

The Diet-Aware Dining Table: Observing Dietary Behaviors over a Tabletop Surface

Keng-hao Chang^a, Shih-yen Liu^b, Hao-hua Chu^{a,c}, Jane Yung-jen Hsu^{a,c},
Cheryl Chen^d, Tung-yun Lin^c, Chieh-yu Chen^a, Polly Huang^e

Department of Computer Science and Information Engineering^a,
Department of Information Management^b,
Graduate Institute of Networking and Multimedia^c,
School and Graduate Institute of Nursing^d,
Department of Electrical Engineering^e,
National Taiwan University
r93018@csie.ntu.edu.tw, b90701219@ntu.edu.tw, {hchu, yjhsu}@csie.ntu.edu.tw;
cheryl.chen@ha.mc.ntu.edu.tw, {r94944009, r94054}@csie.ntu.edu.tw,
phuang@cc.ee.ntu.edu.tw

Abstract. We are what we eat. Our everyday food choices affect our long-term and short-term health. In the traditional health care, professionals assess and weigh each individual's dietary intake using intensive labor at high cost. In this paper, we design and implement a diet-aware dining table that can track what and how much we eat. To enable automated food tracking, the dining table is augmented with two layers of weighing and RFID sensor surfaces. We devise a weight-RFID matching algorithm to detect and distinguish how people eat. To validate our diet-aware dining table, we have performed experiments, including live dining scenarios (afternoon tea and Chinese-style dinner), multiple dining participants, and concurrent activities chosen randomly. Our experimental results have shown encouraging recognition accuracy, around 80%. We believe monitoring the dietary behaviors of individuals potentially contribute to diet-aware healthcare.

1 Introduction

Our dietary habits affect our health in many ways. Research [12] has confirmed that dietary habits are important factors for healthy living and have profound impacts on many chronic illnesses. The vast majority of the population has chronic illnesses [4] such as heart disease, diabetes, hypertension, dyslipidemia, and obesity. A recent Surgeon General Report indicated that approximately 300,000 U.S. deaths are associated with obesity and overweight each year. The total cost attributed to overweight and obesity amounts to \$117 billion in 2000. Proper dietary intake and related interventions are effective in ameliorating symptoms and improving health [5][12][13].

Nutritious dietary is one of the most accessible means for people to prevent illness and to promote well-being [5]. Unlike traditional healthcare in which professionals assess and weigh one's dietary intake and then develop a plan for behavioral changes, ubiquitous healthcare technologies provide an opportunity for individuals effortlessly

to quantify and acknowledge their dietary [5][6] intake. For example, at home patients face the cumbersome need to record everything they eat, a task which can take a minimum of 15-20 minutes per day [2]. Ubiquitous computing technologies provide a means for individuals to proactively monitor their intake and act upon it, leading to better food selection and more sensible eating.

This paper proposes a diet-aware dining table that automatically tracks what and how much each individual eats over the course of a meal. This is in accord with the *vision of disappearing computers* [15], where computing hardware (HW) & software (SW) are hidden into everyday object (i.e., dining table) and remain *invisible* to human users. There are *no* digital access devices (such as cell phones, PDAs, or PCs) needed in order for human users to interact with this digital dietary service. In comparison, traditional dietary tracking software requires human users to recall the amount of food consumed, and then manually enter the data. This is less precise due to mistakes in visual measurement and imperfect memory. More importantly, the traditional method requires explicit human effort to operate digital devices.

We have augmented a dining table with two layers of *sensor surfaces* underneath – the *RFID (Radio Frequency Identification) surface* and the *weighing surface*. By combining the RFID and weighing surfaces, our system can trace the complete *food movement path* from its tabletop container source to other containers, and eventually to the individual. To validate our diet-aware dining table, we have performed experiments, including live dining scenarios (afternoon tea and Chinese-style dinner), multiple dining participants, and concurrent activities chosen randomly. Our experimental results have shown encouraging recognition accuracy around 80%, which is as good as the 80% accuracy of the traditional dietary assessment methods [3].



Fig. 1. Typical Chinese dining table setting.

Our diet-aware dining table supports multiple people sharing a meal on the same dining table. Fig. 1 shows a typical meal setting for a Chinese family – the family members sit around a circular table with the main dishes placed in the center. Individual rice bowls and plates are arranged on the table periphery. Participants first use shared utensils to transfer food servings from the main dishes to their personal plates or rice bowls, and then eat from there. In this dining scenario, multiple table participants are *continuously* and *concurrently* engaging in food transferring and eating motions. This creates *multiple, concurrent person-object interactions* (objects are tabletop objects such as plates, bowls, etc.) from which a single table surface needs to

observe, track, and then infer high level interaction semantics. This is the *main technical challenge* addressed in this paper – how to design a sensor-embedded tabletop surface to track food consumption from each of many table participants.

The remainder of this paper is organized as follows. Section 2 states the design choices, assumptions, and limitations. Section 3 presents our design and implementation. Section 4 describes the experimental set-up and results. Section 5 describes the related work. Finally, Section 6 draws our conclusion and future work.

2 Design Choices, Assumptions, and Limitations

Although the ultimate design objective is to create a *restriction-free*, automated dietary-tracking system that can achieve both *high accuracy* and *precision*, this is a grand challenge requiring extensive future research efforts [9]. We acknowledge this fact, and consider our dietary-tracking system as an early effort to address this problem. Since our work is not yet a perfect solution, we need to state our assumptions, present our design rationale, and discuss our design limitations.

2.1 Why RFID and Weighing Surfaces?

Our diet-aware dining table tracks tabletop interactions such as transferring food among containers and eating food by an individual. To correctly infer individuals' dietary behaviors from their tabletop interactions, our system needs to track how much (weight) and what food items are involved in these interactions. To observe these interactions, a weighing surface and a RFID surface are embedded into an ordinary dining table. Assume that food items are correctly labeled by the RFID tags on food containers, the surface can then be used to identify these RFID-tagged containers. Furthermore, the RFID surface can obtain nutritional information such as calorie count by looking up a food label database indexed by RFID code.

This assumption raises a question as to who would perform the work of inputting the food information for the RFID tags into the database. Three possible scenarios apply: (1) prepared foods (e.g., microwave-ready) are purchased from supermarkets are heated and then placed on the dining table with their original containers and packages containing RFID tags. This is applicable to people who subscribe to a weight-loss dietary program; (2) when the food containers (dishes) are first placed on the dining table, the table explicitly asks users for the food contents through a natural, easy-to-input UI, such as speech interface; and (3) when food is prepared in the kitchen, the cooking person can input the food's content as the food is placed in a serving container.

The weighing surface is used to measure (1) the amount of food transferred across different tabletop containers, as servings of food are transferred between different tabletop containers, and (2) the amount of food consumed by an individual, as personal plates lose weight. More details on how the weight measurements are used to detect food transfer and food consumption events are described in Section 3.

2.2 Complex and Concurrent Interactions Involving Multiple Tabletop Objects.

In a typical family meal setting, there are multiple people dining together on a dining table, and table needs to track multiple, concurrent person-object interactions. In an afternoon tea scenario, if one person is pouring tea to a cup while another one is eating cake, it is impossible to use a single weighing surface to distinguish the amount of tea weight transfer to the cup vs. the amount of cake weight lost through a person's consumption. This scenario is shown in Fig. 2-(a). This is also called the *single-cell-concurrent-interactions problem* where it is impossible to distinguish multiple, concurrent person-object interactions over a single surface using the weight information from only one sensor¹. To address this problem, our solution is to divide the tabletop surface into multiple cells, shown in Fig. 2-(b). When the size of each cell is small enough, it is likely that each tabletop object occupies a different cell. Therefore, our solution uses multiple weighing sensors at different cells to distinguish the weight-change of the tea cup from the weight-change of the cake plate. This idea is generalized as follows: the larger the size of each weighing cell relative to the average size of objects, the higher the likelihood that multiple, concurrent person-object interactions can occur within the same cell, therefore the higher the probability of single-cell-concurrent interactions. To reduce this probability, we divide the weighing surface into cells of an appropriate size that just fit the average size of tabletop food containers, such as plates, bowls, etc.

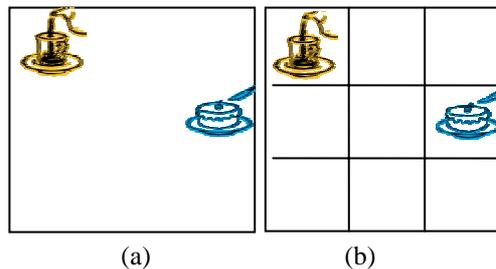


Fig. 2. Surface structure [This illustrates that a multi-cells surface (b) can track multiple person-object interactions whereas a single-cell surface (a) cannot.]

Where *single-cell-concurrent-interactions problem* still occurs, we introduce *common sense semantics* to discern the amount of weight-changes in these concurrent interactions. Consider the situation where a cup and a plate are placed at the same cell X at the same time. When a user pours tea from a tea pot to a cup (leading to weight increase at cell X), we can correctly infer the tea is transferred to a cup rather than to a plate by using common sense in normal dining behavior.

Also, relying only on a weighing surface (i.e., without RFID surface) is insufficient to identify tabletop objects. Distinguishing a tabletop object by its weight is difficult, given that the weights of food containers change over the course of a meal as people transfer food servings. Therefore, we augment the weighing surface with a passive

¹ In the Lancaster's approach [14], the scale is made up of four weighing sensors at four corners of a table.

RFID surface to help identify tabletop objects. Each cell contains a RFID antenna that can read the unique IDs from RFID-tagged objects on that cell.

2.3 Intelligent Surface vs. Intelligent Containers

Early in our design, we faced a fundamental design choice between embedding intelligence into the table or into the food containers. One advantage for choosing the intelligent food containers is that they do not have the *single-cell-concurrent-interactions problem*, because each food container can weight itself and detect its own weight-change events. However, the intelligent containers approach also has many disadvantages. First, it may result in high cost since every food container must have a weight scale and wireless networking module. Second, the smart food containers require battery installments and replacements, whereas the dining table is a piece of stationary furniture that can be plugged into a wall socket. The third disadvantage is that people may buy prepared food items from restaurants that have their own disposable packages and RFID tags. It is inconvenient to have people transfer the food into the intelligent containers every time, in contrast to the convenience of putting tagged packages directly on the intelligent table.

2.4 Assumptions

From the above discussion on design choices, our assumptions for our system in this paper are:

- The dining table, its RFID-tagged tabletop objects (food containers), and table participants form a closed rather an open system. That is, all food transfers can occur only among the tabletop objects and individual mouths. External objects and food sources are not allowed on the table.
- All dining participant have their personal containers (personal plates and cups) that are usually placed in front of their seating. They are used to identify each individual user.
- Food containers must be tagged with RFID tags. We assume that weight, nutrition, and ingredients of the food, as well as, the weight and owners of food containers are known a-priori.
- Tabletop objects are placed within each individual cell. No cross-cell objects are allowed.
- Dining participants avoid leaning their hands and elbows on the table.

3 Design and Implementation

Our system is consisted of HW & SW components. The HW component is made up of the RFID and weighing sensors embedded underneath the table surface shown in Fig. 3. The SW component is made up a rule-based system that aggregates, interprets, and infers tabletop dietary behaviors shown in Fig. 4. The HW component is described first, followed by the SW component.



Fig. 3. Embedded RFID and weighing table surfaces

3.1 Hardware Design and Implementation

Our current table prototype has a dining surface of $90 \times 90 \text{ cm}^2$, which is about the size of a small dining table. To detect multiple, concurrent person-object interactions on the tabletop surface, the tabletop surface is divided into a matrix of 3×3 cells, each with the size of $30 \times 30 \text{ cm}^2$, about the average size of food containers. Each cell contains a weighing sensor and a passive RFID antenna as shown in Fig. 3. The RFID reader is the i-scan MR100 made by Feig. The RFID antennas are connected to the RFID reader through a multiplexer. Each RFID antenna is positioned underneath the table surface such that it has an average, effective read range of 3 cm above the table surface. The weighing sensor is attached to a weight indicator with a resolution of 0.5 gram which can output weight readings through a serial port at a frequency of 8 samples per second.

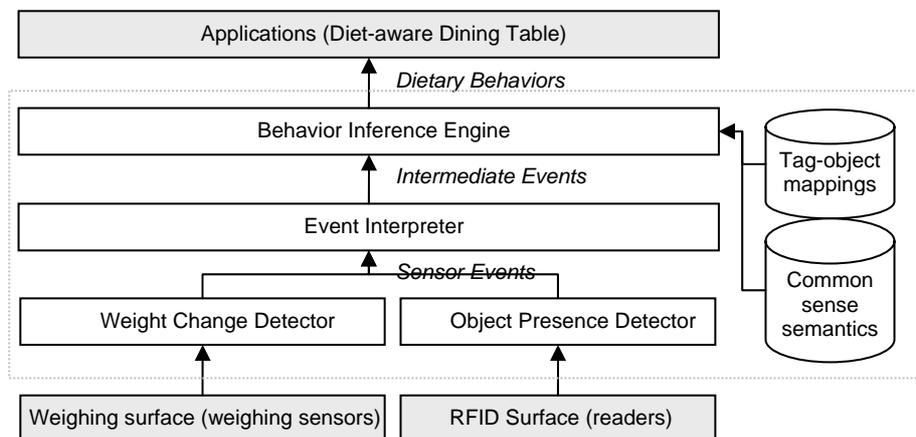


Fig. 4. System architecture

3.2 Software Design and Implementation

We have come up with a rule-based approach that applies our multi-cells weighing and RFID surfaces to detect multiple, concurrent person-object interactions. The system and inference rules are implemented in JESS rule engine [8].

The system architecture is based on a bottom-up event-triggered approach shown in Fig. 4. High level dietary behaviors, such as *pour-tea* and *eat-cake*, can be inferred by interpreting sensor *Cell-Weight* events and *RFID-Presence* events. We describe each software component as follows.

The *weight change detector* performs the following two tasks: (1) it aggregates weight samples collected from each of the 9 weighing sensors; (2) it reports *Cell-Weight* events when the weight has changed by filtering out noises in the stream of weight samples. The *object presence detector* performs similar tasks: (1) it continuously checks for presence and absence of RFID-tagged tabletop objects within each of the 9 RFID reader cells, and reports *RFID-Presence* events as long as the change happens.

The *event interpreter* interprets *intermediate events* shown in Table 1. The event interpreter builds internal states using sensor events from the weight change detector and the object presence detector, and then interprets *Weight-Change* events. Table 2 shows the rules to interpret events. For example, the *Weight-Change(Object_i, Δw)* event represents that the *Object_i*'s weight is changed by Δw , where $Object_i \in \{object\ on\ the\ table\}$ and $Cell_j \in \{cell_{[1-9]}\}$.

Table 1. Intermediate events, sensor events, and internal states

Intermediate Events	Descriptions
<i>Weight-Change(Object_i, Δw)</i>	<i>Object_i</i> 's weight is changed by Δw .
Sensor Events	Descriptions
<i>RFID-Presence(Object_i, Cell_j)</i>	<i>Cell_j</i> detects the presence of <i>Object_i</i> .
<i>Cell-Weight(w, Cell_j)</i>	<i>Cell_j</i> measures weight w .
Internal State	Descriptions
<i>Location(Object_i, Cell_j)</i>	<i>Object_i</i> locates on <i>Cell_j</i> .
<i>Weight(Object_i, w)</i>	<i>Object_i</i> has weight w .

Table 2. Rules for recognizing intermediate events

<p>Event Interpretation Rules</p> <ul style="list-style-type: none"> $Weight(Object_i, w_1) \cap Weight'(Object_i, w_2) \rightarrow Weight-Change(Object_i, w_2-w_1)$ <p>State Update Rules</p> <ul style="list-style-type: none"> $RFID-Presence(Object_i, Cell_j) \rightarrow Location(Object_i, Cell_j)$ $Location(Object_i, Cell_j) \cap Cell-Weight(w, Cell_j) \rightarrow Weight'(Object_i, w)$

The *behavior inference engine* infers *dietary behaviors* initiated by the user u shown in Table 3. Behavior inference engine is essentially the core of the system. It infers food transfer and eating behaviors over the table. In the real world scenarios, there are often different food items on the table, meaning that multiple food sources can be transferred to the same personal container. For example, the weight-increase to a cup may be contributed by pouring of coke, juice, or tea from different bottles and

pots. Moreover, given that there are multiple food transfer interactions happening concurrently, how does the system identify and differentiate the food being transferred from which food source container to which user's personal container?

3.2.1 Weight Matching Algorithm

To track a food movement path from the food's source containers to personal containers, we design a *weight matching algorithm*. The basic idea is to match a weight-decrease from one container to a complementary weight-increase from another container. This matching process can be thought as *a hop of food transfer* from the source food container in the center of the table, to the personal containers on the table periphery. This weight matching model is realized by maintaining a queue of recent *Weight-Change* events. When a new *Weight-Change* event is detected, our model applies a matching function to find a complementary *Weight-Change* event(s) in the waiting queue. A good match is found when the difference between the weight-decrease and the weight-increase pairs is smaller than a chosen *weight matching threshold* value (ϵ). This weight matching model is coded as rules in Table 3. For example, *Transfer(u, w, type)* means that a serving of cake with a weight w has been transferred from the share-plate containing food of *type* to the user u 's personal plate, where *type* is obtained from RFID mappings. This behavior event can be inferred by first observing a weight decrease Δw_1 (<0) in the share-plate, followed by a matching weight increase Δw_2 on the user u 's *Object_{i2}*. A match is found when $|\Delta w_1 + \Delta w_2| < \epsilon$. The *tag-object mappings* provide two relations: *Contains(Object, type)* shows *Object* contains food of the type (*type*), such as cake or tea, and *Owner(Object, u)* means the owner (u) of the *Object*. In addition, *Eat* is inferred if there is a weight-decrease in any personal container.

Table 3. Inference rules for dietary behaviors

Dietary behaviors	Behavior Inference Rules
<i>Transfer(u, w, type)</i>	$Weight\text{-}Change(Object_{i1}, \Delta w_1) \cap (\Delta w_1 < 0)$ $\cap Weight\text{-}Change(Object_{i2}, \Delta w_2) \cap (\Delta w_2 > 0)$ $\cap Contains(Object_{i1}, type) \cap Owner(Object_{i2}, u) \cap (\Delta w_1 + \Delta w_2 < \epsilon)$ $\rightarrow Transfer(u, \Delta w_2, type)$
<i>Eat(u, w, type)</i>	$Weight\text{-}Change(Object_{i1}, \Delta w) \cap (\Delta w < 0) \cap Contains(Object_{i1}, type)$ $\cap Owner(Object_{i1}, u) \rightarrow Eat(u, -\Delta w, type)$

In real world scenarios, there are special, complex interactions that require matching among three or more events. For examples, a person may pour tea from a tea pot to two cups within one continuous motion, or another person may transfer soup from a soup bowl to a personal bowl through multiple scoops. These two examples can be mapped to (1) the amount of one weight decrease matches with the sum of multiple weight increases, or (2) the amount of one weight increase matches with the sum of multiple weight decreases. To address this issue, the weight matching algorithm is extended to match more than two weight transfer events.

3.2.2 Common Sense Semantics

Although dividing the table into cells can reduce the probability of multiple objects on one cell, the situation mentioned in Section 2.2 may still happen. To address this situation, we add *common sense semantics* to extend the inference routines that can disambiguate the multi-objects on one cell problem. For example, if there are one *cup* and one *plate* on the same cell, and the user pours tea from the *pot* to the *cup*; the *Weight-Change* event of $\{cup, plate\}$ will be reported by the *Event Interpreter*. According to the common sense, tea should be poured into the cup rather than the plate. Therefore, the *behavior inference engine* matches the weight-decrease of the *pot* to weight-increase of $\{cup, plate\}$ and generates *Transfer(pot, cup, w)* behavior.

4 Experimental Set-up and Results

We have conducted several experiments to evaluate the accuracy of our dietary tracking table under different dining scenarios. The evaluation metric, *weight accuracy*, measures how well the system can correctly recognize the amount of weight from different food items consumed by the dining participants. It is determined by how well the system can correctly recognize the high-level dietary behaviors: specifically the food *transfer* event and *eat* event. Therefore, the intermediate evaluation metric, *behavior accuracy*, is listed as well. Note that both behaviors are associated with attributes defined in Table 3. The *transfer* event has three attributes (source object, destination object, weight), whereas the *eat* event also has three attributes (user, source object, weight). Correct event recognition is defined as the event's attributes, except the weight attribute, are correctly identified. Since the weight measurements have inherent sensor errors, they are evaluated separately. Specifically, the behavior accuracy is the number of behaviors recognized divided by the number of behaviors conducted by participants. The weight accuracy is the sum of measured weight divided by the sum of actual weight corresponding to dietary behaviors.

The experiments involve three participants. The first two participants are graduate students from our research team who are familiar with our system. The third participant is a graduate student from our department, who is not familiar with our system.

Dining scenarios (# participants, predefined vs. random activity sequences)

We have designed four different dining scenarios. The varying parameters are (1) the number of dining participants and (2) whether dietary behaviors are *predefined* or *random*. As the number of dining participants increases, we expect that they will generate higher number of non-overlapping and concurrent events. Predefined activities mean that the dining participants repeat some pre-arranged sequences of dietary steps which we expect in normal dietary behaviors. The predefined activities may include both sequential and concurrent activities. The exact activity sequences depend on the dining settings described in later subsections. Random activities mean that the dining participants are more or less free to follow their natural eating behaviors within the assumptions of our system defined in Section 2.4.

Dining Settings (afternoon tea vs. dinner)

There are two dining settings: an *afternoon tea* setting and a Chinese-style *dinner* setting. The dinner setting is more complex than the afternoon tea setting since it

involves a larger number of food containers. We describe these two settings in more details as follows. In the afternoon-tea scenario, participant(s) enjoyed an afternoon tea with a cake, a pot of tea, sugar, and milk. The objects (food containers) on the intelligent table are shown in Fig. 5-(a), including a tea pot, a cake plate, a sugar jar, a milk creamer, personal cake plates, and tea cups. The personal cake plates and tea cups are placed on the cells in front of each participant. The cake plate is placed on one center cell. The tea pot, the sugar jar, and the milk creamer are placed together on another center cell. Possible high-level dietary behaviors are *transferring-cake (to a personal plate)*, *pouring-tea (to a personal cup)*, *eating-cake (from a personal plate)*, and *drinking-tea (from a personal cup)*. In the dinner scenario, three dining participants enjoyed a sumptuous dinner with three shared dishes, one shared soup bowl, and a shared rice bowl. The objects (food containers) on the table are shown in Fig. 6-(a), including these shared plates & bowls, as well as personal bowls located on cells in front of each of three participants. Possible high-level dietary behaviors are *transferring-food (to a personal bowl)* and *eating-food (from a personal bowl)*. Note that given the weight and type of the food items consumed, it is relatively straight-forward to compute the calorie count by looking up a nutritional table for these food items.

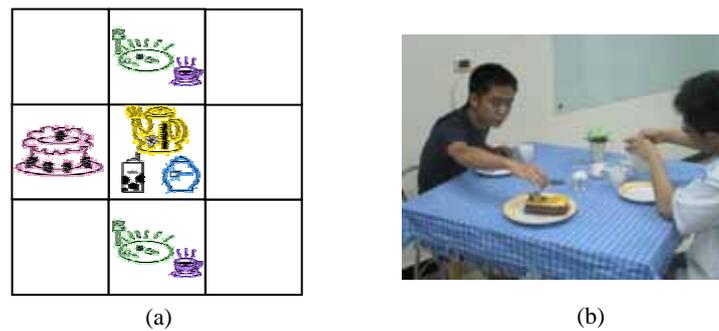


Fig. 5. Afternoon tea scenario showing the placements of table objects and participants

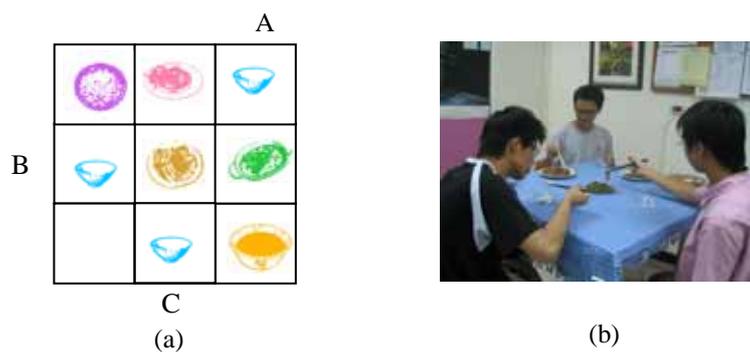


Fig. 6. Chinese-style dinner scenario showing the placements of table objects and participants

We describe the result for each of four dining scenarios as follows. A summary of their experiment results are shown in Table 4.

4.1 Dining Scenario #1: Afternoon Tea – Single User – Predefined Activity Sequence

The first dining scenario involves the afternoon tea setting, single user, and predefined activity sequence. The predefined activity sequence is consisted of the following steps: (1) cut a piece of cake and transfer it to the personal plate; (2) pour tea from the tea pot to the personal cup; (3) add milk to the personal plate from the creamer; (4) eat the piece of cake from the personal plate; (5) drink tea from the personal cup; and (6) add sugar to the personal cup from the sugar jar. This 6-steps sequence is repeated twice during the experiment. The results are shown in Table 4. Based on our measurements, the dietary behavior’s recognition accuracy (i.e., transfer & eat events) is 100%. This result is expected, given that the predefined activity sequence has been anticipated and tested extensively during our prototyping. In addition, this scenario involves only a single user with no concurrent interactions.

Table 4. Experimental results for 4 dining scenarios & their recognition accuracy.

Scenarios			Event Statistics		Results	
Dining Scenarios	# of users	Activity Sequence	Time Duration (seconds)	# of Dietary Behavior	Dietary Behavior Recognition Accuracy	Dietary Weight Recognition Accuracy
#1 Afternoon tea	1	Predefined	73	12	100%	-
#2 Afternoon tea	2	Predefined	162	24	100%	-
#3 Afternoon tea	2	Random	913	78	79.49%	-
#4 Chinese style dinner	3	Random	1811	162	83.33%	82.62%

4.2 Dining Scenario #2: Afternoon Tea – Two users – Predefined Activity Sequence

The second dining scenario also involves the afternoon tea setting and predefined activity sequence, but with two users performing concurrent activities. The predefined activity sequence is consisted of the following steps: (1) A cuts cake and transfers it to A’s personal plate; (2) B pours tea from the tea pot to B’s personal cup; (3) A pours tea to A’s personal cup while B cuts a piece of cake and transfers it to B’s personal plate; (4) A adds sugar from the sugar jar to A’s personal cup while B adds milk from the creamer to B’s personal up; (5) A eats cake and B drinks tea; (6) B eats cake from B’s personal plate while A drinks tea from A’s personal cup; and (7) A pours tea from the tea pot to both A’s and B’s personal cups. This 7-steps predefined activity sequence is repeated twice during the experiment. The results are shown in Table 4.

Based on our measurements, the dietary behavior recognition accuracy is 100%. This result shows that our table is accurate in recognizing concurrent activities from multiple participants.

4.3 Dining Scenario 3: Afternoon Tea - Two Users – Random Activities

The third dining scenario involves the afternoon tea setting and two users, but with random dietary activities. Random activities mean that the table participants are more or less free to perform any impromptu dietary behaviors for 913 seconds over the table within the bound of our assumptions described in Section 2.4. The results are shown in Table 4. Based on our measurements, the recognition accuracy is 79.49%. Table 5 shows the recognition accuracy for each of the two dietary behaviors. The eat events have better recognition accuracy than the transfer events, because they can be directly deduced by personal container’s *Weight-Change* event.

To determine the causes for the misses in activity recognition, we videotaped the afternoon tea scenario. By analyzing the video in combination with the system event logs, we derive four main *leading causes* shown in Table 6. They are described as follows.

Table 5. The accuracy of activity recognition under afternoon tea scenario

Dietary Behavior	# of Actual Events	Recognition Accuracy
Transfer event	41	70.73%
Eat event	37	89.19%

- (c1) *Event interference within the weighing cell’s weight stabilization time*: for activities such as putting down an object on the table, it takes about 1.5 seconds for our weighing sensor to output a stable weight value. If two events occur on the same cell and their time interval is less than the weighing cell’s stabilization time, our system cannot differentiate these two *Weight-Change* events. Instead, our system will incorrectly recognize them as a single *Weight-Change* event. Consider the case where user A puts down the tea pot at cell X while user B immediately (within 1.5 seconds) grabs a sugar cube from the sugar jar on the same cell X . There are actually two *Weight-Change* events of amount (Δw_1) and of amount $(-\Delta w_2)$. However, due to two events interfering with each other within the weight stabilization time, our system can only detect one erroneous *Weight-Change* event of amount $|\Delta w_1 - \Delta w_2|$.

Table 6. Causes of miss recognition in afternoon tea scenario #3. There are 78 activities analyzed from the video log. The number of misses counts both false positives and false negatives.

Causes of misses	# of misses of transfer events	# of misses of eat events	Total
(c1) Event interference within the weighing cell’s weight stabilization time	6	2	8
(c2) Weight matching threshold	2	0	2
(c3) (c3) Slow RFID sample rate	3	0	3
(c4) Noise from weighing cell	1	2	3
Total of misses	12	4	16

- (c2) *Weight matching threshold*: the current threshold value is set to be four grams to filter out noises in the weight readings from weighing cells. However, in some cases, such as transferring one cube of sugar, this threshold value may still be too large. As a result, it may lead to false weight matching involving unrelated weight transfers of similar amounts. Consider the example that user A is removing a cube of sugar from the sugar jar. This results in a *Weight-Change* of approximately four grams in the sugar jar. At the same time, user B is transferring food weighted approximately eight grams. Eight grams is twice as much as four grams, but they are still within the weight matching threshold. Therefore, this leads to false weight matching. To address this issue, we may change the weight matching threshold to be a percentage of transferred weight rather than an absolute value of four grams.
- (c3) *Slow RFID sample rate*: we have found cases when a user picks up a cup and quickly puts it down. This interval is less than the amount of time the RFID reader performs one round of reading over nine antennas. Therefore, a *Weight-Change* event is generated without any corresponding RFID-Presence event. This leads to false inference.
- (c4) *Noises from weighing cells*: although we ask users not to touch the table, some still do during the experiment out of personal habits. This leads to erroneous generation of *Weight-Change* events.

4.4 Dining Scenario 4: Chinese-style dinner - Three users – Random activities

The fourth dining scenario involves the Chinese-style dinner setting, three users, and random dietary activities for 1811 seconds. Similar to the third scenario, three participants perform impromptu dietary behaviors within the bound of our assumptions described in Section 2.4. The results are shown in Table 4. Based on our measurements, the recognition accuracy is 83.33%. Note that increasing number of table participants only slightly increases the activity rate. The reason is that as the number of table participants increases, out of politeness they try to go the dishes less frequently to avoid *in-the-air conflicts* over the dishes.

Table 7 shows the recognition accuracy (for the transfer and eat events) and weight accuracy for each of dietary behaviors. The weight accuracy is computed as the ratio between the *measured* and the *actual* weight transferred or consumed during dietary behaviors. Both the recognition and weight accuracy for the food transferring behaviors are between 80~85%, except for dish A, which is fluid-covered food. The reason for lower accuracy on transferring fluid-covered food is that juices from the fluid-covered food can easily drip from the chopsticks during food transfer (from a very lousy chopstick user). The juice dripping leads to erroneous generation of *Weight-Change* events with both positive and negative values, causing mismatches in the weight matching algorithm. Furthermore, the weight accuracy of transferring dish A is low at 68.42%, because these transfer recognition misses can accumulate to a large weight sum. Similar to the afternoon tea scenario, the eat events have better recognition accuracy because they can be directly deduced from the personal container's *Weight-Change* event.

Table 7. The accuracy of the Chinese-style dinner scenario #4

Dietary Behavior	# of times	Recognition Accuracy	Weight Accuracy
Transfer dish A events	19	73.68%	68.42%
Transfer dish B events	29	79.31%	78.75%
Transfer dish C events	23	82.61%	79.19%
Transfer rice events	12	83.33%	81.88%
Transfer soup events	19	84.21%	80.16%
Eat events	60	88.33%	91.23%
Overall	162	83.33%	82.62%

To determine the causes for the misses in activity recognition, we videotaped the Chinese-style dinner scenario and analyzed the video in combination with the system event logs. We derive five main *leading causes* shown in Table 8. They are described as follows.

- (c1) *Segmented Weight-Change events*: during a lousy food transfer where a user drops a part of food back into the container or on the table, the weight matching algorithm fails because of the difference between weight change values of the container and the personal plate.
- (c2) *Eating before transferring food on personal containers*: this occurs when a user picks up a serving of food from a shared plate. However, before the user completes the transfer to his/her personal plate, he/she eats a bite of food. This violates one of our assumptions in Section 3.4 that eating must come from food in the personal plates. In this case, weight matching method fails to recognize the food transfer event due to the disappearing weight on the intermediate bite. Although the users are told about this restriction, some of them still do it out of personal habits.
- (c3) *Weight matching ambiguity*: weight matching ambiguity occurs when two unrelated *Weight-Change* events of similar weight values are mismatched by our system.
- (c4) *Noises from weighing cells*: the same as (c3) in afternoon tea scenario.
- (c5) *Slow RFID sample rate*: the same as (c4) in afternoon tea scenario.

Table 8. Causes of miss recognition in Chinese-style dinner scenario #4. There are 162 activities analyzed from the video log. The number of misses counts both false positives and false negatives.

Causes of misses	# of misses of transfer events	# of misses of eat events	Total
(c1) Segmented weight-change events	5	0	5
(c2) Eating before transferring food on personal containers	5	5	10
(c3) Weight matching ambiguity	7	0	7
(c4) Noises from weighing cells	3	2	5
(c5) (c3) Slow RFID sample rate	3	0	3
Total of misses	23	7	30

5 Related Work

The related work is organized into the following three categories: traditional dietary assessment methods, ubiquitous dietary tracking systems, and intelligent (tabletop) surfaces. The traditional dietary assessment methods consist of keeping food records, using twenty-four-hour recall, and filling food frequencies questionnaires [16]. In the food record method, food quantities can be either accurately weighed or estimated by household measures before a meal. The twenty-four-hour recall method asks a user to recall the amount of food intakes within the past 24 hours. Food Frequencies Questionnaires (FFQ) list popular food items and ask a user how often and how much these food items are consumed within a defined period, e.g., a week or a month. All traditional assessment methods fail to capture actual energy intakes precisely [3]. Most methods underreport actual energy intake by at least 20 percent. Some of the errors are inevitable because human beings tend to misreport their food intakes. In other words, underreporting errors can be higher (30 percent or more) for certain groups of users. In comparison, our method can achieve 80% plus accuracy, which is as good as the accuracy from these traditional assessment methods.

For the dietary-tracking systems, Mankoff *et al.* [9] has designed a low-cost tracking system based on scanning shopping receipts to estimate what food items people buy and consume. By analyzing the nutritional values of the purchased food items, their system detects missing nutrients and recommends healthier food items to achieve a better nutritional balance. However, their system does not perform individual dietary tracking. The purchased food items in a family setting may be consumed by different household members in different quantities. The household purchased food items can be considered healthy, but the dietary consumption of individual household member can be nutritionally unbalanced due to personal dietary preferences and habits.

Dietary tracking at the individual level has been proposed by Amft *et al.* [1]. Their approach is to place a microphone around a person's inner ear to detect chewing sound from the mouth. Since different types of foods (e.g., potato chips, apples, pasta, etc.) can give different chewing sound, their system can infer what a person is currently eating in his/her mouth. However, different food sources that vary in nutritional contents give out similar chewing sound, e.g., similar sound from drinking water vs. beer. Rather than tracking food intake from chewing sound, this work takes a different approach. It creates a smart dining table, enabling the table to track food transfers among containers and into the individuals' mouths.

The 3rd category of related work is about intelligent surfaces that can infer tabletop human-surface interactions. The closest system to our work is the load sensing table [14] from Lancaster University. They utilized four weighing cells installed at four corners of a rectangular table to acquire the positional information of tabletop objects, and infer interaction events such as adding, removing an object from the surface, or knocking an object over. They demonstrated success with these interaction events. However, their main limitation is recognizing *complex, concurrent interactions involving multiple objects*. For example, their positioning algorithm fails if two or more objects are moved concurrently on the tabletop surface. In comparison, this paper expects such complex, concurrent interactions to be relatively common in family dining scenario; therefore, they are the paper's target.

Other related but less relevant works apply load sensing to derive context information. Smart floor [11] demonstrated that by applying pressure sensors underneath the floors, it is possible to identify users and to track their locations. The posture chair by Selena [10] deployed two matrices of pressure sensors (called *pressure cells*) in a chair to recognize the posture of children, and then infer their affective interest level. To our knowledge, no work that attempts to address complex, concurrent person-object interactions from a load sensing surface. This paper is believed to be the first to augment the load sensing surface with a RFID surface to enable tracking of multiple, concurrent person-object interactions over a tabletop surface.

6 Conclusion and Future Work

We are what we eat. This paper describes the design and implementation of our diet-aware dining table. We have augmented an ordinary dining table with two layers of sensor surfaces underneath the table – the RFID surface and the weighing surface. Given certain assumptions, the diet-aware dining table automatically tracks what and how much each individual eats from the dining table over the course of a meal. We have performed several experiments, including live dining scenarios (afternoon tea and Chinese-style dinner), multiple dining participants, and random concurrent activity sequences. Our experimental results have shown reasonable recognition accuracy of around 80%, which is at least as good as the accuracy of the traditional dietary assessment methods.

Our future work will further improve the recognition accuracy, address some of the main causes of inaccuracy from our experimental results, and relax some of the assumptions and restrictions. Note that some of the restrictions can be solved by making simple design changes. For examples, the current prototype does not allow hands or elbows on the table. To relax this restriction, we can add a slightly protruding frame around the edge of table, so that users can rest their elbows on the frame without affecting our system. We also believe in multi-sensor approach. For example, by deploying a video camera above the table, it is possible to observe events that cannot be detected by RFID and weighing surfaces. Since this table can track tabletop person-food interactions in real time, we hope to build *just-in-time persuasive feedbacks* to encourage better healthy dining behaviors. For examples, users under a dietary program could benefit from weekly warnings of unhealthy food. Patients who cannot consume high amount of sugar (diabetes patients) or water (patients with heart problems) could receive just-in-time notifications when they are over their recommended daily consumption.

ACKNOWLEDGMENTS

This work was partially supported by grants from Taiwan NSC, Quanta Computer, and Intel (#94-2218-E-002-055, #94-2622-E-002-024, #94-2218-E-002-075).

References

- [1] Amft, M. Stäger, P. Lukowicz, G. Tröster, "Analysis of Chewing Sounds for Dietary Monitoring", In Proceedings of Ubicomp 2005, M. Beigl, S. Intille, J. Rekimoto, H. Tokuda (Eds). Lecture Notes in Computer Science, Vol 3660, ISBN 3-540-28760-4; Springer Verlag, Tokyo, Japn, September 2005
- [2] J. Beidler et al. The PNA project. In Proceedings of the sixth annual CCSC northeastern conference on The journal of computing in small colleges, pp. 276–284. The Consortium for Computing in Small Colleges, 2001.
- [3] Dietary Assessment, <http://encyclopedias.families.com/dietary-assessment-516-518-efc>
- [4] Dishman, E. Inventing wellness systems for aging in place. Computer, May 2004, 34-41
- [5] S. Hankinson et al., editors. Healthy Women, Healthy Lives: A Guide to Preventing Disease. A Harvard Medical School book. Simon & Schuster Inc., 2001.
- [6] INSIGHT 19. Beliefs and attitudes of americans toward their diet. USDA Center for Nutrition Policy and Promotion, June 2000.
- [7] INSIGHT 20. Consumption of food group servings: People's perceptions vs. reality. USDA Center for Nutrition Policy and Promotion, October 2000.
- [8] JESS. <http://herzberg.ca.sandia.gov/jess/>
- [9] J. Mankoff, G. Hsieh, H. C. Hung, S. Lee, E. Nitao, "Using Low-Cost Sensing to Support Nutritional Awareness", In Proceedings of Ubicomp 2002, G. Boriello and L.E. Holmquist (Eds). Lecture Notes in Computer Science, Vol 2498, ISBN 3-540-44267-7; Springer Verlag, Gothenburg, Sweden, September 2002
- [10] S. Mota and R. W. Picard (2003), "Automated Posture Analysis for Detecting Learner's Interest Level." Workshop on Computer Vision and Pattern Recognition for Human-Computer Interaction, CVPR HCI, June, 2003.
- [11] R. J. Orr and G. D. Abowd. The Smart Floor: A Mechanism for Natural User Identification and Tracking. GVU Technical Report GIT-GVU-00-02, 2000.
- [12] Rosenberg, I. H. (1996). Nutrition research: an investment in the nation's health. Nutrition Review, 54, s5-s6.
- [13] R. Weinsier and C. Krumdieck. Dairy foods and bone health: examination of the evidence. American Journal of Clinical Nutrition, 72(3):681–689, September 2000.
- [14] Schmidt, M. Strohbach, K. Van Laerhoven, A. Friday and H.-W. Gellersen. "Context Acquisition based on Load Sensing". In Proceedings of Ubicomp 2002, G. Boriello and L.E. Holmquist (Eds). Lecture Notes in Computer Science, Vol 2498, ISBN 3-540-44267-7; Springer Verlag, Gothenburg, Sweden, September 2002
- [15] Mark Weiser, John S. Brown, The Coming Age of Calm Technology, 1996.
- [16] Thompson FE, Byers T. (1994) Dietary assessment resource manual. J Nutr. 124:2245S-2318S
- [17] United States Department of Agriculture. USDA Nutrient Database for Standard Reference. <http://www.nal.usda.gov/fnic/foodcomp/Data/>