Trade-off between Automation and Accuracy in Mobile Photo Recognition Food Logging

Brian Y. Lim, Xinni Chng, Shengdong Zhao
National University of Singapore, Singapore
brianlim@comp.nus.edu.sg, xinni@xinni.co, zhaosd@comp.nus.edu.sg

ABSTRACT
Food logging can help users understand their food choices and encourage healthier eating habits. However, current apps still pose many usability challenges, including tedious manual text entry of food names. Recently, advances in computer vision and deep learning are enabling automatic food recognition for instant and convenient logging. However, as a nascent technology, this suffers from inaccuracy, which may lead to poor adoption or misuse. We investigated the trade-off between accuracy and convenience of automatic photo recognition in comparison to manual search logging. Specifically, we have developed a mobile app prototype that integrates both photo recognition and search logging capabilities, and conducted formative investigations on the usability and usage of automatic photo recognition in food logging in a series of studies: online requirements survey, usability lab study, and 1-week field trial in an Asian country. Participants were interested in convenient, automatic photo logging, but dominantly used manual search logging due to a lack of data coverage and accuracy. We identified reasons for poor accuracy and highlight complications in using inaccurate automatic photo logging. 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 diseases and cancers [41]. For example, consuming too many high-fat and energy dense (high caloric) foods with sedentary lifestyle can cause obesity [43]. The "risk of colorectal cancer could increase by 17% for every 100-gram portion of red meat eaten daily" [42]. Diabetes can be managed through weight loss [44] or carbohydrate counting [40]. Increasing user awareness of 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., [18, 27]). However, food logging apps continue to have many barriers [8], 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. However, this is still a nascent technology with issues in data coverage and accuracy. Will users prefer the convenience of automatic recognition over more predictable and accurate search text entry? In this work, we investigated the trade-off between the accuracy and convenience of automatic photo recognition for food logging with multiple surveys and a formative field trial.

Our contributions are: (1) the development of Nibble, a mobile app with capabilities in both fast, automatic recognition in photo-based food logging and search text logging 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 that participants were willing to try automatic photo logging for its speed and convenience, but preferred manual search logging for its accuracy and reliability. We identify various reasons for inaccuracy and highlight issues that may arise when fully dependent on inaccurate automatic logging. With a deployment in an Asian city, we also discuss implications for localizing mobile food logging for the cultural context.

RELATED WORK
While traditional paper diaries have been recommended by dietitians [13], 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 [8]. 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 human effort. This is performed through expert feedback, peer rating,

Chinese CHI '17, June 7-9, 2017, Guangzhou, China.

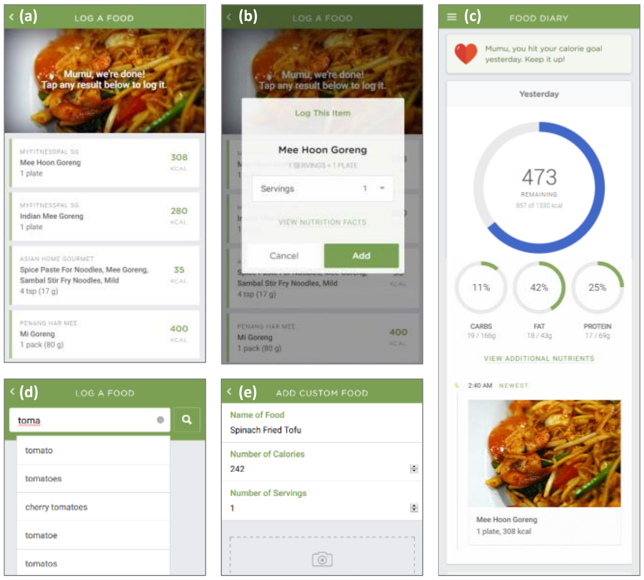

Figure 1. Screenshots of the Nibble app showing how a user logs food: first taking a picture via the smartphone camera (not shown), which gets recognized via the FoodAI API and returns (a) a list of candidate food names in under 1 sec. After choosing, the user (b) specifies the portion size and (c) gets 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.

crowdsourcing, or self-reflection. Professional dietitians can provide credible and accurate expert feedback to users and improve adherence [19, 21, 23], 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) [16], 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 (USD1.40/photo) and labeling is relatively slow (M=94 minutes) [28]. Cordiero et al. provide an interesting alternative where the user reflects on her own diet using her food photos [7], but this suffers from selection bias in users, and require unsustainable reflection effort. Epstein et al. extended this work to support mindfulness in a lightweight mobile app which only requires logging one meal per day [9]. 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 [23], chewing sounds [30] 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, 11, 25, 26]). 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 [7, 20] or mobile apps but with delayed feedback [16, 28]. In this work, our user study provides a formative investigation on the user preference and attitudes towards automatic photo logging in comparison with manual search logging.

NIBBLE MOBILE APP PROTOTYPE

We implemented Nibble, a mobile web app for photo food logging with two primary modes for logging: automated food recognition and manual search entry. 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 FoodAI application programming interface (API) [11] which trained a Convolutional Neural Network (CNN) [22] on the top 100 local foods in Singapore. However, in real-world usage, we did not expect the 100 foods to comprehensively cover foods that users may eat; our focus was to study the impact of inaccuracy in automatic logging. 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 [15] 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* [5] to help motivate users towards and objective. Reminders can mitigate the inconvenience of having to remember to log before eating the food at each meal. Goals can help promote sustained logging.

METHOD

We conducted formative investigations into the usability and usage of automated food recognition in a series of studies. We were primarily interested in evaluating the **trade-off between accuracy and convenience** of automated logging. 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 mockup on a laptop and conducted a scenario-driven usability lab study (5 participants). We took findings from this study to refine and

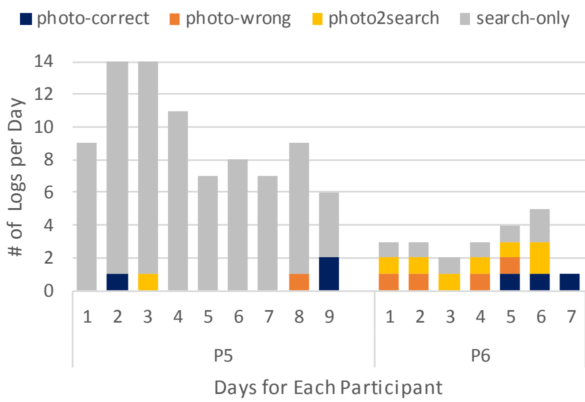

Figure 2. Different logging types that field participants chose for each logging event. Participants could have chosen to (1) only log a dish with automatic photo recognition, which may be correctly or wrongly recognized, (2) initially attempted photo logging, but switched to search, or (3) only used search and avoided taking photos. P5 mostly used search logging and logged many items, while P6 prioritized on photo logging.

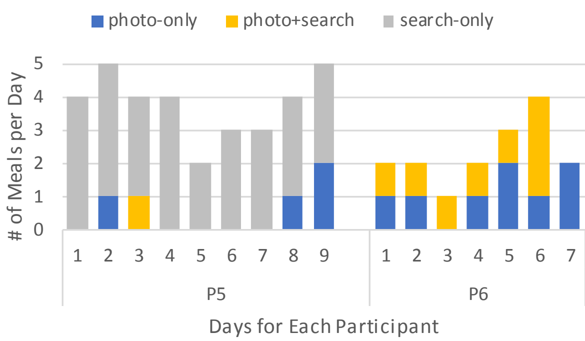

Figure 3. Different logging types used for each meal in a day. P5 logged the most number of meals, but mostly did not use photo logging. P6 tried photo logging in all meals logged.

add features to reduce usability confounds. We then deployed the Nibble mobile app prototype in a 1-week field study (6 university students) to evaluate user acceptance and usage of photo logging with real-time automated food recognition in the wild. To minimize risk to patients suffering from chronic disease, we targeted healthy individuals in this formative study. We instrumented Nibble for interaction logging, conducted pre/post-study interviews and surveys. All survey questions were asked on a 5-point Likert scale (≤–1: Disagree, ≥1: Agree).

FINDINGS AND DISCUSSIONS

We summarize the findings learned from designing and developing Nibble, and our three user studies. We found that participants were interested in convenient, automatic logging, but dominantly use manual search logging due to a lack of data coverage and accuracy. We identify reasons for poor accuracy and raise complications in using inaccurate automatic photo logging.

While we did not have access to the training dataset to know the specific validation accuracy of the FoodAI recognition model, we evaluated the API against 300 images scraped online across the 100 Asian food items represented in the API. We measured an accuracy of 96% for Top-1 labels. With larger validation datasets, on state-of-the art models have achieved (51-79% Top-1 accuracy [4, 26]). Therefore, the API is very accurate for food dishes on which it was trained.

Positive Interest in Automatic Photo Logging

From our online survey, we found that 75% of our survey respondents with mobile food logging experience used text-entry search logging. Although photographing food is popular on mobile social media apps [3], 70% of all survey respondents reported that they never take photos of their meals. Nevertheless, 73% of all respondents were willing to try photo food logging, giving promising results to develop photo logging in Nibble. With evidence indicating user interest in automated photo logging, we conducted our field study to evaluate how users used it in comparison to manual search logging.

Dominant use of Reliable Manual Logging

Our field participants used Nibble for 5-9 days. They logged 242 food items in total and averaged 2-10 items/day per participant. When logging food items, participants could first choose either to log using a *photo* or with *search*. If the photo log is not satisfactory, they could ultimately perform a search log. We consider this a *photo attempt*. 33 logs (14%) were recorded as photo-logs, 30 (12%) as photo attempts which were ultimately recorded as search-logs, and 179 logs (74%) were recorded only using search. The choice for different logging types is more diverse for each meal that the participants logged. We defined consecutive food logs to be of the same meal if they were within 10 minutes of each other. Participants logged 112 meals in total, 1-5 (Median=2) meals per day with 1-8 (Median=2) logs per meal. With multiple logs per meal, participants logged 30 meals (27%) using only photo-logging, 23 meals (21%) with both photo and search logging, and 59 meals (53%) using only search logging.

The lack of usage of photo-logging reflected the lack of trust in its accuracy: "*[I] would rather have NO photo recognition, or very good one. If it's half-baked, it will slow down the entire app*" (P5). P5 is a particularly notable participant whose logs were mostly based on search only (see Figure 2 and Figure 3). Only two field participants agreed that the photo recognition was accurate. 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).

Coverage and Inaccuracy of Automated Photo Logging

As expected with [5, 8, 20], we found database reliability to be a barrier to food logging. 5/6 of our field participants disagreed Nibble's food-nutrition database was complete.

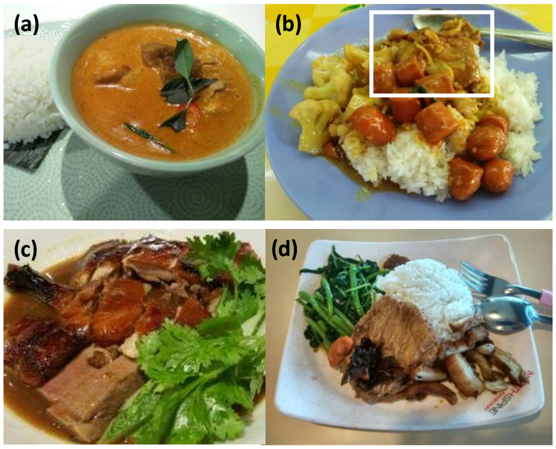

Figure 4. Examples of photos that field participants took indicating how recognition could be correct or wrong: (a) dish correctly recognized as Curry Chicken, (b) mixed rice dish partially correctly recognized as Curry Chicken, though it could have curry chicken (white box added for clarity), (c) dish correctly recognized as Roasted Duck, (d) mixed rice dish incorrectly recognized as Duck Rice.

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 *accuracy* issues in identification. Only two field participants agreed that their food photos were accurately recognized. This led to uncertainty in the outcome of the photo logging, e.g.: "*I have to take the photo and hope that the right recognition comes out*" (P3).

We analyzed the logged photos and found that 21/33 (64% accuracy) photos were correctly classified (Top-1 label). Nibble showed multiple alternative labels for each photo recognition. However, our analysis found that in our field trial, if the Top-1 label was incorrect, then the true label was also not in the Top-5. This is because the recognizer was never trained on the food dishes Therefore, for our results, the Top-5 accuracy was no better than that for Top-1. If we account for photo attempt logs (photo2search) as indicating misrecognitions, then 21/63 (33%) of photos are correctly classified. These are lower accuracies in comparison to evaluations with curated image datasets, per our validation study (96% accuracy) and published state-of the-art benchmarks (51-79% accuracy [4, 26]). This demonstrates an appreciable loss in food recognition accuracy when running in the field with a wider diversity of foods and user behavior, compared to testing with a curated validation dataset or in the lab.

Reasons for Inaccurate Automated Photo Logging

With an initial focus on local foods, international cuisine was not recognizable: "*exotic food such as ramen, were impossible to log correctly*" (P4); some participants (P1 and P5) also logged many Western dishes which were not covered in the image training dataset. 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 that 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 parts.*" (P5). Figure 4 demonstrates some examples of how "mixed rice" dishes can be wrongly classified as other dishes. The prevalence of "mixed" foods demonstrates the need for classifiers that can segment images to detect and identify multiple foods in individual images [25, 26].

Complications using Inaccurate Automated Logging

Interestingly, some participants found that photo-based logging can be inconvenient. 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). By analyzing his images, we found that P1 typically ate mixed foods (e.g., mixed rice, dim sum, western food with lamb chop and pasta), but attempted to isolate each food item in each photo. Supporting "fine grained classification" with image segmentation [26] can allow multiple food items to be recognized in one photo to avoid such tedious photography.

While inaccurate photo recognition may lead some participants (e.g., P5) to prefer manual search logging instead, some participants (e.g., P6, see Figure 2 and Figure 3) attempted to use photo logging often. P6 attempted photo logging for all meals and ultimately kept 7 (50%) of his logs as photo logs (no search). However, our analysis found that four (57%) of his photo logs were wrongly or only partially correctly classified (e.g., Figure 4b). Therefore, the convenience of automated logging may lead to more inaccurate data being logged instead of the user only entering high quality text inputs. This raises a challenge to identifying automatically logged data that the user accepts because they seem "good enough" to the lay end-user but may not be as scientifically or clinically valid (e.g., to a professional dietitian).

CHALLENGES FOR AUTOMATIC FOOD LOGGING

Our results demonstrate the continued need to improve accuracy in automatic photo logging to help user adoption of more convenient logging. We have reiterated issues in food database coverage and classifier accuracy. However, we have also raised issues in usability and overly trusting automated logging that is somewhat accurate. Next, we discuss challenges in deploying food logging in Asia.

Scalability of Nutrition Database and Photo Datasets

Much research into computer vision to automatically recognize foods have been mostly limited to 100-200 food dishes (e.g., [4, 25, 26]). However, in the field, 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 [15]. Typical CNN-based object recognition trains models on 1000 clean images of each item [4, 22, 26]. 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., [30]). 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, 26]).

Localization of Food

Food recognition datasets have mainly been based on western food dishes [4, 26], so this omits many Asian foods. It is important to localize the food image dataset to the location and cuisine culture of the user. Recently, there have been classifiers trained on Japanese [25] and Chinese [5] foods, while our dataset is trained on Singaporean foods [11].

Furthermore, *communal* eating is common in Asian and other ethnic cultures [12]. 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 using Amazon Mechanical Turk) typically has workers based in the United States. Methods to leverage this workforce to recognize ethnic or regional foods may 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% [16].

Expert feedback: many manual food recognition apps use expert dietitians who are on staff or freelancing (e.g., [14, 21, 32]). 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 [33, 35]). Additionally, nutrition science knowledge is generally consistent across countries, but there are slight differences in treatment method (e.g., AND Nutrition Care Process [17] vs. BDA Nutrition and Dietetic Process [13]). Moreover, the provision of actionable recommendations will vary by country and culture [13]. Therefore, using a US-based dietitian freelancer will not be ideal for global consumers.

CONCLUSION AND FUTURE WORK

We conducted formative investigations into the usage and trade-offs in using convenient but inaccurate, automatic photo logging and manual search logging for food journaling. We found that participants were willing to try and occasionally used automatic photo logging for its speed and convenience, but preferred manual search logging for its accuracy and reliability. We identified various reasons for inaccurate recognition and highlighted issues that may arise when depending on inaccurate automatic logging. We discussed challenges in deploying automated food logging in an Asian context.

For future work, we intend to address several of the challenges and deploy a more robust solution for food logging on patients suffering from specific chronic diseases or health risks. To improve photo recognition for foods, we plan to support the recognition of "mixed rice" dishes by performing image segmentation [26] or by recognizing ingredients (e.g., [5]). Moreover, as users photograph foods specific to their habits, this introduces new dishes which we can add to the training dataset to improve the image recognition model through active learning [33].

REFERENCES

1. Academy of Nutrition and Dietetics. Advertising and Sponsorship. http://www.eatrightpro.org/resources/about-us/advertising-and-sponsorship. Retreived 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. Chen, J., and Ngo, C. W. (2016, October). Deep-based ingredient recognition for cooking recipe retrieval. In *Proceedings of the 2016 ACM on Multimedia Conference* (pp. 32-41). ACM.
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. 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.
8. 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.
9. 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.
10. 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..
11. FoodAI. http://foodai.org. Retrieved 21 September 2016.
12. Robin Fox. 2003. Food and eating: an anthropological perspective. *Social issues research center*.
13. Joan Gandy. 2014. *Manual of Dietetic Practice*. John Wiley & Sons.
14. GlycoLeap, diabetes management app by Holmusk. https://glycoleap.com/. Retrieved 21 September 20, 2016.
15. Health Promotion Board. Energy and Nutrition Composition of Food. Singapore. http://focos.hpb.gov.sg/eservices/ENCF/. Retrieved 21 September 2016.
16. 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.
17. Mahan, L. Kathleen, and Janice L. Raymond. *Krause's food & the nutrition care process*. Elsevier Health Sciences, 2016.
18. 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.
19. 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.
20. 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.
21. Kato-Lin, Yi-Chin, Vibhanshu 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).
22. Alex Krizhevsky, Ilya Sutskever and Geoffrey E. Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*.
23. 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.
24. 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
25. Y. Matsuda, H. Hoashi, and K. Yanai. Recognition of Multiple-Food images by detecting candidate regions. In *ICME*, pages 25–30, July 2012.

26. 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
27. MyFitnessPal. 2014. Announcing 75 Million MyFitnessPal Users! http://blog.myfitnesspal.com/announcing-75-million-myfitnesspal-users/. 9 December 2014. Retreived 21 September 2016.
28. 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
29. Nutritionix food database. http://nutritionix.com. Retrieved 21 September 2016.
30. 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.
31. 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.
32. Rise.us nutrition coach mobile app. https://www.rise.us/. Retrieved 21 September 2016.
33. Settles, B. (2010). Active learning literature survey. *University of Wisconsin, Madison*, *52*(55-66), 11.
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 feasability 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.