

Towards a Pervasive Kitchen Infrastructure for Measuring Cooking Competence

J. Wagner, A. van Halteren, J. Hoonhout
Philips Research
Eindhoven, The Netherlands
juergen.wagner@philips.com

T. Ploetz, C. Pham, P. Moynihan, D. Jackson, C. Ladha
K. Ladha, P. Olivier
Culture Lab, Computing Science, Newcastle University, UK
t.ploetz@newcastle.ac.uk

Abstract—Research has demonstrated that a lack of cooking competence can be a significant barrier to healthier eating. We present two studies from which we develop a set of requirements for a pervasive sensor infrastructure that will enable our Ambient Kitchen environment to measure cooking competence in an unobtrusive manner. From the first study we derive key characteristics and potentially measurable aspects of cooking competence. This study also led to the specification and design of a pervasive sensor infrastructure comprising of a set of kitchen utensils equipped with custom-made wireless accelerometers. The second study reports our initial findings from the use of the sensor infrastructure and demonstrates its potential to measure key indicators of cooking competence. Our studies provide initial evidence that cooking competence can be measured automatically using our proposed pervasive kitchen infrastructure.

Keywords—activity recognition, cooking competence, pervasive computing; accelerometer

I. INTRODUCTION

The World Health Organization (WHO) identified obesity as one of the greatest health challenges of the 21st century. Since the 1980s the number of obese people has tripled and this figure is rising at an alarming rate [2]. About a third of the European adult population has a body mass index (BMI) above 30 (the definition of *obese*). Obesity increases the likelihood of medical conditions such as cardiovascular diseases, various types of cancer, diabetes and osteoarthritis, and treating these conditions places significant demands on national healthcare systems. Obesity is mostly attributed to an imbalance between energy intake and energy expenditure. In theory, it should be an easy task to counteract obesity, by either increasing energy expenditure or by reducing energy intake. But achieving permanent changes in diet is well known to be challenging, despite people's awareness of the shortfalls in their diet and the impact this may have on their health and well-being [3]. There are a number of identifiable barriers to successful dietary change and many of these barriers also lead to increased consumption of less healthy energy dense foods, such as convenience meals, fast food or snacks [4]. Food choice is affected by a number of factors [5]. Healthier food and meal choice is in part related to people's levels of cooking competence and confidence [6]. However, many people lack the confidence or skill to plan and prepare meals, to undertake certain meal preparation activities, or to follow recipe

instructions. People with a low level of cooking competence also tend to consume more convenience meals [7].

Our overall goal is to design and deploy digital technologies that can improve people's cooking skills and thereby impact positively on the nutritional quality of their diet. However, to achieve this we must first characterize cooking competence and then develop reliable measures for it, based on sensed activities in the kitchen. We therefore envision a pervasive kitchen infrastructure that can automatically sense and reason about people's food preparation and cooking activities. In order to create such an infrastructure we must first establish the nature of cooking competence and thereby establish its requirements. To this end we have conducted both an exploratory observational study, in people's homes, and as a result we designed and deployed a sensor infrastructure that has the potential to evaluate an individual's cooking competence. We evaluate this infrastructure and provide an overview of the challenges to be overcome to realize our vision of a pervasive kitchen that can truly support users in their food preparation and cooking activities.

II. RELATED WORK

A large number of systems (both research prototypes and consumer products) aim to provide support to users to analyze their diet and make healthier food choices. While traditional digital interventions for healthier eating have relied on users accessing nutritional support and advice at a desktop, there has been a growing recognition that such support needs to be more *situated*, that is, located in the time and place that the actual activities of meal planning and preparation occur. Svensson et al. [8, 9] aimed to develop a system that supports users when they are *uncertain* as to what meal to prepare. They specifically addressed two scenarios, one in which people are uncertain as to what meal to prepare, and another where a meal needs to be planned using a limited set of ingredients. While online databases are a potentially valuable resource for inspiring food and meal choice, the large range of choice can be overwhelming. To support people in these decision-making processes, Svensson et al. created an application that allows users to navigate through a database for a healthy and balanced meal in a simplified manner. Technologies have also been applied to the problem of increasing awareness as to the current nutritional characteristics of an individual's diet (usually with a view to supporting dietary change). Again, such systems can provide appropriate personalized support and guidance for

users. Mankoff et al. [3] aimed to increase a user's understanding of the *amount of nutrients* they actually consume. To assess the current state of a user's diet, shopping receipts were scanned and analyzed, and a shopping list of suggested healthier food alternatives was produced. Ransley et al. [10] also used shopping receipts to analyze the nutritional quality of someone's diet.

A small number of prototype systems aim to help users during meal preparation by both recognizing their activities and guiding them through their tasks. In general this has been with a view to increasing time efficiency and supporting decisions as to what to prepare. Chi et al. [11] developed a nutrition aware kitchen that gives feedback, during the actual preparation of a meal, as to its nutritional value. RFID sensors and scales embedded in the kitchen counter detected and quantified ingredients. Nutritional information was presented to users in an attempt to motivate and help them to make healthier choices. Amft et al. [12] analyzed nutritional behavior directly during actual food intake. By using a range of worn sensor technologies to detect both chewing sounds and swallowing motions, they were able to report accuracies as high as 80% for the identification of food groups consumed. Chang et al. [13] created a dining table using load cells, weight tracking algorithms and RFID technology to measure food intake. While this was a significantly less obtrusive solution (compared with Amft et al's [12]) it required an infrastructure embedded in the user's environment and more configurations (i.e. to define the dishes in the first place). Captured pictures of a meal and storing them in a journal have also been used to promote and support dietary awareness and reflection [14, 15].

A number of projects have sought to provide support for inexperienced cooks through the use of a sensing infrastructure in the kitchen that monitors the users and guides them through meal preparation and cooking activities. Nakauchi et al. [16] developed a system that combines robotic and activity recognition systems. Meal preparation activities are recognized and mapped onto a recipe, and based on this information a robot predicts the next steps and guides a user with both speech and gestures (e.g. points towards the cupboard with the sugar). Related systems include Ju et al.'s [17] rich multimedia experience guides to support users preparing new dishes, and Martins et al.'s [18] cooking guide based on a spoken dialog system which predicts the next activity based on a cooking ontology. Being inexperienced or disorganized can also result in inefficient behavior during meal preparation. To address such problems, systems such as CounterIntelligence [19] have explored the use of various surfaces in the kitchen to support information access for the cook. Bonnani et al. [19] identified the fridge as a particularly valuable resource, and claimed that people's efficiency in meal preparation and meal planning could be significantly improved if they were more aware of what ingredients are in their fridge (e.g. when doing groceries reminding people of what they need to buy). They applied this idea to a range of other kitchen appliances, for example, using displays on cupboards and drawers that indicate to a user where different kitchen utensils were located.

Another approach is to provide appropriate advice in a timely manner, by monitoring people's meal preparation activities. This requires the detection of as many activities as

accurately as possible. In our previous work we embedded sensors, such as accelerometers, into kitchen utensils [20, 21]. The recorded acceleration data can be translated into activities using classifiers trained using machine learning techniques on annotated data sets. An alternative approach is to mount a sensor on the user's wrist, or use a wrist worn RFID reader that can sense passive RFID tags embedded in the environment [22]. While body worn sensor approaches are utensil independent there is a subsequent encumbrance for the user. In addition to detecting the activity, Kranz et al. [34] detected different ingredients. Using a combination of torque sensors in the knife, microphones in the environment, and load cells in the chopping board, they reported being able to distinguish between 6 different ingredients.

Distractions while preparing a meal can also influence the outcome and a user's general motivation to cook. Such distractions can be caused by other family members, phone calls, or unexpected events such as simple mistakes. Tran et al. [23] developed a system that records the user while preparing a meal. When the user wants to continue their meal preparation after being distracted, they can view the recording to see at which point they were distracted, and thus more easily restart their activity. Distractions can also occur while preparing multiple meals at the same time. Hamanda et al. [24] designed a system to address this challenge, claiming that the user can "cook several meals without failure".

Although our overview of related work has identified a number of digital interventions that aim to improve people's nutrition, the approach most of these systems adopt is to support the decision making process, by analyzing a user's diet and/or providing guidance and to help with meal preparation to inexperienced cooks. We can readily see that the problems these systems are mostly seeking to address can in part be linked to a user's cooking competence. Not knowing what to prepare can be the result of a person only having knowledge of a small number of different recipes, or lacking the knowledge as to which ingredients go together. People who regularly cook meals with fresh ingredients generally have a more balanced diet [25]. By being involved in the meal preparation a user is also more aware of the actual composition of the meal. On the other hand, people who consume convenience meals have to rely on the information on their meal's packaging, and have no real opportunity to modify its composition [26]. Furthermore, people with higher levels of cooking competence have a better understanding of the different chemical processes that occur during meal preparation, e.g. temperature treatment or reactions of different ingredients; knowing these processes enables a person to better plan their meal preparation activities. Experienced cooks are generally more time efficient and are certainly able to monitor multiple simultaneous processes.

Our observations about the differences between experienced and inexperienced cooks suggest that increasing an individual's cooking competence may also result in an improvement in the diet of those consuming the food the cooks prepare. Currently, a person's cooking competence is assessed using questionnaires [26]. The obvious shortcoming of such instruments is that they measure subjective self-assessments, and a subject might rate his or her competence incorrectly, for example, based on their opinion of preparing a limited range of

meals that involve only a small range of meal preparation activities. Indeed, the few existing cooking competence questionnaires have low face validity. In response to both this measurement problem, and the link between cooking competence and diet, our first aim is to develop a sensor infrastructure that allows the accurate and objective measurement of cooking competence.

III. OBSERVATIONAL STUDY

Similar to Hudson et al. [27], we chose an observational approach to explore the feasibility and requirements of a pervasive kitchen infrastructure for measuring cooking competence. Instead of inserting different sensors into a kitchen environment we used annotated video observations to simulate and explore possible sensor deployments. Observations were conducted in participants’ home kitchens to access an authentic record of people’s everyday meal preparation activities using their own utensils and appliances. Furthermore, a number of activities that were not directly related to meal preparation were also captured (such as cleaning, and responses to distractions). Another advantage of home-based observation was that our participants prepared meals in a familiar spatial environment in which they knew what appliances are available and where to find them. This allowed “normal” meal preparation strategies and routines to be followed, and enabled us to get a better insight into people’s cooking activities “in the wild”.

Nine participants (5 males and 4 females) were recruited from single-person households, two-person households, and families with young children. We anticipated that this diversity would lead to the observation of a number of different cooking styles, strategies, and levels of competence and confidence. Participants all claimed to be the person responsible for meal planning and meal preparation in their household. The task was to prepare an “every day” meal while being filmed. We placed no constraints on participants as to what meal to prepare, other than that it had to be a self-cooked fresh meal (i.e. ready-made meals were excluded). Designing an appropriate camera setup was a significant challenge, as cameras had to be positioned so as not to interfere with the participant’s meal preparation. Consequently, depending on the kitchen layout, up to 2 cameras were mounted at places that were unobtrusive to the participant (during the cooking process). That is, cameras had to be out of range of moving elements, such as cupboards and drawers. Although the use of small camcorders and gorilla tripods simplified this task, filming in each kitchen required a unique setup and strategy. Figure 1 shows a sample camera setup. Although the cameras were deployed to capture as many activities as possible, in no case was it possible to capture every activity, as we were significantly constrained by the kitchen layout, kitchen size, and available camera mount locations – as a result each video contained a number of “blind spots”. While preparing their meals, participants were mostly alone in the kitchen, although in some cases family members passed through (these sections were removed from the video footage to protect the family members’ privacy). In general, the participants reported feeling comfortable being recorded while cooking; some participants claimed that they only felt monitored for the first few minutes, but after a while forgot

about the cameras and undertook their meal preparation tasks unaffected.



Figure 1. A participant’s kitchen and the camera setup.

The video of the meal preparation sessions was annotated using the ELAN [28] annotation tool. Due to the fact that there were clear differences between the participants, their cooking environment, and the prepared meals, the first step was to define an annotation schema that would be general enough to apply to the different settings and meal preparation behaviors.

A. Annotation Schema

Our initial challenge was to design an annotation schema that would allow us to meaningfully compare different users, in different environments, preparing different meals. There are very few previous examples of approaches to annotating meal preparation activities; and existing annotation schemas are mostly tailored to either specific environments and/or recipes. For example, Spriggs et al. [29] introduced an annotation schema defined to test and evaluate a very specific set of activity recognition algorithms. In developing our schema we had to address three principal challenges. First, we had to define the number of levels of activities that we wished to annotate, i.e. from simply noting presence in the kitchen, to specific motion primitives. Our approach involves two levels: level 1 describes when an object is being *moved* and level 2 describes more specific activities (e.g. as can be found in a cook book). The second challenge was to identify (or define) exactly when an activity starts and ends. Here we took advantage of the two levels of activities, that is, as soon as a utensil is being moved, it is *moving*. This can frame more descriptive activities (e.g. *cutting*). Figure 2 visualizes this concept.

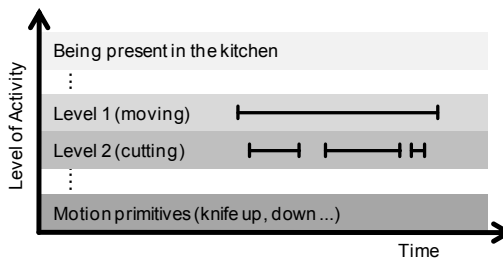
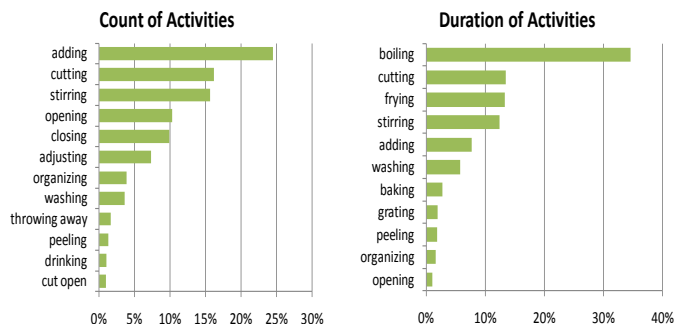


Figure 2. Schematic of different levels of activities and their annotation.


The last challenge was to determine exactly which activities we wanted to annotate. In practice, the selection of these was an iterative process; on the basis of pilot annotation sessions, we found that new activities had to be defined, and related activities grouped together. This was also true for many of the

utensils that were used to perform many of the actions. We decided to treat activities and utensils separately from each other, to allow us to analyze which activities can be performed with which utensils. The final annotation schema consists of 21 activities and 34 utensils and objects. The set of activities contains both *actively* performed activities (e.g. cutting, stirring, grating) as well as *passively* performed activities (e.g. boiling, frying); but also activities that are not related to meal preparation, but that occurred while cooking.

The results of applying this annotation schema can be simplified and used to describe a complex cooking process in a structured manner. The annotation schema (and subsequent results) can be used to get insights as to both the performed activities as well as the utensils used. Figure 3 shows a summary of the most common activities and utensils, in both frequency and duration.



(a) Frequency of the 12 most common activities as a percentage of total number of recorded activities. (b) The 12 activities performed for the longest time as a percentage of the total preparation time.



(c) Frequency of use of the 12 most used utensils/objects as percentage of the total utensil usage. (d) The 12 utensils/objects that were used for the longest time as a percentage of total preparation time.

Figure 3. Most common (frequency & duration) activities & utensils/objects.

B. Cooking Behavior Insights

Although the activity *adding* (to add an ingredient from one container to another) is typically a very short activity it is a key meal preparation activity. Our results for the utensils/objects used show that *ingredients* are manipulated almost twice as often as other objects (both occurrence and duration). Of the other activities, *cutting* and *stirring* were the most performed. These represent about one third of all activities when preparing a meal. They not only occur frequently, but their duration constitutes 30% of the overall cooking process. Importantly, *cutting* and *stirring* activities were performed by every participant. This is something that was also reflected in the use of the utensils; *knives* and *spoons* as well as *chopping boards*

and *pots* and *pans* are all used by participants. This simply confirms the intuitive notion that meal preparation is mostly comprised of the *manipulation* and *adding of ingredients*.

Another pair of activities that occurred very frequently was the *opening* and *closing* of kitchen furniture. These activities were recorded at almost all stages of meal preparation, so long as ingredients were being added or utensils used. *Opening* and *closing* occurred frequently, but they were typically very brief actions. Another unanticipated finding was the frequency of the *washing* activity (of ingredients, utensils or the participant’s hands). This activity most commonly occurred at the beginning or end of a meal preparation task, indicating its potential to be used to frame cooking activities. Furthermore, our observations indicate that there are a range of different meal preparation strategies that depend on personal preference, competence, cooking space, and even cultural background. For example, some participants first gathered together all ingredients for their meal, separated and chopped them and as a very last step cooked them. Others followed a “just-in-time” strategy, where ingredients are gathered and prepared based on the situation. We also observed that the kitchen was not exclusively used to cook; while preparing a meal many other activities took place that were unrelated to meal preparation (e.g. *drinking* or *cleaning*).

The observational study provided us with a number of valuable insights into people’s cooking behavior in their actual home environments. The annotation schema we developed was used as a tool to compare how different meals were prepared (by different participants) but also as a source of information by reference to which we can assess the appropriateness (e.g. in terms of coverage) of a sensor infrastructure for recognizing activities and thereby measuring cooking competence.

C. Cooking Competence

Definitions, or defining qualities, of cooking competence are colored by the context in which a meal is prepared. A chef’s emphasis on taste and presentation would most likely differ from the concerns of a high school domestic science teacher, or a working mother. In our definition, we focus on the activities of meal preparation, rather than the sensory qualities of the resulting meal, and based on the literature and our own observations we have identified a total of seven different factors that characterize cooking competence [1, 30, 31, 5].

(a) *Use of utensils and appliances*: selection and manipulation of utensils and appliances during meal preparation.

This requires the evaluation of both the appropriateness of utensil selection, and a person’s skill in using a utensil [32]. Several different characteristics relate to the level of skill with which a utensil is used (e.g. speed, smoothness, self-similarity of motion, etc.) and these can potentially be measured directly during meal preparation.

(b) *Multitasking*: the ability to perform multiple meal preparation activities simultaneously.

A complex meal, in which a number of non-trivial constituents are separately prepared and then combined, requires a person to undertake multiple activities in parallel. The measurement of multitasking will therefore require the identification of the number of activities taking

place simultaneously. While multitasking can potentially be measured directly during meal preparation, the number of tasks that can be conducted in parallel is clearly a function of the recipe.

- (c) *Monitoring and adaptation*: the ability to accurately perceive the state of the meal being prepared, and to react appropriately in different situations (e.g. novel or stressful situations).

Monitoring and adaptation are particularly important when modifying an existing plan for the preparation of a meal (and they are therefore related to planning). To evaluate this aspect of cooking competence it is necessary to measure meal preparation activities and map them to corresponding recipes. This enables the identification and evaluation of the modification. For example, changing or adapting a meal often involves the addition or replacement of at least one ingredient. The measurement of monitoring and adaptation is therefore likely to require the detection of ingredients.

- (d) *Planning*: the ability to schedule activities prior to, and during, the preparation of a meal.

Planning is required prior to the meal, when sourcing ingredients from the store, and when preparing the meal. To monitor how well a meal was planned, it would again be necessary to measure and track a person's meal preparation activities. Another aspect of planning, the acquisition of ingredients, can be measured by tracking the flow of ingredients in and out of the kitchen. Therefore the measurement of this aspect of planning skill is again likely to require the actual identification (and tracking) of ingredients.

- (e) *Reproduction of recipes*: the ability to reproduce recipes with a consistent quality without following instructions.

The measurement of a person's reproduction of a recipe requires the meal preparation activities to be mapped to the recipe instructions – this (again) implies a need for ingredient identification and tracking. Mapping a meal preparation session closely to a recipe must also take into account a user's meal preparation preferences and strategies. People with experience of preparing a certain recipe are likely to have incorporated adjustments to the constituent ingredients, the tasks, or to the order of the tasks, and such adjustments must also be taken into account.

- (f) *Cognitive skill*: both knowledge of and the ability to reason about basic cooking processes (e.g. the effects of temperature on ingredients and the effects of combining different ingredients).

The measurement of a person's cognitive skill, from direct observations of meal preparation activities, poses a significant challenge.

- (g) *Nutritional knowledge*: an understanding of the nutritional properties of foods, the impact of cooking processes on these properties, and the requirements of a healthy diet.

The degree to which nutritional knowledge is applied, for any individual meal, depends on a wide range of

contextual factors (e.g. the purpose and social context of a meal). Although such factors are unlikely to be inferable from meal preparation activities alone, cumulative evaluation (over relatively long periods of time) of meals planned, ingredients used, and activities performed, may provide a clear indication of a person's nutritional knowledge, or rather, their application of this knowledge.

These seven aspects of cooking competence, combined with the findings of our observational study described earlier, indicate that there is potential to measure a number of aspects of cooking competence from meal preparation activities alone. Our overarching goal is to integrate a sensor infrastructure within the Ambient Kitchen [1] that enables the measurement of cooking competence. As we have already explained in our characterization of the different components of cooking competence, although different measurement modalities may be necessary, a number of these components of cooking competence are (to a significant degree) measurable from a person's meal preparation activities alone. Based on this observation, we established the requirements for a pervasive kitchen infrastructure capable of measuring cooking competence and undertook an initial evaluation of its ability to recognize different meal preparation activities.

IV. SENSOR INFRASTRUCTURE

A. Requirements

There are several requirements for a sensor infrastructure that can be used to measure the aspects of cooking competence that we have identified. As described earlier, the majority of these aspects are closely linked to the user's meal preparation activities. This suggests the principal requirement of the infrastructure is that it enables the measurement and analysis of meal preparation activities. Such an infrastructure should measure activities unobtrusively; it should neither interfere with a person's normal meal preparation routines nor should it give users a feeling of "being monitored". For this reason we have rejected both body-worn and camera-based solutions. Instead, the sensors should be seamlessly deployed in the infrastructure and in particular, in the utensils. In the Ambient Kitchen [1], we have shown how different sensors can be embedded into a kitchen to observe users activities unobtrusively. Another requirement, a commercial imperative, is that the system should be self-contained and deployable within a regular (existing) kitchen. To ensure a reasonable coverage of cooking activities we deployed sensors in utensils that are frequently used during meal preparation. The selection of these utensils is based on the findings in the observational study. The last requirement relates to the cost of the infrastructure, which, as ever, should be minimal.

Previous research in activity recognition has used a large number of different classes of sensor, including accelerometers, radio frequency identification (RFID), video cameras, microphones, load cells, torque sensors, temperature sensors and motion sensors. In our own work accelerometers have proved to be a robust means by which to measure and classify activities for which a knife was used. Pham et al. [20] demonstrated how a 3-axis accelerometer could be used to distinguish between 11 different meal preparation activities performed with a knife. An accelerometer is a relatively cheap

robust sensor, which due to its small size can be easily integrated into everyday objects such as kitchen utensils.

B. Design

Encouraged by our previous experience of 3-axis accelerometer-based activity recognition, we sought to both expand and further explore its application. Tapia et al. [33] have shown that a single sensing technology can be used to recognize a wide variety of activities. We integrated our sensor infrastructure within a number of kitchen utensils: 5 knives, 2 pots and lids, a frying pan, a peeler, a grater, a measuring cup, a sieve, a spoon, a spatula, a ladle, a whisk and a chopping board. Figure 4 (top) shows the complete set of utensils. Each of these utensils was modified to embed a 3-axis accelerometer inside a waterproof case, which also allowed the utensils to be washed. The handles of the original utensils were replaced with new handles designed to (as closely as possible resemble the original design). The new handles were printed in ABS using an FDM rapid prototyping technology. Figure 4 (bottom) shows an example of a (new) modified handle with an embedded sensor. With the accelerometer embedded into a redesigned handle we can unobtrusively (and wirelessly) track utensil use, and use the tracked sensor data to identify aspects of cooking competence. In our evaluation study specific tasks in the overall cooking process were monitored, and the annotation schema developed for the observational study was used to characterize cooking performance.



Figure 4. Overview of the sensor embedded kitchen utensils (top); kitchen knife with an embedded wireless accelerometer (bottom).

C. Preliminary Study

To evaluate the sensor infrastructure a small user study was conducted with the aim of testing the utensils (usability and robustness) and how accurately activity recognition can be performed. We used a more controlled setup than that used in our first study. Five participants, all of whom cook regularly (i.e. most days of the week) were invited to the Experience Lab at Philips Research, Eindhoven. We asked participants to prepare three dishes in the Experience Lab’s kitchen environment using the sensor infrastructure. For this study we used 12 utensils with an embedded accelerometer and 8 regular utensils. In addition, the complete meal preparation activity

was captured using multiple cameras. The recipes were selected to ensure that a wide range of different meal preparation activities were performed, and the recipe instructions encouraged the use of all 12 of the utensils that contained embedded sensors.

After an introduction to both the study and the kitchen, the participants were given the recipes for the three dishes. During the study no restrictions were placed on participants as to how they executed the instructions and in which order they prepared the meals. In addition, they were not required to perform the tasks with any specified utensils (i.e. they were free to choose what they felt was the most appropriate utensils for the tasks at hand). Although the main aim of the study was to evaluate the sensor infrastructure and not the user’s meal preparation, we wanted the participants to engage in the activity in a manner that was as natural as possible. Although each participant had to prepare the same dishes, and had access to the same infrastructure, significant differences in the preparation time were observed ($m=40\text{min}$, $stdv=7\text{min}$). This time difference was most likely due to variations in participant familiarity with the recipes and preparation methods, and their sequencing of the preparation steps (i.e. planning). This observation may indicate that there were significant differences between the levels of cooking competence of our participants (although we did not conduct any separate evaluation of cooking competence).

The physical robustness of the sensor infrastructure is a mundane quality but one that it was important to demonstrate. In particular, it was very important that the remodeled handles did not deteriorate as a result of either contact with water, the high temperatures they were exposed to, or the mechanical stresses they underwent; not only to ensure the user experience but also the safety of our participants. During the experiment all participants were asked to stop mid-task and describe any negative impressions of the utensils.

The main characteristic of the infrastructure that we evaluated was its ability to identify different meal preparation activities. In the past we have evaluated the accuracy of detecting 11 different activities that can be performed with a knife [20]. Our infrastructure needs to be able to recognize meal preparation activities associated with a large number, and wide range, of utensils if we are in future to be able to measure aspects of cooking competence such as multitasking, planning and the reproduction of recipes. In this preliminary study 12 kitchen utensils were selected to cover the majority of activities. While we used a questionnaire to assess users’ perceptions of the infrastructure, we used the annotated sensor data files (annotated by reference to the videoed session) to train activity classifiers, which we evaluated using a 4-fold cross validation procedure on the 5 cooking sessions of our 5 participants.

1) *User perceptions of the sensor infrastructure:* When the participants finished preparing the three dishes we evaluated their experience of the sensor infrastructure. We used a 5-point Likert scale to measure users’ impressions as to the design, functionality and safety of the utensils. From our observations we noted that all participants used the utensils naturally and our mid- and post-task inquiries failed to reveal any concerns

as to their durability. Most of the participants also washed some of the utensils in the sink (at no detriment to the sensors). Utensils used for stirring were rated less positively due to the fact that some of our modified designs (e.g. the whisk) were not well balanced. For example, the additional weight of the sensor in the whisk’s handle occasionally caused the whisk to fall out of the mixing bowl.

2) *Partial data*: Figure 5 shows an example of the error inherent in our pervasive kitchen infrastructure; the limited deployment of sensors meant that it was not possible to detect every activity. Sensors were only embedded in frequently used utensils (and ones in which it was possible to “hide” the sensor, usually inside a handle). Figure 5 shows the average use of measured activities during the entire meal preparation time for Participant 1. It can be seen that between minutes 5 and 12 there are a number of relatively long periods during which no activities were detected. When we compare these results with our record of the order in which the dishes were prepared, it can be seen that the user prepared the salad first. Here the participant mostly used one or more of the 8 utensils that had not been instrumented (e.g. fork, ceramic bowl). A similar pattern was observed for the other participants. Nevertheless, these preliminary results show that the infrastructure allowed the detection of a large number of the activities that occur during meal preparation.

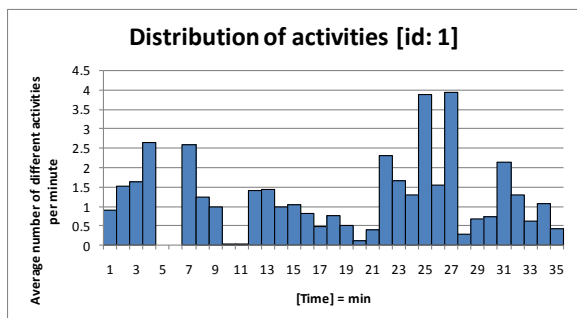


Figure 5 Distribution of measured activities (per minute) for Participant 1.

3) *Activity recognition*: To study the feasibility of a sensor-based approach to assessing cooking competence we employed a standard, instance-based classification approach to automatically analyze the acceleration data recorded while participants prepared their dishes using our instrumented kitchen utensils. Raw sensor data was processed using a sliding window procedure to extract 64-sample frames (50% overlap, 100Hz sampling frequency) from the accelerometer data captured using the sensor-enabled utensils. For every analysis window (frame) a set of statistical features was extracted. In particular, mean, standard deviation, energy, entropy and correlation features were computed for each frame for x-, y-, z-, pitch- and roll-acceleration, resulting in 23-element vectors. These vectors were then classified using a Nearest Neighbor (NN) classifier, which modeled the 34 activities performed using the particular utensils. These 34 activities consist of a set of activities which are a precondition

for meal preparation but do not constitute the actual activity. Activities unrelated to meal preparation (i.e. no action and unknown action) can be performed using every utensil but all the remaining activities related directly to meal preparation. Descriptive activities (such as cutting or stirring) are listed in Table I. The procedure was optimized for classification accuracy within real-time constraints.

The effectiveness of the activity recognition procedure was evaluated using a 4-fold cross validation procedure on 5 cooking sessions for our 5 different participants, where each prepared three different recipes. Given the set of extracted frames (>1 million in total), four sub-sets were extracted, thereby carefully balancing the distributions of the activities within the particular folds. For every experiment four folds were used for training the classifiers, which were then evaluated on the remaining fold. Completely permuting the selection of folds for training and evaluation, this procedure was repeated four times and achieved classification accuracies were averaged. The overall classification accuracies were mostly well beyond 80%. This accuracy includes the context describing actions (no action and unknown action). Table I reports the results for subject-independent activity recognition experiments for the different activities and the utensil it has been performed with. The overall average accuracy of 75% demonstrates the effectiveness of the proposed approach in a relatively unconstrained cooking scenario.

TABLE I RESULTS FOR SUBJECT-INDEPENDENT ACTIVITY RECOGNITION EXPERIMENTS FOR FIVE DIFFERENT UTENSILS

Activity	Utensil	Accuracy	Frames
Cutting	Chef knife	73.09%	996
	Small knife	64.63%	294
Peeling	Peeler	78.31%	332
	Chef knife	66.15%	130
	Small knife	73.91%	92
Grating	Grater	81.07%	1294
Whisking	Whisk	70.72%	461
Stirring	Spoon	75.44%	2896
Average		75.32%	

V. CONCLUSION AND FUTURE WORK

There are numerous different approaches to supporting people to eat more healthily. Our approach is intended to influence people’s diet by providing a “tool” to prepare healthier meals. In our approach this “tool” is cooking competence, the basic ability that enables people to translate the intention of “wanting to change” into actual action. We envision a system that monitors and analyzes a user’s cooking activities and translates this information into a measurement of competences. The measured cooking competence will then afford personalized support and recommendations as well as situated guidance. Such a system has the potential to overcome barriers that hinder people in changing their diet. The intended system can also help to reduce the perceived complexity of a recipe, and increase people’s cooking competence. We have presented an initial step towards a system that provides support and guidance based on a user’s cooking competence. Having identified the requirements of the sensor infrastructure we conducted an exploratory observational study of people’s

cooking behavior in their home kitchen. This involved the development of an annotation schema that allowed us to compare different users, environments and recipes. The outcome of this study provided, on the one hand, an indication that it may be possible to measure certain aspects of cooking competence. On the other hand it provided an insight as to what activities are performed in the kitchen and which utensils are used. These insights, combined with other requirements (i.e. the need for unobtrusive, low cost, and deployable monitoring), provided the foundation for our sensor infrastructure. In a pilot user study we evaluated the sensor infrastructure with respect to: how it was perceived by users; the extent to which it allows us to measure meal preparation; and how accurately we can recognize the constituent activities. The outcome of the second study showed that our sensor infrastructure is perceived as unobtrusive and that activity recognition is likely to be a valuable component of a cooking competence measurement system.

ACKNOWLEDGMENT

Funded by the EU FP7 Marie Curie IAPP project *Balance@Home* and the RCUK Digital Economy Research Hub *Social Inclusion through the Digital Economy*.

REFERENCES

- [1] P. Olivier, G. Xu, A. Monk, and J. Hoey, "Ambient kitchen: designing situated services using a high fidelity prototyping environment," 2nd International Conference on Pervasive Technologies Related to Assistive Environments, 2009.
- [2] F. Branca, H. Nikogosian and T. Lobstein, "The challenge of obesity in the who european region and the strategies for response," World Health Organization, Copenhagen, 2007.
- [3] J. Mankoff, G. Hsieh, H. C. Hung, S. Lee and E. Nitao, "Using low-cost sensing to support nutritional awareness, Ubicomp 2002," Springer-Verlag, 2002, pp. 371-378.
- [4] R. Povey, M. Conner, P. Sparks, R. James and R. Shepherd, "Interpretations of healthy and unhealthy eating, and implications for dietary change," Health Education Research 13 no. 2, 1998, pp. 171-183.
- [5] T. Furst, M. Connors, C. A. Bisogno, J. Sobal and L. Winter Falk, "Food choice: A conceptual model of the process," Appetite 26, no. 3, 1996, pp. 247-266.
- [6] L. Soliah, J. Walter and D. Antosh, "Quantifying the impact of food preparation skills among college women," College Student Journal 40, no. 4, 2006, pp. 729-739.
- [7] J. Levy and G. Auld, "Cooking classes outperform cooking demonstrations for college sophomores," Journal of Nutrition Education and Behavior 36, no. 4, 2004, pp. 197-203.
- [8] M. Svensson, K. Höök and R. Öster, "Designing and evaluating kalas: A social navigation system for food recipes," ACM Transactions on Computer-Human Interaction (TOCHI) 12, no. 3, 2005, pp. 374-400.
- [9] M. Svensson, K. Höök, J. Laakolahti and A. Waern, "Social navigation of food recipes," Proceedings of the SIGCHI conference on Human factors in computing systems, 2001, pp. 341-348.
- [10] J. K. Ransley, J. K. Donnelly, T. N. Khara, H. Botham, H. Arnot, D. C. Greenwood and J. E. Cade, "The use of supermarket till receipts to determine the fat and energy intake in a UK population," Public Health Nutrition 4, no. 6, 2001, pp. 1279-1286.
- [11] P.-Y.P. Chi, J.-H. Chen, H.-H. Chu, and J.-L. Lo, "Enabling Calorie-Aware Cooking in a Smart Kitchen," PERSUASIVE '08: Proceedings of the 3rd international conference on Persuasive Technology, 2008, pp. 116-127.
- [12] O. Amft and G. Tröster, "On-body sensing solutions for automatic dietary monitoring," IEEE Pervasive Computing 8, 2009, pp. 62-70.
- [13] K.-H. C. S.-y. Liu, H.-H. Chu, J. Y.-j. Hsu, C. C.-H. Chen, T.-y. Lin, C.-Y. Chen and P. Huang, "The diet-aware dining table: Observing dietary behaviors over a tabletop surface," Pervasive, 2006.
- [14] B. Brown, M. Chetty, A. Grimes and E. Harmon, "Reflecting on health: A system for students to monitor diet and exercise," CHI '06 extended abstracts on human factors in computing systems, 2006.
- [15] J. Frost and B. K. Smith, "Visualizing health: Imagery in diabetes education," In Proceedings of the conference on Designing for user experiences, 2003.
- [16] Y. Nakauchi, T. Fukuda, K. Noguchi and T. Matsubara, "Intelligent kitchen: Cooking support by LCD and mobile robot with IC-labelled objects," Proceedings of Intelligent Robots and Systems, IEEE Press, 2005, pp. 1911-1916.
- [17] W. Ju, R. Hurwitz, T. Judd and B. Lee, "Counteractive: An interactive cookbook for the kitchen counter," CHI '01 extended abstracts on Human factors in computing systems, 2001.
- [18] F. Martins, J. P. Pardal, L. Franqueira, P. Arez and N. Mamede, "Starting to cook a tutoring dialogue system," IEEE/ACL Workshop on Spoken Language Technology, 2008.
- [19] L. Bonanni, C.-H. Lee and T. Selker, "Counter intelligence: augmented reality kitchen," Proceedings CHI 05, ACM, 2005.
- [20] C. Pham and P. Olivier, "Slice&Dice: Recognizing food preparation activities using embedded accelerometers," Ambient Intelligence, 2009, pp. 34-43.
- [21] T. Ploetz, P. Moynihan, C. Pham and P. Olivier, "Activity recognition and healthier food preparation," In: "Activity recognition in pervasive intelligent environments," Atlantis Press, 2010.
- [22] L. Atallaha and G.-Z. Yang, "The use of pervasive sensing for behavior profiling – a survey," Pervasive and Mobile Computing 5, no. 5, 2009, pp. 447-464.
- [23] Q. T. Tran, G. Calcaterra and E. D. Mynatt, "Cook's collage: Déjà vu display for a home kitchen," Proceedings of HOIT: Home-Oriented Informatics and Telematics 2005, pp. 15-32.
- [24] R. Hamada, J. Okabe, I. Ide, S. i. Satoh, S. Sakai and H. Tanaka, "Cooking Navi: Assistant for daily cooking in kitchen," MULTIMEDIA '05 Proceedings of the 13th annual ACM international conference on Multimedia 2005, pp. 371-374.
- [25] Food Standards Agency, "Eatwell plate", <http://www.eatwell.gov.uk/>, Surrey, 2011.
- [26] M. Caraher, P. Dixon, T. Lang and R. Carr-Hill, "The state of cooking in england: The relationship of cooking skills to food choice," British Food Journal 101, no. 8, 1998, pp. 590-609.
- [27] S. E. Hudson, J. Fogarty, C. G. Atkeson, D. Avrahami, J. Forlizzi, S. Kiesler, L. J. C. and J. Yang, "Predicting human interruptibility with sensors: A wizard of oz feasibility study," In Proceedings of the SIGCHI conference on Human factors in computing systems, ACM, 2003, pp. 257-264.
- [28] Language Archiving Technology, "Elan – linguistic annotator," 2010. <http://www.lat-mpi.eu/tools/elan/>
- [29] E. H. Spriggs, F. D. L. Torre and M. Hebert, "Temporal segmentation and activity classification from first-person sensing," Computer Vision and Pattern Recognition Workshop, 2009, pp. 17-24.
- [30] B. Sidenvall, M. Nydahl and C. Fjellström, "Managing food shopping and cooking: The experiences of older swedish women," Ageing and Society 21, no. 2 2001, pp. 151-168.
- [31] N. I. Larson, C. L. Perry, M. Story and D. Neumark-Sztainer, "Food preparation by young adults is associated with better diet quality," Journal of the American Dietetic Association 106, no. 12 2006, pp. 2001-2007.
- [32] N. Hammerla, T. Ploetz, P. Andras, P. Olivier, "Assessing Motor Performance with PCA." In Proc. Int. Workshop on Frontiers in Activity Recognition using Pervasive Sensing, 2011.
- [33] E. M. Tapia, S. S. Intille and K. Larson, Activity recognition in the home using simple and ubiquitous sensors, Pervasive, 2004, pp. 158-175.
- [34] M. Kranz, A. Schmidt, A. Maldonado, R. Bogdan Rusu, M. Beetz, B. Hoernler, and G. Rigoll, "Context-aware kitchen utilities". In Proceedings of the 1st international conference on Tangible and embedded interaction (TEI '07), pp. 213-214.