging Based Mobile Service

Nina Bhatti^a, Harlyn Baker^a, Hui Chao^a, Scott Clearwater^a, Mike Harville^a, Jhilmil Jain^a,
Nic Lyons^a, Joanna Marguier^b, John Schettino^a, and Sabine Süsstrunk^b

^a Hewlett-Packard Laboratories, 1501 Page Mill Drive, Palo Alto, California 94304 USA;

^b Audiovisual Communications Laboratory (LCAV), EPFL, Lausanne, Switzerland

Abstract

Selecting cosmetics requires visual information and often benefits from the assessments of a cosmetics expert. In this paper we present a unique mobile imaging application that enables women to use their cell phones to get immediate expert advice when selecting personal cosmetic products. We derive the visual information from analysis of camera phone images, and provide the judgment of the cosmetics specialist through use of an expert system. The result is a new paradigm for mobile interactions—image-based information services exploiting the ubiquity of camera phones. The application is designed to work with any handset over any cellular carrier using commonly available MMS and SMS features. Targeted at the unsophisticated consumer, it must be quick and easy to use, not requiring download capabilities or preplanning. Thus, all application processing occurs in the back-end system and not on the handset itself. We present the imaging pipeline technology and a comparison of the services' accuracy with respect to human experts.

Keywords: Image Processing and Computer Vision, Applications and Expert Systems, Applications

1. INTRODUCTION

The first group to experience mobile applications has been the technologically savvy male. In this paper we consider applications that are likely to be used by others—not leading adopters—and, in particular, women. Ubiquity can not be achieved without creating applications that address this group, and their demographics suggest their more mainstream applications are likely to be related to mobile marketing. This trend is due in large part to the decline in effectiveness of traditional media: direct mail, TV, Radio, Print. Brands and retailers see mobile marketing and its personalization as the next great marketing channel. This interest in consumers and mobile devices has led to much experimentation and spending in mobile marketing, and this is expected to grow from \$708 Million to \$2.2 Billion in the US by 2011 [1]. Growth is more anticipated in applications that are fun: “The mobile tide is moving away from ring tones and wall papers to new growth opportunities...” [1]. Looking at the rates of cell phone penetration on a world-wide basis it is clear that for many populations this will be their single communication and computation device. Cell phone penetration in the US is already at 89% and continuing to grow [2]. One billion camera phones are expected to be shipped worldwide in 2009 [3]. In addition, many of the engines of innovation revolve around commercial applications, bringing an accelerating effect to the injection of technology to this space.

We created a mobile imaging-based application in response to a commercial request to help women shop for cosmetics. The goal was to increase women's satisfaction with cosmetics purchases by giving personalized product recommendations. We conducted a series of in depth interviews with women and discovered that they were quite frustrated with their cosmetics shopping experiences even when purchasing at an attended counter. Users also felt that they wanted to be able to use an application at a time of their choosing and that they were unlikely to plan use of the application and therefore unlikely to download software a priori. The technology also needed to be convenient and likely to be with the user when the shopping opportunity appeared, so the mobile phone in the consumer's pocket was the ideal platform. In this paper we will describe the user studies that motivated the creation of the application, the key features required to make it useful, how the application works, and the quality of the product recommendation system.

Women make 85-90% of consumer buying decisions each year [4]. This covers not only the traditional consumer packaged goods, but also durables such as consumer electronics, automobiles, home improvements, etc. [4, 5]. Even when women are not directly responsible for the purchase, they are strong influencers. In this paper, we present a mobile

marketing service targeted towards women to help them with one aspect of their shopping, the selection of cosmetics. Interestingly enough, statistics at project inception indicated that women took more digital photographs per month than men [6] and continued to lead men in camera phone usage. In 2006, 51% of women used their camera feature occasionally to frequently while only 40% of men did so [7]. We considered these conditions and hoped that women would indeed be comfortable in a mobile imaging application.

One of the areas of highest consumer spending is the beauty and cosmetics industry—a \$160 billion market. This industry includes makeup, skin and hair care, fragrances, cosmetic surgery, etc. [8]. The US color cosmetics industry alone is a \$29 billion market [9]. Research shows that foundation makeup is the second most popular product in the color cosmetics category. Foundation is a tinted liquid, cream, or powder substance that is applied to the face before other color cosmetics. Its main purpose is to even the skin tone, hide flaws, and provide a base for further cosmetic application. For this product to work effectively it must be selected to match the skin tone of the user. This requires the accurate perception of the consumer’s skin tone. Our user studies have confirmed what the cosmetics advertisements claim – it is difficult for women to select a foundation that best matches their skin tone, and this often results in purchase dissatisfaction.

Cosmetics selection requires visual ability, color judgment, and knowledge about fashion and trends, which makes for a complex and possibly frustrating situation. We will describe a system that provides these capabilities on any camera phone connected over any service provider. The camera phone will be used as a mobile imaging device with easy-to-use network transmission capability via Multimedia Messaging Service (MMS). A mobile technology in the hands of the consumer is felt to be ideal, as she can interact with an application in a store, at home, and at any time she needs assistance.

We also argue that this is a new paradigm for mobile applications—visual-data based queries. If a picture is worth a thousand words, we suggest in this case a picture may be worth a thousand key strokes. The application only requires the user to take a picture—about the simplest possible UI—and send the photo to the service. The service interprets the photo and sends back a response. If we tried to rely on speech or text for such an application, through the most common forms of cell phone usage, the process would be much more difficult. Instead, we exploit visual information, a photo. We have created additional imaging-based systems designed for specific interpretation of the image, targeted towards home décor, language translation, healthcare, and other uses.

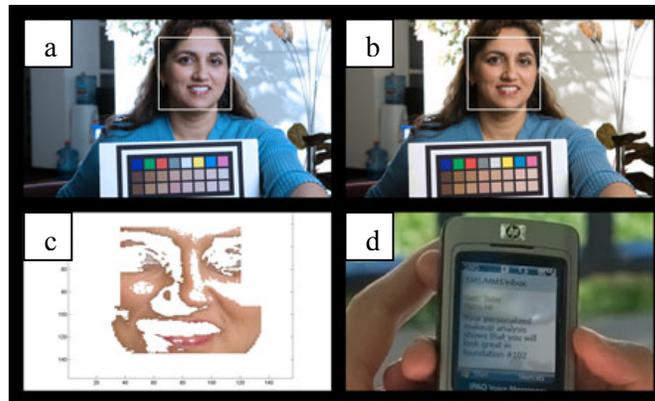


Figure 1. (a) Original image face detection b) color correction, (c) sorted skin pixels, and (d) SMS recommendation.

2. COLOR MATCH ADVISORY SERVICE

To address the issue of makeup selection, we have devised a mobile phone based solution that allows a user to photograph herself using a cell-phone camera while holding a specially designed color chart. This chart is critical because color accuracy is generally impossible to guarantee with uncontrolled illuminant conditions and arbitrary cameras. The color chart allows the image color values to be calibrated. The image is sent by the consumer via MMS to an advisory service host at a backend server. The system uses image processing algorithms to locate the face (Fig. 1a), color science to correct the image color (Fig. 1b), and statistical classifiers to sample the image (Fig. 1c) and determine

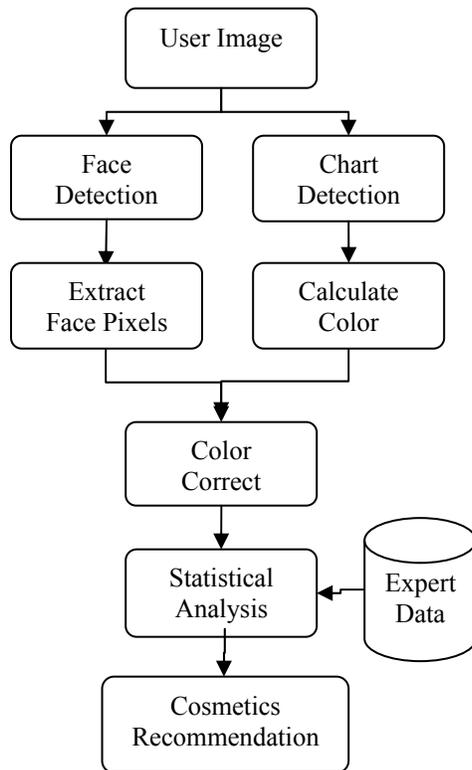


Figure 2. Color Match Service Block Diagram.

the user’s foundation makeup color with accuracy comparable to that of an expert. The photo analysis works regardless of image quality, camera specifications, and lighting conditions. Within seconds of sending the MMS, the consumer receives an SMS (Short Message Service) text message containing the foundation shade recommendation judged best for her complexion (Fig. 1d). Figure 2 sketches the flow of information through the recommendation process.

Our studies have shown that our system gives a distribution of predictions comparable to that found between two cosmetic experts under different lighting conditions and different consumer imaging devices and different cosmetic product lines. The system can be configured to plug in any expert opinion database popular with consumers, a brand-oriented expert, a retailer-sponsored expert, etc.

The user experience is tied to the consumer’s phone’s MMS and SMS facilities. We do not require any software downloads, special camera software, or hardware camera quality. It is designed to work with the lowest quality consumer camera cell-phone and still give good results.

The color chart is required for the color correction of the image to give accurate color information about the users face. For the system to be most effective the chart must be held near the face and roughly in the same plane as the face so that a similar illuminant lands uniformly on the face and the chart. A chart covered with a shadow or a bright light from the side will give poor results. We have ways of detecting this and they are described in the next section. The user could receive a diagnostic SMS reply with errors such “avoid shadows on the chart, please, and try again”, “no chart found in image, try again”. We found that feedback was essential to help novice users with the system.

Our application assumes that a user has access to a camera-ready cell phone; however, they must also have the color chart. The color chart appearance and distribution is an important requirement for pervasive use. Unfortunately, the color accuracy and calibration required makes the chart unsuitable for home printing. Based on conversations with cosmetics companies we have identified the most promising chart distribution schemes:

- Placed in a magazine as a “blow-in”. The most common place that women seek beauty advice is in magazines, so this seems the ideal place to make the chart available with instructions.
- Web based cosmetics and drug retailers include the chart with order shipments. Women already select some of their cosmetics online and there are many new brands that only distribute online.
- Built into a store display. The chart color values need to be visible but the subject can be standing in front of the chart. The chart would be larger and likely have a repeating pattern to deal with individual height variation.
- Available at makeup counters in retail stores for easy pickup to be taken home, or requested and mailed to consumers’ homes.
- For brands that are sold through authorized sales people and “parties,” the representative would make these available during the event.

3. IMAGING SERVICE

The cosmetic recommendation system is composed of image processing, chart detection, color correction, skin pixel extraction, and statistical classification of skin color to a database of expert opinions. In the next sections we will describe these aspects of the application.

Development of such a service is beset by many significant technical challenges. First, while accurate perception of a person's coloration is the most critical factor in determining a proper shade of cosmetic, this perception is affected by the lighting conditions. The captured image is also affected by the imaging device sensor characteristics and camera processing. For our service to work in the uncontrolled lighting of stores and other environments, with any camera, we must find a way to compensate for these effects. In addition, we must devise a robust system to find the user's face in the image, for a wide variety of poses, facial expressions, and distracting backgrounds. Further, once the face is found, we must judiciously sample it to determine the color of the skin and possibly other features, avoiding contamination by shadows, blemishes, sun-damaged areas, and specularities. In our view, successful automated cosmetic recommendation hinges on selection of appropriate facial appearance features and development of a recommendation system that ties these through machine learning to an expert-generated ground truth.

To make recommendations, our system processes an image and delivers a response to the user. Since the imaging technology is computationally intensive, and the goal is a result within a few seconds on any cell phone, computation is performed at the server-end and not on the mobile device.

Color Correction and Skin Color Assessment

In this section, we present consistent, fast techniques for measuring and classifying facial skin color from a single, casually posed digital camera image. The techniques are designed to allow the comparison of skin tones across a wide variety of cameras and environments, for all races of people. We also investigate the effects of varying image quality on our methods. Our aim is to develop methods that are practical for application in many computer vision and image processing tasks, beyond what prior art has enabled.

Reference Color Chart

Color calibration is achieved through the use of the reference color chart which has been specifically designed to calibrate for the color range of skin tones of all races. This allows the system to perform an approximation of a spectral analysis. Using robust computer vision algorithms the color chart is located in the image regardless of its orientation or placement. While no special positioning is required, results are best when the illuminant falling on the face and the chart are the same, so we recommend the chart be held in front at the user's chest area. The chart can take many forms. The simplest is gridded squares, as seen in Figure 1. But we have created several charts more in keeping with the esthetics of consumer brands (see Figure 3). We have found through our studies that we require a minimum of 5 skin-tone colored sampling areas, at least 1 inch in diameter. The appearance of the color values in the chart is used to correct for illuminant conditions and camera differences.

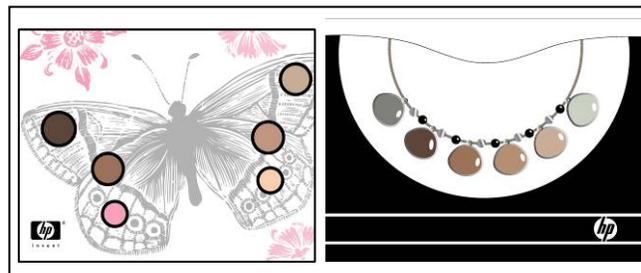


Figure 3: Example color chart designs: incorporated into a butterfly brand symbol, as a necklace.

We have found through our studies that we require a minimum of 5 skin-tone colored sampling areas, at least 1 inch in diameter. The appearance of the color values in the chart is used to correct for illuminant conditions and camera differences.

Validation of Color Correction

A basic question can be posed in addressing the validity of this effort: how effective can automated analysis of casually captured images be at providing the metric information needed for successful color-based product recommendation. Our studies have addressed this in several ways. With a desire to better understand the generality and metric character of our approach, we have submitted our methods to fairly rigorous analytic review. We pose the following three questions:

1. Can a cell-phone imager provide the color discrimination needed for a task that before this had been thought to require expensive colorimeters and spectral analysis?
2. Can a simple printed color pattern serve as a reliable reference palette for repainting an image to true color?
3. How broadly applicable is the technology in the space of inexpensive and varied-resolution cell-phone-type cameras?

We have addressed these issues in great detail in [10, 11], here we present a summary of the service's ability to color correct facial skin from un-calibrated images.

Imagers as spectra-reading substitutes

We first examine the question of adequacy of cell phone imagers for a color correction task that was previously the domain of spectral analysis—noting 3-band RGB imagers only sample a small part of the scene's color, while spectral analysis samples the energy over 40-50 wavelengths through the visible range. Our subjects, holding a copy of our color calibration target, were captured with four different single-chip CCD cameras—an HP 850 (3.9 MPixel), a Nikon D1 (2.7MPixel), an HP iPAQ camera (1.3 MPixel), and a Nokia 6820 CIF-quality cell phone camera (352x288 pixels). Faces were detected and extracted using Viola–Jones face detection [10,12] with only those selected face-region pixels having lightness between 10% and 90% of the distribution being considered (this removes outliers due to hair, contour shadows, eyes, lips, and blemishes—see Fig. 1c).

The mean color-corrected values of qualifying pixels are then used as the skin color estimate. Contrasting with these values, we measured ground truth as the skin reflectance of each subject's face from a uniform area of the cheek using a portable Microflash spectrometer (with 0°/45° measurement geometry). sRGB skin values computed for each of these using the observed skin reflectance spectra were then represented in normalized color coordinates. These values were compared with face colors extracted from the images, color corrected with the transform computed using all skin patches. We determined the correlation between imaged face color values and spectrally derived values to be high. The details of this analysis are published elsewhere [11].

Printed inks as true color reference

Previous work [10] has addressed our assessment of the color gamut and ability to use a linear transform on sampled values to adequately map in this space. The target itself being the only well-defined element in the image, we chose it to evaluate its own correction. The target is composed of color patches, primaries, gray and skin colors. Figure 4 shows the same image corrected with transforms computed using all 16 skin patches, eight skin patches covering the whole range of lightness, and the eight lighter patches only.



Figure 4. Correction with 16, covering-8, and lightest-8 tones.

The number and size of the patches influences the quality of color correction. Our color transform has twelve degrees of freedom, so we must have at least 12 measures—four 3-element RGB patches. We found the mean face color values obtained using four patches forming a convex hull in CIELAB space and the mean face color values obtained using the 16 skin colored patches to be very similar (Fig. 4 left and center)—differing about 1% in mean value and 5% in maximum value—showing sample quantity to be less significant than coverage.

Sensitivity to camera quality

Our test cameras covered the widest range of resolutions available—from megapixel through CIF. We had concerns about our ability to both detect the target in low quality images and to avoid sampling artifacts inherent in cell-phone transmission where providers often permit only much reduced bandwidth data upload. In these conditions, compression can be a major issue.

Larger color patches will be more reliably detected, especially in low-resolution cameras, and reducing the number of patches allows us to increase their size. In the images taken with the Nokia 6820 CIF cell phone, patches can be so small that JPEG artifacts begin to introduce significant errors in the estimation of patch color, and this presents the limiting case for our analysis. Notice the blocking and re-sampling artifacts in the roughly 15x15 CIF patches in Figure 5. While

skin tones can be color corrected using only four patches, for robustness this is not recommended. A shadow projected on just one patch, for example, will cause the resulting color transform to be seriously degraded. More robust methods such as Ransac [13] could ameliorate this, and would be advised for a production system. Figure 6 shows this and several other capture anomalies that could deter correct target recognition and/or color adjustment. As with the shadow handling, future additions to the system will accommodate to these challenges as well.

A detailed analysis using mapped lightness and normalized chromaticity showed that our color correction is coherent across resolutions, with ΔE values ranging from 1.8 to 5 [11]. The highest error came from the CIF imagery, indicating that, while discernible to a human, the difference in correction still provided good classification value over the tested recommendation classes throughout a broad range of camera characteristics. As with the overall color correction issue—where it was not essential that we make everything appear better, just the skin—here we only need performance that permits sufficiently accurate recommendation.

The high correlation across devices indicates that with a solid ground truth, the system can be correctly trained to give sufficiently consistent results for reliable use. While the assessment of color is not perfect, our method requires neither expensive calibrated imaging devices nor a controlled illuminant. It has been demonstrated to perform across a variety of consumer cameras. Our analyses have given us confidence in the system’s ability to meet the demands of wide deployment.



Figure 5. JPEG compression effects.

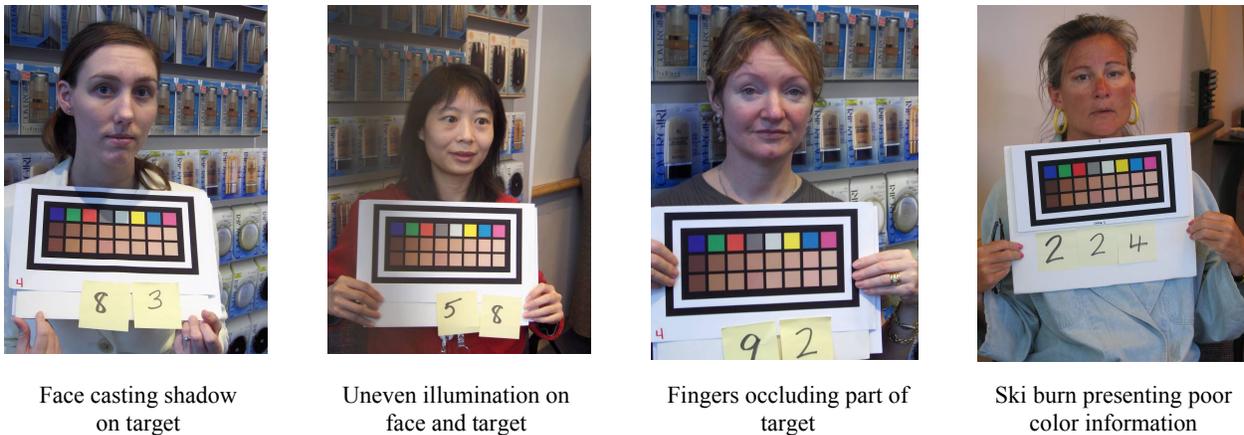


Figure 6. Acquisition anomalies that could deter accurate recognition and color adjustment.

4. CLASSIFICATION OF SKIN COLOR AND PRODUCT RECOMMENDATION

While the imaging service can compensate for lighting and camera system variations it can not provide cosmetics selections, there is one final and critical step to formulating a cosmetics recommendation. A machine learning based approach is used to replicate expert opinion. We have created a statistical classifier that relies on the use of expert generated ground truth.

Cosmetics is not “paint-by-numbers,” that is, detecting exact skin color is not sufficient and color matching to a foundation is not sufficient. To illustrate, if a user is sunburned she is unlikely to want to match this shade, instead the system should recommend a color of foundation that will compensate for this condition. Cosmetics experts know to correct too much red in the skin by applying a yellow toned foundation. We seek to present the same quality of advice. *The goal of our system is to be as good as a professional beauty consultant.* We seek to present the same quality of advice so we need to collect the functional beauty knowledge of products. We accomplish this through the use of systematically collected ground truth and a statistical classifier to replicate cosmetic expert opinion.

Collection of Ground Truth

We have developed a protocol that ensures consistent data for statistical classification and captures the artistry as well as knowledge of cosmetics products. We have performed this protocol with 300 human subjects, several experts, and 6 international brands of cosmetics from low cost to luxury priced brands. The subjects were selected for ethnic diversity so ensure inclusion of all complexion colorings. Our cosmetics experts were instructed to take as much time as needed and that consistency was the critical factor in assigning cosmetic foundation. The data collection is done in the following steps:

- 1 The clean faced subject is photographed (with a color chart). A photometric, full spectral analysis reading of her skin is taken.
- 2 The cosmetic expert is presented with a human female subject with cleansed skin. She judges her complexion and selects the most appropriate foundation for the subject in given brand of cosmetics.
- 3 The cosmetic is applied and the “dry down” effects are observed. Corrections can be made to the selection at this time if the color blending of the cosmetic is not pleasing.
- 4 The foundation color is recorded for each brand in our study.

We analyze the expert’s cosmetic ground truth and locate any odd effects that would affect the consistency of recommendations. That is, we suppose that if the skin color of two subjects is very close (using spectral reflectance and photographic data) it implies that the cosmetic recommendation should be the same. If two subjects are found to have the same photometric coloring and yet the recommendations are quite different, the data is removed from the system. This happened quite rarely but this potential expert “mislabeling” could affect our results. Classifiers do not perform well with inconsistent data. We also check the data to ensure that each cosmetic shade has a minimum number of *exemplars*, subjects for which the foundation shade is correct. This ensures the classifiers ability to identify appropriate wearers of each shade.

Statistical Classification

Several machine learning approaches were evaluated to create classifiers that could predict foundation choices. All rely on high quality ground truth data. One of the more successful methods we use is a Kuiper Statistic method which compares distributions of color values. For image to generate a recommendation for: we compare skin color distribution of R,G,B values (after color correction is applied) to each exemplar from our ground truth set. The exemplar whose face has minimal difference in coloring distribution is used as the basis of product recommendation. The product selected is chosen from our exemplar indexed database which contains the beauty expert’s techniques and judgments. We use the logic that if the subject looks most like an exemplar, then her best cosmetic will be the same as its. In practice, this has led to good results and we presented a blind test between two experts and show that our method results in the same differences as two experts.

Results

To evaluate our performance we invited two experts to classify the same group of women. Expert 1 was a professional makeup artist and educator who classified by observing the subject, selecting a foundation from the line, and testing it on the subjects skin. Expert 2 used a proprietary photometric based classification table. She measured the subject’s skin and indicated the foundation selection. Experts were not aware of the classification by the other expert so no cross influencing was present. To measure error from our foundation recommendation from another, we devised a consistent error metric D, Distance. Shades are not uniformly distributed, e.g. there are several nearly identical shades in the light shades and much greater gaps in the darker shades. The makeup expert assigned a 1 for distance for general one shade difference, .5 for very close shades, and greater than 1 for shades quite far apart.

In table 1 you can see what the error from Expert 1 was from Expert 2 given in Distance. Note the two experts on average have a difference $< 1 D$. That is a little less than a shade off. This establishes the flexibility inherent in expert opinions and a bound on our prediction error.

Table 1 Camera based performance measures

Camera Image Source	Resolution as tested	Average D (distance from Expert 2)
HP 850	4 M / 4	0.764
Canon S400	4 M / 4	0.867
Nokia 6820	0.1 M	0.896
Nikon D1	2.7 M / 4	0.970
HP R707	5 M / 4	0.943
HP R817	5 M / 4	0.702
Expert1	n/a	1.094

So how well does our system do compared to the human experts? We compare our application to the expert opinion. This study was performed on a data set of 73 subjects. The 73 subjects were the basis of the classification and we used a “leave one out” evaluation method on each subject. The subjects were classified by our system after being photographed with the color correction chart. The subjects were also classified by Expert 1 and Expert 2. Note that several photographic devices were tested to see if our results would be too sensitive to photographic devices. The resolutions used are given in table 1 and are typically equal to or less than MMS services typically support. In Figure 7 you can see that the probability of the color match system doing as well as or better than the expert is consistent across all imaging devices. We conclude that our system’s correctness rival’s the quality of expert opinion. This validates not just the classification but the entire application.

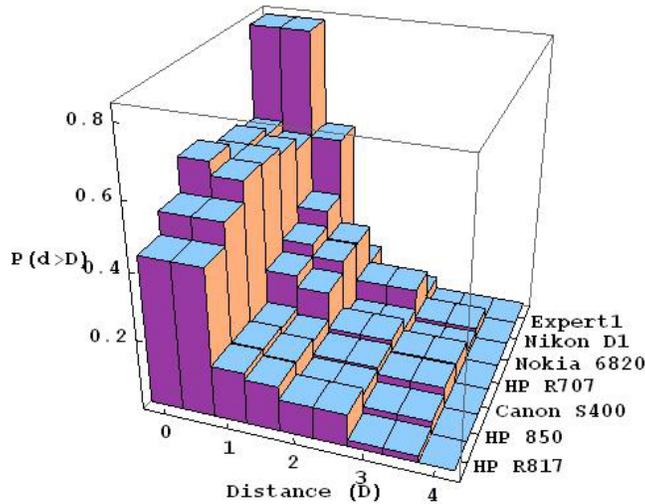


Figure 7: Comparisons of our system to Human Experts.

5. ARCHITECTURE

We selected a no-download style of application because of the casual use requirement, and the specifically non-technical target audience. Research shows that the main factors’ hindering wide-spread usage of mobile services is the customary requirement that additional software be downloaded and installed on mobile devices and the learning curve associated with running the new application. To avoid these drawbacks, our solution is designed so consumers can use the existing MMS/SMS software on their phones and not need to download or install additional software.

This no-download style means that all the intelligence is in the backend system. The semantics of the action taken when the picture is received is defined by the destination address. Additional information is supplied by objects in the image located through computer vision analysis or as additional text messages linked on to a session identified by the cell phone's return address (or phone number). The session definition allows arbitrary additional information however, since each additional interaction may cause users to drop out, this requires caution. So our applications, while potentially supporting sessions, focuses on giving information to users with the least onus possible.

The UI was selected to be as simple as we could manage. The semantics of the action on the image is defined by the message service destination and no other parameter or input values are needed. If additional parameters are needed for example, chart variation, it could be encoded in the chart design itself as a visual code or a variation in the chart layout.

The architecture of the mobile imaging system is divided into transaction manager and imaging servers. The transaction manager does message handling, parsing of messages, dispatching of images to the imaging service, reply or SMS posting, timing of the imaging responses, and error handling. The transaction manager also translates the imaging service results into consumer names for brands and cosmetics. This is done by consulting a marketing database with cosmetics information, and potential special offers and other promotions in real time based on the user's information, product recommendation, and other marketing data criteria. This gives advertisers and brands very close control over the application. Products, color shades, ground truth, and advertising campaigns can be updated at any time.

The Imaging Servers implement the imaging pipeline described in earlier sections of this paper. We have also created a sign translation imaging service and color recommender for décor design. Each of these applications is implemented as separate imaging servers but share the same transaction manager.

The steps to complete a transaction in our architecture are:

1. Cellular transmission of user photo via MMS, destination is email or short code.
2. The cellular carrier receives the MMS and forwards through internet gateway (e.g. SGSN) or SMS aggregator to Internet.
3. SMTP received at Transaction manager.
4. HTTP post of image appropriate to Imaging Service based on email destination or short code ID.
5. HTTP post of image processing results to Transaction Manager.
6. Reply to cell phone via email (SMTP)
7. Cellular carrier receives email destined for handset and converts to SMS
8. Cellular transmission to handset

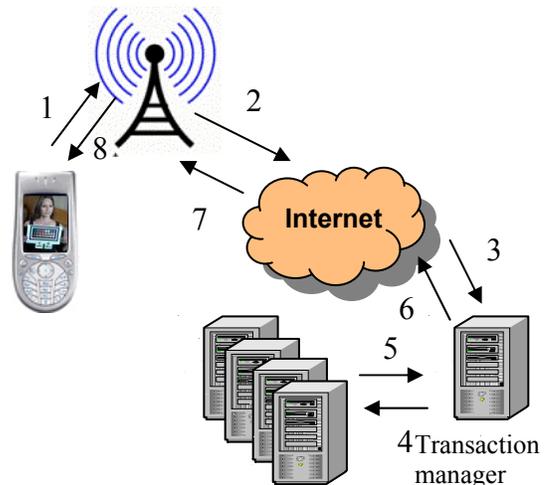


Figure 7: End-to-End System.

6. RELATED WORK

Most past work in mobile imaging has centered on recognizing printed 2D bar codes and responding with a URL or related content [14, 15, 16]. These systems are very popular in Japan where many camera phones come preloaded with bar code recognition software. Neomedia [17] and Scanlife [18] offer this software for download on any phone.

A well established research area of camera phone applications is image based localization. [19, 20] propose using cell phone cameras as indoor localization devices. A user walks inside a building with a cell phone worn around his or her neck, and the phone takes pictures at regular intervals and sends them to a server. Localization is inferred by comparing the query picture to a database. Other image based services have been used for personal memories [20].

Image-based search query [21] is another possible application for mobile phone cameras. Arguing that entering a text query on a mobile phone is too cumbersome, we advocate the use of pictures—supplemented with keystrokes as necessary—for querying search interaction.

Several companies are exploring image recognition for mobile marketing applications, Clic2C [22], Snapnow [23], and Snaptell [24], (now A9, part of Amazon). None of these use images in understanding the consumer’s needs in creating a personalized response.

7. USER EXPERIENCE

For a consumer application to succeed it must be very easy to adopt, require little pre-contemplation before using, and address key unmet user need. For simplicity in delivery, we elected to use the native MMS available on common cell phones. We could have designed an interface to be supported only on a select group of phones, but this would lack the possibility of significant consumer penetration needed for a commercial solution.

The user interface is simple: take an image with the color chart, form an MMS, and send to either a short code or an email destination. The user interface for taking a picture and creating and sending an MMS is dependent on the phone the user has but the user is very likely to know how to take a photo and send the image. It does depend on a subscription to a data plan.

A bit trickier in terms of the user experience is the holding of the chart while taking the photo. A self-portrait experience is sketched in Figure 8. This user experience is much easier when the camera lens and display are on the same side of the phone, otherwise, the user is blindly positioning the chart, taking the photo and then checking if the framing was correct. In practice this is done most easily with the help of a friend. We found in our group interviews that young users, less than 20 years of age, were very likely to enjoy shopping with a friend and would have a companion with which to complete the task. The older woman was more likely to try this at home where she had assistance from younger more “tech-savvy” members of the family. Our target use cases included home use and in-store use of the technology. Again our user research showed that some women were comfortable with public use while others preferred the privacy of the home [25]. The technology also requires a clear photograph of the skin that is free of makeup. Some users preferred the bare skin experience at home while others were comfortable shopping without makeup.

8 CONCLUSIONS AND FUTURE WORK



Figure 8: “O”, (Oprah Magazine) illustration

Recently our system received coverage in *O, The Oprah Magazine*, a leading women’s lifestyle publication [26]. Here and elsewhere we have received interest from a press intrigued with the novelty of a mobile system that can offer advice as good as that of an expert. A predominant element of the feedback from our user studies and cosmetics companies is that *people want rich services that are simple and easy to use*. We feel that addressing a real user need with applications that are simple and intuitive is key to achieving ubiquitous mobile services.

In this paper we have demonstrated a novel mobile application targeted toward women. This application has unique imaging and computer vision capabilities that allow the application to work with any camera phone. We have demonstrated how a camera phone image-based system can provide high quality recommendations comparable to that of an expert who is

viewing the individual. We imagine creating further image-based queries for mobile services for healthcare and other application areas. Work has already begun on a home décor advisory application in which existing fabrics, paint, wall coverings, floor coverings, etc. are analyzed and color recommendations are made to construct a palette with aesthetic design quality comparable to that of an interior design professional.

REFERENCES

- [1] Jupiter Research Mobile Marketing & Media. US Mobile Marketing Forecast 2007–2012. Jan 2008.
- [2] Parks Associates. New Advertising Platforms & Technologies. Feb 2008.
<http://www.ctia.org/advocacy/research/index.cfm/AID/10323>
- [3] Infotrends. Worldwide Camera Phone Forecasts: 2007–2012. July 2008.
- [4] The Six Costliest Mistakes You Can Make in Marketing to Women.
<http://www.inc.com/articles/2003/01/25019.html>
- [5] Barletta, M. Marketing to Women: How to Understand, Reach, and Increase Your Share of the World's Largest Market Segment. Kaplan Business, 2006.
- [6] Infotrends. 2004 Worldwide Camera Phone and Photo Messaging Forecast. March 2004.
- [7] International Communications Research, Cingular Wireless' Annual Fathers Day Survey. May 2006.
- [8] The Beauty Business, The Economist, May 2003.
- [9] Gavenas, M.L. Color Stories: Behind the Scenes of America's Billion-Dollar Beauty Industry, Simon & Schuster, 2002.
- [10] M. Harville, H. Baker, N. Bhatti, and S. Ssstrunk, "Image-based measurement and classification of skin color" Proc IEEE Intl Conf Image Process 2, 374–377, 2005.
- [11] J. Margaier, N. Bhatti, H. Baker, M. Harville and S. Ssstrunk, Assessing human skin color from uncalibrated images, International Journal of Imaging, Systems and Technology, special issue on Applied Color Image Processing, Vol. 17, Nr. 3, 143–151, 2007
- [12] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features", International Journal of Computer Vision, 2001, 511–518.
- [13] M. A. Fischler and R. C. Bolles., "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Comm. ACM* 24, 6, 381–395, 1981.
- [14] Jun Rekimoto and Yuji Ayatsuka, "Cybercode: Designing Augmented Reality Environments with Visual Tags", Proceedings of DARE 2000 on Designing augmented reality environments, 2000.
- [15] <http://www.mobilecodes.org/>
- [16] <http://www.activeprint.org/>
- [17] www.neom.com
- [18] www.scanlife.com
- [19] N. Ravi, P. Shankar, A. Frankel, A. Elgammal, and L. Ifode. Ravi et al, "Indoor Localization Using Camera Phones", 7th IEEE Workshop on Mobile Computing, April 2006.
- [20] Jesus et al., "PhotoNav: A Retrieval Based Interface to Navigate in Physical Spaces Using Memories", WWW2007, May 2007.
- [21] X. Fan, X. Xie, Z. Li, M. Li, and W. Ma. Photo-to-search: using multimodal queries to search the web from mobile devices. In ACM SIGMM international workshop on Multimedia information retrieval, 143–150, Singapore, 2005.
- [22] www.clic2c.com
- [23] www.snapnow.com
- [24] www.snaptell.com
- [25] Jain J., and Bhatti, N. "Imaging-Based Cosmetics Advisory Service", In CHI '09 Extended Abstracts on Human Factors in Computing Systems (CHI 2009). ACM, New York, NY.
- [26] O Magazine, February 2009, Vol. 10, Nr. 2
http://www.oprah.com/slideshow/omagazine/2009_omag_beauty_favorites