

Field Challenges for Mobile Photo Recognition Food Logging

1st Author Name

Affiliation
City, Country
e-mail address

2nd Author Name

Affiliation
City, Country
e-mail address

3rd Author Name

Affiliation
City, Country
e-mail address

ABSTRACT

Food consumption heavily influences one's health and a poor diet can lead to chronic diseases. Food logging using mobile apps can provide a continuous way to longitudinally track dietary behaviors and nutritional intake. This can help users understand their eating habits and encourage healthier food choices. However, current apps still pose many usability challenges, including tedious manual text entry of food names. Fortunately, advances in computer vision and deep learning are enabling automatic food recognition for instant and convenient logging. We have developed a mobile app prototype and conducted formative investigations into the usability and usage of mobile photo recognition food logging in a series of studies: online requirements survey, usability lab study, and 1-week field trial in an Asian country. Our findings reveal patterns and challenges in usage, usability, food dataset coverage and accuracy, and localization issues. We further discuss opportunities for design and technology to address these challenges.

Author Keywords

Food Journals; Food Logging; Image Recognition; Mobile Applications; User Experience; Field Study

ACM Classification Keywords

H.5.2 User Interfaces: Mobile Dietary System

INTRODUCTION

There is an increasing concern for diet-related chronic diseases caused by having an unhealthy diet, such as obesity, heart disease and cancer [41]. For example, consuming too many high-fat and energy dense (high caloric) foods with sedentary living can cause obesity [43], diabetes can be managed through weight loss [44] or carbohydrate counting [40], and the “risk of colorectal cancer could increase by 17% for every 100 gram portion of red meat eaten daily” [42]. Increasing user awareness of

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or re-publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Request permissions from Permissions@acm.org.
Copyright © 2017 ACM

their behaviors can promote health behavior change (e.g., calorie tracking to choose lower-calorie foods [38]). Indeed, many consumers are willing to use mobile apps to monitor their food habits (e.g., [19, 28]). However, food logging apps continue to have many barriers [9], including having tedious manual text entry of food names. Fortunately, advances in computer vision and deep learning are enabling automatic food recognition for instant and convenient logging. In this work, we investigated using such technologies in multiple studies culminating in a field trial.

Our contributions are: (1) the development of Nibble, a mobile app for fast, automatic recognition in photo-based food logging and diet feedback, and (2) formative investigations with an online requirements survey on diet habits and food logging, usability lab study on logging preference and diet feedback, and 1-week field evaluation in an Asian country. We found demand for fast, convenient, accurate photo logging and immediate diet feedback. We describe user experience in using automated photo logging with search logging as a fallback. With a deployment for an Asian city, we also discuss implications for localizing mobile food logging for the cultural context.

RELATED WORK

While traditionally paper diaries have been recommended by dietitians [14], many mobile apps have been developed to support food logging. Many apps simply digitize the text entry by providing a means to enter the food names, and this remains tedious even with search support [9]. Hence, several techniques have been developed to lower the barrier to food logging, such as using mobile phone cameras and other automatic sensing.

Previously, mobile phone cameras were primarily used for data capture, while the interpretation of the image and food recognition was delegated to other people. The human effort is performed through expert feedback, peer rating, crowdsourcing, or self-reflection. Professional dietitians can provide credible and accurate expert feedback to users and improve adherence [20, 22, 24], but relying on experts for each mobile user is expensive. On the other hand, The Eatery app used a free approach by requiring users to rate the healthiness of foods eaten by other users (peers) [17], though, this suffered from low adherence (2.6% active users). A middle-ground approach uses crowdsourcing, which employs cheap online labor to recognize foods in the photos, but the per-image cost is non-trivial

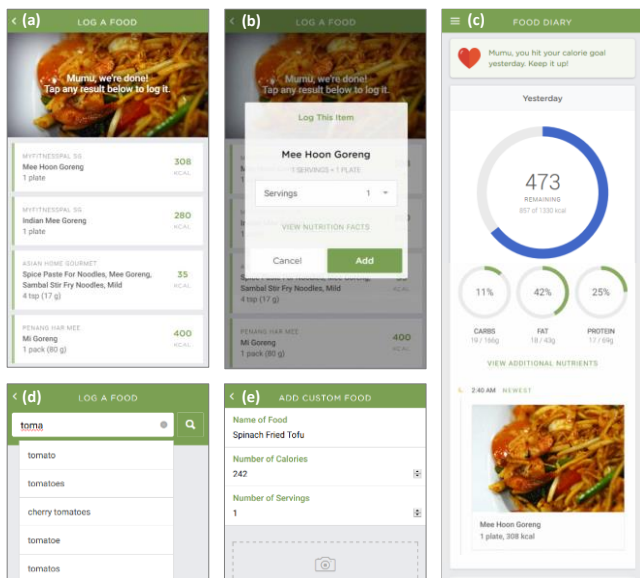


Figure 1. Screenshots of the Nibble app showing how a user logs food: first taking a picture via the smartphone camera (not shown), which gets recognized via the FoodAI API and returns (a) a list of candidate food names in under 1 sec. After choosing, the user (b) specifies the portion size and (c) gets nutrition feedback about calories and macronutrients in a diary. Since photo-based food recognition is nascent and novel, users may face issues in its reliability and trust, so Nibble also includes fallback methods for logging: (d) text search and (e) custom entry creation.

(USD1.40/photo) and labeling is not especially fast ($M=94$ minutes) [29]. Cordiero et al. provide an interesting alternative where the user reflects on her own diet using her food photos [8], but it may suffer from selection bias in users, unsustainable reflection effort. Epstein et al. extend this work to support mindfulness in a lightweight mobile app which only requires logging one meal per day [10]. In our work, we focus on fast, convenient food logging and feedback through automated nutrition analysis.

Automating food recognition can provide a scalable, affordable means for nutrition analysis. Several methods include scanning receipts [24], chewing sounds [31] and ego-centric meal detection [37]. While these use commodity devices, we focus on ubiquitous smartphone cameras for food recognition. Recently, there has been significant research in using computer vision and machine learning for automatic food image recognition (e.g., [4, 12, 26, 27]). Such technology can provide a basis for convenient and fast recognition in photo-based food logging. However, such research has focused on algorithm development and validation on datasets, and it is unclear how end-users will use them. Previous studies with user evaluations have explored the use of website interfaces [8, 21] or mobile apps but with delayed feedback [17, 29]. In this work, our user study provides a formative investigation on the user experience with automated photo recognition food logging *in the field* and elucidates several challenges in data preparation and app operation.

NIBBLE MOBILE APP PROTOTYPE

We implemented Nibble, an mobile web app for photo food logging with automated food recognition. Nibble is a wellness application designed to help users set healthy diet goals, display useful visual summaries and provide effective feedback to guide them towards healthier diets. Figure 1 describes key design features for the food logging.

Food Recognition to Provide Nutrition Feedback

To perform the food recognition, we used the application programming interface (API) of [12] which uses a Convolutional Neural Network (CNN) [23] trained on the top 100 local foods in Singapore. Once Nibble retrieves the food name, it looks up the nutritional information of the food item from two nutritional data sources: a food-nutrition database provided by the Health Promotion Board of Singapore [16] and the Nutritionix API [29]. This nutrition data is presented as feedback to the user in terms of user friendly donut and time-series bar charts indicating calories, macronutrients, and micronutrients. The feedback is provided immediately after the food is logged and identified (post-logging) and through the day (daily).

Other Behavioral Support Features

To aid with usability and minimize confounds, Nibble has other features such as *reminder* triggers [1] to prompt users to log their meals at their typical meal times, and *weight goals* [6] to help motivate users towards and objective.

METHOD

We conducted formative investigations into the usability and usage of mobile food logging with automated food recognition in a series of studies. First, we ran an online survey on diet habits and food logging to gather user requirement and barriers to photo-based food logging (Reddit: 31 global respondents). This helped us to identify key features to implement, such as the popularity of search logging. We then implemented an interactive mock-up on a laptop and conducted a scenario-driven usability lab study (5 participants). We took findings from this study to refine and add features to reduce usability confounds. We then deployed the Nibble mobile app prototype in a 1-week field study (7 university students) to evaluate user acceptance and usage of photo logging with automated food recognition. We instrumented Nibble for interaction logging, conducted pre/post-study interviews and surveys, and gave short surveys about usability and accuracy using the Experience Sampling Method [7] randomly triggered after some food logging events. All survey questions were asked on a 5-point Likert scale (≤ -1 : Disagree, ≥ 1 : Agree).

While the user studies were brief and each with a limited set of participants, together, these form a set of formative user studies which identifies several challenges in designing, developing, and deploying mobile food logging with automated food recognition.

FINDINGS AND DISCUSSIONS

We summarize the findings learned from designing and developing Nibble, and our three user studies.

Photo Recognition Food Logging and Fallbacks

Although photographing food is popular on mobile social media apps [3], 70% of our online survey respondents reported that they never take photos of their meals. Nevertheless, 73% of them were willing to try photo food logging, giving promising results to develop photo logging in Nibble. Noting that 75% of our survey respondents with mobile food logging experience used text-entry search logging, we made sure to include this as a fallback feature.

Our field participants used Nibble for 4-10 days. We analyze data for the first week. In the post survey, they rated food logging *easy to use* (100% strongly / agreed), *fast* (71%), and *useful* (86%). They logged 212 food items in total and averaged 2-10 items/day per participant. 69 logs (27%) were first attempted as photo-logging, however, only 63 meals (13%) were ultimately logged with photos, with the remaining using fallback methods. This suggests 46% acceptability of the food recognition with food logging. We interpret this as a proxy for recognition accuracy. While we do not know the specific validation accuracy of the FoodAI recognition model, it is based on state-of-the-art models (51-79% accuracy [4, 27]). This demonstrates an appreciable loss in food recognition accuracy when running in the field with a wider diversity of foods, compared to testing with a curated validation dataset or in the lab.

The lack of usage of photo-logging is also reflected in the lack of trust expressed in our field study participants: “[I] would rather have NO photo recognition, or very good one. If it’s half-baked, it will slow down the entire app” (P5). Only 29% of field participants agreed that the photo recognition was accurate. This validates the finding in [29] of some users expecting perfect accuracy for photo recognition with crowdsourcing. Instead, they would revert to search logging: “[I] usually choose search logging because [I] forget to take [the] photo. Phone isn’t on me when I’m eating meals” (P4). “I use [search] when the photo logging fails” (P5). Search Logging was the most popular logging method (85%), followed by Photo Logging (13%) and Custom Logging (1%).

Interestingly, some participants had some misconceptions about what was needed to support accurate food recognition. For example, P1 found it too tedious to properly frame food photos for recognition: “When you take pic you have to aim camera. Searching just type in can already” (P1). With the large training set of 1000 images per food item, such precise photography is unnecessary.

To further encourage logging, instead of requiring users to use Nibble to take photos to log their food, with their permission, we could automatically scan their phone’s image gallery, and detect camera usage or social media image uploads to identify and analyze any food images. In contrast, we note that as we seek to make food logging very convenient and seamless, this may hinder the user’s *mindfulness* of their eating which can aid their reflection about their food choices [8, 10].

Coverage and Accuracy: Beyond Database Reliability

As with [6, 9, 21], we found database reliability to be a barrier to food logging; users could sometimes not find some dishes through Nibble. A minority (2/7; 29%) of our field participants agreed that food-nutrition database was complete. This issue of *coverage* affects the completeness of the nutrition database, but also whether the food item was part of the image training dataset, i.e., as a class label.

Even with inclusion into the dataset, the food item may be misrecognized and thus lead to an issue in *accuracy*. We ameliorated this by displaying a list of top potential matches, but this should not be excessively long. Field participants were divided on the *perceived accuracy* of the food-nutrition database (29% disagreed, 29% agreed), and split regarding being able to find their eaten foods (57% agreed). “I have to take the photo and hope that the right recognition comes out” (P3). With an initial focus on local foods, international cuisine was not recognizable: “*exotic food such as ramen, were impossible to log correctly*” (P4). Also, although food recognition can provide educational value to tourists, this puts further demands on the food dataset/database coverage in foreign countries: “I was in *Kuching [Malaysia]. Zichar restaurants sell non-local food*” (P1).

Another difficulty field participants faced was trying to recognize “mixed” or heterogeneous foods in a single dish. “*Things like mixed vegetable rice, economical bee hoon that had multiple varying components*” (P4). “*Mixed Vegetable Rice, Mixed Western food are hard to log. Western food has lamb, coleslaw, hash brown. It takes me 6 minutes to log foods one by one. I’m not really a person to take photos of food. My food becomes cold after taking photos of all the part.*” (P5). The prevalence of “mixed” foods demonstrates the need for classifiers that can segment images to detect and identify multiple foods in individual images [26, 27].

Therefore, the barrier of database reliability should combine both nutrition database coverage and food image model accuracy.

Immediacy of Feedback

Our online survey respondents wanted fast feedback: 33% preferring <1 minute, 10% preferring immediate. Nibble is able to recognize the food image and provide nutrition feedback in <1 sec, and 100% of our field participants agreed that its logging was fast.

Near-instantaneous post-logging feedback can allow for useful nutritional knowledge about the food that the user is about to eat. It does not suffer from recall bias and feedback delay issues in traditional food diaries [14, 39], but may be too late to impact the user’s immediate eating behavior: “I don’t really care about after-logging feedback because I have already eaten the meal, and my next meal will be many hours later, so I would have forgotten the feedback.” (P7). Therefore, it may be more salient to provide feedback

or recommendations at the pre-purchasing decision point. Feedback from a previous meal can be replayed when the next meal time is predicted using historical/contextual cues.

Some participants also preferred getting summary insights faster. We presented daily and weekly diet report interfaces to participants in the usability lab study and found that even though they found weekly reports useful, they preferred the daily summaries: *“I would like more explicit immediate feedback on what to do, like if I had too many calories in the middle of the day, it should give me a warning to cut back on my calories for my next meal. ... [The] weekly report helps on a cognitive level, but it’s difficult to persuade me on the behavioral level because [of] delayed feedback”* (P4). Daily, immediate, actionable feedback could serve as facilitator type triggers [11] to prompt and guide the user to perform the behavior immediately.

Scalability of Nutrition Database and Photo Datasets

While promising, much research into computer vision to automatically recognize foods have been mostly limited to 100-200 food dishes (e.g., [4, 26, 27]). In reality, there are many diverse foods in a given community, especially in cosmopolitan cities. Even in a small country, Singapore, there can be a large diversity of foods: Wikipedia catalogs almost 300 local foods [36], and the government’s Health Promotion Board curates the nutrition of 3531 food items [16]. One can only fathom how many orders of magnitude more of unique foods a food logging app may have to characterize for a user population.

Typical CNN-based object recognition trains models on 1000 clean images of each item [4, 23, 27]. Collecting and filtering images for only 100 foods will require 100,000 images; this is tedious for an individual or small team and typically done via crowdsourcing (e.g., [31]). Furthermore, even as we build a training dataset to support a high variety of food, this reduces the accuracy of the CNN model because of having too many classification classes. One potential remedy is to organize the foods into fewer categories or cuisine types and use a cascade of models, or use contextual features to limit the foods to recognize (e.g., using location to constrain to certain restaurants [3, 27]).

Localization of Food

Food recognition datasets have mainly been based on western food dishes [4, 27], so this omits many Asian foods. It is important to localize the food image dataset to the location and cuisine culture of the user. Matsuda et al. have trained a classifier on Japanese foods [26], while our dataset is trained on Singaporean foods [12].

Furthermore, *communal* eating is common in Asian and other ethnic cultures [13]. Food would be presented at the center of a table for sharing with family portion sizes. P3 described challenges in discriminating what one has eaten from the full shared meal: *“When I am having “Zi Char” [home-cooked meals with multiple dishes] with my family or having... these are hard to log. I’m a bit lazy to log all,*

especially when the dish is hard to find, I don’t log. ... Some dishes I just took a few bites. Those [portion sizes] were a bit hard to estimate.”

Localization of Food Expertise

Automated or semi-automated food logging relies on human intelligence at some point in the data processing. Crowdsourcing methods employ human expertise at logging time, while CNN models leverage on human labeling when creating and curating the training dataset. In all cases, being able to recognize foods depends on the worker’s or user’s familiarity with the cuisine.

Crowdsourcing (commonly Amazon Mechanical Turk) typically has workers based in the United States. Methods to leverage this workforce to recognize ethnic or regional foods will not work due to the lack of cultural familiarity. For example, *Laksa* may be misinterpreted as *curry* or *Mee Goreng* as a tomato-sauce pasta. One potential remedy is to use computer vision to recognize cuisine type and assign to crowdworkers from a specific geography.

Peer rating can be made suitable with a global user base by limiting to the user’s local community who are familiar with her cuisine. However, this method suffers from low user engagement of even as low as 2.6% [17].

Expert feedback: many manual food recognition apps use expert dietitians who are on staff or freelancing (e.g., [15, 22, 33]). This is expensive since registered or accredited dietitians have to be on call. Furthermore, while dietetics is a common profession in the US (100k members in the Academy of Nutrition and Dietetics [1]), there is a scarcity of practitioners in some countries (in Singapore: 52 dietitians and 18 nutritionists in the SNDA [34, 35]). Additionally, nutrition science knowledge is generally consistent across countries, but there are slight differences in treatment method (e.g., AND Nutrition Care Process [18] vs. BDA Nutrition and Dietetic Process [14]). Moreover, the provision of actionable recommendations will vary by country and culture [14]. Therefore, using a US-based dietitian freelancer will not be ideal for global consumers.

CONCLUSION AND FUTURE WORK

We conducted formative investigations into the feasibility, usability and perception of fast, automatic recognition in photo-based food logging. We found that users appreciated the immediate feedback made possible with automated recognition. We demonstrated the value of providing search logging as a fallback method in light of poor recognition accuracy. We discussed challenges in deploying automated food logging in an Asian context.

Even as we seek to make food logging less tedious, most users may not be motivated to log their foods as a quantified-self goal [5]. Instead they may just have the goal to eat more healthily, where food logging is just part of the process. In addition to convenient food logging, a more compelling and actionable diet intervention app should provide *recommendations* of healthy foods [45].

REFERENCES

1. Academy of Nutrition and Dietetics. Advertising and Sponsorship. <http://www.eatrightpro.org/resources/about-us/advertising-and-sponsorship>. Retrieved 21 September 2016.
2. Frank Bentley and Konrad Tollmar. 2013. The power of mobile notifications to increase wellbeing logging behavior. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1095-1098. ACM.
3. Vinay Bettadapura, Edison Thomaz, Aman Parnami, Gregory D. Abowd, and Irfan Essa. 2015. Leveraging context to support automated food recognition in restaurants. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, Hawaii, USA, pages 580–587.
4. L. Bossard, M. Guillaumin, and L. Van Gool. Food-101: Mining discriminative components with random forests. In *ECCV 2014*.
5. Eun Kyoung Choe, Nicole B. Lee, Bongshin Lee, Wanda Pratt, and Julie A. Kientz. 2014. Understanding quantified-selfers' practices in collecting and exploring personal data. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems (CHI '14)*. ACM, New York, NY, USA, 1143-1152.
6. Sunny Consolvo. 2012. Designing for Healthy Lifestyles: Design Considerations for Mobile Technologies to Encourage Consumer Health and Wellness. *Foundations and Trends in Human-Computer Interaction*, 6(3–4), 167-315.
7. Sunny Consolvo and Walker, M. 2003. Using the experience sampling method to evaluate ubicomp applications. *IEEE Pervasive Computing*, 2(2), pp.24-31.
8. Felicia Cordeiro, Elizabeth Bales, Erin Cherry, and James Fogarty. 2015. Rethinking the Mobile Food Journal: Exploring Opportunities for Lightweight Photo-Based Capture. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 3207-3216.
9. Felicia Cordeiro, Daniel A. Epstein, Edison Thomaz, Elizabeth Bales, Arvind K. Jagannathan, Gregory D. Abowd, and James Fogarty. 2015. Barriers and Negative Nudges: Exploring Challenges in Food Journaling. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1159-1162.
10. Daniel A. Epstein, Felicia Cordeiro, James Fogarty, Gary Hsieh, and Sean A. Munson. 2016. Crumbs: Lightweight Daily Food Challenges to Promote Engagement and Mindfulness. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 5632-5644.
11. BJ Fogg. 2009. Creating persuasive technologies: an eight-step design process. In *Proceedings of the 4th International Conference on Persuasive Technology (Persuasive '09)*. ACM, New York, NY, USA, , Article 44 , 6 pages..
12. FoodAI. <http://foodai.org>. Retrieved 21 September 2016.
13. Robin Fox. 2003. Food and eating: an anthropological perspective. *Social issues research center*.
14. Joan Gandy. 2014. *Manual of Dietetic Practice*. John Wiley & Sons.
15. GlycoLeap, diabetes management app by Holmusk. <https://glycoleap.com/>. Retrieved 21 September 20, 2016.
16. Health Promotion Board. Energy and Nutrition Composition of Food. Singapore. <http://focos.hpb.gov.sg/eservices/ENCF/>. Retrieved 21 September 2016.
17. Helander E, Kaipainen K, Korhonen I, Wansink B. Factors related to sustained use of a free mobile app for dietary self-monitoring with photography and peer feedback: retrospective cohort study. *Journal of medical Internet research*. 2014;16(4):e109.
18. Mahan, L. Kathleen, and Janice L. Raymond. *Krause's food & the nutrition care process*. Elsevier Health Sciences, 2016.
19. Leswing Kif. 2015. Under Armour buys MyFitnessPal for \$475M. <https://gigaom.com/2015/02/04/under-armour-buys-myfitnesspal-for-475-million/> 4 February 2015. Retrieved 21 September 2016.
20. Yoojung Kim, Sookyoung Ji, Hyunjeong Lee, Jeong-Whun Kim, Sooyoung Yoo, and Joongseek Lee. 2016. "My Doctor is Keeping an Eye on Me!": Exploring the Clinical Applicability of a Mobile Food Logger. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 5620-5631.
21. K. Kitamura, C. de Silva, T. Yamasaki, and K. Aizawa. 2010. Image processing based approach to food balance analysis for personal food logging. In *Proceedings of the 2010 IEEE International Conference on Multimedia and Expo: ICME '10*, pp. 625–630.
22. Kato-Lin, Yi-Chin, 22hanshu Abhishek, Julie S. Downs, and Rema Padman. "Food for Thought: The Impact of m-Health Enabled Interventions on Eating Behavior." *Available at SSRN 2736792* (2016).

23. Alex Krizhevsky, Ilya Sutskever and Geoffrey E. Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*.
24. Jennifer Mankoff, Gary Hsieh, Ho Chak Hung, Sharon Lee, and Elizabeth Nitao. 2002. Using Low-Cost Sensing to Support Nutritional Awareness. In *Proceedings of the 4th international conference on Ubiquitous Computing (UbiComp '02)*, Gaetano Borriello and Lars Erik Holmquist (Eds.). Springer-Verlag, London, UK, UK, 371-376.
25. Martin, C. K., Han, H., Coulon, S. M., Allen, H. R., Champagne, C. M., & Anton, S. D. (2008, 07). A novel method to remotely measure food intake of free-living individuals in real time: The remote food photography method. *BJN British Journal of Nutrition*, 101(03), 446. doi:10.1017/s0007114508027438
26. Y. Matsuda, H. Hoashi, and K. Yanai. Recognition of Multiple-Food images by detecting candidate regions. In *ICME*, pages 25–30, July 2012.
27. Myers, A., Johnston, N., Rathod, V., Korattikara, A., Gorban, A., Silberman, N., . . . Murphy, K. (2015, 12). Im2Calories: Towards an Automated Mobile Vision Food Diary. *2015 IEEE International Conference on Computer Vision (ICCV)*. doi:10.1109/iccv.2015.146
28. MyFitnessPal. 2014. Announcing 75 Million MyFitnessPal Users! <http://blog.myfitnesspal.com/announcing-75-million-myfitnesspal-users/>. 9 December 2014. Retrieved 21 September 2016.
29. Noronha, J., Hysen, E., Zhang, H., & Gajos, K. Z. (2011). Platemate. *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology - UIST '11*. doi:10.1145/2047196.2047198
30. Nutritionix food database. <http://nutritionix.com>. Retrieved 21 September 2016.
31. Mashfiqui Rabbi, Jean Costa, Fabian Okeke, Max Schachere, Mi Zhang, and Tanzeem Choudhury. 2015. An intelligent crowd-worker selection approach for reliable content labeling of food images. In *Proceedings of the conference on Wireless Health (WH '15)*. ACM, New York, NY, USA, , Article 9 , 8 pages.
32. Rahman, T., Adams, A.T., Zhang, M., Cherry, E., Zhou, B., Peng, H., and Choudhury, T. BodyBeat: A Mobile System for Sensing Non-Speech Body Sounds. *MobiSys 2014*, 2-13.
33. Rise.us nutrition coach mobile app. <https://www.rise.us/>. Retrieved 21 September 2016.
34. Singapore Nutrition and Dietetics Association. Accredited Dietitians. http://www.snda.org.sg/index.php?r=accreditations/accredited_dietitians. Retrieved 21 September 2016.
35. Singapore Nutrition and Dietetics Association. Accredited Nutritionists. http://www.snda.org.sg/index.php?r=accreditations/accredited_nutritionists. Retrieved 21 September 2016.
36. Singaporean Cuisine. *Wikipedia*. https://en.wikipedia.org/wiki/Singaporean_cuisine. Retrieved 21 September 2016.
37. Edison Thomaz. 2016. Automatic Eating Detection in Real-World Settings with Commodity Sensing. Ph.D. Dissertation. School of Interactive Computing, Georgia Tech.
38. C. C. Tsai, G. Lee, F. Raab, G. J. Norman, T. Sohn, W. Griswold, and K. Patrick, "Usability and feasibility of pmeb: A mobile phone application for monitoring real time caloric balance," *Mobile Networks and Applications*, vol. 12, no. 2–3, pp. 173–184, 2007.
39. Walter C. Willett. 1998. *Nutritional Epidemiology*, 2nd ed. Oxford University Press.
40. Hope S. Warshaw and Karmeen Kulkarni. 2001. *Complete guide to carb counting*. American Diabetes Association.
41. World Health Organization. 2003. Diet, nutrition and the prevention of chronic diseases. WHO Technical Report Series, No. 916 (TRS 916). http://apps.who.int/iris/bitstream/10665/42665/1/WHO_TRS_916.pdf. Retrieved 21 September 2016.
42. World Health Organization. 2015. Q&A on the carcinogenicity of the consumption of red meat and processed meat. http://www.iarc.fr/en/media-centre/iarcnews/pdf/Monographs-Q&A_Vol114.pdf. Retrieved 21 September 2016.
43. World Health Organization. 2016. Obesity and overweight. Fact Sheet. June 2016. <http://www.who.int/mediacentre/factsheets/fs311/en/>. Retrieved 21 September 2016.
44. World Health Organization. 2016. Global report on diabetes. http://apps.who.int/iris/bitstream/10665/204871/1/9789241565257_eng.pdf?ua=1. Retrieved 21 September 2016.
45. Longqi Yang, Cheng-Kang Hsieh, Hongjian Yang, Nicola Dell, Serge Belongie, and Deborah Estrin. "Yum-me: Personalized Healthy Meal Recommender System." *arXiv preprint arXiv:1605.07722* (2016).